

On the relationship between predictive coding and backpropagation

Robert Rosenbaum **

Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame, IN, United States of America

* Robert.Rosenbaum@nd.edu



Artificial neural networks are often interpreted as abstract models of biological neuronal networks, but they are typically trained using the biologically unrealistic backpropagation algorithm and its variants. Predictive coding has been proposed as a potentially more biologically realistic alternative to backpropagation for training neural networks. This manuscript reviews and extends recent work on the mathematical relationship between predictive coding and backpropagation for training feedforward artificial neural networks on supervised learning tasks. Implications of these results for the interpretation of predictive coding and deep neural networks as models of biological learning are discussed along with a repository of functions, Torch2PC, for performing predictive coding with PyTorch neural network models.



OPEN ACCESS

Citation: Rosenbaum R (2022) On the relationship between predictive coding and backpropagation. PLoS ONE 17(3): e0266102. https://doi.org/10.1371/journal.pone.0266102

Editor: Gennady S. Cymbalyuk, Georgia State University, UNITED STATES

Received: June 29, 2021

Accepted: March 14, 2022

Published: March 31, 2022

Copyright: © 2022 Robert Rosenbaum. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All code and data to reproduce the results in the paper have been uploaded to figshare (https://doi.org/10.6084/m9.figshare.19387409.v2).

Funding: This work was supported by National Science Foundation grants NSF DMS 1654268, NSF Neuronex DBI 1707400, and the Air Force Office of Scientific Research (ASOFR) award number FA9550-21-1-0223. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Introduction

The backpropagation algorithm and its variants are widely used to train artificial neural networks. While artificial and biological neural networks share some common features, a direct implementation of backpropagation in the brain is often considered biologically implausible in part because of the nonlocal nature of parameter updates: The update to a parameter in one layer depends on activity in all deeper layers. In contrast, biological neural networks are believed to learn largely through local synaptic plasticity rules for which changes to a synaptic weight depend on neural activity local to that synapse. While neuromodulators can have non-local impact on synaptic plasticity, they are not believed to be sufficiently specific to implement the precise, high-dimensional credit assignment required by backpropagation. However, some work has shown that global errors and neuromodulators can work with local plasticity to implement effective learning algorithms [1, 2]. Backpropagation can be performed using local updates if gradients of neurons' activations are passed upstream through feedback connections, but this interpretation implies other biologically implausible properties of the network, like symmetric feedforward and feedback weights. See previous work [3, 4] for a more complete review of the biological plausibility of backpropagation.

Several approaches have been proposed for achieving or approximating backpropagation with ostensibly more biologically realistic learning rules [2-14]. One such approach [11-14] is

Competing interests: The authors have declared that no competing interests exist.

derived from the theory of "predictive coding" or "predictive processing" [15–23]. A relationship between predictive coding and backpropagation was first discovered by Whittington and Bogacz [11] who showed that, when predictive coding is used to train a feedforward neural network on a supervised learning task, it can produce parameter updates that approximate those computed by backpropagation. These original results have since been extended to more general network architectures and to show that modifying predictive coding by a "fixed prediction assumption" leads to an algorithm that produces the exact same parameter updates as backpropagation [12–14].

This manuscript reviews and extends previous work [11-14] on the relationship between predictive coding and backpropagation, as well as some implications of these results on the interpretation of predictive coding and artificial neural networks as models of biological learning. The main results in this manuscript are as follows,

- Accounting for covariance or precision matrices in hidden layers does not affect parameter updates (learning) for predictive coding under the "fixed prediction assumption" used in previous work.
- 2. Predictive coding under the fixed prediction assumption is *algorithmically* equivalent to a direct implementation of backpropagation, which raises the question of whether it should be interpreted as more biologically plausible than backpropagation.
- 3. Empirical results show that the magnitude of prediction errors do not necessarily correspond to surprising features of inputs.

In addition, a public repository of Python functions, Torch2PC, is introduced. These functions can be used to perform predictive coding on any PyTorch Sequential model (see Materials and methods).

Results

A review of the relationship between backpropagation and predictive coding from previous work

For completeness, let us first review the backpropagation algorithm. Consider a feedforward deep neural network (DNN) defined by

$$\hat{v}_0 = x$$
 $\hat{v}_{\ell} = f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell}), \ \ell = 1, \dots, L$
(1)

where each \hat{v}_{ℓ} is a vector or tensor of activations, each θ_{ℓ} is a set of parameters for layer ℓ , and L is the network's depth. In supervised learning, one seeks to minimize a loss function $\mathcal{L}(\hat{y}, y)$ where y is a label associated with input, x, and

$$\hat{y} = f(x; \theta) = \hat{v}_L$$

is the network's output, which depends on parameters $\theta = \{\theta_\ell\}_{\ell=1}^L$. The loss is typically minimized using gradient-based optimization methods with gradients computed using automatic differentiation tools based on the backpropagation algorithm. For completeness, backpropagation is reviewed in the pseudocode below.

Algorithm 1 A standard implementation of backpropagation.

```
Given: Input (x) and label (y) # forward pass \hat{v}_0 = x for \ell = 1, \ldots, L
```

$$\begin{split} \hat{\nu}_{\ell} &= f_{\ell}(\hat{\nu}_{\ell-1}; \theta_{\ell}) \\ \# \text{ backward pass} \\ \delta_{L} &= \frac{\partial \mathcal{L}(\hat{\nu}_{L}, y)}{\partial \hat{\nu}_{L}} \\ \text{for } \ell &= L-1, \ldots, \\ \delta_{\ell} &= \delta_{\ell+1} \frac{\partial f_{\ell+1}(\hat{\nu}_{\ell}; \theta_{\ell+1})}{\partial \hat{\nu}_{\ell}} \\ d\theta_{\ell} &= -\delta_{\ell} \frac{\partial f_{\ell}(\hat{\nu}_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}} \end{split}$$

A direct application of the chain rule and mathematical induction shows that backpropagation computes the gradients,

$$\delta_{\ell} = \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \text{ and } d\theta_{\ell} = -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \theta_{\ell}}.$$

The negative gradients, $d\theta_{\ell}$, are then used to update parameters, either directly for stochastic gradient descent or indirectly for other gradient-based learning methods [24]. For the sake of comparison, I used backpropagation to train a 5-layer convolutional neural network on the MNIST data set (Fig 1A and 1B; blue curves). I next review algorithms derived from the theory of predictive coding and their relationship to backpropagation, as originally derived in previous work [11–14].

A strict interpretation of predictive coding does not accurately compute gradients. I begin by reviewing supervised learning under a strict interpretation of predictive coding. The formulation in this section is equivalent to the one first studied by Whittington and Bogacz [11], except that their results are restricted to the case in which $f_{\ell}(v_{\ell-1}; \theta_{\ell}) = \theta_{\ell}g_{\ell}(v_{\ell-1})$ for some point-wise-applied activation function, g_{ℓ} , and connectivity matrix, θ_{ℓ} . Our formulation extends this formulation to arbitrary vector-valued differentiable functions, f_{ℓ} . For the sake of

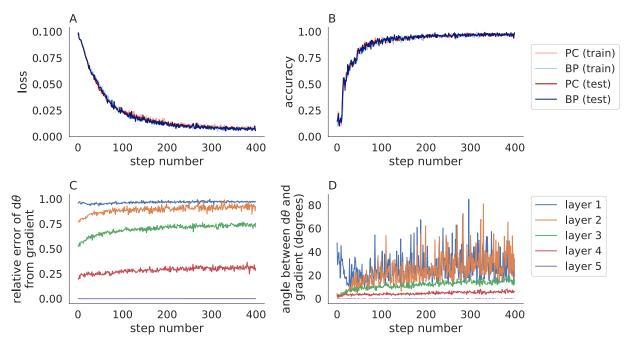


Fig 1. Comparing backpropagation and predictive coding in a convolutional neural network trained on MNIST. A,B) The loss (A) and accuracy (B) on the training set (pastel) and test set (dark) when a 5-layer network was trained using a strict implementation of predictive coding (Algorithm 2 with $\eta = 0.1$ and n = 20; red) and backpropagation (blue). C,D) The relative error (C) and angle (B) between the parameter update, $d\theta$, computed by Algorithm 2 and the negative gradient of the loss at each layer. Predictive coding and backpropagation give similar accuracies, but the parameter updates are less similar.

https://doi.org/10.1371/journal.pone.0266102.g001

continuity with later sections, I also use the notational conventions from [12] which differ from those in [11].

Predictive coding can be derived from a hierarchical, Gaussian probabilistic model in which each layer, ℓ , is associated with a Gaussian random variable, V_{ℓ} , satisfying

$$p(V_{\ell} = \nu_{\ell} | V_{\ell-1} = \nu_{\ell-1}) = \mathcal{N}(\nu_{\ell}; f_{\ell}(\nu_{\ell-1}; \theta_{\ell}), \Sigma_{\ell})$$
(2)

where $\mathcal{N}(\nu; \mu, \Sigma) \propto \exp(-[\nu - \mu]^T \Sigma^{-1} [\nu - \mu]/2)$ is the multivariate Gaussian distribution with mean, μ , and covariance matrix, Σ , evaluated at ν . Following previous work [11–14], I take $\Sigma = I$ to be the identity matrix, but later relax this assumption [21].

If we condition on an observed input, $V_0 = x$, then a forward pass through the network described by Eq (1) corresponds to setting $\hat{v}_0 = x$ and then sequentially computing the conditional expectations or, equivalently, maximizing conditional probabilities,

$$\begin{split} \hat{v}_{\ell} &= E[V_{\ell}|V_{\ell-1} = \hat{v}_{\ell-1}] \\ &= \operatorname*{argmax}_{v_{\ell}} p(V_{\ell} = v_{\ell}|V_{\ell-1} = \hat{v}_{\ell-1}) \\ &= f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell}) \end{split}$$

until reaching an inferred output, $\hat{y} = \hat{v}_L$. Note that this forward pass does not necessarily maximize the global conditional probability, $p(V_L = \hat{y} \mid v_0 = x)$ and it does not account for a prior distribution on V_L , which arises in related work on predictive coding for unsupervised learning [15, 21]. One interpretation of a forward pass is that each \hat{v}_ℓ is the network's "belief" about the state of V_ℓ , when only $V_0 = x$ has been observed.

Now suppose that we condition on both an observed input, $V_0 = x$, and its label, $V_L = y$. In this case, generating beliefs about the hidden states, V_ℓ , is more difficult because we need to account for potentially conflicting information at each end of the network. We can proceed by initializing a set of beliefs, v_ℓ , about the state of each V_ℓ , and then updating our initial beliefs to be more consistent with the observations, x and y, and parameters, θ_ℓ .

