

UC Irvine

UC Irvine Previously Published Works

Title

Neurons learn by predicting future activity

Permalink

<https://escholarship.org/uc/item/5xj8n380>

Journal

Nature Machine Intelligence, 4(1)

ISSN

2522-5839

Authors

Luczak, Artur
McNaughton, Bruce L
Kubo, Yoshimasa

Publication Date

2022

DOI

10.1038/s42256-021-00430-y

Peer reviewed



HHS Public Access

Author manuscript

Nat Mach Intell. Author manuscript; available in PMC 2022 July 25.

Published in final edited form as:

Nat Mach Intell. 2022 January ; 4(1): 62–72. doi:10.1038/s42256-021-00430-y.

Neurons learn by predicting future activity

Artur Luczak^{1,*}, Bruce L. McNaughton^{1,2}, Yoshimasa Kubo¹

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

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

Abstract

Understanding how the brain learns may lead to machines with human-like intellectual capacities. It was previously proposed that the brain may operate on the principle of predictive coding. However, it is still not well understood how a predictive system could be implemented in the brain. Here we demonstrate that the ability of a single neuron to predict its future activity may provide an effective learning mechanism. 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. We show how this mathematically derived learning rule can provide a theoretical connection between diverse types of brain-inspired algorithms, thus, offering a step toward development of a general theory of neuronal learning. We tested this predictive learning rule in neural network simulations and in data recorded from awake animals. Our results also suggest that spontaneous brain activity provides “training data” for neurons to learn to predict cortical dynamics. Thus, the ability of a single neuron to minimize surprise: i.e. the difference between actual and expected activity, could be an important missing element to understand computation in the brain.

Introduction

Neuroscience is at the stage biology was at before Darwin. It has a myriad of detailed observations but no single theory explaining the connections between all of those observations. We do not even know if such a brain theory should be at the molecular level, or at the level of brain regions, or at any scale between. However, looking at Deep Neural Networks, which have achieved remarkable results in tasks ranging from cancer detection to self-driving cars, may provide useful insights. Although such networks may have different

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:

* Luczak@uleth.ca .

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.

Data availability

For testing our model, we used two publicly available datasets: MNIST²⁹ and CIFAR-10³³.

Code availability

Our code is publicly available at: <https://github.com/ykubo82/bioCHL> and at: <https://codeocean.com/capsule/4089503> (<https://doi.org/10.24433/CO.9801818.v1>)⁸⁴

Competing interests: The authors declare no competing interests.

inputs and architectures, most of their impressive behavior can be understood in terms of the underlying common learning algorithm, called backpropagation ¹.

Therefore, a better understanding of the learning algorithm(s) used by the brain could be central to develop a unifying theory of the brain function. There are two main approaches to investigate learning mechanisms in the brain: 1) experimental, where persistent changes in neuronal activity are induced by a specific intervention ²; and 2) computational, where algorithms are developed to achieve specific computational objectives while still satisfying selected biological constraints ^{3,4}. Here, we explore an additional option: 3) theoretical derivation, where a learning rule is derived from basic cellular principles, i.e. from maximizing metabolic energy of a cell. Using this approach, we found that maximizing energy balance by a neuron leads to a predictive learning rule, where a neuron adjusts its synaptic weights to minimize surprise: i.e. the difference between actual and predicted activity. Interestingly, this derived learning rule has a direct relation to some of the most promising biologically inspired learning algorithms, like predictive coding and temporal difference learning (see below), and Hebbian based rules can be seen as a special case of our predictive learning rule (see Discussion). Thus, our approach may provide a theoretical connection between multiple brain-inspired algorithms, and may offer a step toward development of a unified theory of neuronal learning.

There are multiple lines of evidence suggesting that the brain operates as a predictive system ^{5–10}. However, it remains controversial as to how exactly predictive coding could be implemented in the brain ⁴. Most of proposed mechanisms involve specially designed neuronal circuits with “error units” to allow for comparing expected and actual activity ^{11–14}. Whereas those models assume a predictive circuit, we propose an alternative, where there is an internal predictive model within a neuron. As many basic properties of neurons are highly conserved throughout evolution ^{15–17}, therefore, we suggest that a single neuron using predictive learning rule could provide an elementary unit from which a variety of predictive brains may be built.

Interestingly, our predictive learning rule can also be obtained by modifying a temporal difference learning algorithm to be more biologically plausible. Temporal difference learning is one of the most promising ideas of how backpropagation-like algorithms could be implemented in the brain. It is based on using differences in neuronal activity to approximate top-down error signals ^{4,18–24}. A typical example of such algorithms is Contrastive Hebbian Learning ^{25–27}, which was proven to be equivalent to backpropagation under certain assumptions ²⁸. 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:

$$\Delta w_{ij} = \alpha(\hat{x}_i\hat{x}_j - \check{x}_i\check{x}_j), \quad (\text{Eq. 1})$$

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. Our predictive learning rule provides a solution to this problem by predicting free phase steady-state activity, thus eliminating requirement for two separate stimulus presentations.

For manuscript clarity, first we will describe how our predictive learning rule can be obtained by modifying Contrastive Hebbian Learning algorithm. Next, we will validate the predictive learning rule in simulation and in data recorded from awake animals, and we will show how our results shed new light on the function of spontaneous activity. The details of derivation of the learning rule by maximizing neuron energy balance will be presented at the end.

