

Update Compression for Deep Neural Networks on the Edge

Bo Chen¹ Ali Bakhshi² Gustavo Batista² Brian Ng¹ Tat-Jun Chin¹
¹The University of Adelaide ²University of New South Wales

Abstract

An increasing number of artificial intelligence (AI) applications involve the execution of deep neural networks (DNNs) on edge devices. Many practical reasons motivate the need to update the DNN model on the edge device post-deployment, such as refining the model, concept drift, or outright change in the learning task. In this paper, we consider the scenario where retraining can be done on the server side based on a copy of the DNN model, with only the necessary data transmitted to the edge to update the deployed model. However, due to bandwidth constraints, we want to minimise the transmission required to achieve the update. We develop a simple approach based on matrix factorisation to compress the model update—this differs from compressing the model itself. The key idea is to preserve existing knowledge in the current model and optimise only small additional parameters for the update which can be used to reconstitute the model on the edge. We compared our method to similar techniques used in federated learning; our method usually requires less than half of the update size of existing methods to achieve the same accuracy.

1. Introduction

The significant progress in AI has been due in large part to the resurgence of deep neural networks (DNNs), particularly convolutional neural networks (CNNs) and variants thereof. As the complexity and breadth of problems that are solvable in an end-to-end manner by DNNs increase, the DNN architectures (“models”) themselves are becoming deeper [18, 19, 25, 61] and wider [64, 65, 74], which demand greater computational resources to execute.

On the other hand, there is a trend to deploy DNNs in the field through edge computing devices, such as embedded GPUs and FPGAs. This is due to the growing popularity of concepts such as AI on the edge [6, 41, 44, 51], which emphasises processing closer to the source of the data to minimise latency and reduce data transfer, and embodied AI [5, 15, 55], which focusses on building intelligent agents that can explore and interact with the environment.

The gap in computational capability between centralised

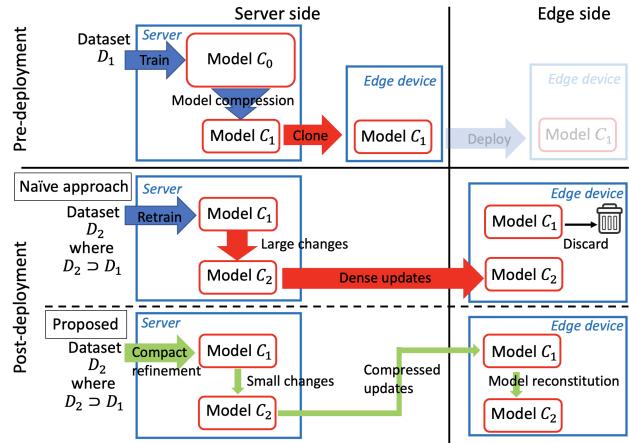


Figure 1. Illustrating update compression to alleviate bandwidth constraints for updating DNN models on the edge.

processing systems and edge computing devices has motivated research in making complex DNNs feasible on the latter platforms. Significant attention has been devoted to compressing DNNs to reduce their memory consumption and/or accelerate their execution. Major categories of such methods include pruning and sparsification [20, 21, 43, 69, 76], quantisation [4, 39, 46, 50, 56], knowledge distillation [7, 22, 53, 68, 73], low-rank optimisation [8, 14, 31, 32, 34, 52], and architecture search [16, 24, 45, 57, 63, 66, 67, 72, 75]. A common observation is that many complex DNNs are over-parametrised, hence allowing substantial compression [9, 29, 54, 60].

Another important requirement is *updating* the DNN on the edge *after* deployment. Numerous practical reasons, such as refining the model using new data, concept drift, or outright changes to the learning task, make updating DNN models unavoidable in machine learning applications. In this paper, we consider the specific scenario where

- The (new) training data and supervision labels are available on the central processing system (the “server”);
- Communication between server and edge can be unreliable or costly, which discourages large data transfers.

Many applications fall under this setting, e.g., AI-enabled sensor networks, field robotics, autonomous driving cars,

intelligent drone fleets, etc. The scenario argues for retraining the DNN on the server based on a clone of the deployed DNN, then transmitting only the new parameter values to update the model on the edge, as illustrated in Fig. 1.

