

Predicting the synaptic information efficacy in cortical layer 5 pyramidal neurons using a minimal integrate-and-fire model

Michael London · Matthew E. Larkum ·
Michael Häusser

Received: 3 July 2008 / Accepted: 23 September 2008
© The Author(s) 2008. This article is published with open access at Springerlink.com

Abstract Synaptic information efficacy (SIE) is a statistical measure to quantify the efficacy of a synapse. It measures how much information is gained, on the average, about the output spike train of a postsynaptic neuron if the input spike train is known. It is a particularly appropriate measure for assessing the input–output relationship of neurons receiving dynamic stimuli. Here, we compare the SIE of simulated synaptic inputs measured experimentally in layer 5 cortical pyramidal neurons in vitro with the SIE computed from a minimal model constructed to fit the recorded data. We show that even with a simple model that is far from perfect in predicting the precise timing of the output spikes of the real neuron, the SIE can still be accurately predicted. This arises from the ability of the model to predict output spikes influenced by the input more accurately than those driven by the background current. This indicates that in this context, some spikes may be more important than others. Lastly we demonstrate another aspect where using mutual information could be beneficial in evaluating the quality of a model, by measuring the mutual information between the model’s output and the neuron’s output. The SIE, thus, could be a useful tool for assessing the quality of models of single neurons in preserving input–output relationship, a property that becomes crucial when we start connecting these reduced models to construct complex realistic neuronal networks.

Keywords Synaptic information efficacy ·
Linear integrate-and-fire model · Predicting every spike ·
Layer 5 cortical pyramidal neuron

1 Introduction

Neurons are widely considered to be the fundamental computing elements of the brain. As such their input–output relationship is of prime importance. This relationship is often quantified by measuring the neuron’s frequency–current curve ($F-I$ curve). This involves injecting current steps of various amplitudes into the neuron and measuring the frequency of the output spikes. With this approach one can also calculate the current threshold (i.e. the minimum current that will cause the neuron to emit an action potential) and the gain (how many more output spikes will be emitted per additional current step). Some types of neurons, most significantly, regular firing cortical pyramidal neurons, typically have a threshold linear $F-I$ curve, a property which many computational models of neural networks rely on [Vogels et al. \(2005\)](#).

However, the frequency–current relationship does not necessarily capture the complete input–output function of the neuron because it assumes that there is no information in the temporal pattern of the input or output. For example, neurons have been shown to also be able to respond reliably not only to the mean current they receive but also to higher frequency components ([Mainen and Sejnowski 1995](#)). Consequently, to correctly characterize the input–output relationship one needs to use a method that also captures the relationship when dynamic stimuli are present and the responses are dynamic as well. Several such methods have been suggested in the past ([Arcas et al. 2003](#); [Chance 2007](#)). Considering the motto of this Special Issue, “predicting every

M. London (✉) · M. Häusser
Department of Physiology, Wolfson Institute for Biomedical
Research, University College London, Gower Street,
London WC1E 6BT, UK
e-mail: m.london@ucl.ac.uk

M. E. Larkum
Department of Physiology, University of Bern,
3012 Bern, Switzerland

spike”, the synaptic information efficacy (SIE) (London et al. 2002) is a particularly appealing approach because it specifically uses prediction power to quantify the input–output relationship. In essence the SIE is the mutual information between an input spike train arriving at a synapse and the output spike train of the postsynaptic neuron. Here we show that we can measure SIE experimentally given a reasonable amount of data. This allows us to answer the question: how well can a reduced model of the neuron capture its input–output relationship?

Clearly if a model fits a real neuron well enough such that given the input current it can predict precisely the output spike train, the SIEs calculated using the real neuron and the model would match perfectly (because the SIE only relies on the input and output spike trains and both would be identical in the experiment and model). However, here we take a slightly different perspective. We are aiming to characterize the input output relationship by measuring the effect of a single synapse on the output of the neuron (this could be extended to the effect of a group of correlated inputs). However, one synapse on its own typically does not operate in a vacuum, rather a neuron is bombarded by inputs arriving at many of its synapses. We assume that the input current into a neuron is composed of a mixture of many synaptic inputs some of which might carry information and some could potentially reflect noise. We are interested only in the information of one of these inputs. We, thus, inject a fluctuating current representing the background synaptic activity as well as one controlled input synapse. We then ask how well can a reduced model capture the SIE of this one, well-defined, input in the presence of this background activity. We show here that even a highly simplified model that is less than perfect in predicting the exact output spikes of the neuron, can capture the SIE very well. This is surprising because the SIE relies only on the relationship between input spikes and output spikes. The explanation is that with respect to the relationship between one input and the output, the output spikes that are affected by this input are more important for the SIE than others, and thus a model succeeding in predicting these spikes is good enough. This demonstrates the point that a model should be constructed in a specific context: under some circumstances it might turn out that there is no need to predict every spike with the same accuracy because some spikes are more significant than others.

