Metabolic cost as an organizing principle for cooperative learning

David Balduzzi[1]**, Pedro A. Ortega**[2]**, Michel Besserve**[2]
[1]Swiss Federal Institute for Technology, ETH Zürich
[2]Max Planck Institute for Intelligent Systems, Tübingen, Germany
david.balduzzi@inf.ethz.ch, {ortega, besserve}@tuebingen.mpg.de

Abstract

This paper investigates how neurons can use metabolic cost to facilitate learning at a population level. Although decision-making by individual neurons has been extensively studied, questions regarding how neurons should behave to cooperate effectively remain largely unaddressed. Under assumptions that capture a few basic features of cortical neurons, we show that constraining reward maximization by metabolic cost aligns the information content of actions with their expected reward. Thus, metabolic cost provides a mechanism whereby neurons encode expected reward into their outputs. Further, aside from reducing energy expenditures, imposing a tight metabolic constraint also increases the accuracy of empirical estimates of rewards, increasing the robustness of distributed learning. Finally, we present two implementations of metabolically constrained learning that confirm our theoretical finding. These results suggest that metabolic cost may be an organizing principle underlying the neural code, and may also provide a useful guide to the design and analysis of other cooperating populations.

1   Introduction

Rational decision making is typically formalized as optimizing a reward function [19, 22, 26]. This paper investigates neuronal learning from an optimization perspective. We assume that both the brain as a whole and individual neurons are rational decision makers optimizing reward functions of some kind.

Since the brain learns from finite samples, it is exposed to a tradeoff between over- and under-fitting: increasing a model's capacity can improve its fit on training data, but potentially worsens performance on future samples [29]. Remarkably, however, the human brain effortlessly handles a wide-range of complex pattern recognition tasks suggesting it both has a large capacity and, paradoxically, also generalizes extremely well. On the basis of these conflicting observations, it has been argued that *useful biases* in the form of "generic mechanisms for representation" must be hardwired into cortex [10]. Our goal in this paper is to propose a bias that is both useful and biologically plausible.

Let us outline the problem. Neurons learn inductively. They can generalize from finite samples and encode estimates of future outcomes (for example, rewards) into their spike-trains [12]. Learning-theoretical results imply that generalizing successfully from small samples requires strong biases [29] or, in other words, specialization. Thus, at any given time some neurons' specialties are more relevant than others. Since most of the data neurons receive are other neurons' outputs, it is essential that they indicate which of their outputs

encode high quality estimates. Downstream neurons should then be biased to specialize on these outputs, thereby reducing the effective search space that neurons explore.

The question we ask is: *How can a* single *neuron near-maximize its* expected *reward and simultaneously help downstream neurons do the same?*

As a partial answer, we propose the following organizing principle. Neurons should consistently label outputs as useful and not useful. Specifically, the optimizations performed by neurons should be designed so that spikes are useful and silences are not. For example, it is well known that, on average, the more spikes a neuron receives, the more it learns by modifying its synapses. Although this fact is often taken for granted, it requires explanation. We suggest that spikes drive learning because they are useful. By useful, we mean an output that (i) predicts high reward with (ii) tight confidence intervals.

Distinguishing useful from irrelevant outputs is helpful to downstream neurons since it reduces the size of the search space they are confronted with. If spikes reliably predict future reward, then neurons should be biased towards learning from spikes. In fact, this fits well with experimental evidence [8, 9, 20]. Moreover, consistently biasing learning toward outputs labeled as useful (i.e. spikes) also provides a principled way to reduce capacity, and thus improve generalization guarantees, without sacrificing empirical performance.

The scope of this paper is limited to showing how neurons may systematically distinguish useful from irrelevant outputs. Fleshing out the implications for learning at the population level is deferred to future work. Furthermore, we only consider excitatory connections in this paper; the role of inhibition is also deferred to future work.

Overview. We explore the consequences of a few basic assumptions regarding spikes, metabolic cost and neuronal rewards, see §2. We consider a minimal model of biological neural network from which we extract guiding principles that may apply, at least approximately, to more complex, biologically realistic models. The most important assumptions are that neurons aim to maximize reward after spiking, and that neurons have to operate within a fixed metabolic budget that constrains how often they can spike in a given time interval.

It turns out that these seemingly innocuous assumptions are the key to distinguishing useful from irrelevant outputs. Our first result is that constrained reward maximization causes neurons to maximize the information they encode into their spikes, see §3. Moreover, if spikes are sufficiently rare, it turns out that spikes dominate the information communicated by neurons, which has interesting implications for credit assignment in cortex, see §3 and [5].

If neurons attempt to maximize empirical reward *after spiking*, it follows that neurons encode reward estimates into spikes. When a neuron produces a spike, it thus signals to downstream neurons that it expects a positive neuromodulatory signal such as dopamine. This raises questions concerning the quality of the reward estimates encoded in spikes, see §4. We show that the more information neurons encode into spikes, the tighter the guarantees tying empirical reward to expected reward.

We conclude by describing some implications of our results for cooperative optimization. Theorems are proved in the appendices.

Related work. Neuronal plasticity and its implications for the neural code have been intensively studied for many years. The work closest in spirit to this paper is Seung's "hedonistic" synapses, which seek to increase average reward [24]. A second related line of research applies the information bottleneck method – an alternate constraint to the one considered here – to neuronal learning [7, 27].

An information-theoretic perspective on synaptic homeostasis that complements the results in this paper is [5]. The consequences of constrained optimization for spike-timing dependent plasticity (STDP) are presented in [2]. A practical implementation of regularized STDP, inspired by the ideas presented here, can be found in [16].

2 A minimal model

The mammalian cortex contains between 10^7 and 10^{11} neurons (depending on the species) that are guided by neuromodulators signaling pleasure, pain and other globally salient events. Neurons communicate with each other via spiketrains – sequences of silences and spikes they receive from and transmit to 10^3 to 10^4 other neurons through connections called synapses. Neurons learn by increasing and decreasing the efficacy of their synapses according to the timing of pre-synaptic (input) and post-synaptic (output) spikes, as well as the presence or absence of neuromodulators [8, 9, 15, 20, 25].