Results

Predictive learning rule and Contrastive Hebbian Learning.

As mentioned earlier, Contrastive Hebbian Learning algorithm requires network to converge to steady-state equilibrium in two separate learning phases, thus exactly the same stimulus has to be presented twice. However, this is unlikely to be the case in the actual brain. Here, we propose to 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). This 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" (Supplementary 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.

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 ($\check{x}_{(1)}, \dots, \check{x}_{(12)}$) 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: $\check{x}_{(120)} \approx \tilde{x} = \lambda_{(1)} * \check{x}_{(1)} + \dots + \lambda_{(12)} * \check{x}_{(12)} + b$, 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}):

$$\Delta w_{ij} = \alpha(\hat{x}_i\hat{x}_j - \tilde{x}_i\tilde{x}_j). \quad (\text{Eq. 2})$$

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 not realistic. To solve this problem, we replaced (\tilde{x}_i) by the actual presynaptic activity in the clamped phase (\hat{x}_i), which we validated in network simulations (see the next section). This change leads to the following simplified synaptic plasticity rule:

$$\Delta w_{ij} = \alpha(\hat{x}_i\hat{x}_j - \hat{x}_i\tilde{x}_j) = \alpha\hat{x}_i(\hat{x}_j - \tilde{x}_j). \quad (\text{Eq. 3})$$

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).

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 (MNIST²⁹; Supplementary Fig. 2; Methods). This network achieved 1.9% error rate, which is similar to neural networks with comparable architecture trained with the backpropagation algorithm²⁹. 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 $R=1\pm 0.0001SD$ (averaged across 1000 hidden neurons in response to 200 randomly selected test images). Note that for we always used a cross-validation approach, where we trained a predictive model for each neuron on a subset of the data and applied that model 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.

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 law³⁰), and that about 80% of cortical neurons are excitatory. The network with this architecture achieved an error rate of 2.66% (Supplementary Fig. 3a). We also tested our algorithm in a network without symmetric weights, which resulted in similar performance as the original network (1.96%, Supplementary Fig. 3b). Moreover, we implemented the predictive learning rule in a network with spiking neurons, which again achieved a similar error rate of 2.46% (Supplementary Fig. 4). Our predictive learning rule was also tested in a deep convolutional network (Fig. 3a), which architecture was shown to resemble neuronal processing in the visual system^{31,32}. Using this convolutional network, we tested our algorithm on a more challenging dataset for biologically inspired algorithms: CIFAR-10³³. This dataset consists of color images representing 10 different classes like: airplanes, cars, birds, cats, etc. We achieved 20.03% error rate, which was comparable with training the same network using backpropagation through time algorithm (Fig. 3b; see Methods for details; code to reproduce those results is available at: <https://github.com/ykubo82/bioCHL/tree/master/conv>). 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 (Supplementary Information). For each of the 6 tones we calculated separately average onset and offset response, giving us 12 different activity profiles for each neuron (Fig. 4a). 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. 4b). Distribution of correlations coefficients for individual neurons were significantly different from 0 (t-test $p<0.0001$; all tests were two-sided; insert in Fig. 4b). This shows that neurons

have predictable dynamics, and from an initial neuronal response its future activity could be estimated.

However, much stronger evidence supporting our learning rule is provided by predicting long-term changes in cortical activity. Specifically, repeated presentation of stimuli over tens of minutes induces long-term changes in neuronal firing rates³⁴, similar as in perceptual learning. Importantly, based on our model, it was possible to infer which individual neurons will increase, and which neurons will decrease their firing rate. To explain it, first let's look at neural network simulations in Fig. 5a. 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. 5a; correlation coefficient $R = 0.35$, $p < 0.0001$; Supplementary Information). 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. 5b, Supplementary Information). Those changes in activity patterns were blocked by an NMDA receptor antagonist, as we showed using this data in³⁴, which provides strong support that this phenomenon depends on synaptic plasticity. Results presented in Fig. 5 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.

Deriving predictive model 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^{35–38}. For example, spontaneous activity is composed of ~50–300 ms long population bursts called packets, which resemble stimulus evoked patterns³⁹. This is illustrated in Figure 6a, where spontaneous activity packets in the auditory cortex are visible before sound presentation^{40,41}. 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; Supplementary Information). 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 (Supplementary Information), and for each neuron we calculated its average activity in each group (Fig. 6b). As in previous analyses in Fig. 4a, 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 (Supplementary Information). This least-square model was then applied to predict future evoked responses from initial evoked activity for all 12 stimuli. Figure 6c 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 consumption⁴², 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^{35–38}, thus spontaneous activity can provide “training data” for neurons to build a predictive model.

Learning rule derivation by maximizing neuron energy.

Interestingly, predictive learning rule in Eq. 3: $\Delta w_{ij} = \alpha \hat{x}_i (\hat{x}_j - \bar{x}_j)$ is not an *ad hoc* algorithm devised to solve a computational problem, but this form of learning rule 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 used⁴³. 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 – 4.8⁴⁴. 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: ⁴⁵). 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 ($k \in \{1, \dots, j, \dots, K\}$); 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 : 1.7–2.7⁴⁴. Note that sum of local population activity $\sum_k x_k$, also includes activity of neuron j : $x_j = \sum_i w_{ij}x_i$, as all local neurons contribute to local neurovascular coupling. Putting all the above terms together, the energy balance of a neuron j could be expressed as:

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

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 (cost). This energy objective of a cell could be paraphrased as “*lazy neuron principle*”: *maximum impact with minimum activity*.