2 Experimental procedures

2.1 Slice preparation and electrophysiology

All experiments were performed using acute parasagittal cortical slices from rat somatosensory cortex, prepared using standard techniques in accordance with institutional and

national guidelines. Sprague–Dawley rats (P18–P25) were anaesthetized with isoflurane (Abbott Labs, Kent, UK) and decapitated. The cortex was surgically dissected and cut sagittally at a slice thickness of 300 μm on a vibratome (VT1000S, Leica Microsystems). After slicing slices were incubated for 30–45 min at 32°C in external solution bubbled with 95% O_2 and 5% CO_2 , and then were kept at room temperature until recording (at 34°C). The standard external recording solution contained (in mM): 125 NaCl, 2.5 KCl, 2 CaCl_2 , 1 MgCl_2 , 25 NaHCO_3 , 1.25 NaH_2PO_4 and 25 D-glucose.

Neurons were visualized with a 40 \times water immersion objective (Zeiss) using IR-DIC optics on an Zeiss upright microscope equipped with an IR camera (C2400-07, Hamamatsu, Tokyo, Japan) connected to a video monitor. Pipettes were prepared from thick-walled filamented borosilicate glass capillaries (Harvard Apparatus, Kent, UK) pulled on a two-stage puller (Model PC-10, Narishige, Tokyo, Japan) with tip resistances between 5 and 6 $\text{M}\Omega$ for whole-cell somatic recording. The composition of the internal solution was (in mM): K gluconate 105, Hepes 10, MgCl_2 2, MgATP_4 2, sodium phosphocreatine 10, GTP 0.3 and KCl 30, with 2 mg/ml biocytin, at pH 7.3.

Recordings were performed using an Axoclamp 2B amplifier (Axon Instruments, CA, USA) connected to a Macintosh computer via an ITC-18 board (Instrutech, Port Washington, NY, USA). Data was acquired with IGOR PRO (version 5.0, Wavemetrics, OR, USA). The current–clamp signal was Bessel filtered at 3 kHz and sampled at 10 kHz. No correction of the liquid junction potential was performed. To generate “background” synaptic input, white noise was convolved with an exponential decay with a time constant of 1 ms. The variance and mean of this background current were adjusted during the experiment to give a spike rate of ~ 10 spikes/s with a CV ~ 0.75 . Data were analyzed using custom macros written for IGOR Pro (Wavemetrics), Matlab (Mathworks) and the CTW algorithm (see below).

2.2 Computer modelling

We have used the simplest version of the leaky integrate-and-fire model. The sub-threshold integration of the model is described by Eq. 1 which is a discrete version of the continuous form, and has three passive parameters: R , membrane resistance in Ω ; τ , membrane time constant in s; and V_{rmp} , resting membrane potential in mV. The model has three additional parameters to describe spikes: θ , voltage threshold in V; V_{AHP} , the voltage to which the membrane potential is reset after each spike; and τ_{ref} , the refractory period during which the membrane potential is clamped to V_{AHP} .

$$v(i+1) = v(i) + dt \left(\frac{1}{\tau} (v(i) - V_{\text{rmp}}) \right) + \frac{RI_{\text{inj}}(i)}{\tau}. \quad (1)$$

We split the process of fitting the model to the data into two stages. In the first stage we fit the sub-threshold response,

and in the second we fit the threshold, AHP and refractory period. V_{rmp} is easily obtainable from the beginning of the voltage trace where no current is injected. However there were often small drifts (<2 mV) in the resting membrane potential due to the length of the experiment and the long period of intensive current injections. Thus V_{rmp} is the only parameter that was adjusted for each trace. All remaining parameters were adjusted once for the training trace and were used for all the rest of the inputs as is. We then identify stretches of the voltage trace which do not contain spikes and use a simplex algorithm to obtain R and τ (Fig. 1b). We then use a spike-triggered average to determine the threshold, AHP and refractory period (Table 1).

2.3 Synaptic information efficacy analysis

Estimation of the synaptic information efficacy was done using the Context Tree Weighting (CTW) algorithm as described in London et al. (2002). The CTW algorithm (Willems et al. 1995) in its simplest form is a universal lossless compression algorithm for binary sequences. For each symbol in the sequence it computes the probability of the symbol to occur based on the previous history of the sequence and uses this probability for compression using arithmetic encoding (Cover and Thomas 1991). Having the probability of each symbol given the history is very useful for computing the entropy of the sequence. Moreover the algorithm can easily be modified to estimate the probability of a symbol given the history of the sequence and the history of another sequence (in our case the input spike train). This enables us to compute the conditional entropy, which is required for computing the mutual information. One advantage for using the CTW algorithm is that its redundancy is bounded above uniformly over all data sequences of arbitrary length. The other major advantage is that increasing the window over which the history is considered does not bias the estimator which means that the sampling catastrophe that often happens with other methods is not an issue with the CTW. Comparisons of the performance of the CTW with other methods (Kennel et al. 2005; Gao et al. 2008) have shown that it can yield more accurate results for both synthetic and neuronal data. We used a custom written JAVA code which integrates with Matlab and is publicly available (<http://www.dendrites.org/~mikilon/Code/CTW>). Empirically we find that we need few hundreds of spikes for the algorithm to converge and hence with 53 s of recording at spike rate of ~ 10 spikes/s the algorithm yields stable estimates. We also found that because the entropy estimation requires only one sequence while the conditional entropy requires two sequences there is a consistent bias which causes independent sequences to have a mutual information of approximately -1 bits/s (while the mutual information should be zero). In

