

BLOCK-LOCAL LEARNING WITH PROBABILISTIC LATENT REPRESENTATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

The ubiquitous backpropagation algorithm requires sequential updates through the network introducing a locking problem. In addition, backpropagation relies on the transpose of forward weight matrices to compute updates, introducing a weight transport problem across the network. Locking and weight transport are problems because they prevent efficient parallelization and horizontal scaling of the training process. We propose a new method to address both these problems and scale up the training of large models. Our method works by dividing a deep neural network into blocks and introduces a feedback network that propagates the information from the targets backwards to provide auxiliary local losses. Forward and backward propagation can operate in parallel and with different sets of weights, addressing the problems of locking and weight transport. Our approach derives from a statistical interpretation of training that treats output activations of network blocks as parameters of probability distributions. The resulting learning framework uses these parameters to evaluate the agreement between forward and backward information. Error backpropagation is then performed locally within each block, leading to “block-local” learning. Several previously proposed alternatives to error backpropagation emerge as special cases of our model. We present results on a variety of tasks and architectures, demonstrating state-of-the-art performance using block-local learning. These results provide a new principled framework for training networks in a distributed setting.

1 INTRODUCTION

Recent developments in machine learning have seen deep neural network architectures scale to billions of parameters (Touvron et al., 2023; Brown et al., 2020). While this has increased the power of these models to unprecedented levels, it has also pushed the computing hardware on which large network models run to its limits. As a result, it has become increasingly important to distribute the model training across many independent computing devices. However, today’s machine learning algorithms are poorly suited for distributed training. The error backpropagation algorithm requires an alternation of interdependent forward and backward phases, each requiring sequential computation. This introduces a locking problem because each phase must wait for the other (Jaderberg et al., 2016). Furthermore, the two phases rely on the same weight matrices to compute updates, which makes it impossible to separate memory spaces. This is referred to as the weight transport problem, see Grossberg (1987); Lillicrap et al. (2014a). Locking and weight transport are problems because they make efficient parallelization and horizontal scaling of large machine learning models across compute nodes extremely difficult.

We propose a new method to address these problems that distributes a globally defined optimization algorithm across a large network of computing devices using only local updates for training. Our approach utilizes a variational inference approach that uses results from probabilistic models to provide auxiliary local targets from a separate feedback network that propagates information from the targets to the input. Thus, messages can be communicated forward and backwards between computational nodes in parallel and include information about extracted features, which are updated using local probabilistic losses calculated using the targets provided by the feedback network. In contrast to previous results, optimizing these local losses does not require a contrastive step where different positive and negative samples are propagated through the network. Within each block, conventional error backpropagation is performed locally (“block local”) both in the forward network

and the backward feedback to adapt parameters during training. Performing forward and backward propagation in parallel mitigates the locking problem, and having a separate feedback network solves the weight transport problem.

The learning model developed here provides a new principled method for distributing the training of networks across multiple computing devices. The solutions emerging from this framework show striking similarities to those of previous models that used random feedback weights to provide local targets (Lee et al., 2015; Meulemans et al., 2020; Lillicrap et al., 2020; Ernoult et al., 2022), but we provide a principled way to train these feedback weights.

In summary, the contribution of this paper is threefold:

1. We provide a theoretical framework for interpreting the representations of deep neural networks as parameters of probability distributions.
2. Based on this probabilistic framework, we derive a new variational bound that allows us to decompose the global log-likelihood loss into a sum of local terms, which provides a principled approach to block-local training of these networks.
3. We show that this probabilistic learning method can achieve state-of-the-art performance on several benchmark classification tasks.

2 RELATED WORK

A number of methods for using local learning in DNNs have been introduced previously. Random feedback alignment (Lillicrap et al., 2016) and related approaches (Akrouf et al., 2019; Nøkland, 2016; Samadi et al., 2017) use fixed random feedback weights to back-propagate errors. Jaderberg et al. (2017) used layer-wise learned predictors of gradients, called “synthetic gradients” to decouple the training of different layers. Target propagation (Lee et al., 2015; Meulemans et al., 2020) has been demonstrated to have competitive performance using random projections for target labels instead of errors (Frenkel et al., 2021; Ernoult et al., 2022; Shibuya et al., 2023). Target Projection Stochastic Gradient Descent (tpSGD) Lomnitz et al. (2022) uses layer-wise SGD and local targets generated via random projections of the labels but does not adapt the backward weights. The network architecture used in our approach is similar to these prior works, but, in addition, provides a principled way to adapt feedback weights.

Some previous methods are based on probabilistic or energy-based cost functions. Jimenez Rezende et al. (2016) used a generative model and a KL-loss for local unsupervised learning of 3D structures. Contrastive learning (Chen et al., 2020; Oord et al., 2019) has been used to construct block-local losses (Xiong et al., 2020; Illing et al., 2021). Equilibrium propagation replaces target clamping with a target nudging phase (Scellier and Bengio, 2017). Another interesting contrastive approach, forward propagation, was recently introduced (Hinton, 2022; Zhao et al., 2023) which needs task-specific negative input examples. In contrast to these methods, our approach does not need separate positive and negative data samples and focuses on block-local learning. A number of methods have been proposed based on predictive coding framework (Millidge et al., 2022; Salvatori et al., 2022) but with a focus on biologically motivated generative models (Ororbia and Mali, 2019).

Other methods (Belilovsky et al., 2019; Löwe et al., 2019) have used greedy local, block- or layer-wise optimization. Notably, (Nøkland and Eidnes, 2019) achieved good results by combining a matching and a local cross-entropy loss. In contrast to our method, they used a similarity matching loss across mini-batches which prevents parallelization across a batch of data samples. Siddiqui et al. (2023) recently used block-local learning based on a contrastive cross-correlation metric over feature embeddings (Zbontar et al., 2021), demonstrating promising performance. Wu et al. (2021) used greedy layer-wise optimization of hierarchical autoencoders for video prediction. Wu et al. (2022) used an encoder-decoder stage for pre-training. In contrast to these methods, we do not rely solely on local greedy optimization but provide a principled way to combine local losses with feedback information without locking and weight transport across blocks and without contrastive learning.

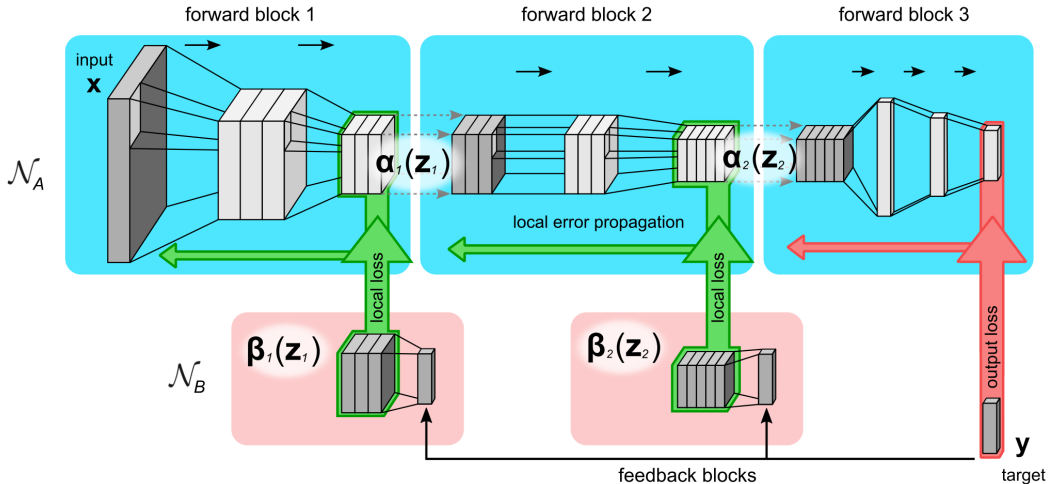


Figure 1: Illustration of use of block-local representations as learning signals on intermediate network layers. A deep neural network architecture \mathcal{N}_A is split into multiple blocks (forward blocks) and trained on an auxiliary local loss. Targets for local losses are provided by a feedback backward network \mathcal{N}_B .

3 A PROBABILISTIC FORMULATION OF DISTRIBUTED LEARNING

In this section we establish a method to partition a deep neural network into blocks by interpreting activations as parameters of probability distributions. We use these intermediate probabilistic representations at each block to derive block-local losses. To do this, we introduce a feedback network that accompanies the feedforward network to compute probabilistic representations. We show that the derived block-local losses and the resulting block-local learning (BLL) can be realized by a posterior bootstrapping mechanism that combines forward and feedback activations.

3.1 USING LATENT REPRESENTATIONS TO CONSTRUCT PROBABILISTIC BLOCK-LOCAL LOSSES

Learning in deep neural networks can be formulated probabilistically (Ghahramani, 2015) in terms of maximum likelihood, i.e. the problem is to minimize the negative log-likelihood $\mathcal{L} = -\log p(\mathbf{x}, \mathbf{y}) = -\log p(\mathbf{y} | \mathbf{x}) - \log p(\mathbf{x})$ with respect to the network parameters θ . For many practical cases where we may not be interested in the prior distribution of the input $p(\mathbf{x})$, we would like to directly minimize $\mathcal{L} = -\log p(\mathbf{y} | \mathbf{x})$.

This probabilistic interpretation of deep learning can be used to define block-local losses and distribute learning across multiple blocks of networks by introducing intermediate latent representations. The idea is illustrated in Fig. 1. A neural network \mathcal{N}_A computing the mapping $\mathbf{x} \rightarrow \mathbf{y}$ takes \mathbf{x} as input and its outputs can be interpreted as the statistical parameters of the conditional distribution $p(\mathbf{y} | \mathbf{x})$. When the network is split at intermediate layers into blocks, training using end-to-end gradient estimation can be replaced by estimators that optimize the blocks $\mathbf{x} \rightarrow \mathbf{z}_1$, $\mathbf{z}_1 \rightarrow \mathbf{z}_2$ \dots $\mathbf{z}_N \rightarrow \mathbf{y}$ separately. To see this, consider the gradient of the log-likelihood loss function

$$-\nabla \mathcal{L} = \nabla \log p(\mathbf{y} | \mathbf{x}), \quad (1)$$

where ∇ is the vector differential operator over parameters θ . For any deep network, it is possible to choose an intermediate activation at an arbitrary layer to represent a latent variable \mathbf{z}_k so that $p(\mathbf{y} | \mathbf{x}) = \mathbb{E}_{p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})} [p(\mathbf{y} | \mathbf{z}_k) p(\mathbf{z}_k | \mathbf{x})]$, where $\mathbb{E}_p[\cdot]$ denotes expectation with respect to p . Therefore, the representations of \mathbf{y} depend on \mathbf{x} only through \mathbf{z}_k , as expected for a feedforward network. Using this conditional independence property, the log-likelihood (1) expands to

$$-\nabla \mathcal{L} = \mathbb{E}_{p(\mathbf{z}_1 \dots \mathbf{z}_N | \mathbf{x}, \mathbf{y})} [\nabla \log p(\mathbf{z}_1 | \mathbf{x}) + \nabla \log p(\mathbf{z}_2 | \mathbf{z}_1) + \dots + \nabla \log p(\mathbf{y} | \mathbf{z}_N)]. \quad (2)$$

The identity in (2) is well known and also exploited in the derivation of the Expectation-Maximization (EM) algorithm (Dempster et al., 1977) (see Sec. S1.3 in the Supplement for a recap).