The error made by a set of beliefs, $\{\nu_\ell\}_{\ell=1}^L$, under parameters, $\{\theta_\ell\}_{\ell=1}^L$, can be quantified by

$$\epsilon_{\ell} = f_{\ell}(\nu_{\ell-1}; \theta_{\ell}) - \nu_{\ell}$$

for $\ell = 1, ..., L - 1$ where $\nu_0 = V_0 = x$ is observed. It is not so simple to quantify the error, ϵ_L , made at the last layer in a way that accounts for arbitrary loss functions. In the special case of a squared-Euclidean loss function,

$$\mathcal{L}(\hat{y}, y) = \frac{1}{2} ||\hat{y} - y||^2,$$

where $||u||^2 = u^T u$. Standard formulations of predictive coding [20, 21] use

$$\epsilon_L = f_L(\nu_{L-1}; \theta_L) - y \tag{3}$$

where recall that y is the label. In this case, ϵ_L satisfies

$$\epsilon_L = \frac{\partial \mathcal{L}(\tilde{v}_L, y)}{\partial \tilde{v}_L} \tag{4}$$

where

$$\tilde{\mathbf{v}}_{\scriptscriptstyle L} = f_{\scriptscriptstyle L}(\mathbf{v}_{\scriptscriptstyle L-1};\theta_{\scriptscriptstyle L}).$$

We use the \tilde{v}_L to emphasize that \tilde{v}_L is different from \hat{v}_L (which is defined by a forward pass starting at $\hat{v}_0 = x$) and is defined in a fundamentally different way from the v_ℓ terms (which do not

necessarily satisfy $v_{\ell} = f_{\ell}(v_{\ell-1}; \theta_{\ell})$). We can then define the total summed magnitude of errors as

$$F = \frac{1}{2} \sum_{\ell=1}^{L} \left\| \epsilon_{\ell} \right\|^{2}$$

More details on the derivation of F in terms of variational Bayesian inference can be found in previous work [12, 16, 20, 21] where F is known as the variational free energy of the model. Essentially, minimizing F produces a model that is more consistent with the observed data. Minimizing F by gradient descent on ν_ℓ and θ_ℓ produce the inference and learning steps of predictive coding, respectively.

Under a more heuristic interpretation, v_ℓ represents the network's "belief" about V_ℓ , and $f_\ell(v_{\ell-1};\theta_\ell)$ is the "prediction" of v_ℓ made by the previous layer. Under this interpretation, ϵ_ℓ is the error made by the previous layer's prediction, so ϵ_ℓ is called a "prediction error." Then F quantifies the total magnitude of prediction errors given a set of beliefs, v_ℓ , parameters, θ_ℓ , and observations, $V_0 = x$ and $V_L = y$.

In predictive coding, beliefs, ν_e , are updated to minimize the error, F. This can be achieved by gradient descent, *i.e.*, by making updates of the form

$$v_{\ell} \leftarrow v_{\ell} + \eta dv_{\ell}$$

where η is a step size and

$$d\nu_{\ell} = -\frac{\partial F}{\partial \nu_{\ell}}$$

$$= \epsilon_{\ell} - \epsilon_{\ell+1} \frac{\partial f_{\ell+1}(\nu_{\ell}; \theta_{\ell+1})}{\partial \nu_{\ell}}$$
(5)

In this expression, $\partial f_{\ell+1}(\nu_\ell;\theta_{\ell+1})/\partial \nu_\ell$ is a Jacobian matrix and $\epsilon_{\ell+1}$ is a row-vector to simplify notation, but a column-vector interpretation is similar. If x is a mini-batch instead of one data point, then ν_ℓ is an $m \times n_\ell$ matrix and derivatives are tensors. These conventions are used throughout the manuscript. The updates in Eq.(5) can be iterated until convergence or approximate convergence. Note that the prediction errors, $\epsilon_\ell = \nu_\ell - f_\ell(\nu_{\ell-1};\theta_\ell)$, should also be updated on each iteration.

Learning can also be phrased as minimizing F with gradient descent on parameters. Specifically,

$$\theta_{\ell} = \theta_{\ell} + \eta_{\theta} d\theta_{\ell}$$

where

$$d\theta_{\ell} = -\frac{\partial F}{\partial \theta_{\ell}}$$

$$= -\epsilon_{\ell} \frac{\partial f_{\ell}(\nu_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}}.$$
(6)

Note that some previous work uses the negative of the prediction errors used here, *i.e.*, they use $\epsilon_{\ell} = \nu_{\ell} - f_{\ell}(\nu_{\ell-1}; \theta_{\ell})$. While this choice changes some of the expressions above, the value of F and its dependence on θ_{ℓ} is not changed because F is defined by the norms of the ϵ_{ℓ} terms. The complete algorithm is defined more precisely by the pseudocode below:

Algorithm 2 A direct interpretation of predictive coding.

Given: Input (x), label (y), and initial beliefs (v_{ℓ})

```
 \begin{tabular}{ll} \# \mbox{ error and belief computation} \\ \begin{tabular}{ll} \begin{tabular}{ll} \begin{tabular}{ll} \# \mbox{ error and belief computation} \\ \begin{tabular}{ll} \
```

Here and elsewhere, n denotes the number of iterations for the inference step. The choice of initial beliefs is not specified in the algorithm above, but previous work [11–14] uses the results from a forward pass, $\nu_\ell = \hat{\nu}_\ell$, as initial conditions and I do the same in all numerical examples.

I tested Algorithm 2 on MNIST using a 5-layer convolutional neural network. To be consistent with the definitions above, I used a mean-squared error (squared Euclidean) loss function, which required one-hot encoded labels [24]. Algorithm 2 performed similarly to backpropagation (Fig 1A and 1B) even though the parameter updates did not match the true gradients (Fig 1C and 1D). Algorithm 2 was slower than backpropagation (31s for Algorithm 2 versus 8s for backpropagation when training metrics were not computed on every iteration) in part because Algorithm 2 requires several inner iterations to compute the prediction errors (n = 20 iterations used in this example). Algorithm 2 failed to converge on a larger model. Specifically, the loss grew consistently with iterations when trying to use Algorithm 2 to train the 6-layer CIFAR-10 model described in the next section. S1 Fig shows the same results from Fig 1 repeated across 30 trials with different random seeds to quantify the mean and standard deviation across trials.

Fig 1C and 1D shows that predictive coding does not update parameters according to the true gradients, but it is not immediately clear whether this would be resolved by using more iterations (larger n) or different values of the step size, η . I next compared the parameter updates, $d\theta_{\ell}$, to the true gradients, $\partial \mathcal{L}/\partial \theta_{\ell}$ for different values of n and η (Fig 2). For the smaller values of η tested ($\eta = 0.1$ and $\eta = 0.2$) and larger values of n (n > 100), parameter updates were similar to the true gradients in the last two layers, but they differed substantially in the first two layers. The largest values of η tested ($\eta = 0.5$ and $\eta = 1$) caused the iterations in Algorithm 2 to diverge.

Some choices in designing Algorithm 2 were made arbitrarily. For example, the three updates inside the inner for-loop over ℓ could be performed in a different order or the outer for-loop over i could be changed to a while-loop with a convergence criterion. For any initial conditions and any of these design choices, if the iterations over i are repeated until convergence or approximate convergence of each v_{ℓ} to a fixed point, v_{ℓ}^* , then the increments must satisfy $dv_{\ell} = 0$ at the fixed point and therefore the fixed point values of the prediction errors, ϵ_{ℓ}^* , must satisfy

$$\epsilon_{\ell}^* = \frac{\partial f_{\ell+1}(\nu_{\ell}^*; \theta_{\ell+1})}{\partial \nu_{\ell}^*} \epsilon_{\ell+1}^* \tag{7}$$

for $\ell = 1, ..., L - 1$. By the definition of ϵ_L , we have

$$\epsilon_L^* = \frac{\partial \mathcal{L}(\tilde{\nu}_L^*, y)}{\partial \tilde{\nu}_I^*}.$$
 (8)

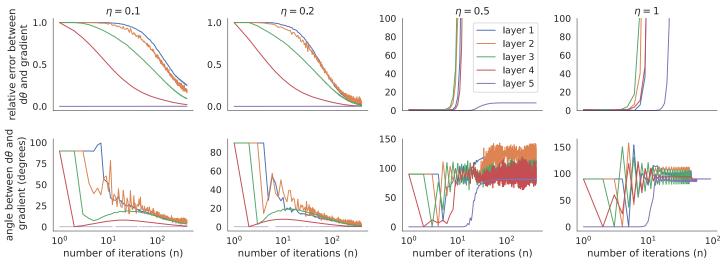


Fig 2. Comparing parameter updates from predictive coding to true gradients in a network trained on MNIST. Relative error and angle between $d\theta_\ell$ produced by predictive coding (Algorithm 2) as compared to the exact gradients, $\partial \mathcal{L}/\partial \theta_\ell$ computed by backpropagation (relative error defined by $\|d\theta_{pc} - d\theta_{bp}\|/\|d\theta_{bp}\|$). Updates were computed as a function of the number of iterations, n, used in Algorithm 2 for various values of the step size, η , using the model from Fig 1 applied to one minibatch of data. Both models were initialized identically to the pre-trained parameter values from the trained model in Fig 1. Parameter updates converge near the gradients after many iterations for smaller values of η , but diverge for larger values.

https://doi.org/10.1371/journal.pone.0266102.g002

where

$$\tilde{\nu}_L^* = f_L(\nu_{L-1}^*; \theta_L).$$

Combining Eqs (7) and (8) gives the fixed point prediction errors of the penultimate layer

$$\epsilon_{L-1}^{*} = \frac{\partial \mathcal{L}(\tilde{v}_{L}^{*}, y)}{\partial \tilde{v}_{L}^{*}} \frac{\partial f_{L}(v_{L-1}^{*}; \theta_{L})}{\partial v_{L-1}^{*}} \\
= \frac{\partial \mathcal{L}(\tilde{v}_{L}^{*}, y)}{\partial v_{L-1}^{*}} \tag{9}$$

where we used the fact that $\tilde{v}_L^* = f_L(v_{L-1}^*; \theta_L)$ and the chain rule. The error in layer L-2 is given by

$$\epsilon_{L-2}^* = \frac{\partial \mathcal{L}(\tilde{v}_L^*, y)}{\partial v_{L-1}^*} \frac{\partial f_{L-1}(v_{L-2}^*; \theta_{L-1})}{\partial v_{L-2}^*}.$$

