
Decoupled Greedy Learning of CNNs

Eugene Belilovsky¹ Michael Eickenberg² Edouard Oyallon³

Abstract

A commonly cited inefficiency of neural network training by back-propagation is the *update locking* problem: each layer must wait for the signal to propagate through the full network before updating. Several alternatives that can alleviate this issue have been proposed. In this context, we consider a simpler, but more effective, substitute that uses minimal feedback, which we call *Decoupled Greedy Learning (DGL)*. It is based on a greedy relaxation of the joint training objective, recently shown to be effective in the context of Convolutional Neural Networks (CNNs) on large-scale image classification. We consider an optimization of this objective that permits us to decouple the layer training, allowing for layers or modules in networks to be trained with a potentially linear parallelization in layers. With the use of a replay buffer we show this approach can be extended to asynchronous settings, where modules can operate with possibly large communication delays. We show theoretically and empirically that this approach converges. Then, we empirically find that it can lead to better generalization than sequential greedy optimization. We demonstrate the effectiveness of DGL against alternative approaches on the CIFAR-10 dataset and on the large-scale ImageNet dataset.

1. Introduction

Jointly training all layers using back-propagation is the standard method for learning neural networks, including the computationally intensive Convolutional Neural Networks (CNNs). Due to the sequential nature of gradient processing, standard back-propagation has several well-known ineffi-

ciencies that prohibit parallelization of the computations of the different constituent modules. Jaderberg et al. (2017) characterize these in order of severity as the *forward*, *update*, and *backward locking* problems. Backward *unlocking* would permit updates of all modules once forward signals have propagated to all subsequent modules, update *unlocking* would permit updates of a module *before* a signal has reached all subsequent modules, and forward *unlocking* would permit a module to operate asynchronously from its predecessor and dependent modules.

Methods addressing backward locking to a certain degree have been proposed in (Huo et al., 2018b;a; Choromanska et al., 2018; Nø kland, 2016). However, update locking is a far more severe inefficiency. Thus Jaderberg et al. (2017) and Czarnecki et al. (2017) propose and analyze Decoupled Neural Interfaces (DNI), a method that uses an auxiliary network to predict the gradient of the backward pass directly from the input. This method unfortunately does not scale well computationally or in terms of accuracy, especially in the case of CNNs (Huo et al., 2018a;b). Indeed, auxiliary networks must predict a weight gradient, usually high dimensional for larger models and input sizes.

A major obstacle to update unlocking is the heavy reliance on the upper modules for feedback. Several works have recently revisited the classic (Ivakhnenko & Lapa, 1965; Bengio et al., 2007) approach of supervised greedy layerwise training of neural networks (Huang et al., 2018a; Marquez et al., 2018). In Belilovsky et al. (2019) it is shown that such an approach, which relaxes the joint learning objective, and does not require global feedback, can lead to high-performance deep CNNs on large-scale datasets. We will show that the greedy learning objective used in these papers can be solved with an alternative optimization algorithm, which permits decoupling the computations and achieves update unlocking. It can be augmented with replay buffers (Lin, 1992) to permit forward unlocking which is a challenge not effectively addressed by any of the prior work. This simpler strategy can be shown to be a superior baseline for parallelizing the training across modules of a neural network.

The paper is structured as follows. In Sec. 2 we propose an optimization procedure for a decoupled greedy learning objective that achieves *update unlocking* and then extend it to

¹MILA ²Center for Computational Mathematics, Flatiron Institute ³CNRS, LIP6. Correspondence to: Eugene Belilovsky <eugene.belilovsky@mila.quebec>, Michael Eickenberg <m-eickenberg@flatironinstitute.org>.

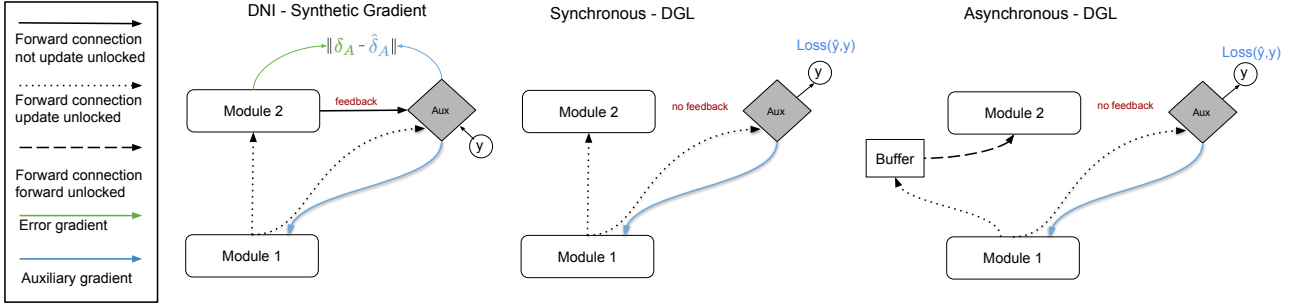


Figure 1. Comparison of DNI, Synchronous and Asynchronous DGL. Note in DGL subsequent modules do not provide feedback thus removing dependencies of the auxiliary network. Asynchronous DGL allows achieving forward unlocking.

an asynchronous setting (async-DGL) using a replay buffer, addressing *forward unlocking*. In Sec. 3 we show that the proposed optimization procedure converges and recovers standard rates of non-convex optimization, motivating empirical observations in the subsequent experimental section. In Sec. 4 we show that DGL can outperform competing methods in terms of scalability to larger and deeper models and stability to optimization hyperparameters and overall parallelism, allowing it to be applied to large datasets such as ImageNet. We extensively study async-DGL and find that it is robust to significant delays. We also empirically study the impact of parallelized training on convergence. Code for experiments is included in the submission.

2. Parallel Decoupled Greedy Learning

In this section we formally define the greedy objective and parallel optimization which we study in both the synchronous and asynchronous setting. We mainly consider the online setting and assume a stream of samples or mini-batches denoted $\mathcal{S} \triangleq \{(x_0^t, y^t)\}_{t \leq T}$, run during T iterations.

2.1. Preliminaries

For comparison purposes, we briefly review the update unlocking approach from DNI (Jaderberg et al., 2017). There, each network module has an associated *auxiliary net* which, given the output activation of the module, predicts the gradient signal from subsequent modules: the module can thus perform an update while modules above are still forward processing. The DNI auxiliary model is trained by using true gradients provided by upper modules when they become available, requiring activation caching. This also means that the auxiliary module can become out of sync with the changing output activation distribution, often requiring slow learning rates. Due to this and the high dimensionality of the predicted gradient which scales with module size, this estimate is challenging. One may ask how well a method that entirely avoids the use of feedback from upper modules would fare given similarly-sized auxiliary networks. We

will show that adapting the objective in (Belilovsky et al., 2019; Bengio et al., 2007) can also allow for update unlock and a degree of forward unlocking, with better properties.

Algorithm 1: Synchronous DGL

Input: Stream $\mathcal{S} \triangleq \{(x_0^t, y^t)\}_{t \leq T}$ of samples or mini-batches.

- 1 **Initialize** Parameters $\{\theta_j, \gamma_j\}_{j \leq J}$.
- 2 **for** $(x_0^t, y^t) \in \mathcal{S}$ **do**
- 3 **for** $j \in 1, \dots, J$ **do**
- 4 $x_j^t \leftarrow f_{\theta_{j-1}}(x_{j-1}^t)$.
- 5 Compute $\nabla_{(\gamma_j, \theta_j)} \hat{\mathcal{L}}(y^t, x_j^t; \gamma_j, \theta_j)$.
- 6 $(\theta_j, \gamma_j) \leftarrow$ Update parameters (θ_j, γ_j) .
- 7 **end**
- 8 **end**

Algorithm 2: Asynchronous DGL with Replay

Input: Stream $\mathcal{S} \triangleq \{(x_0^t, y^t)\}_{t \leq T}$; Distribution of the delay $p = \{p(j)\}_j$; Buffer size M .

