Check for updates

OPEN
Neurons learn by predicting future activity

Artur Luczak [1✉], Bruce L. McNaughton[1,2] and Yoshimasa Kubo[1]

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, whereby 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 algorithm, thus offering a step towards the 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—that is, the difference between actual and expected activity—could be an important missing element to understand computation in the brain.

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

A better understanding of the learning algorithm(s) used by the brain could thus be central to developing a unifying theory of brain function. There are two main approaches to investigating 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]. In this Article we explore an additional option—(3) theoretical derivation—where a learning rule is derived from basic cellular principles, that is, from maximizing the metabolic energy of a cell. Using this approach, we found that maximizing the energy balance by a neuron leads to a predictive learning rule, where a neuron adjusts its synaptic weights to minimize surprise—that is, 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 (Discussion). Thus, our approach may provide a theoretical connection between multiple brain-inspired algorithms and may offer a step towards the 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 the proposed mechanisms involve specially designed neuronal circuits with 'error units' to allow for comparing expected and actual activity[11–14]. Those models assume a predictive circuit, but 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], we suggest that a single neuron using a 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 about 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 (the activity of each neuron achieves a steady-state level). In the second 'clamped phase', in addition to presenting a stimulus to the input, the output neurons are also held clamped at values representing the stimulus category (for example, 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 the synaptic weights (w) according to the equation

$$\Delta w_{ij} = \alpha(\hat{x}_i \hat{x}_j - \check{x}_i \check{x}_j), \tag{1}$$

where i and j are indices of pre- and post-synaptic neurons respectively, and α is a small number representing the 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

[1]Canadian Center for Behavioural Neuroscience, University of Lethbridge, Lethbridge, Alberta, Canada. [2]Center for Neurobiology of Learning and Memory, University of California Irvine, Irvine, CA, USA. ✉e-mail: Luczak@uleth.ca

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 the free-phase steady-state activity, thus eliminating the requirement for two separate stimulus presentations.

For clarity here, first we will describe how our predictive learning rule can be obtained by modifying the 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 the neuron energy balance will be presented at the end.

Results

Predictive learning rule and contrastive Hebbian learning. As mentioned earlier, the contrastive Hebbian learning algorithm requires a network to converge to steady-state equilibrium in two separate learning phases, so 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 (for example, contours). This is then followed by top-down modulation containing more abstract information, such as '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 the predicted free phase and the clamped phase can then be used as a teaching signal. To implement this idea in our model, for each neuron, activity during 12 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, we first presented sample stimuli in the free phase to train a linear model, such that $\check{x}_{(120)} \approx \tilde{x} = \lambda_{(1)}\check{x}_{(1)}, +... + \lambda_{(12)}\check{x}_{(12)} + b$, where $\tilde{x}$ denote predicted activity, λ and b correspond to coefficients and offset term of the least-squares model, and terms in brackets correspond to time steps. Next, a new set of stimuli was used for which the free phase was run 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 the 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 the synaptic weights, in equation (1) we replace the activity in the free phase with predicted activity ($\tilde{x}$):

