Supplementary information

Title: Neurons learn by predicting future activity.

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

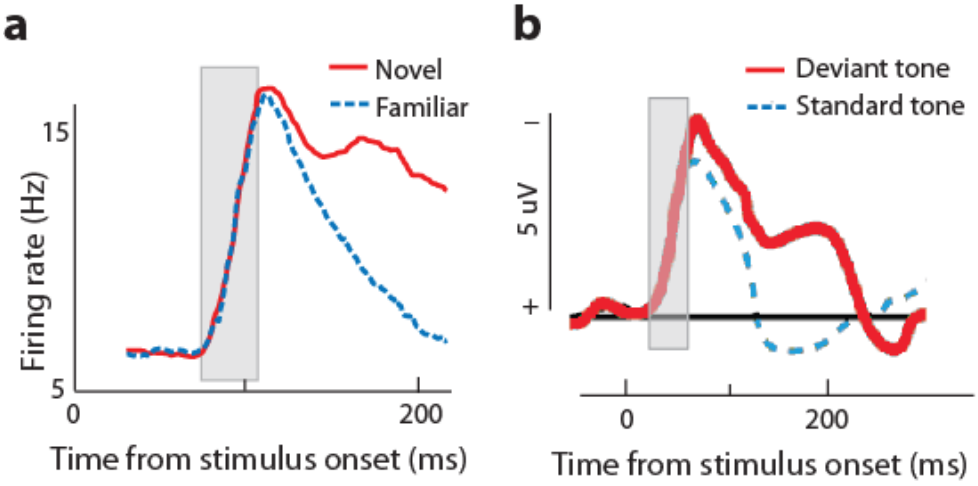

Suppl. Fig. 1. Consistently with our model, experiments in humans and monkeys show that activity corresponding to 'free phase' is followed by equivalent of 'clamped phase'. **a**, Average population activity across 298 inferior temporal cortex neurons in monkeys during passive viewing of novel (red line)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

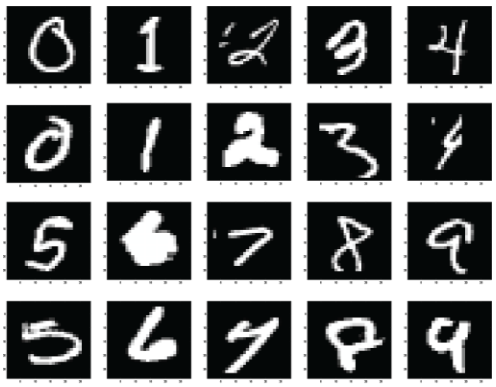

Suppl. Fig. 2. Examples of handwritten digits from MNIST data set (LeCun et al. 1998). Note that classifying such images is a non-trivial task even for humans, as for instance digits at the bottom row (5, 6, 7, 8, 9) could be mistaken for: 0, 4, 4, 9 and 4, respectively.