Computing the expectation with respect to $p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$ corresponds to the E-step and calculating the gradients corresponds to the M-step. The sum inside the expectation separates the gradient estimators into parts: $\mathbf{x} \rightarrow \mathbf{z}_1 \dots \mathbf{z}_N \rightarrow \mathbf{y}$. Importantly, the parts can have separate parameter spaces $\theta_k^{(1)}, \theta_k^{(2)}, \dots, \theta_k^{(N)}$ so that the gradient estimators become independent. This provides the core idea for how to split the training problem into smaller, and potentially more local sub-problems.

However, the E-step is impractical to compute for most interesting applications because of the combinatorial explosion in the state space of \mathbf{z}_k , which renders the expectation in Eq. (2) intractable. To get around this, we use a variational upper bound $\mathcal{F} \geq \mathcal{L}$ (Jordan et al., 1999). We introduce a feedback network with independent parameters (see Fig. 1), that is used to construct an auxiliary distribution $q(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$ to substitute the intractable posterior $p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$. The variational loss \mathcal{F} is then used to jointly minimize \mathcal{L} together with the distance between p and q . We demonstrate that this approach can be used to split gradients in a similar fashion to Eq. (2), yielding a distributed approximate solution to Eq. (1). In the next section, we describe how we construct the variational distribution q .

3.2 AUXILIARY LATENT REPRESENTATIONS

The probabilistic interpretation of hidden layer activity outlined above is valid under relatively mild assumptions, which we will establish here. It is important to note that at no point does the network produce samples of the implicit random variables \mathbf{z}_k ; they are introduced here only to conceptualize the mathematical framework. Instead, at block k , the network outputs the parameters of a probability distribution $\alpha_k(\mathbf{z}_k)$ (e.g., means and variances if α_k is Gaussian). $\alpha_k(\mathbf{z}_k) = p(\mathbf{z}_k | \mathbf{x})$ is the distribution over \mathbf{z}_k for given inputs \mathbf{x} . The network thus translates $\alpha_{k-1} \rightarrow \alpha_k \rightarrow \dots$ by outputting the statistical parameters of the conditional distribution $\alpha_k(\mathbf{z}_k)$ and taking the $\alpha_{k-1}(\mathbf{z}_{k-1})$ parameters as input. More specifically, the network implicitly calculates a marginal distribution

$$\alpha_k(\mathbf{z}_k) = p(\mathbf{z}_k | \mathbf{x}) = \mathbb{E}_{p(\mathbf{z}_{k-1} | \mathbf{x})} [p_k(\mathbf{z}_k | \mathbf{z}_{k-1})] = \mathbb{E}_{\alpha_{k-1}(\mathbf{z}_{k-1})} [p_k(\mathbf{z}_k | \mathbf{z}_{k-1})], \quad (3)$$

where $\mathbb{E}_p[\cdot]$ denotes expectation with respect to the probability distribution p . Consequently, the network realizes a conditional probability distribution $p(\mathbf{y} | \mathbf{x})$ (where \mathbf{x} and \mathbf{y} are network inputs and outputs, respectively). Eq. (3) is an instance of the belief propagation algorithm to efficiently compute conditional probability distributions. If all blocks have a rich enough expressive power (e.g. sufficient number of hidden layers) an accurate representation of the mappings between distributions can be learned in the network weights through error back-propagation. Thus, the distributions $p(\mathbf{z}_k | \mathbf{x})$ in the variational learning framework outlined above are realized simply by propagating inputs \mathbf{x} through the forward network \mathcal{N}_A .

To construct the variational distribution q , we introduce the backward network \mathcal{N}_B , which propagates messages β_k backward. Inferences about the posterior distribution $p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$ for any latent variable \mathbf{z}_k can be made using the belief propagation algorithm, which propagates messages $\alpha_k(\mathbf{z}_k)$ forward through the network using Eq. (3) and messages $\beta_k(\mathbf{z}_k) = q(\mathbf{y} | \mathbf{z}_k)$ backwards from the labels. In Section 4.2 we also experimented with a variant where feedback messages are propagated backward through a multi-layer network. In both cases the variational posterior can be computed up to normalization

$$\rho_k(\mathbf{z}_k) = q(\mathbf{z}_k | \mathbf{x}, \mathbf{y}) \propto p(\mathbf{z}_k | \mathbf{x}) q(\mathbf{y} | \mathbf{z}_k) = \alpha_k(\mathbf{z}_k) \beta_k(\mathbf{z}_k). \quad (4)$$

We make use of the fact that, through Eq. (3), the parameters of a probability distribution $p(\mathbf{z}_k | \mathbf{x})$ are a function of the parameters of $p(\mathbf{z}_i | \mathbf{x})$, for $0 < i < k$, e.g. if α is assumed to be Gaussian we have $(\mu(\alpha_k), \sigma^2(\alpha_k)) = f(\mu(\alpha_i), \sigma^2(\alpha_i))$, where $\mu(\cdot)$ and $\sigma^2(\cdot)$ are the mean and variance of the distribution respectively. Thus, if a network outputs $(\mu(\alpha_i), \sigma^2(\alpha_i))$ on layer i and $(\mu(\alpha_k), \sigma^2(\alpha_k))$ on layer k , a suitable probabilistic loss function will allow the network to learn f from examples. Therefore, the conditional distributions $p_k(\mathbf{z}_k | \mathbf{z}_{k-1})$ and the expectation in Eq. (3) are only implicitly encoded in the network weights. Clearly, the sub-networks that compute the transition from one latent variable to the next can have separated parameter spaces. We will use the exponential family of probability distributions for which this observation can be formalized more thoroughly, as described next.

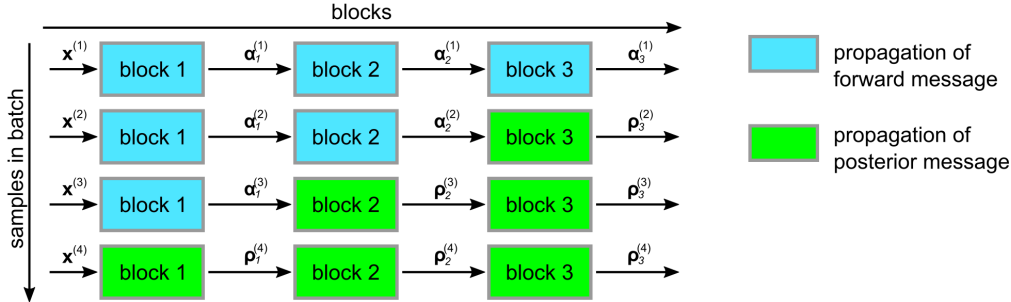


Figure 2: Illustration of posterior bootstrapping. Either the forward message α_k or the posterior message ρ_k is propagated for each sample and block.

3.3 EXPONENTIAL FAMILY DISTRIBUTIONS

To derive concrete losses and update rules for the forward and backward networks, we assume that α_k 's and β_k 's are from the exponential family (EF) of probability distributions, given by

$$\alpha_k(\mathbf{z}_k) = \prod_j \alpha_{kj}(z_{kj}) = \prod_j h(z_{kj}) \exp(T(z_{kj})\phi_{kj} - A(\phi_{kj})) , \quad (5)$$

with base measure h , sufficient statistics T , log-partition function A , and natural parameters ϕ_{kj} . This rich class contains the most common distributions, such as Gaussian, Poisson or Bernoulli, as special cases. For the example of a Bernoulli random variable we have $z_{kj} \in \{0, 1\}$, $T(z_{kj}) = z_{kj}$ and $A(\phi_{kj}) = \log(1 + e^{\phi_{kj}})$ (Koller and Friedman, 2009). One interesting property of the EF is that the Kullback-Leibler (KL-) divergence, to measure the distance between two distributions ρ_k and α_k , with parameters γ_k and ϕ_k , can be expressed using only the means μ and variances σ^2 of the distributions, i.e.

$$-\nabla \mathcal{D}_{KL}(\rho_k | \alpha_k) = \sum_j (\mu(\rho_{kj}) - \mu(\alpha_{kj})) \nabla \phi_{kj} + \sigma^2(\rho_{kj}) (\phi_{kj} - \gamma_{kj}) \nabla \gamma_{kj} . \quad (6)$$

We will exploit this property to construct local learning rules that can be computed efficiently. A network directly implements an EF distribution if the activations a_{kj} at block k encode the natural parameters, $a_{kj} = \phi_{kj}$.

To summarize, a feed-forward DNN $\mathcal{N}_A : \mathbf{x} \rightarrow \mathbf{y}$, can be split into $N + 1$ blocks by introducing implicit latent variables $\mathbf{z}_k : \mathbf{x} \rightarrow \mathbf{z}_k \rightarrow \mathbf{y}$, and generating the respective natural parameters. Blocks can be separated after any arbitrary layer, but some splits may turn out more natural for a particular network architecture. If both distributions α_{kj} and β_{kj} are (assumed to be) members of the EF with natural parameters $\phi_{kj,\alpha}$ and $\phi_{kj,\beta}$, then ρ_{kj} is also EF with parameters $\frac{1}{2}(\phi_{kj,\alpha} + \phi_{kj,\beta})$ ¹.