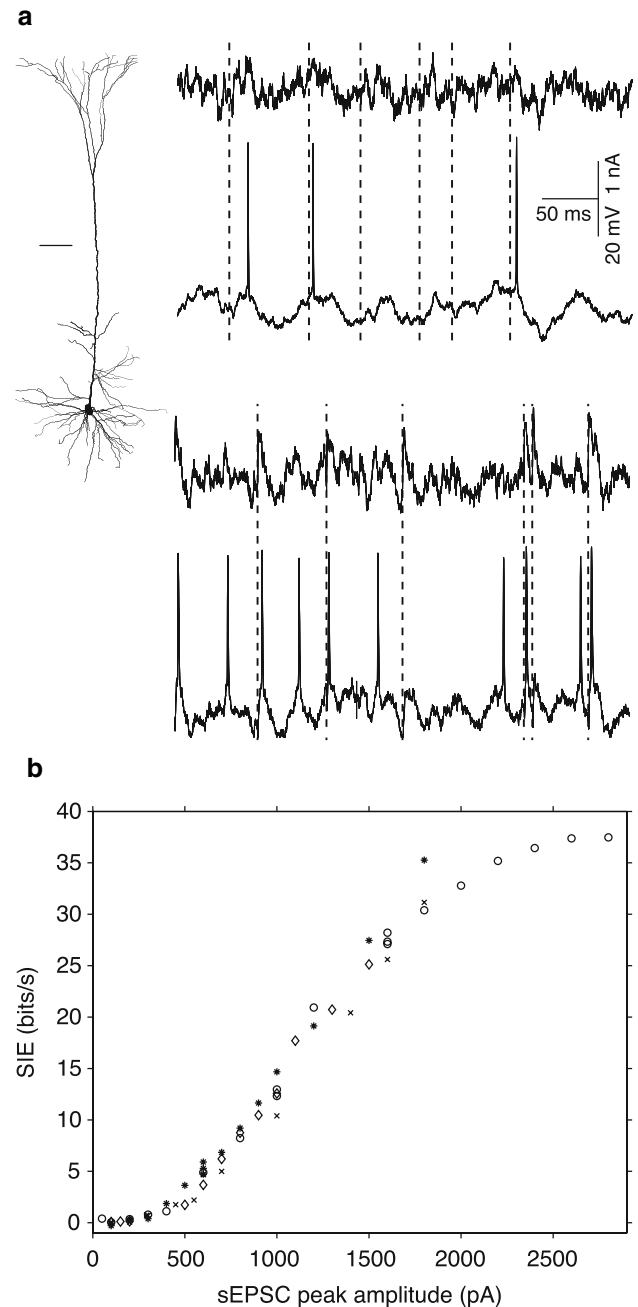


Fig. 1 **a** Whole-cell recordings were made from layer 5 pyramidal cells in rat somatosensory cortex in vitro (scale bar 100 μm). A fluctuating current was injected into the soma and the resulting voltage response was recorded. The current was composed of filtered white noise and a series of sEPSCs. Two short duration examples are shown (top sEPSC peak amplitude 300 pA and bottom 1,600 pA). The times of the sEPSCs are marked by dashed lines. Note that for small sEPSC most events do not cause a spike, while for the larger amplitude they do. **b** The synaptic information efficacy (SIE) computed from the “input spike train” (times of the input sEPSCs) and the neuron’s output spike train is depicted as a function of sEPSC peak amplitude for four cells (four different symbols)

order to overcome this bias we estimated the mutual information as a difference of two conditional entropies as follows:

Table 1 Parameters of the model for the different cells

Cell	R_{input} (M Ω)	τ (ms)	V_{rmp}	Threshold	V_{AHP}	τ_{ref} (ms)
1	48.0	13.0	−65	−42	−52	5
2	74.7	8.5	−65	−41	−51	6
3	43.7	14.5	−76	−58	−68	8
4	54.1	10.3	−66	−40	−43	2

$$\text{SIE}(S_{\text{out}}; S_{\text{in}}) = \hat{H}(S_{\text{out}}|S_{\text{in}}^{\text{shuff}}) - \hat{H}(S_{\text{out}}|S_{\text{in}}) \quad (2)$$

where S_{in} and S_{out} are the input and output spike trains respectively, and $S_{\text{in}}^{\text{shuff}}$ is a surrogate spike train obtained by randomly shuffling the inter-spike intervals of the input spike train such that any temporal correlations between the input and the output are removed. The first term of the right hand side should be equal to the entropy of the output spike train, but has a small positive bias, however in this way the two terms on the right have the same bias and hence the estimated SIE is more accurate.