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

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 may not be 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 sec-
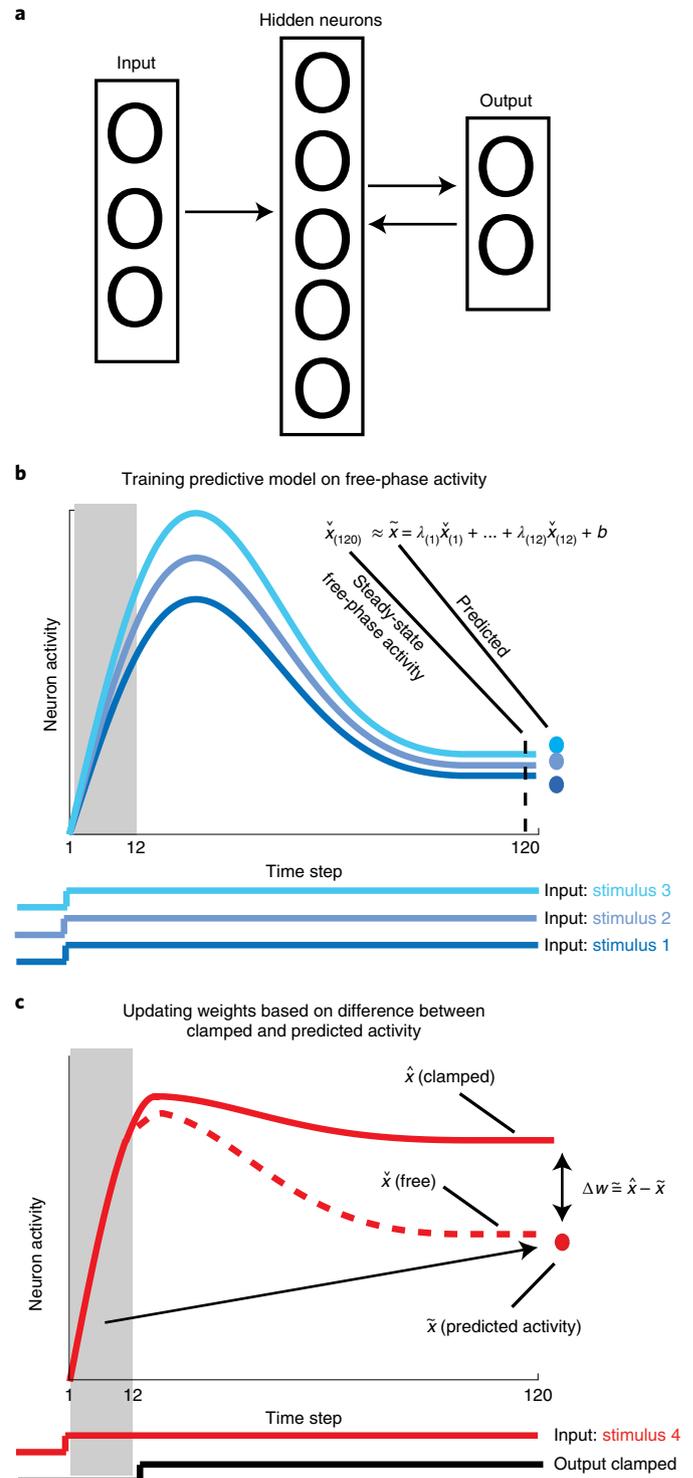
Fig. 1 | Basics of the algorithm. a, Schematic of the network. Note that activity propagates back and forth between hidden and output layers. **b**, Sample neuron activity in the free phase in response to different stimuli (marked with shades of blue). The free-phase responses are used to train a linear model to predict a steady-state activity from the activity at earlier time steps (marked by the shaded area; see main text for details). The bottom traces show the duration of the 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 the input signal (free phase), but, after a few steps, the output signal is also presented (clamped phase, bottom black trace). The red dot represents the steady-state free-phase activity predicted from the initial activity (the 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 the clamped phase ($\hat{x}$) and the predicted free-phase activity ($\tilde{x}$).
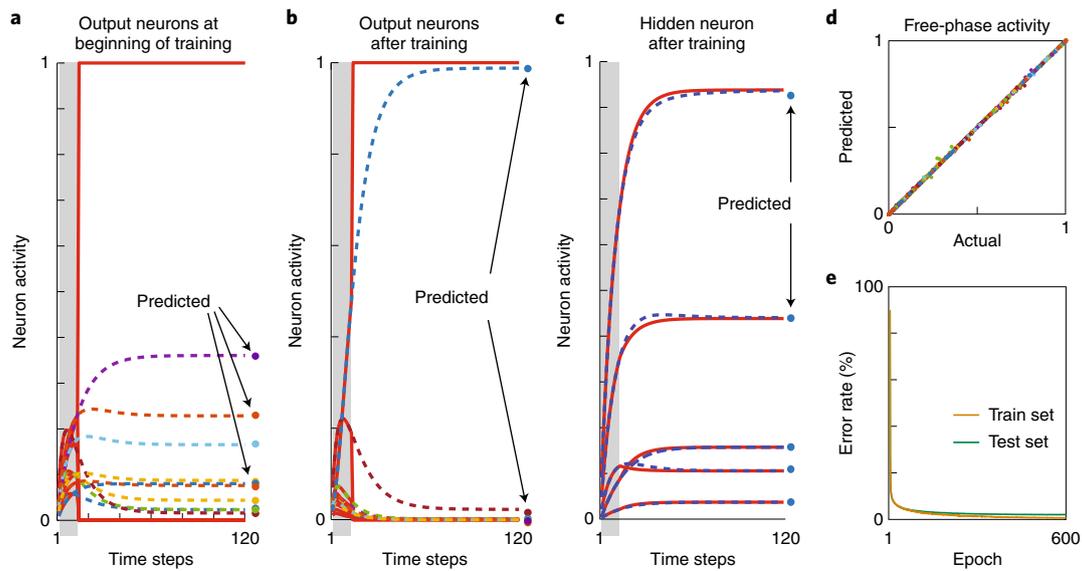
Fig. 2 | Neuron prediction of expected activity. a, Activity of ten output neurons in response to a sample stimulus at the beginning of network training. The grey shaded area indicates the extent of the free phase (time steps 1–12). Solid red lines show activity of the neurons clamped at step 13. For comparison, dashed lines represent the free-phase activity if the output neurons had not been clamped. Dots show the predicted steady-state activity in the free phase based on initial activity (from steps 1–12). **b**, Activity of the same neurons after network training. Note that the free-phase and predicted activities converged to the desired clamped activity. **c**, Activity of a representative neuron in the hidden layer in response to five different stimuli after network training. Solid and dashed lines represent clamped and free phases, respectively, and dots show predicted activity. **d**, Predicted versus actual free-phase activity. For clarity, only every 10th hidden neuron out of 1,000 is shown, in response to 20 sample images. Different colours represent different neurons, but some neurons may share the same colour due to the limited number of colours. The distribution of points along the diagonal shows that the predictions are accurate. **e**, Decrease in error rate across training epochs. Yellow and green lines denote learning curves for the training and test datasets, respectively. Note that, in each epoch, we only used 2% of 60,000 training examples.
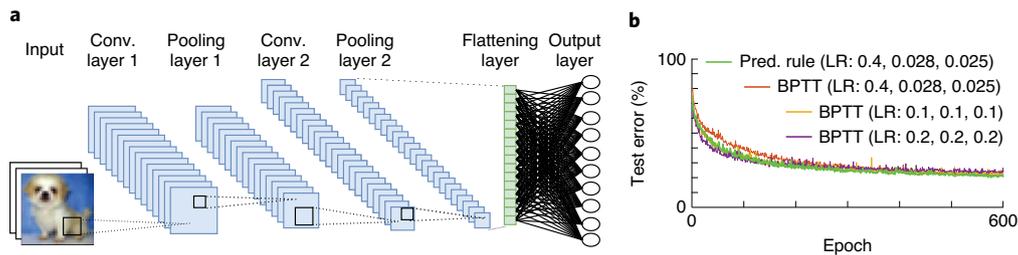

Fig. 3 | Implementation of the predictive learning rule in a multilayer convolutional neuronal network. a, Depiction of our convolutional (Conv.) network architecture (Methods). **b**, Learning curve for the convolutional network trained using the predictive (Pred.) learning rule (green) and, for comparison, learning curves for the same network trained using BPTT. The red line shows a learning curve for BPTT using the same learning rates as in our predictive model (red line; LR: 0.4, 0.028, 0.025), BPTT with a learning rate of 0.1 for all layers (yellow line) and BPTT with a learning rate of 0.2 for all layers (violet line). This shows that, on CIFAR-10, the performance of the deep network using our predictive learning rule was comparable with that of BPTT.

tion). This change leads to the following simplified synaptic plasticity rule (equation (3)):

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

Thus, to modify the 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 the predictive learning rule can be used to solve standard machine learning tasks, we created the following simulation. The neural network had 784 input units, 1,000 hidden units and 10 output units, and it was trained on a handwritten digit recognition task (MNIST[29]; Supplementary Fig. 2 and Methods). This network achieved 1.9%

error rate, which is similar to neural networks with comparable architecture trained with the backpropagation algorithm[29]. This demonstrates that the network with the predictive learning rule can solve challenging nonlinear classification tasks.