3.4 MODULARIZED LEARNING USING LOCAL VARIATIONAL LOSSES AND POSTERIOR BOOTSTRAPPING

We construct and use an upper bound on the actual log-likelihood loss \mathcal{L} for training the model. This upper bound consists only of block-local losses ℓ at all network blocks and is constructed using the forward and feedback networks \mathcal{N}_A and \mathcal{N}_B , respectively, as shown in the Supplement. The local loss ℓ can be written as

$$\ell(p_k, \beta_k | \alpha_{k-1}) = \mathcal{D}_{KL}(q_k | \alpha_k) + \mathcal{H}(p_k | \alpha_{k-1}) , \quad (7)$$

where the first divergence term measures the mismatch between the posterior ρ_k and the forward message α_k , and the second term is an entropy loss that determines the quality of the distributions when propagating data through \mathcal{N}_A , based on the variational posterior (see Supplementary Information S1.2.1 for details).

¹The dependence is linear and can be augmented with an arbitrary constant factor. $\frac{1}{2}$ is conveniently chosen here because it assures that forward messages and posterior parameters are of the same scale.

The loss in Eq. (7) is local in the sense that it is completely determined by the information available at block k , i.e., the local network transfer function specifying p_k , the forward message from the previous block α_{k-1} , and the feedback β_k . Furthermore, the loss is local with respect to learning, i.e. it doesn't require global signals to be communicated to each block. In this sense, our approach differs from previous contrastive methods that need to distinguish between positive and negative samples. In our approach, any sample that passes through a block can be used directly for weight updating and is treated in the same way.

To arrive at this key result we use a new approach that we call ‘‘posterior bootstrapping’’. Posterior bootstrapping combines the information provided by the forward and backward network during learning by propagating one of either the forward message α_k or the parameters to the posterior message ρ_k to the next block. Whether α_k or ρ_k is passed for every sample and every block is determined by a bootstrapping schedule. The optimal schedule is derived in Supplementary Information S1.2.1 and is shown in Fig. 2, where the pattern of posterior propagation forms a block-triangular matrix, giving blocks close to the input a tendency to preferentially propagate forward messages. Computing the posterior in this EF model is computationally very cheap as outlined above, so it introduces no significant overhead. Posterior bootstrapping also does not introduce a locking problem because the backward messages \mathbf{b}_k are simultaneously available at all blocks.

Based on posterior bootstrapping and optimization of ℓ , it is possible to construct an unbiased learning algorithm for the networks \mathcal{N}_A and \mathcal{N}_B . In Supplementary Information S1.2.1 we show the following theorem in detail

Theorem 1. *Let ℓ be the local loss function Eq.(7). Furthermore, let $\alpha_k^{(m)}, \beta_k^{(m)}$ be messages created using the posterior bootstrapping schedule outlined above. Then simultaneous minimization of $\ell(p_k^{(m)}, \beta_k^{(m)} | \alpha_{k-1}^{(m)})$ in all blocks k , minimizes an upper bound on the log-likelihood loss \mathcal{L} .*

The proof of Theorem 1 and additional details are presented in 2 steps in the Supplement. First we show that an upper bound to \mathcal{L} can be constructed by adding loss terms of the form (7). We then show that posterior bootstrapping computes the expectations that are needed to provide the forward and backward messages $\alpha_k^{(m)}$ and $\beta_k^{(m)}$. In simulations we also tested a simpler bootstrapping schedule that just passes forward message α_k through the network.

3.5 GREEDY LOCAL FORWARD AND FEEDBACK NETWORK OPTIMIZATION

We do the overall training of the model using a greedy block-local learning strategy, meaning that we treat the inputs as constants and do not apply the chain rule across block boundaries. We apply a greedy learning strategy to train the feedback network \mathcal{N}_B as well. The role of the feedback network is to propagate information about the labels back to the blocks of the forward network, providing local targets for the losses ℓ . The construction of the feedback network is therefore arbitrary and need not reflect the complexity of the forward network. In this paper, we use the simplest version, where each block of \mathcal{N}_B is given by a single linear layer. This special case is of particular interest because it allows us use a closed form solution to the optimization problem to find the parameters of the feedback network that minimize ℓ . In the Supplement Sec. S1.3 and S1.5 we show the closed-form solution is $\theta_k^{(b)} \stackrel{!}{=} \arg \min_{\theta_k^{(b)}} \ell(p_k, \beta_k | \alpha_{k-1})$ and convergence properties.

3.6 DISTRIBUTED VARIATIONAL LEARNING

In summary, the BLL algorithm is given by Table 1. The two for loops can be interleaved and parallelized by pipelining the propagation of data samples through the network as shown in Fig. 3. Updates can be computed as soon as propagation through a given block is complete. There is no locking, since only the data labels are needed to compute the output of the backward network. Furthermore, there is no weight transport problem since parameter spaces are separated and updates are computed only locally.

BLL shares many similarities with earlier methods. In particular, Direct Feedback Alignment (DFA) propagates targets through random weights to create local learning targets and can therefore be seen as a special case of BLL where feedback weights are kept fixed and the number of blocks equals the number of layers in the model. The loss term that emerges in BLL also shows some similarity

```

for all pairs  $\mathbf{x}, \mathbf{y}$  in the training data set, and learning rate  $\eta$  do
   $\mathbf{a}_0 \leftarrow \mathbf{x}$ 
  for  $1 \leq k \leq N$  do
     $\beta_k \leftarrow g_k(\mathbf{y})$  ▷ Feedback network
     $\alpha_k \leftarrow f_k(\alpha_{k-1})$  ▷ Forward network, computation depends on previous block
     $\theta_k^{(b)} \leftarrow \arg \min_{\theta_k^{(b)}} \ell(p_k, \beta_k | \alpha_{k-1})$ 
     $\theta_k^{(a)} \leftarrow \theta_k^{(a)} + \eta \nabla_{\theta_k^{(a)}} \ell(p_k, \beta_k | \alpha_{k-1})$ 
    if (posterior bootstrapping) then
       $\alpha_k \leftarrow \rho_k$ 
  
```

Table 1: Pseudo code of the BLL training algorithm. The for loops can be interleaved and run in parallel. Colors correspond to the operations in Figure 3

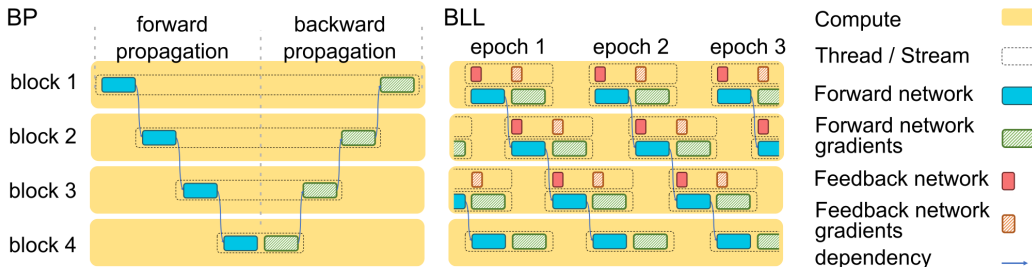


Figure 3: Timeline of execution for error backpropagation (BP) and BLL. BLL presented as a simplest case with forward-only bootstrapping.

with the predictive loss proposed in Nøklund and Eidnes (2019), but losses are derived here from a probabilistic framework and used to simultaneously learn the forward network and local targets.

4 RESULTS

4.1 BLOCK-LOCAL LEARNING OF VISION BENCHMARK TASKS

We evaluated the BLL algorithm on three vision tasks: Fashion-Mnist, CIFAR-10 and Imagenet-1K. Its performance is compared on ResNet18 and ResNet50 architectures with that of Backpropagation (BP), Feedback Alignment (Lillicrap et al., 2014b) (FA) and Local learning using similarity matching loss (Pred-Sim) from (Nøklund and Eidnes, 2019). The ResNet architectures were divided into

Architecture	Algorithm	Fashion-MNIST		CIFAR-10		ImageNet-1K	
		test-1	test-3	test-1	test-3	test-1	test-5
ResNet-18	BLL	94.2	99.3	88.3	98		
	BP	92.7	99.3	95.2	99.3		
	FA	87.9	98.6	70.4	92.5		
	Pred-Sim	93.9	99.3	88	97.7		
ResNet-50	BLL	94.3	99.1	92.6	99.1	53.6	77.1
	BP	93.4	99.4	94	99.2	76.1	92.9
	FA	83.1	97.9	70.3	92		
	Pred-Sim	94.3	99.6	92.4	98.8		

Table 2: Classification accuracy (% correct) on vision tasks. BP: end-to-end backprop, FA: Feedback Alignment, BLL: block local learning, Sim Loss: Local learning with similarity matching loss (Nøklund and Eidnes, 2019). Top-1,3 and 5 accuracies are reported in the respective columns.

four blocks for BLL and Pred-Sim. The splits were introduced after residual layers by grouping subsequent layers into blocks. We also included the predictive loss as suggested in (Nøkland and Eidnes, 2019) in our BLL method (see ablation studies in Supplement to see the role of individual losses in training performance).

Group sizes in the blocks were (4,5,4,5) for ResNet-18 and (12,13,12,13) for ResNet-50. Backward networks for BLL were constructed as linear layers with label size as input and the output size equal to the number of channels in the corresponding ResNet block output. The kernels of ResNet-18/ResNet-50 used by FA architectures during backpropagation were fixed and uniformly initialised following the Kaiming He et al. (2015) initialisation method.

We train ResNet-50 on ImageNet-1K using the standard ImageNet training pipeline from Pytorch (Paszke et al., 2019) without any additional augmentation. We use FFCV (Leclerc et al., 2022) data-loading and training scripts to speed up training. Additional training details and hyperparameters are documented in the Supplement.

The results are summarized in Table 2, top-k test accuracies are shown. Top-3 accuracies count the number of test samples for which the correct class was among the network’s three highest output activations (see Supplement for results over multiple runs). BLL performs slightly better than Pred-Sim overall for all tasks and architectures. It also performs close to end-to-end backpropagation performance except for CIFAR-10 using ResNet18 and ImageNet task, hinting at insufficient information being sent to the blocks through the linear feedback network. Unsurprisingly, FA is outperformed by BLL, the gap becoming wider as the task and model complexity increases (Bartunov et al., 2018).

