



The meaning of spikes from the neuron's point of view: predictive homeostasis generates the appearance of randomness

Christopher D. Fiorillo*, Jaekyung K. Kim and Su Z. Hong

Department of Bio and Brain Engineering, Korean Advanced Institute of Science and Engineering (KAIST), Daejeon, South Korea

Edited by:

Rava Azeredo Da Silveira, Ecole Normale Supérieure, France

Reviewed by:

Benjamin Torben-Nielsen, The Hebrew University of Jerusalem, Israel

Gianluigi Mongillo, Paris Descartes University, France

*Correspondence:

Christopher D. Fiorillo, Department of Bio and Brain Engineering, Korean Advanced Institute of Science and Engineering (KAIST), 291 Daehak-ro, Building 611, Yuseong-gu, Daejeon 305-701, South Korea
e-mail: fiorillo@kaist.ac.kr

The conventional interpretation of spikes is from the perspective of an external observer with knowledge of a neuron's inputs and outputs who is ignorant of the contents of the "black box" that is the neuron. Here we consider a neuron to be an observer and we interpret spikes from the neuron's perspective. We propose both a descriptive hypothesis based on physics and logic, and a prescriptive hypothesis based on biological optimality. Our descriptive hypothesis is that a neuron's membrane excitability is "known" and the amplitude of a future excitatory postsynaptic conductance (EPSP) is "unknown". Therefore excitability is an expectation of EPSP amplitude and a spike is generated only when EPSP amplitude exceeds its expectation ("prediction error"). Our prescriptive hypothesis is that a diversity of synaptic inputs and voltage-regulated ion channels implement "predictive homeostasis", working to insure that the expectation is accurate. The homeostatic ideal and optimal expectation would be achieved when an EPSP reaches precisely to spike threshold, so that spike output is exquisitely sensitive to small variations in EPSP input. To an external observer who knows neither EPSP amplitude nor membrane excitability, spikes would appear random if the neuron is making accurate predictions. We review experimental evidence that spike probabilities are indeed maintained near an average of 0.5 under natural conditions, and we suggest that the same principles may also explain why synaptic vesicle release appears to be "stochastic". Whereas the present hypothesis accords with principles of efficient coding dating back to Barlow (1961), it contradicts decades of assertions that neural activity is substantially "random" or "noisy". The apparent randomness is by design, and like many other examples of apparent randomness, it corresponds to the ignorance of external macroscopic observers about the detailed inner workings of a microscopic system.

Keywords: Bayesian, Jaynes, inference, prediction, neural code, noise, stochastic, random

INTRODUCTION

Within the field of information theory, there is a common though often implicit belief that the information content of a signal depends only on its statistical relationship to the quantity of interest (e.g., Rieke et al., 1997). In this view, we should be able to understand the meaning of a neuron's spikes entirely through an input-output (I-O) analysis, without any consideration of events happening within the neuron. This view of information is a natural consequence of the opinion that probabilities are essentially equivalent to frequencies and are properties of an observed system, rather than being entirely conditional on the knowledge possessed by an observer about that system (Fiorillo, 2012). The faults and limitations of the "frequentist" definition of probability and information within neuroscience have been discussed previously, together with the virtues of the "Bayesian" definition advocated by Jaynes (Fiorillo, 2012). The present work considers a neuron as an observer

that knows only what is in its internal biophysical state. Like Jaynes (2003), we presume that probabilities are always conditional on the local and subjective information of an observer, through the universal and objective principle of logic (Fiorillo, 2012).

At the center of the present work is a very simple idea that has been around at least since the advent of information theory (Shannon, 1948). If an event has two possible outcomes, observation of the outcome will convey the maximum amount of information if the probability of each is equal prior to observation. The output of most neurons is a spike that is virtually all-or-none (unlike a graded output, this allows for reliable communication across distances), and thus the amount of information that could be conveyed by a spike will be maximal when its probability is 1/2. Although this is not controversial, it has nonetheless been associated with considerable confusion. This is because the information conveyed by any event will necessarily depend entirely on

the prior knowledge of the observer, and a signal that is highly informative to one observer could be entirely uninformative to another. Without a clear understanding of how probabilities and information depend on the observer, scientists have often misattributed their information to the neurons they study (Fiorillo, 2012).

Consider the case of selecting questions that must have “yes” or “no” answers (as in the game “20 Questions”). If there is a real number known to be within the range of 0–10, and the goal is to guess it as accurately as possible after receiving the answer to a yes-no question, then the best strategy would be to bisect the range of possible numbers evenly by asking “is the number greater than 5?” This is the optimal question because “5” is the expected value and will maximize the questioner’s uncertainty (entropy) by causing both “yes” and “no” to have equal probabilities of 1/2. The answer will therefore provide the maximal amount of information by eliminating this uncertainty. Any other question would result in less initial uncertainty about the answer, and therefore the answer would convey less information. However, the answer is informative only if one knows the question. A second observer may also know that the answer will be “yes” or “no”, but not know the question. Thus two observers could agree with one another that the probability of “yes” is 1/2, but the one who knows the question could get maximal information, whereas the one who is entirely ignorant of the question would learn nothing useful from the same answer, and might be tempted to declare that the answer was entirely random noise. Here we suggest that in many instances neuroscientists have essentially taken the perspective of the naive observer, being ignorant of the prior information of the neuron that “asks the question” and generates and interprets the spikes.

It is well known that the spike output of neurons often appears “stochastic”, insofar as the same stimulus input causes a variable spike output. Likewise, when an action potential arrives at a synaptic terminal, it may or may not cause release of neurotransmitter. Although it has been noted that this variability in I-O relations could convey information (Softky and Koch, 1993; Softky, 1995), over decades it has repeatedly been attributed to “noise” or “randomness” that is presumed to degrade the ability of neurons to transmit information (Calvin and Stevens, 1968; Tolhurst et al., 1983; Shadlen and Newsome, 1994, 1998; London et al., 2010). Models of how neurons could perform inference have proposed that the variability may serve a function by signifying subjective uncertainty (Pouget et al., 2000; Deneve et al., 2001; Hoyer and Hyvarinen, 2003; Jazayeri and Movshon, 2006; Ma et al., 2006; Beck et al., 2008; Berkes et al., 2011). Whether the variability is noise or a functional signal related to uncertainty, uncertainty (ignorance) itself is not desirable, and both of these viewpoints propose that a more variable output (in response to a fixed input) would correspond to greater uncertainty. In contrast, a central conclusion of the present work is that this same variability is a signature of optimality and knowledge if we consider the perspective of a neuron, which performs work (consumes energy) to produce it.

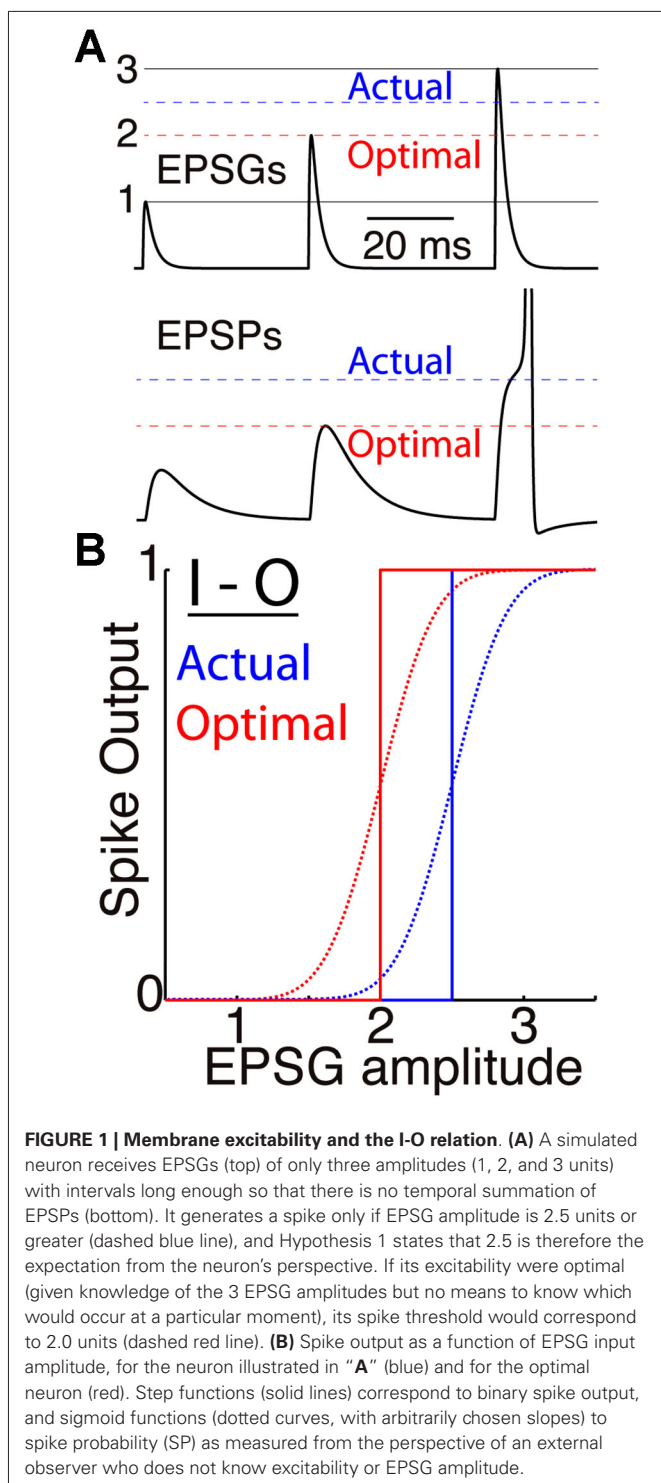
It may seem strange or even unscientific to treat neurons as observers. Indeed, Skinner and other “behaviorists” argued that to be objective and to give psychology a firm scientific basis, we

should not even treat humans and other animals as observers, but that we should instead characterize their behavior solely in terms of I-O relations (environmental inputs and behavioral outputs). Psychology has now mostly rejected that view, and instead seeks to understand mental events and behavior from the first-person perspective of the brain as it observes its world. The probability theory of Jaynes (2003) allows us to describe a subjective state of knowledge in an objective manner (Fiorillo, 2012). Many elegant studies have now demonstrated that aspects of perception, cognition, and motor control that appeared maladaptive can be understood as optimal (rational or Bayesian) if we take the brain’s perspective, including sensory illusions (e.g., Weiss et al., 2002; Niemeier et al., 2003; Yang and Purves, 2003; Körding and Wolpert, 2004).

Illusions would seem to provide clear evidence of the brain’s malfunction. However, the brain only appears to function poorly from the perspective of a third-party observer who knows something that the brain itself does not and could not know. Whereas the brain must infer the sensory stimulus given limited sensory evidence and its own prior knowledge of what is likely, the third party has privileged knowledge of the “true sensory stimulus”. The brain integrates its information as well as it possibly can, but it nonetheless “misperceives” the stimulus on those rare occasions when the stimulus is something unusual and unexpected, like a fast moving object (Weiss et al., 2002). Thus, what once appeared pathological from our perspective as external observers can suddenly be understood as optimal once we consider the unique perspective of the neural observer. We propose that, in a precisely analogous manner, the variability in a neuron’s I-O relationship is optimal rather than pathological once one takes the neuron’s perspective.

INFORMATION THEORY FROM THE NEURON’S PERSPECTIVE

Our theory is closely related to Barlow’s theory of “efficient coding” (Barlow, 1961) and the extensive literature that followed on the relation between neural activity and natural stimulus statistics (e.g., Laughlin, 1981; Srinivasan et al., 1982; Linsker, 1988; Rieke et al., 1997; Stemmler and Koch, 1999; Brenner et al., 2000; Simoncelli and Olshausen, 2001; Hosoya et al., 2005; Sharpee et al., 2006). In particular, the frequency distribution of a neuron’s output (e.g., firing rate) should be “matched” to the distribution of its input intensity, so that if the input distribution is Gaussian, then the I-O relation should ideally follow the form of the corresponding cumulative Gaussian (Laughlin, 1981; Brenner et al., 2000). A neuron’s membrane excitability determines its I-O relationship, and excitability should adapt so that the location (x -offset) and scale (slope) of the I-O curve cause the neuron’s output to be most sensitive to the most probable inputs (the steep and linear portion of the curve) at the expense of improbable inputs (the non-linear and nearly flat portion of the curve). **Figure 1** illustrates this principle as it applies to the cellular level, with the amplitude of excitatory postsynaptic conductance (EPSC) as input and a binary spike or no spike as output. We refer to the principle as “predictive homeostasis”, rather than “predictive” or “efficient” coding, to emphasize that it is implemented by numerous homeostatic mechanisms of the same sort that are well known by biologists to be present in all cells.