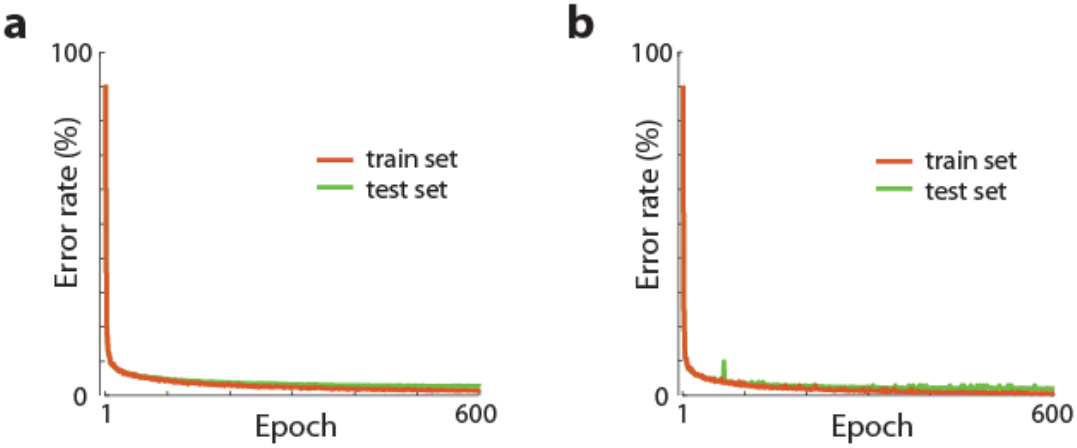

Suppl. Fig. 3. Additional biologically plausible network architectures. **a**, The learning curve for a network with 80% of excitatory and 20% inhibitory neurons in the hidden layer. For each excitatory neuron, all outgoing synaptic connections to output neurons were non-negative ($w \geq 0$), and for inhibitory neurons all outgoing connections to output neurons were negative or zero ($w \leq 0$). The network had the same architecture as our main network (784-1000-10 neurons), and was using the same learning rule (Eq. 3). On the test data set it achieved a comparable error rate of 2.66%. Those results are consistent with other models which showed that implementing Dale's law did not reduce network performance substantially (i.e.: Cornford et al. 2020). **b**, Learning curves for a network with asymmetric connections. The network had the same architecture as our main network (784-1000-10 neurons), and on the same MNIST dataset it achieved an error rate of 1.96%. Typical artificial neural networks trained with the backpropagation algorithm, require that a synaptic weight from neuron i to j (w_{ij}) is exactly the same as a synaptic weight from neuron j to i ($w_{ij}=w_{ji}$ is called symmetric weight). However, it was recently shown that networks without symmetric connections can still converge on a good solution (Lillicrap et al. 2016, Akrout et al. 2019). Similar results were also found by Detorakis et al. (2019) who used asymmetric weights with Contrastive Hebbian Learning. Thus, our results are consistent with that previous work and show that our learning rule also works in such more biologically plausible network configurations.

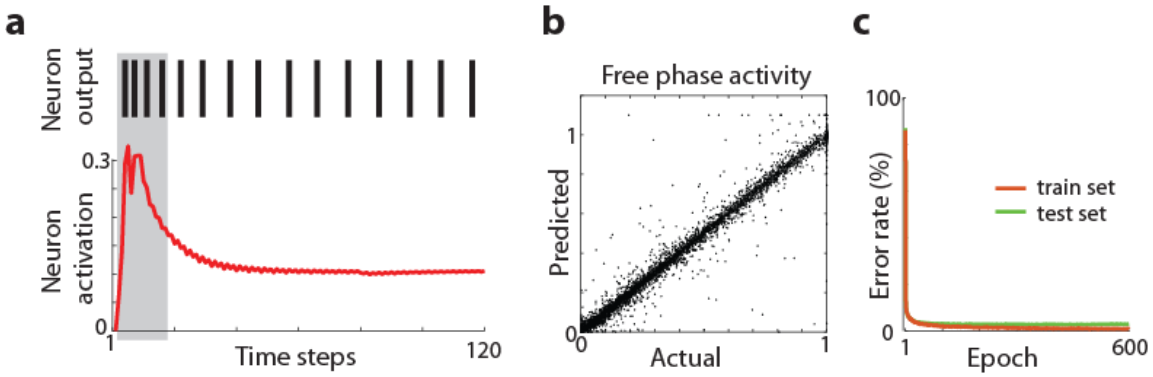

Suppl. Fig. 4. Implementation of predictive learning rule in a spiking neuronal network. **a**, Activity of a sample neuron. Red trace shows neuron internal activation which is a function of synaptic inputs. (Top) Spikes are generated based on the value of internal activation. Gray shaded area marks the extent of free phase which is used to predict steady-state activation. **b**, Actual vs. predicted steady-state free phase activity. Each dot represents the activity of one neuron during presentation of one stimuli. Activity of 1000 hidden neurons during presentation of 200 stimuli from the test dataset is shown. The correlation coefficient between actual and predicted activity was R=0.99 ($p<0.00001$). **c**, Learning curves on MNIST task for training and testing dataset. Error rate of 2.46% on the test set shows that the spiking neuronal network with our learning rule was able to solve the presented task.

Spiking network model.

The code for this network with all implementation details is provided at: https://github.com/ykubo82/bioCHL/tree/master/bioCHLspk . Briefly, our network was based on work by O'Connor et al. (2019) who developed spiking networks using the Equilibrium Propagation algorithm which is an expansion of Contrastive Hebbian Learning (Scellier et al., 2017). The steady state is calculated using the Forward Euler method with each time step t: $A_j^t = (1 - \varepsilon) A_j^{t-1} + \varepsilon p'(A_j^{t-1}) (\sum_i w_{ij}\ p(A_i^{t-1}) + b_j)$, where A is the activation state, $p'(A_j)$ is the derivative of the activation function p, b is the bias, and i and j are neurons indices. ε can be seen as the learning rate of the activations. These states are clipped to range [0, 1]. Each neuron communicates only using binary signals: 0 or 1, to model spiking activity. For generating spikes, an encoder converts the neuron activation state to 1 or 0, which is sent as an output signal (Suppl. Fig. 4a). The binary signals received by a neuron are converted into the activation state by a decoder. This could be seen as converting discrete spikes into a continuous post-synaptic membrane potential in actual neurons. We modified this spiking network by implementing our learning rule and by adding a least-squares model for predicting steady-state activation based on initial activation in steps 1-17. We also used AdaGrad (Duchi et al., 2011) to search for hyperparameters, which resulted in setting the learning rate to 0.01. The network architecture was the same as in our main model with 784-1000-10 neurons. Those results are consistent with other studies which showed that spiking networks can perform as well as non-spiking networks (i.e.: Zenke & Ganguli, 2018; Bellec et al. 2020).