4.2 BLOCK-LOCAL TRANSFORMER ARCHITECTURE FOR SEQUENCE-TO-SEQUENCE LEARNING

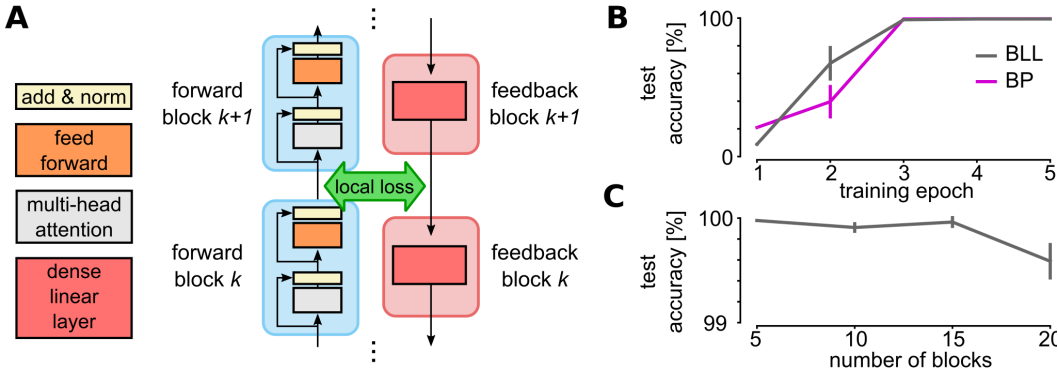


Figure 4: Block local learning of transformer architecture. **A**: Illustration of the transformer forward and feedback network. **B**: Learning curves of block local (BLL) and end-to-end backpropagation (BP) training. **C**: Test accuracy vs. number of blocks in the transformer model. Error bars show standard deviations over 5 runs.

Transformer architectures are well suited for distributed computing due to their repetitive network structure. We demonstrate a proof-of-concept result on training a transformer with BLL. We used a transformer model with 20 self-attention blocks with a single attention head each. Block local losses were added after each block and trained locally. The feedback network was constructed here as a multi-layer network by projecting targets through dense layers and used the local loss for training. See Fig. 4 A for an illustration. The transformer was trained for 5 epochs on a sequence-to-sequence task, where a random permutation of numbers 0..9 was presented and had to be re-generated in reverse order.

BLL achieves a convergence speed that is comparable to that of end-to-end BP on this task. Fig. 4 B shows learning curves of BLL and BP. Both algorithms converge after around 3 epochs to nearly perfect performance. BLL also achieved good performance for a wide range of network depths. Fig. 4 C shows the performance after 5 epochs for different transformer architectures. Using only 5 transformer blocks yields performance of around 99.9% (average over five independent runs). The

test accuracy on this task for the 20 block transformer was 99.6%. These results suggest that BLLod is equally applicable to transformer architectures.

5 DISCUSSION

We have demonstrated a probabilistic framework for rigorously defining block-local losses for deep architectures. Our method represents the parameters of probability distributions using the network activations and introduces a feedback network that propagates information backwards from the targets to the input to provide targets for intermediate layers. These targets can be interpreted as prototypical representations that each block must achieve in order to solve the overall classification task. The forward network and the backward feedback can work in parallel and with different sets of weights, solving the locking problem and the weight transport problem. We have shown that our block-local training approach outperforms existing local training approaches and approaches the task performance of backprop in some cases.

While we used linear layers for the feedback network in most of this work, which scales to mid-sized learning problems, in Section 4.2 we demonstrated a proof of concept of using more complex feedback structures. It will be interesting to explore potentially biologically realistic feedback structures for future work as well.

We also showed that our method can scale up to ImageNet and work on different architectures, including transformers. Both of these results on complex tasks and network structures suggest that BLL can scale up to very large models. Our method not only provides a novel way of performing distributed training of large models but also hints at new paradigms of self-supervised training that are biologically plausible.

The proposed method may also help further blur the boundary between deep learning and probabilistic models. Several previous models have shown that DNNs can represent probability distribution (Abdar et al., 2021; Pawlowski et al., 2017; Tran et al., 2019; Malinin and Gales, 2019). Unlike these previous methods, our method does not require Monte Carlo sampling or contrastive training. Instead, it exploits the log-linear structure of exponential family distributions to propagate probabilistic messages efficiently. In fact, we found that combining our local loss with the information-theoretic predictive loss proposed in (Nøkland and Eidnes, 2019) gave the best results. Although BLL was derived using a probabilistic approach, it also shares interesting similarities with earlier non-probabilistic method, such as DirectFeedback Alignment.

Overall, this work addresses an important open problem of modern ML: How can ML models be efficiently distributed and horizontally scaled over many compute nodes for training models too large to fit on one node. Doing so may also allow us to train large models more efficiently, since it would allow us to distribute computation over many smaller, energy-efficient devices rather than a large power-hungry device. This would also make our method especially well suited for new energy efficient hardware for ML, such as neuromorphic devices. The energy consumption and resulting carbon footprint of ML is becoming a major concern and the proposed training method may provide a new direction to reduce the impact of ML.

REPRODUCIBILITY

We ensure that the results presented in this paper are easily reproducible using just the information provided in the main text as well as the supplement. Details of the models used in our simulations are presented in the main paper and further elaborated in the supplement. We provide additional details and statistics over multiple runs in the supplement section S2. We use publicly available libraries and datasets in our simulations. We will further provide the source code to the reviewers and ACs in an anonymous repository once the discussion forums are opened. This included code will also contain “readme” texts to facilitate easy reproducibility. The theoretical analysis provided in Section 3 is derived in the supplement.

REFERENCES

- Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76:243–297, 2021.
- Mohamed Akrouf, Collin Wilson, Peter Humphreys, Timothy Lillicrap, and Douglas B Tweed. Deep learning without weight transport. *Advances in neural information processing systems*, 32, 2019.
- Sergey Bartunov, Adam Santoro, Blake A. Richards, Luke Marris, Geoffrey E. Hinton, and Timothy P. Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 9390–9400, Red Hook, NY, USA, 2018. Curran Associates Inc.
- Eugene Belilovsky, Michael Eickenberg, and Edouard Oyallon. Greedy layerwise learning can scale to ImageNet. In *Proceedings of the 36th International Conference on Machine Learning*, pages 583–593. PMLR, 2019. URL <https://proceedings.mlr.press/v97/belilovsky19a.html>.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. *arXiv:2005.14165 [cs]*, July 2020. URL <http://arxiv.org/abs/2005.14165>.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1): 1–22, 1977. ISSN 00359246. doi: 10.1111/j.2517-6161.1977.tb01600.x. URL <https://onlinelibrary.wiley.com/doi/10.1111/j.2517-6161.1977.tb01600.x>.
- Maxence M Ernout, Fabrice Normandin, Abhinav Moudgil, Sean Spinney, Eugene Belilovsky, Irina Rish, Blake Richards, and Yoshua Bengio. Towards scaling difference target propagation by learning backprop targets. In *International Conference on Machine Learning*, pages 5968–5987. PMLR, 2022.
- Charlotte Frenkel, Martin Lefebvre, and David Bol. Learning without feedback: Fixed random learning signals allow for feedforward training of deep neural networks. *Frontiers in Neuroscience*, 15, 2021. ISSN 1662-453X. URL <https://www.frontiersin.org/articles/10.3389/fnins.2021.629892>.
- Zoubin Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553): 452–459, 2015.
- Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science*, 11(1):23–63, 1987.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, 2015.
- Geoffrey Hinton. The forward-forward algorithm: Some preliminary investigations. *arXiv preprint arXiv:2212.13345*, 2022.
- Bernd Illing, Jean Ventura, Guillaume Bellec, and Wulfram Gerstner. Local plasticity rules can learn deep representations using self-supervised contrastive predictions. In *Advances in Neural Information Processing Systems*, volume 34, pages 30365–30379. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/feadeld2047977cd0cefdafc40175a99-Abstract.html>.

- Max Jaderberg, Wojciech Marian Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, and Koray Kavukcuoglu. Decoupled Neural Interfaces using Synthetic Gradients. *arXiv:1608.05343 [cs]*, August 2016. URL <http://arxiv.org/abs/1608.05343>.
- Max Jaderberg, Wojciech Marian Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, and Koray Kavukcuoglu. Decoupled neural interfaces using synthetic gradients. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1627–1635. PMLR, 2017. URL <https://proceedings.mlr.press/v70/jaderberg17a.html>.
- Danilo Jimenez Rezende, S. M. Ali Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/hash/1d94108e907bb8311d8802b48fd54b4a-Abstract.html>.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37:183–233, 1999.
- Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. pages 32–33, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- Guillaume Leclerc, Andrew Ilyas, Logan Engstrom, Sung Min Park, Hadi Salman, and Aleksander Madry. ffcv. <https://github.com/libffcv/ffcv/>, 2022. commit xxxxxx.
- Dong-Hyun Lee, Saizheng Zhang, Asja Fischer, and Yoshua Bengio. Difference target propagation. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part I 15*, pages 498–515. Springer, 2015.
- Timothy P. Lillicrap, Daniel Cownden, Douglas B. Tweed, and Colin J. Akerman. Random feedback weights support learning in deep neural networks. *arXiv:1411.0247 [cs, q-bio]*, November 2014a. URL <http://arxiv.org/abs/1411.0247>.
- Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random feedback weights support learning in deep neural networks. *arXiv preprint arXiv:1411.0247*, 2014b.
- 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): 13276, 2016. ISSN 2041-1723. doi: 10.1038/ncomms13276. URL <https://www.nature.com/articles/ncomms13276>.
- Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton. Back-propagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020. ISSN 1471-0048. doi: 10.1038/s41583-020-0277-3. URL <https://www.nature.com/articles/s41583-020-0277-3>.
- Michael Lomnitz, Zachary Daniels, David Zhang, and Michael Piacentino. Learning with local gradients at the edge, 2022. URL <http://arxiv.org/abs/2208.08503>.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts, 2017.
- Sindy Löwe, Peter O’ Connor, and Bastiaan Veeling. Putting an end to end-to-end: Gradient-isolated learning of representations. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/851300ee84c2b80ed40f51ed26d866fc-Abstract.html>.
- Andrey Malinin and Mark Gales. Reverse kl-divergence training of prior networks: Improved uncertainty and adversarial robustness. *Advances in Neural Information Processing Systems*, 32, 2019.