Update compression To mitigate the communication bottleneck, it is vital to *compress the model update* such that the data can be transferred as bandwidth-efficiently as possible. Note that this aim differs from compressing the model itself; indeed, even if the deployed DNN is already a compressed model, it is still worthwhile to seek to transfer as little as possible to the edge to accomplish the update, particularly when recurrent updates are required.

A simple solution is to apply lossless compression techniques [58] on the model changes. However, this misses the potential gains via using more strategic model refinement that directly compresses the update during retraining.

Our contributions Towards the above end, we inject the goal of update compression in the model refinement step itself. We propose a *compact refinement* technique based on re-parameterisation of DNN layers, which allows to express the model changes in terms of a small set of new parameters and to preserve a large part of model knowledge in the form of frozen buffers. On the edge side, a *reconstitution* step integrates the learnt update package with the preserved knowledge and reconfigures the model. The proposed method is illustrated in Fig. 1 and 2.

We examined the trade-off between update size and impact on inference accuracy to illustrate the superiority of our method over existing update compression techniques. In our classification experiment our method lifts the test accuracy of the initial model by $\approx 9\%$ with an update package equivalent to 0.5% of the model parameters, while competing methods require 2 to 40 times bigger update packages to achieve the same lift. Moreover, our experiments show that our model is less reliant on parameter redundancy than existing techniques, which is more suitable for compact models deployed on resource limited edge devices.

2. Related Work

In this section, we briefly survey the related literature. We focus on three main research directions: *Model compression* techniques that create small-scale versions of large DNNs. *Update compression* methods that allow efficient transmission of model parameters in client-server architectures. *Split computing* approaches that divide a DNN model into two parts that execute in the client and server cooperatively.

2.1. Model Compression

The execution of deep-learning models on edge devices has attracted much attention in the last few years [6]. Edge devices are often limited by memory, processing and power

constraints that motivate the development of compressed DNN models.

Compressing a model can benefit memory footprint and processing time, with often a minor decrease in classification performance [10]. There are six main categories of model compression techniques: quantisation, network slimming, weight pruning, low-rank factorisation, knowledge distillation and neural architecture search. We list below a few prominent works from each category.

Quantisation reduces the number of bits used to represent the network weights [4, 39, 46, 50, 56]. Network slimming removes the least important channels from convolutional layers [21, 69, 76]. Weight pruning removes the least significant connections in the model [20]. Low-rank factorisation use matrix/tensor decomposition to estimate the informative parameters of the DNN model [8, 14, 31, 32, 34, 52]. Knowledge distillation trains a more compact neural network to reproduce the output of a larger network [7, 22, 53, 68, 73]. Neural architecture search uses a search procedure to automatically design efficient models [16, 24, 45, 57, 63, 66, 67, 72, 75].

While the above mentioned techniques are capable of effectively compress the model size and, consequently, the update size, they are likely to be less efficient for update compression because the scenario is different: for update compression it does not require the whole model parameters to be updated which means the model can store aside useful information and only change what is necessary.

2.2. Model Update Compression

Besides hardware constraints, edge devices often rely on network communication to exchange data, model parameters and predictions. However, network communication can be costly, unreliable and energy-hungry. Improving communication efficiency is a research topic that has been explored in various client-server DNN architectures such as Federated Learning (FL) [49] and, more generally, distributed SGD [71].

Our work assumes that training data is available at the server, which has the computational power to retrain the DNN model. In contrast, Federated Learning requires that training data, including class labels, be available at the client devices. These devices have the computational power to update the model locally. The server receives the model updates from multiple clients and averages these updates.

The main approaches for model update compression can be divided in three groups: low-rank update, random masking and gradient-based techniques.

The low-rank update was used as an update compression method in FL by [33]. It consists in decomposing the gradient tensor in two matrices, L and R , with a rank r . L is generated randomly, and R is optimised during training. After training, L is compressed as a random seed and trans-

mitted with R to the server.

Random masking defines a random sparsity matrix and requires the training to update the non-null entries in such a sparse matrix [33]. Like the low-rank update method, a random seed defines the sparse pattern transmitted to the server with the non-zero entries.