- 1 **Initialize:** Buffers $\{B_j\}_j$; params $\{\theta_j, \gamma_j\}_j$.
- 2 **while training do**
- 3 Sample j in $\{1, \dots, J\}$ following p .
- 4 **if** $j = 1$ **then**
- 5 $(x_0, y) \leftarrow \mathcal{S}$
- 6 **else**
- 7 $(x_{j-1}, y) \leftarrow B_{j-1}$.
- 8 **end**
- 9 $x_j \leftarrow f_{\theta_{j-1}}(x_{j-1})$.
- 10 Compute $\nabla_{(\gamma_j, \theta_j)} \hat{\mathcal{L}}(y, x_j; \gamma_j, \theta_j)$.
- 11 $(\theta_j, \gamma_j) \leftarrow$ Update parameters (θ_j, γ_j) .
- 12 **if** $j < J$ **then** $B_j \leftarrow (x^j, y)$.
- 13 **end**

2.2. Optimization for Greedy Objective

Let \mathbf{X}_0 and Y be the data and labels, \mathbf{X}_j be the output representation for module j . We will denote the per-module objective function $\hat{\mathcal{L}}(\mathbf{X}_j, Y; \theta_j, \gamma_j)$, where the parameters θ_j correspond to the module parameter (i.e. $\mathbf{X}_{j+1} = f_{\theta_j}(\mathbf{X}_j)$). Here γ_j represents parameters of a auxiliary networks used to predict the final target and compute the local objective. $\hat{\mathcal{L}}$ in our case will be the empirical risk with a cross-entropy loss. The greedy training objective is thus given recursively

by defining P_j :

$$\min_{\theta_j, \gamma_j} \hat{\mathcal{L}}(\mathbf{X}_j, Y; \theta_j, \gamma_j), \quad (P_j)$$

where $\mathbf{X}_j = f_{\theta_{j-1}^*}(\mathbf{X}_{j-1})$ and θ_{j-1}^* is the minimizer of Problem (P_{j-1}) . A natural way to solve the optimization problem for J modules, (P_j) , is thus by sequentially solving the problems $\{P_j\}_{j \leq J}$ starting with $j = 1$. This is the approach taken in e.g. Marquez et al. (2018); Huang et al. (2018a); Bengio et al. (2007); Belilovsky et al. (2019). Here we consider an alternative procedure for optimizing the same objective, which we refer to as Sync-DGL. It is outlined in Alg 1. In Sync-DGL individual updates of each set of parameters are performed in parallel across the different layers. Each layer processes a sample or mini-batch, then passes it to the next layer, while simultaneously performing an update based on its own local loss. Note that at line 5 the subsequent layer can already begin computing line 4. Therefore, this algorithm achieves update unlocking. Once x_j^t has been computed, subsequent layers can begin processing. Sync-DGL can also be seen as a generalization of the biologically plausible learning method proposed in concurrent work (Nøkland & Eidnes, 2019). Appendix D also gives an explicit version of an equivalent multi-worker pseudo-code. Fig. 1 illustrates the decoupling compared to how samples are processed in the DNI algorithm.

In this work we solve the sub-problems P_j by backpropagation, but we note that any iterative solver available for P_j will be applicable (e.g. (Choromanska et al., 2018)). Finally we emphasize that unlike the sequential solvers of (e.g. Bengio et al. (2007); Belilovsky et al. (2019)) the distribution of inputs to each sub-problem solver changes over time, resulting in a learning dynamic whose properties have never been studied nor contrasted with sequential solvers.

2.3. Asynchronous DGL with Replay

We can now extend this framework to address *forward unlocking* (Jaderberg et al., 2017). DGL modules already do not depend on their successors for updates. We can further reduce dependency on the previous modules such that they can operate asynchronously. This is achieved via a replay buffer that is shared between adjacent modules, enabling them to reuse older samples. Scenarios with communication delays or substantial variations in speed between layers/modules benefit from this. We study one instance of such an algorithm that uses a replay buffer of size M , shown in Alg. 2 and illustrated in Fig. 1.

Our minimal distributed setting is as follows. Each worker j has a buffer that it writes to and that worker $j + 1$ can read from. The buffer uses a simple read/write protocol. A buffer B_j lets layer j write new samples. When it reaches capacity it overwrites the oldest sample. Layer $j + 1$ requests samples from the buffer B_j . They are selected by a last-in-

first-out (LIFO) rule, with precedence for the least reused samples. Alg. 2 simulates potential delays in such a setup by the use of a probability mass function (pmf) $p(j)$ over workers, analogous to typical asynchronous settings such as (Leblond et al., 2017). At each iteration, a layer is chosen at random according to $p(j)$ to perform a computation. In our experiments we limit ourselves to pmfs that are uniform over workers except for a single layer which is chosen to be selected less frequently on average. Even in the case of a uniform pmf, asynchronous behavior will naturally arise, requiring the reuse of samples. Alg. 2 permits a controlled simulation of processing speed discrepancies and will be used over settings of p and M to demonstrate that training and testing accuracy remain robust in practical regimes. Appendix D also provides pseudo-code for implementation in a parallel environment.

Unlike common data-parallel asynchronous algorithms (Zhang et al., 2015), the asynchronous DGL does not rely on a master node and requires only local communication similar to recent decentralized schemes (Lian et al., 2017). Contrary to decentralized SGD, DGL nodes only need to maintain and update the parameters of their local module, permitting much larger modules. Combining asynchronous DGL with distributed synchronous SGD for sub-problem optimization is a promising direction. For example it can alleviate a common issue of the popular distributed synchronous SGD in deep CNNs, which is the often limiting maximum batch size (Goyal et al., 2017).

2.4. Auxiliary and Primary Network Design

Like DNI our procedure relies on an auxiliary network to obtain update signal, both methods thus require auxiliary network design in addition to the main CNN architecture. Belilovsky et al. (2019) have shown that spatial averaging operations can be used to construct a scalable auxiliary network for the same objective as used in Sec 2.2. However, they did not directly consider the parallel training use case, where additional care must be taken in the design: The primary consideration is the relative speed of the auxiliary network with respect to its associated main network module. We will use primarily FLOP count in our analysis and aim to restrict our auxiliary networks to be 5% of the main network.

Although auxiliary network design might seem like an additional layer of complexity in CNN design and may require invoking slightly different architecture principles, this is not inherently prohibitive since architecture design is often related to training (e.g., the use of residuals is originally motivated by optimization issues inherent to end-to-end backprop (He et al., 2016)).

Finally, we note that although we focus on the distributed learning context, this algorithm and associated theory for greedy objectives is generic and has other potential appli-

cations. For example greedy objectives have recently been used in (Haarnoja et al., 2018; Huang et al., 2018a) and even with a single worker DGL reduces memory.

3. Theoretical Analysis

We now study the converge results of DGL. Since we do not rely on any approximated gradients, we can derive stronger properties than DNI (Czarnecki et al., 2017), such as a rate of convergence in our non-convex setting. To do so, we analyze Alg. 1 when the update steps are obtained from stochastic gradient methods. We show convergence guarantees (Bottou et al., 2018) under reasonable assumptions. In standard stochastic optimization schemes, the input distribution fed to a model is fixed (Bottou et al., 2018). In this work, the input distribution to each module is time-varying and dependent on the convergence of the previous module. At time step t , for simplicity we will denote all parameters of a module (including auxiliary) as $\Theta_j^t \triangleq (\theta_j^t, \gamma_j^t)$, and samples as $Z_j^t \triangleq (X_j^t, Y^t)$, which follow the density $p_j^t(z)$. For each auxiliary problem, we aim to prove the strongest existing guarantees (Bottou et al., 2018; Huo et al., 2018a) for the non-convex setting despite time-varying input distributions from prior modules. Proofs are given in the Appendix.