- Alexander Meulemans, Francesco Carzaniga, Johan Suykens, João Sacramento, and Benjamin F Grewe. A theoretical framework for target propagation. *Advances in Neural Information Processing Systems*, 33:20024–20036, 2020.
- Beren Millidge, Tommaso Salvatori, Yuhang Song, Rafal Bogacz, and Thomas Lukasiewicz. Predictive coding: towards a future of deep learning beyond backpropagation? *arXiv preprint arXiv:2202.09467*, 2022.
- Radford M. Neal and Geoffrey E. Hinton. *A View of the Em Algorithm that Justifies Incremental, Sparse, and other Variants*, pages 355–368. Springer Netherlands, Dordrecht, 1998. ISBN 978-94-011-5014-9. doi: 10.1007/978-94-011-5014-9_12. URL https://doi.org/10.1007/978-94-011-5014-9_12.
- Arild Nøkland. Direct Feedback Alignment Provides Learning in Deep Neural Networks. *arXiv:1609.01596 [cs, stat]*, September 2016. URL <http://arxiv.org/abs/1609.01596>.
- Arild Nøkland and Lars Hiller Eidnes. Training neural networks with local error signals. In *International conference on machine learning*, pages 4839–4850. PMLR, 2019.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2019. URL <http://arxiv.org/abs/1807.03748>.
- Alexander G Ororbia and Ankur Mali. Biologically motivated algorithms for propagating local target representations. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 4651–4658, 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Nick Pawłowski, Andrew Brock, Matthew CH Lee, Martin Rajchl, and Ben Glocker. Implicit weight uncertainty in neural networks. *arXiv preprint arXiv:1711.01297*, 2017.
- Tommaso Salvatori, Luca Pinchetti, Beren Millidge, Yuhang Song, Tianyi Bao, Rafal Bogacz, and Thomas Lukasiewicz. Learning on arbitrary graph topologies via predictive coding. *Advances in neural information processing systems*, 35:38232–38244, 2022.
- 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, January 2017. ISSN 0899-7667. doi: 10.1162/NECO.a.00929. URL <https://doi.org/10.1162/NECO.a.00929>.
- Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.
- Tatsukichi Shibuya, Nakamasa Inoue, Rei Kawakami, and Ikuro Sato. Fixed-weight difference target propagation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 9811–9819, 2023.
- Shoaib Ahmed Siddiqui, David Krueger, Yann LeCun, and Stéphane Deny. Blockwise self-supervised learning at scale, 2023. URL <http://arxiv.org/abs/2302.01647>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. LLaMA: Open and Efficient Foundation Language Models, February 2023. URL <http://arxiv.org/abs/2302.13971>.

- Dustin Tran, Mike Dusenberry, Mark Van Der Wilk, and Danijar Hafner. Bayesian layers: A module for neural network uncertainty. *Advances in neural information processing systems*, 32, 2019.
- Bohan Wu, Suraj Nair, Roberto Martin-Martin, Li Fei-Fei, and Chelsea Finn. Greedy hierarchical variational autoencoders for large-scale video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2318–2328, 2021.
- Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. TinyViT: Fast pretraining distillation for small vision transformers, 2022. URL <http://arxiv.org/abs/2207.10666>.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Yuwen Xiong, Mengye Ren, and Raquel Urtasun. LoCo: Local contrastive representation learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 11142–11153. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/7fa215c9efebb3811a7ef58409907899-Abstract.html>.
- Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction, 2021. URL <http://arxiv.org/abs/2103.03230>.
- Gongpei Zhao, Tao Wang, Yidong Li, Yi Jin, Congyan Lang, and Haibin Ling. The cascaded forward algorithm for neural network training. *arXiv preprint arXiv:2303.09728*, 2023.

SUPPLEMENTARY INFORMATION

S1 A PROBABILISTIC FORMULATION OF DISTRIBUTED LEARNING

S1.1 MARKOV CHAIN MODEL

Here, we provide additional details to the learning model presented in Section 3 of the main text. To establish these results, we consider the Markov chain model $\mathbf{x} \rightarrow \mathbf{z}_1 \rightarrow \mathbf{z}_2 \rightarrow \dots \rightarrow \mathbf{y}$ of a DNN split into $N + 1$ blocks, with inputs \mathbf{x} , outputs \mathbf{y} and intermediate representations \mathbf{z}_k at block k . To simplify the notation we will define the input $\mathbf{z}_0 := \mathbf{x}$ and output $\mathbf{z}_{N+1} := \mathbf{y}$, and $\mathbf{z} = \{\mathbf{z}_k\}, 1 \leq k \leq N$, the auxiliary latent variables. The DNN suggests a conditional independence structure given by the first-order Markov chain of random variables \mathbf{z}_k

$$p(\mathbf{y}, \mathbf{z} | \mathbf{x}) = p(\mathbf{z}_1 \dots \mathbf{z}_{N+1} | \mathbf{z}_0) = \prod_{k=1}^{N+1} p_k(\mathbf{z}_k | \mathbf{z}_{k-1}, \boldsymbol{\theta}_k), \quad (\text{S1})$$

where $p_k(\mathbf{z}_k | \mathbf{z}_{k-1}, \boldsymbol{\theta}_k)$ is the input-output mapping of the k -th block subject to block-local network parameters $\boldsymbol{\theta}_k$. If it is clear from the context, we will omit the explicit mention of the parameter vectors $\boldsymbol{\theta}_k$. The computation of messages α_k comes naturally in a feed-forward neural network as the flow of information follows the canonical form, input \rightarrow output. Every block of the network thus translates $\alpha_{k-1} \rightarrow \alpha_k$ by outputting the statistical parameters of the conditional distribution $p(\mathbf{z}_k | \mathbf{x})$ and takes $p(\mathbf{z}_{k-1} | \mathbf{x})$ as input. This interpretation is valid for a suitable split of any DNN into $N + 1$ blocks (N is the number of splits) that fulfills a mild set of conditions (see Section S1.4 for details). It is important to note that the random variables $(\mathbf{z}_1, \mathbf{z}_2, \dots)$ are only implicit. The network generates the parameters to the probability distribution and at no points needs to sample values for these random variables.

S1.2 USING LATENT REPRESENTATIONS TO CONSTRUCT PROBABILISTIC BLOCK-LOCAL LOSSES

Many commonly used loss functions in deep learning have a probabilistic interpretation, e.g., the cross entropy loss of a binary classifier is identical to the Bernoulli log-likelihood, and the mean squared error corresponds to the log-likelihood of a Gaussian with constant variance. In this formulation, the outputs of the DNN are interpreted as the statistical parameters to a conditional probability distribution (e.g., the mean of a Gaussian) and the loss function measures the support of observed data samples \mathbf{x} and \mathbf{y} .

To introduce intermediate block-local representations \mathbf{z}_k in the network, we consider a variational upper bound \mathcal{F} to the log-likelihood loss \mathcal{L}

$$\mathcal{F} = -\log p(\mathbf{y} | \mathbf{x}) + \frac{1}{N} \sum_{k=1}^N \mathcal{D}_{KL}(q_k | p_k) \geq \mathcal{L}, \quad (\text{S2})$$

where p_k and q_k are true and variational posterior distributions over latent variables $p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$ and $q(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$, respectively. Using the Markov property (S1), assuming a fully factorized distribution, implies the conditional independence

$$p(\mathbf{y}, \mathbf{z}_k | \mathbf{x}) = p(\mathbf{y} | \mathbf{z}_k) p(\mathbf{z}_k | \mathbf{x}), \quad (\text{S3})$$

for any k . Using this property, we can rewrite Eq. (S2)

$$\begin{aligned} \mathcal{F} &= -\log p(\mathbf{y} | \mathbf{x}) + \frac{1}{N} \sum_{k=1}^N \mathcal{D}_{KL}(q_k | p_k) = \frac{1}{N} \sum_{k=1}^N \mathbb{E}_{q_k} \left[\log \frac{q(\mathbf{z}_k | \mathbf{x}, \mathbf{y})}{p(\mathbf{y}, \mathbf{z}_k | \mathbf{x})} \right] \\ &= \frac{1}{N} \sum_{k=1}^N \mathbb{E}_{q_k} \left[\log \frac{q(\mathbf{z}_k | \mathbf{x}, \mathbf{y})}{p(\mathbf{z}_k | \mathbf{x})} - \log p(\mathbf{y} | \mathbf{z}_k) \right]. \end{aligned}$$

Finally, the $\frac{1}{N}$ term can be dropped since it is a constant factor with respect to the network parameters, merely scaling the loss and thus ineffective in learning. To arrive at a block-local formulation

of the loss, we separate the generation of the forward and backward messages, and the computation of the local losses. Using this, we write Eq. (S4) in the form

$$\mathcal{F} = \sum_{k=1}^N \ell^{(\alpha)}(p_k, \beta_k | \alpha_{k-1}) - \mathbb{E}_{q_k} [\log p(\mathbf{y} | \mathbf{z}_k)] , \quad (\text{S4})$$

with local losses given by

$$\ell^{(\alpha)}(p_k, \beta_k | \alpha_{k-1}) = \mathcal{D}_{KL}(g_k(\alpha_{k-1}, \beta_k) | f_k(\alpha_{k-1})) = \mathcal{D}_{KL}(q_k | \alpha_k) ,$$

where we defined the mapping $f_k(\alpha_{k-1}) = \mathbb{E}_{\alpha_{k-1}} [p_k(\mathbf{z}_k | \mathbf{z}_{k-1})] = \alpha_k$ and $g_k(\alpha_{k-1}, \beta_k) = \mathbb{E}_{\alpha_{k-1}} [\frac{1}{\mathcal{Z}} p_k(\mathbf{z}_k | \mathbf{z}_{k-1}) \beta(\mathbf{z}_k, \mathbf{y})] = \mathbb{E}_{\alpha_{k-1}} [q(\mathbf{z}_k | \mathbf{z}_{k-1}, \mathbf{y})] = q_k$, with normalization \mathcal{Z} . Eq. (S4) is an upper bound on the log-likelihood loss $\mathcal{L} = -\log p(\mathbf{y} | \mathbf{x}) \leq \mathcal{F}$. Since \mathcal{L} is strictly positive, minimizing \mathcal{F} to zeros implies that also \mathcal{L} becomes zero (Jordan et al., 1999). We can also add any positive auxiliary loss $\tilde{\ell}(p_k, \beta_k | \alpha_{k-1}) \geq 0$ to $\ell^{(\alpha)}$, which results in a new upper bound for \mathcal{L} . Therefore, in Eq. (S4) we can also use the augmented loss

$$\ell(p_k, \beta_k | \alpha_{k-1}) = \mathcal{D}_{KL}(q_k | \alpha_k) + \tilde{\ell}(p_k, \beta_k | \alpha_{k-1}) , \quad (\text{S5})$$

instead of $\ell^{(\alpha)}(p_k, \beta_k | \alpha_{k-1})$ directly. Importantly, all terms of the loss ℓ can be computed locally through the forward propagation of the k -th block to realize f_k and the computation of the posterior g_k .