We can calculate such required changes in synaptic weights w 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}

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

The appearance of x_i in the last term in (Eq. 5) comes from the fact that $\sum_k x_k$, includes x_j which is function of $w_{ij}x_i$, as explained above. Thus, 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 moving x_i in front of brackets and after switching order of terms we obtain:

$$\Delta w_{ij} = x_i \left(\beta_2 b_2 \bar{x}^{\beta_2 - 1} - \beta_1 b_1 x_j^{\beta_1 - 1} \right). \quad (\text{Eq. 6})$$

In case that $\beta_1 = 2$ and $\beta_2 = 2$, this formula simplifies from exponential to linear. However, even if β_1 and β_2 are anywhere in the range: $1.7 < \beta_1 < 4.8$ and $1.7 < \beta_2 < 2.7$ ⁴⁴, the expression $(\bar{x}^{\beta_2 - 1} - x_j^{\beta_1 - 1})$ still is well approximated by its linearized version: $(\bar{x} - x_j)$ for typical values of x in range 0–1 (Supplementary Fig. 5). After also denoting that $\alpha_1 = \beta_1 b_1$ and $\alpha_2 = \frac{\beta_2 b_2}{\alpha_1}$ and after taking α_1 in front of brackets, we obtain:

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

Although in this derivation we used linear model of a neuron, including a non-linear neural model like ReLU: $f(x) = x^+ = \max(0, x)$ leads to similar expression (Supplementary Information). Moreover, if we use the same derivation steps but to maximize neuron energy balance *in the future*, then Eq. 7, changes to (Eq. S7): $\Delta w_{ij} = \alpha_3 x_{i,t} (\alpha_4 \tilde{\bar{x}} - \tilde{x}_j)$ (see Supplementary Information for details of derivation)

Note that Eq. S7 has the same form as the predictive learning rule in Eq. 3: $\Delta w_{ij} = \alpha \hat{x}_i (\hat{x}_j - \tilde{x}_j)$. Here, $\tilde{\bar{x}}$ 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 from Eq. 7, became here future predicted activity: \tilde{x}_j . Thus, this derivation 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^{46–48}. 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.

This is supported by a theoretical derivation showing that the predictive learning rule provides an optimal strategy for maximizing metabolic energy of a neuron. To our knowledge, this is the first time where a synaptic learning rule has been derived from basic cellular principles, i.e. from maximizing energy of a cell. This provides a more solid theoretical basis over previous biologically-inspired algorithms, which were developed *ad hoc* to solve specific computational tasks while still satisfying selected biological constraints. However, it should be emphasized that many of those previous algorithms provided novel and insightful ideas which enabled the development of our model. Importantly, our derived learning rule provides a theoretical connection between those diverse types of brain-inspired algorithms, as discussed below.

One of the most influential ideas about brain's learning algorithm was proposed by Donald Hebb, which is based on correlated firing: a.k.a. 'cells that fire together wire together'⁴⁹. This could be written as: $\Delta w_{ij} \propto x_i x_j$, where Δw_{ij} is change in synaptic weight between neurons i and j , \propto denotes proportionality, and x_i and x_j represents pre- and post- synaptic activity, respectively. Note that this is a special case of our predictive learning rule: $\Delta w_{ij} \propto x_i(x_j - \tilde{x}_j)$ when $\tilde{x}_j = 0$, i.e. when a neuron does not make any prediction (note that here x_i and x_j represent actual activity as it is the case in clamped phase (i.e. \hat{x}_i and \hat{x}_j in Eq. 3), thus for comparison clarity, hat symbol ^ can be omitted here). Despite its influential role, the original Hebb's rule was shown to be unstable, as the synaptic weights will tend to increase or decrease exponentially. To solve this problem, a BMC theory was proposed⁵⁰, which can be expressed in a simplified form as: $\Delta w_{ij} \propto x_i(x_j - \theta_j)x_j$, where θ_j can be considered as the average activity of neuron j across all input patterns. Note that if in our equation: $\Delta w_{ij} \propto x_i(x_j - \tilde{x}_j)$, we would use the simplest predictive model: always predicting the average activity, then $\tilde{x}_j = \theta_j$ and our predictive rule becomes equivalent to the core part of BCM rule and could be seen as a linearized version of the full BCM rule. However, it was noted that networks trained using the BCM rule do not achieve the same level of accuracy as other learning rules⁵¹. It is consistent with our experience that performance of our algorithm deteriorated when we used average activity of each neuron for predictions. We interpret from this, that dynamically adjusting predictions based on most recent activity allows for more precise weight adjustments.

Moreover, we described in the results section, how our predictive learning rule directly relates to Contrastive Hebbian Learning, which belongs to class of temporal difference learning algorithms. Our algorithm is also similar to other predictive algorithms (see Introduction). The main difference is that we propose that neurons can internally calculate their predictions, rather than relying on specialized neuronal circuits, as pointed out by the Reviewer. We already mentioned earlier that organisms with simpler neuronal systems may not have predictive circuits as proposed to exist in the cortex^{12,14}. Thus, a predictive learning rule at the level of a single neuron may provide more basic description of the learning process across different brains. However, our model should not be taken as precluding possibility that in more complex brains, in addition to intracellular predictions,

neurons may form predictive circuits to enhance predictive abilities of an organism. Our model is also closely related to the work of^{52–54} where depolarization of basal dendrites serves as a prediction of top down signals from apical dendrites in pyramidal neurons. Again, our derived model could be seen as generalization of those ideas as it is not constrained to any specific cell type. The other interesting 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 plasticity⁵⁵. Considering all the above, we suggest that our plasticity rule derived from basic metabolic principles could serve as a common denominator for diverse types of biologically inspired learning algorithms, and as such, it may offer a step toward development of a unified neuronal learning theory.

