Credit Assignment Through Broadcasting a Global Error Vector

Anonymous Author(s)
Affiliation
Address
email

Abstract

Backpropagation (BP) is remarkably effective in training deep neural networks (DNNs). At the same time, biological neural circuits appear to perform credit assignment, but cannot implement BP, implying the existence of other powerful learning algorithms. Here, we explore the extent to which a single, globally broadcast learning signal enables training of DNNs. We present both a learning rule, called global error-vector broadcasting (GEVB), and a class of DNNs, called vectorized nonnegative networks (VNNs), in which this rule operates. VNNs have vector-valued units and nonnegative weights past the first layer. The GEVB learning rule generalizes three-factor Hebbian learning, updating each weight by an amount proportional to the inner product of the presynaptic activation and a globally broadcast error vector when the postsynaptic unit is active. We show that these weight updates are matched in sign to the gradient, enabling accurate credit assignment. Moreover, at initialization, these updates are exactly proportional to the gradient in the limit of infinite network width. GEVB matches the performance of BP in VNNs, and in some cases outperforms direct feedback alignment (DFA) applied in conventional networks. Unlike DFA, GEVB successfully trains convolutional layers. Altogether, our theoretical and empirical results point to a surprisingly powerful role for a global learning signal in training DNNs.

1 Introduction

Deep neural networks (DNNs) trained using backpropagation (BP) have achieved breakthroughs on a myriad of tasks [1, 2]. The power of BP lies in its ability to discover intermediate feature representations by following the gradient of a loss function. In a manner analogous to DNN training, synapses in multilayered cortical circuits undergo plasticity that modulates neural activity several synapses downstream to improve performance on behavioral tasks [3, 4]. However, neural circuits cannot implement BP, implying that evolution has found another algorithm, or collection thereof [5]. This observation has motivated biologically plausible alternatives to BP [6, 7].

Credit assignment algorithms based on broadcasting a global learning signal, such as node perturbation [8], are attractive due to their biological plausibility, as the global signal could correspond to a neuromodulator that influences local synaptic plasticity [9, 10]. However, perturbation methods are of little practical relevance as the variance in their gradient estimates is prohibitively large [6, 11]. Instead, recent research has focused on methods that compute or estimate the gradient by transmitting detailed, unit-specific information to hidden layers through top-down feedback. The foremost example of such a method is BP. BP's use of the feedforward weights in the feedback pathway, a property called weight symmetry, is not biologically plausible – this is the weight transport problem [12]. Another method of this class, feedback alignment (FA), operates identically to BP but uses fixed, random matrices in the feedback pathway, thereby circumventing the weight transport

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

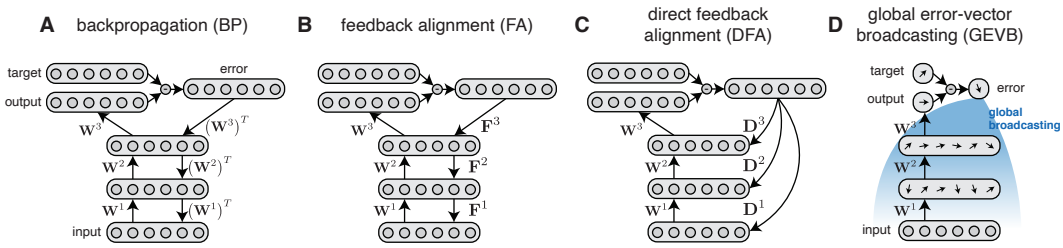

Figure 1: Credit assignment methods. **(A)** Backpropagation (BP) transmits the error vector backwards layer-by-layer using the transposes of the feedforward weights. **(B)** Feedback alignemnt (FA) transmits the error vector backwards layer-by-layer using fixed, random feedback matrices. **(C)** Direct feedback alignment (DFA) delivers the error vector directly to each hidden layer using fixed, random feedback matrices. **(D)** Global error-vector broadcasting (GEVB) conveys the full error vector to all hidden units without unit-specific feedback. Thus, there are no feedback parameters. GEVB operates in networks in which each hidden unit is vector-valued.

problem [13]. FA inspired a further method called direct feedback alignment (DFA), which delivers random projections of the output error vector directly to hidden units [14, 15]. DFA approaches the performance of BP in certain modern architectures [16], however its performance still lags behind that of BP in many cases [14, 16–19]. A particularly striking shortcoming of DFA is its inability to train convolutional layers [16, 17, 19, 20]. To be implemented in neural circuits, both FA and DFA require a biologically unrealistic "error network" to compute top-down learning signals, though models based on segregated dendrites aim to lift this requirement [21, 22].

Here, we show that DNNs can be trained to performance on par with BP by broadcasting a single, global learning signal to all hidden units and applying local, Hebbian-like updates to the weights. Our proposed learning rule, called global error-vector broadcasting (GEVB), is not perturbation-based, but instead distributes information about the output error throughout the network. Unlike DFA, which delivers a unique random projection of the output error vector to each hidden unit, GEVB broadcasts the full error vector to all hidden units. This learning rule operates in a new class of DNNs called vectorized nonnegative networks (VNNs), which have vector-valued units and nonnegative weights past the first layer. While our use of vector-valued units is neuroscientifically speculative (see Section 7 for possible neural implementations), nonnegative connections are motivated by the fact that cortical projection neurons are excitatory [3]. The GEVB learning rule updates each weight by an amount proportional to the inner product of the presynaptic activation and the global error vector when the postsynaptic unit is active, a form of three-factor Hebbian learning with a vector-valued global third factor [23, 24]. Our experimental results show that this form of global error-based learning is surprisingly powerful, performing on par with BP in VNNs and overcoming DFA's inability to train convolutional layers.