Acoustic stimuli

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

Changes in neuronal activity across periods

For each animal we divided the duration of recording into two halves. For each neuron we calculated the mean stimulus evoked firing rate in 30-40 ms time window, averaged over all stimulus presentations in each half of the experiment. Similarly, for each neuron we calculated the average predicted activity for the 30-40 ms time window, averaged across all stimuli in the 1st half of the experiment. We then calculated the difference between activities in the 2nd half minus the 1st half, and we correlated it with the difference between activity in the 1st half minus predicted activity in the 1st half (Fig. 5b). As described in the Methods section, only neurons with average stimulus evoked firing rates higher than 3 SD above pre-stimulus baseline were used in our analyses. Dividing the data in different proportions e.g. 60%-40% gave the same conclusions.

For comparison, analogous analyses were done on artificial neurons (Fig. 5a). For each neuron in the hidden layer the steady-state clamped activity was averaged over all 1200 stimulus presentations in a single learning epoch. Similarly, for each neuron in the hidden layer, the predicted activity was averaged across all stimuli presented in an epoch. We then calculated the difference between clamped activities in 2 epoch: $\Delta C_i = \langle \hat{x}_{i,M} \rangle - \langle \hat{x}_{i,N} \rangle$, where <> denotes average, i is index of neuron, and M and N are indexes of epochs with $M > N$. This was then correlated with difference between average clamped and predicted activity in epoch N: $\Delta P_i = \langle \hat{x}_{i,N} \rangle - \langle \tilde{x}_{i,N} \rangle$. We found that the correlation between ΔC and ΔP was strongest in the earliest epochs where the learning curve was most changing (Fig. 2e). For this simulation we used neural network as described in Methods with 784-1000-10 neurons, and we used the learning rule described in Eq. 3. However, in order to reproduce the repeated presentation of the same stimuli as in animal experiments, we used the same 1200 stimulus in each learning epoch. Presenting different stimuli at each epoch as in our original simulation, gave qualitatively similar results.

Note that usually the activity of a neuron in the 1st half of recording is correlated with its activity in the 2nd half, which typically is a trivial experimental observation. However, it is nontrivial to correctly predict (only from activity in the 1st half) if that neuron activity will increase or if it will decrease in the 2nd half. This is exactly what we are presenting in Fig. 5b, that a change in average firing rate from the 1st to 2nd half can be correctly predicted based on our learning rule. To state it differently, the activity of a neuron in the 1st half (X_1) is usually correlated with its activity in the 2nd half (X_2), and this can be described as: $X_2 = X_1 + eps$, where *eps* is a small positive or negative number representing change in activity from the 1st to the 2nd half. Thus, in

Fig. 5b, we are not showing that X_1 is similar (correlated) to X_2, but rather that *eps* correlates with 'surprise': the difference between actual and predicted activity in the 1st half, which provides a novel insight about neuronal plasticity.

Predicting evoked dynamics from spontaneous activity

For predicting evoked dynamics from spontaneous, first we analyzed 'types' of spontaneous activity patterns. For this we divided spontaneous packets into 10 groups based on principal component analysis (PCA). Specifically, each packet was represented as a NxT matrix, where N is number of neurons and T=30 is number of 3ms long time bins. This matrix was then converted into a single vector V with length of NxT. This was repeated for each packet, thus we obtained a VxP matrix, where P is the number of spontaneous packets (P for each animal was: 293, 855, 1215 and 2472). We then ran PCA on the VxP matrix and based on values of the 1st PC we divided packets into 10 groups. Note that the distribution of packets in PCA space was continuous, thus from this analysis it should not be interpreted that there are distinct types of packets. Dividing packets into 5-20 groups and including the 2nd and 3rd PC gave consistent results with those presented in the main text.