Biological neurons have a variety of cellular mechanisms which operate on time scales of 1~100ms suitable for implementing predictions^{56–60}. The most likely mechanism appears to be calcium signaling. For example, when a neuron is activated, it leads to corresponding elevation of somatic calcium for tens of ms⁶¹. This time period with elevated calcium could indicate that a certain level of new input is expected to arrive in that time window. For instance, if bottom-up visual stimulus triggers multiple spikes in a neuron, then the resulting proportional increase in calcium concentration may signal that a higher level of follow up activity is expected, which could correspond to predicting a higher level of e.g. top-down modulation. This would be consistent with our experimental data, where higher activity at stimulus onset is correlated with higher activity ~20 ms later (Fig. 4; see Suppl. Information for more details on the plausibility of the predictive mechanism implementation and on proposed experiments to test it more directly). Interestingly, the core prediction of BCM and our model that synaptic weights should increase/decrease if neuron is stimulated above/below expected activity, is supported by experimental evidence where applying strong/weak electrical stimulation to cells in CA1 induced LTP/LTD, respectively⁶², which also involves calcium-dependent mechanisms⁶³. There are also other possible cellular properties which could support predictive mechanisms. For instance, it was shown that neurons can preferentially respond to inputs arriving at specific resonance frequencies (range: ~1–50 Hz)^{64,65}. This provides another example that neurons do have cellular mechanisms to ‘remember’ and to ‘act’ accordingly based on their past activity tens of ms earlier⁵⁸. Therefore, considering the cellular mechanisms listed above and the consistency of our model with experimental data presented in Figs 4–6, altogether it shows that neurons are at least capable of implementing the predictive learning rule.

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)⁶⁶. Also the later part of the packet can encode if there is discrepancy with expectation (e.g. this is a novel stimulus^{67,68}; Supplementary Fig. 1). This is likely because only the latter part of the packet can receive top-down modulation after information about that stimulus is exchanged between other brain areas, which is the case even during passive stimulus presentation^{69,70}. 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.

When recording neuronal activity in the cortex, the slowest oscillations (<10Hz) are by far the most dominant^{41,71}, and it is one of the biggest questions in neuroscience: what is the function of those oscillations⁷². Therefore it is worth noticing how learning rule derived from the basic cellular principles may relate to packets which are the main part of slow oscillations^{39,73,74}. As described above, dividing information into discrete packets, could provide an effective mechanism to improve neuronal predictions. It could allow for easier differentiation of feed-forward signals arriving during the initial wave of a packet from predicted top-down information arriving later. Another big question in neuroscience is about the function of spontaneous brain activity⁴². For example, why would the brain spend so much energy to generate packets even during e.g. sleep? Interestingly, similarly as in the brain where most energy is consumed by spontaneous activity⁴², in our model most energy (i.e. computational time) is used for free phase network activity, which allows intracellular predictive model to learn network dynamics in an unsupervised way. Thus, free phase activity in our model suggests that the function of spontaneous packets could be to provide neurons with diverse training data to improve the robustness of the predictive model, as supported by results presented in Fig. 6. Moreover, note that free phase activity may also be used for unsupervised learning. For example, if new input is present in the free phase, neurons can still calculate if such evoked activity is consistent with internal model predictions. If not, then weights can be modified to get free phase activity evoked by new stimuli to be closer to the prediction (it is the same mechanism as we use in clamped phase during supervised learning). This is a similar idea to unsupervised pre-training⁷⁵, however more future work is needed to investigate it.

Limitations

While the present study proposes a novel theoretical perspective on neuronal learning, this also comes with caveats that should be taken into account. Due to limits in current technology, parts of our model could not yet be properly validated experimentally. The major caveat in our model is the assumption of a cellular mechanism for predicting future activity. Although, neurons do have activity dependent calcium signaling⁶¹, there is no direct evidence that neurons use it to predict expected activity. The data that we present in Figs 4 & 6, shows that neurons have predictable dynamics, and this should be interpreted as only demonstrating that the main prerequisites for the predictive learning rule have been met, but it does not prove that neurons use it to make predictions. Also, for computational simplicity, in our model we present only one stimulus at a time to the network. In contrast, brains receive a constant stream of sensory stimuli, and new sensory inputs can arrive at the same time as top-down signals, which is not the case in our model. However, new sensory stimuli arriving during already in progress neuronal packets were shown to