Gradient-based techniques aim to reduce the communication burdens during training via gradient quantisation [3, 59], sparsification [33, 62], or both [12, 17]. However, compressing communication during training differs from compressing communication for model update, as model training requires many iterations whereas the latter only allows for a one-off data transmission.

2.3. Split Computation

Split Computing (SC) [48] is a framework that divides the DNN model into head and tail models, which are executed in the edge device and server, respectively. SC is attractive when compressed models for edge devices cannot achieve the same level of accuracy as their full counterpart models. SC has two main limitations: the inference time becomes the sum of inference time on client and server plus the communication latency, and the need for a “bottleneck” layer in the first model layers. The bottleneck layer provides a cutpoint for the head and tail models with a compact representation that reduces communication overhead.

The straightforward splitting of the DNN as suggested by [28, 30, 36] can result in either transferring a large part of the processing burden to the edge device or transmitting a larger volume of data on the network. Distilling the head section of DNN and introducing a bottleneck within the distilled head model as suggested in [47] can mitigate this problem and decrease the computational cost and the required bandwidth considerably.

3. Problem setting

We consider the problem of performing bandwidth efficient updating for a DNN on an edge device deployed in the field. As depicted in Fig. 1, the setting consists of a central server and an edge device, between which there is a communication bandwidth limit. We assume that the central server has no restriction in computational resources while the edge device has limited resources and can only perform inference and light-duty computation.

Pre-deployment For concreteness, we focus on classifier models $C(\cdot|\theta)$ with parameters θ . Also, since our work concentrates on update compression and not model compression (see Fig. 1 on the distinction), we assume that $C(\cdot|\theta)$ is already a lightweight model suitable for edge devices.

Let $D_1 = \{(\mathbf{X}_i, Y_i)\}_{i=1}^{N_1}$ denote the initial training dataset. At the pre-deployment stage, training C with D_1

results in model parameters θ_1

$$\theta_1 = \arg \min_{\theta} \ell(\theta, D_1), \quad (1)$$

where ℓ is some loss function for training C , such as cross-entropy loss. The classifier $C(\cdot|\theta_1)$ is then loaded onto the edge device and deployed to the field.

Post-deployment After deployment, the central server has accumulated a larger dataset $D_2 = \{(\mathbf{X}_i, Y_i)\}_{i=1}^{N_2}$, where $N_2 > N_1$ and $D_2 \supset D_1$. The goal is to update C on the edge device, currently with parameters θ_1 , with the new dataset D_2 , subject to a data transmission bottleneck.

Let Δ denote the update package to be sent from the server to the edge device. A basic method is to retrain C using D_2 , and let Δ be the set of all model parameters, *i.e.*,

$$\Delta = \arg \min_{\theta} \ell(\theta, D_2).$$

At the edge upon receiving Δ , θ_1 is discarded and we simply equate $\theta_2 = \Delta$ and insert it into $C(\cdot|\theta_2)$ to complete the model update. However, in this case Δ will always be the set of all parameters even for a small incremental update. The key to generating more compact updates lies in reusing the learnt knowledge embedded in θ_1 , as we will show next.

4. Update compression

Instead of densely retraining the whole model, we would like a *compact refinement* algorithm that produces a compact update package Δ using the current model $C(\cdot|\theta_1)$ and new dataset D_2

$$\Delta = \text{compact-refine}(C(\cdot|\theta_1), D_2), \quad (2)$$

so that it can be efficiently transmitted to the edge device. At the edge side, upon receiving the compact update package Δ , the current model $C(\cdot|\theta_1)$ incorporates Δ and reconstitutes the updated model

$$C(\cdot|\theta_2) = \text{reconstitute}(C(\cdot|\theta_1), \Delta). \quad (3)$$

The overall pipeline of such a update compression scheme is demonstrated in Fig. 2.

4.1. Compact refinement

To achieve compact refinement, we first re-parameterise the current network in a way such that

1. the new learnable parameters are much fewer than the original ones; and
2. the model can be effectively updated by refining the new parameters with the new dataset D_2 .

To this end, we focus on the network’s convolutional (Conv) layers and fully connected (FC) layers, as mainstream DNNs largely consist of these two types of layers. Other types of layers, such as batch normalisation [26], are not targeted for update compression in our method.