Note that we cannot apply the chain rule to reduce this product (like we did for Eq.(9)) because it is not necessarily true that $v_{L-1}^* = f_{L-1}(v_{L-2}^*; \theta_{L-1})$. I revisit this point below. We can continue this process to derive

$$\epsilon_{L-3}^* = \frac{\partial \mathcal{L}(\tilde{v}_L^*, y)}{\partial v_{L-1}^*} \frac{\partial f_{L-1}(v_{L-2}^*; \theta_{L-1})}{\partial v_{L-2}^*} \frac{\partial f_{L-2}(v_{L-3}^*; \theta_{L-2})}{\partial v_{L-3}^*}$$

and continue for $\ell = L - 4, ..., 1$. In doing so, we see (by induction) that ϵ_{ℓ}^* can be written as

$$\epsilon_{\ell}^* = \frac{\partial \mathcal{L}(\tilde{v}_L^*, y)}{\partial \tilde{v}_{L-1}^*} \prod_{\ell'=\ell}^{L-2} \frac{\partial f_{\ell'+1}(v_{\ell'}^*; \theta_{\ell'+1})}{\partial v_{\ell'}^*}. \tag{10}$$

for $\ell = 1, ..., L - 2$. Therefore, if the inference loop converges to a fixed point, then the

subsequent parameter update obeys

$$d\theta_{\ell} = -\frac{\partial \mathcal{L}(\tilde{v}_{L}^{*}, y)}{\partial \tilde{v}_{L-1}^{*}} \left[\prod_{\ell'=\ell}^{L-2} \frac{\partial f_{\ell'+1}(v_{\ell'}^{*}; \theta_{\ell'+1})}{\partial v_{\ell'}^{*}} \right] \frac{\partial f_{\ell}(v_{\ell-1}^{*}; \theta_{\ell})}{\partial \theta_{\ell}}$$
(11)

by Eq (6). It is not clear whether there is a simple mathematical relationship between these parameter updates and the negative gradients, $d\theta_\ell = -\partial \mathcal{L}/\partial \theta_\ell$, computed by backpropagation.

It is tempting to assume that $v_\ell^* = f_\ell(v_{\ell-1}^*; \theta_\ell)$, in which case the product terms would be reduced by the chain rule. Indeed, this assumption would imply that $v_\ell^* = \hat{v}_\ell$ and $\tilde{v}_L^* = \hat{v}_L$ and, finally, that $\epsilon_\ell = \partial \mathcal{L}/\partial \hat{v}_\ell$ and $d\theta_\ell = -\partial \mathcal{L}/\partial \theta_\ell$, identical to the values computed by backpropagation. However, we cannot generally expect to have $v_\ell^* = f_\ell(v_{\ell-1}^*; \theta_\ell)$ because this would imply that $\epsilon_\ell^* = 0$ and therefore $\partial \mathcal{L}/\partial v_\ell^* = \partial \mathcal{L}/\partial \theta_\ell = 0$. In other words, Algorithm 2 is only equivalent to backpropagation in the case where parameters are at a critical point of the loss function, so all updates are zero. Nevertheless, this thought experiment suggests a modification to Algorithm 2 for which the fixed points do represent the true gradients [11, 12]. I review that modification in the next section.

Note also that the calculations above rely on the assumption of a Euclidean loss function, $\mathcal{L}(\hat{y},y) = \|\hat{y} - y\|^2/2$. If we want to generalize the algorithm to different loss functions, then Eqs (3) and (4) could not both be true, and therefore Eqs (7) and (8) could not both be true. This leaves open the question of how to define ϵ_L when using loss functions that are not proportional to the squared Euclidean norm. If we were to define ϵ_L by (3), at the expense of losing (4), then the algorithm would not account for the loss function at all, so it would effectively assume a Euclidean loss, *i.e.*, it would compute the same values that are computed by Algorithm 2 with a Euclidean loss. If we instead were to define ϵ_L by Eq. (4) at the expense of (3), then Eqs (5) and (7) would no longer be true for $\ell = L - 1$ and Eq. (6) would no longer be true for $\ell = L$. Instead, all three of these equations would involve second-order derivatives of the loss function, and therefore the fixed point Eqs (10) and (11) would also involve second order derivatives. The interpretation of the parameter updates is not clear in this case. One might instead try to define ϵ_L by the result of a forward pass,

$$\epsilon_L = f_L(\hat{v}_{L-1}; \theta_L) - y$$
$$= \hat{v}_L - y$$

but then ϵ_L would be a constant with respect to ν_{L-1} , so we would have $\partial \epsilon_L / \partial \nu_{L-1} = 0$, and therefore Eq.(5) at $\ell = L - 1$ would become

$$dv_{L-1} = -\frac{\partial F}{\partial v_{L-1}}$$
$$= \epsilon_{L-1}$$

which has a fixed point at $\epsilon_{L-1}^*=0$. This would finally imply that all the errors converge to $\epsilon_\ell^*=0$ and therefore $d\theta_\ell=0$ at the fixed point.

I next discuss a modification of Algorithm 2 that converges to the same gradients computed by backpropagation, *and* is applicable to general loss functions [11, 12].

Predictive coding modified by the fixed prediction assumption converges to the gradients computed by backpropagation. Previous work [11, 12] proposed a modification of the predictive coding algorithm described above called the "fixed prediction assumption" which I now review. Motivated by the considerations in the last few paragraphs of the previous section, we can selectively substitute some terms of the form v_{ℓ} and $f_{\ell}(v_{\ell-1}; \theta_{\ell})$ in Algorithm 2 with \hat{v}_{ℓ}

(or, equivalently, $f_\ell(\hat{v}_{\ell-1};\theta_\ell)$) where \hat{v}_ℓ are the results of the original forward pass starting from $\hat{v}_0=x$. Specifically, the following modifications are made to the quantities computed by Algorithm 2

$$\epsilon_{\ell} = \hat{v}_{\ell} - v_{\ell}
\epsilon_{L} = \frac{\partial \mathcal{L}(\hat{v}_{L}, y)}{\partial \hat{v}_{L}}
dv_{\ell} = \epsilon_{\ell} - \epsilon_{\ell+1} \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}}
d\theta_{\ell} = -\epsilon_{\ell} \frac{\partial f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}}$$
(12)

for $\ell = 1, ..., L - 1$. This modification can be interpreted as "fixing" the predictions at the values computed by a forward pass and is therefore called the "fixed prediction assumption" [11, 12]. Additionally, the initial conditions of the beliefs are set to the results from a forward pass, $\nu_{\ell} = \hat{\nu}_{\ell}$ for $\ell = 1, ..., L - 1$. The complete modified algorithm is defined by the pseudocode below:

Algorithm 3 Supervised learning with predictive coding modified by the fixed prediction assumption. Adapted from the algorithm in [12] and similar to the algorithm from [11].

```
Given: Input (x) and label (y)
# forward pass
\hat{v}_0 = x
for \ell = 1, \ldots, L
     \hat{\mathbf{v}}_{\ell} = f_{\ell}(\hat{\mathbf{v}}_{\ell-1}; \theta_{\ell})
     v_{\ell} = \hat{v}_{\ell}
# error and belief computation
\epsilon_{\scriptscriptstyle L} = {\partial \mathcal{L}(\hat{\mathbf{v}}_{\scriptscriptstyle L},\,\mathbf{y}) \over \partial \hat{\mathbf{v}}_{\scriptscriptstyle I}}
for i = 1, ..., n
      for \ell = L - 1, ..., 1
           \epsilon_\ell = 
u_\ell - \hat{
u}_\ell
           dv_{\ell} = \epsilon_{\ell} - \epsilon_{\ell}
dv_{\ell} = \epsilon_{\ell} - \epsilon_{\ell+1} \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}}
           v_{\ell} = v_{\ell} + \eta dv_{\ell}
# parameter update computation
for \ell = 1, ..., L
     d	heta_\ell = -\epsilon_\ell rac{\partial f_\ell(\hat{v}_{\ell-1};\,	heta_\ell)}{\partial 	heta_\ell}
```

Note, again, that some choices in Algorithm 3 were made arbitrarily. The three updates inside the inner for-loop over ℓ could be performed in a different order or the outer for loop over i could be changed to a while-loop with a convergence criterion. Regardless of these choices, the fixed points, ϵ_{ℓ}^* , can again be computed by setting $dv_{\ell} = 0$ to obtain

$$\epsilon_{\ell}^* = \epsilon_{\ell+1}^* rac{\partial f_{\ell+1}(\hat{m{v}}_{\ell}; m{ heta}_{\ell+1})}{\partial \hat{m{v}}_{\ell}}.$$

Now note that ϵ_L is fixed, so

$$\epsilon_{\scriptscriptstyle L}^* = \frac{\partial \mathcal{L}(\hat{v}_{\scriptscriptstyle L}, y)}{\partial \hat{v}_{\scriptscriptstyle L}}$$

and we can combine these two equations to compute

$$\begin{split} \epsilon_{L-1}^* &= \frac{\partial \mathcal{L}(\hat{v}_L, y)}{\partial \hat{v}_L} \frac{\partial f_L(\hat{v}_{L-1}; \theta_L)}{\partial \hat{v}_{L-1}} \\ &= \frac{\partial \mathcal{L}(\hat{v}_L, y)}{\partial \hat{v}_{L-1}} \end{split}$$

where we used the chain rule and the fact that $\hat{v}_{\ell} = f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})$. Continuing this approach we have,

$$\begin{split} \epsilon_{\ell}^* &= \epsilon_{\ell+1}^* \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}} \\ &= \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \end{split}$$

for all $\ell = 1, ..., L$ (where recall that $\hat{y} = \hat{v}_L$ is the output from the feedfoward pass). Combining this with the modified definition of $d\theta_\ell$, we have

$$d\theta_{\ell} = -\epsilon_{\ell}^* \frac{\partial f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}}$$

$$= -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \frac{\partial \hat{v}_{\ell}}{\partial \theta_{\ell}}$$

$$= -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \theta_{\ell}}$$

where we use the chain rule and the fact that $\hat{v}_\ell = f_\ell(\hat{v}_{\ell-1};\theta_\ell)$. We may conclude that, if the inference step converges to a fixed point $(dv_\ell = 0)$, then Algorithm 3 computes the same values of $d\theta_\ell$ as backpropagation and also that the prediction errors, ϵ_ℓ , converge to the gradients, $\delta_\ell = \partial \mathcal{L}/\partial \hat{v}_\ell$, computed by backpropagation. As long as the inference step *approximately* converges to a fixed point $(dv_\ell \approx 0)$, then we should expect the parameter updates from Algorithm 3 to *approximate* those computed by backpropagation. In the next section, I extend this result to show that a special case of the algorithm computes the true gradients in a fixed number of steps.