Model neurons. The cortex can be modeled as a population of K neurons. Neuron n_k follows policy π_k, where $\pi_k(a|s)$ is a Markov matrix specifying the probability the neuron picks action, or output, $a \in \mathcal{A} = \{0, 1\}$ upon encountering situation $s \in \mathcal{S}$. A situation, or input, is a vector $s = (a_1, \ldots, a_K)$ whose entries are the actions of all the neurons in the brain at the previous time step. Since each neuron is only exposed to a small fraction of the brain, it ignores most entries in the vector. Thus, the policy of neuron k is $\pi_k(a|s) = \pi_k(a|\Sigma_k)$ for some subset $\Sigma_k \subset \{a_1, \ldots, a_K\}$. Let $P(S)$ denote the prior over situations. We will assume the prior is i.i.d. when providing guarantees on estimates, see Remark 4 for a brief discussion.

Neurons are exposed to a global *neuromodulatory signal* $\nu \in \mathcal{N}$ that signals the performance of the population as a whole. Neuromodulators are drawn with probability $P(\nu|s)$.

We make the following assumptions.

Assumption 1 (reward maximization).
Neurons maximize a reward function that depends on neuromodulatory signals, input spikes and output spikes:

$$\widehat{\pi}_k = \arg\max_{\pi \in \mathcal{M}} \sum_{i=1}^{N} R_k(s_i, a_i, \nu_i), \tag{1}$$

where a_i is the output chosen by π in response to input s_i.

The set $\mathcal{M}$ is the set of possible neuronal mechanisms, a subset of the set Markov matrices on $S \times A$. Two examples of $\mathcal{M}$ that are relevant to our discussion are discrete threshold neurons

$$\mathcal{M} = \left\{ H(\langle \mathbf{w}, \mathbf{s} \rangle - \vartheta) \,\middle|\, \mathbf{w} \in \mathbb{R}^K \right\}$$

where $H(\bullet)$ is the Heaviside function, see [2], and the full set of Markov matrices, see discussion of Q-learning below.

Note that, since different neurons are exposed to different subsets of the total brain activity, they have different reward structures, even for the same neuromodulators: in general $R_k \neq R_j$ for $j \neq k$ since $\Sigma_k \neq \Sigma_j$.

Since we focus on the behavior of a single neuron, we will often drop the subscript k from the notation below.

Assumption 2 (spikes gate rewards).
The reward function is gated by synaptic outputs:

$$R(s, a, \nu) = R(s, \nu) \cdot \mathbb{I}_{a=1}, \quad \text{where } \mathbb{I}_{a=1} = \begin{cases} 1 & a = 1 \\ 0 & else \end{cases} \quad \text{is the indicator function.} \tag{2}$$

Neurophysiological evidence suggests that neurons only potentiate or depotentiate their synapses shortly before or after producing spikes [8,9,15,25]. This encourages specialization: neurons search for a small set of inputs that reliably predict future reward signals – other inputs are ignored.

For simplicity we assume neurons have two outputs, spikes and silence. We use notations a_0 or $a = 0$ for silence and a_1 or $a = 1$ for spikes.

Assumption 3 (metabolic budget).
Neurons have a fixed metabolic budget that determines the maximum frequency of spiking over some (unspecified) time period:

$$\pi(a_1) \leq \rho \tag{3}$$

where $\pi(a_1) = \sum_s p(s) \cdot \pi(a_1|s)$ is the spiking frequency under policy π.

Although a soft constraint is more biologically plausible, imposing a hard constraint simplifies the exposition without significantly altering the conclusions. The main effect of softening the constraints is to allow the information carried by spikes and the capacity of neurons to vary, thereby softening Theorems 2 and 3 below.

Since neurons only modify their synapses when they spike, it follows that neurons that spike very infrequently learn very little if at all. Thus, we expect that there are mechanisms in place ensuring that neurons not only stay within their metabolic budget, but also that they come close to using all of it.

Spikes and silences are not abstract, interchangeable symbols. Spiking and responding to spikes carries a much higher metabolic cost than not spiking [13]. This cost is significant since the nervous system consumes a disproportionate share of an organism's total energy budget [1]. It has been hypothesized that a function of sleep is to homeostatically regulate synaptic strengths, so as to control metabolic expenditures associated with action potentials, see [11, 14, 28, 30].

Unlike "agents in the wild", individual neurons have negligible impact on what happens next: a single neuron has little effect on the neurons it targets – since each receives inputs from thousands of other neurons. The inability of individual neurons to manipulate their environment simplifies the optimization problems they face by stripping out the recursive Bellman aspect [6]:

Assumption 4 (disempowerment).
An individual neuron has no immediate influence on its environment

$$p\Big(s(t+1,\ldots)\,\Big|\,s(t),a_k(t)\Big) = p\Big(s(t+1,\ldots)\,\Big|\,s(t)\Big), \tag{4}$$

where t refers to time. I.e. conditioning on the neuron's output makes no difference to the distribution on subsequent actions by the rest of the brain.

The future situations a neuron encounters are essentially unaffected by its output. This assumption fails at the population level – populations of neurons necessarily affect the organism's actions. Nevertheless Eq. (4) is a reasonable assumption at the individual neuron level since, for example, destroying a single neuron makes essentially no difference to brain function.

3 Encoding useful information in spikes

This section considers the implications of our assumptions for the information content of spikes. Mutual information provides a formal method for quantifying the information content of a channel. We present a related measure, effective information, that measures the information content of a single output in terms of how much that output reduces prior uncertainty regarding the set of inputs. Importantly, we show that under the above assumptions, spikes not only reduce uncertainty, but also signal that the neuron received an input that was historically followed by high rewards. The information encoded in spikes is thus indicative of expected future rewards.

Information. It is useful to consider neurons as *communication channels* mapping situations to actions. The average information communicated by a neuron is then the mutual information $I_\pi(\mathcal{S}, \mathcal{A})$. However, since we are interested in the information communicated by specific actions (in particular, spikes), we introduce effective information[1] [3, 4] which quantifies the information that a single output encodes about the input.

---

[1]We extend the definition of effective information in [3] to allow arbitrary priors instead of restricting to the uniform distribution.