3 Results

3.1 Measuring SIE in layer 5 pyramidal cells

Using somatic whole-cell recordings from layer 5 pyramidal cells, we injected a fluctuating current to mimic background synaptic activity, causing the neuron to spike at approximately 10 spikes/s. On top of the fluctuating current we added a sequence of EPSC-shaped currents (sEPSC) with a fixed time course (see Sect. 2, Fig. 1) and varied their peak amplitude (the range of spike output rates in these experiments was 8–20 spikes/s). Figure 1 illustrates the injection protocol and the resulting voltage response of the neuron. We then used the sequence of times of the input sEPSC as an input “spike train” and computed the mutual information between this spike train and the output spike train emitted by the neuron to obtain the SIE for every current amplitude. The SIE as a function of current amplitude is depicted in Fig. 1b for four different neurons. For a small input (e.g. 0–400 pA) the SIE is typically very small (less than 2 bits/s). Such an input is illustrated in the top example trace (Fig. 1a: 300 pA). In this case the input sEPSC is smaller than the standard deviation of the background current and cannot be identified by inspection. Moreover, for a given output action potential it is sometimes difficult to tell whether it was caused by the sEPSC or not. Nevertheless even such a small current sometimes clearly influences spike generation (see second and fifth inputs in the example) and changes the probability of the neuron to emit an action potential. This effect, though too small to be noticeable by eye, is detected by the SIE. For larger input amplitudes the SIE grows linearly with input

size. When the current is large enough to rise above the noise an output spike will frequently follow the input sEPSC. For the example shown in Fig. 1a (1,600 pA), 85% of sEPSCs are followed by an output spike within 4 ms. Typically a failure to trigger a spike occurs because the input just followed an output spike while the neuron was refractory, or the input followed a random large hyperpolarization (see third and fifth inputs in the example). For even larger sEPSCs (> 2 nA) the SIE saturates at approximately 35–40 bits/s (for the bin size of 3 ms we use here). This saturation occurs because the information is bounded by the minimum of the entropies of the input and the output. Intuitively one cannot gain more information than there is in the output or than is supplied by the input (see also London et al. 2002, Fig. 2). The reason that only one cell shows this saturation is that we did not drive the neurons with input amplitudes large enough to show this effect (this is generally detrimental to the recording and the health of the neuron).

The relationship between SIE and input size is therefore nonlinear, and follows a sigmoidal trajectory. This relationship qualitatively agrees very well with earlier predictions based on simulations (London et al. 2002) where its properties are discussed in more detail. This non-linear shape is in contrast to the linear relationship expected between cross-correlation and input size as described in Herrmann and Gerstner (2001) and is primarily due to the nonlinear properties of the mutual information measure. In the following sections we would like to explore how well a simplified model can capture the relationship between SIE and input amplitude described above.

3.2 Constructing simplified integrate and fire models to match experimental data

Many successful sophisticated methods for constructing simplified models of spiking neurons are described in this issue, and in the literature (Brunel and Latham 2003; Fourcaud-Trocme et al. 2003; Rauch et al. 2003; Keren 2005; Jolivet et al. 2006; Badel et al. 2008). Nevertheless we deliberately choose here to take the most simplistic view and fit the recordings with a leaky integrate-and-fire model in order to demonstrate the power of the SIE approach. The model has three passive parameters (R membrane resistance in Ω , τ membrane time constant in s and V_{rmp} resting membrane potential in V).

The model has three additional parameters to describe spikes: θ , voltage threshold in V; V_{AHP} , the voltage to which the membrane potential is reset after each spike; and τ_{ref} , the refractory period during which the membrane potential is clamped to V_{AHP} . Figure 2 describes the fitting process (see Sect. 2). We split the process of fitting the model into two stages. As the current–voltage relationship for these cells is linear over a wide range of subthreshold potentials