Our theoretical results are organized as follows, followed by experimental results and discussion. **Section 2** describes credit assignment in conventional DNNs via BP, FA, and DFA. We then identify an obstacle to performing credit assignment in the absence of weight symmetry, namely, computing the sign of the gradient. **Section 3** introduces VNNs and the GEVB learning rule. We prove that GEVB produces weight updates matched in sign to the gradient under appropriate choice of the VNN nonlinearity, overcoming the obstacle described in Section 2. **Section 4** shows that random nonnegative initializations place VNNs in a regime of exact gradient alignment in the limit of infinite network width. Finally, **Section 5** presents a method for initializing VNNs based on ON/OFF cells that ensures well-behaved forward propagation.

2 Credit assignment in conventional networks

Consider a fully connected network of L layers with pre-activations h_i^ℓ and activations a_i^ℓ, where $i = 1, \ldots, n_\ell$ and $\ell = 0, \ldots, L$. Let w_{ij}^ℓ denote the weights, b_i^ℓ the biases, and ϕ the pointwise

nonlinearity. The forward dynamics are

$$h_i^\ell = \sum_{j=1}^{n_{\ell-1}} w_{ij}^\ell a_j^{\ell-1} + b_i^\ell, \quad a_i^\ell = \phi\left(h_i^\ell\right), \tag{1}$$

where ϕ is implicitly the identity function at the output layer L. The network inputs and outputs are a_i^0 and a_i^L, respectively. We place a loss $\mathcal{L}$ on the output, and define $\partial\mathcal{L}/\partial a_i^L = e_i$, where e_i is the output error vector. On a single training example, a negative-gradient weight update can be written

$$\Delta w_{ij}^\ell = -\delta_i^\ell \phi'\left(h_i^\ell\right) a_j^{\ell-1} \quad \text{where} \quad \delta_i^\ell = \frac{\partial \mathcal{L}}{\partial a_i^\ell} = \sum_{k=1}^{n_L} \frac{\partial a_k^L}{\partial a_i^\ell} e_k. \tag{2}$$

BP, FA, and DFA are different methods of computing or approximating δ_i^ℓ. BP computes δ_i^ℓ exactly using the recurrence relation

$$\delta_i^\ell = \sum_{k=1}^{n_{\ell+1}} w_{ki}^{\ell+1} \phi'(h_k^{\ell+1}) \delta_k^{\ell+1} \tag{3}$$

with the initial condition $\delta_i^L = e_i$ (Fig. 1A). FA approximates δ_i^ℓ by performing a backward pass identical to Eq. 3, but using fixed, random feedback weights f_{ik}^ℓ in place of the transposed feedforward weights w_{ki}^ℓ (Fig. 1B). This backward pass is given by

$$\delta_i^\ell = \sum_{k=1}^{n_{\ell+1}} f_{ik}^{\ell+1} \phi'(h_k^{\ell+1}) \delta_k^{\ell+1} \tag{4}$$

with the initial condition $\delta_i^L = e_i$. Learning occurs in FA due to automatic partial alignment of w_{ki}^ℓ with f_{ik}^ℓ [13]. Finally, DFA approximates δ_i^ℓ by setting the derivative $\partial a_k^L/\partial a_i^\ell$ in the formula for δ_i^ℓ in Eq. 2 to a fixed, random matrix d_{ik}^ℓ (Fig. 1C). This yields

$$\delta_i^\ell = \sum_{k=1}^{n_L} d_{ik}^\ell e_k. \tag{5}$$

In a similar manner to FA, learning occurs in DFA due to automatic partial alignment of $\partial a_k^L/\partial a_i^\ell$ with d_{ik}^ℓ [19, 20]. The learning rule we propose further simplifies Eqns. 3, 4, and 5 by setting the quantity analogous to δ_i^ℓ to a constant for all units (Section 3.2).

A fundamental obstacle to accurate credit assignment in the absence of weight symmetry is that the sign of the gradient at each weight depends on detailed information about downstream weights to which the feedback pathway does not have access, even in the presence of partial alignment of feedforward and feedback weights. For example, suppose a DNN confidently predicts class c' instead of the target class c, yielding $e_i \approx \delta_{ic'} - \delta_{ic}$. Then, the gradient at w_{ij}^ℓ depends on $\delta_i^\ell \approx \partial a_{c'}^L/\partial a_i^\ell - \partial a_c^L/\partial a_i^\ell$, which can be positive or negative depending on a difference in the strengths of complex polysynaptic interactions between unit i in layer ℓ and units c and c' in layer L.

Intriguingly, this obstacle dissolves in networks with scalar output, positive weights past the first layer, and $\phi' > 0$. In this case, we have $\delta_i^\ell = (\partial a_1^L/\partial a_i^\ell)e_1$ with $\partial a_1^L/\partial a_i^\ell > 0$. Thus, as per Eq. 2, the gradient sign at w_{ij}^ℓ is equal to the sign of $\phi'(h_i^\ell)a_j^{\ell-1}e_1$, a three-factor Hebbian rule. This simplification to credit assignment arising from sign-constrained weights in scalar-output networks was observed by both Balduzzi et al. [25], who proposed the Kickback algorithm, and Lechner [18], who used sign-constrained weights in DFA. However, the restriction to scalar-output (e.g., binary classification) networks is highly limiting. The GEVB learning rule, described in the next section, generalizes this observation to vector-output (e.g., multi-way classification) networks.