Efficient coding by a neuron should ideally have the effect of maximizing the Shannon information in its spike output (I_S) about its excitatory synaptic input (E_{EPSP}). This information would be equal to the reduction in entropy (H) caused by observation of the spike output (S).

$$I_S(E_{EPSP}|E_XS) = H(E_{EPSP}|E_X) - H(E_{EPSP}|E_XS) \quad (1)$$

The “prior entropy” [$H(E_{EPSP}|E_X)$] would be the width or uncertainty of the probability distribution, which would be reduced by observing spike output (S). We equate a neuron's membrane excitability with the prior information (E_X). Unfortunately we are unable at this time to derive the prior probability distribution [$p(E_{EPSP}|E_X)$], as discussed previously (Fiorillo, 2012). However, since spike output (S) is binary, its information content is maximized when its prior probability [$p(S|E_X)$] is 1/2. This corresponds to that case that it indicates whether the input is more or less than the prior expectation ($\langle E_{EPSP} \rangle | E_X$), so that observation of spike output reduces the prior entropy [$H(E_{EPSP}|E_X)$] by 1/2 (assuming the prior distribution is symmetric). We propose below that this is precisely the meaning of a spike from the neuron's point of view.

There are at least three respects in which the present hypothesis goes beyond previous conceptions of this principle. First and of greatest importance, we consider the neuron's point of view, by which we mean that the expectation of interest is entirely conditional on the neuron's prior information as inherent in its membrane excitability (E_X). This contrasts with the conventional approach in which probabilities are conditional on the knowledge of an external observer (e.g., Rieke et al., 1997). Although the importance of “taking the neuron's point of view” has been recognized previously (Rieke et al., 1997), it has virtually never been done for reasons described elsewhere (Fiorillo, 2012). It requires that we set aside whatever knowledge we might have about the neuron's input (such as knowledge of the frequency distribution of its inputs, commonly referred to as “the stimulus ensemble”) and its statistical relationship to the neuron's output. Instead, we consider only the information that we believe to be contained within the physical structure of the neuron or other observer.

A second distinction from most previous work is in considering the millisecond timescale, in which spike output is necessarily binary. Most previous work on efficient coding, including our own (Fiorillo, 2008), considered output to be analog firing rate or membrane voltage (Laughlin, 1981; Linsker, 1988; Stemmler and Koch, 1999; Brenner et al., 2000; but see Deneve, 2008a). However, all the relevant membrane properties evolve on the scale of 1 ms or even less (Softky, 1995), including synaptic conductance and the action potential itself. There are many slower processes as well, some of which effectively average over spikes, and these are also incorporated within the theory (Fiorillo, 2008). But to adequately capture the informational dynamics, and especially the proximal cause of spikes, we must consider the fastest time scale in which membrane properties vary. Our approach rests on the intuitive notion that information is inseparable from matter, energy, and causation, so that if the charge distribution across a neuron's membrane is changing within a millisecond, so is the information that it contains.

Third, like Stemmler and Koch (1999), we focus on a single neuron at the cellular level, so that input and output are both within the neuron. This contrasts with most previous work on efficient coding in which input corresponds to external sensory stimuli, and thus multiple neurons intervene between input and spike output. The cellular approach has several advantages. First, as a simpler system, analysis of a single neuron better allows us

to identify mechanisms by which neurons predict their input and adapt their I-O relation (e.g., Hong et al., 2013). Second, whereas efficient coding is often thought of as a “sensory problem”, the cellular analysis emphasizes the general nature of the problem, since all spiking neurons are faced with the problem of using a binary output to efficiently represent an analog input. Third, the inference made by a neuron or any observer must be based entirely on local information, as must all learning. Compared to a network of neurons, it is relatively easy to imagine how learning or natural selection might optimize I-O relations of single neurons (Fiorillo, 2008). If a single neuron can optimize its local I-O relation, a feedforward series of such neurons, each adapted to the local statistics of its input, can be expected to approximately optimize the systems-level I-O relation (between external sensory stimuli and spike output).

A PREVIOUS MODEL OF PREDICTION BY A SINGLE NEURON

A previous theory had the goal of providing a general informational account of the nervous system (Fiorillo, 2008). Although narrower in scope, the present work builds on the previous framework. In both cases, a neuron’s membrane compares the intensity of its sensory related input to its expectation, and its output signals “prediction error”. As in other models of neuronal prediction error (e.g., Schultz et al., 1997; Izhikevich, 2006; Frémaux et al., 2010), the homeostatic null point was presumed to correspond to “baseline firing rate” or an intermediate membrane voltage “just below” spike threshold (Fiorillo, 2008). The null point and its neural basis was not precisely defined or justified. Indeed, the notion of a “firing rate” is itself problematic, since it is not a physical entity that exists at any particular point in space and time. The present work provides a solution by considering the timescale of a single spike and proposing that the threshold voltage for a spike corresponds to the homeostatic point at which the prediction error is 0.

Information necessarily has both a qualitative and quantitative component, and the earlier theory dealt with both (Fiorillo, 2008). Like Shannon (1948), we deal here only with the quantitative aspect. Since EPSPs are the proximal cause of spikes, the set of excitatory synapses will determine the qualitative nature of the information within spikes by specifying the neuron’s receptive field, or “stimulus” (for example, light of certain wavelengths in some region of space). Previous work suggested how learning could shape a receptive field to insure that it contributes information of biological importance (e.g., Schultz et al., 1997; Izhikevich, 2006; Fiorillo, 2008; Frémaux et al., 2010). Here we address only the issue of how membrane excitability can enable spike output to best preserve (maximize) information about EPSP input (the “coding strategy”), regardless of the meaning and importance of that input (and “what it should code”).

THEORY OF PREDICTIVE HOMEOSTASIS

After specifying a simple biophysical model of membrane excitability, we proceed to describe its information content based only on its physical properties (Hypothesis 1), and then to prescribe excitability based on biological optimality (Hypothesis 2).

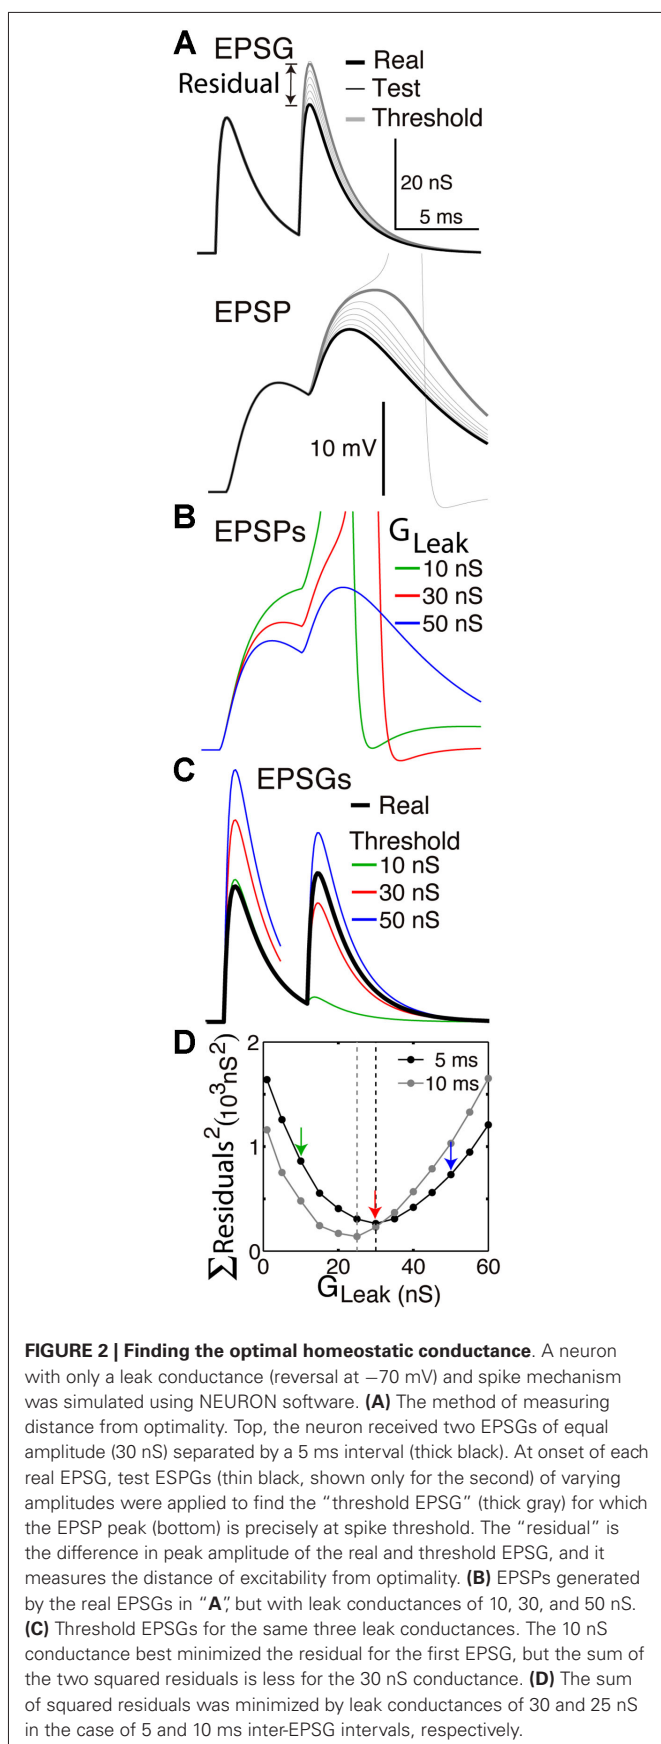
THE BIOPHYSICAL MODEL

We assume a generic neuron that follows established biophysics. It has a spike generating mechanism as well as two conductances, an EPSP and an “intrinsic” conductance. In general, each of these will vary over time, and will be the sum of multiple individual conductances (synapses or types of ion channel). EPSPs are discrete events consisting of one or more unitary EPSPs, where a unitary EPSP results from a spike in a single presynaptic neuron (which releases at least one vesicle of neurotransmitter, and usually more). Unitary EPSPs vary in amplitude over time, even in the case that they derive from the same presynaptic neuron (due to variable numbers of vesicles being released across multiple release sites). In addition, unitary EPSPs and their associated EPSPs will display spatial and temporal summation. Although, they are discrete events and may occur only sparsely over time, they convey information about continuously varying external variables, such as light or sound intensity. We focus on the amplitude of EPSPs because although they are internal to a neuron, their timing and amplitude depend almost entirely on variables external to the cell.

Whether an EPSP arriving at a synapse causes a spike to be initiated near the soma depends on many well known factors, including initial membrane voltage, axial resistance, and membrane capacitance and conductance. Voltage-regulated ion channels change membrane conductance to amplify or suppress an EPSP, even during its brief rise time (~ 1 ms). The summed effect of all of these factors is encapsulated by the term “excitability”. We define it here as the energy barrier (E_X) which separates an EPSP from spike threshold, although to be consistent with common usage we would say that a lower energy barrier (a smaller distance from spike threshold) corresponds to a higher excitability. A spike will occur only if the energy in an EPSP (E_{EPSP}) exceeds the barrier (E_X). Thus excitability is the amount of work that an EPSP must perform to cause a spike, and this corresponds to the amount of charge that must be transferred from the synapse to the site of spike initiation. This could be calculated in principle given a detailed model, but it can be found more easily by injecting “test EPSPs” of incrementally increasing amplitude until spike threshold is reached (**Figure 2A**). The amplitude of the “threshold EPSP” that is found in this way has precisely the energy E_X that it is needed to reach spike threshold.

DESCRIBING A NEURON’S EXCITABILITY AND EXPECTATION

Here we equate “the observer” with “membrane excitability”. The excitability at time t will determine whether an EPSP with onset at time t will cause a spike. What is the expectation of EPSP amplitude given only excitability? Some would answer that this question cannot be answered unless one first observes EPSP amplitudes and thus has knowledge of the frequency of various amplitudes. However, we follow the probability theory of Jaynes (commonly referred to as “Bayesian”), according to which one “sample” of any quantity can provide an expectation of another (through logic, as expressed in the principle of maximum entropy) (Jaynes, 2003; Fiorillo, 2012). From this we presume that knowledge of one mass can be used to estimate another mass, one energy can be used to estimate another energy, etc. More information is always better



for a biological observer, but an observer simply “knows what it knows”.

The EPSP performs work to drive membrane voltage towards spike threshold, and excitability works against the EPSP. In the absence of any information beyond excitability itself, the probability that an EPSP will cause a spike is $1/2$ (based on logic, the maximum entropy principle), and thus the expectation ($\langle E_{EPSP} \rangle$) must be equal to the excitability (E_X) (equation 2) (by the mathematical definition of “expectation”).