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

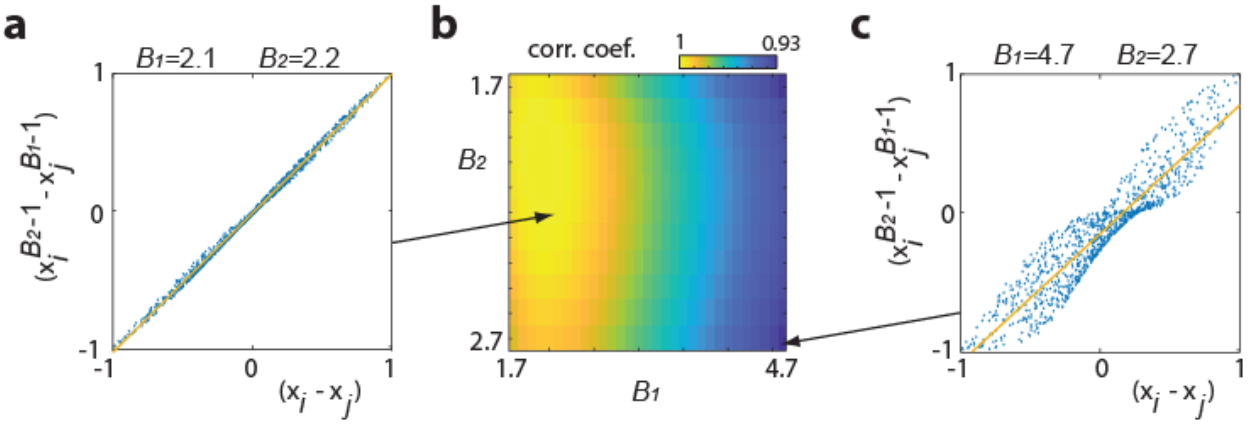

Suppl. Fig. 5. Testing effect of different values of β_1 and β_2. In Eq. 6 we used: $\beta_1 = 2$ and $\beta_2 = 2$, which allowed to reduce that formula from exponential to linear. Here we illustrate that resulted linearized expression in form: $(x_i - x_j)$ is a reasonable approximation of non-linear version: $({x_i}^{\beta_2 - 1} - {x_j}^{\beta_1 - 1})$ for β_1 and β_2 in range observed experimentally: $1.7 < \beta_1 < 4.8$ and $1.7 < \beta_2 < 2.7$ (Devor et al., 2003), and for typical values of x in range between 0-1. **a**, Relation between $(x_i - x_j)$ and $({x_i}^{\beta_2 - 1} - {x_j}^{\beta_1 - 1})$ for $\beta_1 = 2.1$ and $\beta_2 = 2.2$. Here we generated 1000 random values for x_i and separately for x_j, uniformly distributed on interval 0-1. Distribution of points along diagonal shows close to linear relation between results of both formulas. Least square regression line is shown in yellow. Correlation coefficient between points obtained with linear and non-linear formula was R=0.9988. **b**, Similarly, we calculated correlation coefficients between both formulas for all combinations of β_1 and β_2 in range: $1.7 \leq \beta_1 \leq 4.8$ and $1.7 \leq \beta_2 \leq 2.7$, using step size of 0.1. The lowest correlation coefficient was R = 0.9335 for $\beta_1 = 4.7$ and $\beta_2 = 2.7$. **c**, The same plot as in panel **a**, but for the "worst case" scenario with $\beta_1 = 4.7$ and $\beta_2 = 2.7$. Note that even in this case, the relation between both formulas is close to linear.

Maximizing future energy balance.

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

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

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

In the main text we show in simulations (Fig. 2) and in experimental data (Fig. 4), that for small *n*, activity of neuron *j* at time *t+n* could be approximated by a linear function of its activity at earlier time step *t*: $x_{j,t+n} = \lambda_j x_{j,t}$, where λ is a regression coefficient. Thus Eq. S1 can be rewritten as

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

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

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

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

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

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

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

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

Considering results from Suppl. Fig. 5, we can use $\beta_1 = 2$ and $\beta_2 = 2$, which allows to simplify and reorganize terms as follows:

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

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

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

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

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

Note also that this derivation of learning rule from metabolic principles is obtained simply by only considering input and output energy to the cell, without need for including specific metabolic interactions. This provides an important theoretical contribution by showing how complex relations between cell metabolism and plasticity could be described more simply.