To verify that the neurons could correctly predict future free-phase activity, we took a closer look at sample neurons. Figure 2a presents the activity of all ten 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 the free phase. At time step 13, the output neurons were clamped, with the activity of nine neurons set to 0 and the activity of one neuron, representing the 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 the free phase, the network achieves a steady state,
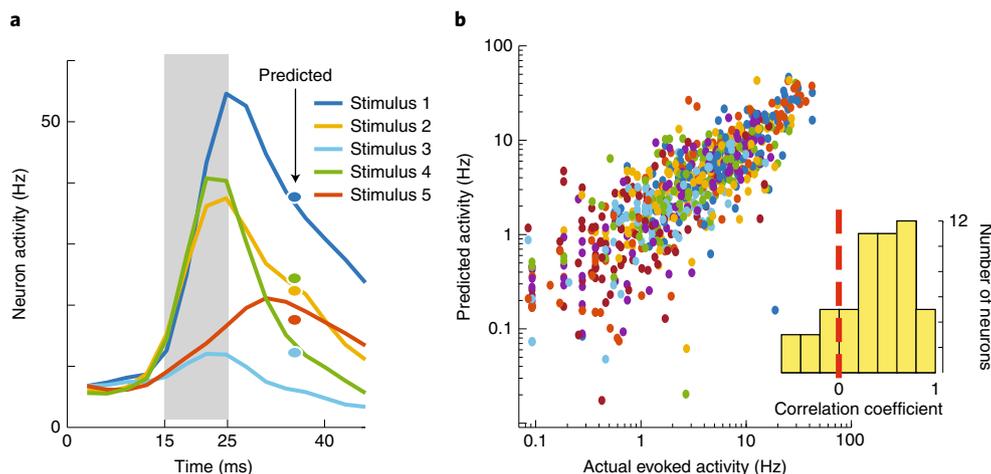
a

b

Fig. 4 | Predicting the future activity of cortical neurons. a, Response of a representative neuron to different stimuli. For visualization, only 5 out of 12 responses are shown. The grey shaded area indicates the time window that was used to predict future activity. Dots show the predicted average activity in the 30–40 ms time window. Colours correspond to different stimuli. **b**, Actual versus predicted activity for 55 cells from four animals in response to 12 stimuli. Different colours represent different neurons, but some neurons may share the same colour due to the limited number of colours. Inset: histogram of correlation coefficients for individual neurons. The skewness of the distribution to the right shows that, for most neurons, the correlation between actual and predicted response was positive.
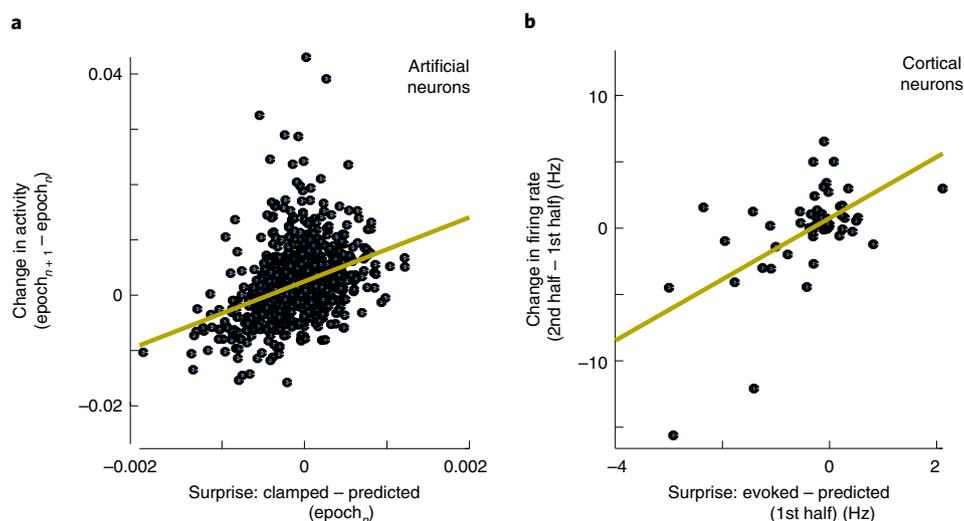
a

b

Fig. 5 | Long-term changes in neuronal activity in our model and in cortical neurons. a, Average change in clamped steady-state activity between two 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. The regression line is shown in yellow. **b**, Average change in firing rate between the first and second half of our experiment with repetitive auditory stimulation. This firing rate change correlates with the difference between stimulus-evoked and predicted activity during the first half of the experiment (Supplementary Information). Each dot represents the activity of one neuron averaged across stimuli. The similar behaviour 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.

with predicted activity closely matching. When the network is fully trained, it still takes about 50 steps for the network dynamics in the free phase to converge to a steady state (Fig. 2b). Note that, although all units initially increase their 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 the free phase closely matched the actual steady-state activity. The hidden units also converged to a steady state after about 50 steps. Figure 2c illustrates the response of one representative hidden neuron to five 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 Fig. 2d. The average correlation coefficient between predicted and actual free-phase activity was $R = 1 \pm 0.0001$ s.d. (averaged across 1,000 hidden neurons in response to 200 randomly selected test images). Note that 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 the weights (Methods). Thus, neurons were able to successfully generalize their predictions to new unseen stimuli. The
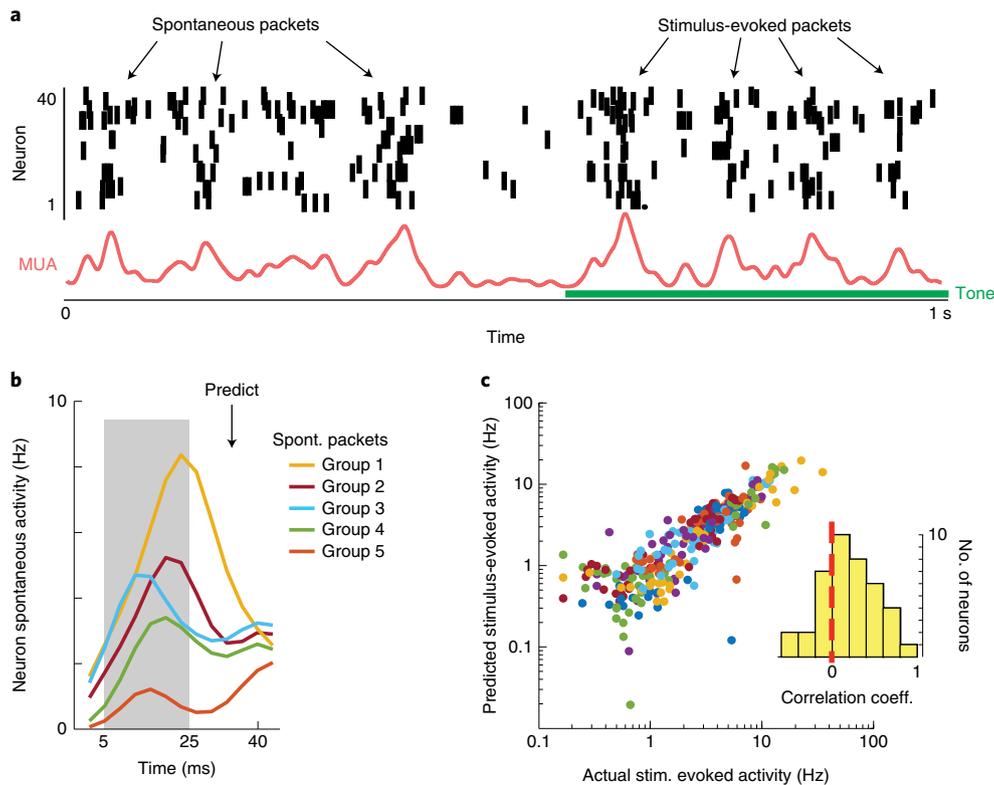
Fig. 6 | Predicting stimulus-evoked responses from spontaneous activity dynamics. a, Sample spiking activity in the auditory cortex before and during tone presentation. Note that spontaneous activity is not continuous, but rather composed of bursts called packets, which are similar to tone-evoked packets. The red trace shows the smoothed multiunit activity (MUA)—the summed activity of all neurons. Panel adapted with permission from ref. [41], Society for Neuroscience. **b**, Spontaneous (Spont.) packets were divided into ten groups based on population activity patterns. The activity of a single neuron in five different spontaneous packet groups is shown. The grey shaded area indicates the time window used for predicting future average activity within the 30–40 ms time window (marked by an arrow). This predictive model derived from spontaneous activity was then applied to predict future evoked activity based on the initial evoked response. **c**, Actual versus predicted tone-evoked activity. The plot convention is the same as in Fig. 4. The skewness of the histogram to the right shows that, for most neurons, the evoked dynamics can be estimated based on spontaneous neuron activity.