The variational posterior q is given by Eq. (4). Alternatively we can also use a multi-layer feedback network, that propagates messages β_k backward

$$\begin{aligned} \beta_k(\mathbf{z}_k) &= p(\mathbf{y} | \mathbf{z}_k) = \mathbb{E}_{\mathbf{z}_{k+1}} [p_k(\mathbf{y} | \mathbf{z}_{k+1}) p_k(\mathbf{z}_{k+1} | \mathbf{z}_k)] \\ &= \mathbb{E}_{\mathbf{z}_{k+1}} [\beta_{k+1}(\mathbf{z}_{k+1}) p_k(\mathbf{z}_{k+1} | \mathbf{z}_k)] . \end{aligned} \quad (\text{S6})$$

This is the method that was using in Section 4.2. This method re-introduces a locking problem but may work better in some scenarios where more complex feedback messages are required. Also here the feedback cannot be computed in closed form so we resort to a gradient-based method.

S1.2.1 ESTIMATING THE LOG-LIKELIHOOD LOSS THROUGH POSTERIOR BOOTSTRAPPING

Next, we show how the remaining term $\mathbb{E}_{q_k} [\log p(\mathbf{y} | \mathbf{z}_k)]$ in Eq. (S4) can be estimated locally. The intuition behind this result is that $-\mathbb{E}_{q_k} [\log p(\mathbf{y} | \mathbf{z}_k)]$ is of a similar form as the log-likelihood loss (Eq. (1) of the main text), i.e., the likelihood of the data labels \mathbf{y} of the residual network $\mathbf{z}_k \rightarrow \mathbf{y}$. Thus, treating \mathbf{z}_k as block-local input data and minimizing the augmented ELBO loss from layer $\mathbf{z}_k \rightarrow \mathbf{y}$ minimizes another upper bound on the global loss \mathcal{L} . To formalize this observation we introduce the recursive short-hand notation $q_{k \rightarrow j} = \mathbb{E}_{q_{k \rightarrow (j-1)}} [q(\mathbf{z}_j | \mathbf{z}_{j-1}, \mathbf{y})]$, with $q_{k \rightarrow k} = q_k$. Using this, we find the following chain of inequalities

$$\begin{aligned} & -\mathbb{E}_{q_k} [\log p(\mathbf{y} | \mathbf{z}_k)] \\ & \leq -\mathbb{E}_{q_k} [\log p(\mathbf{y} | \mathbf{z}_k)] + \mathcal{D}_{KL}(\mathbb{E}_{q_k} [q(\mathbf{z}_{k+1} | \mathbf{z}_k, \mathbf{y})] | \mathbb{E}_{q_k} [p(\mathbf{z}_{k+1} | \mathbf{z}_k, \mathbf{y})]) \end{aligned} \quad (\text{S7})$$

$$\leq \mathbb{E}_{q_{k \rightarrow (k+1)}} [\log \mathbb{E}_{q_k} [q(\mathbf{z}_{k+1} | \mathbf{z}_k, \mathbf{y})] - \mathbb{E}_{q_k} [\log p(\mathbf{z}_{k+1} | \mathbf{z}_k, \mathbf{y}) + \log p(\mathbf{y} | \mathbf{z}_k)]] \quad (\text{S8})$$

$$= \mathbb{E}_{q_{k \rightarrow (k+1)}} [\log \mathbb{E}_{q_k} [q(\mathbf{z}_{k+1} | \mathbf{z}_k, \mathbf{y})] - \mathbb{E}_{q_k} [\log p_{k+1}(\mathbf{z}_{k+1} | \mathbf{z}_k)]] - \mathbb{E}_{q_{k \rightarrow (k+1)}} [\log p(\mathbf{y} | \mathbf{z}_{k+1})]$$

$$\begin{aligned} & \leq \mathcal{D}_{KL}(\mathbb{E}_{q_k} [q(\mathbf{z}_{k+1} | \mathbf{z}_k, \mathbf{y})] | \mathbb{E}_{q_k} [p_{k+1}(\mathbf{z}_{k+1} | \mathbf{z}_k)]) \\ & \quad - \mathbb{E}_{q_k, q_{k \rightarrow (k+1)}} [\log p_{k+1}(\mathbf{z}_{k+1} | \mathbf{z}_k)] - \mathbb{E}_{q_{k \rightarrow (k+1)}} [\log p(\mathbf{y} | \mathbf{z}_{k+1})] \end{aligned} \quad (\text{S9})$$

$$= \ell_{k, k+1}^{(\rho)} - \mathbb{E}_{q_{k \rightarrow (k+1)}} [\log p(\mathbf{y} | \mathbf{z}_{k+1})] , \quad (\text{S10})$$

with local loss

$$\begin{aligned} \ell_{k, l}^{(\rho)} &= \mathcal{D}_{KL}(\mathbb{E}_{q_{k \rightarrow l}} [q(\mathbf{z}_{l+1} | \mathbf{z}_l, \mathbf{y})] | \mathbb{E}_{q_{k \rightarrow l}} [p_{l+1}(\mathbf{z}_{l+1} | \mathbf{z}_l)]) - \mathcal{H}(p_{l+1} | q_{k \rightarrow l}) \\ & \quad \text{with } \mathcal{H}(p_{l+1} | q_{k \rightarrow l}) = \mathbb{E}_{q_{k \rightarrow l}, q_{k \rightarrow (l+1)}} [\log p_{l+1}(\mathbf{z}_{l+1} | \mathbf{z}_l)] \end{aligned} \quad (\text{S11})$$

and where $\mathbb{E}_{q_{k \rightarrow l}} [\cdot]$ denotes expectation with respect to $q_{k \rightarrow l}$. In (S8) we used Jensen's inequality, $-\mathbb{E} [\log p(X)] \geq -\log \mathbb{E} [p(X)]$, and in (S9) we used $-\log \mathbb{E}_{q_k} [\alpha_{k+1}(\mathbf{z}_k)] \geq 0$, i.e., the negative log expectation of a probability distribution is always positive or zero.

Next we generalize the inequality (S9) in the following theorem:

Theorem S1.1. Let $p(\mathbf{y}, \mathbf{z} | \mathbf{x})$ be a probabilistic model subject to the conditional independence properties over $N + 1$ blocks, given by Eq. (S1). Let $\ell_k^{(\rho)}$ be the local loss function Eq.(S11). Furthermore, let $\ell_k^{(\omega)} = -\mathbb{E}_{q_{k \rightarrow N}}[\log p(\mathbf{y} | \mathbf{z}_N)]$. Then the log-likelihood loss Eq. (1) is bounded from above by the sum of losses $\mathcal{F}_N \geq \mathcal{L}$

$$\mathcal{F}_N = \sum_{k=1}^N \ell_k^{(\alpha)} + \ell_k^{(\omega)} + \sum_{l=k+1}^N \ell_{k,l}^{(\rho)} \quad (\text{S12})$$

Proof. We prove Theorem 1 by induction over block i . Let $\mathcal{L} = -\log p(\mathbf{y} | \mathbf{x})$, and $\mathcal{F}_1 = \mathcal{F}$ as in Eq. (S2). We show a transition $\mathcal{F}_{i-1} \rightarrow \mathcal{F}_i$, with $\mathcal{F}_{i-1} \leq \mathcal{F}_i$, and where \mathcal{F}_N recovers Eq. (S12). This result implies a hierarchy of loss functions $0 \leq \mathcal{L} \leq \mathcal{F}_1 \leq \mathcal{F}_2 \leq \dots \leq \mathcal{F}_N$, and $\mathcal{L} \leq \mathcal{F}_N$ follows, which completes the proof.

We define:

$$\mathcal{F}_i = \sum_{k=1}^N \ell_k^{(\alpha)} + \sum_{l=k+1}^i \ell_{k,l}^{(\rho)} - \mathbb{E}_{q_{k \rightarrow j}}[\log p(\mathbf{y} | \mathbf{z}_j)] \Big|_{j=\max(i,k)} \quad (\text{S13})$$

For $i = 1$ we recover Eq. (S4) and thus $\mathcal{L} \leq \mathcal{F}_1$ holds as established before. Using the result (S9) we can take the inductive step $\mathcal{L}_{i-1} \rightarrow \mathcal{L}_i$

$$\begin{aligned} \mathcal{F}_{i-1} &= \sum_{k=1}^N \ell_k^{(\alpha)} + \sum_{l=k+1}^{i-1} \ell_{k,l}^{(\rho)} - \mathbb{E}_{q_{k \rightarrow j}}[\log p(\mathbf{y} | \mathbf{z}_j)] \Big|_{j=\max(i-1,k)} \\ &= \sum_{k=1}^{i-1} \ell_k^{(\alpha)} + \sum_{l=k+1}^{i-1} \ell_{k,l}^{(\rho)} - \mathbb{E}_{q_{k \rightarrow (i-1)}}[\log p(\mathbf{y} | \mathbf{z}_{i-1})] + \sum_{k'=i}^N \ell_{k'}^{(\alpha)} - \mathbb{E}_{q_{k'}}[\log p(\mathbf{y} | \mathbf{z}_{k'})] \\ &\leq \sum_{k=1}^{i-1} \ell_k^{(\alpha)} + \sum_{l=k+1}^{i-1} \ell_{k,l}^{(\rho)} + \ell_{k,i}^{(\rho)} - \mathbb{E}_{q_{k \rightarrow i}}[\log p(\mathbf{y} | \mathbf{z}_i)] \\ &\quad + \ell_i^{(\alpha)} + \ell_{i,i+1}^{(\rho)} - \mathbb{E}_{q_{i \rightarrow (i+1)}}[\log p(\mathbf{y} | \mathbf{z}_{i+1})] + \sum_{k'=i+1}^N \ell_{k'}^{(\alpha)} - \mathbb{E}_{q_{k'}}[\log p(\mathbf{y} | \mathbf{z}_{k'})] \\ &= \sum_{k=1}^N \ell_k^{(\alpha)} + \sum_{l=k+1}^i \ell_{k,l}^{(\rho)} - \mathbb{E}_{q_{k \rightarrow j}}[\log p(\mathbf{y} | \mathbf{z}_j)] \Big|_{j=\max(i,k)} = \mathcal{F}_i. \end{aligned} \quad (\text{S14})$$