I next tested Algorithm 3 on MNIST using the same 5-layer convolutional neural network considered above. I used a cross-entropy loss function, but otherwise used all of the same parameters used to test Algorithm 2 in Fig 1. The modified predictive coding algorithm (Algorithm 3) performed similarly to backpropagation in terms of the loss and accuracy (Fig 3A and 3B). Parameter updates computed by Algorithm 3 did not match the true gradients, but pointed in a similar direction and provided a closer match than Algorithm 2 (compare Fig 3C and 3D to Fig 1C and 1D). Algorithm 3 was similar to Algorithm 2 in terms of training time (29s for Algorithm 3 versus 31s for Algorithm 2 and 8s for backpropagation). S2 Fig shows the same results from Fig 3 repeated across 30 trials with different random seeds to quantify the mean and standard deviation across trials.

I next compared the parameter updates computed by Algorithm 3 to the true gradients for different values of n and η (Fig 4). When η < 1, the parameter updates, $d\theta_\ell$, appeared to converge, but did not converge exactly to the true gradients. This is likely due to numerical floating point errors accumulated over iterations. When η = 1, the parameter updates at each layer remained constant for the first few iterations, then immediately jumped to become very near the updates from backpropagation. In the next section, I provide a mathematical analysis of

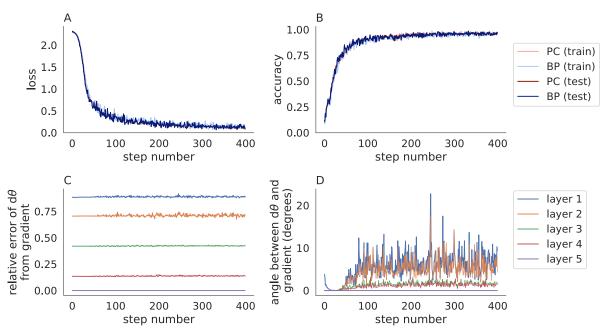


Fig 3. Predictive coding modified by the fixed prediction assumption compared to backpropagation in a convolutional neural network trained on MNIST. Same as Fig 1 except Algorithm 3 was used (with $\eta=0.1$ and n=20) in place of Algorithm 2. The accuracy of predictive coding with the fixed prediction assumption is similar to backpropagation, but the parameter updates are less similar for these hyperparameters.

https://doi.org/10.1371/journal.pone.0266102.g003

this behavior and show that when $\eta = 1$, Algorithm 3 computes the true gradients in a fixed number of steps.

To see how well these results extend to a larger model and more difficult benchmark, I next tested Algorithm 3 on CIFAR-10 [25] using a six-layer convolutional network. While the

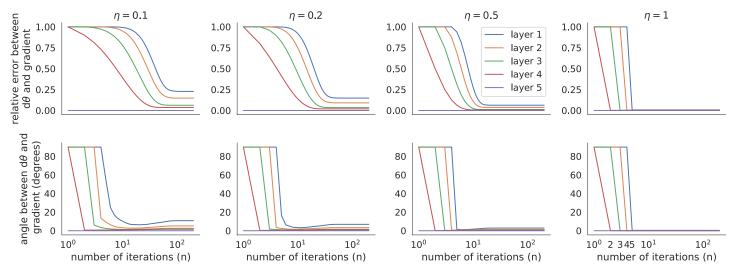


Fig 4. Comparing parameter updates from predictive coding modified by the fixed prediction assumption to true gradients in a network trained on MNIST. Relative error and angle between $d\theta$ produced by predictive coding modified by the fixed prediction assumption (Algorithm 3) as compared to the exact gradients computed by backpropagation (relative error defined by $\|d\theta_{pc} - d\theta_{bp}\|/\|d\theta_{bp}\|$). Updates were computed as a function of the number of iterations, n, used in Algorithm 3 for various values of the step size, η , using the model from Fig 3 applied to one mini-batch of data. Both models were initialized identically to the pre-trained parameter values from the backpropagation-trained model in Fig 3. In the rightmost panels, some lines are not visible where they overlap at zero. Parameter updates quickly converge to the true gradients when η is larger.

https://doi.org/10.1371/journal.pone.0266102.g004

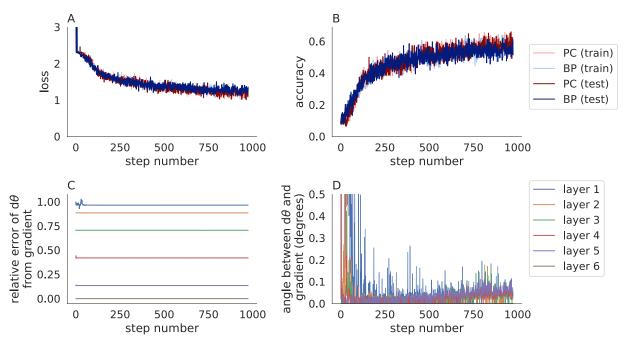


Fig 5. Predictive coding modified by the fixed prediction assumption compared to backpropagation in convolutional neural networks trained on CIFAR-10. Same as Fig 3 except a larger network was trained on the CIFAR-10 data set. The accuracy of predictive coding with the fixed prediction assumption is similar to backpropagation and parameter updates are similar to the true gradients.

https://doi.org/10.1371/journal.pone.0266102.g005

network only had one more layer than the MNIST network used above, it had 141 times more parameters (32,695 trainable parameters in the MNIST model versus 4,633,738 in the CIFAR-10 model). Algorithm 3 performed similarly to backpropagation in terms of loss and accuracy during learning (Fig 5A and 5B) and produced parameter updates that pointed in a similar direction, but still did not match the true gradients (Fig 5C and 5D). Algorithm 3 was substantially slower than backpropagation (848s for Algorithm 3 versus 58s for backpropagation when training metrics were not computed on every iteration).

Predictive coding modified by the fixed prediction assumption using a step size of $\eta = 1$ computes exact gradients in a fixed number of steps. A major disadvantage of the approach outlined above—when compared to standard backpropagation—is that it requires iterative updates to v_{ℓ} and ϵ_{ℓ} . Indeed, previous work [12] used n = 100-200 iterations, leading to substantially slower performance compared to standard backpropagation. Other work [11] used n = 20 iterations as above. In general, there is a tradeoff between accuracy and performance when choosing n, as demonstrated in Fig 4. However, more recent work [13, 14] showed that, under the fixed prediction assumption, predictive coding can compute the exact same gradients computed by backpropagation in a fixed number of steps. That work used a more specific formulation of the neural network which can implement fully connected layers, convolutional layers, and recurrent layers. They also used an unconventional interpretation of neural networks in which weights are multiplied outside the activation function, *i.e.*, $f_{\ell}(x; \theta_{\ell}) = \theta_{\ell}g_{\ell}(x)$, and inputs are fed into the last layer instead of the first. Next, I show that their result holds for arbitrary feedforward neural networks as formulated in Eq (1) (with arbitrary functions, f_{ℓ}) and this result has a simple interpretation in terms of Algorithm 3. Specifically, the following theorem shows that taking a step size of $\eta = 1$ yields an exact computation of gradients using just n = L iterations (where L is the depth of the network).

Theorem 1. If Algorithm 3 is run with step size $\eta = 1$ and at least n = L iterations then the algorithm computes

$$\epsilon_{\ell} = \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}}$$

and

$$d heta_\ell = -rac{\partial \mathcal{L}(\hat{y}, y)}{\partial heta_\ell}$$

for all $\ell = 1, ..., L$ where $\hat{v}_{\ell} = f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})$ are the results from a forward pass with $\hat{v}_0 = x$ and $\hat{y} = \hat{v}_L = f(x; \theta)$ is the output.

Proof. For the sake of notational simplicity within this proof, define $\delta_\ell = \partial \mathcal{L}(\hat{v}_L, y)/\partial \hat{v}_\ell$. Therefore, we first need to prove that $\epsilon_\ell = \delta_\ell$. First, rewrite the inside of the error and belief loop from Algorithm 3 while explicitly keeping track of the iteration number in which each variable was updated,

$$egin{array}{ll} \epsilon^i_\ell &=
u^{i-1}_\ell - \hat{
u}_\ell \ & d
u^i_\ell &= \epsilon^i_\ell - \epsilon^i_{\ell+1} rac{\partial f_{\ell+1}(\hat{
u}_\ell; heta_{\ell+1})}{\partial \hat{
u}_\ell} \ &
u^i_\ell &=
u^{i-1}_\ell + d
u^i_\ell. \end{array}$$

Here, v_ℓ^i , ϵ_ℓ^i , and dv_ℓ^i denote the values of v_ℓ^i , ϵ_ℓ^i , and dv_ℓ^i respectively at the end of the ith iteration, $v_\ell^0 = \hat{v}_\ell$ corresponds to the initial value, and all terms without superscripts are constant inside the inference loop. There are some subtleties here. For example, we have v_ℓ^{i-1} in the first line because v_ℓ is updated after ϵ_ℓ in the loop. More subtly, we have $\epsilon_{\ell+1}^i$ in the second equation instead of $\epsilon_{\ell+1}^{i-1}$ because the for loop goes backwards from $\ell=L-1$ to $\ell=1$, so $\epsilon_{\ell+1}$ is updated before ϵ_ℓ . First note that

$$\epsilon^1_{\ell} = 0$$

for $\ell=1,\ldots,L-1$ because $v_\ell^0=\hat{v}_\ell$. Now compute the change in ϵ_ℓ across one step,

$$\begin{split} \epsilon_\ell^{i+1} - \epsilon_\ell^i &= \nu_\ell^i - \nu_\ell^{i-1} \\ &= d\nu_\ell^i \\ &= \epsilon_\ell^i - \epsilon_{\ell+1}^i \frac{\partial f_{\ell+1}(\hat{\nu}_\ell; \theta_{\ell+1})}{\partial \hat{\nu}_\ell}. \end{split}$$

Note that this equation is only valid for $i \ge 1$ due to the i-1 term (ν_{ℓ}^{-1} is not defined). Adding ϵ_{ℓ}^{i} to both sides of the resulting equation gives

$$\epsilon_\ell^{i+1} = \epsilon_{\ell+1}^i rac{\partial f_{\ell+1}(\hat{m{v}}_\ell; m{ heta}_{\ell+1})}{\partial \hat{m{v}}_\ell}.$$

We now use induction to prove that $\epsilon_\ell = \delta_\ell$ after n = L iterations. Indeed, we prove a stronger claim that $\epsilon_\ell^i = \delta_\ell$ at $i = L - \ell + 1$. First note that $\epsilon_L^i = \delta_L$ for all i because ϵ_L^i is initialized to δ_L and then never changed. Therefore, our claim is true for the base case $\ell = L$.

