Title: Neurons learn by predicting future activity.

Authors: Artur Luczak[1] *, Bruce L. McNaughton[1,2], Yoshimasa Kubo[1]

* E-mail: Luczak@uleth.ca

Affiliations:

[1] Canadian Center for Behavioural Neuroscience; University of Lethbridge, AB, Canada.

[2] Center for Neurobiology of Learning and Memory, University of California Irvine, USA.

Significance statement

Understanding how the brain learns may lead to machines with human-like intellectual capacities. Donald Hebb proposed the influential idea that the brain's learning algorithm is based on correlated firing: a.k.a. 'cells that fire together wire together'. However, Hebb's rule and other biologically inspired learning algorithms are still likely missing some important components needed to replicate brain learning mechanisms. Here we provide evidence for a predictive learning rule: a neuron predicts its future activity and adjusts its incoming weights to minimize surprise: i.e. the difference between actual and expected activity. Interestingly, we show that such a rule is equivalent to maximizing the neuron's energy balance, which could be paraphrased as: cells adjust their weights to achieve maximum impact with minimum activity.

Abstract

Plasticity mechanisms in the brain are still not well understood. Here we demonstrate that the ability of a neuron to predict its future activity may provide an effective mechanism for learning in the brain. We show that comparing a neuron's predicted activity with the actual activity provides a useful learning signal for modifying synaptic weights. Interestingly, this predictive learning rule can be derived from a metabolic principle, where neurons need to minimize their own synaptic activity (cost), while maximizing their impact on local blood supply by recruiting other neurons. This reveals an unexpected connection that learning in neural networks could result from simply maximizing the energy balance by each neuron. We validated this predictive learning rule in neural network simulations and in data recorded from awake animals. We found that in the sensory cortex it is indeed possible to predict a neuron's activity ~10-20ms into the future. Moreover, in response to stimuli, cortical neurons changed their firing rate to minimize surprise: i.e. the difference between actual and expected activity, as predicted by our model. Our results also suggest that spontaneous brain activity provides "training data" for neurons to learn to predict cortical dynamics. Thus, this work demonstrates that the ability of a neuron to predict its future inputs could be an important missing element to understand computation in the brain.

Introduction

Artificial neural networks have shown remarkable performance in a multitude of difficult tasks, such as cancer detection, speech recognition, or self-driving cars (LeCun et al., 2015; Wang et al., 2019; Zhao et al., 2019). However, the expertise of such neural networks is restricted to particular domains. This suggests that the human brain, which can solve a wide variety of tasks, e.g. driving a car while discussing development of new cancer tests, may be using more powerful computational algorithms than those used in artificial neural networks. Currently, the best performing neural networks are trained using the backpropagation algorithm (Rumelhart et al., 1986). However, networks that use backpropagation have multiple properties which are difficult to reconcile with biological networks. For example: (1) they require separate phases for sending a bottom-up 'sensory' signal and for receiving a top-down error signal, whereas in the brain, top-down feedback is combined with incoming sensory information; (2) the error signal propagates backwards through 'axons'; (3) artificial neuron models are biologically unrealistic as they generally do not include, for example: spiking activity, internal neuron dynamics, or Dale's law (either excitatory or inhibitory outputs), among many other simplifications. This question of how to bridge the gaps between biological and artificial learning algorithms is a subject of rapidly growing research at the intersection of neuroscience and computer science (Lillicrap et al., 2016; Kuśmierz et al., 2017; Krotov and Hopfield, 2019).

One of the most promising ideas of how backpropagation-like algorithms could be implemented in the brain is based on using temporal difference in neuronal activity to approximate top-down error signals (Ackley et al., 1985; Hinton and McClelland, 1988; Dayan et al., 1995; Hinton et al., 1995; O'Reilly, 1996; Scellier and Bengio, 2017; Lillicrap et al., 2020). A typical example of such algorithms is Contrastive Hebbian Learning (Pineda, 1987; Almeida, 1990; Baldi and Pineda, 1991), which was proven to be equivalent to backpropagation under certain assumptions (Xie and Seung, 2003). Contrastive Hebbian Learning requires networks to have reciprocal connections between hidden and output layers, which allows activity to propagate in both directions (Fig. 1A). The learning consists of two separate phases. First, in the 'free phase', a sample stimulus is continuously presented to the input layer and the activity propagates through the network until the dynamics converge to an equilibrium (activity of each neuron achieves steady-state level). In the second 'clamped phase', in addition to presenting stimulus to the input, the output neurons are also held clamped at values representing stimulus category (e.g.: 0 or 1), and the network is again allowed to converge to an equilibrium. For each neuron, the difference between activity in the clamped ($\hat{x}$) and free ($\check{x}$) phases is used to modify synaptic weights (w) according to the equation (Eq. 1): $\Delta w_{ij} = \alpha(\hat{x}_i\hat{x}_j - \check{x}_i\check{x}_j)$, where i and j are indices of pre- and post- synaptic neurons respectively, and α is a small number representing learning rate. Intuitively, this can be seen as adjusting weights to push each neuron's activity in the free phase, closer to the desired activity represented by the clamped phase. The obvious biological plausibility issue with this algorithm is that it requires the neuron to experience exactly the same stimulus twice in two separate phases, and that the neuron needs to 'remember' its activity from the previous phase.

Here, we solve this problem by combining both activity phases into one, which is inspired by sensory processing in the cortex. For example, in visual areas, when presented with a new picture, there is initially bottom-up driven activity containing mostly visual attributes of the stimulus (e.g. contours), which is then followed by top-down modulation containing more abstract information, e.g., 'this object is a member of category x', or 'this object is novel ' (Suppl. Fig. 1). Accordingly, our algorithm first runs only the initial part of the free phase, which represents

bottom-up stimulus driven activity, and then, after a few steps the network output is clamped, corresponding to top-down modulation.

Results

Synaptic learning rule to minimize prediction error

The novel insight here is that the initial bottom-up activity is enough to allow neurons to predict the steady-state part of the free phase activity, and the mismatch between predicted free phase and clamped phase can then be used as a teaching signal. To implement this idea in our model, for each neuron, activity during twelve initial time steps of the free phase ($\check{x}_{(1)}$,.., $\check{x}_{(12)}$) was used to predict its steady-state activity at time step 120: $\check{x}_{(120)}$ (Fig. 1B). Specifically, first, we presented sample stimuli in free phase to train linear model, such that: $\check{x}_{(120)} \approx \tilde{x} = \lambda_{(1)} * \check{x}_{(1)}, + \ldots + \lambda_{(12)} * \check{x}_{(12)} + b$, where $\tilde{x}$ denote predicted activity, λ and b correspond to coefficients and offset terms of the least-squares model, and terms in brackets correspond to time steps. Next, a new set of stimuli was used for which free phase was ran only for the first 12 steps, and from step 13 the network output was clamped (Fig. 1C). The above least-squares model was then applied to predict free phase steady-state activity for each neuron, and the weights were updated based on the difference between predicted and clamped activity (Methods). Thus, to modify synaptic weights, in Eq. (1) we replaced activity in free phase with predicted activity ($\tilde{x}$): (Eq. 2) $\Delta w_{ij} = \alpha(\hat{x}_i\hat{x}_j - \tilde{x}_i\tilde{x}_j)$. However, the problem is that this equation implies that a neuron needs also to know the predicted activity of all its presynaptic neurons ($\tilde{x}_i$), which could be problematic. To solve this problem, we found that ($\tilde{x}_i$) could be replaced by the actual presynaptic activity in the clamped phase ($\hat{x}_i$), which leads to the following simplified synaptic plasticity rule (Eq. 3): $\Delta w_{ij} = \alpha(\hat{x}_i\hat{x}_j - \hat{x}_i\tilde{x}_j) = \alpha\hat{x}_i(\hat{x}_j - \tilde{x}_j)$. Thus, to modify synaptic weights, a neuron only compares its actual activity ($\hat{x}_j$) with its predicted activity ($\tilde{x}_j$), and applies this difference in proportion to each input contribution ($\hat{x}_i$).