Let us fix a depth j , such that $j > 1$ and consider the converged density of the previous layer, $p_{j-1}^*(z)$. We define the following distance: $c_{j-1}^t \triangleq \int |p_{j-1}^t(z) - p_{j-1}^*(z)| dz$. Denoting ℓ the composition of the non-negative loss function and the network, we will study the expected risk $\mathcal{L}(\Theta_j) \triangleq \mathbb{E}_{p_{j-1}^*}[\ell(Z_{j-1}; \Theta_j)]$. We will now state several standard assumptions we use.

Assumption 1 (L -smoothness). \mathcal{L} is differentiable and its gradient is L -Lipschitz.

We consider the SGD scheme with learning rate $\{\eta_t\}_t$:

$$\Theta_j^{t+1} = \Theta_j^t - \eta_t \nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t), \quad (1)$$

where $Z_{j-1}^t \sim p_{j-1}^t$.

Assumption 2 (Robbins-Monro conditions). The step sizes satisfy $\sum_t \eta_t = \infty$ yet $\sum_t \eta_t^2 < \infty$.

Assumption 3 (Finite variance). There exists $G > 0$ such that $\forall t, \Theta_j, \mathbb{E}_{p_{j-1}^t}[\|\nabla_{\Theta_j} \ell(Z_{j-1}; \Theta_j)\|^2] \leq G$.

The Assumptions 1, 2 and 3 are standard (Bottou et al., 2018; Huo et al., 2018a), and we show in the following that our proof of convergence leads to similar rates, up to a multiplicative constant. The following assumption is specific to our setting where we consider a time-varying distribution:

Assumption 4 (Convergence of the previous layer). We assume that $\sum_t c_{j-1}^t < \infty$.

Lemma 3.1. Under Assumption 3 and 4, for all Θ_j , one has $\mathbb{E}_{p_{j-1}^*}[\|\nabla_{\Theta_j} \ell(Z_{j-1}; \Theta_j)\|^2] \leq G$.

We are now ready to prove the core statement for the convergence results in this setting:

Lemma 3.2. Under Assumptions 1, 3 and 4, we have:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\Theta_j^{t+1})] &\leq \mathbb{E}[\mathcal{L}(\Theta_j^t)] + \frac{LG}{2} \eta_t^2 \\ &\quad - \eta_t (\mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] - \sqrt{2G} c_{j-1}^t). \end{aligned}$$

The expectation is taken over each random variable. Also, note that without the temporal dependency (i.e. $c_j^t = 0$), this becomes analogous to Lemma 4.4 in (Bottou et al., 2018). Naturally it follows, that

Proposition 3.1. Under Assumptions 1, 2, 3 and 4, each term of the following equation converges:

$$\begin{aligned} \sum_{t=0}^T \eta_t \mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] &\leq \mathbb{E}[\mathcal{L}(\Theta_j^0)] \\ &\quad + G \sum_{t=0}^T \eta_t \left(\sqrt{2c_{j-1}^t} + \frac{L\eta_t}{2} \right). \end{aligned}$$

Thus the DGL scheme converges in the sense of (Bottou et al., 2018; Huo et al., 2018a). We can also obtain the following rate:

Corollary 3.1. The sequence of expected gradient norm accumulates around 0 at the following rate:

$$\inf_{t \leq T} \mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] \leq \mathcal{O} \left(\frac{\sum_{t=0}^T \sqrt{c_{j-1}^t} \eta_t}{\sum_{t=0}^T \eta_t} \right). \quad (2)$$

Thus compared to the sequential case, the parallel setting adds a delay that is controlled by $\sqrt{c_{j-1}^t}$.

4. Experiments

We conduct experiments that empirically show that DGL optimizes the greedy objective well, showing it is favorable against recent state-of-the-art proposals for decoupling training of deep network modules. We show that unlike previous decoupled proposals it can still work on a large-scale dataset (ImageNet) and that it can, in some cases, generalize better than standard back-propagation. We then extensively evaluate the asynchronous DGL, simulating large delays. For all experiments we use architectures taken from prior works and standard optimization settings.

4.1. Other Approaches and Auxiliary Network Designs

This section presents experiments evaluating DGL with the CIFAR-10 dataset (Krizhevsky, 2009) and standard data augmentation. We first use a setup that permits us to compare against the DNI method and which also highlights the

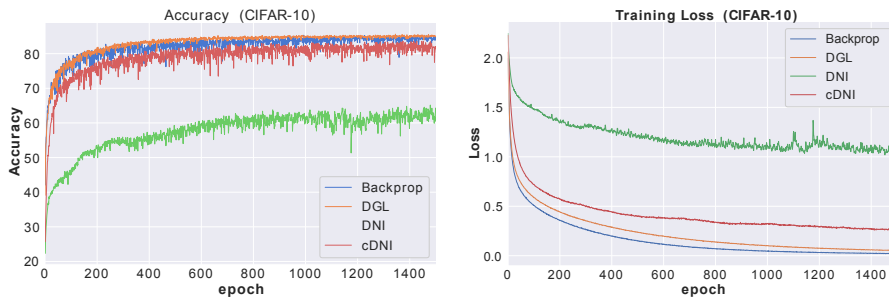


Figure 2. Comparison of DNI, cDNI, and DGL in terms of training loss and test accuracy for experiment from (Jaderberg et al., 2017). DGL converges better than cDNI and DNI with the same auxiliary net. and generalizes better than backprop.

generality and scalability of DGL. We then consider the design of a more efficient auxiliary network which will help to scale to the ImageNet dataset. We will also show that DGL is effective at optimizing the greedy objective compared to a naive sequential algorithm.

Comparison to DNI We reproduce the CIFAR-10 CNN experiment described in (Jaderberg et al., 2017), Appendix C.1. This experiment utilizes a 3 layer network with auxiliary networks of 2 hidden CNN layers. We compare our reproduction to the DGL approach. Instead of the final synthetic gradient prediction for the DGL we apply a final projection to the target prediction space. Here, we follow the prescribed optimization procedure from (Jaderberg et al., 2017), using Adam with a learning rate of 3×10^{-5} . We run training for 1500 epochs and compare standard backprop, DNI, context DNI (cDNI) (Jaderberg et al., 2017) and DGL. Results are shown in Fig. 2. Details are included in the Appendix. The DGL method outperforms DNI and the cDNI by a substantial amount both in test accuracy and training loss. Also in this setting, DGL can generalize better than standard backprop and obtains a close final training loss.

We also attempted DNI with the more commonly used optimization settings for CNNs (SGD with momentum and step decay), but found that DNI would diverge when larger learning rates were used, although DGL sub-problem optimization worked effectively with common CNN optimization strategies. We also note that the prescribed experiment uses a setting where the scalability of our method is not fully exploited. Each layer of the primary network of (Jaderberg et al., 2017) has a pooling operation, which permits the auxiliary network to be small for synthetic gradient prediction. This however severely restricts the architecture choices in the primary network to using a pooling operation at each layer. In DGL, we can apply the pooling operations in the auxiliary network, thus permitting the auxiliary network to be negligible in cost even for layers without pooling (whereas synthetic gradient predictions often have to be as costly as the base network). Overall, DGL is more scalable, accurate and robust to changes in optimization hyper-parameters than DNI.