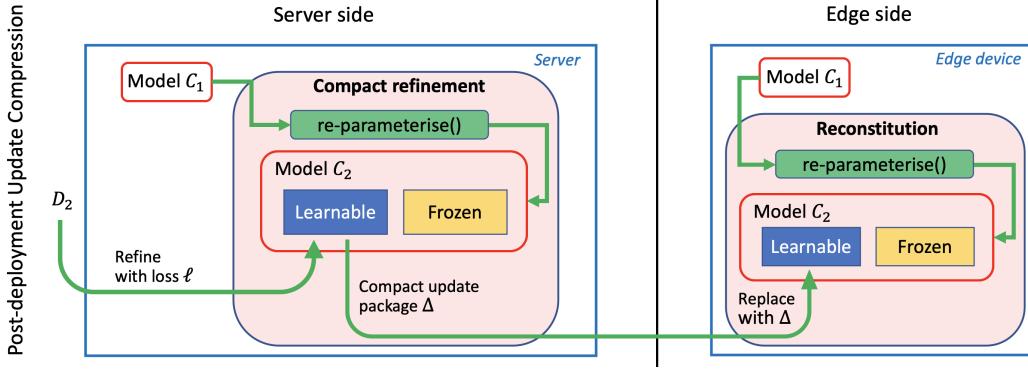


Figure 2. A schematic pipeline of the proposed post-deployment update compression framework. To achieve compactness of the update package, we propose to re-parameterise the model into a learnable part and a frozen part. The learnable parameters facilitates the compact refinement and the frozen data enables the edge device to recycle knowledge from existing model.

Low Rank Approximation (LRA) We first describe the compact refinement algorithm using a baseline re-parameterisation method LRA before we introduce the proposed techniques. Let $\phi_l \subset \theta$ denote the weight tensor of layer l , which is either a Conv or FC layer. Note that ϕ_l does not include the bias terms b_l of the layer. If the layer has bias terms, they are not aimed for update compression in our method and remain as learnable parameters. We firstly re-arrange ϕ_l to a matrix of size $o \times i$. For Conv layers, o is the number of output channels and i is the product of the number of input channels and the size of the kernel along each dimension, *e.g.*, the kernel height and the kernel width for the 2D case. For FC layers, o is simply the number of output features and i the input features.

The matrix ϕ_l is then decomposed using SVD:

$$\phi_l = U \cdot \text{diag}(s) \cdot V^T, \quad (4)$$

where $U \in \mathbb{R}^{o \times m}$, $V \in \mathbb{R}^{i \times m}$, $m = \min(o, i)$, $s \in \mathbb{R}^{m \times 1}$ is the vector of singular values sorted in descending order, and $\text{diag}(s)$ is the $m \times m$ square matrix of 0s except for the diagonal entries which are specified by s .

For either Conv or FC layers, the arrangement of ϕ_l can be viewed as o rows of filters with length i . Consequently, the decomposed U , V and s can be viewed as a mapping matrix, a matrix of base filters, and a weight vector respectively, the latter indicating the importance of each base filter.

The decomposition of ϕ_l can result in even more parameters if U , s , and V are directly used as new parameters. To achieve compactness, we adopt LRA to encode the original parameters. Let $U_{1:r}$ denote the $o \times r$ matrix which consists of the first r columns of U , and $s_{1:r}$ denote the $r \times 1$ vector of the first r entries in s . We then define a new set of learnable parameters for the layer

$$\phi'_l = \{L, R\}, \quad (5)$$

where

$$L = U_{1:r} \cdot \text{diag}(s_{1:r}), \quad R = V_{1:r}^T. \quad (6)$$

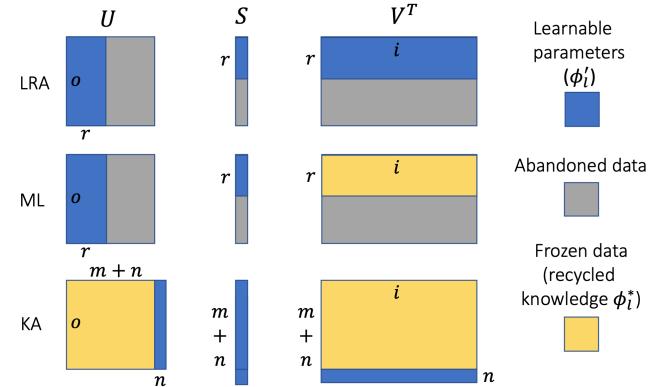


Figure 3. Illustration of different re-parameterisation methods.