Definition 1 (effective information).
Given a neuron with policy $\pi(a|s)$ and prior $P(\mathcal{S})$ on situations, the effective information *generated by action $a \in \mathcal{A}$ is*

$$ei(\pi, a) := D\big[\pi(S|a) \,\big\|\, P(S)\big] \quad where \; \pi(s|a) := \frac{\pi(a|s)}{\pi(a)} \cdot P(s) \tag{5}$$

is computed via Bayes' rule and $D[\bullet\|\bullet]$ is the Kullback-Leibler divergence $D[p\|q] := \sum p_i \log \frac{p_i}{q_i}$.

An interesting special case is when the prior on situations is the uniform distribution and the policy is deterministic. It follows that

$$ei(\pi, a) = -\log \frac{|\pi^{-1}(a)|}{|\mathcal{S}|}, \tag{6}$$

where $|\bullet|$ denotes cardinality and (since the policy is deterministic), π is a function $\pi : \mathcal{S} \to \mathcal{A}$. In Eq (6), effective information quantifies the *selectivity* of an output: the fraction of inputs causing the policy to output a. The smaller the fraction, or alternatively the more sensitive output a is to perturbations in the input, the higher effective information [5].

Remark 1. Suppose we have model $P_\mathcal{M}(d|h)$ that specifies the probability of observing data given a hypothesis. Further suppose we have prior distribution $P(h)$ on hypotheses. If we observe data d, how much have we learned about the hypotheses? The *Bayesian information gain* is

$$D\big[P_\mathcal{M}(H|d) \,\big\|\, P(H)\big]. \tag{7}$$

If we consider a neuron's policy as a model, with inputs as hypotheses and outputs as evidence, then effective information quantifies the Bayesian information gained about the inputs given an output.

Remark 2. The *expectation* of ei is mutual information: $\mathbb{E}_{\pi(a)}\big[ei(\pi, a)\big] = I_\pi(\mathcal{S}, \mathcal{A})$.

Information aligns with rewards (theory). We show that neurons implementing constrained reward maximization from Assumption 1 also maximize the effective information of their spikes $ei(\pi, a_1)$, that we will call *information per spike*.

Definition 2 (empirical reward).
Given a finite sample of situations, actions and neuromodulators $(s_i, a_i, \nu_i)_{i=1}^N$, let the empirical reward *observed after performing action a in situation s be*

$$\widehat{R}(s, a) := \frac{1}{N} \sum_{\{i|s_i=s, a_i=a\}} R(s_i, a_i, \nu_i).$$

We also introduce the empirical reward after spiking, $\frac{1}{N} \sum_{\{i|a_i=1\}} R\big(s_i, a_1, \nu_i\big)$ *and the* empirical reward per spike $\frac{1}{|\{i|a_i=1\}|} \sum_{\{i|a_i=1\}} R\big(s_i, a_1, \nu_i\big)$.

Theorem 1 (maximizing reward/spike maximizes information/spike).
Assume that situations yield different empirical rewards, i.e. $\widehat{R}(s, a_1) \neq \widehat{R}(s', a_1)$ for all $s, s' \in \mathcal{S}$. Recall that by Assumptions 1 and 3 an optimal policy $\widehat{\pi}$ satisfies

$$\widehat{\pi} = \underset{\{\pi \in \mathcal{M} | \pi(a_1) \leq \rho\}}{\arg\max} \sum_{i=1}^N R(s_i, a_i, \nu_i),$$

If we also apply constraint $\pi(a_1) \geq \rho$, i.e. ρ is both upper and lower bound, then the optimal policy $\widehat{\pi}$ maximizes information per spike. More precisely, the optimal policy $\widehat{\pi}$ satisfies

$$ei(\widehat{\pi}, a_1) \geq ei(\pi, a_1) \; for \; all \; \pi \; such \; that \; \pi(a_1) = \rho.$$

Remark 3. The optimal policy will satisfy $\widehat{\pi}(a_1) = \rho$ if there are enough inputs s satisfying $\mathbb{E}_\nu[R(s, \nu, a_1)] > 0$. In other words, if there are enough situations where spiking, on average, is followed by positive reward. Imposing the lower bound in the theorem means we compare the optimal policy with alternate policies that spike with the same frequency.

Thus, the optimal policy necessarily maximizes both the empirical reward after spiking and the effective information encoded in spikes. We illustrate this result with simulations below.

Information aligns with rewards (experiments). A learning algorithm implementing constrained reward maximization is a modification of Q-learning [31]:

Example 1 (Metabolically constrained Q-learning)**.**
If a neuron chooses action a in situation s and subsequently receives neuromodulator ν, then it updates the Q-matrix by

$$Q(s,a) \leftarrow Q(s,a) + \alpha \cdot \left[\widehat{R}(s,a,\nu) - Q(s,a) \right],$$

where α controls the rate. After updating Q, the neuron constructs new policy

$$\pi(a|s) = \mathcal{M}^n \left(e^{Q(s,a)} \right).$$

Operation $\mathcal{M}(\bullet)$ renormalizes the policy twice: first by $Z(a)$ chosen such that $\sum_{s \in \mathcal{S}} \pi(a|s) P(s) = P(a)$ for all a, and then by $Z(s)$ chosen such that $\sum_{a \in \mathcal{A}} \pi(a|s) = 1$ for all s. Setting $n = 3$ yields a policy that approximately implements the metabolic constraint.

Figure 1 shows how effective information and empirical reward after spiking covary as neurons Q-learn. We initialized 5000 neurons randomly and applied metabolic constraints $\rho \in \{0.1, 0.3, 0.5\}$. Rewards are drawn randomly. As the neurons adapt, their policies become both more deterministic and more likely to spike in situations yielding higher rewards, so as neurons adapt they both encode more information into their spikes and predict higher rewards after spiking. The tighter the metabolic constraint (i.e. the lower ρ), the higher the empirical reward after spiking. Thus, the information encoded in spikes provides a reliable guide to the empirical reward after spiking.

Spikes dominate information content. Theorem 1 shows information per spike is maximized by constrained reward maximization. Theorem 2 below consider how much of the total information communicated by a neuron is carried by spikes.