3 Credit assignment in vectorized nonnegative networks

3.1 Vectorized nonnegative networks

Here, we introduce VNNs, the class of DNNs in which our proposed learning rule operates. Given a task with output dimension K, such as K-way classification, each VNN unit is a vector of dimension

K. Connections obey "vectorized weight sharing" such that each vector unit computes a weighted linear combination of presynaptic vector units. Our notational convention is to use Latin letters to index different vector units and Greek letters to index the components of a single vector unit. Thus, Greek indices always run over $1, \ldots, K$. Boldface variables represent all K components of a vector unit. For $\ell = 1 \ldots, L$, let $h^\ell_{i\mu}$ denote the pre-activations and $a^\ell_{i\mu}$ the activations. The inputs, a^0_i, are not vectorized and hence lack a μ subscript. The vector units in the first hidden layer compute a representation of the input using the weights $w^1_{i\mu j}$. For $\ell > 1$, the weights are denoted by w^ℓ_{ij}. Crucially, these $\ell > 1$ weights are nonnegative, consistent with excitatory cortical projection neurons. For $\ell = 1, \ldots, L$, the biases are denoted by $b^\ell_{i\mu}$. Finally, the nonlinearity is denoted by $\Phi : \mathbb{R}^K \to \mathbb{R}^K$. The forward dynamics are

$$h^\ell_{i\mu} = \begin{cases} \sum_{j=1}^{n_0} w^1_{i\mu j} a^0_j + b^1_{i\mu} & \ell = 1 \\ \sum_{j=1}^{n_{\ell-1}} w^\ell_{ij} a^{\ell-1}_{j\mu} + b^\ell_{i\mu} & \ell > 1 \end{cases} \qquad a^\ell_{i\mu} = \Phi_\mu\left(\boldsymbol{h}^\ell_i\right) \tag{6}$$

where Φ is implicitly the identity function at the output layer L. We assume that there is a single vector output unit $a^L_{1\mu}$, denoted by a^L_μ for brevity, on which we place a loss $\mathcal{L}$. We define $\partial\mathcal{L}/\partial a^L_\mu = e_\mu$, where e_μ is the output error vector. We will show that our proposed learning rule matches the sign of the gradient when the nonlinearity has the form

$$\Phi_\mu(\boldsymbol{h}) = G(\boldsymbol{h}) h_\mu, \tag{7}$$

with $G(\boldsymbol{h}) \geq 0$ and piecewise constant. While there are many options for this function, we choose

$$G(\boldsymbol{h}) = \Theta\left(\boldsymbol{t} \cdot \boldsymbol{h}\right), \tag{8}$$

where Θ is the Heaviside step function and $\boldsymbol{t}$ is a gating vector that differs across units. In the $K = 1$ case, Φ reduces to the rectified linear unit (ReLU) function if $t > 0$ since $\Phi(h) = \Theta(th)h = \Theta(h)h = \text{ReLU}(h)$. We sample the gating vectors uniformly over $\{-1, 1\}^K$.

While we described the inputs as n_0 scalar units with all-to-all connections with the first hidden layer, the inputs can be described equivalently as Kn_0 vector units with vectorized weight-shared connections with the first hidden layer, unifying the $\ell = 1$ and $\ell > 1$ cases of the learning rule (see Appendix). We use this convention henceforth.

3.2 Global error-vector broadcasting learning rule

The GEVB learning rule works by globally broadcasting the output error vector e_μ to all hidden units:

$$\textbf{GEVB:} \quad \Delta w^\ell_{ij} = -G\left(\boldsymbol{h}^\ell_i\right) \sum_\mu a^{\ell-1}_{j\mu} e_\mu. \tag{9}$$

Thus, presynaptic units that are aligned or anti-aligned with the output error vector have their weight onto the postsynaptic unit decreased or increased, respectively, when the postsynaptic unit is active. Note that this rule has no feedback parameters. The GEVB learning rule can be interpreted as a three-factor Hebbian rule with a vector-valued global third factor e_μ [23, 24]. Our experimental results in Section 6 show that this simple learning rule is surprisingly powerful.

3.3 Global error-vector broadcasting matches the sign of the gradient

We now show that by choosing Φ as in Eq. 7, with $G(\boldsymbol{h}) \geq 0$ and piecewise constant, GEVB weight updates are matched in sign to the gradient. The gradient on a single training example is

$$\frac{\partial \mathcal{L}}{\partial w^\ell_{ij}} = G(\boldsymbol{h}^\ell_i) \sum_{\mu,\nu} e_\mu \frac{\partial a^L_\mu}{\partial a^\ell_{i\nu}} a^{\ell-1}_{j\nu}. \tag{10}$$

The K-by-K Jacobian $\partial a^L_\mu / \partial a^\ell_{i\nu}$ in Eq. 10 describes how component ν of vector unit i in layer ℓ polysynaptically influences component μ of the vector output unit. This Jacobian can be computed iteratively in a BP-like manner,

$$\frac{\partial a^L_\mu}{\partial a^\ell_{i\nu}} = \sum_{k=1}^{n_{\ell+1}} w^{\ell+1}_{ki} G(\boldsymbol{h}^{\ell+1}_k) \frac{\partial a^L_\mu}{\partial a^{\ell+1}_{k\nu}} \tag{11}$$

with the initial condition $\partial a^L_\mu / \partial a^L_{1\nu} = I_{\mu\nu}$, where $I_{\mu\nu}$ is the identity matrix. These recursions ensure that $\partial a^L_\mu / \partial a^\ell_{i\nu}$ is proportional to $I_{\mu\nu}$ for all ℓ. Thus, we write

$$\frac{\partial a^L_\mu}{\partial a^\ell_{i\nu}} = \hat{\delta}^\ell_i I_{\mu\nu} \quad \text{where} \quad \hat{\delta}^\ell_i = \frac{\partial a^L_\mu}{\partial a^\ell_{i\mu}}. \tag{12}$$

Substitution of Eq. 12 into Eq. 11 gives

$$\hat{\delta}^\ell_i = \sum_{k=1}^{n_{\ell+1}} w^{\ell+1}_{ki} G(\boldsymbol{h}^{\ell+1}_k) \hat{\delta}^{\ell+1}_k \tag{13}$$

with the initial condition $\hat{\delta}^L_i = 1$. The nonnegativity of the $\ell > 1$ weights and of $G(\boldsymbol{h})$ ensure that $\hat{\delta}^\ell_i \geq 0$ for all ℓ. Under a mild additional assumption (see Appendix), we have strict positivity, $\hat{\delta}^\ell_i > 0$. Altogether, substituting Eq. 12 into Eq. 10, the gradient is

$$\frac{\partial \mathcal{L}}{\partial w^\ell_{ij}} = \hat{\delta}^\ell_i G(\boldsymbol{h}^\ell_i) \sum_\mu a^{\ell-1}_{j\mu} e_\mu, \tag{14}$$

where $\hat{\delta}^\ell_i$ is backpropagated according to Eq. 13. Since $\hat{\delta}^\ell_i > 0$, omitting the $\hat{\delta}^\ell_i$ term in Eq. 14 yields an expression with the same sign as the gradient. This is precisely the GEVB learning rule of Eq. 9.