network error rates for the training and test datasets are 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 law[30]), and that ~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 a performance similar to 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 further tested in a deep convolutional network (Fig. 3a), the architecture of which has been 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-10[33]. This dataset consists of colour images representing ten different classes (for example, aeroplanes, cars, birds and cats). We achieved an error rate of 20.03%, which was comparable with that achieved training the same network using a

backpropagation through time (BPTT) algorithm (Fig. 3b; details are provided in the Methods and code to reproduce the 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 analysed neuronal recordings from the auditory cortex in awake rats (Methods). As stimuli we presented six tones, each 1 s long and interspersed by 1 s of silence, repeated continuously for over 20 min. (Supplementary Information). For each of the six tones we separately calculated the average onset and offset response, giving us 12 different activity profiles for each neuron (Fig. 4a). For each stimulus, the activity in the 15–25 ms time window was used to predict average future activity within the 30–40 ms window. We used 12-fold cross-validation, whereby responses from 11 stimuli were used to train the least-squares model, which was then applied to predict neuron activity for the one remaining stimulus. This procedure was repeated 12 times for each neuron. The average correlation coefficient between actual and predicted activity was $R = 0.36 \pm 0.05$ s.e.m. (averaged across 55 cells from four animals; Fig. 4b). The distributions of correlation coefficients for individual neurons were significantly different from 0 (t-test $P < 0.0001$; all tests were two-sided; inset, Fig. 4b). This shows that neurons have predictable dynamics and, from an initial neuronal response, their future activity can 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 rates[34], similar to that 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 this, first let us look at the neural network simulation results in Fig. 5a. This 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 first to the second half of the experiment was positively correlated with differences between evoked and predicted activity during the first half of the experiment ($R = 0.58$, $P < 0.0001$; Fig. 5b and Supplementary Information). Those changes in activity patterns were blocked by an NMDA (N-methyl-D-aspartate) receptor antagonist, as we showed, using this data, in ref. [34], which provides strong support that this phenomenon depends on synaptic plasticity. The results presented in Fig. 5 could be understood in terms of equation (3): if the actual activity is higher than predicted, then the synaptic weights are increased, thus leading to higher activity of that neuron in the next epoch. The similar behaviour of artificial and cortical neurons, where firing rate changes to minimize 'surprise' (the difference between actual and predicted activity), thus provides a strong evidence in support of the predictive learning rule presented here.

Deriving the 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 patterns[39]. This is illustrated in Fig. 6a, where spontaneous activity packets in the auditory cortex are visible before sound presentation[40,41]. In our experiments, each 1-s-long tone presentation was interspersed with 1 s of silence, and the activity during 200–1,000 ms after each tone was considered as spontaneous (animals were in a soundproof chamber; Supplementary Information). The individual spontaneous packets were extracted to estimate the neuronal dynamics (Methods), then the spontaneous packets were divided into ten groups based on similarity in principal component analysis (PCA) space (Supplementary Information), and for each neuron we calculated its average activity in each group (Fig. 6b). As in the previous analyses in Fig. 4a, the initial activity in time window 5–25 ms was used to derive the least-squares model to predict future spontaneous activity in the 30–40 ms time window (Supplementary Information). This least-squares model was then applied to predict future evoked responses from initial evoked activity for all 12 stimuli. Figure 6c shows actual versus predicted evoked activity for all neurons and stimuli (correlation coefficient $R = 0.2 \pm 0.05$ s.e.m., averaged over 40 cells from four animals; the inset 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[42], but the function of this activity remains a mystery. The foregoing results offer a new insight: because neuronal dynamics during spontaneous activity is similar to that during evoked activity[35–38], spontaneous activity can provide 'training data' for neurons to build a predictive model.

Learning rule derivation by maximizing the neuron energy. Interestingly, the predictive learning rule in equation (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 this form of learning rule arises naturally as a consequence of minimizing 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 used adenosine triphosphate (ATP)[43]. 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 presynaptic neuron i, w represents synaptic weights, b_1 is a constant to match energy units and β_1 describes a nonlinear relation between neuron activity and energy usage, which is estimated to be between 1.7 and 4.8 (ref. [44]). The remaining ~30% of neuron energy is consumed by 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 (for a review on neurovascular coupling see ref. [45]). This activity-dependent energy supply can be expressed as $+b_2 (\sum_k x_k)^{\beta_2}$, where x_k represents the spiking activity of neuron k from a local population of K neurons ($k \in \{1, …, j, …, K\}$), b_2 is a constant and β_2 reflects the exponential relation between activity and blood volume increase, which is estimated to be in the range of 1.7–2.7 (ref. [44]). Note that the sum of local population activity $\sum_k x_k$, also includes the 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

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

This formulation shows that, to maximize the 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 the 'lazy neuron principle': maximum impact with minimum activity.