be suppressed⁷⁶, which could serve to largely reduce interference between stimuli, as assumed in our model. The biological validity of this model assumption should be more directly tested. It is also important for our model that all data presented in the free phase to train predictive model have the same statistical distribution as data presented in the clamped phase. If only noise inputs were presented to the network in the free phase, then performance of our model would likely deteriorate. As mentioned earlier, numerous studies showed that spontaneous brain activity is not like a random noise, but rather it has similar statistical properties to stimulus evoked patterns^{35–38}. That, together with experimental results presented in Fig. 6, provide rationale for our network to use data with similar distributions during free and clamped phase. Moreover, there are other open questions about this model. For instance, consistently with our model, individual neurons can respond to novel stimuli with higher or lower firing rate as compared to familiar stimuli^{67,77}. However, on average, neurons recorded in cortex show typically higher firing rate to novel stimuli^{67,77}, which is not explained by our model. This discrepancy could be due to inherent sampling bias in electrophysiological recordings toward most active cells⁷⁸. It also may suggest additional network level predictive mechanisms which could explain the elevated response to novel stimuli, as proposed in^{13,14}. More future work is needed to answer those questions. It should also be noted that although our analytical derivation of the synaptic learning rule provides an important first step to link predictive learning models to metabolic activity, it required to largely simplify description of metabolic processes to only the few most important variables. Biological accuracy of this simplified description still needs to be investigated. Future work should also explore if implementing a non-linear predictive model within neurons could further improve performance of our network. Nevertheless, considering that presented model provides theoretical connection between diverse types of brain-inspired algorithms, this work could lead to a better understanding of neuronal principles⁷⁹.

Methods

Neural Network (MNIST dataset)

The code for our network with predictive learning rule, which we used to produce results presented in Fig. 2 is available at <https://github.com/ykubo82/bioCHL>, which contains all implementation details. Briefly, the base network has the architecture: 784–1000–10 with sigmoidal units, and with symmetric connections (see Supplementary Fig. 3–4 for more biologically plausible network architectures which we also tested). Neuron activity dynamics in hidden layer is described as in standard network with Contrastive Hebbian Learning⁸⁰:

$$x_{j,t} = S\left(\sum_p w_{p,j} x_{p,t-1} + \gamma \sum_o w_{o,j} x_{o,t-1} + b_j\right)h + x_{j,t-1}(1-h)$$

where $w_{p,j}$ denotes weight from neuron p in input layer to neuron j in hidden layer, $w_{o,j}$ denotes weight from output layer neuron to hidden layer neuron j , b is a bias, t is a time step and S is a sigmoid activation function. Parameter $h=0.1$ is the Euler method's time-step commonly used to improve computational stability. However, changing h to 0.2 or 1 resulted in similar network performance here. In the standard implementation of Contrastive Hebbian

Learning, all top-down connections $w_{o,j}$ are also multiplied by a small number γ (~ 0.1)⁸⁰. This different treatment of feed-forward and feedback connections could be biologically questionable as many brain circuits are highly recurrent and e.g. granule cells do not seem to have specific dendrites for receiving feedback signals. Therefore, to make our network more biologically plausible we set this feedback gain factor γ to 1, thus allowing our network to learn by itself what should be the contribution of each input. For output layer, term $\sum_o w_{o,j} x_{o,t-1}$ is set to 0 as there are no top-down connections to that layer. Neurons in the input layer do not have any dynamics as their activity is set to a value corresponding to pixel intensity in the presented image. To accelerate training, we used AdaGrad⁸¹, and we applied a learning rate of 0.03 to the hidden layer and 0.02 for the output layer. Synaptic weights for neurons in hidden and output layers were modified 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 ($\check{x}_{(1), \dots, \check{x}_{(12)}}$). This can be expressed as:

$$\tilde{x}_{(120)} = \lambda_{(1)} * \check{x}_{(1)} + \dots + \lambda_{(12)} * \check{x}_{(12)} + b, \quad (\text{Eq. 8})$$

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 & 3 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.

Convolutional Neural Network (CIFAR-10 dataset)

The convolutional network has an input layer of size: $32 \times 32 \times 3$, corresponding to size of a single image with 3 color channels in CIFAR-10 dataset (this dataset consists of 5000 training and 1000 test images for each of 10 classes³³). The network has two convolutional and pooling layers followed by one fully connected output layer (Fig. 3a). The filter size for all the convolutional layers is 3×3 with stride 1, and the number of filters is 256 and 512 for first and second convolutional layers, respectively. We did not use zero-padding. For pooling, we used the max pooling with 2×2 filters and stride 2. The activation function for the convolutional and the fully connected layers was the hard-sigmoid activation function, $S(x) = (1 + \text{hardtanh}(x-1)) * 0.5$, as implemented in²⁴. The learning rates were: 0.4, 0.028, and 0.025 for the first, second convolutional layer, and for the fully connected output layer, respectively. The Euler method's time-step h was set to 1. Considering that clamping

output neurons at only two extreme values: 0 or 1, may not be the most accurate model of top-down signals in the brain, thus here we implemented weak clamping as proposed in ²³. Shortly, instead of setting the value of the output neuron to 0 or 1 during clamped phase, output neurons were only slightly nudged toward required values. For example, if an output neuron should have value of 1, then it was clamped at value: $\check{x} + \epsilon$, where \check{x} is free phase steady-state activity of that output neuron, and ϵ is a small nudging factor toward 1. To calculate nudging for each neuron we used a clamping factor of 0.01 as described in ²³. This network with our predictive learning rule achieved 20.03% accuracy on CIFAR-10 dataset. Using the original ‘hard’ clamping, changing h to 0.1 or increasing number of neurons to 326 in the first layer gave similar results. We also directly compared predictive learning rule with backpropagation through time (BPTT) on the same convolutional network (Fig. 3). We selected BPTT as it uses a roll-out through time which is more comparable to our model. To ensure generality of presented results, we repeated the training with BPTT three times using different learning rates for each simulation. Using BPTT with the same learning rates as in our predictive model (0.4, 0.028, 0.025), the error rate was 20.88%. For BPTT with learning rate of 0.1 for all layers, the error rate was 21.23%, and 22.77% for a learning rate of 0.2 (Fig. 3b). Code for convolutional network was adopted from ⁸², which we modified to include our predictive learning rule. To reproduce our results, our code for the convolutional network with all implementation details is available at: <https://github.com/ykubo82/bioCHL/tree/master/conv>. Altogether, those results show that our predictive learning rule can also be successfully implemented in deeper networks and on more challenging tasks.