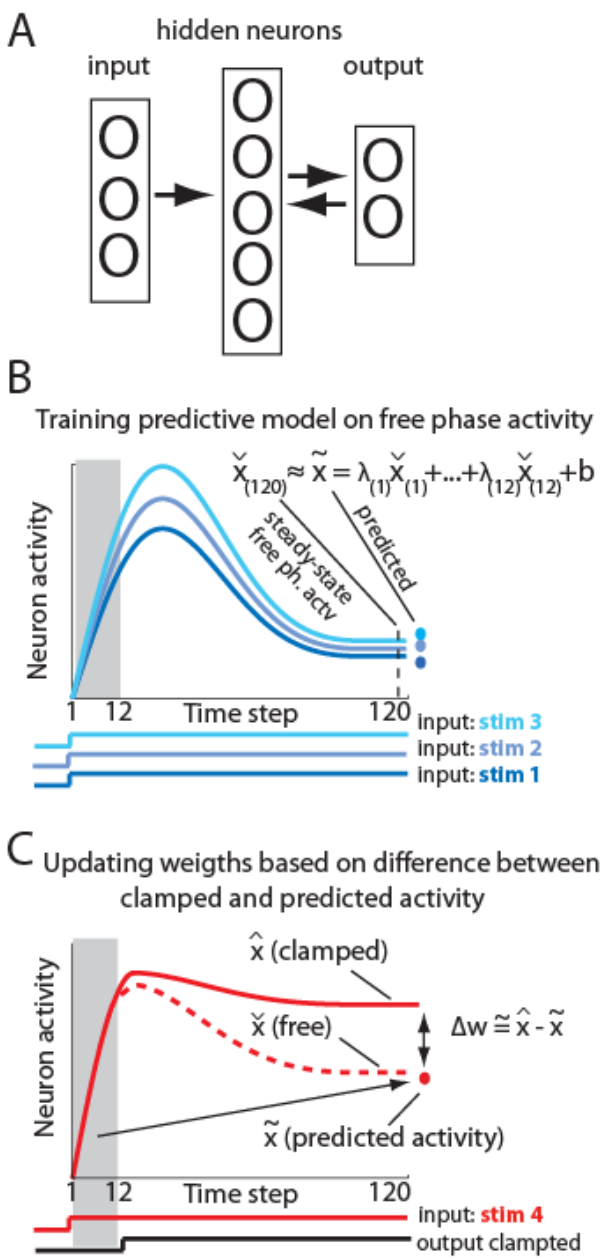

Fig. 1. Basics of the algorithm. (A) Schematic of the network. Note that activity propagates back-and-forth between hidden and output layers. (B) Sample neuron activity in the free phase in response to different stimuli (marked with shades of blue). The free phase responses were used to train a linear model to predict a steady-state activity from the activity at earlier time steps (marked by shaded area; see main text for details). Bottom traces show duration of inputs, and dots represent predicted activity. (C) Activity of a neuron in response to a new stimulus with the network output clamped. Initially the network receives only input signal (free phase), but after a few steps the output signal is also presented (clamped phase; bottom trace). The red dot represents steady-state free phase activity predicted from initial activity (shaded region). For comparison, the dashed line shows a neuron's activity in the free phase if the output is not clamped. Synaptic weights (w) are adjusted in proportion to the difference between steady-state activity in clamped phase ($\hat{x}$) and predicted free phase activity ($\tilde{x}$).

Predictive learning rule validation in neural network simulations

To test if predictive learning rule can be used to solve standard machine learning tasks, we created the following simulation. The neural network had 784 input units, 1000 hidden units, and 10 output units, and it was trained on a hand-written digit recognition task (MNIST(LeCun et al., 1998); Suppl. Fig. 2; Methods). This network achieved 1.9% error rate, which is at human level accuracy on this task (human error rate ~1.5–2.5%; (Simard et al., 1993)). This accuracy is also similar to neural networks with comparable architecture trained with the backpropagation algorithm (LeCun et al., 1998). This demonstrates that the network with predictive learning rule can solve challenging non-linear classification tasks.

To verify that neurons could correctly predict future free phase activity, we took a closer look at sample neurons. Figure 2A illustrates the activity of all 10 output neurons in response to an image of a sample digit after the first epoch of training. During time steps 1-12, only the input signal was presented and the network was running in free phase. At time step 13, the output neurons were clamped, with activity of 9 neurons set to 0 and the activity of one neuron representing correct image class set to 1. For comparison, this figure also shows the activity of the same neurons without clamped outputs (free phase). It illustrates that, after about 50 steps in free phase, the network achieves steady-state, with predicted activity closely matching. When the network is fully trained, it still takes about 50 steps for the network dynamics in free phase to converge to steady-state (Fig. 2B). Note that, although all units initially increase activity at the beginning of the free phase, they later converge close to 0, except the one unit representing the correct category. Again, predictions made from the first 12 steps during free phase closely matched the actual steady-state activity. The hidden units also converged to steady-state after about 50 steps. Figure 2C illustrates the response of one representative hidden neuron to 5 sample stimuli. Because hidden units experience the clamped signal only indirectly, through synapses from output neurons, their steady-state activity is not bound to converge only to 0 or 1, as in the case of output neurons. Actual and predicted steady-state activity for hidden neurons is presented in Figure 2D. The average correlation coefficient between predicted and actual free phase activity was $R=1\pm0.0001SD$ (averaged across 1000 hidden neurons in response to 200 randomly selected test images). Note that for all predictions we used a cross-validation approach, where we trained a predictive model for each neuron on a subset of the data and applied it to new examples, which were then used for updating weights (Methods). Thus, neurons were able successfully to generalize their predictions to new unseen stimuli. The network error rate for the training and test dataset is shown in Fig. 2E. This demonstrates that the predictive learning rule worked well, and each neuron accurately predicted its future activity.

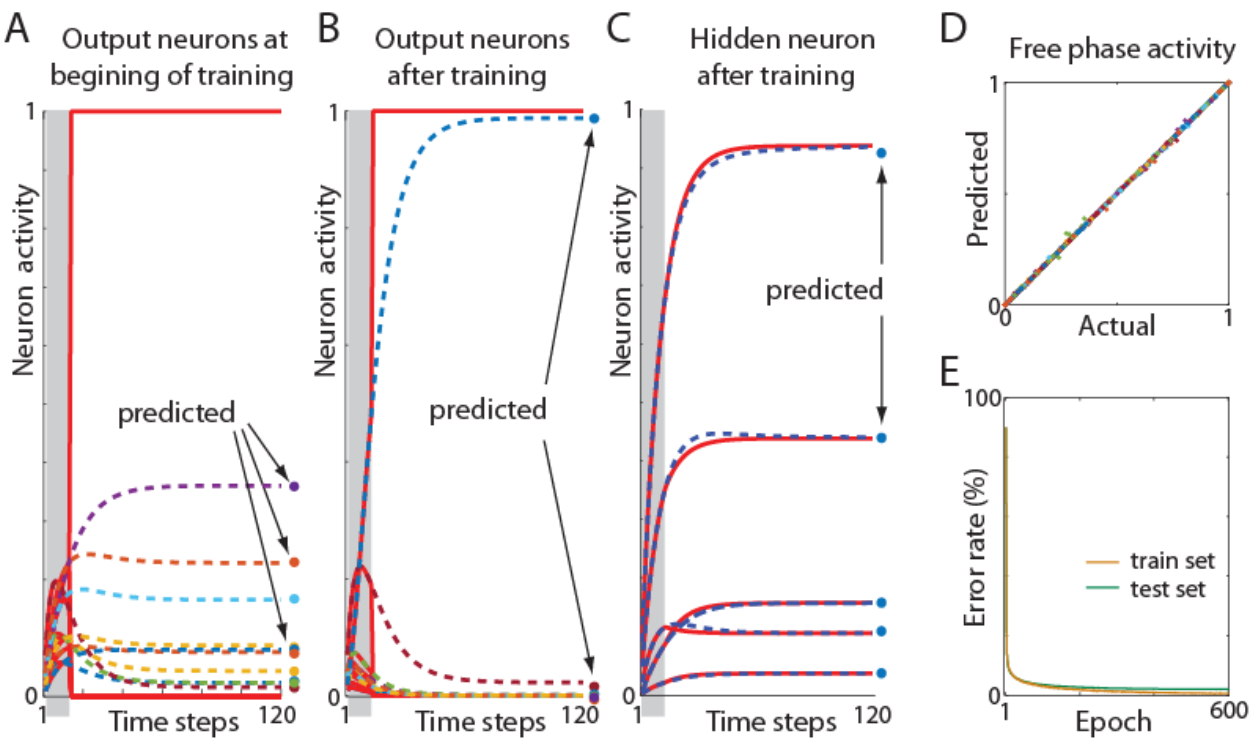

Fig. 2. Neuron prediction of expected activity. (A) Activity of 10 output neurons in response to a sample stimulus at the beginning of the network training. Gray area indicates the extent of the free phase (time steps 1-12). Solid red lines show activity of the neurons clamped at step 13. For comparison, dashed lines represent free phase activity if output neurons had not been clamped. Dots show predicted steady-state activity in free phase based on initial activity (from steps 1-12). (B) Activity of the same neurons after network training. Note that free phase and predicted activity converged to desired clamped activity. (C) Activity of a representative neuron in the hidden layer in response to 5 different stimuli after network training. Solid and dashed lines represent clamped and free phase respectively, and dots show predicted activity. (D) Predicted vs. actual free phase activity. For visualization clarity only every tenth hidden neuron out of 1000 is shown, in response to 20 sample images. Different colors represent different neurons, but some neurons may share the same color due to the limited number of colors. Distribution of points along the diagonal shows that predictions are accurate. (E) Decrease in error rate across training epochs. Yellow and green lines denote learning curves for training and test data set respectively. Note that in each epoch we only used 2% out of 60,000 training examples.

Biologically motivated network architectures

We also tested the predictive learning rule in multiple other network architectures, which were designed to reflect additional aspects of biological neuronal networks. First, we introduced a constraint that 80% of the hidden neurons were excitatory, and the remaining 20% had only inhibitory outputs. This follows observations that biological neurons release either excitatory or inhibitory neurotransmitters, not both (Dale's law (Eccles et al., 1954)), and that about 80% of cortical neurons are excitatory. The network with this architecture achieved an error rate of 2.66%

(Suppl. Fig. 3A), which again is comparable with human error rate on this task (Simard et al., 1993). We also tested our algorithm in a network with two hidden layers, in a network without symmetric weights, and in a network with all-to-all connectivity within the hidden layer. We found that all of those modified networks similarly converged towards the solution (Suppl. Fig. 3B & 4). We also implemented predictive learning rule in a network with spiking neurons, which again achieved a similar error rate of 2.46% (Suppl. Fig. 5). Altogether, this shows that our predictive learning rule performs well in a variety of biologically motivated network architectures.