This shows that the global loss can be decomposed into a sum of local losses. Setting $i = N$ in Eq. (S13), the proof of Theorem S1.1 follows. \square

What remains to be shown is that posterior bootstrapping allows us to compute the required terms. To arrive at this result we further study the local loss (S11). Note, that this expression can be written in the form (S5) as a function of the forward network transfer $p(\mathbf{z}_{l+1} | \mathbf{z}_l)$ given the distribution $q_{k \rightarrow l}(\mathbf{z}_l)$ for $1 \leq k \leq l \leq N$, i.e.

$$\begin{aligned} \ell_{k,l}^{(\rho)} &= \ell(p_{l+1}, \beta_{l+1} | q_{k \rightarrow l}) = \\ &= \mathcal{D}_{KL}(g_k(q_{k \rightarrow l}, \beta_{l+1}) | f_k(q_{k \rightarrow l})) - \mathcal{H}(p_{l+1} | q_{k \rightarrow l}), \end{aligned} \quad (\text{S15})$$

where we used the mappings f_k and g_k as in Eq. (S5). Thus by choosing $\hat{\ell}(p_{l+1}, \beta_{l+1} | q_{k \rightarrow l}) = \mathcal{H}(p_{l+1} | q_{k \rightarrow l})$, we can spell out the local loss in the exact same way as Eq. (S5), but passing the posterior messages $q_{k \rightarrow l}$ instead of the forward pass α_k . The last block with index $N + 1$ directly optimizes $\ell_k^{(\omega)} = \mathbb{E}_{q_{k \rightarrow N}}[\log p(\mathbf{y} | \mathbf{z}_N)]$, which is also local to that block.

We thus propose to realize the sum over k, l in (S14) using a posterior bootstrapping schedule. Instead of passing only the forward messages, blocks may be selected to compute the variational posterior distribution q locally and pass that to the next block instead. The optimal posterior bootstrapping schedule, according to (S14), is the one that computes all N^2 combinations of passing α and q messages giving rise to the structure in Fig. 2. We are now ready to prove Theorem 1, which we reverbate here for completeness

Theorem S1.2. *Let $p(\mathbf{y}, \mathbf{z} | \mathbf{x})$ be a probabilistic model subject to the conditional independence properties over $N + 1$ blocks, given by Eq. (S1). Let ℓ be the local loss function Eq.(S5). Furthermore, let $\alpha_k^{(m)}, \beta_k^{(m)}$ be messages created using the optimal posterior bootstrapping schedule outlined above. Then, the simultaneous minimization of $\ell(p_l^{(m)}, \beta_l^{(m)} | \alpha_k^{(m)})$ in all blocks k , minimizes an upper bound on the log-likelihood loss \mathcal{L} .*

Proof. The proof of Theorem S1.2 follows directly from Theorem S1.1 by substituting the loss terms in the sums with (S15). Importantly, all messages generated by the bootstrapping schedule are treated the same, so there is no contrastive step or need for global information to signal a network-wide learning phase. In simulations we also experimented with different bootstrapping schedules, other than the optimal one. \square

S1.3 RELATIONSHIP TO EM AND CONVERGENCE PROPERTIES

As outlined above the model can be closely linked to the EM algorithm. The split of gradient estimators using the Markov assumption is a key property of algorithms derived from EM, and also the key property exploited in BLL. EM makes use of the identity (Dempster et al., 1977)

$$\begin{aligned} -\nabla \mathcal{L} &= \nabla \log p(\mathbf{y} | \mathbf{x}) = \frac{1}{p(\mathbf{y} | \mathbf{x})} \nabla p(\mathbf{y} | \mathbf{x}) \\ &= \frac{1}{p(\mathbf{y} | \mathbf{x})} \nabla \mathbb{E}_{\mathbf{z}_k} [p(\mathbf{y} | \mathbf{z}_k) p(\mathbf{z}_k | \mathbf{x})] = \mathbb{E}_{p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})} [\nabla \log p(\mathbf{y} | \mathbf{z}_k) + \nabla \log p(\mathbf{z}_k | \mathbf{x})], \end{aligned}$$

where in the last step we used that $p(\mathbf{y} | \mathbf{x})$ is constant under the expectation and $\frac{p(\mathbf{y} | \mathbf{z}_k) p(\mathbf{z}_k | \mathbf{x})}{p(\mathbf{y} | \mathbf{x})} = p(\mathbf{z}_k | \mathbf{x}, \mathbf{y})$ (Bayes' rule).

We use a variational approach where the posterior is replaced by q . It has been established in prior work that, similar to the EM algorithm, the variational loss \mathcal{L} can be minimized by alternating two optimization steps (Jordan et al., 1999; Neal and Hinton, 1998)

$$\text{E-step: } q^{(t)} = \arg \min_q \mathcal{F}(q, \boldsymbol{\theta}^{(t-1)}) \quad (\text{S16})$$

$$\text{M-step: } \boldsymbol{\theta}^{(t)} = \arg \min_{\boldsymbol{\theta}} \mathcal{F}(q^{(t)}, \boldsymbol{\theta}). \quad (\text{S17})$$

In Neal and Hinton (1998) it was shown that this approach also works if gradient descent is used for the optimization of (some of) the parameters. We use here the variant where parameters of the forward network are optimized via gradient descent whereas the loss with respect to the feedback network parameters is directly optimized.

S1.4 GENERAL EXPONENTIAL FAMILY DISTRIBUTION

To arrive at a result for the gradient of the first (KL-divergence) term ℓ_k in Eq. (S4) we seek distributions for which the marginals can be computed in closed form. We assume forward messages α and posterior ρ be given by general exponential family distributions

$$\alpha_k(\mathbf{z}_k) = \prod_j \alpha_{kj}(z_{kj}) = \prod_j h(z_{kj}) \exp(T(z_{kj}) \phi_{kj} - A(\phi_{kj})) \quad (\text{S18})$$

$$\rho_k(\mathbf{z}_k) = \prod_j \rho_{kj}(z_{kj}) = \prod_j h(z_{kj}) \exp(T(z_{kj}) \gamma_{kj} - A(\gamma_{kj})) \quad (\text{S19})$$

with base measure h , sufficient statistics T , log-partition function A , and natural parameters ϕ_{kj} and γ_{kj} . Using this the KL loss becomes

$$\mathcal{D}_{KL}(\rho_k | \alpha_k) = \sum_j \mathbb{E}_{\rho_{kj}} [T(z_{kj}) (\phi_{kj} - \gamma_{kj}) - A(\phi_{kj}) + A(\gamma_{kj})], \quad (\text{S20})$$

and thus

$$-\nabla \mathcal{D}_{KL}(\rho_k | \alpha_k) = \sum_j \left(\mathbb{E}_{\rho_{kj}} [T(z_{kj})] - \mathbb{E}_{\alpha_{kj}} [T(z_{kj})] \right) \nabla \phi_{kj} + \underbrace{\left(\mathbb{E}_{\rho_{kj}} [T(z_{kj})^2] - \mathbb{E}_{\rho_{kj}} [T(z_{kj})]^2 \right)}_{\sigma^2(\rho_{kj})} (\phi_{kj} - \gamma_{kj}) \nabla \gamma_{kj}, \quad (\text{S21})$$

which by defining $\mu(p) = \mathbb{E}_p [T(z_{kj})]$ can be written in the compact form

$$-\nabla \mathcal{D}_{KL}(\rho_k | \alpha_k) = \sum_j (\mu(\rho_{kj}) - \mu(\alpha_{kj})) \nabla \phi_{kj} + \sigma^2(\rho_{kj}) (\phi_{kj} - \gamma_{kj}) \nabla \gamma_{kj}.$$

This is the result Eq. (6) of the main text.

S1.4.1 GAUSSIAN RANDOM VARIABLES WITH KNOWN VARIANCE

Throughout the numerical simulations we use the network to represent Gaussian distributions with known variance. For this distribution we have $T(z_{kj}) = z_{kj}$, $\mathbb{E}_{\rho_{kj}} [T(z_{kj})] = \phi_{kj}$, and furthermore $\sigma^2(\rho_{kj}) = \sigma^2 (= \text{const})$. We get

$$-\nabla \ell_k = \sum_j (\gamma_{kj} - \phi_{kj}) \nabla \phi_{kj} + \sigma (\phi_{kj} - \gamma_{kj}) \nabla \gamma_{kj} \quad (\text{S22})$$

Using the parameterization $\phi_{kj} = a_{kj}$ and $\gamma_{kj} = \frac{1}{2}(a_{kj} + b_{kj})$, we further get

$$-\nabla_a \ell_k = \left(\frac{\sigma}{2} - 1 \right) \sum_j (a_{kj} - b_{kj}) \nabla a_{kj}. \quad (\text{S23})$$

This is the KL loss that was used to minimize the distance between forward and feedback features.