Recent theoretical work has shown that gradient sign information is sufficient for attaining or improving stochastic gradient descent convergence rates in non-convex optimization [26]. A caveat is that these rates assume that the gradient sign is computed on mini-batches, whereas GEVB matches the gradient sign on individual examples. In Section 6, we show that training VNNs using batched GEVB updates yields performance on par with BP. Whether non-batched GEVB updates are particularly effective by virtue of matching the gradient sign exactly is a question for future study.

4 Gradient alignment beyond sign agreement

The GEVB learning rule approximates the gradient by setting $\hat{\delta}^\ell_i$ in Eq. 14 to a positive constant for all units i in layer ℓ. This approximation is accurate if the empirical distribution of $\hat{\delta}^\ell_i$ in layer ℓ is tightly concentrated about a positive value, and poor if this distribution is diffuse. The level of concentration can be quantified as the relative standard deviation $r^\ell = \sigma_{\hat{\delta}^\ell} / \mu_{\hat{\delta}^\ell}$, where $\mu_{\hat{\delta}^\ell}$ and $\sigma^2_{\hat{\delta}^\ell}$ denote the mean and variance of the empirical distribution of $\hat{\delta}^\ell_i$ in layer ℓ. Since $\hat{\delta}^\ell_i > 0$, we have $r^\ell > 0$. When $r^\ell \ll 1$, this distribution is close to a delta function at a positive value and $\hat{\delta}^\ell_i$ is accurately approximated as a positive constant. The relative standard deviation is closely related to the alignment angle θ^ℓ between GEVB weight updates and the gradient, a metric commonly used to measure the quality of weight updates, according to $\theta^\ell = \tan^{-1} r^\ell$ (see Appendix for proof) [13, 19].

Intriguingly, in randomly initialized VNNs, r^ℓ scales inversely with n_ℓ in sufficiently deep layers so that, in the limit of infinite network width, $r^\ell, \theta^\ell \to 0$. Concretely, if the elements of $\boldsymbol{W}^\ell$ are sampled i.i.d. from a distribution with mean $\mu_{w^\ell} > 0$ and variance $\sigma^2_{w^\ell}$, with $\sigma_{w^\ell}/\mu_{w^\ell}$ smaller than order $\sqrt{n_\ell}$, r^ℓ obeys the recurrence relation

$$r^\ell = \sqrt{\frac{2}{n_{\ell+1}}} \frac{\sigma_{w^{\ell+1}}}{\mu_{w^{\ell+1}}} \sqrt{1 + (r^{\ell+1})^2} \tag{15}$$

with the initial condition $r^{L-1} = \sigma_{w^L}/\mu_{w^L}$ (see Appendix for derivation; layer-L GEVB updates are equal to the gradient by definition). Since $\sigma_{w^\ell}/\mu_{w^\ell}$ is smaller than order $\sqrt{n_\ell}$, r^ℓ scales inversely with n_ℓ after sufficiently many recursions. Our experiments in VNNs used a non-i.i.d. initialization described in Section 5 whose backward-pass behavior is qualitatively similar to using i.i.d. weights with $\sigma_{w^\ell}/\mu_{w^\ell} \sim 1$. Our measurements of GEVB alignment angles at initialization (Section 6; Fig. 2A) were in agreement with Eq. 15. Specifically, in layer $L-1$, we observed $\theta^{L-1} \sim 45°$, corresponding to $r^{L-1} \sim 1$; in layers $\ell < L-1$, we observed small θ^ℓ, corresponding to $r^\ell \ll 1$.

The limit of infinite network width has been shown to dramatically simplify the dynamics of gradient descent in DNNs, an idea embodied in the Neural Tangent Kernel (NTK) [27, 28]. Our finding that credit assignment is simplified in the same limit suggests a potentially fruitful connection between the NTK regime and biologically plausible DNN training.

5 Initializing nonnegative networks with ON/OFF cells

GEVB can be applied in networks initialized with zero weights, a method sometimes used with FA and DFA. However, BP is incompatible with zero initialization as the gradient vanishes. To enable fair comparisons between GEVB and BP, we used the same nonzero initialization for both training methods. Random nonnegative initializations have the added benefit of placing VNNs in a regime in which GEVB weight updates are highly gradient-aligned (Section 4). Unfortunately, nonnegative i.i.d. initializations are unsuitable for initializing DNNs as they tend to produce exploding forward passes. Specifically, if the weights in layer ℓ are sampled i.i.d. from a distribution with mean ~ 1 and variance $\sim 1/n_{\ell-1}$, the weight matrix has an outlier eigenvalue of size $\sim\sqrt{n_{\ell-1}}$, and the projection of the pre-activations onto this mode grows across layers. Here, we present a nonnegative initialization for VNNs based on ON/OFF cells that yields well-behaved forward propagation.