Predictive learning rule validation in awake animals

To test whether real neurons could also predict their future activity, we analyzed neuronal recordings from the auditory cortex in awake rats (Methods). As stimuli we presented 6 tones, each 1s long and interspersed by 1s of silence, repeated continuously for over 20 minutes (Suppl. Materials). For each of the 6 tones we calculated separately average onset and offset response, giving us 12 different activity profiles for each neuron (Fig. 3A). For each stimulus, the activity in the time window 15-25ms was used to predict average future activity within the 30-40ms window. We used 12-fold cross-validation, where responses from 11 stimuli were used to train the least-square model, which was then applied to predict neuron activity for the 1 remaining stimulus. This procedure was repeated 12 times for each neuron. The average correlation coefficient between actual and predicted activity was $R = 0.36 \pm 0.05$ SEM (averaged across 55 cells from 4 animals, Fig. 3B). Distribution of correlations coefficients for individual neurons were significantly different from 0 (t-test $p<0.0001$; insert in Fig. 3B). This shows that neurons have predictable dynamics, and from an initial neuronal response its future activity can be estimated.

Repeated presentation of stimuli over tens of minutes also induced long-term changes in neuronal firing rates (Bermudez Contreras et al., 2013), similar as in perceptual learning. Importantly, based on our model, it was possible to infer which individual neurons will increase or decrease their firing rate. To explain it, first let's look at neural network simulations in Fig. 3C. It shows that, for a neuron, the average change in activity from one learning epoch to the next, depends on the difference between clamped (actual) activity and predicted (expected) activity, in the previous learning epoch (Fig. 3C; correlation coefficient $R = 0.35$, $p<0.0001$; Suppl. Materials). Similarly, for cortical neurons, we found that the change in firing rate from the 1st to the 2nd half of the experiment was positively correlated with differences between evoked and predicted activity during the 1st half of experiment ($R = 0.58$, $p<0.0001$; Fig. 3D, Suppl. Materials). This could be understood in terms of Eq. 3, where if actual activity is higher than predicted, then synaptic weights are increased, thus leading to higher activity of that neuron in the next epoch. Therefore, similar behavior of artificial and cortical neurons, where firing rate changes to minimize 'surprise': difference between actual and predicted activity, provides a strong evidence in support of the predictive learning rule presented here.

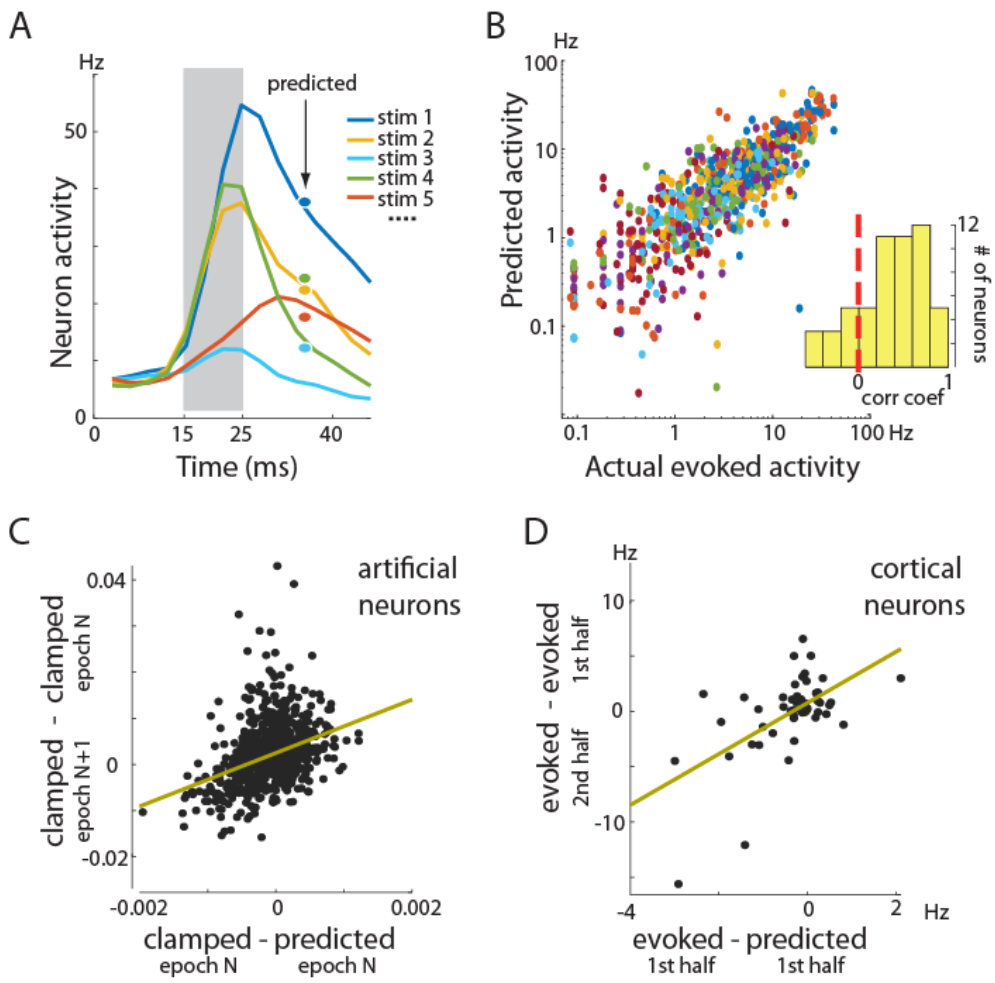

Fig. 3. **Predicting future activity of cortical neurons**. (A) Response of a representative neuron to different stimuli. For visualization only 5 out of 12 responses are shown. Gray area indicates the time window which was used to predict future activity. Dots show predicted average activity in the 30-40ms time window. Colors correspond to different stimuli. (B) Actual vs. predicted activity for 55 cells from 4 animals in response to 12 stimuli. Different colors represent different neurons, but some neurons may share the same color due to the limited number of colors. Insert: histogram of correlation coefficients for individual neurons. Skewness of the distribution to the right shows that for most neurons the correlation between actual and predicted response was positive. (C) Average change in clamped steady-state activity between 2 consecutive learning epochs in our network model. This change relates to the difference between clamped and predicted activity in the earlier epoch (*N*=7; Suppl. Materials). Each dot represents one neuron. Regression line is shown in yellow. (D) Average change in firing rate between 1st and 2nd half of our experiment with repetitive auditory stimulation. This firing rate change correlates to the difference between stimulus evoked and predicted activity during the 1st half of the experiment (Suppl. Materials). Each dot represents the activity of one neuron averaged across stimuli. The similar behavior of cortical and artificial neurons suggests that both may be using essentially the same learning rule. Note that, although the results in panel A, showing that firing rate is consistent from early to later phase

of response, may be expected, the results in panel D, that neurons change firing rate depending on value of predicted activity, provides novel insight about neuronal plasticity.

Deriving predictive model parameters from spontaneous activity

Next, we tested whether spontaneous brain activity could also be used to predict neuronal dynamics during stimulus presentation. Spontaneous activity, such as during sleep, is defined as an activity not directly caused by any external stimuli. However, there are many similarities between spontaneous and stimulus evoked activity (Kenet et al., 2003; MacLean et al., 2005; Berkes et al., 2011; Luczak and MacLean, 2012). For example, spontaneous activity is composed of ~50-300 ms long population bursts called packets, which resemble stimulus evoked patterns (Luczak et al., 2015). This is illustrated in Figure 4A, where spontaneous activity packets in the auditory cortex are visible before sound presentation (Luczak et al., 2009; Luczak et al., 2013). In our experiments, each 1s long tone presentation was interspersed with 1s of silence, and the activity during 200-1000 ms after each tone was considered as spontaneous (animals were in soundproof chamber; Suppl. Materials). The individual spontaneous packets were extracted to estimate neuronal dynamics (Methods). Then the spontaneous packets were divided into 10 groups based on similarity in PCA space (Suppl. Materials), and for each neuron we calculated its average activity in each group (Fig. 4B). As in previous analyses in Fig 3A, the initial activity in time window 5-25ms was used to derive the least-square model to predict future spontaneous activity in the 30-40ms time window (Suppl. Materials). This least-square model was then applied to predict future evoked responses from initial evoked activity for all 12 stimuli. Figure 4C shows actual vs predicted evoked activity for all neurons and stimuli (correlation coefficient $R = 0.2 \pm 0.05$ SEM, averaged over 40 cells from 4 animals; the insert shows the distribution of correlation coefficients of individual neurons; $p=0.0008$, t-test). Spontaneous brain activity is estimated to account for over 90% of brain energy consumption (Raichle and Mintun, 2006), however the function of this activity still is a mystery. The foregoing results offer a new insight: because neuronal dynamics during spontaneous activity is similar to evoked activity, thus spontaneous activity can provide “training data” for neurons to build a predictive model.