Now suppose that $\epsilon_{\ell+1}^i = \delta_{\ell+1}$ for $i = L - (\ell+1) + 1 = L - \ell$. We need to show that $\epsilon_{\ell}^{i+1} = \delta_{\ell}$. From above, we have

$$\begin{split} \epsilon_{\ell}^{i+1} &= \epsilon_{\ell+1}^{i} \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}} \\ &= \delta_{\ell+1} \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}} \\ &= \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}} \\ &= \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell+1}} \frac{\partial \hat{v}_{\ell+1}}{\partial \hat{v}_{\ell}} \\ &= \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \frac{\partial \hat{v}_{\ell+1}}{\partial \hat{v}_{\ell}} \\ &= \delta_{\ell}. \end{split}$$

This completes our induction argument. It follows that $\epsilon_\ell^i = \delta_\ell$ at iteration $i = L - \ell + 1$ at all layers $\ell = 1, ..., L$. The last layer to be updated to the correct value is $\ell = 1$, which is updated on iteration number i = L - 1 + 1 = L. Hence, $\epsilon_\ell = \delta_\ell$ for all $\ell = 1, ..., L$ after n = L iterations. This proves the first statement in our theorem. The second statement then follows from the definition of $d\theta_\ell$,

$$\begin{split} d\theta_{\ell} &= -\epsilon_{\ell} \frac{\partial f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}} \\ &= -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \frac{\partial f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}} \\ &= -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}} \frac{\partial \hat{v}_{\ell}}{\partial \theta_{\ell}} \\ &= -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \theta_{\ell}}. \end{split}$$

This completes the proof.

This theorem ties together the implementation and formulation of predictive coding from [12] (*i.e.*, Algorithm 3) to the results in [13, 14]. As noted in [13, 14], this result depends critically on the assumption that the values of ν_ℓ are initialized to the activations from a forward pass, $\nu_\ell = \hat{\nu}_\ell$ initially. The theoretical predictions from Theorem 1 are confirmed by the fact that all of the errors in the rightmost panels of Fig 4 converge to zero after n = L = 5 iterations.

To further test the result empirically, I repeated Figs 3 and 5 using $\eta = 1$ and n = L (in contrast to Figs 3 and 5 which used $\eta = 0.1$ and n = 20). The loss and accuracy closely matched those computed by backpropagation (Figs 6A and 6B and 7A and 7B). More importantly, the parameter updates closely matched the true gradients (Figs 6C and 6D and 7C and 7D), as predicted by Theorem 1. The differences between predictive coding and backpropagation in Fig 6 were due floating point errors and the non-determinism of computations performed on GPUs. For example, similar differences to those seen in Fig 6A and 6B were present when the same training algorithm was run twice with the same random seed. The smaller number of iterations (n = L in Figs 6 and 7 versus n = 20 in Figs 3 and 5) resulted in a shorter training time (13s for MNIST and 300s for CIFAR-10 for Figs 6 and 7, compare to 29s and 848s in Figs 3 and 5, and compare to 8s and 58s for backpropagation).

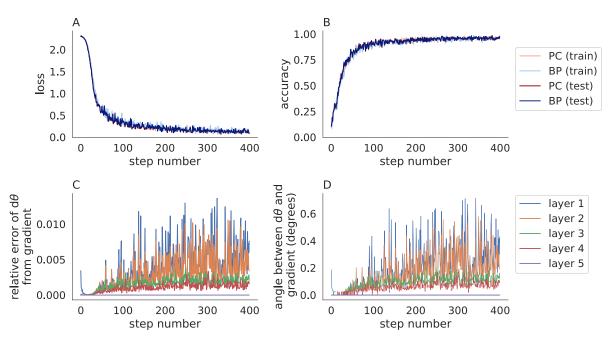


Fig 6. Predictive coding modified by the fixed prediction assumption with $\eta = 1$ compared to backpropagation in convolutional neural networks trained on MNIST. Same as Fig 3 except $\eta = 1$ and n = L. Predictive coding with the fixed prediction assumption approximates true gradients accurately when $\eta = 1$.

https://doi.org/10.1371/journal.pone.0266102.g006

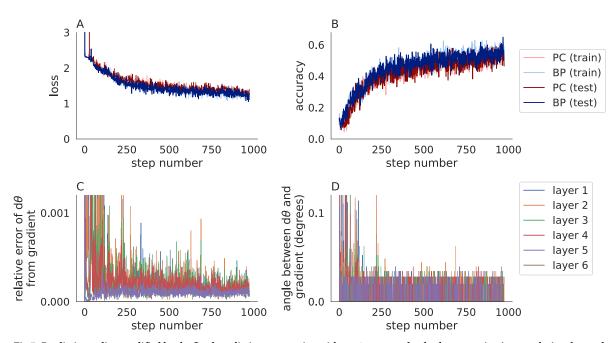


Fig 7. Predictive coding modified by the fixed prediction assumption with $\eta = 1$ compared to backpropagation in convolutional neural networks trained on CIFAR-10. Same as Fig 5 except $\eta = 1$ and n = L. Predictive coding with the fixed prediction assumption approximates true gradients accurately when $\eta = 1$.

https://doi.org/10.1371/journal.pone.0266102.g007

In summary, a review of the literature shows that a strict interpretation of predictive coding (Algorithm 2) does not converge to the true gradients computed by backpropagation. To compute the true gradients, predictive coding must be modified by the fixed prediction assumption (Algorithm 2). Further, I proved that Algorithm 2 computes the exact gradients when $\eta = 1$ and $n \ge L$, which ties together results from previous work [12–14].

Predictive coding with the fixed prediction assumption and $\eta = 1$ is functionally equivalent to a direct implementation of backpropagation

The proof of Theorem 1 and the last panel of Fig 4 give some insight into a how Algorithm 3 works. First note that the values of ν_ℓ in Algorithm 3 are only used to compute the values of ϵ_ℓ and are not otherwise used in the computation of $d\theta_\ell$ or any other quantities. Therefore, if we only care about understanding parameter updates, $d\theta_\ell$, we can ignore the values of ν_ℓ and only focus on how ϵ_ℓ is updated on each iteration, i. Secondly, note that when $\eta=1$, each ϵ_ℓ is updated only once: $\epsilon_\ell^i=0$ for $i< L-\ell+1$ and $\epsilon_\ell^i=\epsilon_{\ell+1}^i\partial f_{\ell+1}(\hat{\nu}_\ell;\theta_{\ell+1})/\partial\hat{\nu}_\ell$ for $i\geq L-\ell+1$, so ϵ_ℓ is only changed on iteration number $i=L-\ell+1$. In other words, the error computation in Algorithm 3 when $\eta=1$ and $\eta=L$ is equivalent to

```
\label{eq:epsilon} \begin{split} & \# \text{ error computation } \\ & \epsilon_L = \frac{\partial \mathcal{L}(\hat{v}_L, y)}{\partial \hat{v}_L} \\ & \text{ for } i = 1, \dots, L \\ & \text{ for } \ell = L-1, \dots, 1 \\ & \text{ if } \ell = L-i+1 \\ & \epsilon_\ell = \epsilon_{\ell+1} \frac{\partial f_{\ell+1}(\hat{v}_\ell; \, \theta_{\ell+1})}{\partial \hat{v}_\ell} \end{split}
```

The two computations are equivalent in the sense that they compute the same values of the errors, ϵ_ℓ^i , on every iteration. The formulation above makes it clear that the nested loops are unnecessary because for each value of i, ϵ_ℓ is only updated at one value of ℓ . Therefore, the nested loops and if-statement can be replaced by a single for-loop. Specifically, the error computation in Algorithm 3 when $\eta=1$ is equivalent to

```
\begin{split} & \text{\# error computation} \\ & \epsilon_L = \frac{\partial \mathcal{L}(\hat{v}_L, \, y)}{\partial \hat{v}_L} \\ & \text{for } \ell = L - 1, \, \dots, \, 1 \\ & \epsilon_\ell = \epsilon_{\ell+1} \frac{\partial f_{\ell+1}(\hat{v}_\ell; \, \theta_{\ell+1})}{\partial \hat{v}_\ell} \end{split}
```

This is *exactly* the error computation from the standard backpropagation algorithm, *i.e.*, Algorithm 1. Hence, if we use $\eta = 1$, then Algorithm 3 is just backpropagation with extra steps and these extra steps do not compute any non-zero values. If we additionally want to compute the fixed point beliefs, then they can still be computed using the relationship

$$v_{\ell} = \epsilon_{\ell} + \hat{v}_{\ell}.$$

We may conclude that, when $\eta=1$, Algorithm 3 can be replaced by an exact implementation of backpropagation without any effect on the results or effective implementation of the algorithm. This raises the question of whether predictive coding with the fixed prediction assumption should be considered any more biologically plausible than a direct implementation of backpropagation.

Accounting for covariance or precision matrices in hidden layers does not affect learning under the fixed prediction assumption

Above, I showed that predictive coding with the fixed prediction assumption is functionally equivalent to backpropagation. However, the predictive coding algorithm was derived under an assumption that covariance matrices in the probabilistic model are identity matrices, $\Sigma_{\ell} = I$.

This raises the question of whether relaxing this assumption could generalize backpropagation to account for the covariances, as suggested in previous work [11, 12, 26].

We can account for covariances by returning to the calculations starting from the probabilistic model in Eq. (2) and omit the assumption that $\Sigma_{\ell} = I$. To this end, it is helpful to define the precision-weighted prediction errors [20, 21, 26],

$$\tilde{\epsilon}_{\ell} = \epsilon_{\ell} P_{\ell}$$

for $\ell = 1, ..., L - 1$ where $P_{\ell} = \Sigma_{\ell}^{-1}$ is the inverse of the covariance matrix of V_{ℓ} , which is called "precision matrix." Recall that we treat ϵ_ℓ as a row-matrix, which explains the right-multiplication in this definition.