First, we group the hidden units into pairs with equal-and-opposite gating vectors $\boldsymbol{t}$, defining the structure of ON and OFF cells. We then initialize the weights such that the units in each pair have equal-and-opposite activations. The subnetwork of ON cells has the same activations as a network of half the size with i.i.d. mixed-sign weights, and thus the ON/OFF initialization exhibits well-behaved forward propagation insofar as the underlying mixed-sign initialization does. To initialize $\boldsymbol{W}^\ell$, we sample an i.i.d. mixed-sign weight matrix $\tilde{\boldsymbol{W}}^\ell$ of half the size, then construct $\boldsymbol{W}^\ell$ according to

$$\boldsymbol{W}^1 = \begin{pmatrix} +\tilde{\boldsymbol{W}}^1 \\ -\tilde{\boldsymbol{W}}^1 \end{pmatrix} \qquad \boldsymbol{W}^\ell = \begin{pmatrix} \left[+\tilde{\boldsymbol{W}}^\ell\right]^+ & \left[-\tilde{\boldsymbol{W}}^\ell\right]^+ \\ \left[-\tilde{\boldsymbol{W}}^\ell\right]^+ & \left[+\tilde{\boldsymbol{W}}^\ell\right]^+ \end{pmatrix} \quad \ell > 1 \tag{16}$$

where $[\cdot]^+$ denotes positive rectification. These weights are mixed-sign for $\ell = 1$ and nonnegative for $\ell > 1$, consistent with the definition of VNNs. The manner in which ON/OFF initialization avoids exploding forward propagation can be understood as follows. Let $(\lambda, \boldsymbol{v})$ denote an eigenvalue/vector pair of $\tilde{\boldsymbol{W}}^\ell$ and $(\lambda^+, \boldsymbol{v}^+)$ such a pair of abs$(\tilde{\boldsymbol{W}}^\ell)$. Then, $(\boldsymbol{v}, -\boldsymbol{v})/\sqrt{2}$ and $(\boldsymbol{v}^+, \boldsymbol{v}^+)/\sqrt{2}$ are eigenvectors of $\boldsymbol{W}^\ell$ with eigenvalues λ and λ^+, respectively. Thus, the spectrum of $\boldsymbol{W}^\ell$ is the union of the spectra of $\tilde{\boldsymbol{W}}^\ell$ and abs$(\tilde{\boldsymbol{W}}^\ell)$, which contains a large outlier eigenvalue from abs$(\tilde{\boldsymbol{W}}^\ell)$, a nonnegative i.i.d. matrix. However, each pre-activation vector has the form $\boldsymbol{h}^\ell = (\tilde{\boldsymbol{h}}^\ell, -\tilde{\boldsymbol{h}}^\ell)$, and the projection of $\boldsymbol{h}^\ell$ onto the eigenvector of the outlier eigenvalue is zero since $(\boldsymbol{v}^+, \boldsymbol{v}^+) \cdot (\tilde{\boldsymbol{h}}, -\tilde{\boldsymbol{h}}) = 0$. In BP, the backpropagated signal lacks the ON/OFF structure of the forward-propagated signal and therefore possesses a component along the eigenvector of the outlier eigenvalue, resulting in a growing backward pass. Our experiments used an optimizer with a normalizing effect, preventing large weight updates in deep layers when using BP [29]. Importantly, when using GEVB, well-behaved forward propagation is sufficient for well-behaved weight updates.

6 Experimental results

Here, we show that GEVB performs well in practice. To disentangle the impact on network performance of vectorization and nonnegativity, we trained vectorized and conventional networks, with and without a nonnegativity constraint. Vectorized networks were trained using GEVB and BP, and conventional networks were trained using DFA and BP. The nonlinearity in conventional networks was the scalar case of the VNN nonlinearity with $t = \pm 1$. Mixed-sign networks were initialized using He initialization, and nonnegative networks were initialized using ON/OFF initialization with an underlying He initialization [30]. In conventional nonnegative networks, the DFA feedback matrices were nonnegative [18]. We trained fully connected, convolutional, and locally connected architectures. Locally connected networks have the same receptive field structure as convolutional networks, but lack weight sharing [17]. We trained models on MNIST [31] and CIFAR-10 [32], using wider and deeper networks for CIFAR-10. We used ADAM for a fixed number of epochs (namely, 200), stopping early at zero training error. For each experiment, we performed five random initializations. Training lasted $\sim$10 days using five NVIDIA GeForce GTX 1080 Ti GPUs (see Appendix for details).

Our experimental results are summarized in Tables 1 and 2 (see Appendix for learning curves). We first examined the impact of vectorization and nonnegativity on model performance by considering errors under BP training. Across tasks and architectures, imposing vectorization or nonnegativity typically increased BP test and train errors, and imposing both increased the errors more than imposing either feature on its own. However, these increases in error were modest. One exception

Table 1: MNIST test errors (%). Train errors are shown in parentheses if greater than 0.005. Errors corresponding to GEVB in VNNs are shown in color. For each type of architecture, the smallest GEVB or DFA test error is **bold**.

	Vectorized networks			
	Nonnegative		Mixed-sign	
	GEVB	BP	GEVB	BP
Fully connected	**1.87 (0.06)**	1.93 (0.05)	2.32 (0.33)	1.84 (0.07)
Convolutional	2.33 (1.03)	1.3 (0.19)	1.83 (0.83)	0.8
Locally connected	1.78	1.64	1.84 (0.05)	1.44

	Conventional networks			
	Nonnegative		Mixed-sign	
	DFA	BP	DFA	BP
Fully connected	2.2 (0.3)	1.36	2.09 (0.19)	1.29
Convolutional	**1.56 (0.67)**	0.71	1.64 (0.42)	0.65
Locally connected	1.98 (0.32)	1.21	**1.48 (0.09)**	1.07

Table 2: CIFAR-10 test errors (%). Conventions are the same as in Table 1

	Vectorized networks			
	Nonnegative		Mixed-sign	
	GEVB	BP	GEVB	BP
Fully connected	**47.62 (1.25)**	47.03 (0.72)	48.86 (2.02)	45.98 (0.78)
Convolutional	**33.74 (28.83)**	30.85 (16.29)	38.43 (20.59)	30.54 (1.71)
Locally connected	41.08 (0.83)	41.01 (0.34)	40.11 (0.83)	38.77 (0.36)

	Conventional networks			
	Nonnegative		Mixed-sign	
	DFA	BP	DFA	BP
Fully connected	48.69 (3.28)	45.42 (0.96)	49.54 (2.88)	45.69 (0.73)
Convolutional	54.18 (48.43)	32.13 (0.23)	44.07 (17.25)	28.8 (0.15)
Locally connected	41.18 (10.62)	35.51	**39.41 (2.94)**	32.32

was locally connected networks trained on CIFAR-10, whose test error increased by 8.69% upon imposition of vectorization and nonnegativity. As the train error increased by only 0.34%, these features primarily diminished the generalization ability, rather than the capacity, of the model.