Theorem 2 (spikes carry essentially all information)**.**
Suppose a neuron has two actions (silence a_0 and spike a_1) and produces spikes infrequently: $\pi(a_1) \ll 1$. Then the total information communicated by the neuron is approximately the information it communicates using spikes alone:

$$I_\pi(\mathcal{S}; \mathcal{A}) = \pi(a_1) \cdot ei(\pi, a_1) + O\left(\pi(a_1)^2 \right). \tag{8}$$

If the metabolic constraint is tight, meaning ρ and so $\pi(a_1)$ are small, then for the optimal policy encodes a lot of information into spikes. In this setting, Theorem 2 implies that the information communicated by a neuron is (up to first order) carried by spikes alone.

Eq (6) and Theorem 2 together have interesting implications for credit assignment, see [5] for details. In particular, if spikes carry most of the information in cortex, then spiking neurons and synapses should reinforced in response to positive global signals such as dopamine, and conversely for negative global signals. Neurons and synapses that are silent contribute little to the information generated by cortex, and so should be neither potentiated nor depotentiated. This fits neurophysiological evidence suggesting that spikes play a distinguished role in synaptic potentiation and depotentiation [8, 9, 15, 20].

Encoding reward estimates in spikes. Finally, we briefly illustrate the effect of the metabolic constraint ρ on the empirical reward per spike. Suppose we have sampled empirical rewards $\{\widehat{R}(s, a_1) | s \in \mathcal{S}\}$. The optimal policy is constructed as follows. First, rank states by their empirical reward. Let $\mathcal{S}_\rho$ denote the states in the top ρ^{th} percentile. Define

$$\widehat{\pi}(a_1|s) = \begin{cases} 1 & s \in \mathcal{S}_\rho \\ 0 & \text{else.} \end{cases}$$

The optimal policy can be visualized as moving a window of fixed size and variable shape over the input space, such that the underlying configuration maximizes reward.

Figure 2 shows an example. Situations are ranked according to their empirical reward. The metabolic constraint is set at $\rho = 15\%$, so the optimal policy spikes for the 15% of situations with highest reward. The policy picking these situations results in a reward per spike of 0.89. Situations which do not cause spikes receive *no reward*; if spikes and silences would be exchanged, this reward would be .21 (we call it reward after silence).

We make two observations. First, tightening the metabolic constraint, so that the policy spikes for $\ll 15\%$ of situations, increases the reward per spike. Second, the variance in reward per spike is much lower than after not spiking, and typically decreases with ρ.

As a general rule of thumb, tightening the metabolic constraint by decreasing ρ both increases the empirical reward per spike and reduces the variance in empirical reward per spike.

4 Guarantees on reward estimates

In a dynamically changing environment like the cortex, it is important that neurons reliably represent high reward. This section shows that the reliability of spikes depends on how much information is encoded in them.

Theoretical guarantees. We say that a neuron's spikes reliably represent reward when there is a low variability in the empirical mean reward when the neuron spikes. For a given policy π. The expected and empirical rewards per spike are

$$R_\pi := \mathbb{E}\Big[R(S, a_1, N)\Big] := \sum_{s,n} \pi(s|a_1) \cdot P(\nu|s) \cdot R(s, a_1, \nu) \text{ and}$$

$$\widehat{R}_\pi := \frac{1}{T_1} \sum_{\{(s_t, \nu_t)|\pi(s_t)=1\}} R\big(s_t, a_1, \nu_t\big) \text{ respectively,}$$

where T_1 counts spikes produced by the neuron during $[1, T]$. We measure sample size in terms of spikes (rather than spikes and silences) because spikes are metabolically expensive, so only spikes count towards the cost of collecting a finite sample.

Let us introduce notation for computations with respect to the *uniform prior* $U(s) = \frac{1}{|\mathcal{S}|}$, where $|\mathcal{S}|$ is the total number of possible situations. Let $\pi_u(a) = \sum_s \pi(a|s) \cdot U(s)$, $\pi_u(s|a) = \pi(a|s)\frac{U(s)}{\pi(a)}$, and $ei_u(\pi, a) = \sum_s \pi_u(s|a) \log \frac{\pi_u(s|a)}{\pi_u(a)}$. Note that ei_u recovers the original notion of effective information in [3], which Definition 1 generalizes.

The following theorem is proved using a version of Occam's razor [23].

Theorem 3 (error bound for empirical reward).
Suppose the neuron chooses a deterministic policy π under the constraint that it spikes for a fixed fraction of situations: $\pi_u(a_1) = const$. Further, suppose that situations are sampled i.i.d. Without loss of generality,[2] assume that rewards lie in $[0, b]$. Then with probability at least $1 - \delta$,

$$\Big|R_\pi - \widehat{R}_\pi\Big| \leq b \cdot \sqrt{|\mathcal{S}| \cdot \frac{ei_u(\pi, a_1) + 1}{2T_1 \cdot e^{ei_u(\pi, a_1)}} + \frac{\log \frac{2}{\delta}}{2T_1}} \tag{9}$$

Guarantees improve as ei_u increases since $\frac{x+1}{e^x}$ decreases as x increases.

Note that ei and ei_u covary since increasing the number of situations where a neuron spikes decreases *both* ei and ei_u; similarly, decreasing the number of situations where a neuron spikes increases both ei and ei_u.

Thus, tightening the metabolic constraint in Assumption 3, by choosing low ρ, yields policies that have better guarantees on their reward estimates.

---

[2]Since only relative rewards affect the choice of optimal policy, it follows that negative rewards can be stripped out of the optimization problem by introducing an additive constant.

Remark 4. The assumption that inputs are i.i.d. is not realistic for cortical neurons. We make two remarks. First, if rewards are only non-zero in the presence of neuromodulatory signals, then the assumption states that situations directly preceding neuromodulator release are i.i.d, which is more reasonable. Second, similar results have recently been obtained in non-i.i.d. scenarios using more sophisticated PAC-Bayes methods [21] – and these may be applicable to our setting.

Guarantees in practice. Figure 3 plots information encoded in spikes (x-axis) against the difference between the normalized empirical and expected reward (y-axis). Rewards were drawn randomly and the expected and empirical error for 16,000 deterministic policies with $|\mathcal{S}| = 50$, $P(\mathcal{S})$ uniform and $T_1 = 20$ were computed. Policies were sampled randomly with k, the number of situations causing the policy to spike, varying uniformly across $[1, 25]$. The figure shows that both normalized error and the standard error of the error decrease as ei increases. Figure 3 confirms that the bound in Theorem 3 is a reasonable guide to performance in practice.