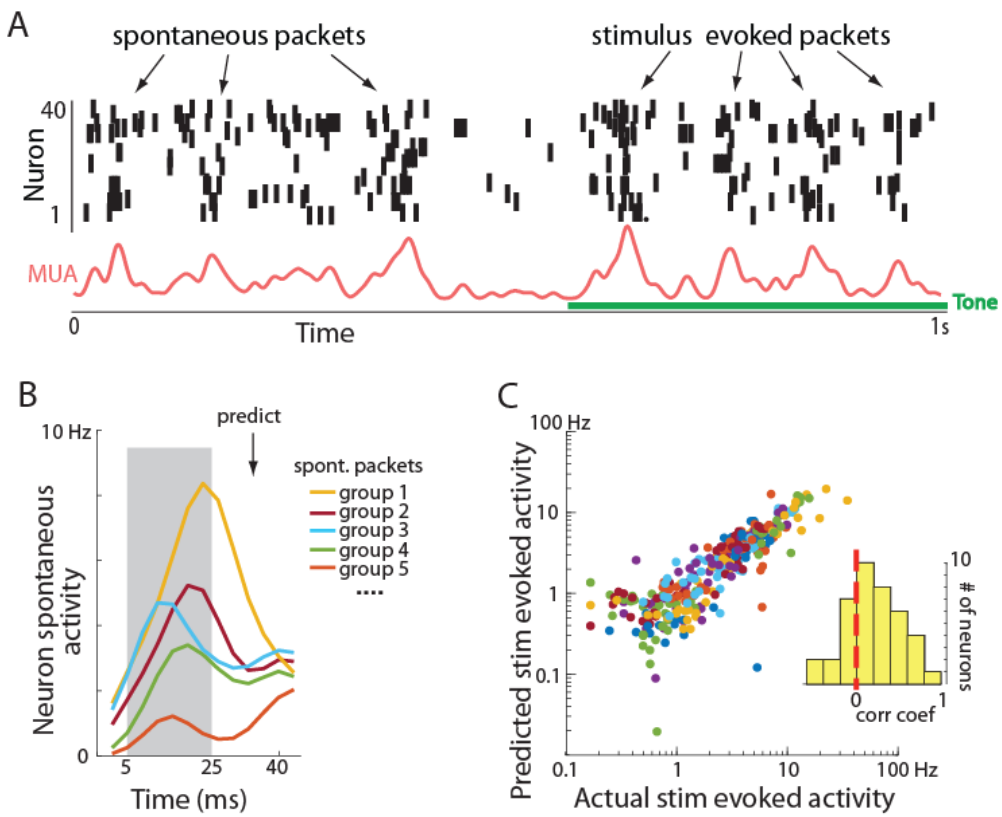

Fig 4 Predicting stimulus evoked responses from spontaneous activity dynamics. (A) Sample spiking activity in the auditory cortex before and during tone presentation. Note that spontaneous activity is not continuous, but rather composed of bursts called packets which are similar to tone evoked packets. The red trace shows smoothed multiunit activity: summed activity of all neurons (adopted with permission from (Luczak et al., 2013)). (B) Spontaneous packets were divided into 10 groups based on population activity patterns. Activity of a single neuron in 5 different spontaneous packet groups is shown. Gray area indicates the time window used for predicting future average activity within the 30-40ms time window (marked by arrow). This predictive model derived from spontaneous activity was then applied to predict future evoked activity based on initial evoked response. (C) Actual vs. predicted tone evoked activity. Plot convention is the same as in Fig. 3B. Skewness of the histogram to the right shows that for most neurons the evoked dynamics can be estimated based on spontaneous neuron's activity.

Predictive learning rule derivation by maximizing neuron energy balance.

Importantly, predictive learning rule in Eq. 3: $\Delta w_{ij} = \alpha \hat{x}_i(\hat{x}_j - \tilde{x}_j)$ is not an *ad hoc* algorithm devised to solve a computational problem, but is a form of learning rule that arises naturally as a consequence of minimizing a metabolic cost by a neuron. Most of the energy consumed by a neuron is for electrical activity, with synaptic potentials accounting for ~50% and action potential for ~20% of ATP used (Harris et al., 2012). Using a simplified linear model of neuronal activity, this energy consumption for a neuron j can be expressed as $-b_1(\sum_i w_{ij}x_i)^{\beta_1}$, where x_i represents the activity of pre-synaptic neuron i, w represents synaptic weights, b_1 is a

constant to match energy units, and β_1 describes a non-linear relation between neuron activity and energy usage, which is estimated to be between 1.7 - 4.8 (Devor et al., 2003). The remaining ~30% of neuron energy is consumed on housekeeping functions, which could be represented by a constant $-\varepsilon$. On the other hand, the increase in neuronal population activity also increases local blood flow leading to more glucose and oxygen entering a neuron (see review on neurovascular coupling: (Sokoloff, 2008)). This activity dependent energy supply can be expressed as: $+b_2(\sum_k x_k)^{\beta_2}$, where x_k represents spiking activity of neuron k from a local population of K neurons ($k \in \{1, \dots, j, \dots K\}$), with activity of neuron j: $x_j = \sum_i w_{ij} x_i$; b_2 is a constant and β_2 reflects the exponential relation between activity and blood volume increase, which is estimated to be in range β_2: 1.7-2.7 (Devor et al., 2003). Putting all those terms together, the energy balance of a neuron j could be expressed as

(Eq. 4): $E_j = -\varepsilon - b_1(\sum_i w_{ij} x_i)^{\beta_1} + b_2(\sum_k x_k)^{\beta_2}$.

This formulation shows that to maximize energy balance, a neuron has to minimize its electrical activity (be active as little as possible), but at the same time, it should maximize its impact on other neurons' activities to increase blood supply (be active as much as possible). Thus, weights have to be adjusted to strike a balance between two opposing demands: maximizing the neuron's downstream impact and minimizing its own activity. This energy objective could be paraphrased as: *achieving the cell's maximum impact with minimum activity (cost to itself)*. We can calculate such required changes in synaptic weights Δw that will maximize neuron's energy E_j by using gradient ascent method. For that, we need to calculate derivative of E_j with respect to w_{ij}

(Eq. 5): $\Delta w_{ij} = \frac{\partial E_j}{\partial w_{ij}} = 0 - x_i \beta_1 b_1 (\sum_i w_{ij} x_i)^{\beta_1 - 1} + x_i \beta_2 b_2 (\sum_k x_k)^{\beta_2 - 1}$.

If we denote population activity as: $\bar{x} = \sum_k x_k$, and considering that $\sum_i w_{ij} x_i = x_j$, then after switching order of terms we obtain (Eq. 6): $\Delta w_{ij} = x_i(\beta_2 b_2 \bar{x}^{\beta_2 - 1} - \beta_1 b_1 x_j^{\beta_1 - 1})$.

If we also denote that $\alpha_1 = \beta_1 b_1$ and $\alpha_2 = \frac{\beta_2 b_2}{\alpha_1}$, then we can take α_1 in front of brackets; and considering that β_1 and $\beta_2 > 1.7$ (see (Devor et al., 2003)), we may assume that $\beta_1 \approx 2$ and $\beta_2 \approx 2$, which simplifies Eq. 6 to

(Eq. 7): $\Delta w_{ij} = \alpha_1 x_i(\alpha_2 \bar{x} - x_j)$.

Note that Eq. 7 has the same form as the predictive learning rule in (Eq. 3): $\Delta w_{ij} = \alpha \hat{x}_i (\hat{x}_j - \tilde{x}_j)$.

Moreover, if we assume that neuron adapts to maximize energy balance *in the future*, then Eq. 7, changes to (Eq. S7): $\Delta w_{ij} = \alpha_3 x_{i,t} (\alpha_4 \tilde{\bar{x}} - \tilde{x}_j)$, where $\tilde{\bar{x}}$ represents population recurrent activity, which can be thought of as top-down modulation, similar to $\hat{x}$. Also note that activity of neuron j: x_j, becomes future predicted activity: $\tilde{x}_j$ (see Suppl. Materials for details of derivation). Thus, this analysis shows that the best strategy for a neuron to maximize future energy resources requires predicting its future activity. Altogether this reveals an unexpected connection, that learning in neural networks could result from simply maximizing energy balance by each neuron.

Discussion

Here we present theoretical, computational and biological evidence that the basic principle underlying single neuron learning may rely on minimizing future surprise: a difference between actual and predicted activity. Thus, a single neuron is not only performing summation of its inputs, but it also predicts the expected future, which we propose is a crucial component of the brain's learning mechanism. Note that a single neuron has complexity similar to single cell organisms, which were shown to have 'intelligent' adaptive behaviors, including predicting consequences of its action in order to navigate toward food and away from danger (Tero et al., 2010; Kaiser, 2013; Boisseau et al., 2016). This suggests that typical neuronal models used in machine learning may be too simplistic to account for the essential computational properties of biological neurons. Our work suggests that a predictive mechanism may be an important computational element within neurons, which could be crucial to understand learning mechanisms in the brain.

There are multiple lines of evidence that the brain operates as a predictive system (Bar, 2007; Clark, 2013; Schwartenbeck et al., 2015; O'Reilly et al., 2017; Buzsáki; Gordon et al., 2019). However, it remains controversial how exactly predictive coding could be implemented in the brain (Lillicrap et al., 2020). Most of proposed mechanisms involve specially designed neuronal circuits to allow for comparing expected and actual activity (Rao and Ballard, 2005; Bastos et al., 2012; Whittington and Bogacz, 2017; Sacramento et al., 2018). Although those models are to some degree biologically motivated, they require precise network configuration, which could be difficult to achieve, considering that the brain connectivity is highly variable within and between brain areas. Here we propose that this problem can be solved by implementing predictions within a single neuron. Biological neurons have a variety of cellular mechanisms which operate on time scales of 1~100ms suitable for implementing predictions (Gutfreund et al., 1995; Stuart and Sakmann, 1995; Koch et al., 1996; Larkum et al., 1999; Ha and Cheong, 2017). For instance, it was proposed that depolarization of the soma by basal dendrites in pyramidal cells could serve as an expected signal for comparison with top-down modulation arriving from apical dendrite (Hawkins and Ahmad, 2016; Guerguiev et al., 2017). The other interesting biological aspect of our model is that it belongs to the category of energy-based models, for which it was shown that synaptic update rules are consistent with spike-timing-dependent plasticity (Bengio et al., 2017). All of this demonstrates compatibility of our model with neurophysiology.