Surgery, recording and neuronal data

The experimental procedures for the awake, head-fixed experiment have been previously described ^{40,41} 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 200 μ m, 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)⁸³. 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 ^{40,73}.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

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 A. Gruber for sharing computational resources, K. Ali, L. Grasse, M. Klassen, E. Chalmers and R. Torabi for help, and we thank P. Bartho for sharing data.

References

1. Rumelhart DE, Durbin R, Golden R & Chauvin Y Backpropagation: The basic theory. Backpropagation: Theory, architectures and applications, 1–34 (1995).
2. Magee JC & Grienberger C Synaptic plasticity forms and functions. *Annual review of neuroscience* 43, 95–117 (2020).
3. Shouval HZ Models of synaptic plasticity. *Scholarpedia* 2, 1605 (2007).
4. Lillicrap TP, Santoro A, Marris L, Akerman CJ & Hinton G Backpropagation and the brain. *Nature Reviews Neuroscience*, 1–12 (2020). [PubMed: 31796912]
5. Schwartenbeck P, et al. Evidence for surprise minimization over value maximization in choice behavior. *Scientific reports* 5, 16575 (2015). [PubMed: 26564686]
6. Gordon N, Tsuchiya N, Koenig-Robert R & Hohwy J Expectation and attention increase the integration of top-down and bottom-up signals in perception through different pathways. *PLoS biology* 17, e3000233 (2019). [PubMed: 31039146]
7. Bar M The proactive brain: using analogies and associations to generate predictions. *Trends in cognitive sciences* 11, 280–289 (2007). [PubMed: 17548232]
8. Clark A Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences* 36, 181–204 (2013). [PubMed: 23663408]
9. Buzsáki G *The brain from inside out*, (Oxford University Press., 2019).
10. O'Reilly RC, Wyatte DR & Rohrlich J Deep predictive learning: a comprehensive model of three visual streams. *arXiv preprint arXiv:1709.04654* (2017).
11. Bastos AM, et al. Canonical microcircuits for predictive coding. *Neuron* 76, 695–711 (2012). [PubMed: 23177956]
12. Rao RP & Ballard DH Probabilistic models of attention based on iconic representations and predictive coding. in *Neurobiology of attention* 553–561 (Elsevier, 2005).
13. Whittington JC & Bogacz R An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural computation* 29, 1229–1262 (2017). [PubMed: 28333583]
14. Sacramento J, Costa RP, Bengio Y & Senn W Dendritic cortical microcircuits approximate the backpropagation algorithm. in *Advances in neural information processing systems* 8721–8732 (2018).
15. Gomez M, et al. Ca²⁺ signaling via the neuronal calcium sensor-1 regulates associative learning and memory in *C. elegans*. *Neuron* 30, 241–248 (2001). [PubMed: 11343658]
16. Roberts AC & Glanzman DL Learning in *Aplysia*: looking at synaptic plasticity from both sides. *Trends in neurosciences* 26, 662–670 (2003). [PubMed: 14624850]
17. Kandel ER, Schwartz JH & Jessell TM *Principles of neural science*, (McGraw-Hill New York, 2000).
18. O'Reilly RC Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural computation* 8, 895–938 (1996).
19. Ackley DH, Hinton GE & Sejnowski TJ A learning algorithm for Boltzmann machines. *Cognitive science* 9, 147–169 (1985).
20. Hinton GE & McClelland JL Learning representations by recirculation. in *Neural information processing systems* 358–366 (1988).