We can calculate the required changes in synaptic weights Δw that will maximize a neuron's energy E_j by using the gradient ascent method. For this, we need to calculate the derivative of E_j with respect to w_{ij}:

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

The appearance of x_i in the last term in equation (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 the brackets and after switching the order of terms, we obtain

$$\Delta w_{ij} = x_i (\beta_2 b_2 \bar{x}^{\beta_2 - 1} - \beta_1 b_1 x_j^{\beta_1 - 1}). \qquad (6)$$

In the case where $\beta_1 = 2$ 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 $1.7 < \beta_2 < 2.7$, respectively[44], the expression $(\bar{x}^{\beta_2 - 1} - x_j^{\beta_1 - 1})$ is still well approximated by its linearized version, $(\bar{x} - x_j)$, for typical values of x in the 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 the brackets, we obtain

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

Although in this derivation we used a linear model of a neuron, including a nonlinear 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 the neuron energy balance in the future, then equation (7) changes to $\Delta w_{ij} = \alpha_3 x_{i,t}(\alpha_4 \overbrace{x} - \tilde{x}_j)$ (Supplementary equation (7); details about its derivation are provided in the Supplementary Information). Note that the above Supplementary equation (7) has the same form as the predictive learning rule in equation (3): $\Delta w_{ij} = \alpha \hat{x}_i(\hat{x}_j - \tilde{x}_j)$, where, $\overbrace{x}$ represents population recurrent activity, which can be thought of as top-down modulation, similar to $\hat{x}$. Also note that the activity of neuron j, x_j from equation (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 each neuron simply maximizing the energy balance.

Discussion

We have presented theoretical, computational and biological evidence that the basic principle underlying single neuron learning may rely on minimizing future surprise: the 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 have been shown to have 'intelligent' adaptive behaviours, including predicting the consequences of their actions so as to navigate towards 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 understanding learning mechanisms in the brain.

This is supported by a theoretical derivation showing that the predictive learning rule provides an optimal strategy for maximizing the metabolic energy of a neuron. To our knowledge, this is the first time a synaptic learning rule has been derived from basic cellular principles, that is, 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 that enabled the development of our model. Importantly, our derived learning rule provides a theoretical connection between those diverse brain-inspired algorithms, as discussed below.

One of the most influential ideas about the brain's learning algorithm was proposed by Donald Hebb, based on correlated firing and also known as 'cells that fire together wire together'[49]. This could be written as $\Delta w_{ij} \propto x_i x_j$, where Δw_{ij} is the 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 $\tilde{x}_j = 0$, that is, when a neuron does not make any prediction (note, here, that x_i and x_j represent actual activity as is the case in the clamped phase, that is, $\hat{x}_i$ and $\hat{x}_j$ in equation (3), so for comparison clarity, the 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 BCM theory was proposed[50] that 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 the 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 rules[51]. This is consistent with our experience that the performance of our algorithm deteriorated when we used the average activity of each neuron for predictions. From this, we interpret that dynamically adjusting predictions based on the most recent activity allows for more precise weight adjustments.

Moreover, we described in the Results how our predictive learning rule directly relates to contrastive Hebbian learning, which belongs to the class of temporal difference learning algorithms. Our algorithm is also similar to other predictive algorithms. The main difference is that we propose that neurons can internally calculate their predictions, rather than relying on specialized neuronal circuits. We mentioned earlier that organisms with simpler neuronal systems may not have the predictive circuits that are proposed to exist in the cortex[12,14]. Thus, a predictive learning rule at the level of a single neuron may provide a more basic description of the learning process across different brains. However, our model should not be taken as precluding the possibility that, in more complex brains, in addition to intracellular predictions, neurons may form predictive circuits to enhance the predictive abilities of an organism. Our model is also closely related to the work in refs. [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 a 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 has been shown that synaptic update rules are consistent with spike-timing-dependent plasticity[55]. 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 algorithm and, as such, it may offer a step towards the development of a unified neuronal learning theory.

Biological neurons have a variety of cellular mechanisms that operate on timescales of ~10–100 ms, suitable for implementing predictions[56–60]. The most likely mechanism appears to be calcium signalling. For example, when a neuron is activated, this leads to a corresponding elevation of somatic calcium for tens of milliseconds[61]. This time period with elevated calcium could indicate that a certain level of new input is expected to arrive in that time window. For example, if a 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, for example, 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; the Supplementary Information provides 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 a neuron is stimulated above/below the expected activity is supported by experimental evidence from applying strong/weak electrical stimulation inducing long-term potentiation (LTP)/long-term depression (LTD), respectively[62], which also involves calcium-dependent mechanisms[63]. There are also other possible cellular properties that could support predictive mechanisms. For example, it has been shown that neurons can preferentially respond to inputs arriving at specific resonance frequencies (range, ~1–50 Hz)[64,65]. This is another example suggesting that neurons do have cellular mechanisms to 'remember' and to 'act' accordingly based on their past activity

tens of milliseconds earlier[58]. Accordingly, considering the cellular mechanisms listed above and the consistency of our model with the experimental data presented in Figs. 4–6 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 milliseconds. We call such population bursts packets because they have a stereotypical structure, with neurons active at the beginning conveying bottom-up sensory information (for example, this is a face) and, later in the packet, representing additional higher-order information (for example, this is a happy face of that particular friend)[66]. Also, the later part of the packet can encode if there is a discrepancy with expectation (for example, this is a novel stimulus[67,68]; Supplementary Fig. 1). This is probably because only the later 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 per second, as only after evaluating if a given image is consistent with expectation can the next image be processed by the next packet, which takes ~50 ms. 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 (<10 Hz) are by far the most dominant[41,71], and one of the biggest questions in neuroscience is the function of those oscillations[72]. It is thus worth noticing how a learning rule derived from basic cellular principles may relate to packets that 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 feedforward signals arriving during the initial wave of a packet from predicted top-down information arriving later during the same packet. Another big question in neuroscience is about the function of spontaneous brain activity[42]. For example, why would the brain spend so much energy to generate packets even during sleep, for example? Interestingly, as in the brain, where most energy is consumed by spontaneous activity[42], in our model most energy (that is, computational time) is used for free-phase network activity, which allows the 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 the results presented in Fig. 6. Moreover, note that free-phase activity may also be used for unsupervised learning. For example, if a new input is present in the free phase, neurons can still calculate whether such evoked activity is consistent with internal model predictions. If not, then weights can be modified to get the free-phase activity evoked by new stimuli to be closer to the prediction (this is the same mechanism as we use in the clamped phase during supervised learning). This is a similar idea to unsupervised pre-training[75], but more future work is needed to investigate it.