Hypothesis 1: A spike is generated only if EPSP amplitude is greater than expected given the information in membrane excitability (equation 2).

$$\langle E_{EPSP,t+r} \rangle | E_{X,t} = E_{X,t} \quad (2)$$

Equation 2 refers to the expectation (at time t) of the peak amplitude of an EPSP if one were to start at that same time (which would reach its peak at time $t + r$, where r refers to the stereotyped rise time of an EPSP, typically about 0.5 ms). Thus, we are referring to an expectation of a potential future event. When excitability is low, the energy barrier (E_X) is high, and a large EPSP is expected. The expectation is just the “center” of a probability distribution, and we are not speculating here about the uncertainty (width) of the distribution. Because spike threshold corresponds to the expectation, a spike conveys maximal information (relative to any other relation between expectation and spike threshold, assuming that the prior probability distribution is symmetrical). A spike would not convey any information beyond that in membrane voltage at the spike initiation site at the time of its generation, but it would reliably communicate that information throughout the entire neuron.

Given an EPSP with onset at time t , it is useful to denote membrane excitability at time t as “prior information” to distinguish it from the new information in the EPSP. Hypothesis 1 is essentially just that a spike signals “prediction error”. Prediction errors are known to be efficient and useful signals, but there is not much intelligence in a prediction error if the prediction itself is not accurate. Our use of “prediction error” is merely descriptive and could be applicable to a large variety of physical entities.

A traditional balance scale provides a useful analogy. It consists of an arm that rotates around a central joint depending on a known reference weight on the left (excitability) and an unknown weight on the right (the EPSP). The arm rotates continuously (over some range) as a function of the difference between the two weights (membrane voltage). The scale generates a binary output to signal which weight is greater (right side up or down, analogous to a spike). Prior to placement of the unknown weight on the right side, the expectation of the unknown weight would be equal to the known weight, and “right side up” (a spike) would indicate that the unknown weight was greater than the expectation.

Our challenge here is merely to describe the information in a physical entity based on physics and logic. Unlike most of biology, physics is purely descriptive insofar as it is agnostic to “how the world should work”. In describing the knowledge within a balance scale or a neuron, we do not presume that the reference weight or the excitability is well suited for measuring the unknown weight

or EPSG. Rather we presume that it is the only information possessed by the observer.

CRITIQUE OF HYPOTHESIS 1

We believe that an observer is something remarkably simple. Excitability is simple insofar as it is just a single number at a single moment in time. We presume that knowledge of excitability does not include knowledge of its constituents, not even membrane voltage and conductance. Furthermore, excitability is entirely internal to a neuron, and therefore it could be said that we are “taking the neuron’s point of view”. However, we also believe that a single observer should correspond to information that is physically integrated in space and time. Membrane voltage fits this definition, summing currents over a local region of membrane. By contrast, excitability includes multiple factors in addition to membrane voltage. Even though those factors are all present in a small and well defined part of a neuron, they are not being physically summed together. Therefore we infer that excitability corresponds to information distributed across more than one observer, and Hypothesis 1 falls short of our ideal of “describing the knowledge of an observer”. It remains for future work to describe the knowledge in membrane voltage about EPSG input.

Despite its limitations, we conclude that Hypothesis 1 represents a useful default hypothesis by virtue of its simplicity. For the expectation of EPSG amplitude given excitability to be other than excitability itself, and the prior probability of a spike to be other than 0.5, would require additional information that we have not assumed. We assume that our observer has knowledge, but it is an absolute minimal level of knowledge. The prescriptive hypothesis presented below does not depend upon the veracity of this descriptive hypothesis. However, the simple logic of this descriptive hypothesis suggests that a neuron does not need any particular design in order for it to have an expectation, and for its spikes to be maximally informative from its perspective. Hypothesis 1 states that a known is the best guess of an unknown, whereas Hypothesis 2 states that biology works towards the ideal in which the known is in fact equal to the unknown.

PRESCRIBING A NEURONS EXCITABILITY AND EXPECTATION

Hypothesis 2 expresses our belief that the excitability of a healthy adult neuron will tend to accurately predict EPSG amplitude, because it has been shaped by both natural selection over generations and associative learning rules. An accurate analog expectation of an analog variable will be too high in half of cases, and too low in the other half.

Hypothesis 2: Expectations will be accurate under natural conditions, and therefore half of EPSGs will cause spikes.

The perfect match of EPSG amplitude to its expectation is not attainable, and even if it were, the spike output must be 0 or 1, corresponding to a negative or positive prediction error. The optimal expectation would therefore result in a spike in response to precisely half of EPSGs (Figure 1). The spike probability (SP) of 1/2 given the neuron’s prior information (Equation 2) would then match the frequency distribution of spikes given EPSGs. Although this optimal neuron would have low subjective uncertainty (which

we do not attempt to quantify here), its spikes would appear to be random to an external observer who is ignorant of variations in excitability and EPSG amplitude.

Hypothesis 1 proposes that spikes are maximally informative simply as a consequence of logic and physics, regardless of the prior information in excitability. Hypothesis 2 implies that accurate prediction (or minimizing total entropy) is an important biological goal, and it will be aided by accurate prior information (a low prior entropy, $[H(E_{EPSG}|E_X)]$). As a contrived example, neuron “A” has prior knowledge that EPSG amplitude must be between 0 and 10, and neuron “B” knows that it must be between 4 and 6. If the expectation of each neuron is 5, then a spike would tell neuron “A” that the actual amplitude must be between 5 and 10, whereas it would indicate to neuron “B” that it is between 5 and 6. Neuron “B” would have more total information given its prior information and a spike [less “posterior” entropy, $H(E_{EPSG}|E_XS)$].

MEASURING THE ACCURACY OF A NEURON’S PREDICTIONS

Here we propose methods by which we can measure the accuracy of a neuron’s expectation. The neuron itself also needs to assess the accuracy of its predictions, but we address this in a later section (Prediction Errors in Learning). We use “accuracy” to refer to a distance between two expectations, the more certain of which could be called “reality” or the “actual” or “true” value, and the less certain “the expectation”. Thus “accuracy” is not the same as uncertainty (entropy). We have not attempted to quantify uncertainty, but greater uncertainty would naturally tend to be associated with less accuracy.

RESIDUALS

The best method for us to assess the accuracy of expectations will depend on our ability to know and control the neuron. If we have complete knowledge and control, as in a Hodgkin-Huxley model, then it is easy to measure excitability by injecting “test EPSGs” of varying amplitude to find the “threshold EPSG” amplitude that causes an EPSP to reach precisely to spike threshold (Figure 2A). This “threshold EPSG” is our measure of excitability (E_X) and corresponds to the neuron’s expectation of EPSG amplitude (Hypothesis 1). The “residual” EPSG amplitude (E_{res}) measures the difference between the neuron’s expectation (E_X) and the actual EPSG (E_{EPSG}) (Figure 2A).

$$E_{res} = E_{EPSG} - E_X \quad (3)$$

Indices for time have been omitted for simplicity. In principle, excitability could be measured at each instant (“real time”). However, it is not an expectation of EPSG likelihood, but of EPSG amplitude if one were to occur. Thus the expectation is only put to the test when an EPSG occurs. A high level of inhibitory conductance would correspond to expectation of a large EPSG, but if none occurs, the cost would be only metabolic and not informational. Thus we measure excitability only at the time of real EPSGs (onset of test and real EPSGs are always synchronous) (Figure 2A).

A residual (E_{res}) will be positive if the EPSG causes a spike, and will indicate the extent to which the EPSG was greater than

expected and excitability was too high (E_X was too low). It will be negative if EPSP amplitude was less than expected, indicating that excitability was too low (E_X was too high). If a neuron is making accurate predictions, the residuals will be small with an average of 0, and spike generation will be highly sensitive to small variations in EPSP amplitude, as experiments have demonstrated in some cases (Carandini, 2004; Kuenzel et al., 2011).

The residual is an error of sorts, and an objective of learning would be to modify homeostatic conductances so that residuals are minimized. However, we are not suggesting that a neuron has any direct means of knowing the residuals, and we reserve the term “prediction error” to describe errors from the neuron’s point of view (as signaled by spikes). Thus the residuals cannot drive learning, and we present them only as a means for us to assess the accuracy of a neuron’s predictions (its distance from optimality) (Figure 2).

SPIKE PROBABILITY FROM THE EXPERIMENTER’S POINT OF VIEW

Hypothesis 2 states that predictions should be accurate under natural conditions in a normally functioning brain. Even with *in vivo* intracellular recording (Figure 3) it is challenging to estimate EPSP amplitudes and excitability. It is therefore useful to assess the accuracy of a neuron’s expectations given only knowledge of the time of EPSP inputs (or presynaptic spikes) and spike outputs. For this purpose we define *spike probability* (SP):

Spike probability (SP): The probability that at least one output spike is caused by an EPSP at a “driving” synapse.

SP is defined to incorporate knowledge that at least one input spike and associated EPSP has occurred at a “driving” synapse (see Driving EPSPs as the Cause of Spikes), and that if the EPSP causes a spike, it will occur within a certain delay (generally less than 5 ms in a typical neuron). A unitary EPSP could be the sole cause of a spike or a partial cause together with other EPSPs that are nearly synchronous. SP does not include knowledge of excitability or the specific amplitude of the EPSP (since that information is seldom available from experimental data). Thus, our uncertainty about whether or not a spike will occur is not because there is anything intrinsically “random” about the occurrence of spikes, but because we are ignorant of the “black box” that is EPSP amplitude and membrane excitability.

SP of 1/2 is the homeostatic ideal and will cause spike output to be maximally sensitive to small gradations in EPSP amplitude (Figure 1). However, we presume that it is a virtually impossible goal to achieve, for the same reason that predictions of the future market value of stocks are almost never precisely accurate. We expect that SP varies dynamically from near 0 to 1 under natural conditions, but with an average over time that is near 1/2.

IMPLEMENTATION OF PREDICTIVE HOMEOSTASIS

DRIVING EPSPs AS THE CAUSE OF SPIKES

We presume that information follows causation, and that the EPSP is (and should be) the proximal cause of spikes. The neuron’s knowledge is substantially limited to its proximal causes, for the same reason that our conscious perception corresponds to proximal sensory causes in the immediate past, which we


can infer with relatively little uncertainty, but not the innumerable distal causes further in the past. In a typical neuron the proximal cause of a spike would be an EPSP with onset about 0.5–5 ms earlier, mediated primarily by AMPA-type glutamate receptors (Figure 3). We propose that slow membrane depolarization would not be the proximal cause of spikes under natural conditions, including that initiated by offset of synaptic inhibition, activation of G-protein-coupled receptors, and perhaps even AMPA-type glutamate receptors located on distal dendrites (Branco and Häusser, 2011). Excitatory synapses have been classified as either “drivers” or “modulators”, which are distinguished from one another by dendritic location and a variety of other features (Sherman and Guillery, 1998). “Drive” means “cause” in this case, and the drivers alone cause spikes and determine a neuron’s receptive field. In the lateral geniculate nucleus of the thalamus (LGN), “feedforward” retinogeniculate AMPA EPSPs on proximal dendrites are the drivers of spikes, whereas “feedback” corticogeniculate AMPA EPSPs on distal dendrites are “modulatory” (Sherman and Guillery, 1998). The latter could provide “value information” that modulates the strength of driving synapses and thereby changes the shape of a neuron’s functional receptive field (Fiorillo, 2008), but we do not consider its particular function here. “The EPSP” to which we refer is exclusively from “driving synapses”. Slow excitation is naturally more predictable, and predictive homeostasis should work to insure that only fast and large EPSPs cause spikes.

MECHANISMS AND TIMESCALES

We propose that there are a vast array of homeostatic ion channels and synapses functioning simultaneously on numerous timescales within a single neuron. Each would make a prediction based on a particular recurring pattern of synaptic excitation or inhibition, like those in Figure 3. Homeostatic synaptic inhibition would be dedicated to spatial patterns of activity across neurons (as in the case of “surround inhibition”), whereas voltage-regulated ion channels would be dedicated to temporal patterns (Fiorillo, 2008). To illustrate and test this idea, we need to relate specific homeostatic mechanisms to specific patterns of synaptic drive.

Figure 1 illustrates the case of a neuron that receives EPSPs of only three amplitudes (1, 2, and 3) that are sufficiently separated in time so that there is no temporal summation. Knowing those three amplitudes but nothing else, the optimal expectation is “2” and an EPSP of amplitude “2” should cause an EPSP that reaches precisely to spike threshold (Figure 1). Figure 2 illustrates the case of a neuron that receives excitatory drive from a single presynaptic neuron, with each presynaptic spike causing a unitary EPSP of identical amplitude. Despite the simplification of identical unitary amplitudes, this example remains interesting since temporal summation will cause EPSP and EPSP amplitude to vary with presynaptic inter-spike interval (ISI). This is in fact similar to some real synapses, as discussed further below.