Note that this does not change the network’s architecture. After a Conv or FC layer is re-parameterised, for each forward propagation during the refinement process, the layer’s weight tensor ϕ_l is firstly computed with ϕ'_l using a recover-weight() routine

$$\phi_l \leftarrow \text{recover-weight}(\phi'_l) \quad (7)$$

$$= L \cdot R. \quad (8)$$

The layer then performs normal forward processing. For backward propagation in such layers, only gradients with respect to ϕ'_l are computed instead of ϕ_l , since ϕ_l is now an intermediate variable rather than a leaf variable in the computational graph. The LRA re-parameterisation method is conceptually illustrated in Fig. 3 row 1.

By re-parameterising all Conv and FC layers, the network’s new learnable parameters become

$$\theta' = (\theta \setminus \{\phi_l\}_{l=1}^L) \cup \{\phi'_l\}_{l=1}^L, \quad (9)$$

where L is the total number of Conv and FC layers in the network. Note that θ' includes the network’s original pa-

rameters that are not targeted for compact refinement, *e.g.*, bias terms, batch-norm parameters.

Finally, the update package Δ is obtained via refining θ' from its current values with D_2

$$\Delta = \arg \min_{\theta'} \ell(\theta', D_2). \quad (10)$$

The LRA approach could use a small rank r to achieve compactness

$$|\Delta| = |\theta'| \ll |\theta|. \quad (11)$$

However, since it does not reuse any existing knowledge from θ_1 it can result in severe underfitting if r is too small. LRA can be viewed as a network compression method which is similar to previous works [11, 27, 32, 37, 42]. It is thus considered as a baseline method in our experiments.

Mapping Learning (ML) To further reduce the number of learnable parameters used in refinement while performing effective update, we propose to recycle part of previously learnt knowledge embedded in θ_1 . We obtain L and R the same way as Eq. (6) with rank r . The matrix R consists of the most important r rows of base filters learnt from D_1 . It is thus possible to utilise these base filters and only learn a new mapping for the model to adapt to dataset D_2 .

Therefore, the re-parameterisation of ϕ_l should produce a set of learnable parameters ϕ'_l and a set of frozen data ϕ^*_l which stores the recycled knowledge:

$$\phi'_l = L, \quad \phi^*_l = R. \quad (12)$$

The frozen data ϕ^*_l is stored in the network as a buffer, which participates in forward computation but does not require gradient computation in back-propagation. Consequently, the recover-weight() routine is given by

$$\text{recover-weight}(\phi'_l, \phi^*_l) = \phi'_l \cdot \phi^*_l = L \cdot R. \quad (13)$$

Knowledge Augmentation (KA) To maximally preserve previously learnt knowledge in order to minimise the update size, we propose to freeze both U and V in Eq. (4), and only refine on s as free parameters. However, the base filters V and its mappings U learnt from D_1 are likely to generalise poorer than they could have if they were learnt from D_2 , because a larger training set is usually beneficial as it is less likely to cause overfitting.

To let the model learn to adapt to D_2 whilst preserving current knowledge U and V , we propose to augment U and V with additional columns for capturing the “knowledge drift” between D_1 and D_2 . Let $U' \in \mathbb{R}^{o \times n}$ denote the augmenting vectors and $[U, U'] \in \mathbb{R}^{o \times (m+n)}$ denote the augmented mapping matrix. Similarly let $V' \in \mathbb{R}^{i \times n}$ be the augmenting vectors and $[V, V'] \in \mathbb{R}^{i \times (m+n)}$ be the augmented filter matrix. Note that both U and V are frozen and only U' and V' are learnable parameters.