Limitations. Although the present study proposes a novel theoretical perspective on neuronal learning, this also comes with caveats that should be taken into account. Because of the limits of current technology, parts of our model cannot 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 demonstrate activity-dependent calcium signalling[61], there

is no direct evidence that neurons use it to predict expected activity. The data that we present in Figs. 4 and 6 show 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 they do 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. Brains, in contrast, 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 neuronal packets already in progress have been shown to be suppressed[76], 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 the 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 the performance of our model would probably deteriorate. As mentioned earlier, numerous studies have shown 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 the experimental results presented in Fig. 6, provide a rationale for our network to use data with similar distributions during the free and clamped phases. Moreover, there are other open questions about this model. For example, consistent 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 the cortex show a 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 towards the most active cells[78]. It also may suggest the existence of additional network-level predictive mechanisms that could explain the elevated response to novel stimuli, as proposed in refs. [13,14]. More future work is needed to answer these 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 us to largely simplify the description of metabolic processes to only the few most important variables. The biological accuracy of this simplified description still needs to be investigated. Future work should also explore whether implementing a nonlinear predictive model within neurons could further improve the performance of our network. Nevertheless, considering that the presented model provides a theoretical connection between diverse brain-inspired algorithms, this work could lead to a better understanding of neuronal principles[79].

Methods

Neural network (the MNIST dataset). The code for our network with the predictive learning rule that we used to produce the results presented in Fig. 2 is available at https://github.com/ykubo82/bioCHL, which contains all implementation details. In brief, the base network has the following architecture: 784–1000–10 with sigmoidal units, and with symmetric connections (Supplementary Figs. 3 and 4 provide more biologically plausible network architectures that we also tested). The neuron activity dynamics in the hidden layer is described as in a standard network with contrastive Hebbian learning[80]:

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

where $w_{p,j}$ denotes the weight from neuron p in the input layer to neuron j in the hidden layer, $w_{o,j}$ denotes the weight from the 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) (ref. [80]). This different treatment of feedforward and feedback connections could be biologically questionable, as many brain circuits are highly recurrent; for example,

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 the contribution of each input should be. For the 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 AdaGrad[81], and we applied a learning rate of 0.03 to the hidden layer and 0.02 to the output layer. Synaptic weights for neurons in the hidden and output layers were modified as described in equation (3).

Future activity prediction. For all the predictions we used a cross-validation approach. Specifically, in each training cycle, we ran the free phase on 490 examples, which were used to derive the 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 ($\tilde{x}_{(1)}, ..., \tilde{x}_{(12)}$). This can be expressed as

$$\tilde{x}_{(120)} = \lambda_{(1)}\tilde{x}_{(1)} + ... + \lambda_{(12)}\tilde{x}_{(12)} + b, \qquad (8)$$

where terms in brackets correspond to time steps, and λ and b correspond to coefficients and offset term found by the least-squares method. Next, ten new examples were taken for which the free phase was run only for 12 steps, then the above derived least-squares model was applied to predict the free-phase steady-state activity for each of the ten 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 ten new examples. This process was repeated 120 times in each training epoch. From the MNIST dataset we used 60,000 examples for the above described training and 10,000 additional examples that were only used for testing. For all plots in Figs. 2 and 3 we only used test examples that the network had never seen during training. This demonstrates that each neuron can accurately predict its future activity even for novel stimuli that 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 the size of a single image with three colour channels in the CIFAR-10 dataset (this dataset consists of 5,000 training and 1,000 test images for each of ten classes[33]). 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 the 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 fully connected layers was the hard-sigmoid activation function, $S(x) = (1 + \text{hardtanh}(x - 1)) * 0.5$, as implemented in ref. [24]. The learning rates were 0.4, 0.028 and 0.025 for the first and second convolutional layers 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, here we implemented weak clamping as proposed in ref. [23]. In brief, instead of setting the value of the output neuron to 0 or 1 during the clamped phase, output neurons were only slightly nudged towards the required values. For example, if an output neuron should have a value of 1, then it was clamped at value $\tilde{x} + \varepsilon$, where $\tilde{x}$ is the free-phase steady-state activity of that output neuron and ε is a small nudging factor towards 1. To calculate nudging for each neuron we used a clamping factor of 0.01 as described in ref. [23]. This network with our predictive learning rule achieved 20.03% accuracy on the CIFAR-10 dataset. Using the original 'hard' clamping, changing h to 0.1 or increasing the number of neurons to 326 in the first layer gave similar results. We also directly compared the predictive learning rule with 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 the generality of the 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 a 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). The code for the convolutional network was adopted from ref. [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 described previously[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 (300–500 g) under ketamine–xylazine anaesthesia, 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 anaesthetized with isoflurane, the dura was resected and, after a recovery period, recording began. For recording we used silicon microelectrodes (Neuronexus Technologies) consisting of eight or four shanks spaced by 200 μm, 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 s.d. above the pre-stimulus baseline were used in analysis, resulting in 9, 12, 12 and 22 neurons from each rat. To predict evoked activity from spontaneous activity, we also required that neurons must have a 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 in refs. [40,73].

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this Article.

Data availability
For testing our model, we used two publicly available datasets: MNIST[29] and CIFAR-10[33].