The neuron of Figure 2 is the simplest possible insofar as its only homeostatic conductance is a “leak” conductance, which is voltage-insensitive and constant by definition. If the leak conductance is optimal, it will accurately predict the long-term average EPSP and EPSP amplitude, which vary due to temporal summation. Thus half of unitary EPSPs will exceed average



amplitude and cause a spike, and this will only occur when there is sufficient temporal summation (Figures 2A, B). When the neuron’s excitability is “at rest”, the first EPSP in a pair will be of less than expected amplitude and cause no spike, but it will increase excitability so that a second EPSP 5 ms later will cause a spike. We note that even when EPSPs cease and excitability approaches a stable resting level, excitability still provides an expectation of EPSP amplitude, and this could be called the neuron’s “default prior”. Since the amplitude of EPSPs provides information about the intensity of external stimuli, like light intensity, the neuron would always have an expectation of both EPSPs and external stimuli.

Although accurate prediction of the long-term average is important, it would generally be highly inaccurate over shorter periods. Prediction using only a very slow ion channel, like a leak conductance, is analogous to predicting the market value of a stock using only knowledge of its trajectory over the last 10 or 100 years. Better predictions can be made given additional knowledge of any regular pattern, including short term trends. There are a remarkable diversity of voltage-regulated ion channels with a spectrum of kinetic properties, and each of these is proposed to be specialized for making a prediction based on a particular recurring pattern (Fiorillo, 2008). There are obviously many patterns that involve changes in EPSP rate. An increased rate of EPSPs and spikes would activate K⁺ channels with moderately slow kinetics (such as K_v 7), which would then mediate “spike frequency adaptation” and restore SP towards 0.5 over periods of tens of milliseconds or more.

The rising phases of EPSPs and action potentials are the fastest changes in excitatory drive. Excitability will provide an expectation of an EPSP’s amplitude prior to its onset, but A-type K⁺ channels activate during the rising phase of an EPSP and thus contribute to excitability as well (in many types of neurons). In principle, the amplitude of an EPSP can be predicted by the slope of its rising phase. If some subtypes of A-type channels inactivate within a millisecond or two (Migliore et al., 1999), they could conceivably exploit this pattern by preferentially suppressing EPSPs with fast rising phases that would otherwise be excessively large (suprathreshold), while inactivating and thereby causing less suppression of slowing rising EPSPs that would otherwise be too small (subthreshold). The amplitude of an individual EPSP could also be predicted given knowledge of the amplitude of a related EPSP (sharing the same cause) in another neuron. This prediction would be mediated by homeostatic synaptic inhibition via GABA_A chloride channels, which typically has an onset about 1 ms after EPSP onset (e.g., Blitz and Regehr, 2005), and derives

from a neuron with a similar and overlapping set of excitatory drivers (a similar receptive field) (e.g., Wang et al., 2011). GABA_A and A-type channels both provide fast homeostatic mechanisms that may “update expectations” within the brief period between EPSP onset and EPSP peak. In addition, it is likely that during a large EPSP (with or without a spike), these channels open and thereby decrease excitability in anticipation of an additional EPSP, since otherwise the sum of the two would greatly exceed spike threshold (Figures 2A–C). It should be noted in this respect that homeostatic mechanisms should be finely tuned not only to excitatory drive but to other homeostatic mechanisms as well. For example, the inactivation of A-type K⁺ channels could conceivably implement a predictive compensation for GABA_A inhibition, since simultaneous activation of both conductances could cause an excessive decrease in excitability.

FINDING THE OPTIMAL HOMEOSTATIC CONDUCTANCE

Figure 2 illustrates how we can find the optimal leak conductance for a neuron that has no other homeostatic conductances and receives two EPSPs of equal amplitude that are separated by a 5 ms interval. Because of temporal summation (Figures 2A, B), it is not possible with only a leak conductance to have two EPSPs with both peaks at spike threshold. The best that can be done is for the first to be too small and the second too large (Figure 2C). We presume that the optimal leak conductance is the one that minimizes the sum of the two squared residuals (Figure 2D). If the EPSP interval is 10 rather than 5 ms, there is less temporal summation and thus the optimal leak conductance is smaller (Figure 2D). The optimal conductance in a real neuron would depend on the entire ensemble of EPSP amplitudes that the neuron experiences, and a conductance that is optimal for one pattern will often be suboptimal for another. It is apparent from Figure 2 that if inhibition were increased by an appropriate amount following the first EPSP it would help to restore excitability and thereby make the neuron better prepared for addition of the second EPSP. A neuron with a suitable dynamic homeostatic conductance (like GABA_A or A-type channels) would therefore be closer to optimal than any neuron with only a leak conductance.

TWO CLASSES OF ION CHANNEL

Our earlier work distinguished two classes of ion channel (Fiorillo, 2008). Class 1 ion channels were proposed to contribute “current” sensory related information, which generally corresponds to the EPSP, whereas class 2 inputs contribute “prior information” that is “the prediction” in “prediction error”. The defining and categorical distinction was that class 1 ion channels

(and synapses) have been selected through evolution or Hebbian learning to drive membrane voltage (or firing rate) away from the homeostatic level where the error is 0, whereas class 2 ion channels have been selected (through evolution or anti-Hebbian learning) to drive voltage towards its homeostatic level. The distinction is categorical (rather than merely denoting statistical tendencies) since implementation of one learning rule rather than its opposite would require that ion channel proteins are somehow “tagged” in a way that would determine whether the channel is to be regulated by either a Hebbian or anti-Hebbian rule (Fiorillo, 2008). The proposal that voltage-regulated ion channels can be regulated by associative learning rules remains unverified.

The present work provides further insight by specifying that the homeostatic level is spike threshold. GABA_A-mediated IPSPs can be class 1 or 2, as illustrated by the two types of GABA synapse on thalamocortical neurons of LGN, both of which derive from local interneurons (Blitz and Regehr, 2005; Wang et al., 2007, 2011). Retinogeniculate excitation of an ON-type neuron provides evidence for high light intensity (or positive contrast), whereas excitation has the opposite meaning in an OFF-type neuron. In an ON-type neuron, inhibition from an OFF-type interneuron (with an almost identical receptive field) contributes evidence against high light intensity (and against the EPSP), as shown in **Figure 3** (Wang et al., 2007). Sometimes there will be evidence both for and against light in the receptive field at the same time, but generally these will be anti-correlated (Mastrorarde, 1989). Thus OFF-type inhibition will usually occur in the absence of EPSPs and drive hyperpolarization away from spike threshold (**Figure 3**). However, an ON-type neuron also receives synaptic inhibition from an ON-type interneuron (with a similar receptive field). This “same-sign” inhibition naturally occurs at nearly the same time as retinogeniculate excitation (typically delayed by 1 ms), and is caused in a feedforward manner by the same retinal ganglion neuron (Blitz and Regehr, 2005), and thus it is homeostatic and prevents excessive excitability. We did not incorporate opponent inhibition into our framework above, but we could do so by proposing that homeostatic mechanisms should predict the sum of new evidence for and against a neuron’s stimulus, and a neuron generates a spike when new evidence for its stimulus (an EPSP) exceeds both new evidence against it (opponent inhibition) as well as old evidence for it (homeostatic inhibition).

Sodium channels may also be class 1 or 2. Axonal action potentials are the “all-or-none” output of a neuron, and the sodium channels that cause them clearly drive voltage away from homeostasis and would be class 1. By contrast, it is generally believed that dendritic sodium channels do not consistently contribute to causing axonal action potentials (axonal spikes are not initiated in dendrites). If they usually amplify dendritic EPSPs towards threshold for an axonal spike under natural conditions, they would be homeostatic, whereas if they usually drive depolarization beyond threshold for an axonal spike, they would oppose homeostasis and be in class 1. It is interesting to note that class 1 and class 2 sodium channels could have identical structure at the level of individual channels, but their opposite relationship to

homeostasis could be due solely to their sub-cellular localization and density.

SENSORY VERSUS MOTOR NEURONS

The principle of predictive homeostasis should apply to all neurons, and even all cells. However, we suspect that there is an important difference between sensory and motor neurons, and that above we have described sensory neurons. We do not try to justify the distinction here, but a rough and intuitive summary is that sensory neurons use a binary spike output to efficiently convey information about “shades of gray” in their sensory input (EPSP amplitude), whereas motor neurons use it to convert a shade of gray into a “black or white decision”. As partially summarized elsewhere (Fiorillo, 2013a), neurons in motor systems are proposed to express voltage-regulated ion channels (such as L- and T-type calcium channels) that drive voltage away from spike threshold (class 1 ion channels). This is known to be the case in some motor neurons (Marder and Bucher, 2007), but we propose it to be a general feature of motor neurons. In the case that these channels are excitatory, they would tend to be coactive with EPSPs, and together with EPSPs they would cause spikes. The goal of predictive homeostasis (class 2 ion channels) in motor neurons would be to maintain excitability by predicting the sum of EPSPs and these pattern-amplifying intrinsic conductances. Our definition of SP does not incorporate knowledge of the state of these intrinsic “drivers”, and we suspect that they would have the effect of driving SP towards the extremes (0 or 1) so that it is rarely near 0.5. Nonetheless, effective homeostatic mechanisms should still cause long-term average SP to be near 0.5. Furthermore, since EPSPs would not be the sole cause of spikes in motor neurons, spike count may outnumber EPSP count during some excitatory events.

IMPLICATIONS OF PREDICTIVE HOMEOSTASIS

METABOLIC COST OF PREDICTIVE HOMEOSTASIS AND SPIKES

The importance of predictive homeostasis is evident in the large amount of energy that it consumes. Sodium, potassium, and chloride conductances all tend to increase in synchrony, allowing these ions to flow down their electrochemical gradients and counteract one another electrically while causing only modest changes in membrane voltage (Pouille and Scanziani, 2001; Shu et al., 2003; Wehr and Zador, 2003; Blitz and Regehr, 2005; Hosoya et al., 2005; Wilent and Contreras, 2005; Berg et al., 2007; Okun and Lampl, 2008; Wang et al., 2011; Sengupta et al., 2013). The high cost that neurons pay demonstrates the biological value of the information that they gain from predictive homeostasis.

In proposing that SP of 0.5 is optimal, we are ignoring the metabolic cost of action potentials (Attwell and Laughlin, 2001). If their cost were high relative to the value of the information they convey, the optimal SP should be lower than 0.5 to conserve energy. Although we do not consider metabolic costs here, we do review experimental evidence below that SP is maintained moderately close to 0.5, consistent with the possibility that metabolic costs do not have a large impact on SP.

THE MEANING OF SPIKE OUTPUT TO SUB-CELLULAR OBSERVERS

A spike exerts its effect on sub-cellular observers (e.g., proteins) throughout the neuron. A spike can be approximated as a digital or binary signal, but it acts on virtually all neural observers as an analog signal, perhaps most notably in its effects on calcium concentration. If we were to analyze what these observers know about EPSP amplitude and what information is provided by the spike, we would follow the same framework given above. By predicting some local quantity like calcium concentration, they would also be predicting the more distal EPSP amplitude, which in turn corresponds to information about more distal quantities including a stimulus external to the nervous system.

Will the optimal “encoding strategy” of our neuron benefit these other observers, even though they are ignorant of the membrane excitability? The answer is “yes” because the more effective the predictive homeostasis, the less patterned the spike output, which would be “temporally decorrelated” relative to synaptic input (e.g., Wang et al., 2003; Deneve, 2008a). A spiking neuron with any other encoding strategy, or which makes inaccurate predictions, would have a less variable and more predictable output and thus convey less information from the perspective of most observers. This also applies to an observer who counts spikes over a longer period of time (a “rate code”). The fact that the firing rate of neurons is “analog” and not merely binary depends on the same sorts of homeostatic mechanisms, such as the potassium channels that are activated following a spike and create the “relative refractory period”. Without a relative refractory period, a neuron’s firing rate would be either 0 or a fixed number of spikes in response to any input (the number depending only on the absolute refractory period).

The benefit of maximizing the variance of input or output can be a source of confusion. It may seem that if the goal of an observer is to minimize its uncertainty (posterior entropy), it would be best to select inputs that have minimum variance and are thus more predictable. This issue has been called “the dark room problem”, since the uncertainty of neurons in the visual system would appear to be minimized if the animal simply avoids light (Friston et al., 2012). The reason that unpredictable inputs are better, other things being equal, is that the uncertainty that ultimately matters is not about EPSP amplitude, or the amount of light that enters the eye, or even about the identity of distant objects, but about future value. This explains the fact that animals explore uncertain environments, and the utility of Hebbian learning, which maximizes positive prediction errors (Fiorillo, 2008).