Thus, the more information encoded in spikes, the better a neuron's empirical estimate of its expected reward per spike. The metabolic constraint in Assumption 3 controls the quality of a neuron's empirical estimates of its expected reward.

5 Discussion

The space of possible policies that the cortex as a whole could implement is *vast*: it consists in choosing synaptic weights of millions or billions of neurons each receiving inputs from thousands or tens of thousands of other neurons. However, the space of policies that makes sense biologically is probably much smaller due to biological as well as learning-theoretic considerations. This paper has shown that metabolically constrained reward maximization provides a biologically plausible way for neurons to distinguish useful outputs from those that are not. It follows from our assumptions that spikes:

- are responsible for most of the information communicated by neurons;
- signal when, based on empirical estimates, neurons predict high reward; and
- come equipped with performance guarantees that increase as the metabolic constraint is tightened, and so more information is encoded in spikes.

This suggests that neurons should *privilege spikes* during learning; and indeed there is a large body of experimental evidence that this is exactly what occurs [8, 9, 15, 25]: synaptic plasticity is triggered by pre- and post-synaptic spikes, with the decision to potentiate or depotentiate depending on their precise timing.

An important unanswered question is setting the metabolic constraint ρ. Clearly, if ρ is too low, then neurons will barely fire at all, which is not desirable. Conversely, if ρ is too high then information and empirical reward encoded in spikes, as well as the quality of the empirical estimates, all degrade, which is also to be avoided. In this paper, we have simply highlighted the metabolic constraint ρ as an important lever that neurons may actively manipulate. Finding an optimal value or range of values for ρ is outside the scope of this paper.

Biasing neuronal mechanisms. Q-learning is not a practical learning rule; it simply constructs a giant lookup table. If we think of the set of inputs causing a neuron to fire as a window of varying shape, then the metabolic constraint fixes the size of the window and Q-learning places no other constraints on its shape.

However, if upstream neurons systematically encode information and reward in spikes, then it makes sense to bias shape of downstream neuronal "firing windows" to take the asymmetry between spikes and silences into account. This is exactly what we occurs in cortex. The vast majority of synapses are *excitatory*: the more input spikes neurons receive, the more likely they are to spike themselves.

Cortical neurons are biased toward firing for more spikes, which makes sense if spikes are reliable predictors of future reward. Thus, neuronal "firing windows" are shaped such that if a firing pattern causes a neuron to spike, then so does any firing pattern containing strictly more spikes. Moreover, sufficiently many pre-synaptic spikes will (essentially) always cause a neuron to fire.

Thus, neurons aggregate evidence for high reward (spikes from upstream neurons), and modify their synaptic strengths to maximize their empirical reward. Fine-tuning is necessary since no two neurons have exactly the same connectivity, and therefore no two neurons use the same data to predict global neuromodulatory signals.

Inhibitory (GABA) synapses do not fit this picture. Extending our framework for plasticity to include both types of neurons is challenging. However, we conjecture this incompatibility can be overcome, with inhibition playing a complementary role; possibly centered on imposing sparse activity in the brain [17, 18] and selecting competing neural assemblies and structures.

Global versus overlapping local optimizations. The results in this paper depend on specific assumptions and model choices. A particularly important assumption is that all neurons attempt to maximize the same global neuromodulatory reward signal. In this scenario, since neurons differ in their connectivity, they have access to different subsets of brain activity representing different environmental features and events, and therefore specialize on different sources of reward.

The reality is more complicated with multiple overlapping neuromodulatory signals including dopamine, noradrenaline, acetylcholine and others. Moreover, neurons involved in, say, early visual processing may not require neuromodulatory guidance; rather, they may search for stable invariants over short time frames (hundreds of milliseconds). There is likely a diverse array of reward functions implemented across cortex.

It is thus unclear whether the brain can be accurately described as optimizing a single well-defined reward function. Nevertheless, until decisively shown to be false, we believe the optimization perspective to be a fruitful working hypothesis. Even if it does not apply to the brain as a whole, it may nevertheless provide insight into how populations of neurons in specific brain areas converge on useful behaviors.

A spiking currency. Finally, it is interesting to speculate on an analogy between spikes and paper currency. Money plays many overlapping roles in an economy, including: (i) focusing attention; (ii) stimulating activity; and (iii) providing a quantitative lingua franca for tracking revenues and expenditures. Note that, as for the brain, it is unclear whether an economy as a whole can be reduced to optimizing a single well-defined function.

Spikes may play similar roles in cortex to those of paper currency in an economy. Spikes focus attention: STDP and other proposed learning rules are particularly sensitive to spikes and spike timing. Spikes stimulate activity: input spikes cause output spikes. Finally, spikes leave trails of (Calcium) traces that are used to reinforce and discourage neuronal behaviors in response to neuromodulatory signals.

Neither money nor spikes are intrinsically valuable. Currency can be devalued by inflation. Similarly, the information content and guarantees associated with spikes can be eroded by overpotentiating synapses which reduces their selectivity (potentially leading to epileptic seizures in extreme cases). Regulating the information content of spikes is therefore essential. Assumption 3 provides a simple constraint that can be approximately imposed by regulating synaptic weights. Indeed, there is evidence that one of the functions of sleep is precisely this [11, 14, 28, 30].

Spikes with high information content are valuable because they come with strong guarantees on their estimates. They are therefore worth paying attention to, worth responding to, worth keeping track of, and worth learning from.

References

[1] Attwell, D. and Laughlin, S. B., An energy budget for signaling in the grey matter of the brain, *J Cereb Blood Flow Metab* **21** (2001) 1133–45.

[2] Balduzzi, D. and Besserve, M., Towards a learning-theoretic analysis of spike-timing dependent plasticity, *Advances in Neural Information Processing Systems (NIPS)* (2012).

[3] Balduzzi, D. and Tononi, G., Integrated Information in Discrete Dynamical Systems: Motivation and Theoretical Framework., *PLoS Comput Biol* **4** (2008) e1000091.

[4] Balduzzi, D. and Tononi, G., Qualia: the geometry of integrated information, *PLoS Comput Biol* **5** (2009) e1000462.