21. Hinton GE, Dayan P, Frey BJ & Neal RM The "wake-sleep" algorithm for unsupervised neural networks. *Science* 268, 1158–1161 (1995). [PubMed: 7761831]
22. Dayan P, Hinton GE, Neal RM & Zemel RS The helmholtz machine. *Neural computation* 7, 889–904 (1995). [PubMed: 7584891]
23. Scellier B & Bengio Y Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience* 11, 24 (2017). [PubMed: 28522969]
24. Laborieux A, et al. Scaling equilibrium propagation to deep convnets by drastically reducing its gradient estimator bias. *Frontiers in neuroscience* 15, 129 (2021).
25. Baldi P & Pineda F Contrastive learning and neural oscillations. *Neural computation* 3, 526–545 (1991). [PubMed: 31167332]
26. Almeida LB A learning rule for asynchronous perceptrons with feedback in a combinatorial environment. in *Artificial neural networks: concept learning* 102–111 (1990).
27. Pineda FJ Generalization of back-propagation to recurrent neural networks. *Physical review letters* 59, 2229 (1987). [PubMed: 10035458]
28. Xie X & Seung HS Equivalence of backpropagation and contrastive Hebbian learning in a layered network. *Neural computation* 15, 441–454 (2003). [PubMed: 12590814]
29. LeCun Y, Bottou L, Bengio Y & Haffner P Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86, 2278–2324 (1998).
30. Eccles JC, Fatt P & Koketsu K Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurons. *The Journal of physiology* 126, 524 (1954). [PubMed: 13222354]
31. LeCun Y & Bengio Y Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks* 3361, 1995 (1995).
32. Yamins DL, et al. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences* 111, 8619–8624 (2014).
33. Krizhevsky A & Hinton G Learning multiple layers of features from tiny images (2009).
34. Bermudez Contreras EJ, et al. Formation and reverberation of sequential neural activity patterns evoked by sensory stimulation are enhanced during cortical desynchronization. *Neuron* 79, 555–566 (2013). [PubMed: 23932001]
35. MacLean JN, Watson BO, Aaron GB & Yuste R Internal dynamics determine the cortical response to thalamic stimulation. *Neuron* 48, 811–823 (2005). [PubMed: 16337918]
36. Berkes P, Orbán G, Lengyel M & Fiser J Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331, 83–87 (2011). [PubMed: 21212356]
37. Kenet T, Bibitchkov D, Tsodyks M, Grinvald A & Arieli A Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956 (2003). [PubMed: 14586468]
38. Luczak A & MacLean JN Default activity patterns at the neocortical microcircuit level. *Frontiers in integrative neuroscience* 6(2012).
39. Luczak A, McNaughton BL & Harris KD Packet-based communication in the cortex. *Nature Reviews Neuroscience* (2015).
40. Luczak A, Barthó P & Harris KD Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62, 413–425 (2009). [PubMed: 19447096]
41. Luczak A, Barthó P & Harris KD Gating of sensory input by spontaneous cortical activity. *The Journal of Neuroscience* 33, 1684–1695 (2013). [PubMed: 23345241]
42. Raichle ME & Mintun MA Brain work and brain imaging. *Annu. Rev. Neurosci* 29, 449–476 (2006). [PubMed: 16776593]
43. Harris JJ, Jolivet R & Attwell D Synaptic energy use and supply. *Neuron* 75, 762–777 (2012). [PubMed: 22958818]
44. Devor A, et al. Coupling of total hemoglobin concentration, oxygenation, and neural activity in rat somatosensory cortex. *Neuron* 39, 353–359 (2003). [PubMed: 12873390]
45. Sokoloff L The physiological and biochemical bases of functional brain imaging. in *Advances in Cognitive Neurodynamics ICCN 2007* 327–334 (Springer, 2008).
46. Boisseau RP, Vogel D & Dussutour A Habituation in non-neural organisms: evidence from slime moulds. *Proceedings of the Royal Society B: Biological Sciences* 283, 20160446 (2016).

47. Kaiser AD Are myxobacteria intelligent? *Frontiers in microbiology* 4, 335 (2013). [PubMed: 24273536]
48. Tero A, et al. Rules for biologically inspired adaptive network design. *Science* 327, 439–442 (2010). [PubMed: 20093467]
49. Hebb DO *The organization of behavior: A neuropsychological theory*, (Wiley, New York, 1949).
50. Bienenstock EL, Cooper LN & Munro PW Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* 2, 32–48 (1982). [PubMed: 7054394]
51. Krotov D & Hopfield JJ Unsupervised learning by competing hidden units. *Proceedings of the National Academy of Sciences* 116, 7723–7731 (2019).
52. Hawkins J & Ahmad S Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in neural circuits* 10, 23 (2016). [PubMed: 27065813]
53. Guerguiev J, Lillicrap TP & Richards BA Towards deep learning with segregated dendrites. *Elife* 6, e22901 (2017). [PubMed: 29205151]
54. Payeur A, Guerguiev J, Zenke F, Richards BA & Naud R Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature neuroscience*, 1–10 (2021).
55. Bengio Y, Mesnard T, Fischer A, Zhang S & Wu Y STDP-compatible approximation of backpropagation in an energy-based model. *Neural computation* 29, 555–577 (2017). [PubMed: 28095200]
56. Stuart G & Sakmann B Amplification of EPSPs by axosomatic sodium channels in neocortical pyramidal neurons. *Neuron* 15, 1065–1076 (1995). [PubMed: 7576650]
57. Koch C, Rapp M & Segev I A brief history of time (constants). *Cerebral Cortex* 6, 93–101 (1996). [PubMed: 8670642]
58. Gutfreund Y, Yarom Y & Segev I Subthreshold Oscillations and Resonant-Frequency in Guinea-Pig Cortical-Neurons - Physiology and Modeling. *J Physiol-London* 483, 621–640 (1995). [PubMed: 7776248]
59. Larkum ME, Zhu JJ & Sakmann B A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398, 338–341 (1999). [PubMed: 10192334]
60. Ha GE & Cheong E Spike Frequency Adaptation in Neurons of the Central Nervous System. *Exp Neurol* 26, 179–185 (2017). [PubMed: 28912640]
61. Ali F & Kwan AC Interpreting in vivo calcium signals from neuronal cell bodies, axons, and dendrites: a review. *Neurophotonics* 7, 011402 (2019). [PubMed: 31372367]
62. Dudek S & Bear M Homosynaptic Long-Term Depression in Area CA1 of Hippocampus and Effects of N-methyl-D-aspartate Receptor Blockade. *Proceedings of the National Academy of Sciences of the United States of America* 89, 4363–4367 (1992). [PubMed: 1350090]
63. Bear MF Mechanism for a sliding synaptic modification threshold. *Neuron* 15, 1–4 (1995). [PubMed: 7619513]
64. Llinas RR, Grace AA & Yarom Y In vitro neurons in mammalian cortical layer 4 exhibit intrinsic oscillatory activity in the 10-to 50-Hz frequency range. *Proceedings of the National Academy of Sciences* 88, 897–901 (1991).
65. Hutcheon B & Yarom Y Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends in neurosciences* 23, 216–222 (2000). [PubMed: 10782127]
66. Sugase Y, Yamane S, Ueno S & Kawano K Global and fine information coded by single neurons in the temporal visual cortex. *Nature* 400, 869–873 (1999). [PubMed: 10476965]
67. Freedman DJ, Riesenhuber M, Poggio T & Miller EK Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. *Cerebral Cortex* 16, 1631–1644 (2006). [PubMed: 16400159]
68. Sams M, Paavilainen P, Alho K & Naatanen R Auditory frequency discrimination and event-related potentials. *Electroencephalography and clinical neurophysiology* 62, 437–448 (1985). [PubMed: 2415340]
69. Roland PE, et al. Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. *Proceedings of the National Academy of Sciences* 103, 12586–12591 (2006).