S1.5 CLOSED FORM SOLUTION OF BACKWARD NETWORK

Here we show the closed form solution for optimizing the backward network. Over a set of M training samples we seek to solve

$$\theta_k^{(b)} \stackrel{!}{=} \arg \min_{\theta_k^{(b)}} \sum_{m=1}^M \ell_k(\rho_k^{(m)}, \alpha_k^{(m)}) = \arg \min_{\theta_k^{(b)}} \sum_{m=1}^M \mathcal{D}_{KL}(\rho_k^{(m)} | \alpha_k^{(m)}) + \mathcal{H}(\rho_k^{(m)}, \alpha_k^{(m)}).$$

As in the remainder of this paper, we treat the inputs to the block k as constants. The second cross-entropy term only depends on the parameters of the forward network. Taking the gradient with respect to backward network parameters γ_{kj} thus yields

$$\begin{aligned} \nabla_{\gamma_{kj}} \ell_k &= \nabla_{\gamma_{kj}} \mathcal{D}_{KL}(\rho_k | \alpha_k) \\ &= \sum_{m=1}^M \nabla \mu(\gamma_{kj}^{(m)}) (\gamma_{kj}^{(m)} - \phi_{kj}^{(m)}) \nabla \gamma_{kj}^{(m)} + \mu(\gamma_{kj}^{(m)}) \nabla \gamma_{kj}^{(m)} - \nabla A(\gamma_{kj}^{(m)}) \nabla \gamma_{kj}^{(m)} \stackrel{!}{=} 0 \\ \Leftrightarrow \sum_{m=1}^M \nabla \mu(\gamma_{kj}^{(m)}) (\gamma_{kj}^{(m)} - \phi_{kj}^{(m)}) + \mu(\gamma_{kj}^{(m)}) - \nabla A(\gamma_{kj}^{(m)}) &\stackrel{!}{=} 0 \end{aligned}$$

Assuming Gaussian with known variance $\mu(\gamma_{kj}^{(m)}) = \sigma \gamma_{kj}^{(m)}$, $\nabla \mu(\gamma_{kj}^{(m)}) = \sigma$, $A(\gamma_{kj}^{(m)}) = \frac{(\gamma_{kj}^{(m)})^2}{2}$ and $\nabla A(\gamma_{kj}^{(m)}) = \gamma_{kj}^{(m)}$ gradient with respect to γ_{kj}

$$\begin{aligned} \Leftrightarrow \sum_{m=1}^M \nabla \mu(\gamma_{kj}^{(m)}) (\gamma_{kj}^{(m)} - \phi_{kj}^{(m)}) + \mu(\gamma_{kj}^{(m)}) - \nabla A(\gamma_{kj}^{(m)}) &\stackrel{!}{=} 0 \\ \Leftrightarrow \sum_{m=1}^M \sigma (\gamma_{kj}^{(m)} - \phi_{kj}^{(m)}) + \sigma \gamma_{kj}^{(m)} - \gamma_{kj}^{(m)} &\stackrel{!}{=} 0 \\ \Leftrightarrow \sum_{m=1}^M \gamma_{kj}^{(m)} (2\sigma - 1) - \sigma \phi_{kj}^{(m)} &\stackrel{!}{=} 0 \end{aligned}$$

Hyperparameter	Value
weight of KL loss	0.70
weight of entropy loss \mathcal{H}	0.014
magnitude of added noise to estimate \mathcal{H}	0.013
weight of correlation loss	0.70
predictive loss scaling	0.1
weight of output CE loss	0.49
posterior bootstrapping	optimal
batch size	256

Table S1: Hyperparameters used for training ResNet-50 on FashionMNIST task.

$$\begin{aligned}
\gamma_{kj}^{(m)} &= \frac{1}{2} (a_{kj}^{(m)} + b_{kj}) \\
\leftrightarrow \sum_{m=1}^M \frac{1}{2} (a_{kj}^{(m)} + b_{kj}) (2\sigma - 1) - \sigma a_{kj}^{(m)} &\stackrel{!}{=} 0 \\
\leftrightarrow \frac{M}{2} (2\sigma - 1) b_{kj} - \frac{1}{2} \sum_{m=1}^M a_{kj}^{(m)} &\stackrel{!}{=} 0 \\
\leftrightarrow b_{kj} &\stackrel{!}{=} \frac{1}{M(2\sigma - 1)} \sum_{m=1}^M a_{kj}^{(m)} = \frac{c_1}{M} \sum_{m=1}^M a_{kj}^{(m)},
\end{aligned}$$

with constant c_1 . The optimal parameters for the backward network is thus given by the class-specific mean over the forward messages.

S2 NUMERICAL SIMULATIONS

We assessed the models results variability over five runs for each model and each task, using different random seeds. We used 5 different losses derived from our theoretical framework or previously established: the BLL *KL loss*, the *entropy loss* \mathcal{H} , the *prediction loss* as in (Nøkland and Eidnes, 2019), a *correlation loss* that punishes high auto-correlation of features within a batch and the output cross-entropy (CE) loss, that is only effective at the last block. Each loss was assigned a weight to scale it relative to the other losses and then combined to block-local losses that were optimized individually.

S2.0.1 FASHIONMNIST CLASSIFICATION TASK

FashionMNIST is a freely available dataset consisting of 60k training grayscale images and 10k grayscale test images of fashion items published under the MIT License (MIT) (Xiao et al., 2017). The images were normalized to have mean 0 and stds 1 and augmented with random horizontal flips during training. The BLL networks for FashionMNIST experiments used the same ResNet architectures but augmented with the feedback blocks. For the forward network we used the Adam optimizer with a learning rate of 0.03 without weight decay, a Cosine annealing learning rate (LR) scheduler (Loshchilov and Hutter, 2017) with max iterations set to 140. We used the direct closed form optimization for the feedback network on every batch but applied it with a rate of only 0.9 to account for missing classes. The remaining hyperparameters used are given in Table S1.

S2.0.2 CIFAR10 CLASSIFICATION TASK

CIFAR10 is a freely available dataset consisting of 50k training images and 10k test images from (Krizhevsky, 2009). We used the same data augmentation, optimizers and hyperparameters used for FashionMNIST to train CIFAR10 (see Table S1).

		test-1 (mean±std)	test-3 (mean±std)
ResNet-18	BLL	94±0.2	99.3±0.04
	BP	92.7±0.1	99.2±0.7
	FA	88.2±0.3	98.7±0.2
	Pred-Sim	93.7 ±0.2	99.4±0.1
ResNet-50	BLL	94.1±0.24	99.1±0.1
	BP	93.4±0.6	99.4±0.05
	FA	86.6±0.7	98.6±0.1
	Pred-Sim	94.2 ±0.2	99.4±0.08

Table S2: As in Table 2. Classification accuracy (% correct) for 5 runs on FashionMNIST vision tasks.

		test-1 (mean±std)	test-3 (mean±std)
ResNet-18	BLL	88.1±0.2	97.9±0.1
	BP	92.5±1.5	98.3±0.3
	FA	72.0±0.6	92.8±0.1
	Pred-Sim	87.8 ±0.2	97.9±0.1
ResNet-50	BLL	92.3±0.2	98.9±0.1
	BP	91.1±1.1	98.7±0.2
	FA	62.5±0.4	88.2±0.2
	Pred-Sim	92.1 ±0.2	98.8±0.1

Table S3: As in Table 2. Classification accuracy (% correct) for 5 runs on CIFAR10 task.

S2.1 IMAGENET CLASSIFICATION TASK

We train ResNet-50 on Imagenet-1K using FFCV (Leclerc et al., 2022) library. Standard hyperparameters of ImageNet were not changed from the FFCV baseline results. We reused most of the hyperparameters specific to our method from the CIFAR-10 task. See Table S4 for full details on the hyperparameters used for Imagenet-1K.

S2.2 ABLATION STUDY

We performed ablation studies to assess the importance of the different losses in BLL: Correlation loss, predictive loss and KL loss. To this end, we disabled one loss at a time and trained on CIFAR-10. In addition we also disabled all local losses, effectively only training the last block directly at the

Hyperparameter	Value
weight of KL loss	0.25
weight of entropy loss \mathcal{H}	1.0
magnitude of added noise to estimate \mathcal{H}	0.01
weight of correlation loss	0.7
feedback network LR	0.9
weight of output CE loss	0.43
predictive loss scaling	0.5
posterior bootstrapping	disabled
batch size	512
feedback optimizer momentum	0.01

Table S4: Hyperparameters used for training ResNet-50 on ImageNet task. All other hyperparameters relating to the forward network training are not modified from the baseline FFCV training script

modification	performance
w/o correlation loss	91.4±0.2
w/o predictive loss	91.5±0.2
w/o KL loss	91.8±0.3
w/o all local losses	52.1±13.3
with simplified bootstrapping	92.4±0.1
benchmark BLL	92.3±0.2

Table S5: BLL CIFAR-10 test accuracy with modified algorithm

targets. Furthermore, we assess the importance of the optimal posterior bootstrapping by using only the simplified bootstrapping schedule where only forward messages are propagated while keeping all losses enabled. The results are presented in Table S5.

A decrease in performance is observed whenever one of the local losses is removed, but removing them all drastically reduces the performance as expected. No local losses means training only the block layer while freezing the remaining blocks, thus the decrease in performance. Using a simplified bootstrapping scheme gives comparable performance as augmenting forward messages with backward messages. In this case the message augmentation doesn't provide sensible advantage. Exploring different augmentation methods on more tasks and architectures might give better insight.

We study the effect of splitting the network into blocks in Figure S1. The performance decreases as the number introduced in the network increases. This effect is more pronounced as the difficulty of the task increases. We compare this to the effect of adding additional losses to the training method without splitting the network. This network is trained end-to-end with backpropagation, these additional losses introduced slight performance degradation.

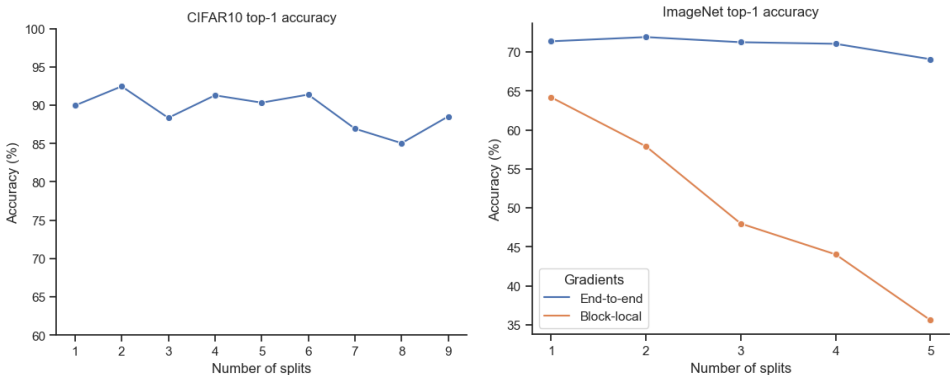


Figure S1: Top-1 classification accuracy on CIFAR-10 (Left) and ImageNet (Right) across number of splits in the ResNet-50. ImageNet performance at 30 epochs of training is compared to backpropagation training while keeping additional losses introduced in our method.

S2.3 HARDWARE AND SOFTWARE DETAILS

ResNet18 and ResNet50 models and experiments were implemented in PyTorch (Paszke et al., 2019). Most of our experiments were run on NVIDIA A100 GPUs and some initial evaluations and the MINST experiments were conducted on NVIDIA V100 and Quadro RTX 5000 GPUs. In total we used about 190,000 core hours for training and hyper-parameter searches.