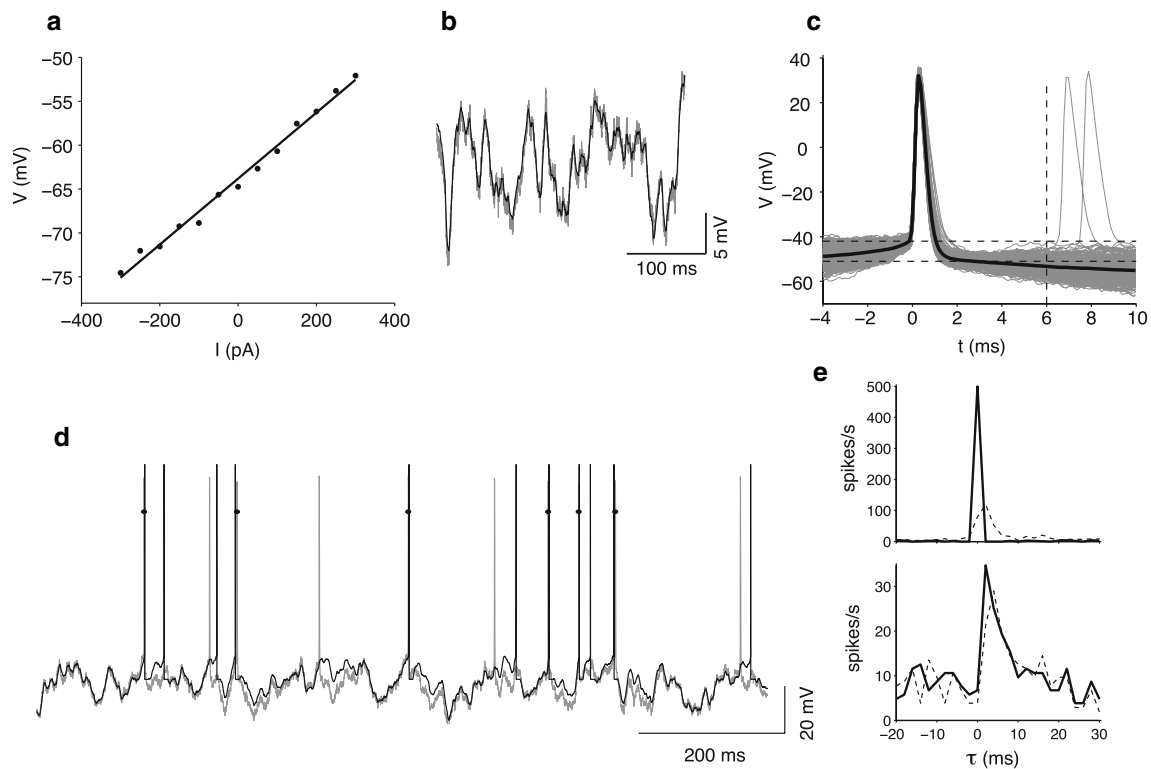


Fig. 2 Fitting the data with a linear integrate and fire model. **a** Current voltage relationship showing a linear relationship over wide range of membrane potential. **b** Fitting a subthreshold voltage trace using a simplex algorithm (grey data, solid black model) yields the parameters R_{in} , τ_m and V_{rmp} . The algorithm typically results in a very good fit to the data. **c** All spikes in the training trace are aligned, shown in grey. Their average (spike triggered average) is shown in solid black. The voltage threshold (θ), V_{AHP} and the refractory period τ_{ref} are obtained from this data. **d** An example of a recorded voltage trace (grey) and the models prediction (solid black). Markers depict time points in which model

successfully predicted output spike within 2 ms precision. The model, however, has a significant number of misses and false positives. **e** The autocorrelogram of the output spike train (top panel solid) of the neuron compared to the crosscorrelogram of between the output of the neuron and the output of the model (top panel dashed line), shows that indeed the model has many misses. The crosscorrelogram between the input spike train and either the output spike train of the neuron (solid bottom panel) or the output spike train of the model (bottom panel dashed line) shows that the model predicts the output spikes that are correlated with the input very well

(Fig. 2a), we choose to use a standard minimization procedure to fit the subthreshold response in the first stage. In the second step we fit the threshold, V_{AHP} and refractory period based on the spike triggered average curve (see also Sect. 2).

We used one trace with a small sEPSC amplitude (<300 pA) for training. We then used this model on the rest of the input currents (with different amplitudes). Predicting every spike with such a simple LIF model is impossible. The gamma coincidence factor (GCF) (Kistler et al. 1997) measures how many spikes of the model coincide with the real spikes of the neuron (with the required tight condition of 2 ms accuracy), normalized to chance level (i.e. if the two were spiking in their mean rate but completely independent of each other). For the input currents used in this study the model achieved a GCF of 0.5 ± 0.1 . This score is expected for the standard linear integrate-and-fire model but is lower than that of adaptive integrate-and-fire models presented in this Special Issue which achieve a GCF of ~ 0.8 .

3.3 Comparing the effect of synaptic input in the experiment and in the model

As the SIE is measuring the efficacy of the synaptic input using only the input and output spike trains, clearly, as pointed out above, if the model could predict the output of the cell perfectly for any given input, then the SIE for the model and the real neuron would be identical. However, as we see in Fig. 2 for the integrate and fire model some of the original spikes are missed, some are spurious and some are shifted compared to the real cell.

Figure 3a depicts the SIE as a function of sEPSC current amplitude for one of the neurons presented in Fig. 1b. The model shown in Fig. 2 was used to integrate the same input currents, which were used in the experiments. The SIE computed for the model shows a remarkable similarity to the SIE computed for the data. Similar results were obtained for two of the other cells presented in Fig. 1 (RMS—2.3, 3.1, 3.13 bits/s for cells 1, 2 and 3, respectively). The explanation for

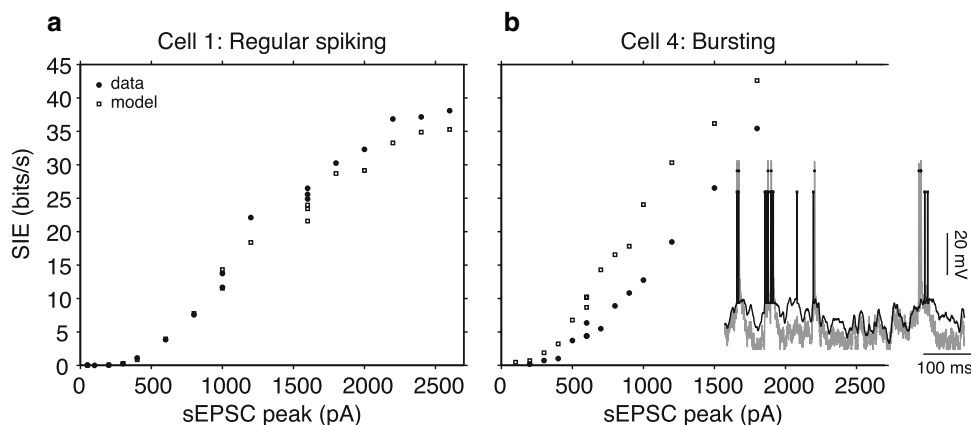


Fig. 3 Comparing the SIE of the model and the data. **a** The SIE of the data of one cell as a function of sEPSC amplitude and the corresponding SIE computed for the LIF model shown in Fig. 2. Very close agreement between the two is achieved. This model was a regular firing cell. Note that the model was constructed using one trace only and then used with