	Relative FLOPS	Acc.
CNN-aux	200%	92.2
MLP-aux	0.7%	90.6
MLP-SR-aux	4.0%	91.2

Table 1. Comparison of auxiliary networks on CIFAR. CNN-aux applied in previous work is inefficient w.r.t. the primary module. We report flop count of the aux net relative to the largest module. MLP-aux and MLP-SR-aux applied after spatial averaging operations are far more effective with min. acc. loss.

Auxiliary Network Design We consider different auxiliary networks for CNNs. As a baseline we use convolutional auxiliary layers as in (Jaderberg et al., 2017) and (Belilovsky et al., 2019). For distributed training application this approach is sub-optimal as the auxiliary network can be substantial compared to the base network, leading to poorer parallelization gains. We note however that even in those cases (that we don’t study here) where the auxiliary network computation is potentially on the order of the primary network, it can still give advantages for parallelization for very deep networks and many available workers.

The primary network architecture we use for this study is a simple CNN similar to VGG family models (Simonyan & Zisserman, 2014) and those used in (Belilovsky et al., 2019). It consists of 6 convolutions of size 3×3 , batchnorm and shape preserving padding, with 2×2 maxpooling at layers 1 and 3. The width of the first layer is 128 and is doubled at each downsampling operation. The final layer does not have an auxiliary model– it is followed by a pooling and 2-hidden layer fully connected network, for all experiments. Two alternatives to the CNN auxiliary of (Belilovsky et al., 2019) are explored (Tab. 1).

The baseline auxiliary strategy based on (Belilovsky et al., 2019) and (Jaderberg et al., 2017) applies 2 CNN layers followed by a 2×2 averaging and projection, denoted as *CNN-aux*. First, we explore a direct application of the spatial averaging to 2×2 output shape (regardless of the resolution) followed by a 3-layer MLP (of constant width). This is denoted *MLP-aux* and drastically reduces the FLOP count with minimal accuracy loss compared to *CNN-aux*. Finally, we study a staged spatial resolution, first reducing the spatial resolution by $4 \times$ (and total size $16 \times$), then applying 3




Figure 3. Comparison of sequential and parallel training. Parallel catches up rapidly to sequential.

1×1 convolutions followed by a reduction to 2×2 and a 3 layer MLP, that we denote as *MLP-SR-aux*. These latter two strategies that leverage the spatial averaging produce auxiliary networks that are less than 5% of the FLOP count of the primary network even for large spatial resolutions as in real world image datasets. We will show that MLP-SR-aux is still effective even for the large-scale ImageNet dataset. We note that these more effective auxiliary models are not easily applicable in the case of DNI’s gradient prediction.

Sequential vs. Parallel Optimization of Greedy Objective

We briefly compare the sequential optimization of the greedy objective (Belilovsky et al., 2019; Bengio et al., 2007) to the DGL (Alg. 1). We use a 6 layer CIFAR-10 network with an MLP-SR-aux auxiliary model. In parallel we train the layers together for 50 epochs and in the sequential training we train each layer for 50 epochs before moving to the subsequent one. Thus the difference to DGL lies only in the input received at each layer (fully converged previous layer versus not fully converged previous layer). The rest of the optimization settings are identical. Fig. 3 shows comparisons of the learning curves for sequential training and DGL at layer 4 (layer 1 is the same for both as the input representation is not varying over the training period). DGL quickly catches up with the sequential training scheme and appears to sometimes generalize better. Like Oyallon (2017), we also visualize the dynamics of training per layer in Fig. 4, which demonstrates that after just a few epochs the individual layers build a dynamic of progressive improvement with depth.

Multi-Layer modules We have so far mainly considered the setting of layer-wise decoupling. This approach however can easily be applied to generic modules. Indeed, approaches such as DNI (Jaderberg et al., 2017) often consider decoupling entire multi-layer modules. Furthermore the propositions for backward unlocking (Huo et al., 2018b;a) also rely on and report they can often only decouple 100 layer networks into 2 or 4 blocks before observing optimization issues or performance losses and require that the number of parallel modules be much lower than the network




Figure 4. Per layer-loss on CIFAR: after few epochs, the layers build a dynamic of progressive improvement in depth.

Backprop	DDG	DGL
93.53	93.41	93.5 ± 0.1

Table 2. ResNet-110($K = 2$) for Backprop and DDG method from (Huo et al., 2018b). DGL is run for 3 trials to compute variance. They give the same acc. with DGL being update unlocked, DDG only backward unlocked. DNI is reported to not work in this setting (Huo et al., 2018b).

depth for the theoretical guarantees to hold. As in those cases, using multi-layer decoupled modules can improve performance and is natural in the case of deeper networks. We now use such a multi-layer approach to directly compare to the backward unlocking of (Huo et al., 2018b) and then subsequently we will apply this on deep networks for ImageNet. From here on we will denote K the number of total modules a network is split into.

Comparison to DDG Huo et al. (2018b) propose a solution to the backward locking (less efficient than solving update-locking, see discussion in Sec 5). We show that even in this situation the DGL method can provide a strong baseline for work on backward unlocking. We take the experimental setup from (Huo et al., 2018b), which considers a ResNet-110 parallelized into $K = 2$ blocks. We use the auxiliary network MLP-SR-aux which has less than 0.1% the FLOP count of the primary network. We use the exact optimization and network split points as in (Huo et al., 2018b).

To assess variance in CIFAR-10 accuracy, we perform 3 trials. Tab. 2 shows that the accuracy is the same across the DDG method, backprop, and our approach. DGL achieves better parallelization because it is update unlocked. We use the parallel implementation provided by (Huo et al., 2018b) to obtain a direct wall clock time comparison. We note that there are multiple considerations for comparing speed across these methods (see Appendix C).

Wall Time Comparison We compare to the parallel implementation of (Huo et al., 2018b) using the same communication protocols and run on the same hardware. We find for $K = 2, 4$ GPU gives a 5%, 18% respectively speedup

over DDG. With DDG $K = 4$ giving approximately $2.3\times$ speedup over standard backprop on same hardware (close to results from (Huo et al., 2018b)).

4.2. Large-scale Experiments

Existing methods considering update or backward locking have not been evaluated on large image datasets as they are often unstable or already show large losses in accuracy on smaller datasets. Here we study the optimization of several well-known architectures, mainly the VGG family (Simonyan & Zisserman, 2014) and the ResNet (He et al., 2016), with DG on the ImageNet dataset. In all our experiments we use the MLP-SR-aux auxiliary net which scales well from the smaller CIFAR-10 to the larger ImageNet. The final module has no auxiliary network. For all optimization of auxiliary problems and for end-to-end optimization of reference models we use the shortened optimization schedule prescribed in (Xiao et al., 2019). Results are shown in Tab. 3. We see that for all the models DGL can perform as well and sometimes better than the end-to-end trained models, while permitting parallel training. In all these cases the auxiliary networks are negligible (see Appendix Table 4 for more details). For the VGG-13 architecture we also evaluate the case where the model is trained layer by layer ($K = 10$). Although here performance is slightly degraded, we find it is surprisingly high given that no backward communication is performed. We conjecture that improved auxiliary models and combinations with methods such as (Huo et al., 2018a) to allow feedback on top of the local model, may further improve performance. Also for the settings with larger potential parallelization, slower but more performant auxiliary models could potentially be considered as well.

The synchronous DGL has also favorable memory usage compared to DDG and to the DNI method, DNI requiring to store larger activations and DDG having high memory compared to the base network even for few splits (Huo et al., 2018a). Although not our focus, the single worker version of DGL has favorable memory usage compared to standard backprop training. For example, the ResNet-152 DGL $K = 2$ setting can fit 38% more samples on a single 16GB GPU than the standard end-to-end training.