[5] Balduzzi, D. and Tononi, G., What can neurons do for their brain? Communicate selectivity with spikes, *Theory in Biosciences* (2012).

[6] Bellman, R., *Dynamic Programming* (Princeton University Press, 1957).

[7] Buesing, L. and Maass, W., Simplified rules and theoretical analysis for information bottleneck optimization and PCA with spiking neurons, in *Adv in Neural Information Processing Systems (NIPS)* (2007).

[8] Dan, Y. and Poo, M.-M., Spike timing-dependent plasticity of neural circuits., *Neuron* **44** (2004) 23–30.

[9] Dan, Y. and Poo, M.-M., Spike timing-dependent plasticity: from synapse to perception., *Physiol Rev* **86** (2006) 1033–1048.

[10] Geman, S., Bienenstock, E., and Doursat, R., Neural Networks and the Bias/Variance Dilemma, *Neural Comp* **4** (1992) 1–58.

[11] Gilestro, G. F., Tononi, G., and Cirelli, C., Widespread changes in synaptic markers as a function of sleep and wakefulness in Drosophila, *Science* **324** (2009) 109–12.

[12] Gottfried, J. A., O'Doherty, J., and Dolan, R. J., Encoding predictive reward value in human amygdala and orbitofrontal cortex, *Science* **301** (2003) 1104–7.

[13] Hasenstaub, A., Otte, S., Callaway, E., and Sejnowski, T. J., Metabolic cost as a unifying principle governing neuronal biophysics, *Proc Natl Acad Sci U S A* **107** (2010) 12329–34.

[14] Maret, S., Faraguna, U., Nelson, A. B., Cirelli, C., and Tononi, G., Sleep and waking modulate spine turnover in the adolescent mouse cortex., *Nat Neurosci* **14** (2011) 1418–1420.

[15] Markram, H., Lübke, J., Frotscher, M., and Sakmann, B., Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps, *Science* **275** (1997) 213–5.

[16] Nere, A., Olcese, U., Balduzzi, D., and Tononi, G., A neuromorphic architecture for object recognition and motion anticipation using burst-STDP, *PLoS One* **7** (2012) e36958.

[17] Olshausen, B. A. and Field, D. J., Sparse coding with an overcomplete basis set: a strategy employed by v1?, *Vision Res* **37** (1997) 3311–25.

[18] Olshausen, B. A. and Field, D. J., Sparse coding of sensory inputs, *Curr Opin Neurobiol* **14** (2004) 481–7.

[19] Ortega, P. A. and Braun, D. A., Information, utility and bounded rationality, *The fourth conference on artificial general intelligence* (2011) 269–274.

[20] Pawlak, V., Wickens, J. R., Kirkwood, A., and Kerr, J. N. D., Timing is not everything: neuromodulation opens the STDP gate, *Front. Syn. Neurosci* **2** (2010).

[21] Rubin, J., Shamir, O., and Tishby, N., Trading Value and Information in MDPs, in *Decision Making with Imperfect Decision Makers*, eds. Guy, T. V., Kárný, M., and Wolpert, D. (Springer, 2011).

[22] Russell, S. and Norvig, P., *Artificial Intelligence: A Modern Approach*, 3rd edn. (Prentice Hall, 2009).

[23] Seldin, Y. and Tishby, N., Multi-Classification by Categorical Features via Clustering, in *Proceedings of the 25th International Conference on Machine Learning* (2008).

[24] Seung, H. S., Learning in Spiking Neural Networks by Reinforcement of Stochastic Synaptic Transmission, *Neuron* **40** (2003).

[25] Song, S., Miller, K. D., and Abbott, L. F., Competitive Hebbian learning through spike-timing-dependent synaptic plasticity, *Nature Neuroscience* **3** (2000).

[26] Sutton, R. S. and Barto, A. G., *Reinforcement Learning: An Introduction* (MIT Press, 1998).

[27] Tishby, N., Pereira, F., and Bialek, W., The information bottleneck method, in *Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing*, eds. Hajek, B. and Sreenivas, R. (1999).

[28] Tononi, G. and Cirelli, C., Sleep and synaptic homeostasis: a hypothesis, *Brain Res. Bull.* **62** (2003) 143–150.

[29] Vapnik, V., *Estimation of Dependencies Based on Empirical Data* (Springer, 1982).

[30] Vyazovskiy, V. V., Olcese, U., Lazimy, Y., Faraguna, U., Esser, S. K., Williams, J. C., Cirelli, C., and Tononi, G., Cortical firing and sleep homeostasis, *Neuron* **63** (2009) 865–78.

[31] Watkins, C. and Dayan, P., Q-learning, *Machine Learning* **8** (1992) 279–292.

A.1   Proof of Theorem 1

Proof. First we show that deterministic policies maximize the information encoded in spikes, then we show that deterministic policies maximize reward per spike.

Deterministic policies maximize effective information subject to $\pi(a_1) \leq \rho$. Observe that

$$ei(\pi, a) = \sum_{s \in \mathcal{S}} \pi(s|a) \log \frac{\pi(s|a)}{P(s)}$$

$$= \sum_{s \in \mathcal{S}} \pi(s|a) \log \frac{\pi(a|s)}{\pi(a)}$$

$$= \underbrace{- \log \pi(a)}_{\text{log frequency of output } a} + \underbrace{\sum_{s \in \mathcal{S}} \pi(s|a) \log \pi(a|s)}_{\text{stochasticity}}.$$

The log-frequency, or surprise, term is nonnegative and the stochasticity term is non-positive. It is easy to see that the stochasticity term is maximized at 0 if and only if output a is chosen deterministically – i.e. $\pi(a|s)$ is either 0 or 1 for all $s \in \mathcal{S}$.