Next, we compared GEVB to BP in vectorized networks. In fully and locally connected VNNs, GEVB achieved test and train errors similar to those of BP (0.14% and 0.59% maximum discrepancies in test error on MNIST and CIFAR-10, respectively). In convolutional VNNs, we observed a performance gap between GEVB and BP (1.33% and 2.89% discrepancies in test error on MNIST and CIFAR-10, respectively). One possible reason for this gap is that weight sharing in convolutional networks breaks GEVB's gradient sign match guarantee. Nevertheless, on CIFAR-10, GEVB in convolutional VNNs yielded substantially lower test error than all other convolutional experiments using GEVB or DFA, a feature we examine in detail below. Finally, in vectorized mixed-sign networks, the performance gap between GEVB and BP was larger than in VNNs across tasks and architectures. Rather than enjoying guaranteed gradient alignment by virtue of nonnegative-constrained weights, learning in these networks relied on the tendency of GEVB to generate a bias toward positive weights, a special case of the feedback alignment effect [13].

As indicated by the **bold** entries in Tables 1 and 2, GEVB in vectorized networks in some cases outperformed DFA in conventional networks. In particular, GEVB had lower test error than DFA for fully connected architectures on MNIST, and for fully connected and convolutional architectures on CIFAR-10. This is particularly notable in light of the fact that vectorized networks have higher errors than conventional networks under BP training. Moreover, GEVB tended to produce considerably lower train errors than DFA.

To gain insight into our experimental results, we measured the alignment of both GEVB and DFA weight updates with the gradient over the course of training [13, 19]. Consistent with the theoretical result of Section 4, GEVB alignment angles at initialization were around $\sim 45^\circ$ for layer $L-1$, and small (typically under 20°) for layers $\ell < L-1$ (Fig. 2A). Over the course of training, these angles

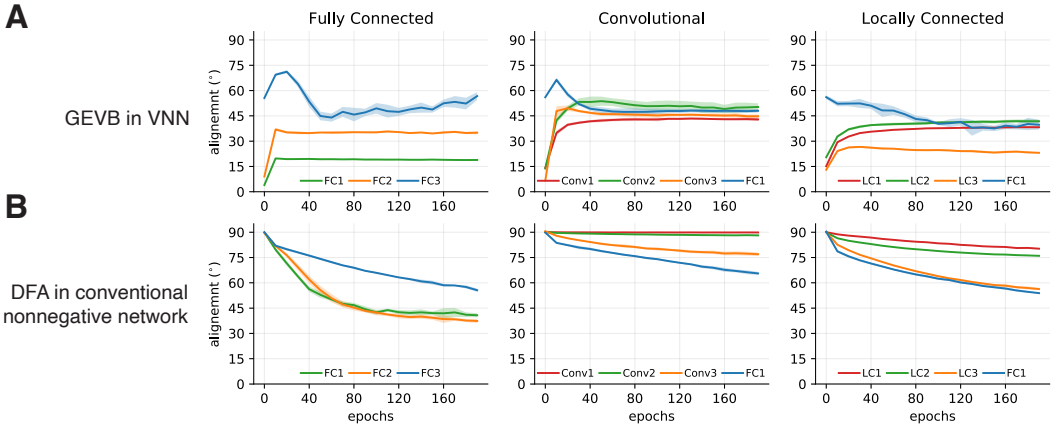

Figure 2: Gradient alignment angles over the course of training on CIFAR-10. In the legend, layers progress from deep to shallow from left to right. **(A)** GEVB in VNNs. **(A)** DFA in conventional nonnegative networks. Error bars are standard deviations across five runs. See Appendix for corresponding plots for mixed-sign networks and MNIST.

increased, plateauing around or below 45°. By contrast, DFA exhibited alignment angles ∼90° at initialization, and these angles dropped over the course of training due to the feedback alignment effect (Fig. 2B). The plateaued alignment angles for GEVB were typically smaller than those for DFA. In convolutional networks, the weight updates of DFA failed to align with the gradient in convolutional layers, consistent with large errors as well as prior studies [16, 17, 19, 20]. By contrast, GEVB exhibited plateaued alignment angles around 45° for convolutional layers.

Finally, we examined the ability of GEVB to train convolutional layers by studying the learned representations on CIFAR-10 at the output of the convolutional layers. t-SNE embeddings of these representations before and after training revealed that GEVB, but not DFA, improved the level of class clustering (Fig. 3A; see Appendix for t-SNE details) [33]. We quantified this improvement using a measure of cluster quality defined as $1 - (\frac{1}{K}\sum_{c=1}^{K}\overline{r_c^2})/\overline{r^2}$, where $\overline{r_c^2}$ and $\overline{r^2}$ denote the average class-c and average overall squared pairwise distance, respectively, in the t-SNE embedding. For all experiments, the cluster quality was slightly greater than zero at initialization and typically improved during training (Fig. 3B). DFA largely failed to improve cluster quality. By contrast, GEVB not only matched, but surpassed BP's improvement in cluster quality in vectorized networks.