the same parameters for all other current amplitudes. **b** Same as **a**, only here the cell was firing bursts of action potentials in response to the current injection. The LIF cannot properly fit a bursting cell, and tends to burst more for large inputs (see *inset*). The result is that the SIE of the model overestimates the SIE of the real neuron (see text)

this close agreement is quite interesting. Both the fluctuating input current and the input spike trains are random and thus the neuron and the model have statistically similar output spike trains (even though they do not overlap on the millisecond time scale). This makes their entropies similar. The LIF model, while poor in predicting the spike caused by the background fluctuating current (e.g. GCF = 0.5 when the sEPSC is very small; see Fig. 2) is pretty good in predicting spikes that are driven by the input sEPSC. This makes the conditional entropy of the model and the neuron quite similar as well. All in all the SIE which is the difference between the entropy and the conditional entropy corresponds closely between the model and the experiments. In that sense the SIE is less sensitive to the precise timing of the spikes arising from the “background” current, and more sensitive to the timing of spike in relation to the input sEPSC. For cell 4 the agreement between the SIE of the model and of the experiments was not as good as in the other cells shown in Fig. 3b (RMS: 6.7 bits/s). Interestingly, this cell was firing bursts (typically of two spikes) in response to current injection at the soma, in contrast to the other three neurons. An example trace of recording from this neuron is shown in the inset, together with the corresponding model trace. Clearly it is impossible to describe this behavior with such a simple model. With our fitting strategy the model has the disadvantage that V_{AHP} is relatively high (-43 mV) and τ_{ref} is very short (2 ms) in order to account for the bursting activity. However, in contrast to the real neuron the model has no memory, and thus produces many more spikes per burst than the real neuron. Moreover, the real neuron tended to be quiet after a burst, while the model does not. The result shows that the model often reacts to the input with more spikes than the real cell, which increases the entropy of the output. However when the input is known (to the estimation algorithm), it accounts for these

spikes, which reduces the conditional entropy. The increase in entropy and reduction in conditional entropy compared to the real data, results in an overestimate of the mutual information, hence the overestimate of the SIE when computed with the model.

3.4 Using mutual information to evaluate the predictive power of the model

Till now we computed the mutual information between the input and the output of both the experimental data and of a model to evaluate how well the model captures the input–output relationship of the real cell. We can also use the mutual information in a different way: to evaluate directly how well the model predicts the output. The traditional way this is done (especially in this Special Issue) is using the GCF (Kistler et al. 1997). Similarly, we can measure the mutual information between the output spike train of the model and that of the real neuron. If the two are completely independent we should obtain zero mutual information. If, on the other hand, the model output is precisely that of the neuron, then the mutual information will be the entropy of that output spike train. Normalizing the mutual information by the entropy of the neuron’s output will yield a measure between 0 and 1, similar to the GCF.

It is important to understand that there are differences between the two measures. On one hand, if one is interested in a model that predicts every spike precisely, then of course the GCF provides a good measure to compare models. On the other hand, if the model predicts the exact output spike train but with a time shift of, say, 3 ms, then the GCF will be very small even though the output is by almost any measure identical. Similarly, if the model predicts a pair of output spikes for every single output spike of the neuron, again, the two

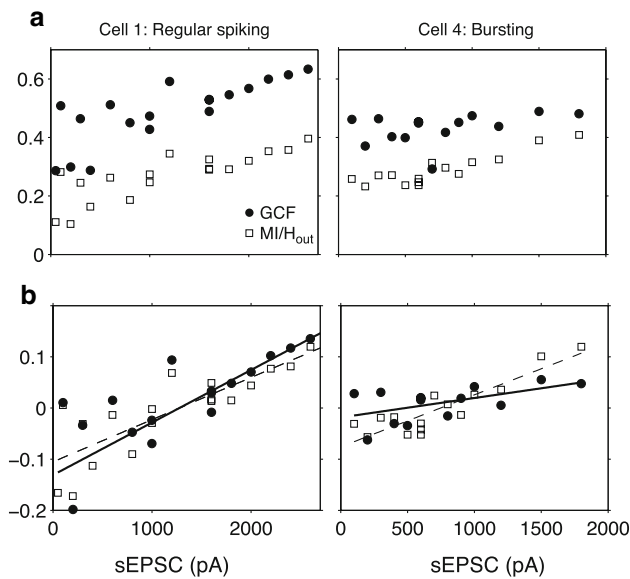


Fig. 4 Using mutual information to evaluate the prediction power of the model. **a** The GCF between the model output and the neuron’s output is depicted as a function of sEPSC amplitude (*filled circles*). The mutual information between the model’s output spike train and the output of the neuron, normalized by the output entropy of the neuron is also shown (*open squares*). **b** Same data as in **a** is depicted when the mean of each group is subtracted so that the trend can be compared. Data is fitted with linear regression line (GCF, *solid line*; normalized mutual information, *dashed line*). Note the good match for the regular spiking cell and the disagreement between the two measures for the bursting cell

spike trains would carry almost exactly the same amount of information, but the GCF will give a rather poor score for the performance of the model (because only ~50% of the spikes are predicted). The mutual information is, however, agnostic about such transformations, and thus suitable as a complementary method to evaluate the quality of models.