Modifying the definition of ϵ_L to account for covariances is not so simple because the Gaussian model for V_{ℓ} is not justified for non-Euclidean loss functions such as categorical loss functions. Moreover, it is not clear how to define the covariance or precision matrix of the output layer when labels are observed. As such, I restrict to accounting for precision matrices in hidden layers only, and leave the question of accounting for covariances in the output layer for future work with some comments on the issue provided at the end of this section. To this end, let us not modify the last layer's precision and instead define

$$\tilde{\epsilon}_L = \epsilon_L = \frac{\partial L(\hat{y}, y)}{\partial \hat{y}}.$$

The free energy is then defined as [20, 21]

$$F = \frac{1}{2} \sum_{\ell=1}^{L} \|\tilde{\epsilon}_{\ell}\|^{2}.$$

Performing gradient descent on F with respect to v_{ℓ} therefore gives

$$d v_\ell = ilde{\epsilon}_\ell - ilde{\epsilon}_{\ell+1} rac{\partial f_{\ell+1}(v_\ell; heta_{\ell+1})}{\partial v_\ell}$$

and performing gradient descent on F with respect to θ_{ℓ} gives

$$d heta_\ell = - ilde{\epsilon}_\ell rac{\partial f_\ell(
u_{\ell-1}; heta_\ell)}{\partial heta_\ell}.$$

These expressions are identical to Eqs (5) and (6) derived above except that $\tilde{\epsilon}_{\ell}$ takes the place

The precision matrices themselves can be learned by performing gradient descent on F with respect to P_{ℓ} or, as suggested in other work [21], by parameterizing the model in terms of $\Sigma_\ell = P_\ell^{-1}$ and performing gradient descent with respect to Σ_ℓ . Alternatively, one could use techniques from the literature on Gaussian graphical models to learn a sparse or low-rank representation of P_{ℓ} . I circumvent the question of estimating P_{ℓ} by instead just asking how an estimate of P_{ℓ} (however it is obtained) would affect learning. I do assume that P_{ℓ} is symmetric. I also simplify the calculations by restricting the analysis to predictive coding with the fixed prediction assumption, leaving the analysis of fixed point prediction errors and parameter updates under strict predictive coding with precisions matrices for future work. Some analysis has been performed in this direction [21], but not for the supervised learning scenario considered here.

Putting this together, predictive coding under the fixed prediction assumption while accounting for precision matrices in hidden layers is defined by the following equations

$$\begin{split} \tilde{\epsilon}_{\ell} &= [\hat{v}_{\ell} - v_{\ell}] P_{\ell} \\ \tilde{\epsilon}_{L} &= \frac{\partial \mathcal{L}(\hat{v}_{L}, y)}{\partial \hat{v}_{L}} \\ dv_{\ell} &= \tilde{\epsilon}_{\ell} - \tilde{\epsilon}_{\ell+1} \frac{\partial f_{\ell+1}(\hat{v}_{\ell}; \theta_{\ell+1})}{\partial \hat{v}_{\ell}} \\ d\theta_{\ell} &= -\tilde{\epsilon}_{\ell} \frac{\partial f_{\ell}(\hat{v}_{\ell-1}; \theta_{\ell})}{\partial \theta_{\ell}} \end{split}$$

The only difference between these equations and $\underline{\text{Eq (12)}}$ is that they use $\tilde{\epsilon}_{\ell}$ in place of $\epsilon_{\ell} = \hat{\nu}_{\ell} - \nu_{\ell}$. Following the same line of reasoning, therefore, if the updates to ν_{ℓ} are repeated until convergence, then fixed point precision-weighted prediction errors satisfy

$$\tilde{\epsilon}_{\ell}^* = \tilde{\epsilon}_{\ell+1}^* \frac{\partial f_{\ell+1}(\hat{\nu}_{\ell}; \theta_{\ell+1})}{\partial \hat{\nu}_{\ell}}.$$

Notably, this is the same equation derived for ϵ_{ℓ} under the fixed prediction assumption with $\Sigma_{\ell} = I$, so fixed point precision-weighted prediction errors are also the same,

$$\tilde{\epsilon}_{\ell}^{*} = \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{v}_{\ell}}$$

and, therefore, parameter updates are the same as well,

$$d\theta_{\ell} = -\frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \theta_{\ell}}.$$

In conclusion, accounting for precision matrices in hidden layers does not affect learning under the fixed prediction assumption. Fixed point parameter updates are still the same as those computed by backpropagation. This conclusion is independent of how the precision matrices are estimated, but it does rely on the assumption that fixed points for v_{ℓ} exist and are unique.

Above, we only considered precision matrices in the hidden layers because accounting for precision matrices in the output layer is problematic for general loss functions. The use of a precision matrix in the output implies the use of a Gaussian model for the output layer and labels, which is inconsistent with some types of labels and loss functions. If we focus on the case of a squared-Euclidean loss function,

$$\mathcal{L}(\hat{y}, y) = \frac{1}{2} ||\hat{y} - y||^2,$$

then the use of precision matrices in the output layer is more parsimonious and we can define

$$\tilde{\epsilon}_L = [\hat{v}_L - y]P_L = \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{y}}P_L$$

in place of the definition above (recalling that $\hat{y} = \hat{v}_L$). Following the same calculations as above, gives fixed points of the form

$$\tilde{\epsilon}_{\ell}^{*} = \frac{\partial \mathcal{L}(\hat{y}, y)}{\partial \hat{y}} P_{L} \frac{\partial \hat{y}}{\partial \hat{v}_{\ell}}$$

and, therefore, weight updates take the form

$$d heta_\ell = -rac{\partial \mathcal{L}(\hat{y},y)}{\partial \hat{y}} P_L rac{\partial \hat{y}}{\partial heta_\ell}.$$

at the fixed point. Hence, accounting for precision matrices at the output layer can affect learning by re-weighting the gradient of the loss function according to the precision matrix of the output layer. Note that the precision matrices of the hidden layers still have no effect on learning in this case. Previous work relates the inclusion of the precision matrix in output layers with the use of natural gradients [26, 27].

Prediction errors do not necessarily represent surprising or unexpected features of inputs

Deep neural networks are often interpreted as abstract models of cortical neuronal networks. To this end, the activations of units in deep neural networks are compared to the activity (typically firing rates) of cortical neurons [3, 28, 29]. This approach ignores the representation of errors within the network. More generally, the activations in one particular layer of a feedforward deep neural network contain no information about the activations of deeper layers, the label, or the loss. On the other hand, the activity of cortical neurons can be modulated by downstream activity and information believed to be passed upstream by feedback projections. Predictive coding provides a precise model for the information that deeper layers send to shallower layers, specifically prediction errors.

Under the fixed prediction assumption (Algorithm 3), prediction errors in a particular layer are approximated by the gradients of that layers' activations with respect to the loss function, $\epsilon_\ell = \delta_\ell = \frac{\partial \mathcal{L}}{\partial \hat{\nu}_\ell}$, but under a strict interpretation of predictive coding (Algorithm 2), prediction errors do not necessarily reflect gradients. We next empirically explored how the representations of images differ between the activations from a feedforward pass, $\hat{\nu}_\ell$, the prediction errors under the fixed prediction assumption, $\epsilon_\ell = \delta_\ell$, as well as the beliefs, ν_ℓ , and prediction errors, ϵ_ℓ , under a strict interpretation of predictive coding (Algorithm 2). To do so, we computed each quantity in VGG-19 [30], which is a large, feedforward convolutional neural network (19 layers and 143,667,240 trainable parameters) pre-trained on ImageNet [31].

The use of convolutional layers allowed us to visualize the activations and prediction errors in each layer. Specifically, we took the Euclidean norm of each quantity across all channels and plotted them as two-dimensional images for layers $\ell=1$ and $\ell=10$ and for two different input images (Fig 8). For each image and each layer (each row in Fig 8), we computed the Euclidean norm of four quantities. First, we computed the activations from a forward pass through the network (\hat{v}_{ℓ} , second column). Under predictive coding with the fixed prediction assumption (Algorithm 3), we can interpret the activations, \hat{v}_{ℓ} , as "beliefs" and the gradients, δ_{ℓ} , as "prediction errors." Strictly speaking, there is a distinction between the beliefs, \hat{v}_{ℓ} , from a feedforward pass and the beliefs, $v_{\ell}=\hat{v}_{\ell}+\epsilon_{\ell}$, when labels are provided. Either could be interpreted as a "belief." However, we found that the difference between them was negligible for the examples considered here.

Next, we computed the gradients of the loss with respect to the activations (δ_ℓ , third column in Fig 8). The theory and simulations above and from previous work confirms that these gradients approximate the prediction errors from predictive coding with the fixed prediction assumption (Algorithm 3). Indeed, for the examples considered here, the differences between the two quantities were negligible. Next, we computed the beliefs (ν_ℓ , fourth column in Fig 8) computed by strict predictive coding (Algorithm 2). Finally, we computed the prediction errors (ϵ_ℓ , last column in Fig 8) computed by strict predictive coding (Algorithm 2).

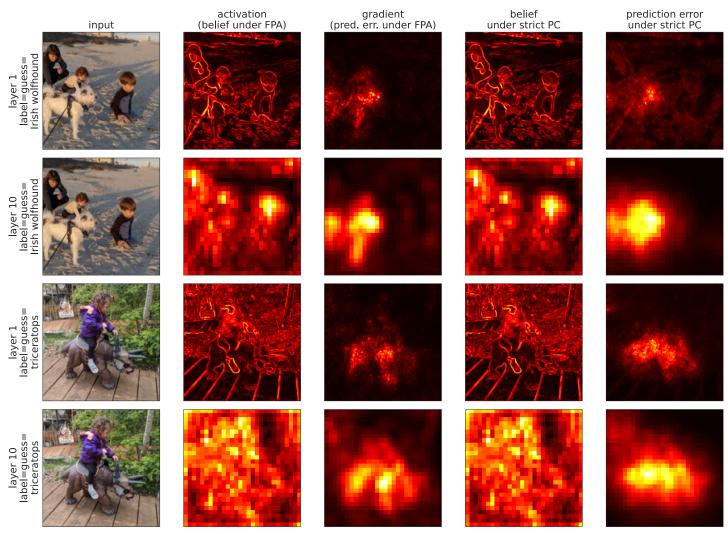


Fig 8. Magnitude of activations, beliefs, and prediction errors in a convolutional neural network pre-trained on ImageNet. The Euclidean norm of feedforward activations (\hat{v} , interpreted as beliefs under the fixed prediction assumption), gradients of the loss with respect to activations ($\delta_{\ell} = \partial \mathcal{L}/\partial \hat{v}$, interpreted as prediction errors under the fixed prediction assumption), beliefs (v) under strict predictive coding, and prediction errors (ϵ_{ℓ})) under strict predictive coding computed from the VGG-19 network [30] pre-trained on ImageNet [31] with two different photographs as inputs at two different layers. The vertical labels on the left ("triceratops" and "Irish wolfhound") correspond to the guessed label which was also used as the "true" label (y) used to compute the gradients.

https://doi.org/10.1371/journal.pone.0266102.g008

Note that we used a VGG-19 model that was pre-trained using backpropagation. Hence, the weights were not necessarily the same as the weights that would be obtained if the model were trained using predictive coding, particularly strict predictive coding (Algorithm 2) which does not necessarily converge to the true gradients. Training a large ImageNet model like VGG-19 with predictive coding is extremely computationally expensive. Regardless, future work should address the question of whether using pre-trained weights (versus weights trained by predictive coding) affects the conclusions reached here.