4.3. Asynchronous DGL with Replay

We now study the effectiveness of Alg. 2 w.r.t delays. We use a 5 layer CIFAR-10 network with the MLP-aux and with all other architecture and optimization settings as in the auxiliary network experiments of Sec. 4.1. Each layer is equipped with a buffer of size M . At each iteration, a layer is chosen according to the pmf $p(j)$, and a batch selected from buffer B_{j-1} . One layer is slowed down by decreasing its selection probability in the pmf $p(j)$ by a factor S . We evaluate different slowdown factors (up to

$S = 2.0$). Accuracy versus S is shown in Fig. 5. For this experiment we use a buffer of size $M = 50$. We run separate experiments with the slowdown applied at each layer of the network as well as 3 random seeds for each of these settings (thus 18 experiments per data point). We show the evaluations for 10 values of S . To ensure a fair comparison we stop updating layers once they have completed 50 epochs, ensuring an identical number of gradient updates for all layers in all experiments.

In practice one could continue updating until all layers are trained. In Fig. 5 we compare to the synchronous case. First, observe that the accuracy of the synchronous algorithm is maintained in the setting where $S = 1.0$ and the pmf is uniform. Note that even this is a non-trivial case, as it will mean that layers inherently have random delays (as compared to Alg. 1). Secondly, observe that accuracy is maintained until approximately $1.2\times$ and accuracy losses after that the difference remains small. Note that even case $S = 2.0$ is somewhat drastic: for 50 epochs, the slowed-down layer is only on epoch 25 while those following it are at epoch 50.

We now consider the performance with respect to the buffer size. Results are shown in Fig. 6. For this experiment we set $S = 1.2\times$. Observe that even a tiny buffer size can yield only a slight loss in performance accuracy. Building on this demonstration there are multiple directions to improve Async DGL with replay. For example improving the efficiency of the buffer (Oyallon et al., 2018), by including data augmentation in feature space (Verma et al., 2018), mixing samples in batches, or improved batch sampling, among others.

5. Related work

To the best of our knowledge (Jaderberg et al., 2017) is the first to directly consider the update or forward locking problems in deep feed-forward networks. Other works (Huo et al., 2018a;b) study the backward locking problem. Furthermore, a number of backpropagation alternatives (Choromanska et al., 2018; Lee et al., 2014; Nø kland, 2016) can address backward locking. However, update locking is a more severe inefficiency. Consider the case where each layer’s forward processing time is T_F and is equal across a network of L layers. Given that the backward pass is a constant multiple in time of the forward, in the most ideal case the backward unlocking will still only scale as $\mathcal{O}(LT_F)$ with L parallel nodes, while update unlocking could scale as $\mathcal{O}(T_F)$.

One class of alternatives to standard back-propagation aims to avoid its biologically implausible aspects, most notably the weight transport problem (Bartunov et al., 2018; Nø kland, 2016; Lillicrap et al., 2014; Lee et al., 2014; Ororb-ia et al., 2018; Ororb-ia & Mali, 2019). Some of these




Figure 5. Evaluation of Async DGL. A single layer is slowed down on average over others, with negligible losses of accuracy at even substantial delays.




Figure 6. Buffer size vs. Acc. for Async DGL. Smaller buffers produce only small loss in acc.

methods (Lee et al., 2014; Nøkland, 2016) can also achieve backward unlocking as they permit all parameters to be updated at the same time, but only once the signal has propagated to the top layer. However, they do not solve the update or forward locking problems. Target propagation uses a local auxiliary network as in our approach, for propagating backward optimal activations computed from the layer above. Feedback alignment replaces the symmetric weights of the backward pass with random weights. Direct feedback alignment extends the idea of feedback alignment passing errors from the top to all layers, potentially enabling simultaneous updates. These approaches have also not been shown to scale to large datasets (Bartunov et al., 2018), obtaining only 17.5% top-5 accuracy on ImageNet (reference model achieving 59.8%). On the other hand, greedy learning has been shown to work well on this task (Belilovsky et al., 2019). We also note concurrent work in the context of biologically plausible models by (Nøkland & Eidnes, 2019) which improves on results from (Mostafa et al., 2018), showing an approach similar to a specific instantiation of the synchronous version of DGL. This work however does not consider the applications to unlocking nor asynchronous training and cannot currently scale to ImageNet.

Another line of related work inspired by optimization methods such as *Alternating Direction Method of Multipliers* (ADMM) (Taylor et al., 2016; Carreira-Perpinan & Wang, 2014; Choromanska et al., 2018) use auxiliary variables to break the optimization into sub-problems. These approaches are fundamentally different from ours as they optimize for

Model (training method)	Top-1	Top-5
VGG-13 (DGL per Layer, $K = 10$)	64.4	85.8
VGG-13 (DGL $K = 4$)	67.8	88.0
VGG-13 (backprop)	66.6	87.5
VGG-19 (DGL $K = 4$)	69.2	89.0
VGG-19 (DGL $K = 2$)	70.8	90.2
VGG-19 (backprop)	69.7	89.7
ResNet-152 (DGL $K = 2$)	74.5	92.0
ResNet-152 (backprop)	74.4	92.1

Table 3. ImageNet results using training schedule of (Xiao et al., 2019) for DGL and standard e2e backprop. DGL with VGG and ResNet obtains similar or better accuracies, while enabling parallelization and reduced memory.

the joint training objective, the auxiliary variables providing a link between a layer and its successive layers, whereas we consider a different objective where a layer has no dependence on its successors. None of these methods can achieve update or forward unlocking. However, some (Choromanska et al., 2018) are able to have simultaneous weight updates (backward unlocked). Another issue with ADMM methods is that most of the existing approaches except for (Choromanska et al., 2018) require standard (“batch”) gradient descent and are thus difficult to scale. They also often involve an inner minimization problem and have thus not been demonstrated to work on large-scale datasets. Furthermore, none of these have been combined with CNNs.

Distributed optimization based on data parallelism is a popular area in machine learning beyond deep learning models and often studied in the convex setting (Leblond et al., 2018). For deep network optimization the predominant method is distributed synchronous SGD (Goyal et al., 2017) and variants, as well as asynchronous (Zhang et al., 2015) variants. Our work is closer to a form of model parallelism rather than data parallelism, and can be easily combined with many data parallel methods (e.g. distributed synchronous SGD). Finally, recent proposals for “pipelining” (Huang et al., 2018b) consider systems level approaches to optimize latency times. These methods do not address the update, forward, locking problems (Jaderberg et al., 2017) which are algorithmic constraints of the learning objective and backpropagation. Pipelining can be seen as attempting to work around these restrictions, with the fundamental limitations remaining. Removing and reducing update, backward, forward locking would simplify the design and efficiency of such systems-level machinery. Finally, tangential to our work Lee et al. (2015) considers auxiliary objectives but with a joint learning objective, which is not capable of addressing any of the problems considered in this work.

6. Conclusion

We have analyzed and introduced a simple and strong baseline for parallelizing per layer and per module computations in CNN training. This work matches or exceeds state-of-the-art approaches addressing these problems and is able to scale to much larger datasets than others. Future work can develop improved auxiliary problem objectives and combinations with delayed feedback.

Acknowledgements

EO acknowledges NVIDIA for its GPU donation. EB acknowledges funding from IVADO. We would like to thank John Zarka, Louis Thiry, Georgios Exarchakis, Fabian Pedregosa, Maxim Berman, Amal Rannen, Kyle Kastner, Aaron Courville, and Nicolas Pinto for helpful discussions.