Our work also suggests that packets could be basic units of information processing in the brain. It is well established that sensory stimuli evoke coordinated bursts (packets) of neuronal activity lasting from tens to hundreds of ms. We call such population bursts packets, because they have stereotypical structure, with neurons active at the beginning conveying bottom-up sensory information (e.g. this is a face), and later in the packet neurons represent additional higher order information (e.g. this is a happy face of that particular friend)(Sugase et al., 1999). Also the later part of the packet can encode if there is discrepancy with expectation (e.g. this is a novel stimulus (Sams et al., 1985; Freedman et al., 2006); Suppl. Fig. 1). This is likely because only the later part of the packet can receive top-down modulation after information about that stimulus is exchanged between other brain areas (Roland et al., 2006; Xu et al., 2007). Thus, our work suggests that the initial part of the packet can be used to infer what the rest of the brain may 'think' about this stimulus, and the difference from this expectation can be used as a learning mechanism to modify synaptic connections. This could be the reason why, for example, we cannot process visual information faster than ~20 frames/s, as only after evaluating if a given image is consistent with expectation, can the next image be processed by the next packet, which takes ~50ms. Our

predictive learning rule thus implies, that sensory information is processed in discrete units and each packet may represent an elementary unit of perception.

Methods

Neural Network

The code to reproduce our network with all the implementation details is available at https://github.com/ykubo82/bioCHL . Briefly, the base network has the architecture: 784-1000-10 with sigmoidal units, and with symmetric connections (see Suppl. Fig. 3-5 for more biologically plausible network architectures which we also tested). To accelerate training, we used AdaGrad (Duchi et al., 2011) , and we applied a learning rate of 0.03 to the hidden layer and 0.02 for the output layer. In the standard implementation of Contrastive Hebbian Learning, the learning rate α in Eq. 1 is also multiplied by a small number (~0.1) for all top-down connections (e.g. from output to hidden layer) (Detorakis et al., 2019). This different treatment of feed-forward and feedback connections could be biologically questionable as cortical circuits are highly recurrent. Therefore, to make our predictive learning rule more biologically plausible we discarded this feedback gain factor as we wanted our network to learn by itself what should be the contribution of each input, as described in Eq. 3.

Future activity prediction

For all the predictions we used a cross-validation approach. Specifically, in each training cycle we ran free phase on 490 examples, which were used to derive least-squares model for each neuron to predict its future activity at time step 120 ($\tilde{x}_{(120)}$), from its initial activity at steps 1-12 ($\breve{x}_{(1)}$,.., $\breve{x}_{(12)}$). This can be expressed as: $\tilde{x}_{(120)} = \lambda_{(1)} * \breve{x}_{(1)}, + \ldots + \lambda_{(12)} * \breve{x}_{(12)} + b$, where terms in brackets correspond to time steps, and λ and *b* correspond to coefficients and offset terms found by least-squares method. Next, 10 new examples were taken, for which free phase was ran only for 12 steps, then the above derived least-squares model was applied to predict free phase steady-state activity for each of 10 examples. From step 13 the network output was clamped. The weights were updated based on the difference between predicted and clamped activity calculated only from those 10 new examples. This process was repeated 120 times in each training epoch. Moreover, the MNIST dataset has 60,000 examples which we used for the above described training, and 10,000 additional examples which were only used for testing. For all plots in Figure 2 we only used test examples which network never saw during training. This demonstrates that each neuron can accurately predict its future activity even for novel stimuli which were never presented before.

Surgery, recording and neuronal data

The experimental procedures for the awake, head-fixed experiment have been previously described (Luczak et al., 2009; Luczak et al., 2013) and were approved by the Rutgers University Animal Care and Use Committee, and conformed to NIH Guidelines on the Care and Use of Laboratory Animals. Briefly, a headpost was implanted on the skull of four Sprague-Dawley male rats (300-500g) under ketamine-xylazine anesthesia, and a craniotomy was performed above the auditory cortex and covered with wax and dental acrylic. After recovery the animal was trained for 6-8 days to remain motionless in the restraining apparatus. On the day of the surgery, the animal

was briefly anesthetized with isoflurane, the dura was resected, and after a recovery period, recording began. For recording we used silicon microelectrodes (Neuronexus technologies, Ann Arbor MI) consisting of 8 or 4 shanks spaced by 200µm, with a tetrode recording configuration on each shank. Electrodes were inserted in layer V in the primary auditory cortex. Units were isolated by a semiautomatic algorithm (klustakwik.sourceforge.net) followed by manual clustering (klusters.sourceforge.net)(Harris et al., 2000). Only neurons with average stimulus evoked firing rates higher than 3 SD above pre-stimulus baseline were used in analysis, resulting in 9, 12, 12, and 22 neurons from each rat. For predicting evoked activity from spontaneous, we also required that neurons must have mean firing rate during spontaneous packets above said threshold which reduced the number of neurons to 40. The spontaneous packet onsets were identified from the spiking activity of all recorded cells as the time of the first spike marking a transition from a period of global silence (30 ms with at most one spike from any cell) to a period of activity (60 ms with at least 15 spikes from any cells), as described before in (Luczak et al., 2007; Luczak et al., 2009). The data presented here are available from the corresponding author upon reasonable request.

Acknowledgments: This work was supported by Compute Canada, NSERC and CIHR grants to AL, and DARPA #HR0011-18-2-0021 and NIH #MH125557 to BLM.. We thank Karim Ali, Lukas Grasse, Mikhail Klassen and Reza Torabi for help, and we thank Peter Bartho for sharing data.

Author contributions: AL conceived the project, analyzed data, performed computer simulations and wrote the manuscript; BLM engaged in theoretical discussions and commented extensively on the manuscript. YK performed computer simulations and contributed to writing manuscript.

Competing interests: The authors declare no competing interests.

References

Ackley DH, Hinton GE, Sejnowski TJ (1985) A learning algorithm for Boltzmann machines. Cognitive science 9:147-169.

Almeida LB (1990) A learning rule for asynchronous perceptrons with feedback in a combinatorial environment. In: Artificial neural networks: concept learning, pp 102-111.

Baldi P, Pineda F (1991) Contrastive learning and neural oscillations. Neural computation 3:526-545.

Bar M (2007) The proactive brain: using analogies and associations to generate predictions. Trends in cognitive sciences 11:280-289.

Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ (2012) Canonical microcircuits for predictive coding. Neuron 76:695-711.

Bengio Y, Mesnard T, Fischer A, Zhang S, Wu Y (2017) STDP-compatible approximation of backpropagation in an energy-based model. Neural computation 29:555-577.

Berkes P, Orbán G, Lengyel M, Fiser J (2011) Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science 331:83-87.

Bermudez Contreras EJ, Schjetnan AGP, Muhammad A, Bartho P, McNaughton BL, Kolb B, Gruber AJ, Luczak A (2013) Formation and reverberation of sequential neural activity patterns evoked by sensory stimulation are enhanced during cortical desynchronization. Neuron 79:555-566.

Boisseau RP, Vogel D, Dussutour A (2016) Habituation in non-neural organisms: evidence from slime moulds. Proceedings of the Royal Society B: Biological Sciences 283:20160446.
Buzsáki G (2019) The brain from inside out: Oxford University Press.
Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and brain sciences 36:181-204.
Dayan P, Hinton GE, Neal RM, Zemel RS (1995) The helmholtz machine. Neural computation 7:889-904.
Detorakis G, Bartley T, Neftci E (2019) Contrastive Hebbian learning with random feedback weights. Neural Networks 114:1-14.
Devor A, Dunn AK, Andermann ML, Ulbert I, Boas DA, Dale AM (2003) Coupling of total hemoglobin concentration, oxygenation, and neural activity in rat somatosensory cortex. Neuron 39:353-359.
Duchi J, Hazan E, Singer Y (2011) Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. Journal of machine learning research 12:2121-2159.
Eccles JC, Fatt P, Koketsu K (1954) Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurones. The Journal of physiology 126:524.
Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2006) Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. Cerebral Cortex 16:1631-1644.
Gordon N, Tsuchiya N, Koenig-Robert R, Hohwy J (2019) Expectation and attention increase the integration of top-down and bottom-up signals in perception through different pathways. PLoS biology 17:e3000233.
Guerguiev J, Lillicrap TP, Richards BA (2017) Towards deep learning with segregated dendrites. Elife 6:e22901.
Gutfreund Y, Yarom Y, Segev I (1995) Subthreshold Oscillations and Resonant-Frequency in Guinea-Pig Cortical-Neurons - Physiology and Modeling. J Physiol-London 483:621-640.
Ha GE, Cheong E (2017) Spike Frequency Adaptation in Neurons of the Central Nervous System. Exp Neurobiol 26:179-185.
Harris JJ, Jolivet R, Attwell D (2012) Synaptic energy use and supply. Neuron 75:762-777.
Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsáki G (2000) Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. Journal of neurophysiology 84:401-414.
Hawkins J, Ahmad S (2016) Why neurons have thousands of synapses, a theory of sequence memory in neocortex. Frontiers in neural circuits 10:23.
Hinton GE, McClelland JL (1988) Learning representations by recirculation. In: Neural information processing systems, pp 358-366.
Hinton GE, Dayan P, Frey BJ, Neal RM (1995) The" wake-sleep" algorithm for unsupervised neural networks. Science 268:1158-1161.
Kaiser AD (2013) Are myxobacteria intelligent? Frontiers in microbiology 4:335.
Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. Nature 425:954-956.
Koch C, Rapp M, Segev I (1996) A brief history of time (constants). Cerebral Cortex 6:93-101.
Krotov D, Hopfield JJ (2019) Unsupervised learning by competing hidden units. Proceedings of the National Academy of Sciences 116:7723-7731.
Kuśmierz Ł, Isomura T, Toyoizumi T (2017) Learning with three factors: modulating Hebbian plasticity with errors. Current opinion in neurobiology 46:170-177.