Figure 4 shows the GCF as well as the normalized mutual information for the two cells presented in Fig. 3a, b. For the first cell the two measures have very similar shape, only the mutual information gives lower scores for the model compared to the GCF. Very similar matches were obtained for the other two regularly spiking cells. For the bursting cell, however, the two methods gave different results. The GCF is relatively independent of the input (i.e. sEPSC) amplitude. The mutual information, on the other hand, grows linearly with the input size. Curiously, this is the model that performed least well in capturing the input–output relationship of the neuron by overestimating the SIE. This result hints that indeed the mutual information can pick up the relationship between the prediction of the model and the real data that the GCF ignores. However, in this case there is also the complication that the model that performs worst, yields better mutual information. There is however quite a simple explanation for this. In Fig. 3b we estimated the mutual information between the input and the output. As we showed, the model overes-

timated this information because for large inputs it tends to “over burst” compared to the real neuron (i.e. fire more spikes in a burst per input). Thus, when we know the input, we can reduce the entropy of the model’s output even more than we can do for the real neuron. Here when we compute the mutual information between the model’s output and the neuron’s output, the input is not known, but, for large enough sEPSCs, we already know that the model almost certainly produces a burst per each input. So to some extent the model output acts in a similar way as the input given to the synapse. In other words, bursts of the real neuron are well predicted by bursts of the model, even though there are more spikes in each burst, and this increases the mutual information between the model’s output and the neuron’s output.

4 Discussion

Understanding brain dynamics requires good network models. Such models must possess two properties: first, the elements of the model should be representative of real neurons in the sense that when given the right input they will produce a similar output to that of the real cell; second, the connectivity of the network should be known. Constructing simplified models of neurons that are able to “predict every spike” is obviously an important step, especially when realizing that the properties of the elements of the network might have a big effect on the network dynamics. However, simplifying the neuronal elements of the network imposes a problem on the connectivity, because of the difficulty of bridging between the knowledge of the biophysical properties of synaptic communication between neurons (Häusser and Roth 1997; Markram et al. 1997; Thomson and Lamy 2007) and its implications for the input–output properties of the neuron (Fetz and Gustafsson 1983; Douglas et al. 1996; Barbour et al. 2007). Succeeding in constructing a network out of simplified neuronal elements means that the effect of activating a connection between two real neurons on the output of the postsynaptic neuron will be the same in the model. This could be a very complicated task, especially as there exist non-linear interactions that may take place between inputs arriving at dendrites (Schiller et al. 1997; Larkum et al. 1999; Häusser and Mel 2003; Larkum et al. 2004; Polsky et al. 2004; London and Häusser 2005; Losonczy et al. 2008). It is especially important to be able to assess whether the connectivity being used indeed preserves the efficacy of the connections in the model. The SIE might provide a way to quantify this. We are interested in ways to quantify the efficacy of the synapse within the context of network activity so that we can assess whether the “anatomical connectivity” in the model also represents the “effective connectivity” in the real network.

Here we have used the synaptic information efficacy to quantify the ability of simple models to capture the somatic input–output relationship of layer 5 pyramidal neurons. We measured the SIE for real neurons and compared it for the SIE obtained from simple leaky integrate and fire models (LIF) that were fit to the data recorded from these neurons. Surprisingly, even models as simple as LIF can capture the input–output relationship very well, despite the fact that only about 50% of the output spikes of the neuron are accurately predicted by the model. The SIE is hence a good candidate for evaluation of additional models and in other contexts.

An advantage of the SIE is that it considers higher order correlations both within the input and output and between them. Indeed, as we demonstrate this has little effect if the neuron's response is as simple as predicted by an LIF. Nevertheless, most layer 5 pyramidal neurons do show more complicated, higher-order behaviour, such as bursts, especially when dendritic inputs are involved (Reuveni et al. 1993; Schiller et al. 1997; Larkum et al. 1999; Williams and Stuart 1999; Stuart and Häusser 2001; Larkum et al. 2004). An example of an extreme case is presented in Fig. 3b where the neuron responded with burst even in the case of somatic current injection. The SIE can be used in such cases without any additional modifications, and thus is a natural candidate for evaluating models in a context that goes beyond a simple input–output relationship. Moreover, we can use the mutual information between the model's output and the neuron's output (when identical input current is used for both) to evaluate the performance of the model. This has the advantage of considering the natural variability of the neuron when the same input is injected, as well as considering effects such as time shift delay and change in representations (i.e. replacing a burst by a single spike).

The overall conclusion is that when we pursue the goal of constructing a reduced model of a network, some aspects of mapping synapses onto their targets are relatively insensitive to how perfect the single model neuron is. Hence, in parallel to perfecting the reduced single neuron model, we can start addressing the complexities of connectivity in neural networks using simple models. We should start thinking about how to connect the neurons in the network such that we preserve their effective connections as they are in the real network, and the SIE could be a useful tool to achieve that.