We favor the view that the flow of information follows the flow of force and energy and causation, implying that the absence of one event (a spike) is never the cause of another event, and therefore conveys no information. However, in the context of a force and energy that creates an expectation (e.g., a potassium current), the absence of an event (e.g., an EPSP) could allow the expectation itself to modify the observer’s knowledge (e.g., hyperpolarize the membrane). The occurrence of an EPSP would be known to some dendritic proteins, and the EPSP would cause the expectation of a spike. The absence of a spike in the context of an EPSP would indicate that the EPSP was smaller than expected. By contrast, an axon terminal (or its postsynaptic targets) does

not experience an EPSP. Such an observer would have no means to distinguish whether the absence of a spike corresponds to no EPSP or an EPSP of less than expected amplitude, and thus absence of a spike would convey little or no information.

OPPONENT (ON–OFF) REPRESENTATIONS

If the absence of a spike conveys little or no information to an axon terminal and its postsynaptic targets, and yet a system is to be sensitive to decrements as well as increments, there must be “opponent” neurons that represent the same stimulus dimension (like light intensity) but with opposite polarity. The ON and OFF neurons in the visual system are a well known example (Wang et al., 2011), and it is likely that the neural representation of reward value follows the same principle (Fiorillo et al., 2013; Fiorillo, 2013b). The value of opponent representations has long been recognized on the grounds that firing rates are highly “rectified”, but single spikes are “fully rectified”. Thus a spike in an ON neuron would be evidence of “high”, a spike in an OFF neuron evidence of “low”, and absence of a spike is not evidence at all.

PREDICTION ERRORS IN LEARNING

A spike functions as a “teaching signal”, and what the “student observer” learns will depend on whether or not that observer knows that an EPSP has occurred. Observers within a dendritic spine can and should have information about the occurrence and timing of three events: a local EPSP at that synapse, a global EPSP, and a spike. Given this information and the Hebbian objective at an excitatory synapse of maximizing positive prediction errors (spikes) (Fiorillo, 2008), if an excitatory synapse was active just prior to a spike, it contributed to a positive error and should be strengthened. If a local EPSP at that synapse failed to cause a spike, then it contributed to a negative error (or alternatively, we can say that it was not associated with a positive error), and thus it should be weakened. If the local synapse was inactive, then it bears no responsibility for a spike or its absence, and its strength should not be changed substantially. This proposed rule is spike-timing dependent and shares an essential feature of the rule proposed by Bienenstock, Cooper, and Munro (the “BCM rule”), in which synaptic activation paired with only modest postsynaptic activity causes depression of synaptic strength (Bienenstock et al., 1982). The BCM rule has useful computational properties and is supported by experimental evidence (Cooper and Bear, 2012). The BCM-like rule proposed here would allow a neuron to selectively strengthen those excitatory synapses that are least predictable and thus provide the neuron with the most new information (Fiorillo, 2008).

Whereas observers in the dendrite can gain information from a “spike failure” (EPSP + no spike), downstream observers in the axon terminal or its postsynaptic target neuron would not. For example, midbrain dopamine neurons signal a “reward prediction error” that is thought to teach reward value to downstream neurons (e.g., Schultz et al., 1997; Izhikevich, 2006; Fiorillo, 2008; Frémaux et al., 2010). Whereas past models used both increments and decrements in firing rate (positive and negative prediction errors relative to “spontaneous, baseline firing rate”), here we conclude that teaching should be performed by spikes alone, not their absence. Indeed, recent experiments have shown that decrements

in firing of dopamine neurons are not useful for learning about “negative reward value” (Fiorillo et al., 2013; Fiorillo, 2013b). Rather, spikes from an opponent (“reward-OFF”) neuron would be essential for “bidirectional” learning.

COMPARISON TO OTHER INTERPRETATIONS OF SPIKE VARIABILITY

Whereas virtually all previous models have proposed that spike variability is indicative of greater uncertainty of neurons about their sensory inputs, we propose the opposite. Our optimal neuron fulfills Softky’s definition of an “efficient neuron”, having membrane excitability that is exquisitely fine-tuned to its synaptic input (Softky, 1995). Our ignorance of the neuron’s excitability (its internal knowledge) causes output that is maximally informative from the neuron’s perspective to appear random to us (Softky, 1995; Fiorillo, 2012).

The characterization of spike variability depends on probabilities, and we have previously criticized the field in general for failing to adequately distinguish conditional probabilities from event frequencies, and for failing to take the neuron’s point of view (Fiorillo, 2012). Spike variability is conventionally interpreted as “noise” within the literature on “efficient coding” (EC; Barlow, 1961; Rieke et al., 1997). Although we have come to the opposite conclusion, our interpretation is a straightforward application of information theory and it is generally in the spirit of EC.

We view our work, and the EC literature in general, as very distinct from a viewpoint that might be called “probabilistic computation” (PC), which is characterized by the assumption that neurons need to “represent and compute probabilities”. We agree with this viewpoint with respect to the general premise that the function of neurons is related to “Bayesian inference”, but we have a very different view of what this means (Fiorillo, 2012). First, opposite to our conclusion, a popular view within PC is that spike variability (in response to fixed inputs) increases with the subjective uncertainty of neurons (Pouget et al., 2000; Deneve et al., 2001; Hoyer and Hyvarinen, 2003; Jazayeri and Movshon, 2006; Ma et al., 2006; Beck et al., 2008; Berkes et al., 2011). Second, PC presumes that inference requires probabilities to be “calculated” by the brain, and thus neurons may perform methods like Expectation-Maximization (presuming that exact integration in calculus may be beyond the brain’s abilities) (Deneve, 2008a,b; Friston, 2010; Clark, 2013). Equations that describe reason, like Bayes’s Theorem, are viewed as prescriptive of what the brain should do. This is in sharp contrast to our view, and that of the EC literature and physics in general, in which probabilities are merely descriptive. We presume that, at least in principle, probabilities can be used by scientists to describe information that is physically integrated in both space and time (as in a neuron’s membrane voltage). This is precisely analogous to the way that calculus describes motion in physics. A moving object does not “perform calculus”, and the methods of calculus do not prescribe the motion of an object. The descriptive view is that it is sufficient to have the right information in the right place at the right time, without any additional need for “performing mathematics”. Biophysical models are purely descriptive, but PC and other theories that prescribe Bayesian principles for the brain have generally not addressed the fundamental issue of biophysical description (Fiorillo, 2012).

A particularly comparable example of PC is Deneve’s “Bayesian spiking neuron”, which signals prediction error (Deneve, 2008a,b). This model presumes that probabilities are a property of the neuron’s external stimulus, and that they are synaptically conveyed to a neuron. Thus a neuron’s membrane voltage and spike output represents probabilities (frequencies) of the stimulus, whereas our neuron represents the stimulus itself (its physical intensity). Beyond this fundamental difference, Deneve’s optimal neuron produces identical output in response to repeated presentations of identical input, whereas ours produces variable output (it would be impossible for any neuron to know if repeated inputs are identical, and we doubt that they ever can be identical in reality). The presumption of Deneve was that the null point, where the error is 0, should lie some distance below spike threshold (specified by the free and constant parameter “ g_0 ” in the model of Deneve). Thus if spike threshold is -50 mV, a perfect expectation might result in a membrane voltage of -55 mV, and would therefore never cause a spike in Deneve’s neuron. In our neuron, a perfect expectation would cause membrane voltage to be at spike threshold, thereby generating maximally variable output (SP = 0.5, similar to Deneve’s model if g_0 is 0).

EXPERIMENTAL EVIDENCE FOR PREDICTIVE HOMEOSTASIS

The theory predicts that homeostatic conductances should counterbalance EPSPs to maintain excitability, and that this should happen on all timescales, including the millisecond scale. The predictions of the theory can be made quantitatively precise given a particular EPSP pattern (Figure 2), as discussed above. In support of the theory, a remarkably precise balance of excitation and inhibition has been found in brain slices (Pouille and Scanziani, 2001; Shu et al., 2003; Blitz and Regehr, 2005) as well as *in vivo* in both sensory neurons (Wehr and Zador, 2003; Hosoya et al., 2005; Wilentz and Contreras, 2005; Higley and Contreras, 2006; Okun and Lampl, 2008) and spinal motoneurons (Berg et al., 2007), even on a millisecond timescale (e.g., Wehr and Zador, 2003). Since this phenomenon is moderately well known we do not review it here, but instead focus on I-O relations.

TEMPORAL DECORRELATION AND SPIKE STATISTICS

If a neuron makes accurate predictions and signals prediction error, its spike output should have less temporal pattern than its input, since the predicted patterns have been removed from the output. Multiple studies have demonstrated “temporal decorrelation” within sensory systems (e.g., Dan et al., 1996; Brenner et al., 2000), including at the level of single neurons (e.g., Wang et al., 2003, 2010b). Many additional studies have described spike statistics without such analyses of I-O relations. The present theory predicts that the more accurate a neuron’s predictions, the more random its spikes will appear (SP will be closer to 0.5). Thus neurons presented with dynamically varying inputs or complex patterns that are inherently more difficult to predict should have less random output than neurons exposed to more static conditions or simpler patterns. It is well known that many neurons have a relatively random, “Poisson-like” pattern of spike output under “resting conditions” (Tolhurst et al., 1983; Softky and Koch, 1993; Shadlen and Newsome, 1998; Churchland et al., 2010), which are relatively static by definition. By contrast, stimuli

that are difficult to predict due to rapid temporal dynamics cause spike counts and intervals to be extremely regular across repeated trials (Mainen and Sejnowski, 1995; Bair and Koch, 1996; Berry et al., 1997; de Ruyter van Steveninck et al., 1997; DeWeese et al., 2003; Avissar et al., 2007; Butts et al., 2007). When the same stimuli are constant or more slowly varying, and thus more predictable, the same neurons show much greater variability in their spike statistics (Mainen and Sejnowski, 1995; de Ruyter van Steveninck et al., 1997; Butts et al., 2007).

THALAMOCORTICAL NEURONS IN LGN

Thalamocortical neurons in LGN offer several advantages for testing the theory. First, as part of the early visual system, they have been studied extensively both *in vivo* and *in vitro*, with respect to information theory as well as mechanisms. Second, they have a spiking input and output, like most neurons in the brain but distinct from some early sensory neurons such those in the retina. Third, they have one dominant presynaptic excitatory “driver” that generates large unitary retinogeniculate EPSPs (Usrey et al., 1998; **Figure 3**). This provides an important technical advantage, since the large size of unitary EPSPs allows them to be recorded simultaneously with spikes in a single neuron with a single extracellular electrode (Kaplan and Shapley, 1984). It also narrows the range of possible EPSP patterns, since there is little spatial summation. Furthermore, each unitary EPSP is of similar amplitude *in vivo* in LGN (Carandini et al., 2007; Wang et al., 2007), as well as in the case of other powerful synapses (Lorteije et al., 2009; Borst, 2010; Kuenzel et al., 2011). Thus we can approximate the input pattern by considering that it consists only of unitary EPSPs of equal amplitude, so the total synaptic drive is determined solely by temporal summation (**Figure 2**).

Table 1 and **Figure 4** show output and input rates from studies in LGN and elsewhere in which both were recorded simultaneously. Most studies reported only averages of SP over periods of seconds or minutes. We sometimes report averages of SP across populations of cells, but this is done only to provide a concise summary, or because greater detail was not reported in some studies.

Retinal spikes are the cause of all spikes in LGN (Kaplan and Shapley, 1984), and it is well established that average SP is near 0.5 at the retinogeniculate synapse *in vivo* (Kaplan et al., 1987; Movshon et al., 2005; Carandini et al., 2007; Casti et al., 2008), including when naturalistic stimuli are presented (Sincich et al., 2007) and in the awake state (Weyand, 2007). Casual inspection of whole-cell voltage recordings suggests that average SP has intermediate values even over periods of about a second (**Figure 3**; Wang et al., 2007). At least 95% of all spikes in LGN neurons were caused by retinogeniculate EPSPs (Sincich et al., 2007; Weyand, 2007). When an EPSP followed a long period without one, it rarely caused a spike ($SP < 5\%$). However, following a first EPSP, SP at the time of a second rose to a population average of about 85% at 2 ms, and then declined back to its low “resting level” over 40 ms (**Figure 5**; Sincich et al., 2007). Weyand (2007) found similar results. There is strong evidence that SP over this time scale in LGN is influenced by fast homeostatic mechanisms (Blitz and Regehr, 2005), but in principle these data can be explained by temporal summation alone (Carandini et al., 2007;

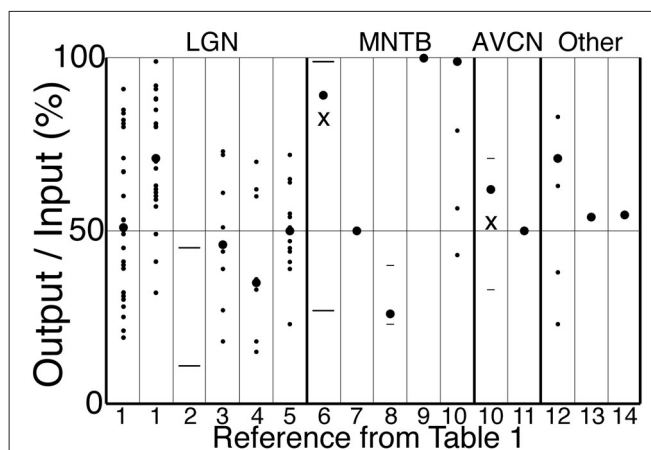


FIGURE 4 | Output to input ratios from the references in Table 1.