Larkum ME, Zhu JJ, Sakmann B (1999) A new cellular mechanism for coupling inputs arriving at different cortical layers. Nature 398:338-341.
LeCun Y, Bengio Y, Hinton G (2015) Deep learning. Nature 521:436.
LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. Proceedings of the IEEE 86:2278-2324.
Lillicrap TP, Cownden D, Tweed DB, Akerman CJ (2016) Random synaptic feedback weights support error backpropagation for deep learning. Nature communications 7:1-10.
Lillicrap TP, Santoro A, Marris L, Akerman CJ, Hinton G (2020) Backpropagation and the brain. Nature Reviews Neuroscience:1-12.
Luczak A, MacLean JN (2012) Default activity patterns at the neocortical microcircuit level. Frontiers in integrative neuroscience 6.
Luczak A, Barthó P, Harris KD (2009) Spontaneous events outline the realm of possible sensory responses in neocortical populations. Neuron 62:413-425.
Luczak A, Bartho P, Harris KD (2013) Gating of sensory input by spontaneous cortical activity. The Journal of Neuroscience 33:1684-1695.
Luczak A, McNaughton BL, Harris KD (2015) Packet-based communication in the cortex. Nature Reviews Neuroscience.
Luczak A, Barthó P, Marguet SL, Buzsáki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. Proceedings of the National Academy of Sciences 104:347-352.
MacLean JN, Watson BO, Aaron GB, Yuste R (2005) Internal dynamics determine the cortical response to thalamic stimulation. Neuron 48:811-823.
O'Reilly RC (1996) Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. Neural computation 8:895-938.
O'Reilly RC, Wyatte DR, Rohrlich J (2017) Deep predictive learning: a comprehensive model of three visual streams. arXiv preprint arXiv:170904654.
Pineda FJ (1987) Generalization of back-propagation to recurrent neural networks. Physical review letters 59:2229.
Raichle ME, Mintun MA (2006) Brain work and brain imaging. Annu Rev Neurosci 29:449-476.
Rao RP, Ballard DH (2005) Probabilistic models of attention based on iconic representations and predictive coding. In: Neurobiology of attention, pp 553-561: Elsevier.
Roland PE, Hanazawa A, Undeman C, Eriksson D, Tompa T, Nakamura H, Valentiniene S, Ahmed B (2006) Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. Proceedings of the National Academy of Sciences 103:12586-12591.
Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. Nature 323:533-536.
Sacramento J, Costa RP, Bengio Y, Senn W (2018) Dendritic cortical microcircuits approximate the backpropagation algorithm. In: Advances in neural information processing systems, pp 8721-8732.
Sams M, Paavilainen P, Alho K, Naatanen R (1985) Auditory frequency discrimination and event-related potentials. Electroencephalography and clinical neurophysiology 62:437-448.
Scellier B, Bengio Y (2017) Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. Frontiers in computational neuroscience 11:24.

Schwartenbeck P, FitzGerald TH, Mathys C, Dolan R, Kronbichler M, Friston K (2015) Evidence for surprise minimization over value maximization in choice behavior. Scientific reports 5:16575.
Simard P, LeCun Y, Denker JS (1993) Efficient pattern recognition using a new transformation distance. In: Advances in neural information processing systems, pp 50-58.
Sokoloff L (2008) The physiological and biochemical bases of functional brain imaging. In: Advances in Cognitive Neurodynamics ICCN 2007, pp 327-334: Springer.
Stuart G, Sakmann B (1995) Amplification of EPSPs by axosomatic sodium channels in neocortical pyramidal neurons. Neuron 15:1065-1076.
Sugase Y, Yamane S, Ueno S, Kawano K (1999) Global and fine information coded by single neurons in the temporal visual cortex. Nature 400:869-873.
Tero A, Takagi S, Saigusa T, Ito K, Bebber DP, Fricker MD, Yumiki K, Kobayashi R, Nakagaki T (2010) Rules for biologically inspired adaptive network design. Science 327:439-442.
Wang F, Casalino LP, Khullar D (2019) Deep learning in medicine—promise, progress, and challenges. JAMA internal medicine 179:293-294.
Whittington JC, Bogacz R (2017) An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. Neural computation 29:1229-1262.
Xie X, Seung HS (2003) Equivalence of backpropagation and contrastive Hebbian learning in a layered network. Neural computation 15:441-454.
Xu W, Huang X, Takagaki K, Wu J-y (2007) Compression and reflection of visually evoked cortical waves. Neuron 55:119-129.
Zhao Z-Q, Zheng P, Xu S-t, Wu X (2019) Object detection with deep learning: A review. IEEE transactions on neural networks and learning systems 30:3212-3232.

Suppl. Information

Title: Neurons learn by predicting future activity.

Authors: Artur Luczak, Bruce L. McNaughton, Yoshimasa Kubo

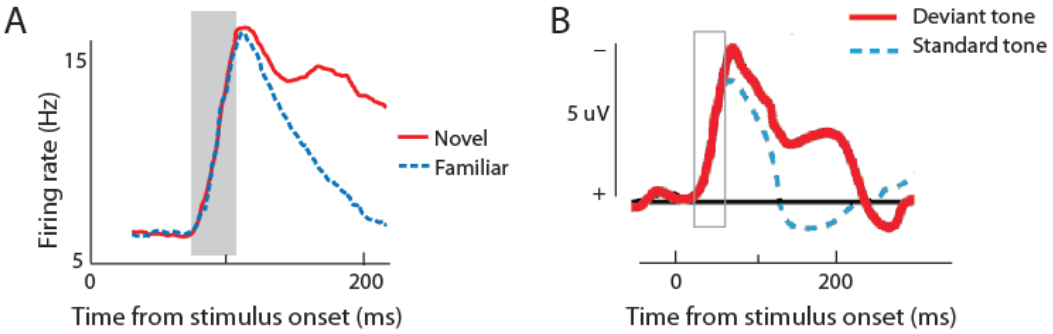

Suppl. Fig. 1. Consistently with our model, experiments in humans and monkeys show that activity corresponding to 'free phase' is followed by equivalent of 'clamped phase'. (A) Average population activity across 298 inferior temporal cortex neurons in monkeys during passive viewing of novel (red line) and familiar stimuli (dashed blue line; adopted from Freedman et al. 2006). Note that consistently with our model, for both types of stimuli the neuronal response are initially the same (denoted by gray area). However, later response to novel stimuli diverges from the familiar, likely due to top-down modulation from other brain regions. In the context of our work, response to familiar stimuli could be seen as expected activity equivalent to 'free phase', and response to novel stimuli could be interpreted as activity with additional top-down modulation analogous to 'clamped phase' in our model. Our learning rule suggests that the late response to novel stimuli which deviates from expected activity can provide training signal for synaptic update. (B) Average event-related potential recorded with EEG electrode located above the frontal cortex in healthy human adults in response to 1000 Hz standard tones and to 1032 Hz deviant tones (20% probability; adopted from Sams et al. 1985). This is a typical example from a rich body of scientific literature on the mismatch negativity phenomenon (a brain response to violations of expected stimulus rule). The similarity of neuronal responses at early phase and divergence from expected response for unexpected stimuli only in the later phase provides another biological example as justification for combining free and clamped phase in our model.

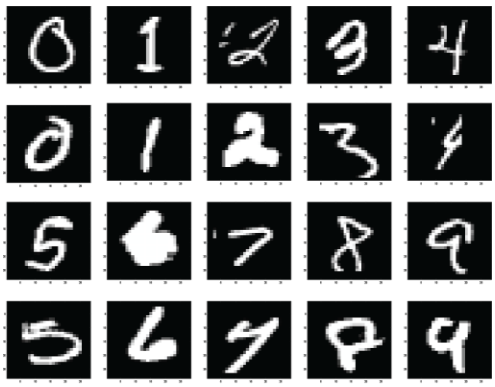

Suppl. Fig. 2. Examples of handwritten digits from MNIST data set (LeCun et al. 1998). Note that classifying such images is a non-trivial task even for humans, as for instance digits at the bottom row (5, 6, 7, 8, 9) could be mistaken for: 0, 4, 4, 9 and 4, respectively.