Activation function

To simplify presented derivation of learning rule we used linear model of a neuron. However, we can obtain the same expression even if we use non-linear neuronal model with activation function like ReLU: $f(x) = x^+ = \max(0, x)$ (Glorot et al. 2011). This is because for $x \geq 0$: $\mathrm{ReLU}(x) = x$, which results in the same formulas as presented in our derivation. The only difference is that if we use ReLU then we need to write a separate expression for $x < 0$, to specify that in such case $\mathrm{ReLU}(x) = 0$. For sigmoid activation function, showing equivalence of the formulas could be more complicated. However, there are multiple arguments that ReLU could be a better model of biological neurons than logistic sigmoid neurons (Glorot et al. 2011), and our neural network simulations performed similarly for both neuron models.

Biological plausibility of predictive mechanism implementation within a neuron.

In this study, we propose that each neuron implements a linear predictive model, which can be written in general form as: $\tilde{x}_{(t)} = \lambda_{(1)} * x_{(t-1)} + \ldots + \lambda_{(n)} * x_{(t-n)} + constant$, where $x_{(t-n)}$ is past activity at time step t-n, n is a maximum number of time steps contributing to the prediction, and $\lambda_{(n)}$ is a coefficient describing how past activity at time step t-n contributes to predicted activity at time t. This may appear to be a challenging computation for a cell, considering that it requires storing values of many λ coefficients and to multiply past activity by a specific λ at each time step. Here we explain how this equation could be implemented by known intracellular processes, like regulation of ion concentration.

In neurons, an action potential generates a temporal increase in calcium concentration. This variation in calcium can be approximated by convolution of the action potential with the waveform of the unitary calcium transient (Yaksi et al. 2006). Note that convolution is a linear operation and can be mathematically expressed similarly to our predictive model. Thus, calcium concentration at time t can be described as: $y_{(t)} = \lambda_{(1)} * x_{(t-1)} + \ldots + \lambda_{(n)} * x_{(t-n)}$, where coefficients λ correspond to values of unitary calcium transient at time steps 1 ... n, and x is the past activity. For instance, if $\lambda_{(1)}$ has the highest value, then the above equation means that neuron activity at time t-1 is the biggest predictor of calcium concentration at time t. Consequently, decay of λ values for larger n, tells us how quickly calcium ions are cleaned from the cell at those time steps after the action potential. Therefore, values of λ are directly related to the physical properties of calcium channels, such as number of channels and ion transport efficiency, among other factors. This illustrates that our equation for the predictive model could be implemented within each neuron, simply as a physical process of activity-dependent changes in ions concentration.

Our model proposes that, when a neuron receives e.g. larger than predicted top-down modulation, then not only will synapses be modified as described in Eq. 3, but also predictive model parameters λ will be updated to reduce prediction error. As explained above, this could be accomplished by e.g. changing the efficiency or number of ion channels, which would modify the calcium dynamics used for predictions. Although this is only a conjecture, this offers an experimentally testable hypothesis, that the waveform of the unitary calcium transient should change, depending on the frequency at which a neuron was stimulated earlier.

Supplementary references

Akrout M, Wilson C, Humphreys PC, Lillicrap T, Tweed D. Deep learning without weight transport. arXiv preprint arXiv:1904.05391. (2019)

Bellec G, Scherr F, Subramoney A, Hajek E, Salaj D, Legenstein R, Maass W. A solution to the learning dilemma for recurrent networks of spiking neurons. Nature communications. (2020)

Cornford J, Kalajdzievski D, Leite M, Lamarquette A, Kullmann DM, Richards BA. Learning to live with Dale's principle: ANNs with separate excitatory and inhibitory units. ICLR 2021.

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

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

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

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

Glorot X, Bordes A, Bengio Y. Deep sparse rectifier neural networks. In Proceedings of the fourteenth international conference on artificial intelligence and statistics 2011 Jun 14 (pp. 315-323).

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

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

Lim S, McKee JL, Woloszyn L, Amit Y, Freedman DJ, Sheinberg DL, Brunel N. Inferring learning rules from distributions of firing rates in cortical neurons. Nature neuroscience. 2015 Dec;18(12):1804-10.

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

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

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

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

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

Woloszyn L, Sheinberg DL. Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. Neuron. 2012 Apr 12;74(1):193-205.

Yaksi E, Friedrich RW. Reconstruction of firing rate changes across neuronal populations by temporally deconvolved Ca 2+ imaging. Nature methods. 2006 May;3(5):377-83.

Zenke F, Ganguli S. Superspike: Supervised learning in multilayer spiking neural networks. Neural computation. 2018 Jun;30(6):1514-41.