References

- Bartunov, S., Santoro, A., Richards, B., Marris, L., Hinton, G. E., and Lillicrap, T. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In *Advances in Neural Information Processing Systems*, pp. 9389–9399, 2018.
- Belilovsky, E., Eickenberg, M., and Oyallon, E. Greedy layerwise learning can scale to imagenet. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- Bengio, Y., Lamblin, P., Popovici, D., and Larochelle, H. Greedy layer-wise training of deep networks. In *Advances in neural information processing systems*, pp. 153–160, 2007.
- Bottou, L., Curtis, F. E., and Nocedal, J. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- Carreira-Perpinan, M. and Wang, W. Distributed optimization of deeply nested systems. In *Artificial Intelligence and Statistics*, pp. 10–19, 2014.
- Choromanska, A., Kumaravel, S., Luss, R., Rish, I., Kingsbury, B., Tejwani, R., and Bouneffouf, D. Beyond backprop: Alternating minimization with co-activation memory. *arXiv preprint arXiv:1806.09077*, 2018.
- Czarnecki, W. M., Swirszcz, G., Jaderberg, M., Osindero, S., Vinyals, O., and Kavukcuoglu, K. Understanding synthetic gradients and decoupled neural interfaces. *CoRR*, abs/1703.00522, 2017. URL <http://arxiv.org/abs/1703.00522>.
- Goyal, P., Dollár, P., Girshick, R. B., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677, 2017. URL <http://arxiv.org/abs/1706.02677>.
- Haarnoja, T., Hartikainen, K., Abbeel, P., and Levine, S. Latent space policies for hierarchical reinforcement learning. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1851–1860, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/haarnoja18a.html>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Huang, F., Ash, J., Langford, J., and Schapire, R. Learning deep resnet blocks sequentially using boosting theory. *International Conference on Machine Learning (ICML)*, 2018a.
- Huang, Y., Cheng, Y., Chen, D., Lee, H., Ngiam, J., Le, Q. V., and Chen, Z. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *arXiv preprint arXiv:1811.06965*, 2018b.
- Huo, Z., Gu, B., and Huang, H. Training neural networks using features replay. *Advances in Neural Information Processing Systems*, 2018a.
- Huo, Z., Gu, B., qian Yang, and Huang, H. Decoupled parallel backpropagation with convergence guarantee. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2098–2106, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018b. PMLR. URL <http://proceedings.mlr.press/v80/huo18a.html>.
- Ivakhnenko, A. G. and Lapa, V. G. *Cybernetic Predicting Devices*. CCM Information Corporation., 1965.
- Jaderberg, M., Czarnecki, W. M., Osindero, S., Vinyals, O., Graves, A., Silver, D., and Kavukcuoglu, K. Decoupled neural interfaces using synthetic gradients. *International Conference of Machine Learning*, 2017.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Leblond, R., Pedregosa, F., and Lacoste-Julien, S. Asaga: Asynchronous parallel saga. In *20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017*, 2017.

- Leblond, R., Pedregosa, F., and Lacoste-Julien, S. Improved asynchronous parallel optimization analysis for stochastic incremental methods. *Journal of Machine Learning Research*, 19(81):1–68, 2018. URL <http://jmlr.org/papers/v19/17-650.html>.
- Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z., and Tu, Z. Deeply-supervised nets. In *Artificial intelligence and statistics*, pp. 562–570, 2015.
- Lee, D., Zhang, S., Biard, A., and Bengio, Y. Target propagation. *CoRR*, abs/1412.7525, 2014. URL <http://arxiv.org/abs/1412.7525>.
- Lian, X., Zhang, W., Zhang, C., and Liu, J. Asynchronous decentralized parallel stochastic gradient descent. *arXiv preprint arXiv:1710.06952*, 2017.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. Random feedback weights support learning in deep neural networks. *CoRR*, abs/1411.0247, 2014. URL <http://arxiv.org/abs/1411.0247>.
- Lin, L.-J. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3-4):293–321, 1992.
- Marquez, E. S., Hare, J. S., and Niranjana, M. Deep cascade learning. *IEEE Transactions on Neural Networks and Learning Systems*, 29(11):5475–5485, Nov 2018. ISSN 2162-237X. doi: 10.1109/TNNLS.2018.2805098.
- Mostafa, H., Ramesh, V., and Cauwenberghs, G. Deep supervised learning using local errors. *Frontiers in neuroscience*, 12:608, 2018.
- Nøkland, A. Direct feedback alignment provides learning in deep neural networks. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 1037–1045. 2016.
- Nøkland, A. and Eidnes, L. H. Training neural networks with local error signals. *arXiv preprint arXiv:1901.06656*, 2019.
- Ororbia, A. G. and Mali, A. Biologically motivated algorithms for propagating local target representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4651–4658, 2019.
- Ororbia, A. G., Mali, A., Kifer, D., and Giles, C. L. Conducting credit assignment by aligning local representations. *arXiv preprint arXiv:1803.01834*, 2018.
- Oyallon, E. Building a regular decision boundary with deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5106–5114, 2017.
- Oyallon, E., Belilovsky, E., Zagoruyko, S., and Valko, M. Compressing the input for cnns with the first-order scattering transform. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 301–316, 2018.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Taylor, G., Burmeister, R., Xu, Z., Singh, B., Patel, A., and Goldstein, T. Training neural networks without gradients: A scalable admm approach. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2722–2731, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/taylor16.html>.
- Verma, V., Lamb, A., Beckham, C., Najafi, A., Courville, A., Mitliagkas, I., and Bengio, Y. Manifold mixup: Learning better representations by interpolating hidden states. 2018.
- Xiao, W., Chen, H., Liao, Q., and Poggio, T. A. Biologically-plausible learning algorithms can scale to large datasets. *International Conference on Learning Representations*, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- Zhang, S., Choromanska, A. E., and LeCun, Y. Deep learning with elastic averaging sgd. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*, pp. 685–693. Curran Associates, Inc., 2015.

A. Proofs

Lemma 4.1. *Under Assumption 3 and 4, one has: $\forall \Theta_j, \mathbb{E}_{p_{j-1}^*} [\|\nabla_{\Theta_j} \ell(Z_{j-1}; \Theta_j)\|^2] \leq G$.*

Proof. First of all, observe that under Assumption 4 and via Fubini's theorem:

$$\sum_t c_{j-1}^t = \sum_t \int |p_{j-1}^t(z) - p_{j-1}^*(z)| dz = \int \sum_t |p_{j-1}^t(z) - p_{j-1}^*(z)| dz < \infty \quad (3)$$

thus, $\sum_t |p_j^t - p_j^*|$ is convergent a.s. and $|p_j^t - p_j^*| \rightarrow 0$ a.s. as well. From Fatou's lemma, one has:

$$\int p_{j-1}^*(z) \|\nabla_{\Theta_j} \ell(z; \Theta_j)\|^2 dz = \int \liminf_t p_{j-1}^t(z) \|\nabla_{\Theta_j} \ell(z; \Theta_j)\|^2 dz \quad (4)$$

$$\leq \liminf_t \int p_{j-1}^t(z) \|\nabla_{\Theta_j} \ell(z; \Theta_j)\|^2 dz \leq G \quad (5)$$

□

Lemma 4.2. *Under Assumptions 1, 3 and 4, we have:*

$$\mathbb{E}[\mathcal{L}(\Theta_j^{t+1})] \leq \mathbb{E}[\mathcal{L}(\Theta_j^t)] - \eta_t (\mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] - \sqrt{2G} c_{j-1}^t) + \frac{LG}{2} \eta_t^2,$$

Observe that the expectation is taken over each random variable.