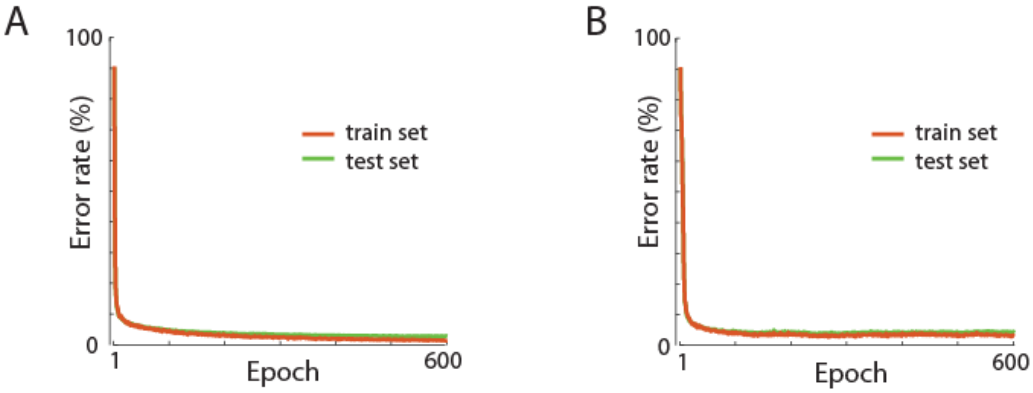

Suppl. Fig. 3 (A) The learning curve for a network with 80% of excitatory and 20% inhibitory neurons in the hidden layer. For each excitatory neuron, all outgoing synaptic connections to output neurons were non-negative ($w \geq 0$), and for inhibitory neurons all outgoing connections to output neurons were negative or zero ($w \leq 0$). The network had the same architecture as our main network (784-1000-10 neurons), and was using the same learning rule (Eq. 3). On the test data set it achieved a comparable error rate of 2.66%. (B) The learning curve for a network with 2 hidden layers (784-1000-512-10 neurons). The network had all the same learning rule and hyperparameters as our main network. However because in the deeper network it took a longer time for the network to settle to steady-state, we extended the free phase to 26 steps and started the clamped phase at step 27. This network solved the MNIST task with similar accuracy (error rate on test data: 3.68%), showing that our learning rule can be also applied to deeper networks.

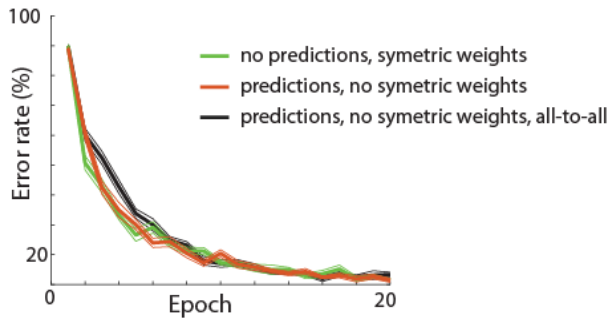

Suppl. Fig. 4. Learning curves for a network with our learning rule but with asymmetric connections (orange), and for the same network with additional asymmetric lateral connections within the hidden layer (all-to-all connectivity in hidden layer; black line). To increase the efficiency of model testing, here we used simpler networks with only 50 hidden neurons, and trained it only on a fraction of MNIST examples. For comparison, the learning curve for a network with the same number of neurons using original Contrastive Hebbian Learning algorithm with symmetric connections is shown in green. Similarity of learning curves suggests that our learning rule also works in more biologically plausible network configurations.

Note on symmetric and recurrent connections: Typical artificial neural networks trained with the backpropagation algorithm, require that a synaptic weight from neuron i to j (w_{ij}) is exactly the same as a synaptic weight from neuron j to i ($w_{ij}=w_{ji}$ is called symmetric weight). However, it was recently shown that networks without symmetric connections can still converge on a good solution (Lillicrap et al. 2016). Similar results were also found by Detorakis et al. (2019) who used asymmetric weights with Contrastive Hebbian Learning. Thus, our results are consistent with that previous work.

More interesting is our result that a fully recurrent network with lateral connections also converged similarly. This type of connectivity more closely reflects biological conditions as cortical neurons form extremely recurrent networks with many lateral connections. The main interest here is that the backpropagation algorithm is not well suited for networks with highly recurrent connectivity. Thus, the algorithm presented here with predictive neurons may provide a promising solution to training recurrent networks, but more future work is needed to address this question.

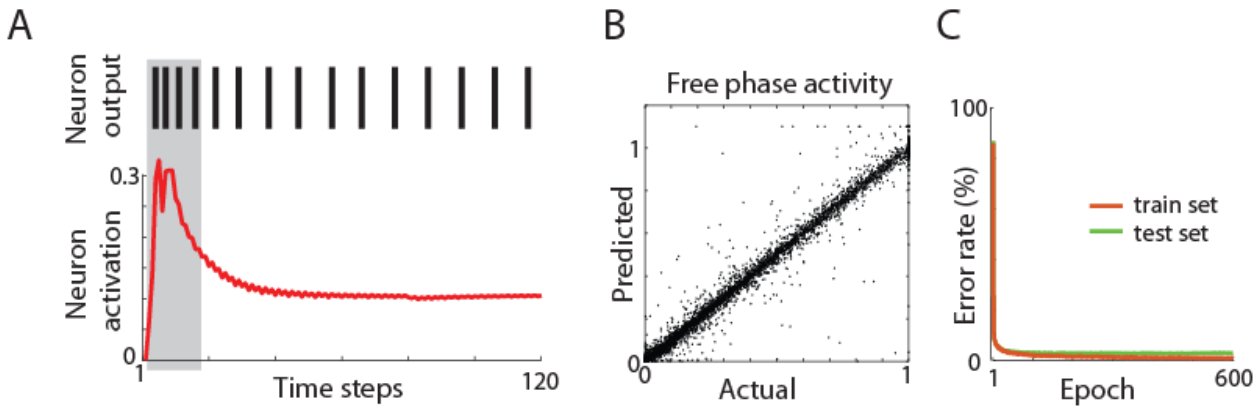

Suppl. Fig. 5. Implementation of our learning rule in a spiking neuronal network. (A) Activity of a sample neuron. Red trace shows neuron internal activation which is a function of synaptic inputs. (Top) Spikes are generated based on the value of internal activation. Gray shaded area marks the extent of free phase which is used to predict steady-state activation. (B) Actual vs. predicted steady-state free phase activity. Each dot represents the activity of one neuron during presentation of one stimuli. Activity of 1000 hidden neurons during presentation of 200 stimuli from the test dataset is shown. The correlation coefficient between actual and predicted activity was R=0.99 ($p<0.00001$). (C) Learning curves on MNIST task for training and testing dataset. Error rate of 2.46% on the test set shows that the spiking neuronal network with our learning rule was able to solve the presented task.

Spiking network model.

The code for this network with all implementation details is provided at: https://github.com/ykubo82/bioCHL/tree/master/bioCHLspk . Briefly, our network was based on work by O'Connor et al. (2019) who developed spiking networks using the Equilibrium Propagation algorithm which is an expansion of Contrastive Hebbian Learning (Scellier et al., 2017). The steady state is calculated using the Forward Euler method with each time step t: $A_j^t = (1 - \varepsilon) A_j^{t-1} + \varepsilon p'(A_j^{t-1}) (\sum_i w_{ij}\ p(A_i^{t-1}) + b_j)$, where A is the activation state, $p'(A_j)$ is the derivative of the activation function p, b is the bias, and i and j are neurons indices. ε can be seen as the learning rate of the activations. These states are clipped to range [0, 1]. Each neuron communicates only using binary signals: 0 or 1, to model spiking activity. For generating spikes, an encoder converts the neuron activation state to 1 or 0, which is sent as an output signal (panel A). The binary signals received by a neuron are converted into the activation state by a decoder. This could be seen as converting discrete spikes into a continuous post-synaptic membrane potential in actual neurons. We modified this spiking network by implementing our learning rule and by adding a least-squares model for predicting steady-state activation based on initial activation in steps 1-17. We also used AdaGrad (Duchi et al., 2011) to search for hyperparameters, which resulted in setting the learning rate to 0.01. The network architecture was the same as in our main model with 784-1000-10 neurons.

Acoustic stimuli

In animal experiments, as stimuli we used 1 s long pure tones (2, 3.3, 5.4, 8.9, 15 and 24 kHz at 60dB), interleaved with 1s periods of silence, as described in (Luczak et al. 2009 & 2013). Activity occurring >200ms after stimulus offset and before the next stimulus onset was regarded as spontaneous. The stimuli were continuously presented for 20-108 min, depending on how long the animal appeared to be comfortable during the experiment. All stimuli were tapered at the beginning and end with a 5ms cosine window. Experiments took place in a single-walled sound isolation chamber (IAC, Bronx, NY) with sounds presented free field (RP2/ES1, Tucker-Davis, Alachua, FL). To compensate for the transfer function of the acoustic chamber, tone amplitudes were calibrated prior to the experiment using a condenser microphone placed next to the animal's head (7017, ACO Pacific, Belmont CA) and a MA3 microphone amplifier (Tucker-Davis).

Changes in neuronal activity across periods

For each animal we divided the duration of recording into two halves. For each neuron we calculated the mean stimulus evoked firing rate in 30-40 ms time window, averaged over all stimulus presentations in each half of the experiment. Similarly, for each neuron we calculated the average predicted activity for the 30-40 ms time window, averaged across all stimuli in the 1st half of the experiment. We then calculated the difference between activities in the 2nd half minus the 1st half, and we correlated it with the difference between activity in the 1st half minus predicted activity in the 1st half (Fig. 3D). As described in the Methods section, only neurons with average stimulus evoked firing rates higher than 3 SD above pre-stimulus baseline were used in our analyses. Dividing the data in different proportions e.g. 60%-40% gave the same conclusions.