Overall, the activations, \hat{v}_{ℓ} , from a feedforward pass were qualitatively very similar to the beliefs, v_{ℓ} , computed under a strict interpretation of predictive coding (Algorithm 2). To a slightly lesser degree, the gradients, δ_{ℓ} , from a feedforward pass were qualitatively similar to the prediction errors computed under a strict interpretation of predictive coding (Algorithm 2). Since \hat{v}_{ℓ} and δ_{ℓ} approximate beliefs and prediction errors under the fixed prediction

assumption, these observations confirmed that the fixed prediction assumption does not make large qualitative changes to the representation of beliefs and errors in these examples. Therefore, in the discussion below, we used "beliefs" and "prediction errors" to refer to the quantities from both models.

Interestingly, prediction errors were non-zero even when the image and the network's "guess" was consistent with the label (no "mismatch"). Indeed, the prediction errors were largest in magnitude at pixels corresponding to the object predicted by the label, *i.e.*, at the most predictable regions. While this observation is an obvious consequence of the fact that prediction errors are approximated by the gradients, $\delta_\ell = \frac{\partial \mathcal{L}}{\partial \hat{v}_\ell}$, it is contradictory to the heuristic or intuitive interpretation of prediction errors as measurements of "surprise" in the colloquial sense of the word [16].

As an illustrative example from Fig 8, it is not surprising that an image labeled by "triceratops" contains a triceratops, but this does not imply a lack of prediction errors because the space of images containing a triceratops is large and any one image of a triceratops is not wholly representative of the label. Moreover, the pixels to which the loss is most sensitive are those pixels containing the triceratops. Therefore those pixels give rise to larger values of $\epsilon_\ell \approx \delta_\ell = \partial \mathcal{L}/\partial \hat{v}_\ell$. Hence, in high-dimensional sensory spaces, predictive coding models do not necessarily predict that prediction error units encode "surprise" in the colloquial sense of the word.

In both examples in Fig 8, we used an input, y, that matched the network's "guessed" label, *i.e.*, the label to which the network assigned the highest probability (argmax(\hat{y})). Prediction errors are often discussed in the context of mismatched stimuli in which top-down input is inconsistent with bottom-up predictions [32–37]. Mismatches can be modeled by taking a label that is different from the network's guess. In Fig 9, we visualized the prediction errors in response to matched and mismatched labels. The network assigned a probability of p = 0.9991

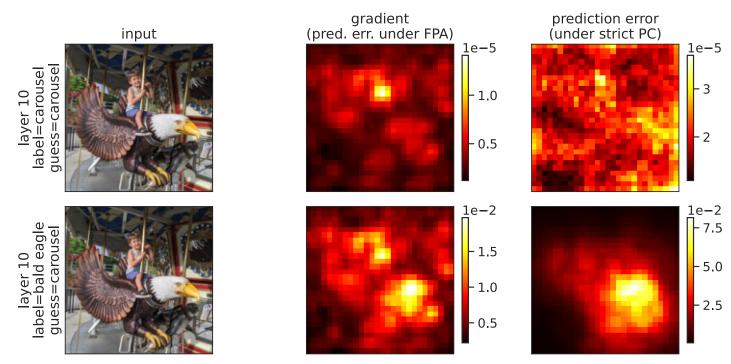


Fig 9. Magnitude of activations, beliefs, and prediction errors in response to matched and mismatched inputs and labels. Same as Fig 8, but for the bottom row the label did not match the network's guess.

https://doi.org/10.1371/journal.pone.0266102.g009

to the label "carousel" and a probability of $p = 3.63 \times 10^{-8}$ to the label "bald eagle". The low probability assigned to "bald eagle" is, at least in part, a consequence of the network being trained with a softmax loss function, which implicitly assumes one label per input. When we applied the mismatched label "bald eagle," prediction errors were larger in pixels that are salient for that label (*e.g.*, the bird's white head, which is a defining feature of a bald eagle). Moreover, the prediction errors as a whole are much larger in magnitude in response to the mismatched label (see the scales of the color bars in Fig 9).

In summary, the relationship between prediction errors and gradients helped demonstrate that prediction errors sometimes, but do not always conform to their common interpretation as unexpected features of a bottom-up input in the context of a top-down input. Also, beliefs and prediction errors were qualitatively similar with and without the fixed prediction assumption for the examples considered here.

Discussion

We reviewed and extended previous work [11–14] on the relationship between predictive coding and backpropagation for learning in neural networks. Our results demonstrated that a strict interpretation of predictive coding does not accurately approximate backpropagation, but is still capable of learning (Figs 1 and 2). Previous work proposed a modification to predictive coding called the "fixed prediction assumption" which causes predictive coding to converge to the same parameter updates produced by backpropagation, under the assumption that the predictive coding iterations converge to fixed points. Hence, the relationship between predictive coding and backpropagation identified in previous work relies critically on the fixed prediction assumption. Formal derivations of predictive coding in terms of variational inference [20] do not produce the fixed prediction assumption. It is possible that an alternative probabilistic model or alternative approaches to the variational formulation could help formalize a model of predictive coding under the fixed prediction assumption.

We proved analytically and verified empirically that taking a step size of $\eta = 1$ in the modified predictive coding algorithm computes the exact gradients computed by backpropagation in a fixed number of steps (modulo floating point numerical errors). This result is consistent with similar, but slightly less general, results in previous work [13, 14].

A closer inspection of the the fixed prediction assumption with $\eta=1$ showed that it is algorithmically equivalent to a direct implementation of backpropagation. As such, any potential neural architecture and machinery that could be to implement predictive coding with the fixed prediction assumption could also implement backpropagation directly. This result calls into question whether predictive coding with the fixed prediction assumption is any more biologically plausible than a direct implementation of backpropagation.

Visualizing the beliefs and prediction errors produced by predictive coding models applied to a large convolutional neural network pre-trained on ImageNet showed that beliefs and prediction errors were activated by distinct parts of input images, and the parts of the images that produced larger prediction errors were not always consistent with an intuitive interpretation of prediction errors as representing surprising or unexpected features of inputs. These observations are consistent with the fact that prediction errors approximate gradients of the loss function in backpropagation [11–14]. Gradients are large for input features that have a larger impact on the loss. While surprising features can have a large impact on the loss, unsurprising features can as well. We only verified this finding empirically on a few examples. The reader can try additional examples by inserting the URL of any image into the file PredErrsFromurrimage.ipynb contained in the directories linked in Materials and Methods, and can also be accessed directly at https://bit.ly/3JwGUM9. Future work should attempt to quantify

the relationship between prediction errors and surprising features more systematically across many inputs. In addition, prediction errors could be computed for learning tasks associated with common experimental paradigms so they can be used to make experimentally testable predictions.

When interpreting artificial deep neural networks as models of biological neuronal networks, it is common to compare activations in the artificial network to biological neurons' firing rates [28, 29]. However, under predictive coding models and other models in which errors are propagated upstream by feedback connections, many biological interpretations posit the existence of "error neurons" that encode the errors sent upstream. In most such models (including predictive coding), error neurons reflect or approximate the gradient of the loss function with respect to artificial neurons' activations, δ_ℓ . Any model that hypothesizes the neural representation of backpropagated errors would predict that some recorded neural activity should reflect these errors. Therefore, if we want to draw analogues between artificial and biological neural networks, the activity of biological neurons should be compared to both the activations *and* the gradients of artificial neurons.

Following previous work [11, 12], we took the covariance matrices underlying the probabilistic model to be identity matrices, $\Sigma_\ell = I$, when deriving the predictive coding model. We also showed that relaxing this assumption by allowing for arbitrary precision matrices in hidden layers does not affect learning under the fixed prediction assumption. Future work should consider the utility of accounting for covariance (or precision) matrices in models without the fixed prediction assumption (*i.e.*, under the "strict" model) and accounting for precisions or covariances in the output layer. Moreover, precision matrices could still have benefits in other settings such as recurrent network models, unsupervised learning, or active inference.

Predictive coding and deep neural networks (trained by backpropagation) are often viewed as competing models of brain function. Better understanding their relationship can help in the interpretation and implementation of each algorithm as well as their mutual relationships to biological neuronal networks.

Materials and methods

All numerical examples were performed on GPUs using Google Collaboratory with custom-written PyTorch code. The networks trained on MNIST used two convolutional and three fully connected layers with rectified linear activation functions using 2 epochs, a learning rate of 0.002, and a batch size of 300. The networks trained on CIFAR-10 used three convolutional and three fully connected layers with rectified linear activation functions using 5 epochs, a learning rate of 0.01, and a batch size of 256. All networks were trained using the Adam optimizer with gradients replaced by the output of the respective algorithm. All of the code to produce the figures in the manuscript can be found at https://doi.org/10.6084/m9.figshare.19387409.v2 A Google Drive folder with Colab notebooks that produce all figures in this text can be found at https://drive.google.com/drive/folders/1m_y0G_sTF-pV9pd2_sysWt1nvRvHYzX0 An additional copy of the same code is also stored at https://github.com/RobertRosenbaum/PredictiveCodingVsBackProp Full details of the neural network architectures and metaparameters can be found in this code.

Torch2PC software for predictive coding with PyTorch models

The figures above were all produced using PyTorch [38] models combined with custom written functions for predictive coding. Functions for predictive coding with PyTorch models are collected in the Github Repository Torch2PC. Currently, the only available functions are

intended for models built using the Sequential class, but more general functions will be added to Torch2PC in the future. The functions can be imported using the following commands

```
!git clonehttps://github.com/RobertRosenbaum/Torch2PC.git
from Torch2PC import TorchSeq2PC as T2PC
```

The primary function in TorchSeq2PC is PCInfer, which performs one predictive coding step (computes one value of $d\theta$) on a batch of inputs and labels. The function takes an input ErrType, which is a string that determines whether to use a strict interpretation of predictive coding (Algorithm 2; ErrType="Strict"), predictive coding with the fixed prediction assumption (Algorithm 3; "FixedPred"), or to compute the gradients exactly using backpropagation (Algorithm 1; "Exact"). Algorithm 2 can be called as follows,

```
vhat,Loss,dLdy,v,epsilon=
T2PC.PCInfer(model,LossFun,X,Y,"Strict",eta,n,vinit)
```

where model is a Sequential PyTorch model, LossFun is a loss function, X is a minibatch of inputs, Y is a minibatch of labels, eta is the step size, n is the number of iterations to use, and vinit is the initial value for the beliefs. If vinit is not passed, it is set to the result from a forward pass, vinit = vhat. The function returns a list of activations from a forward pass at each layer as vhat, the loss as Loss, the gradient of the output with respect to the loss as dLdy, a list of beliefs, v_{ℓ} , at each layer as v, and a list of prediction errors, ϵ_{ℓ} , at each layer as epsilon. The values of the parameter updates, $d\theta_{\ell}$, are stored in the grad attributes of each parameter, model.param.grad. Hence, after a call to PCInfer, gradient descent could be implemented by calling

```
with torch.no_grad():
    for p in modelPC.parameters():
        p-=eta*p.grad
Alternatively, an arbitrary optimizer could be used by calling
optimizer.step()
where optimizer is an optimizer created using the PyTorch optim class, e.g., by calling
```