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

References
1. Rumelhart, D. E., Durbin, R., Golden, R. & Chauvin, Y. in *Backpropagation*: *Theory, Architectures and Applications* (eds Chauvin, Y. & Rumelhart, D. E.) 1–34 (Psychology Press, 1995).
2. Magee, J. C. & Grienberger, C. Synaptic plasticity forms and functions. *Ann. Rev. Neurosci.* **43**, 95–117 (2020).
3. Shouval, H. Z. Models of synaptic plasticity. *Scholarpedia* **2**, 1605 (2007).
4. Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J. & Hinton, G. Backpropagation and the brain. *Nat. Rev. Neurosci.* **21**, 335–346 (2020).
5. Schwartenbeck, P. et al. Evidence for surprise minimization over value maximization in choice behavior. *Sci. Rep.* **5**, 16575 (2015).
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 Biol.* **17**, e3000233 (2019).
7. Bar, M. The proactive brain: using analogies and associations to generate predictions. *Trends Cogn. Sci.* **11**, 280–289 (2007).
8. Clark, A. Whatever next? Predictive brains, situated agents and the future of cognitive science. *Behav. Brain Sci.* **36**, 181–204 (2013).
9. Buzsáki, G. *The Brain from Inside Out* (Oxford Univ. Press, 2019).
10. O'Reilly, R. C., Wyatte, D. R. & Rohrlich, J. Deep predictive learning: a comprehensive model of three visual streams. Preprint at https://arxiv.org/abs/1709.04654 (2017).
11. Bastos, A. M. et al. Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
12. Rao, R. P. & Ballard, D. H. in *Neurobiology of Attention* (eds Itti, L. et al.) 553–561 (Elsevier, 2005).
13. Whittington, J. C. & Bogacz, R. An approximation of the error backpropagation algorithm in a predictive coding network with local Hebbian synaptic plasticity. *Neural Comput.* **29**, 1229–1262 (2017).
14. Sacramento, J., Costa, R. P., Bengio, Y. & Senn, W. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems* 8721–8732 (NIPS, 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).
16. Roberts, A. C. & Glanzman, D. L. Learning in aplysia: looking at synaptic plasticity from both sides. *Trends Neurosci.* **26**, 662–670 (2003).
17. Kandel, E. R., Schwartz, J. H. & Jessell, T. M. *Principles of Neural Science* (McGraw-Hill, 2000).
18. O'Reilly, R. C. Biologically plausible error-driven learning using local activation differences: the generalized recirculation algorithm. *Neural Comput.* **8**, 895–938 (1996).
19. Ackley, D. H., Hinton, G. E. & Sejnowski, T. J. A learning algorithm for Boltzmann machines. *Cogn. Sci.* **9**, 147–169 (1985).
20. Hinton, G. E. & McClelland, J. L. Learning representations by recirculation. In *Neural Information Processing Systems* 358–366 (NIPS, 1988).

21. Hinton, G. E., Dayan, P., Frey, B. J. & Neal, R. M. The 'wake-sleep' algorithm for unsupervised neural networks. *Science* **268**, 1158–1161 (1995).

22. Dayan, P., Hinton, G. E., Neal, R. M. & Zemel, R. S. The Helmholtz machine. *Neural Comput.* **7**, 889–904 (1995).

23. Scellier, B. & Bengio, Y. Equilibrium propagation: bridging the gap between energy-based models and backpropagation. *Front. Comput. Neurosci.* **11**, 24 (2017).

24. Laborieux, A. et al. Scaling equilibrium propagation to deep ConvNets by drastically reducing its gradient estimator bias. *Front. Neurosci.* **15**, 129 (2021).

25. Baldi, P. & Pineda, F. Contrastive learning and neural oscillations. *Neural Comput.* **3**, 526–545 (1991).

26. Almeida, L. B. A learning rule for asynchronous perceptrons with feedback in a combinatorial environment. In *Artificial Neural Networks*: *Concept Learning* (ed. Diederich, J.) 102–111 (ACM, 1990).

27. Pineda, F. J. Generalization of back-propagation to recurrent neural networks. *Phys. Rev. Lett.* **59**, 2229–2232 (1987).

28. Xie, X. & Seung, H. S. Equivalence of backpropagation and contrastive Hebbian learning in a layered network. *Neural Comput.* **15**, 441–454 (2003).

29. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).

30. Eccles, J. C., Fatt, P. & Koketsu, K. Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurones. *J. Physiol.* **126**, 524–562 (1954).

31. LeCun, Y. & Bengio, Y. Convolutional networks for images, speech, and time-series. In *The Handbook of Brain Theory and Neural Networks* (ed. Arbib, M. A.) 3361 (MIT Press, 1995).

32. Yamins, D. L. et al. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proc. Natl Acad. Sci. USA* **111**, 8619–8624 (2014).

33. Krizhevsky, A. & Hinton, G. *Learning Multiple Layers of Features from Tiny Images* Technical Report TR-2009 (Univ. Toronto, 2009).

34. Bermudez Contreras, E. J. et al. Formation and reverberation of sequential neural activity patterns evoked by sensory stimulation are enhanced during cortical desynchronization. *Neuron* **79**, 555–566 (2013).

35. MacLean, J. N., Watson, B. O., Aaron, G. B. & Yuste, R. Internal dynamics determine the cortical response to thalamic stimulation. *Neuron* **48**, 811–823 (2005).

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).

37. Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A. & Arieli, A. Spontaneously emerging cortical representations of visual attributes. *Nature* **425**, 954–956 (2003).

38. Luczak, A. & MacLean, J. N. Default activity patterns at the neocortical microcircuit level. *Front. Integrative Neurosci.* **6**, 30 (2012).

39. Luczak, A., McNaughton, B. L. & Harris, K. D. Packet-based communication in the cortex. *Nat. Rev. Neurosci.* **16**, 745–755 (2015).

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

41. Luczak, A., Bartho, P. & Harris, K. D. Gating of sensory input by spontaneous cortical activity. *J. Neurosci.* **33**, 1684–1695 (2013).

42. Raichle, M. E. & Mintun, M. A. Brain work and brain imaging. *Annu. Rev. Neurosci.* **29**, 449–476 (2006).

43. Harris, J. J., Jolivet, R. & Attwell, D. Synaptic energy use and supply. *Neuron* **75**, 762–777 (2012).

44. Devor, A. et al. Coupling of total hemoglobin concentration, oxygenation, and neural activity in rat somatosensory cortex. *Neuron* **39**, 353–359 (2003).

45. Sokoloff, L. The physiological and biochemical bases of functional brain imaging. In *Advances in Cognitive Neurodynamics ICCN 2007* (eds Wang, R. et al.) 327–334 (Springer, 2008).

46. Boisseau, R. P., Vogel, D. & Dussutour, A. Habituation in non-neural organisms: evidence from slime moulds. *Proc. R. Soc. B* **283**, 20160446 (2016).

47. Kaiser, A. D. Are myxobacteria intelligent? *Front. Microbiol.* **4**, 335 (2013).

48. Tero, A. et al. Rules for biologically inspired adaptive network design. *Science* **327**, 439–442 (2010).

49. Hebb, D. O. *The Organization of Behavior*: *A Neuropsychological Theory* (Wiley, 1949).

50. Bienenstock, E. L., Cooper, L. N. & Munro, P. W. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* **2**, 32–48 (1982).

51. Krotov, D. & Hopfield, J. J. Unsupervised learning by competing hidden units. *Proc. Natl Acad. Sci. USA* **116**, 7723–7731 (2019).