Symbols indicate mean across cells (large circles), individual cells (small circles), s.d. or quantiles (short lines), and range (long lines). For references 6 and 10, the “X” (mean) and other symbols correspond to a subset of cells in which O/I was less than 1, whereas large circles indicate our estimates of the total population mean. The strongest stimulus condition is shown for those references for which more than one condition is reported in **Table 1**. Reference 1 has data for cat (left) and monkey (right). A single cell with O/I of 150% is not shown for “12”.

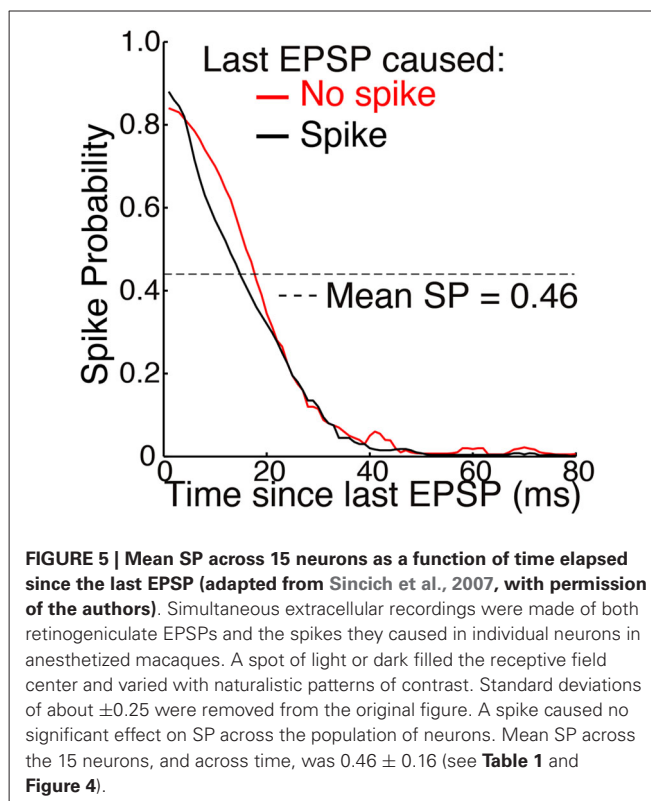


FIGURE 5 | Mean SP across 15 neurons as a function of time elapsed since the last EPSP (adapted from Sincich et al., 2007, with permission of the authors). Simultaneous extracellular recordings were made of both retinogeniculate EPSPs and the spikes they caused in individual neurons in anesthetized macaques. A spot of light or dark filled the receptive field center and varied with naturalistic patterns of contrast. Standard deviations of about ± 0.25 were removed from the original figure. A spike caused no significant effect on SP across the population of neurons. Mean SP across the 15 neurons, and across time, was 0.46 ± 0.16 (see **Table 1** and **Figure 4**).

Casti et al., 2008) if cells have a nearly optimal leak conductance (**Figure 2**).

Sincich et al. (2007) also found that SP was not altered following a spike (**Figure 5**). This supports our theory, according

Table 1 | Time averaged input and output rates and ratios from studies that measured both simultaneously over seconds or more in cells known to receive driving excitation from just one or a few presynaptic neurons.

Reference	Cells	Stimulus	N	Input (Hz)	Output (Hz)	O/I (%)
1. Kaplan et al. (1987)	Cat LGN	Contrast 64–80%	37	NA	NA	51 ± 20, 19–91
	Monkey LGN	Contrast 64–80%	26	NA	NA	71 ± 17, 32–99
2. Movshon et al. (2005)	Monkey LGN	Contrast	7	NA	NA	11–45
3. Sinich et al. (2007)	Monkey LGN	Contrast	15	29 ± 9 (n = 13)	13*	46 ± 16
4. Casti et al. (2008)	Cat LGN	Contrast	14	NA	NA	35 ± 19, 7–70
5. Weyand (2007)	Cat LGN	Contrast	12	50 ± 16	24 ± 8	50 ± 13
6. Lorteije et al. (2009)	Mouse MNTB	Spontaneous	11	71 ± 11, 0.4–174	62*	87, 32–99
		Monaural tone, 80 dB	15	352 ± 34	289*	82, 27–99
7. Hermann et al. (2007)	Gerbil MNTB	300 Hz Elec. Stim.	18	300	150*	50*
8. Kopp-Scheinflug et al. (2003)	Gerbil MNTB	Spontaneous	54	49 [17–79] 1–308	21 [6–45] 0–189	48 [36–71] 5–94
		One tone		355 [212–486]	194 [126–230]	55 [38–74]
				48–868	18–334	23–95
		Two tone		NA	NA	26 [23–40]
9. McLaughlin et al. (2008)	Cat MNTB	Binaural tone	49	100–500*	100–500*	100
10. Englitz et al. (2009)	Gerbil MNTB	Spontaneous	3	NA	NA	79, 39, 43
	Gerbil AVCN	Spontaneous	27	NA	NA	76 ± 17
		1 and 2 tone monaural	34	NA	NA	52 ± 19
11. Kuenzel et al. (2011)	Gerbil AVCN	Spontaneous	39	79*	56 ± 23	71 ± 26
		Monaural Tone,	39	NA	NA	50*
		51 ± 11 dB				
12. Chadderton et al. (2004)	Rat cerebellar granule cells	Tactile, Ipsilateral	5	3.2* EPSCs	2.2* spikes	71, 23–150*
13. Bratton et al. (2010)	Rat lumbar SG	Spontaneous	38	5.4	2.9 ± 0.3	54*
14. McAllen et al. (2011)	Rat cardiac SG	Spontaneous	6	4.1 ± 0.7	0.2–4.3	55*

Portions of these data are illustrated graphically in **Figure 4**. The output to input ratio (O/I) should match average SP in principle, but could differ due to a technical failure to detect an input, especially in cases of high frequency bursts of input. Only “12” reported a single cell in which output exceeded input, but “12” was also the only study in which input and output counts were not measured simultaneously. None of these studies reported a significant incidence of either spikes in the absence of input, or more than one output spike resulting from a single input spike. Data are number of neurons (N), population mean ± sd (median and quantiles for “8”) and range, except for 3 single cell’s in MNTB for “10”. Asterisks (*) indicate our estimates based either on our calculations or visual inspection of figures. NA: not available. “Spontaneous” indicates that sensory conditions were not specified and may not have been controlled. Studies were in anesthetized animals except “5” in awake cats, and “7” and “14” in vitro during efforts to mimic natural conditions. Data was obtained with extracellular recordings, except intracellular recordings in 6–7 and 12–14. In “6” and “10,” O/I are only presented for those neurons that displayed “failures” (O/I < 1.0), whereas O/I was reported as 1.0 in the remainder of 24 total neurons in “6,” and in 63, 21, and 9 other cells in “10” in the cases of spontaneous activity in MNTB and AVCN, and sound evoked activity in AVCN, respectively.

to which opening of potassium and other channels following an action potential helps to maintain homeostatic excitability rather than causing hyperpolarization (see Mechanisms and Timescales).

Given sufficient knowledge of natural patterns, the theory should allow us to predict membrane excitability and the properties of specific ion channels. In work presented in preliminary form and currently under review for publication, we provide evidence in brain slices of LGN that T-type Ca^{2+} channels promote predictive homeostasis by restoring SP towards 0.5 during naturally occurring periods of hyperpolarization like that illustrated in **Figure 3** (Hong et al., 2013). Without predictive amplification of EPSPs by T-type channels, SP would be near 0 for extended periods and spike output would be temporarily “blind” to visual stimuli. In support of theory, we found that the amplification of retinogeniculate EPSPs by T-type channels typically results in just 0 or 1 spike. Previous reports of T-type driven bursts of multiple spikes were most likely due to the unnaturally high membrane

excitability that is typical of experimental conditions in brain slices.

EVIDENCE FROM OTHER TYPES OF NEURONS

In the medial nucleus of the trapezoid body (MNTB) in the early auditory system, the calyx of Held is an exceptionally powerful synapse in which a single presynaptic terminal wraps almost entirely around the postsynaptic soma. It was long believed to be a reliable “fail-safe relay” with SP very near 1. However, the existence of powerful synaptic inhibition suggests that SP should sometimes be lower (Awatramani et al., 2004). SP has been found to be intermediate (Kopp-Scheinflug et al., 2003) or almost 1.0 (McLaughlin et al., 2008; Englitz et al., 2009) based on extracellular recordings, but the more definitive technique of whole-cell patch recording found SP across cells to be near 0.9 *in vivo* with a monaural single tone (Lorteije et al., 2009). When attempts were made to mimic *in vivo* conditions within brain slices, SP varied across a range of about 0.2–1.0 centered near 0.5 (Hermann et al., 2007). None of the *in vivo* studies attempted to examine natural

patterns of sound, and most used monaural or binaural pure tones (**Table 1**). Pure tones and monaural stimulation are rare events in the natural world. Sound intensities are usually similar across tones (“broadband”) and in both ears, and this likely provides the basis for homeostatic synaptic “surround” inhibition that suppresses SP, as found in many other neurons. Indeed, SP in response to the neuron’s preferred tone was suppressed (below 0.4) by simultaneous presentation of a second tone (**Table 1**; Kopp-Scheinflug et al., 2003).

The endbulb of Held in the anteroventral cochlear nucleus (AVCN) is another powerful axosomatic synapse, and SP has been found to be near 0.5 during acoustic stimulation (Englitz et al., 2009; Kuenzel et al., 2011). Further support for the theory comes from the observation that as EPSP amplitude declined as a function of inter-EPSP interval over a period of milliseconds (apparently due to presynaptic decay of paired-pulse facilitation), excitability (and SP) increased in a corresponding manner so that spike generation adapted its sensitivity to the smaller range of EPSP amplitudes (Kuenzel et al., 2011).

That even the most powerful axosomatic synapses appear to have average SP nearer to 0.5 than 1.0 under natural conditions supports the theory, and makes it doubtful whether average SP is near 1.0 at any synapse. The climbing fiber to Purkinje cell synapse in the cerebellum is reliable (SP = 1) under standard conditions in brain slices, but the strength of other powerful synapses has been found to be much weaker under natural conditions *in vivo*, due to both pre- and postsynaptic factors (see Borst, 2010, for a review of presynaptic factors). We did not find any reports of stable, simultaneous recordings of pre- and postsynaptic activity *in vivo* at the climbing fiber to Purkinje cell synapse.

A typical cerebellar granule cell has four excitatory synapses of approximately equal strength, whereas sympathetic ganglion neurons are autonomic motoneurons that receive synaptic excitation from about 2–5 cholinergic preganglionic neurons, one of which is typically the main driver. Intracellular recordings have found SP to be near 0.5 in each of these cell types (**Table 1**; **Figure 4**; Chadderton et al., 2004; Bratton et al., 2010; McAllen et al., 2011).

In cells with larger numbers of drivers, such as cortical neurons, unitary EPSPs in the soma are typically not more than 1 mV, and significant spatial summation across synapses is required to evoke a spike. According to the theory, small unitary EPSPs should not be common under natural conditions, since that would cause SP to be very low. Indeed, there can be a high degree of synchronous excitation causing large EPSPs in at least some cortical neurons (DeWeese and Zador, 2006; Poulet and Petersen, 2008; Okun and Lampl, 2008; Gentet et al., 2010; Yu and Ferster, 2010), as suggested by simulations and theory (Softky and Koch, 1993; Wang et al., 2010a). In auditory cortex, extracellular recordings of spikes (DeWeese et al., 2003) and intracellular recordings of EPSPs with spikes blocked (DeWeese and Zador, 2006), were each performed under otherwise similar conditions. Taken together, these latter two studies imply that large EPSPs cause spikes with intermediate SP.

We have summarized above all of the highly relevant data that we found from simultaneous recordings of synaptic drive and spike output, and it supports our hypothesis that average SP is near 0.5 under natural conditions. However, of the modest

number of *in vivo* studies that are most relevant to our hypothesis, few reported direct and detailed analyses of SP, and few attempted to study natural patterns of synaptic excitation.