The observation that GEVB not only trains convolutional layers, but surpasses BP's improvement in cluster quality, suggests that GEVB applies more pressure than BP on deep layers to extract task-relevant features. Understanding how the representations generated by GEVB differ from those generated by BP is an avenue for future study. Such an understanding could enable disambiguation of cortical learning algorithms on the basis of neural recordings.

7 Discussion

Biological implementation of vectorization: The GEVB learning rule raises the question of how vector units could be implemented in neural circuits. Here, we suggest two solutions. First, each vector unit might be single neuron, with vectorization unfolding in time. In each of K consecutive time bins, a different set of input neurons are active, producing a single component of the network output. This output and the corresponding component of the target serve as excitatory and inhibitory inputs to a neuromodulatory system, which broadcasts the resulting component of the error vector. Then, synapses undergo conventional (i.e., non-vectorized) three-factor Hebbian plasticity, $\Delta w_{ij}^{\ell} = -G(\boldsymbol{h}_i^{\ell})a_{j\mu}^{\ell-1}e_{\mu}$. Integrating these weight updates over all K time bins performs the sum over μ in Eq. 9, implementing the GEVB learning rule. As biases are vector-valued, different currents must be injected into each neuron at each time bin. One subtlety is that each neuron must be active or inactive at all time bins according to the sign of $\boldsymbol{t} \cdot \boldsymbol{h}$, and computing this inner product requires the full temporal input, violating causality. A simple workaround is to make $\boldsymbol{t}$ sparse in all but its first component, so that each neuron is gated based on its input during the first time bin. Alternatively,

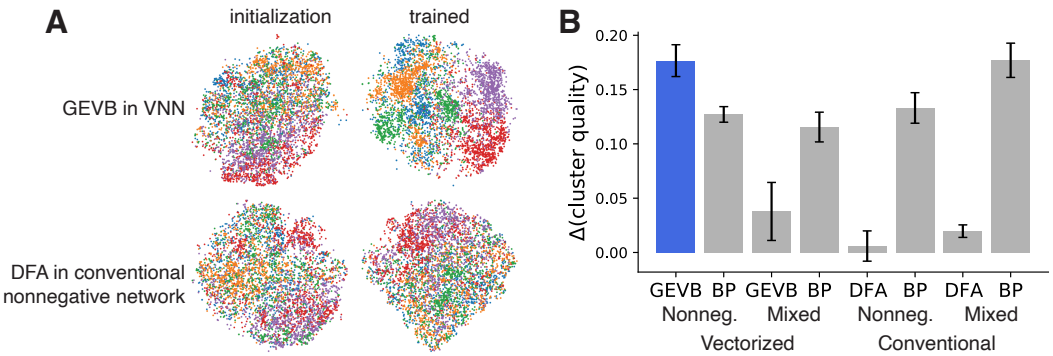

Figure 3: Comparison of learned representations at the output of the convolutional layers in networks trained on CIFAR-10. **(A)** t-SNE embeddings of five image classes before and after training for GEVB in VNNs (top) and DFA in conventional nonnegative networks (bottom). **(B)** Change in cluster quality due to training for all eight convolutional experiments on CIFAR-10. See main text for the definition of cluster quality. GEVB in VNNs is shown in color. Error bars are standard deviations across five runs.

gating could be performed based on an external signal independent of the pre-activations [34]. Second, each vector unit might correspond to a group of K neurons, where neuron μ in each presynaptic group projects only to neuron μ in all postsynaptic groups. As per vectorized weight sharing, the K synapses connecting corresponding neurons in a presynaptic and a postsynaptic group must share a common weight. This is possible if all K synapses undergo the same plasticity. However, this plasticity is non-local as the update to each synapse depends on the activity of all K neurons in the presynaptic group. Thus, we find the temporal interpretation of vectorization more plausible. Other implementations of vectorization are conceivable, such as solutions based on multiplexing information in the frequency domain or complex neuronal morphologies.

Scaling to high-dimensional outputs: Using a large output dimension K in VNNs is cumbersome as the computational complexity of the forward pass scales in K. One solution is to use target vectors of size $K \sim \log C$ for C-way classification with binary, rather than one-hot, encodings of the target class index. This is especially appropriate if the binary encoding is congruent with hierarchical structure in the data. For example, such structure is present in ImageNet [35]. If vectorization unfolds in time, it is natural to produce more coarse-grained classifications of the input first.

Prior uses of vector units in machine learning: Other works have used vector units for computational, rather than credit-assignment, purposes. Capsule Neural Networks use vector-valued capsules to represent object parts [36]. In [37], three-dimensional vector units with vectorized weight sharing are used to achieve equivariance to the rotation group $\mathrm{SO}(3)$. In the Gated Linear Network framework, each unit stores a vector of parameters describing a probability distribution that predicts the target [34, 38]. These networks use a form of vectorized weight sharing so that each unit represents a weighted geometric mixture of the distributions represented by the units in the previous layer. These weights were constrained to be nonnegative in [38].

Limitations and future directions: First, while we focused on MNIST and CIFAR-10, more challenging tasks are likely to elicit greater performance discrepancies between GEVB and other learning algorithms [17]. Second, while the GEVB learning rule is biologically sensible, we apply it in networks that are only loosely inspired by neural circuits. Training spiking networks using GEVB would be an initial step in assessing the capabilities of GEVB in more realistic neural circuits. Third, we focused on feedforward models. Extending GEVB to recurrent networks would provide a biologically plausible alternative to backpropagation through time. Finally, the GEVB learning rule was designed to be aligned with the gradient, but gradient-based learning might not be desirable in scenarios such as few-shot or continual learning. Future work should consider training algorithms that are both biologically plausible and overcome shortcomings of gradient-based learning [39].

Potential negative societal impacts: The computational complexity of the forward pass in VNNs is larger by a factor of K than in conventional networks. Thus, widespread adoption of VNNs in large-scale applications, however unlikely, could result in significant energy consumption.