Proof. By L -smoothness:

$$\mathcal{L}(\Theta_j^{t+1}) \leq \mathcal{L}(\Theta_j^t) + \nabla \mathcal{L}(\Theta_j^t)^T (\Theta_j^{t+1} - \Theta_j^t) + \frac{L}{2} \|\Theta_j^{t+1} - \Theta_j^t\|^2 \quad (6)$$

Substituting $\Theta_j^{t+1} - \Theta_j^t$ on the right:

$$\mathcal{L}(\Theta_j^{t+1}) \leq \mathcal{L}(\Theta_j^t) - \eta_t \nabla \mathcal{L}(\Theta_j^t)^T \nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t) + \frac{L\eta_t^2}{2} \|\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)\|^2 \quad (7)$$

Taking the expectation w.r.t. Z_{j-1}^t which has a density p_{j-1}^t , we get:

$$\mathbb{E}_{p_{j-1}^t} [\mathcal{L}(\Theta_j^{t+1})] \leq \mathcal{L}(\Theta_j^t) - \eta_t \nabla \mathcal{L}(\Theta_j^t)^T \mathbb{E}_{p_{j-1}^t} [\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)] + \frac{L\eta_t^2}{2} \mathbb{E}_{p_{j-1}^t} [\|\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)\|^2]$$

From Assumption 3, we obtain that:

$$\frac{L\eta_t^2}{2} \mathbb{E}_{p_{j-1}^t} [\|\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)\|^2] \leq \frac{L\eta_t^2 G}{2} \quad (8)$$

Then, as a side computation, observe that:

$$\|\mathbb{E}_{p_{j-1}^t} [\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)] - \nabla \mathcal{L}(\Theta_j^t)\| = \left\| \int \nabla \ell(z, \Theta_j^t) p_{j-1}^t(z) dz - \int \nabla \ell(z, \Theta_j^t) p_{j-1}^*(z) dz \right\| \quad (9)$$

$$\leq \int \|\nabla \ell(z, \Theta_j^t)\| |p_{j-1}^t(z) - p_{j-1}^*(z)| dz \quad (10)$$

$$= \int (\|\nabla \ell(z, \Theta_j^t)\| \sqrt{|p_{j-1}^t(z) - p_{j-1}^*(z)|}) \sqrt{|p_{j-1}^t(z) - p_{j-1}^*(z)|} dz \quad (11)$$

$$(12)$$

Let us apply the Cauchy-Swchartz inequality, we obtain:

$$\|\mathbb{E}_{p_{j-1}^t}[\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)] - \nabla \mathcal{L}(\Theta_j^t)\| \leq \sqrt{\int \|\nabla \ell(z, \Theta_j^t)\|^2 |p_{j-1}^t(z) - p_{j-1}^*(z)| dz} \sqrt{\int |p_{j-1}^t(z) - p_{j-1}^*(z)| dz} \quad (13)$$

$$= \sqrt{\int \|\nabla \ell(z, \Theta_j^t)\|^2 |p_{j-1}^t(z) - p_{j-1}^*(z)| dz} \sqrt{c_{j-1}^t} \quad (14)$$

Then, observe that:

$$\int \|\nabla \ell(z, \Theta_j^t)\|^2 |p_{j-1}^t(z) - p_{j-1}^*(z)| dz \leq \int \|\nabla \ell(z, \Theta_j^t)\|^2 (p_{j-1}^t(z) + p_{j-1}^*(z)) dz \quad (15)$$

$$= \mathbb{E}_{p_{j-1}^t}[\|\nabla \ell(Z_{j-1}, \Theta_j^t)\|^2] + \mathbb{E}_{p_{j-1}^*}[\|\nabla \ell(Z_{j-1}, \Theta_j^t)\|^2] \quad (16)$$

$$\leq 2G \quad (17)$$

The last inequality follows from Lemma 4.1 and Assumption 3.

Then, using again Cauchy-Schwartz inequality:

$$\left| \|\nabla \mathcal{L}(\Theta_j^t)\|^2 - \nabla \mathcal{L}(\Theta_j^t)^T \mathbb{E}_{p_{j-1}^t}[\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)] \right| = \left| \nabla \mathcal{L}(\Theta_j^t)^T (\nabla \mathcal{L}(\Theta_j^t) - \mathbb{E}_{p_{j-1}^t}[\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)]) \right| \quad (18)$$

$$\leq \|\nabla \mathcal{L}(\Theta_j^t)\| \|\mathbb{E}_{p_{j-1}^t}[\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)] - \nabla \mathcal{L}(\Theta_j^t)\| \quad (19)$$

$$\leq \|\nabla \mathcal{L}(\Theta_j^t)\| \sqrt{2Gc_{j-1}^t} \quad (20)$$

Then, taking the expectation leads to

$$|\mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2 - \nabla \mathcal{L}(\Theta_j^t)^T \mathbb{E}_{p_{j-1}^t}[\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)]]| \leq \mathbb{E}[|\|\nabla \mathcal{L}(\Theta_j^t)\|^2 - \nabla \mathcal{L}(\Theta_j^t)^T \mathbb{E}_{p_{j-1}^t}[\nabla_{\Theta_j} \ell(Z_{j-1}^t; \Theta_j^t)]]|] \quad (21)$$

$$\leq \mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|] \sqrt{2Gc_{j-1}^t} \quad (22)$$

$$\leq \sqrt{\mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2]} \sqrt{2Gc_{j-1}^t} \quad (23)$$

$$(24)$$

However, observe that by Lemma 4.1 and Jensen inequality:

$$\|\nabla \mathcal{L}(\Theta_j^t)\|^2 = \|\mathbb{E}_{p_j^*}[\nabla_{\Theta_j} \ell(Z, \Theta_j^t)]\|^2 \leq \mathbb{E}_{p_j^*}[\|\nabla_{\Theta_j} \ell(Z, \Theta_j^t)\|^2] \leq G \quad (25)$$

Combining this inequality and Assumption 3, we get:

$$\mathbb{E}[\mathcal{L}(\Theta_j^{t+1})] \leq \mathbb{E}[\mathcal{L}(\Theta_j^t)] - \eta_t (\mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] - \sqrt{2Gc_{j-1}^t}) + \frac{LG}{2} \eta_t^2,$$

□

Proposition 4.2. *Under Assumptions 1, 2, 3 and 4, each term of the following equation converges:*

$$\sum_{t=0}^T \eta_t \mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] \leq \mathbb{E}[\mathcal{L}(\Theta_j^0)] + G \sum_{t=0}^T \eta_t (\sqrt{2c_{j-1}^t} + \frac{L\eta_t}{2}) \quad (26)$$

Proof. Applying Lemma 3.2 for $t = 0, \dots, T - 1$, we obtain (observe the telescoping sum), for our non-negative loss:

$$\sum_{t=0}^T \eta_t \mathbb{E}[\|\nabla \mathcal{L}(\Theta_j^t)\|^2] \leq \mathbb{E}[\mathcal{L}(\Theta_j^0)] - \mathbb{E}[\mathcal{L}(\Theta_j^{T+1})] + \sqrt{2}G \sum_{t=0}^T \sqrt{c_j^t} \eta_t + \frac{LG}{2} \sum_{t=0}^T \eta_t^2 \quad (27)$$

$$\leq \mathbb{E}[\mathcal{L}(\Theta_j^0)] + \sqrt{2}G \sum_{t=0}^T \sqrt{c_j^t} \eta_t + \frac{LG}{2} \sum_{t=0}^T \eta_t^2 \quad (28)$$

$$(29)$$

Yet, $\sum \sqrt{c_j^t} \eta_t$ is convergent, as $\sum c_j^t$ and $\sum_t \eta_t^2$ are convergent, thus the right term is bounded. \square

B. Additional Descriptions of Experiments

Here we provide some additional details of the experiments. Code for experiments is provided along with the supplementary materials.