VESICLE RELEASE PROBABILITY AT PRESYNAPTIC TERMINALS

When a spike arrives at a single release site in a presynaptic axon terminal, it typically causes release of just 0 or 1 vesicle. Release is commonly said to be “stochastic”, but it is a highly regulated process (Branco and Staras, 2009), and the appearance of “randomness” reflects our ignorance of the internal machinery that determines whether or not a vesicle is released (see Fiorillo, 2012). It seems almost inconceivable that biology could not have achieved reliable transmission if it were desirable. We propose that vesicle release is analogous to spike generation, and that it only occurs when spike input exceeds its expected value. Since a neuron’s spike conveys information that is qualitatively identical to its EPSP input (it shares the same receptive field and thus concerns the same part of the world), it may seem redundant and pointless to repeat in the axon terminal the same sort of process that occurs in the dendrites and soma. However, an axon terminal is generally far from the soma, and we presume that there is information delivered to the terminal from local neurons (for example, via GABA). In principle this information could be conveyed to the soma, but that would consume time and energy, and it is thus more efficient to use it locally to predict incoming spikes, as previously suggested (Fiorillo, 2008).

Spikes entering a terminal from the axon would presumably be the sole driver and proximal cause of almost all vesicle release. They act via a transient increase in calcium conductance (analogous to an EPSP) and cytosolic concentration (analogous to an EPSP). The free calcium concentration varies from spike to spike and depends on a variety of factors, including the state of potassium channels (that influence action potential duration), calcium channels, and calcium buffers (Meinrenken et al., 2003). A vesicle that is “docked” and “ready” will be released in the event that calcium concentration exceeds a threshold (created by the requirement for binding of multiple Ca^{2+} ions to proteins that drive vesicle exocytosis). “Vesicle releasability” (analogous to membrane excitability) would be defined as the energy barrier separating the state of the release site from release threshold. In analogy to Hypothesis 1, it would correspond to the expectation at the onset of a spike of the peak calcium conductance (which would typically be reached a fraction of a millisecond later). Vesicle release would therefore only occur when the expectation is exceeded. Since we cannot directly observe releasability, we define release probability (RP, more commonly abbreviated “ P_r ”) as the probability that an action potential will cause release of at least one vesicle at a single release site. As in the case of SP, RP reflects our ignorance of the internal state of the synaptic terminal.

If we could measure RP in real time at a single release site, then we would expect it to have the general features of SP, with numerous processes of excitation and inhibition (“facilitation and depression”) acting on numerous timescales. The goal of predictive homeostasis would be to make accurate predictions of calcium conductance, with the result that vesicle release is maximally

sensitive to small variations in the amplitude of calcium events (conductance and concentration). A perfectly accurate expectation would result in RP of 0.5. Like SP, we would expect RP to vary dynamically from nearly 0 to 1 with a homeostatic average of 0.5.

Average RP (across release sites and over time) has been found to fall within an intermediate range (~ 0.1 – 0.9) with an “average of averages” near 0.5 across a diversity of synapses (Branco and Staras, 2009). However, measures of RP are typically performed in brain slices in the presence of unnaturally high extracellular calcium and low firing rates (Hermann et al., 2007; Borst, 2010). A more general obstacle is that synaptic terminals are relatively poorly understood, mostly due to their small size, and this limits our ability to “take the release site’s point of view”. If we had a model of vesicle release comparable to a Hodgkin-Huxley model of membrane voltage, and the ability to monitor local calcium concentration at a single release site on a millisecond timescale, then we would not need to rely so heavily on the crude and indirect measure of RP. Typical measures of RP are less direct than what we have presented above for SP, since they are usually averages over many release sites, sometimes involving multiple presynaptic neurons.

We would like to know RP as a function of ISI, analogous to SP in **Figure 5**. However, studies often track the dynamics of RP (through its influence on EPSP amplitude) without quantifying its actual value. Because calcium concentration depends strongly on ISIs, and predictions should be accurate in the case of commonly occurring ISIs, we would expect that RP should be particularly sensitive to variations near these typical intervals. Indeed, in neurons of the fish electrosensory system that have average firing rates near 200 Hz, RP is highly sensitive to ISI variations near the average of 5 ms (Khanbabaie et al., 2010).

A unitary EPSP usually involves multiple release sites, and the powerful synapses discussed above have hundreds. Thus a presynaptic spike will almost always cause an EPSP, and its amplitude will be determined by the average RP across multiple release sites. Because homeostatic mechanisms at each release site are all working independently to “pull” RP towards 0.5, and because each release site experiences the same pattern of spikes, average RP across a presynaptic neuron’s synapses (and the corresponding postsynaptic EPSP amplitude) should not show large variations over time (depression and facilitation should be modest). This has indeed been found to be the case *in vivo* (Borst, 2010). However, maintaining average RP near 0.5 across release sites would maximize the variance (over time and release sites) in RP and unitary EPSP amplitude relative to any other average RP. The resulting variations in EPSP amplitude should convey information, and effective homeostasis of excitability in the postsynaptic neuron would maximize the sensitivity of spike output to these variations (and thus to presynaptic RP). In support of this theory, it has been found that although variations in presynaptic RP and postsynaptic EPSP amplitude are modest, spike output adapts to maintain high sensitivity to them (Kuenzel et al., 2011).

CONCLUSION

We believe that taking the neuron’s point of view is essential if we are to understand the relation between biophysical mechanisms

and their information content. Perhaps the most striking insight provided by our approach is its simple explanation of why spikes appear so “random”. What appears noisy and even pathological from the standard perspective of a physiologist can be readily understood as optimal from the perspective of a neuron. Taking the neuron’s point of view is a radical departure from past efforts to understand the neural “code”, and it could provide the foundation for a general theory of brain function that is able in principle to explain and predict the diversity of neural mechanisms (Fiorillo, 2008, 2012; Hong et al., 2013).

AUTHOR CONTRIBUTIONS

Christopher D. Fiorillo conceived the ideas and wrote the manuscript. Jaekyung K. Kim performed the simulations shown in **Figures 1** and **2**, and Jaekyung K. Kim and Christopher D. Fiorillo designed **Figures 1** and **2**. Su Z. Hong and Christopher D. Fiorillo reviewed the experimental literature on SP and designed **Table 1** and **Figure 4**.

ACKNOWLEDGMENTS

We thank Sunil Kim and Fleur Zeldenrust for helpful comments on the manuscript, and Lawrence Sincich for advice on describing and reproducing his data (**Figure 5**). Research was supported by a grant to Christopher D. Fiorillo from the Basic Science Research Program (2012R1A1A2006996) of the National Research Foundation of Korea (NRF) funded by the Ministry of Education, Science and Technology.

REFERENCES

- Attwell, D., and Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metab.* 21, 1133–1145. doi: 10.1097/00004647-200110000-00001
- Avissar, M., Furman, A. C., Saunders, J. C., and Parsons, T. D. (2007). Adaptation reduces spike-count reliability, but not spike-timing precision, of auditory nerve responses. *J. Neurosci.* 27, 6461–6472. doi: 10.1523/jneurosci.5239-06.2007
- Awatramani, G. B., Turecek, R., and Trussell, L. O. (2004). Inhibitory control at a synaptic relay. *J. Neurosci.* 24, 2643–2647. doi: 10.1523/jneurosci.5144-03.2004
- Bair, W., and Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* 8, 1185–1202. doi: 10.1162/neco.1996.8.6.1185
- Barlow, H. B. (1961). “Possible principles underlying the transformation of sensory messages,” in *Sensory Communication*, ed W. A. Rosenblith (Cambridge, MA: MIT Press), 217–234.
- Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., et al. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142–1152. doi: 10.1016/j.neuron.2008.09.021
- Berg, R. W., Alaburda, A., and Hounsgaard, J. (2007). Balanced inhibition and excitation drive spike activity in spinal half-centers. *Science* 315, 390–393. doi: 10.1126/science.1134960
- Berkes, P., Orbán, G., Lengyel, M., and Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331, 83–87. doi: 10.1126/science.1195870
- Berry, M. J., Warland, D. K., and Meister, M. (1997). The structure and precision of retinal spike trains. *Proc. Natl. Acad. Sci. U S A* 94, 5411–5416. doi: 10.1073/pnas.94.10.5411
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* 2, 32–48.
- Blitz, D. M., and Regehr, W. G. (2005). Timing and specificity of feed-forward inhibition within the LGN. *Neuron* 45, 917–928. doi: 10.1016/j.neuron.2005.01.033