References

[1] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

[2] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[3] Eric R Kandel, James H Schwartz, Thomas M Jessell, Steven Siegelbaum, A James Hudspeth, and Sarah Mack. *Principles of neural science*, volume 4. McGraw-hill New York, 2000.

[4] Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, et al. A deep learning framework for neuroscience. *Nature neuroscience*, 22(11):1761–1770, 2019.

[5] Francis Crick. The recent excitement about neural networks. *Nature*, 337(6203):129–132, 1989.

[6] Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.

[7] James CR Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in cognitive sciences*, 23(3):235–250, 2019.

[8] Peter Latham. Node perturbation in vanilla deep networks, Feb 2019. URL `http://www.gatsby.ucl.ac.uk/~pel/tn/notes/node_perturbation.pdf`.

[9] Pietro Mazzoni, Richard A Andersen, and Michael I Jordan. A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences*, 88(10): 4433–4437, 1991.

[10] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

[11] Justin Werfel, Xiaohui Xie, and H Sebastian Seung. Learning curves for stochastic gradient descent in linear feedforward networks. *Neural computation*, 17(12):2699–2718, 2005.

[12] Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science*, 11(1):23–63, 1987.

[13] Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1):1–10, 2016.

[14] Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. *arXiv preprint arXiv:1609.01596*, 2016.

[15] Arash Samadi, Timothy P Lillicrap, and Douglas B Tweed. Deep learning with dynamic spiking neurons and fixed feedback weights. *Neural computation*, 29(3):578–602, 2017.

[16] Julien Launay, Iacopo Poli, François Boniface, and Florent Krzakala. Direct feedback alignment scales to modern deep learning tasks and architectures. *arXiv preprint arXiv:2006.12878*, 2020.

[17] Sergey Bartunov, Adam Santoro, Blake A Richards, Luke Marris, Geoffrey E Hinton, and Timothy Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. *arXiv preprint arXiv:1807.04587*, 2018.

[18] Mathias Lechner. Learning representations for binary-classification without backpropagation. In *International Conference on Learning Representations*, 2019.

[19] Julien Launay, Iacopo Poli, and Florent Krzakala. Principled training of neural networks with direct feedback alignment. *arXiv preprint arXiv:1906.04554*, 2019.

[20] Maria Refinetti, Stéphane d'Ascoli, Ruben Ohana, and Sebastian Goldt. The dynamics of learning with feedback alignment. *arXiv preprint arXiv:2011.12428*, 2020.

[21] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Towards deep learning with segregated dendrites. *ELife*, 6:e22901, 2017.

[22] Blake A Richards and Timothy P Lillicrap. Dendritic solutions to the credit assignment problem. *Current opinion in neurobiology*, 54:28–36, 2019.

[23] Nicolas Frémaux and Wulfram Gerstner. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in neural circuits*, 9:85, 2016.

[24] Łukasz Kuśmierz, Takuya Isomura, and Taro Toyoizumi. Learning with three factors: modulating hebbian plasticity with errors. *Current opinion in neurobiology*, 46:170–177, 2017.

[25] David Balduzzi, Hastagiri Vanchinathan, and Joachim Buhmann. Kickback cuts backprop's red-tape: Biologically plausible credit assignment in neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.

[26] Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. signsgd: Compressed optimisation for non-convex problems. In *International Conference on Machine Learning*, pages 560–569. PMLR, 2018.

[27] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.

[28] Jaehoon Lee, Lechao Xiao, Samuel S Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *arXiv preprint arXiv:1902.06720*, 2019.

[29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[31] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[32] A Krizhevsky. Learning multiple layers of features from tiny images. *Master's thesis, University of Tronto*, 2009.

[33] David M Chan, Roshan Rao, Forrest Huang, and John F Canny. t-sne-cuda: Gpu-accelerated t-sne and its applications to modern data. In *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pages 330–338. IEEE, 2018.

[34] Joel Veness, Tor Lattimore, David Budden, Avishkar Bhoopchand, Christopher Mattern, Agnieszka Grabska-Barwinska, Eren Sezener, Jianan Wang, Peter Toth, Simon Schmitt, et al. Gated linear networks. *arXiv preprint arXiv:1910.01526*, 2019.

[35] Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu. Hd-cnn: hierarchical deep convolutional neural networks for large scale visual recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 2740–2748, 2015.

[36] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. *arXiv preprint arXiv:1710.09829*, 2017.

[37] Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulenard, Andrea Tagliasacchi, and Leonidas Guibas. Vector neurons: A general framework for so (3)-equivariant networks. *arXiv preprint arXiv:2104.12229*, 2021.

[38] David Budden, Adam Marblestone, Eren Sezener, Tor Lattimore, Greg Wayne, and Joel Veness. Gaussian gated linear networks. *arXiv preprint arXiv:2006.05964*, 2020.

[39] Jack Lindsey and Ashok Litwin-Kumar. Learning to learn with feedback and local plasticity. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21213–21223. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/file/f291e10ec3263bd7724556d62e70e25d-Paper.pdf.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See heading "Limitations and future directions" in Section 7.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See heading "Potential negative societal impacts" in Section 7.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] A minor assumption on which the proof of GEVB's gradient sign match property relies is given in the Appendix. The assumptions for the proof of Eq. 15 are stated in in Section 4.
   (b) Did you include complete proofs of all theoretical results? [Yes] We prove GEVB's gradient sign match property in Section 3.3. We prove Eq. 15 in the Appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We include code for training VNNs using both GEVB and BP in the supplemental material and describe these files in the Appendix.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] The training details are given in Section 6 and further details are given in the Appendix.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We ran each experiment five times and error bars are reported in all plots. While error bars are not reported in Tables 1 and 2 due to space constraints, this information is available in the plots of learning curves included in the Appendix
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Information about GPUs and training time is given in Section 7.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We used the standard MNIST and CIFAR-10 datasets, and we cite both.
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]