Deterministic policies maximize reward per spike. Suppose there are N situations ordered according to their empirical reward, so that $\widehat{R}(s_1) < \cdots < \widehat{R}(s_N)$. It is clear that a deterministic policy spiking only for the $\rho \cdot N$ policies with highest empirical reward maximizes empirical reward after spiking. $\square$

A.2   Proof of Theorem 2

Proof. Observe that

$$P(s|a_0) = \frac{P(s) - \pi(s|a_1) \cdot \pi(a_1)}{1 - \pi(a_1)}$$

$$= P(s) + \pi(a_1)\big(P(s) - \pi(s|a_1)\big) + O\big(\pi(a_1)^2\big).$$

Thus to first order in $\pi(a_1)$, $\pi(s|a_0)$ is of the form $p + \delta p$ where $\int \delta p = 0$. We can then compute

$$D\big[p + \delta p \,\|\, p\big] = \int (p + \delta p) \log_2 \frac{p + \delta p}{p}$$

$$= \alpha \int (p + \delta p) \frac{\delta p}{p} \left(1 - \frac{\delta p}{2p}\right) + O\big((\delta p)^3\big) = \alpha \int \frac{(\delta p)^2}{2p} + O\big((\delta p)^3\big),$$

11

where $\alpha = \frac{1}{\ln 2}$.

Substituting $p = P(s)$ and $\delta p = \pi(a_1)\big(P(s) - \pi(s|a_1)\big)$ gives

$$D\big[\pi(S|a_0) \,\|\, P(S)\big] = D\big[p + \delta p \,\|\, p\big] + O\big(\pi(a_1)^2\big)$$
$$= \alpha \cdot \pi(a_1)^2 \int \frac{\big(P(s) - \pi(s|a_1)\big)}{2P(s)} + O\big(\pi(a_1)^2\big).$$

Thus $I(S; A) = \pi(a_1) D\big[\pi(S|a_1) \,\|\, P(S)\big] + O\big(\pi(a_1)^2\big).$ $\square$

A.3   Proof of Theorem 3

Occam's razor can be paraphrased to say that the simplest hypothesis should be preferred. Suppose we have a setof hypotheses $\mathcal{H}$ with prior distribution $P(h)$ on $\mathcal{H}$. Let $-\log P(h)$ denote the complexity of hypothesis h. Let $L : \mathcal{X} \times \mathcal{H} \to [0, b]$ be a loss function. Then

Theorem 4 (Occam's razor).
For any data generating distribution on $\mathcal{X}$ and any prior distribution $P(h)$ over $\mathcal{H}$, with a probability greater than $1 - \delta$ over drawing an i.i.d. sample from $\mathcal{X}$ of size T, for all $h \in \mathcal{H}$:

$$\left| L(h) - \widehat{L}(h) \right| \leq b \cdot \sqrt{\frac{-\log P(h) + \log \frac{2}{\delta}}{2T}}.$$

Proof. See [23]. $\square$

Let $\mathcal{H} = \{\pi : \mathcal{S} \to \mathcal{A}\}$ denote the set of deterministic policies, where $\mathcal{A} = \{a_0, a_1\} = \{0, 1\}$. Define loss function

$$L : \big(\mathcal{S} \times \mathcal{N}\big) \times \mathcal{H} \longrightarrow \mathbb{R} : (s, \nu) \times \pi \mapsto R(s, \pi(s), \nu).$$

Further, set probability distribution $P(s, \nu) = P(s) \cdot P(\nu|s)$ on $\mathcal{S} \times \mathcal{N}$. Theorem 4 holds for any sampling distribution. In particular we may use the policy π to restrict samples to situations that cause the neuron to spike to obtain $P(s, \nu|\pi(s) = a_1) = \pi(s|a_1) \cdot P(\nu|s)$. It follows that

$$L(\pi) = R_\pi = \mathbb{E}\Big[R(s, a_1, \nu) \,\Big|\, \pi(s|a_1) \cdot P(\nu|s)\Big] \text{ is the expected reward and}$$

$$\widehat{L}(\pi) = \widehat{R}_\pi = \frac{1}{T_1} \sum_{\{(s_t, \nu_t)|\pi(s_t) = a_1\}} R\big(s_t, a_1, \nu_t\big) \text{ is the empirical reward,}$$

where T_1 is the number of spike produced by the neuron during $[1, T]$.

Theorem 3.
Let $\mathcal{H} \supset \mathcal{H}_k = \big\{\pi : \mathcal{S} \to \mathcal{A} \text{ s.t. } |\pi^{-1}(a_1)| = k\big\}$ denote policies that spike for exactly k situations. Given the setup above, with probability at least $1 - \delta$,

$$\left| R_\pi - \widehat{R}_\pi \right| \leq b \cdot \sqrt{|\mathcal{S}| \cdot \frac{ei_u(\pi, a_1) + 1}{2T_1 \cdot e^{ei_u(\pi, a_1)}} + \frac{\log \frac{2}{\delta}}{2T_1}}$$
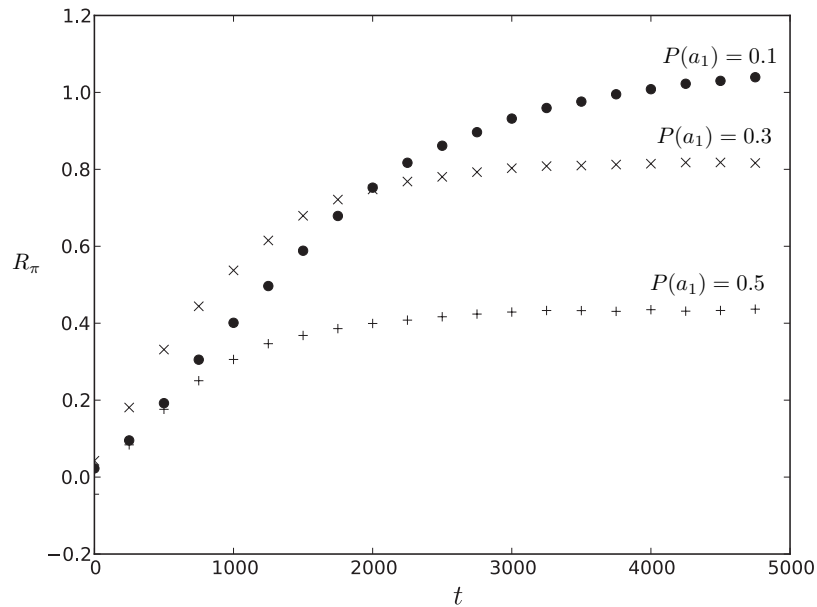
Proof. Let $N = |\mathcal{S}|$ denote the number of possible situations. We put the uniform prior on $\mathcal{H}_k$, so $P(\pi) = \frac{1}{\binom{N}{k}}$. By Stirling's approximation, $\log \binom{N}{k} \leq k \log \big(\frac{N \cdot e}{k}\big)$, and it follows that

$$-\log P(\pi) = \log \binom{N}{k} \leq k \log \frac{N \cdot e}{k} = N \cdot \frac{k}{N} \cdot \Big(\log \frac{N}{k} + 1\Big).$$