70. Xu W, Huang X, Takagaki K & Wu J-y. Compression and reflection of visually evoked cortical waves. *Neuron* 55, 119–129 (2007). [PubMed: 17610821]
71. Buzsáki G & Draguhn A Neuronal oscillations in cortical networks. *Science* 304, 1926–1929 (2004). [PubMed: 15218136]
72. Buzsaki G *Rhythms of the Brain*, (Oxford University Press, 2006).
73. Luczak A, Barthó P, Marguet SL, Buzsáki G & Harris KD Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences* 104, 347–352 (2007).
74. Luczak A Packets of Sequential Neural Activity in Sensory Cortex. in *Analysis and Modeling of Coordinated Multi-neuronal Activity* 163–182 (Springer, 2015).
75. Hinton GE & Salakhutdinov RR Reducing the dimensionality of data with neural networks. *Science* 313, 504–507 (2006). [PubMed: 16873662]
76. Sachdev RN, Ebner FF & Wilson CJ Effect of subthreshold up and down states on the whisker-evoked response in somatosensory cortex. *Journal of neurophysiology* 92, 3511–3521 (2004). [PubMed: 15254074]
77. Woloszyn L & Sheinberg DL Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron* 74, 193–205 (2012). [PubMed: 22500640]
78. Harris KD, Quiroga RQ, Freeman J & Smith SL Improving data quality in neuronal population recordings. *Nature neuroscience* 19, 1165–1174 (2016). [PubMed: 27571195]
79. Luczak A & Kubo Y Predictive neuronal adaptation as a basis for consciousness. *Front. Syst. Neurosci* (submitted) PsyArXiv <https://psyarxiv.com/538fw/> (2021).
80. Detorakis G, Bartley T & Neftci E Contrastive Hebbian learning with random feedback weights. *Neural Networks* 114, 1–14 (2019). [PubMed: 30831378]
81. Duchi J, Hazan E & Singer Y Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of machine learning research* 12, 2121–2159 (2011).
82. Ernout M, Grollier J, Querlioz D, Bengio Y & Scellier B Updates of Equilibrium Prop Match Gradients of Backprop Through Time in an RNN with Static Input. *Advances in neural information processing systems*, 7079–7089 (2019).
83. Harris KD, Henze DA, Csicsvari J, Hirase H & Buzsáki G Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *Journal of neurophysiology* 84, 401–414 (2000). [PubMed: 10899214]
84. Luczak A, McNaughton BL & Kubo Y Neurons learn by predicting future activity. *CodeOcean* (2021).




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 black 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}).




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.




Fig. 3. Implementation of predictive learning rule in a multilayer convolutional neuronal network.

a. Depiction of our convolutional network architecture (see Methods). **b.** Learning curves for convolutional network trained using a predictive learning rule (green) and for comparison the same network trained using backpropagation through time (BPTT). The red line shows a learning curve for BPTT using the same learning rates as in our predictive model (LR: 0.4, 0.028, 0.025); yellow line: BPTT with learning rate of 0.1 for all layers, and violet line: BPTT with learning rate of 0.2 for all layers. It illustrates that on CIFAR-10, performance of the deep network using our predictive learning rule was comparable with BPTT.




Fig. 4. 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.




Fig. 5. Long-term changes in neuronal activity in our model and in cortical neurons.

a. Average change in clamped steady-state activity between 2 consecutive learning epochs in our network model. This change relates to ‘surprise’: the difference between actual (clamped) and predicted activity, in the earlier epoch ($n=7$; Supplementary Information). Each dot represents one neuron. Regression line is shown in yellow. **b.** Average change in firing rate between 1st and 2nd half of our experiment with repetitive auditory stimulation. This firing rate change correlates with the difference between stimulus evoked and predicted activity during the 1st half of the experiment (Supplementary Information). 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. Thus, this evidence that a neuronal change in firing rate relates to ‘surprise’, provides a novel insight about neuronal plasticity.




Fig. 6. 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 ⁴¹).

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. 4. Skewness of the histogram to the right shows that for most neurons the evoked dynamics can be estimated based on spontaneous neuron's activity.