The input model should be a PyTorch Sequential model. Each layer is treated as a single predictive coding layer. Multiple functions can be included within the same layer by wrapping them in a separate call to Sequential. For example the following code:

optimizer = optim.Adam (model.parameters()) before the call to T2PC.

```
model = nn.Sequential(
    nn.Conv2d(1,10,3),
    nn.ReLU(),
    nn.MaxPool2d(2),
    nn.Conv2d(10,10,3),
    nn.ReLU())
```

PCInfer.

will treat each item as its own layer (5 layers in all). To treat each "convolutional block" as a separate layer, instead do

```
model = nn.Sequential(
    nn.Sequential(
    nn.Conv2d(1,10,3),
    nn.ReLU(),
    nn.MaxPool2d(2)),
    nn.Sequential(
    nn.Conv2d(10,10,3),
        nn.ReLU()))
which has just 2 layers.
```

```
Algorithm 3 can be called as follows,
  vhat, Loss, dLdy, v, epsilon=
    T2PC.PCInfer (model, LossFun, X, Y, "FixedPred", eta, n)
  The input vinit is not used for Algorithm 3, so it does not need to be passed in. The exact
values computed by backpropagation can be obtained by calling
  vhat, Loss, dLdy, v, epsilon=
    T2PC.PCInfer(model,LossFun,X,Y,"Exact")
  The inputs vinit, eta, and n are not used for computing exact gradients, so they do not
need to be passed in. Theorem 1 says that
  T2PC.PCInfer(model,LossFun,X,Y,"FixedPred",eta = 1,n = len
(model))
  computes the same values as
  T2PC.PCInfer (model, LossFun, X, Y, "Exact")
  up to numerical floating point errors. The inputs eta, n, and vinit are optional. If they
are omitted by calling
  T2PC.PCInfer (model, LossFun, X, Y, ErrType)
  then they default to eta=.1, n = 20, vinit = None which produces vinit = vhat
when
  ErrType="Strict". More complete documentation and a complete example is pro-
```

vided as SimpleExample.ipynb in the GitHub repository and in the code accompanying this

paper. More examples are provided by the code accompanying each figure above.

Supporting information

S1 Fig. Comparing backpropagation and predictive coding in a convolutional neural network trained on MNIST across multiple trials. Same as Fig 1 except the model was trained 30 times with different random seeds. Dark curves show the mean values and shaded regions show ± one standard deviation across trials. (EPS)

S2 Fig. Comparing backpropagation and predictive coding modified by the fixed prediction assumption in a convolutional neural network trained on MNIST across multiple trials. Same as Fig 3 except the model was trained 30 times with different random seeds. Dark curves show the mean values and shaded regions show ± one standard deviation across trials. (EPS)

Author Contributions

Conceptualization: Robert Rosenbaum. Formal analysis: Robert Rosenbaum. Funding acquisition: Robert Rosenbaum. **Investigation:** Robert Rosenbaum. Methodology: Robert Rosenbaum. Software: Robert Rosenbaum. Visualization: Robert Rosenbaum.

Writing - original draft: Robert Rosenbaum.

Writing - review & editing: Robert Rosenbaum.

References

- Izhikevich EM. Solving the distal reward problem through linkage of STDP and dopamine signaling. Cerebral cortex. 2007; 17(10):2443-2452. https://doi.org/10.1093/cercor/bhl152 PMID: 17220510
- Clark DG, Abbott L, Chung S. Credit Assignment Through Broadcasting a Global Error Vector. arXiv preprint arXiv:210604089. 2021;.
- Lillicrap TP, Santoro A, Marris L, Akerman CJ, Hinton G. Backpropagation and the brain. Nature Reviews Neuroscience. 2020; 21(6):335-346. https://doi.org/10.1038/s41583-020-0277-3 PMID:
- Whittington JC, Bogacz R. Theories of error back-propagation in the brain. Trends in Cognitive Sciences. 2019; 23(3):235-250. https://doi.org/10.1016/j.tics.2018.12.005 PMID: 30704969
- Urbanczik R, Senn W. Learning by the dendritic prediction of somatic spiking. Neuron. 2014; 81 (3):521-528. https://doi.org/10.1016/j.neuron.2013.11.030 PMID: 24507189
- Lillicrap TP. Cownden D, Tweed DB, Akerman CJ. Random synaptic feedback weights support error backpropagation for deep learning. Nature Communications. 2016; 7(1):1-10. https://doi.org/10.1038/ ncomms13276 PMID: 27824044
- Scellier B, Bengio Y. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. Frontiers in computational neuroscience. 2017; 11:24. https://doi.org/10.3389/fncom. 2017.00024 PMID: 28522969
- Aljadeff J, D'amour J, Field RE, Froemke RC, Clopath C. Cortical credit assignment by Hebbian, neuromodulatory and inhibitory plasticity. arXiv preprint arXiv:191100307. 2019;.
- Kunin D, Nayebi A, Sagastuy-Brena J, Ganguli S, Bloom J, Yamins D. Two routes to scalable credit assignment without weight symmetry. In: International Conference on Machine Learning. PMLR; 2020. p. 5511-5521.
- 10. Payeur A, Guerquiev J, Zenke F, Richards BA, Naud R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. Nature Neuroscience. 2021; p. 1-10.
- Whittington JC, Bogacz R. An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. Neural Computation. 2017; 29(5):1229-1262. https:// doi.org/10.1162/NECO a 00949 PMID: 28333583
- Millidge B, Tschantz A, Buckley CL. Predictive coding approximates backprop along arbitrary computation graphs. arXiv preprint arXiv:200604182. 2020;.
- Song Y, Lukasiewicz T, Xu Z, Bogacz R. Can the brain do backpropagation?—exact implementation of backpropagation in predictive coding networks. Advances in Neural Information Processing Systems. 2020: 33:22566. PMID: 33840988
- Salvatori T, Song Y, Lukasiewicz T, Bogacz R, Xu Z. Predictive Coding Can Do Exact Backpropagation on Convolutional and Recurrent Neural Networks. arXiv preprint arXiv:210303725. 2021;.
- Rao RP, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extraclassical receptive-field effects. Nature Neuroscience. 1999; 2(1):79-87. https://doi.org/10.1038/4580 PMID: 10195184
- Friston K. The free-energy principle: a unified brain theory? Nature Reviews Neuroscience. 2010; 11 (2):127-138. PMID: 20068583
- Huang Y, Rao RP. Predictive Coding. Wiley Interdisciplinary Reviews: Cognitive Science. 2011; 2 (5):580-593. PMID: 26302308
- Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ. Canonical microcircuits for predictive coding. Neuron. 2012; 76(4):695-711. https://doi.org/10.1016/j.neuron.2012.10.038 PMID: 23177956
- Clark A. Surfing uncertainty: Prediction, action, and the embodied mind. Oxford University Press; 2015. 19.
- Buckley CL, Kim CS, McGregor S, Seth AK. The free energy principle for action and perception: A mathematical review. Journal of Mathematical Psychology. 2017; 81:55-79. https://doi.org/10.1016/j.jmp. 2017.09.004
- Bogacz R. A tutorial on the free-energy framework for modelling perception and learning. Journal of 21. Mathematical Psychology. 2017; 76:198-211. https://doi.org/10.1016/j.jmp.2015.11.003 PMID: 28298703

- Spratling MW. A review of predictive coding algorithms. Brain and cognition. 2017; 112:92–97. https://doi.org/10.1016/j.bandc.2015.11.003 PMID: 26809759
- Keller GB, Mrsic-Flogel TD. Predictive processing: a canonical cortical computation. Neuron. 2018; 100 (2):424–435. https://doi.org/10.1016/j.neuron.2018.10.003 PMID: 30359606
- 24. Goodfellow I, Bengio Y, Courville A, Bengio Y. Deep Learning. MIT press Cambridge; 2016.
- 25. Krizhevsky A, Hinton G, et al. Learning multiple layers of features from tiny images. Citeseer. 2009;.
- Millidge B, Seth A, Buckley CL. Predictive Coding: a Theoretical and Experimental Review. arXiv preprint arXiv:210712979. 2021;.
- 27. Amari Si. Information geometry of the EM and em algorithms for neural networks. Neural networks. 1995; 8(9):1379–1408. https://doi.org/10.1016/0893-6080(95)00003-8
- Schrimpf M, Kubilius J, Hong H, Majaj NJ, Rajalingham R, Issa EB, et al. Brain-Score: Which Artificial Neural Network for Object Recognition is most Brain-Like? bioRxiv preprint. 2018;.
- Schrimpf M, Kubilius J, Lee MJ, Murty NAR, Ajemian R, DiCarlo JJ. Integrative Benchmarking to Advance Neurally Mechanistic Models of Human Intelligence. Neuron. 2020;. https://doi.org/10.1016/j.neuron.2020.07.040 PMID: 32918861
- Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:14091556. 2014;.
- Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, et al. Imagenet large scale visual recognition challenge. International Journal of Computer Vision. 2015; 115(3):211–252. https://doi.org/10.1007/s11263-015-0816-y
- **32.** Hertäg L, Sprekeler H. Learning prediction error neurons in a canonical interneuron circuit. Elife. 2020; 9:e57541. https://doi.org/10.7554/eLife.57541 PMID: 32820723
- **33.** Gillon CJ, Pina JE, Lecoq JA, Ahmed R, Billeh Y, Caldejon S, et al. Learning from unexpected events in the neocortical microcircuit. bioRxiv. 2021;.
- Keller GB, Bonhoeffer T, Hübener M. Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. Neuron. 2012; 74(5):809–815. https://doi.org/10.1016/j.neuron.2012.03.040 PMID: 22681686
- Zmarz P, Keller GB. Mismatch receptive fields in mouse visual cortex. Neuron. 2016; 92(4):766–772. https://doi.org/10.1016/j.neuron.2016.09.057 PMID: 27974161
- Attinger A, Wang B, Keller GB. Visuomotor coupling shapes the functional development of mouse visual cortex. Cell. 2017; 169(7):1291–1302. https://doi.org/10.1016/j.cell.2017.05.023 PMID: 28602353
- 37. Homann J, Koay SA, Glidden AM, Tank DW, Berry MJ. Predictive coding of novel versus familiar stimuli in the primary visual cortex. BioRxiv. 2017; p. 197608.
- Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. Advances in Neural Information Processing Systems. 2019; 32:8026–8037.