Acknowledgments We thank Arnd Roth for helpful discussions and for his helpful comments on the manuscript. This work was supported by grants from the Wellcome Trust, the Medical Research Council and the Gatsby Charitable Foundation. Michael London was supported by a Long-Term Fellowship of the HFSP.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Arcas BA, Fairhall AL, Bialek W (2003) Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput* 15(8):1715–1749
- Badel L, Lefort S, Brette R, Petersen CC, Gerstner W, Richardson MJ (2008) Dynamic I–v curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *J Neurophysiol* 99:656–666
- Barbour B, Brunel N, Hakim V, Nadal JP (2007) What can we learn from synaptic weight distributions. *Trends Neurosci* 30:622–629
- Brunel N, Latham PE (2003) Firing rate of the noisy quadratic integrate-and-fire neuron. *Neural Comput* 15(10):2281–2306
- Chance FS (2007) Receiver operating characteristic (ROC) analysis for characterizing synaptic efficacy. *J Neurophysiol* 97(2):1799–1808
- Cover TM, Thomas JA (1991) Elements of information theory. Wiley, New York. ISBN 0-471-06259-6
- Douglas RJ, Mahowald M, Martin KA, Stratford KJ (1996) The role of synapses in cortical computation. *J Neurocytol* 25:893–911
- Fetz EE, Gustafsson B (1983) Relation between shapes of post-synaptic potentials and changes in firing probability of cat motoneurons. *J Physiol (Lond)* 341:387–410
- Fourcaud-Trocme N, Hansel D, van Vreeswijk C (2003) How spike generation mechanisms determine the neuronal response to fluctuating inputs. *J Neurosci*
- Gao Y, Kontoyiannis I, Bienenstock E (2008) Estimating the entropy of binary time series: methodology, some theory and a simulation study. *Entropy* 10:71–99. doi:10.3390/entropy-e10020071
- Häusser M, Mel B (2003) Dendrites: bug or feature? *Curr Opin Neurobiol* 13:372–383
- Häusser M, Roth A (1997) Estimating the time course of the excitatory synaptic conductance in neocortical pyramidal cells using a novel voltage jump method. *J Neurosci* 17:7606–7625
- Herrmann A, Gerstner W (2001) Noise and the PSTH response to current transients: I. General theory and application to the integrate-and-fire neuron. *J Comput Neurosci* 11:135–151
- Jolivet R, Rauch A, Lüscher HR, Gerstner W (2006) Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *J Comput Neurosci* 21:35–49
- Kennel M, Shlens J, Abarbanel HD, Chichilnisky EJ (2005) Estimating entropy rates with Bayesian confidence intervals. *Neural Comput* 17:1531–1576
- Keren N (2005) Constraining compartmental models using multiple voltage recordings and genetic algorithms. *J Neurophysiol* 94:3730–3742
- Kistler WM, Gerstner W, Hemmen JL (1997) Reduction of the Hodgkin–Huxley equations to a single-variable threshold model. *Neural Comput* 9:1015–1045
- Larkum ME, Zhu JJ, Sakmann B (1999) A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398:338–341
- Larkum ME, Senn W, Lüscher HR (2004) Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cereb Cortex* 14:1059–1070
- London M, Häusser M (2005) Dendritic computation. *Annu Rev Neurosci* 28:503–532
- London M, Schreiner A, Häusser M, Larkum ME, Segev I (2002) The information efficacy of a synapse. *Nat Neurosci* 5:332–340
- Losonczy A, Makara JK, Magee JC (2008) Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature* 452:436–441
- Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506
- Markram H, Lübke J, Frotscher M, Roth A, Sakmann B (1997) Physiology and anatomy of synaptic connections between thick

- tufted pyramidal neurones in the developing rat neocortex. *J Physiol (Lond)* 500(Pt 2):409–440
- Polsky A, Mel BW, Schiller J (2004) Computational subunits in thin dendrites of pyramidal cells. *Nat Neurosci* 7:621–627
- Rauch A, La Camera G, Luscher HR, Senn W, Fusi S (2003) Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *J Neurophysiol* 90(3):1598–1612
- Reuveni I, Friedman A, Amitai Y, Gutnick MJ (1993) Stepwise repolarization from Ca^{2+} plateaus in neocortical pyramidal cells: evidence for nonhomogeneous distribution of HVA Ca^{2+} channels in dendrites. *J Neurosci* 13:4609–4621
- Schiller J, Schiller Y, Stuart G, Sakmann B (1997) Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *J Physiol (Lond)* 505(Pt 3):605–616
- Stuart GJ, Häusser M (2001) Dendritic coincidence detection of EPSPs and action potentials. *Nat Neurosci* 4:63–71
- Thomson AM, Lamy C (2007) Functional maps of neocortical local circuitry. *Front Neurosci* 1(1):19–42
- Vogels TP, Rajan K, Abbott LF (2005) Neural network dynamics. *Annu Rev Neurosci* 28:357–376
- Willems FMJ, Shtarkov YM, Tjalkens T (1995) The context-tree weighting method: basic properties. *IEEE Trans Inf Theory*, pp 653–664
- Williams SR, Stuart GJ (1999) Mechanisms and consequences of action potential burst firing in rat neocortical pyramidal neurons. *J Physiol (Lond)* 521(Pt 2):467–482