For comparison, analogous analyses were done on artificial neurons (Fig. 3C). For each neuron in the hidden layer the steady-state clamped activity was averaged over all 1200 stimulus presentations in a single learning epoch. Similarly, for each neuron in the hidden layer, the predicted activity was averaged across all stimuli presented in an epoch. We then calculated the difference between clamped activities in 2 epoch: $\Delta C_i = \langle\hat{x}_{i,M}\rangle - \langle\hat{x}_{i,N}\rangle$, where <> denotes average, i is index of neuron, and M and N are indexes of epochs with $M > N$. This was then correlated with difference between average clamped and predicted activity in epoch N: $\Delta P_i = \langle\hat{x}_{i,N}\rangle - \langle\tilde{x}_{i,N}\rangle$. We found that the correlation between ΔC and ΔP was strongest in the earliest epochs where the learning curve was most changing (Fig. 2E). For this simulation we used neural network as described in Methods with 784-1000-10 neurons, and we used the learning rule described in Eq. 3. However, in order to reproduce the repeated presentation of the same stimuli as in animal experiments, we used the same 1200 stimulus in each learning epoch. Presenting different stimuli at each epoch as in our original simulation, gave qualitatively similar results.

Predicting evoked dynamics from spontaneous activity

For predicting evoked dynamics from spontaneous, first we analyzed 'types' of spontaneous activity patterns. For this we divided spontaneous packets into 10 groups based on principal component analysis (PCA). Specifically, each packet was represented as a NxT matrix, where N is number of neurons and T=30 is number of 3ms long time bins. This matrix was then converted into a single vector V with length of NxT. This was repeated for each packet, thus we

obtained a VxP matrix, where P is the number of spontaneous packets (P for each animal was: 293, 855, 1215 and 2472). We then ran PCA on the VxP matrix and based on values of the 1st PC we divided packets into 10 groups. Note that the distribution of packets in PCA space was continuous, thus from this analysis it should not be interpreted that there are distinct types of packets. Dividing packets into 5-20 groups and including the 2nd and 3rd PC gave consistent results with those presented in the main text.

For predicting evoked dynamics from spontaneous, we faced the problem that defining the onset of spontaneous packets is less accurate than the onset of evoked responses. Thus, to ensure a similar precision of alignment, during each stimulus presentation we detected the onset of evoked responses using the same criteria as for spontaneous packets. Then, the evoked responses were realigned according to that of the detected onset. This allowed for estimating the spontaneous and evoked activity with the same precision. Using the original onset times instead of the realigned gave qualitatively similar results.

Maximizing future energy balance.

Intuitively, it makes sense that planning, i.e. making predictions, can improve success of an organisms in accessing more energy resources. In this section we show that this holds true even for a single neuron, where maximizing future energy balance is best achieved by predicting future activity. For that, first we will write equation for energy balance for a neuron at time $t+n$, where t represents current time and n is a small time increment. Using the same logic and notation as in Eq. 4, energy balance of a neuron j at time $t+n$, can be expressed as a function of: housekeeping processes ($-\varepsilon$), cost of electrical activity (which for simplified linear neurons can be written as sum of its synaptic inputs: $x_{j,t+n} = \sum_i w_{ij} x_{i,t+n}$), and energy supply from local blood vessels controlled by combined local activity of neurons ($\sum_k x_{k,t+n}$):

(Eq. S1): $E_{j,t+n} = -\varepsilon - b_1(\sum_i w_{ij} x_{i,t+n})^{\beta_1} + b_2(\sum_k x_{k,t+n})^{\beta_2}$

where β_1 and β_2 describe a non-linear relation between activity and energy (Devor et al., 2003).

In the main text we show in simulations (Fig. 2) and in experimental data (Fig. 3), that for small n, activity of neuron j at time $t+n$ could be approximated by a linear function of its activity at earlier time step t: $x_{j,t+n} = \lambda_j x_{j,t}$, where λ is a regression coefficient. Thus Eq. S1 can be rewritten as

(Eq. S2): $E_{j,t+n} = -\varepsilon - b_1(\lambda_j \sum_i w_{ij} x_{i,t})^{\beta_1} + b_2(\sum_{k\neq j} x_{k,t+n} + \lambda_j x_{j,t})^{\beta_2}$

Using gradient assent method, we can calculate change in weights to maximize future energy balance:

(Eq. S3): $\Delta w_{ij} = \frac{\partial E_{j,t+n}}{\partial w_{ij}} = -x_{i,t}\lambda_j\beta_1 b_1(\lambda_j \sum_i w_{ij} x_{i,t})^{\beta_1 - 1} +$

$$+ x_i \lambda_j \beta_2 b_2(\sum_{k\neq j} x_{k,t+n} + \lambda_j \sum_i w_{ij} x_{i,t})^{\beta_2 - 1}$$

Note that in Eq. S3: $\lambda_j \sum_i w_{ij} x_{i,t} = \lambda_j x_{j,t} \approx \tilde{x}_j$, thus this term corresponds to predicted future activity: $\tilde{x}_j$. We will also denote a population activity $\sum_{k \neq j} x_{k,t+n}$ as: $\tilde{\tilde{x}}$, which simplifies Eq. S3 to:

(Eq. S4): $\Delta w_{ij} = -x_{i,t}\lambda_j \beta_1 b_1 \tilde{x}_j^{\beta_1 - 1} + x_i \lambda_j \beta_2 b_2 (\tilde{\tilde{x}} + \tilde{x}_j)^{\beta_2 - 1}$

after factoring out $x_{i,t}\lambda_j$ and switching order of terms, we get:

(Eq. S5): $\Delta w_{ij} = x_{i,t}\lambda_j (\beta_2 b_2 \left(\tilde{\tilde{x}} + \tilde{x}_j \right)^{\beta_2 - 1} - \beta_1 b_1 \tilde{x}_j^{\beta_1 - 1})$

Considering that β_1 and $\beta_2 > 1.7$ (Devor et al., 2003), we may assume that $\beta_1 \approx 2$ and $\beta_2 \approx 2$, which allows to simplify and reorganize terms as follows:

(Eq. S6): $\Delta w_{ij} = x_{i,t}\lambda_j (2b_2 \left(\tilde{\tilde{x}} + \tilde{x}_j \right) - 2b_1 \tilde{x}_j) = x_{i,t} 2\lambda_j (b_2 \tilde{\tilde{x}} + b_2 \tilde{x}_j - b_1 \tilde{x}_j)$

$= x_{i,t} 2\lambda_j (b_2 \tilde{\tilde{x}} - \tilde{x}_j (b_1 - b_2)) = x_{i,t} 2\lambda_j (b_1 - b_2)(\frac{b_2}{(b_1 - b_2)} \tilde{\tilde{x}} - \tilde{x}_j)$

after denoting constant terms as $\alpha_3 = 2\lambda_j (b_1 - b_2)$, and as $\alpha_4 = \frac{b_2}{(b_1 - b_2)}$, we obtain:

(Eq. S7): $\Delta w_{ij} = \alpha_3 \, x_{i,t} \, (\alpha_4 \tilde{\tilde{x}} - \tilde{x}_j)$

This shows that maximizing future energy balance requires neuron to predict its future activity: $\tilde{x}_j$. Also note that in the brain, networks are highly recurrent, thus population activity $\tilde{\tilde{x}}$ can be seen as providing similar role as top-down modulation: $\hat{x}_i$. Altogether, this suggests, that type of a learning rule as in Eq. 3: $\Delta w_{ij} = \alpha \hat{x}_i (\hat{x}_j - \tilde{x}_j)$, may be a necessity to allow neuron for its highly energy intensive operation.

Supplemental references

Detorakis, G., Bartley, T. & Neftci, E. Contrastive Hebbian learning with random feedback weights. Neural Networks 114, 1-14, doi:10.1016/j.neunet.2019.01.008 (2019).

Devor A, Dunn AK, Andermann ML, Ulbert I, Boas DA, Dale AM (2003) Coupling of total hemoglobin concentration, oxygenation, and neural activity in rat somatosensory cortex. Neuron 39:353-359.

Duchi, J., Hazan, E. & Singer, Y. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. Journal of machine learning research 12, 2121-2159 (2011).

Freedman, D. J., Riesenhuber, M., Poggio, T. & Miller, E. K. Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. Cerebral Cortex 16, 1631-1644 (2006).

LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. Proceedings of the IEEE 86, 2278-2324 (1998).

Lillicrap, T. P., Cownden, D., Tweed, D. B. & Akerman, C. J. Random synaptic feedback weights support error backpropagation for deep learning. Nature communications 7, 1-10 (2016).

Luczak, A., Barthó, P. & Harris, K. D. Spontaneous events outline the realm of possible sensory responses in neocortical populations. Neuron 62, 413-425 (2009).

Luczak, A., Bartho, P. & Harris, K. D. Gating of sensory input by spontaneous cortical activity. The Journal of Neuroscience 33, 1684-1695 (2013).

Luczak A, Hackett TA, Kajikawa Y, Laubach M. Multivariate receptive field mapping in marmoset auditory cortex. Journal of neuroscience methods. 15;136(1):77-85 (2004).

Luczak A, Narayanan NS. Spectral representation—analyzing single-unit activity in extracellularly recorded neuronal data without spike sorting. Journal of neuroscience methods. 15;144(1):53-61 (2005)

O'Connor P, Gavves E, Welling M. Training a spiking neural network with equilibrium propagation. In The 22nd International Conference on Artificial Intelligence and Statistics 2019 Apr 11 (pp. 1516-1523).

Sams, M., Paavilainen, P., Alho, K. & Naatanen, R. Auditory frequency discrimination and event-related potentials. Electroencephalography and clinical neurophysiology 62, 437-448, (1985).

Scellier, B. & Bengio, Y. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. Frontiers in computational neuroscience 11, 24 (2017).