52. Hawkins, J. & Ahmad, S. Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Front. Neural Circuits* **10**, 23 (2016).

53. Guerguiev, J., Lillicrap, T. P. & Richards, B. A. Towards deep learning with segregated dendrites. *eLife* **6**, e22901 (2017).

54. Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A. & Naud, R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat. Neurosci.* **24**, 1010–1019 (2021).

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

56. Stuart, G. & Sakmann, B. Amplification of EPSPs by axosomatic sodium channels in neocortical pyramidal neurons. *Neuron* **15**, 1065–1076 (1995).

57. Koch, C., Rapp, M. & Segev, I. A brief history of time (constants). *Cerebral Cortex* **6**, 93–101 (1996).

58. Gutfreund, Y., Yarom, Y. & Segev, I. Subthreshold oscillations and resonant-frequency in guinea-pig cortical-neurons—physiology and modeling. *J. Physiol.* **483**, 621–640 (1995).

59. Larkum, M. E., Zhu, J. J. & Sakmann, B. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* **398**, 338–341 (1999).

60. Ha, G. E. & Cheong, E. Spike frequency adaptation in neurons of the central nervous system. *Exp. Neurobiol.* **26**, 179–185 (2017).

61. Ali, F. & Kwan, A. C. Interpreting in vivo calcium signals from neuronal cell bodies, axons and dendrites: a review. *Neurophotonics* **7**, 011402 (2019).

62. Dudek, S. & Bear, M. Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proc. Natl Acad. Sci. USA* **89**, 4363–4367 (1992).

63. Bear, M. F. Mechanism for a sliding synaptic modification threshold. *Neuron* **15**, 1–4 (1995).

64. Llinas, R. R., Grace, A. A. & Yarom, Y. In vitro neurons in mammalian cortical layer 4 exhibit intrinsic oscillatory activity in the 10- to 50-Hz frequency range. *Proc. Natl Acad. Sci. USA* **88**, 897–901 (1991).

65. Hutcheon, B. & Yarom, Y. Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends Neurosci.* **23**, 216–222 (2000).

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).

67. 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).

68. Sams, M., Paavilainen, P., Alho, K. & Naatanen, R. Auditory frequency discrimination and event-related potentials. *Electroencephalogr. Clin. Neurophysiol.* **62**, 437–448 (1985).

69. Roland, P. E. et al. Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. *Proc. Natl Acad. Sci. USA* **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).

71. Buzsáki, G. & Draguhn, A. Neuronal oscillations in cortical networks. *Science* **304**, 1926–1929 (2004).

72. Buzsaki, G. *Rhythms of the Brain* (Oxford Univ. Press, 2006).

73. Luczak, A., Barthó, P., Marguet, S. L., Buzsáki, G. & Harris, K. D. Sequential structure of neocortical spontaneous activity in vivo. *Proc. Natl Acad. Sci. USA* **104**, 347–352 (2007).

74. Luczak, A. in *Analysis and Modeling of Coordinated Multi-Neuronal Activity* (ed. Tatsuno, M.) 163–182 (Springer, 2015).

75. Hinton, G. E. & Salakhutdinov, R. R. Reducing the dimensionality of data with neural networks. *Science* **313**, 504–507 (2006).

76. Sachdev, R. N., Ebner, F. F. & Wilson, C. J. Effect of subthreshold up and down states on the whisker-evoked response in somatosensory cortex. *J. Neurophysiol.* **92**, 3511–3521 (2004).

77. Woloszyn, L. & Sheinberg, D. L. Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron* **74**, 193–205 (2012).

78. Harris, K. D., Quiroga, R. Q., Freeman, J. & Smith, S. L. Improving data quality in neuronal population recordings. *Nat. Neurosci.* **19**, 1165–1174 (2016).

79. Luczak, A. & Kubo, Y. Predictive neuronal adaptation as a basis for consciousness. *Front. Syst. Neurosci.* **15**, 767461 (2021).

80. Detorakis, G., Bartley, T. & Neftci, E. Contrastive Hebbian learning with random feedback weights. *Neural Netw.* **114**, 1–14 (2019).

81. Duchi, J., Hazan, E. & Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.* **12**, 2121–2159 (2011).

82. Ernoult, 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. In *Advances in Neural Information Processing Systems* 7079–7089 (NIPS, 2019).

83. Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H. & Buzsáki, G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.* **84**, 401–414 (2000).

84. Luczak, A., McNaughton, B. L. & Kubo, Y. Neurons learn by predicting future activity. *CodeOcean* https://doi.org/10.1101/2020.09.25.314211 (2021).

Author contributions

A.L. conceived the project, analysed data, performed computer simulations and wrote the manuscript. B.L.M. engaged in theoretical discussions and commented extensively on the manuscript. Y.K. performed computer simulations and contributed to writing the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s42256-021-00430-y.

Correspondence and requests for materials should be addressed to Artur Luczak.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Corresponding author(s): Artur Luczak

Last updated by author(s): Oct 31, 2021

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. F, t, r) with confidence intervals, effect sizes, degrees of freedom and P value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☒	☐	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Our code is publicly available at: https://github.com/ykubo82/bioCHL.
Data analysis	Our code is publicly available at: https://github.com/ykubo82/bioCHL.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:
- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

For testing our model, we used two publicly available datasets: MNIST [ref 29] and CIFAR-10 [ref 33].

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	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 (see Methods for detail).
Data exclusions	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.
Replication	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.
Randomization	All animals belong to the same group.
Blinding	Blinding was not relevant for this study. All animals belong to the same group.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☐	☒ Animals and other organisms
☒	☐ Human research participants
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Animals and other organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research

Laboratory animals	Sprague-Dawley male rats (300-500g)
Wild animals	*Provide details on animals observed in or captured in the field; report species, sex and age where possible. Describe how animals were caught and transported and what happened to captive animals after the study (if killed, explain why and describe method; if released, say where and when) OR state that the study did not involve wild animals.*
Field-collected samples	*For laboratory work with field-collected samples, describe all relevant parameters such as housing, maintenance, temperature, photoperiod and end-of-experiment protocol OR state that the study did not involve samples collected from the field.*
Ethics oversight	The experimental data have been previously published [ref 40,41] and all procedures were approved by the Rutgers University Animal Care and Use Committee, and conformed to NIH Guidelines on the Care and Use of Laboratory Animals

Note that full information on the approval of the study protocol must also be provided in the manuscript.