The theorem follows since $\pi_u(a_1) = \frac{k}{N}$ and $ei_u(\pi, a_1) = \log \frac{N}{k}$. $\square$
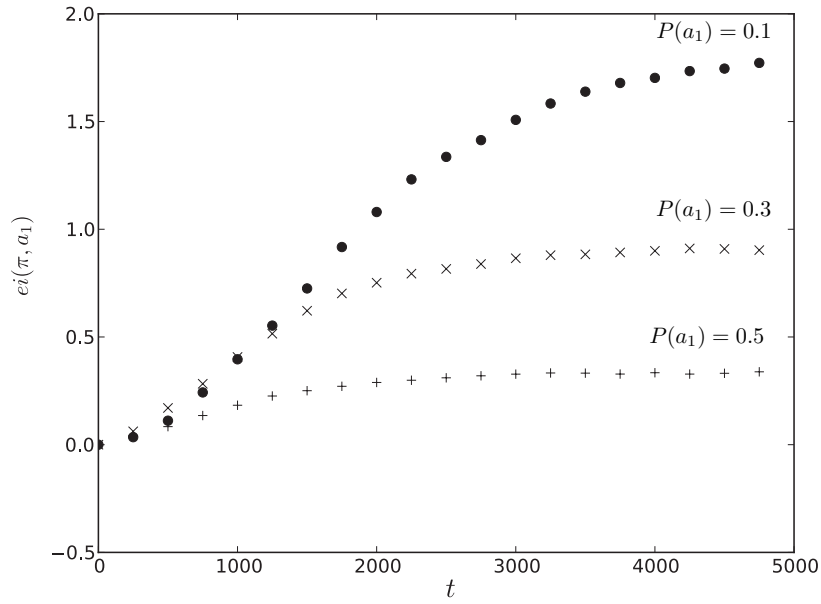
(a) Effective information from spikes

(b) Empirical reward after spiking

Figure 1: **Metabolically constrained Q-learning.** As neurons learn, effective information and empirical reward increase in qualitatively the same way. Tighter metabolic constraints yield both higher effective information and greater rewards.
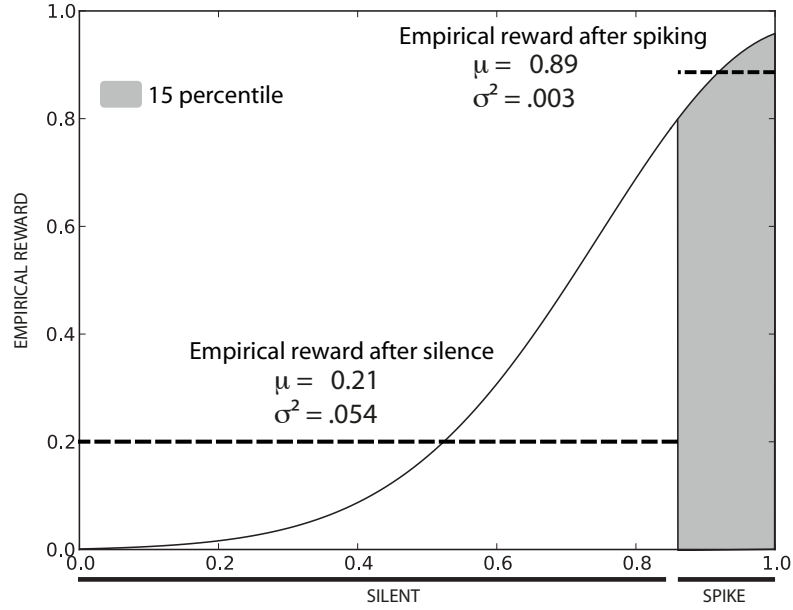
Figure 2: **Concentration of reward.** The x-axis lists situations, ranked by the empirical reward the neuron would receive if it spiked. Situations are grouped into two categories: the top 15%, which cause spikes, and the rest, which do not. The average and variance of the empirical reward in each category is displayed.
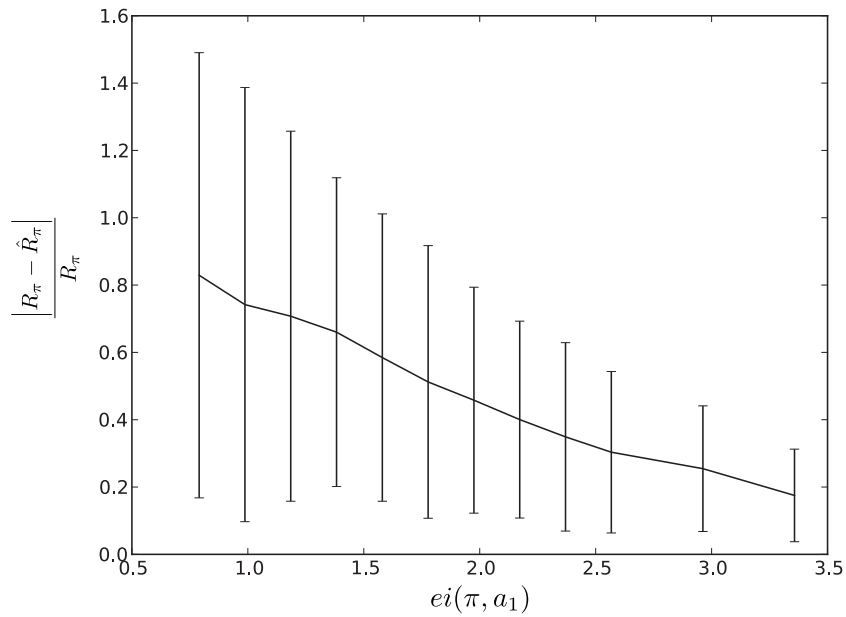

Figure 3: **Empirical versus expected reward.** The normalized difference between expected and empirical reward, plotted against effective information.

14