- Borst, J. G. G. (2010). The low synaptic release probability in vivo. *Trends Neurosci.* 33, 259–266. doi: 10.1016/j.tins.2010.03.003
- Branco, T., and Häusser, M. (2011). Synaptic integration gradients in single cortical pyramidal cell dendrites. *Neuron* 69, 885–892. doi: 10.1016/j.neuron.2011.02.006
- Branco, T., and Staras, K. (2009). The probability of neurotransmitter release: variability and feedback control at single synapses. *Nat. Rev. Neurosci.* 10, 373–383. doi: 10.1038/nrn2634
- Bratton, B., Davies, P., Jänig, W., and McAllen, R. (2010). Ganglionic transmission in a vasomotor pathway studied in vivo. *J. Physiol.* 588, 1647–1659. doi: 10.1113/jphysiol.2009.185025
- Brenner, N., Bialek, W., and de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. *Neuron* 26, 695–702. doi: 10.1016/s0896-6273(00)81205-2
- Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M., et al. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature* 449, 92–95. doi: 10.1038/nature06105
- Calvin, W. H., and Stevens, C. F. (1968). Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J. Neurophysiol.* 31, 574–587.
- Carandini, M. (2004). Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol.* 2:e264. doi: 10.1371/journal.pbio.0020264
- Carandini, M., Horton, J. C., and Sincich, L. C. (2007). Thalamic filtering of retinal spike trains by postsynaptic summation. *J. Vis.* 7, 20–20. doi: 10.1167/7.14.20
- Casti, A., Hayot, F., Xiao, Y., and Kaplan, E. (2008). A simple model of retina-LGN transmission. *J. Comput. Neurosci.* 24, 235–252. doi: 10.1007/s10827-007-0053-7
- Chadderton, P., Margrie, T. W., and Häusser, M. (2004). Integration of quanta in cerebellar granule cells during sensory processing. *Nature* 428, 856–860. doi: 10.1038/nature02442
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378. doi: 10.1038/nn.2501
- Clark, A. (2013). Whatever next? Predictive brains, situated agents and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204. doi: 10.1017/s0140525x12000477
- Cooper, L. N., and Bear, M. F. (2012). The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nat. Rev. Neurosci.* 13, 798–810. doi: 10.1038/nrn3353
- Dan, Y., Atick, J. J., and Reid, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* 16, 3351–3362.
- Deneve, S. (2008a). Bayesian spiking neurons I: inference. *Neural Comput.* 20, 91–117. doi: 10.1162/neco.2008.20.1.91
- Deneve, S. (2008b). Bayesian spiking neurons II: learning. *Neural Comput.* 20, 118–145. doi: 10.1162/neco.2008.20.1.118
- Deneve, S., Latham, P. E., and Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nat. Neurosci.* 4, 826–831. doi: 10.1038/90541
- de Ruyter van Steveninck, R. R., Lewen, G. D., Strong, S. P., Koberle, R., and Bialek, W. (1997). Reproducibility and variability in neural spike trains. *Science* 275, 1805–1808. doi: 10.1126/science.275.5307.1805
- DeWeese, M. R., Wehr, M., and Zador, A. M. (2003). Binary spiking in auditory cortex. *J. Neurosci.* 23, 7940–7949.
- DeWeese, M. R., and Zador, A. M. (2006). Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *J. Neurosci.* 26, 12206–12218. doi: 10.1523/jneurosci.2813-06.2006
- Englitz, B., Tolnai, S., Typlt, M., Jost, J., and Rübsamen, R. (2009). Reliability of synaptic transmission at the synapses of held in vivo under acoustic stimulation. *PLoS One* 4:e7014. doi: 10.1371/journal.pone.0007014
- Fiorillo, C. D. (2008). Towards a general theory of neural computation based on prediction by single neurons. *PLoS One* 3:e3298. doi: 10.1371/journal.pone.0003298
- Fiorillo, C. D. (2012). Beyond Bayes: on the need for a unified and Jaynesian definition of probability and information within neuroscience. *Information* 3, 175–203. doi: 10.3390/info3020175
- Fiorillo, C. D. (2013a). Prediction by single neurons. *United States Patent* 8504502.
- Fiorillo, C. D. (2013b). Two dimensions of value: dopamine neurons represent reward but not aversiveness. *Science* 341, 546–549. doi: 10.1126/science.1238699
- Fiorillo, C. D., Song, M. R., and Yun, S. R. (2013). Multiphasic temporal dynamics in responses of midbrain dopamine neurons to appetitive and aversive stimuli. *J. Neurosci.* 33, 4710–4725. doi: 10.1523/jneurosci.3883-12.2013
- Frémaux, N., Sprekeler, H., and Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. *J. Neurosci.* 30, 13326–13337. doi: 10.1523/jneurosci.6249-09.2010
- Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138. doi: 10.1038/nrn2787
- Friston, K. J., Thornton, C., and Clark, A. (2012). Free-energy minimization and the dark-room problem. *Front. Psychol.* 3:130. doi: 10.3389/fpsyg.2012.00130
- Gentet, L. J., Avermann, M., Matyas, F., Staiger, J. F., and Petersen, C. C. H. (2010). Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice. *Neuron* 65, 422–435. doi: 10.1016/j.neuron.2010.01.006
- Hermann, J., Pecka, M., von Gersdorff, H., Grothe, B., and Klug, A. (2007). Synaptic transmission at the calyx of held under in vivo-like activity levels. *J. Neurophysiol.* 98, 807–820. doi: 10.1152/jn.00355.2007
- Higley, M. J., and Contreras, D. (2006). Balanced excitation and inhibition determine spike timing during frequency adaptation. *J. Neurosci.* 26, 448–457. doi: 10.1523/jneurosci.3506-05.2006
- Hong, S., Kim, H. R., Yun, S. R., and Fiorillo, C. D. (2013). “T-type calcium channels promote predictive homeostasis in thalamocortical neurons of LGN,” in *Computational and Systems Neuroscience 2013 (COSYNE)* (Salt Lake City, UT, USA).
- Hosoya, T., Baccus, S. A., and Meister, M. (2005). Dynamic predictive coding by the retina. *Nature* 436, 71–77. doi: 10.1038/nature03689
- Hoyer, P. O., and Hyvarinen, A. (2003). “Interpreting neural response variability as Monte Carlo sampling of the posterior,” in *Advances in Neural Information Processing Systems 15*, eds S. Becker, et al. (Cambridge, MA: MIT Press), 277–284.
- Izhikevich, E. M. (2006). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* 17, 2443–2452. doi: 10.1093/cercor/bhl152
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge, UK: Cambridge University Press.
- Jazayeri, M., and Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nat. Neurosci.* 9, 690–696. doi: 10.1038/nn1691
- Kaplan, E., Purpura, K., and Shapley, R. M. (1987). Contrast affects the transmission of visual information through the mammalian lateral geniculate nucleus. *J. Physiol.* 391, 267–288.
- Kaplan, E., and Shapley, R. (1984). The origin of the S (slow) potential in the mammalian lateral geniculate nucleus. *Exp. Brain Res.* 55, 111–116. doi: 10.1007/BF00240504
- Khanbabaie, R., Nesse, W. H., Longtin, A., and Maler, L. (2010). Kinetics of fast short-term depression are matched to spike train statistics to reduce noise. *J. Neurophysiol.* 103, 3337–3348. doi: 10.1152/jn.00117.2010
- Kopp-Scheinpflug, C., Lippe, W. R., Dorrscheidt, G. J., and Rübsamen, R. (2003). The medial nucleus of the trapezoid body in the gerbil is more than a relay: comparison of pre- and postsynaptic activity. *J. Assoc. Res. Otolaryngol.* 4, 1–23. doi: 10.1007/s10162-002-2010-5
- Körding, K. P., and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature* 427, 244–247. doi: 10.1038/nature02169
- Kuenzel, T., Borst, J. G. G., and van der Heijden, M. (2011). Factors controlling the input-output relationship of spherical bushy cells in the gerbil cochlear nucleus. *J. Neurosci.* 31, 4260–4273. doi: 10.1523/jneurosci.5433-10.2011
- Laughlin, S. B. (1981). A simple coding procedure enhances a neuron’s information capacity. *Z. Naturforsch. C* 36, 910–912.
- Linsker, R. (1988). Self-organization in a perceptual network. *Computer* 21, 105–117. doi: 10.1109/2.36
- London, M., Roth, A., Beeren, L., Häusser, M., and Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* 466, 123–127. doi: 10.1038/nature09086
- Lorteije, J. A. M., Rusu, S. I., Kushmerick, C., and Borst, J. G. G. (2009). Reliability and precision of the mouse calyx of Held synapse. *J. Neurosci.* 29, 13770–13784. doi: 10.1523/jneurosci.3285-09.2009

- Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438. doi: 10.1038/nn1790
- Mainen, Z. F., and Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science* 268, 1503–1506. doi: 10.1126/science.7770778
- Marder, E., and Bucher, D. (2007). Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annu. Rev. Physiol.* 69, 291–316. doi: 10.1146/annurev.physiol.69.031905.161516
- Mastrorarde, D. N. (1989). Correlated firing of retinal ganglion cells. *Trends Neurosci.* 12, 75–80. doi: 10.1016/0166-2236(89)90140-9
- McAllen, R. M., Salo, L. M., Paton, J. F. R., and Pickering, A. E. (2011). Processing of central and reflex vagal drives by rat cardiac ganglion neurons: an intracellular analysis. *J. Physiol.* 589, 5801–5818. doi: 10.1113/jphysiol.2011.214320
- McLaughlin, M., van der Heijden, M., and Joris, P. X. (2008). How secure is in vivo synaptic transmission at the calyx of Held? *J. Neurosci.* 28, 10206–10219. doi: 10.1523/JNEUROSCI.2735-08.2008
- Meinrenken, C. J., Borst, J. G. G., and Sakmann, B. (2003). Local routes revisited: the space and time dependence of the Ca²⁺ signal for phasic transmitter release at the rat calyx of Held. *J. Physiol.* 547, 665–689. doi: 10.1113/jphysiol.2002.032714
- Migliore, M., Hoffman, D. A., Magee, J. C., and Johnston, D. (1999). Role of an A-type K⁺ conductance in the back-propagation of action potentials in the dendrites of hippocampal pyramidal neurons. *J. Comput. Neurosci.* 7, 5–15. doi: 10.1023/A:1008906225285
- Movshon, J. A., Kiorpes, L., Hawken, M. J., and Cavanaugh, J. R. (2005). Functional maturation of the macaque's lateral geniculate nucleus. *J. Neurosci.* 25, 2712–2722. doi: 10.1523/jneurosci.2356-04.2005
- Niemeier, M., Crawford, J. D., and Tweed, D. B. (2003). Optimal transsaccadic integration explains distorted spatial perception. *Nature* 422, 76–80. doi: 10.1038/nature01439
- Okun, M., and Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* 11, 535–537. doi: 10.1038/nn.2105
- Pouget, A., Dayan, P., and Zemel, R. (2000). Information processing with population codes. *Nat. Rev. Neurosci.* 1, 125–132. doi: 10.1038/35039062
- Pouille, F., and Scanziani, M. (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293, 1159–1163. doi: 10.1126/science.1060342
- Poulet, J. F. A., and Petersen, C. C. H. (2008). Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454, 881–885. doi: 10.1038/nature07150
- Rieke, F., Warland, D., de Ruyter van Steveninck, R. R., and Bialek, W. (1997). *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. doi: 10.1126/science.275.5306.1593
- Sengupta, B., Laughlin, S. B., and Niven, J. E. (2013). Balanced excitatory and inhibitory synaptic currents promote efficient coding and metabolic efficiency. *PLoS Comput. Biol.* 9:e1003263. doi: 10.1371/journal.pcbi.1003263
- Shadlen, M. N., and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4, 569–579. doi: 10.1016/0959-4388(94)90059-0
- Shadlen, M. N., and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J. Neurosci.* 18, 3870–3896.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.* 27, 379–423. doi: 10.1002/j.1538-7305.1948.tb01338.x
- Sharpee, T. O., Sugihara, H., Kurgansky, A. V., Rebrik, S. P., Stryker, M. P., and Miller, K. D. (2006). Adaptive filtering enhances information transmission in visual cortex. *Nature* 439, 936–942. doi: 10.1038/nature04519
- Sherman, S. M., and Guillery, R. W. (1998). On the actions that one nerve cell can have on another: distinguishing “drivers” from ‘modulators’. *Proc. Natl. Acad. Sci. U S A* 95, 7121–7126. doi: 10.1073/pnas.95.12.7121
- Shu, Y., Hasenstaub, A., and McCormick, D. A. (2003). Turning on and off recurrent balanced cortical activity. *Nature* 423, 288–293. doi: 10.1038/nature01616
- Simoncelli, E. P., and Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci.* 24, 1193–1216. doi: 10.1146/annurev.neuro.24.1.1193
- Sincich, L. C., Adams, D. L., Economides, J. R., and Horton, J. C. (2007). Transmission of spike trains at the retinogeniculate synapse. *J. Neurosci.* 27, 2683–2692. doi: 10.1523/jneurosci.5077-06.2007
- Softky, W. R. (1995). Simple codes versus efficient codes. *Curr. Opin. Neurobiol.* 5, 239–247. doi: 10.1016/0959-4388(95)80032-8
- Softky, W. R., and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13, 334–350.
- Srinivasan, M. V., Laughlin, S. B., and Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B Biol. Sci.* 216, 427–459. doi: 10.1098/rspb.1982.0085
- Stemmler, M., and Koch, C. (1999). How voltage-dependent conductances can adapt to maximize the information encoded by neuronal firing rate. *Nat. Neurosci.* 2, 521–527. doi: 10.1038/9173
- Tolhurst, D. J., Movshon, J. A., and Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 23, 775–785. doi: 10.1016/0042-6989(83)90200-6
- Usrey, W. M., Reppas, J. B., and Reid, R. C. (1998). Paired-spike interactions and synaptic efficacy of retinal inputs to the thalamus. *Nature* 395, 384–387. doi: 10.1038/26487
- Wang, X., Hirsch, J. A., and Sommer, F. T. (2010b). Recoding of sensory information across the retinthalamic synapse. *J. Neurosci.* 30, 13567–13577. doi: 10.1523/JNEUROSCI.0910-10.2010
- Wang, H. P., Spencer, D., Fellous, J. M., and Sejnowski, T. J. (2010a). Synchrony of thalamocortical inputs maximizes cortical reliability. *Science* 328, 106–109. doi: 10.1126/science.1183108
- Wang, X., Sommer, F. T., and Hirsch, J. A. (2011). Inhibitory circuits for visual processing in thalamus. *Curr. Opin. Neurobiol.* 21, 726–733. doi: 10.1016/j.conb.2011.06.004
- Wang, X.-J., Liu, Y., Sanchez-Vives, M. V., and McCormick, D. A. (2003). Adaptation and temporal decorrelation by single neurons in the primary visual cortex. *J. Neurophysiol.* 89, 3279–3293. doi: 10.1152/jn.00242.2003
- Wang, X., Wei, Y., Vaingankar, V., Wang, Q., Koepsell, K., Sommer, F. T., et al. (2007). Feedforward excitation and inhibition evoke dual modes of firing in the cat's visual thalamus during naturalistic viewing. *Neuron* 55, 465–478. doi: 10.1016/j.neuron.2007.06.039
- Wehr, M., and Zador, A. M. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426, 442–446. doi: 10.1038/nature02116
- Weiss, Y., Simoncelli, E. P., and Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nat. Neurosci.* 5, 598–604. doi: 10.1038/nn858
- Weyand, T. G. (2007). Retinogeniculate transmission in wakefulness. *J. Neurophysiol.* 98, 769–785. doi: 10.1152/jn.00929.2006
- Wilent, W. B., and Contreras, D. (2005). Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. *Nat. Neurosci.* 8, 1364–1370. doi: 10.1038/nn1545
- Yang, Z., and Purves, D. (2003). A statistical explanation of visual space. *Nat. Neurosci.* 6, 632–640. doi: 10.1038/nn1059
- Yu, J., and Ferster, D. (2010). Membrane potential synchrony in primary visual cortex during sensory stimulation. *Neuron* 68, 1187–1201. doi: 10.1016/j.neuron.2010.11.027

Conflict of Interest Statement: We have cited a patent that discusses the distinction between sensory and motor neurons within its specification (Fiorillo, 2013a). The patented material itself concerns application of Hebbian learning rules to voltage-regulated ion channels, and the patent was not cited with respect to issues of learning.

Received: 04 November 2013; accepted: 02 April 2014; published online: 29 April 2014.
 Citation: Fiorillo CD, Kim JK and Hong SZ (2014) The meaning of spikes from the neuron's point of view: predictive homeostasis generates the appearance of randomness. *Front. Comput. Neurosci.* 8:49. doi: 10.3389/fncom.2014.00049
 This article was submitted to the journal *Frontiers in Computational Neuroscience*.
 Copyright © 2014 Fiorillo, Kim and Hong. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.