Comparisons to DNI The comparison to DNI attempts to directly replicate the Appendix C.1 (Jaderberg et al., 2017). Although the baseline accuracies for backprop and cDNI are close to those reported in the original work, those of DNI are worse than those reported in (Jaderberg et al., 2017), which could be due to minor differences in the implementation. We utilize a popular pytorch DNI implementation available and source code will be provided.

Auxiliary network study We use SGD with momentum of 0.9 and weight decay 5×10^{-4} (Zagoruyko & Komodakis, 2016) and a short schedule of 50 epochs and decay factor of 0.2 every 15 epochs (Belilovsky et al., 2019).

Sequential vs Greedy optimization experiments We use the same architecture and optimization as in the Auxiliary network study

Imagenet We use the shortened optimization schedule prescribed in (Xiao et al., 2019). It consists of training for 50 epochs with mini-batch size 256, uses SGD with momentum of 0.9, weight decay of 10^{-4} , and a learning rate of 0.1 reduced by a factor 10 every 10 epochs.

C. Detailed Discussion of Relative Speed of Competing Methods

Here we describe in more detail the elements governing differences between methods such as DNI(Jaderberg et al., 2017), DDG/FA(Huo et al., 2018a), and the simpler DGL. We will argue that if we take the assumption that each approach runs for the same number of epochs or iterations and applies the same splits of the network then DGL is by construction faster than the other methods which rely on feedback. The relative speeds of these methods are governed by the following:

1. Computation besides forward and backward passes on primary network modules (e.g. auxiliary networks forward and backward passes)
2. Communication time of sending activations from one module to the next module
3. Communication time of sending feedback to the previous module
4. Waiting time for signal to reach final module

As discussed in the text our auxiliary modules which govern (1) for DGL are negligible thus the overhead of (1) is negligible. DNI will inherently have large auxiliary models as it must predict gradients, thus (1) will be much greater than in DGL. (2) should be of equal speed across all methods given the same implementation and hardware. (3) does not exist for the case of DGL but exists for all other cases. (4) applies only in the case of backward unlocking methods (DDG/FA) and does not exist for DNI or DGL as they are update unlocked.

Thus we observe that DGL by construction is faster than the other methods. We note that for use cases in multi-GPU settings communication would need to be well optimized for use of any of these methods. Although we include a parallel

Decoupled Greedy Learning of CNNs

	Flops Net	Flops Aux
VGG-13 ($K = 4$)	13 GFLOPs	0.2 GFLOP
VGG-19 ($K = 4$)	20 GFLOPs	0.2 GFLOP
ResNet-152 ($K = 2$)	11 GFLOP	0.02 GFLOP

Table 4. ImageNet comparisons of FLOPs for auxiliary model in major models trained. Auxiliary networks are negligible.

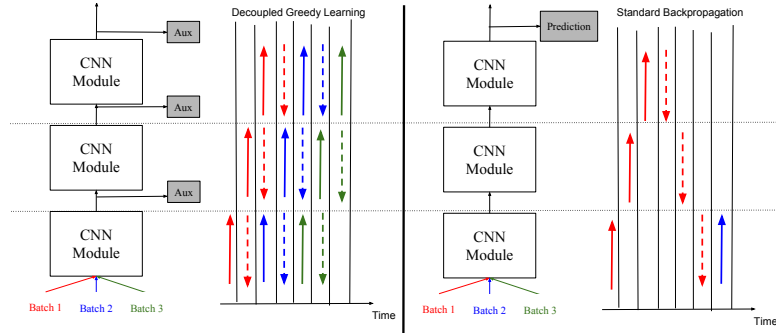


Figure 7. We illustrate the signal propagation for three mini-batches processed by standard back-propagation and with decoupled greedy learning. In each case a module can begin processing forward and backward passes as soon as possible. For illustration we assume same speed for forward and backward passes, and discount the auxiliary network computation (negligible here).

implementation based on the software from (Huo et al., 2018b), an optimized distributed implementations of the ideas presented here and related works is outside of the scope of this work.

C.1. Auxiliary Network Sizes and FLOP comparisons on ImageNet

We briefly illustrate the sizes of auxiliary networks. Lets take as an example the ImageNet experiments for VGG-13. At the first layer the output is $224 \times 224 \times 64$. The MLP-aux here would be applied after averaging to $2 \times 2 \times 64$, and would consists of 3 fully connected layers of size 256 ($2 * 2 * 64$) followed by a projection to 1000 image categories. The MLP-SR-aux network used would first reduce to $56 \times 56 \times 64$ and then apply 3 layers of 1×1 convolutions of width 64. This is followed by reduction to 2×2 and 3 FC layers as in the MLP-aux network. As mentioned in Sec. 4.2 the auxiliary networks are negligible in size. We further illustrate this in 4.

D. Additional pseudo-code

To illustrate the parallel implementations of the Algorithms we show a different pseudocode implementation with an explicit behavior for each worker specified. The following Algorithm 3 is equivalent to Algorithm 1 in terms of output but directly illustrates a parallel implementation. Similarly 4 illustrates a parallel implementation of the algorithm described in Algorithm 2. The probabilities used in Algorithm 4 are not included here as they are derived from communication and computation speed differences. Finally we illustrate the parallelism compared to backprop in 7

Algorithm 3: DGL Parallel Implementation

Input: Stream $\mathcal{S} \triangleq \{(x_0^t, y^t)\}_{t \leq T}$ of samples or mini-batches;

- 1 **Initialize** Parameters $\{\theta_j, \gamma_j\}_{j \leq J}$
- 2 Worker 0:
- 3 **for** $x_0^t \in \mathcal{S}$ **do**
- 4 $x_1^t \leftarrow f_{\theta_0^t}(x_0^t)$
- 5 Send x_0^t to worker 1
- 6 Compute $\nabla_{(\gamma_1, \theta_1)} \hat{\mathcal{L}}(y^t, x_0^t; \gamma_0^t, \theta_0^t)$
- 7 $(\theta_0^{t+1}, \gamma_0^{t+1}) \leftarrow$ Step of parameters (θ_0^t, γ_0^t)
- 8 **end**
- 9 Worker j :
- 10 **for** $t \in 0 \dots T$ **do**
- 11 Wait until x_{j-1}^t is available
- 12 $x_j^t \leftarrow f_{\theta_{j-1}^t}(x_{j-1}^t)$
- 13 Compute $\nabla_{(\gamma_j, \theta_j)} \hat{\mathcal{L}}(y^t, x_j^t; \gamma_j^t, \theta_j^t)$
- 14 Send x_j^t to worker x_{j+1}^t
- 15 $(\theta_j^{t+1}, \gamma_j^{t+1}) \leftarrow$ Step of parameters (θ_j^t, γ_j^t)
- 16 **end**

Algorithm 4: DGL Async Buffer Parallel Impl.

Input: Stream $\mathcal{S} \triangleq \{(x_0^t, y^t)\}_{t \leq T}$; Distribution of the delay $p = \{p_j\}_j$; Buffer size M

- 1 **Initialize:** Buffers $\{B_j\}_j$ with size M ; params $\{\theta_j, \gamma_j\}_j$
- 2 Worker j :
- 3 **while training do**
- 4 **if** $j = 1$ **then** $(x_0, y) \leftarrow \mathcal{S}$
- 5 **else** $(x_{j-1}, y) \leftarrow B_{j-1}$
- 6
- 7 $x_j \leftarrow f_{\theta_{j-1}}(x_{j-1})$
- 8 Compute $\nabla_{(\gamma_j, \theta_j)} \hat{\mathcal{L}}(y, x_j; \gamma_j, \theta_j)$
- 9 $(\theta_j, \gamma_j) \leftarrow$ Step of parameters (θ_j, γ_j)
- 10 **if** $j < J$ **then** $B_j \leftarrow (x^j, y)$
- 11 **end**