Let $s' \in \mathbb{R}^{(m+n) \times 1}$ denote the augmented weight vector where the whole vector is treated as learnable parameters. The layer is thus re-parameterised as

$$\phi'_l = \{U', s', V'\}, \quad \phi^*_l = \{U, V\}, \quad (14)$$

and the recover-weight() routine is

$$\text{recover-weight}(\phi'_l, \phi^*_l) = [U, U'] \cdot \text{diag}(s') \cdot [V, V']^T. \quad (15)$$

An illustration of this re-parameterisation method is given in Fig. 3 row 3. The motivation of this design is to maximally reuse previous knowledge by freezing U and V , and at the same time allowing them to be updated by learning the augmenting vectors and incorporating them using the re-learned weights s' .

After re-parameterisation we initialise ϕ'_l in a way such that the refinement process starts from the current accuracy level. The refinement process starts with a weight tensor

$$\phi_l \leftarrow [U, U'] \cdot \text{diag}(s') \cdot [V, V']^T \quad (16)$$

which could be abruptly different from the original ϕ_l because U' , s' and V' are initialised randomly. To avoid a significant drop in accuracy at the beginning of refinement, we let $s'_{1:m} = s$ and U', V' and $s'_{m+1:m+n}$ be small random values that are close to 0 for initialisation.

Algorithm 1 compact-refine($C(\cdot|\theta)$, D)

```

1:  $\theta' \leftarrow \theta$ 
2: for  $l$  in  $\{1, \dots, \mathcal{L}\}$  do
3:    $\phi_l \leftarrow$  layer  $l$  weight tensor in  $\theta_1$ 
4:    $\phi'_l, \phi^*_l \leftarrow$  re-parameterise( $\phi_l$ )
5:    $\theta' \leftarrow (\theta' \setminus \phi_l) \cup \phi'_l$ 
6: end for
7:  $\Delta \leftarrow \arg \min_{\theta'} \ell(\theta', D)$ 
8: return  $\Delta$ 

```

The overall procedure of compact refinement is summarised in Algorithm 1, with the re-parameterise() routine and the recover-weight() routine depending on the specific method employed.

4.2. Model reconstitution

Once the edge device receives Δ , it then incorporates Δ to update its current model. This process is described in Algorithm 2. For all non Conv or FC layers, the edge device updates their parameters with new values from Δ . For Conv and FC layers, their weights are recovered using the new values of the learnable parameters ϕ'_l from Δ and the values of the frozen data ϕ^*_l obtained by decomposing the current model weights θ_1 .

An additional desirable feature of the proposed KA method is that, once the edge device reconstitutes the updated model the augmenting vectors are integrated into the

Algorithm 2 reconstitute($C(\cdot|\theta_1)$, Δ)

```
1:  $\theta_2 \leftarrow \{\text{all non Conv/FC layers in } \Delta\}$ 
2: for  $l$  in  $\{1, \dots, \mathcal{L}\}$  do
3:    $\phi_l \leftarrow \text{layer } l \text{ weight tensor in } \theta_1$ 
4:    $\phi_l^t, \phi_l^* \leftarrow \text{re-parameterise}(\phi_l)$ 
5:    $\phi_l^\Delta \leftarrow \text{layer } l \text{ free parameters in } \Delta$ 
6:    $\phi_l \leftarrow \text{recover-weight}(\phi_l^\Delta, \phi_l^*)$ 
7:    $b_l \leftarrow \text{layer } l \text{ bias in } \Delta$ 
8:    $\theta_2 \leftarrow \theta_2 \cup \{\phi_l, b_l\}$ 
9: end for
10: return  $C(\cdot|\theta_2)$ 
```

weight tensor by the recover-weight() routine and thus can be discarded. This prevents the edge device to store more and more augmenting vectors as a result of repetitive updates, which eliminates the requirement of maintenance over time for expanding its memory capacity.

5. Experiments

In this section we conduct various experiments to evaluate the performance of update compression with the proposed methods, as well as competing methods.

5.1. Competitor methods

We compared the proposed update compression algorithms to two techniques used in Federated Learning, Random Mask (RM) [33] and Low-Rank Update (LRU) [33], as well as a compression method Filter Pruning (FP) [40].

RM For the RM method we generate a random binary mask for the entire model parameters. The proportion of parameters that are selected by the mask is controlled by a hyper-parameter P . We then refine the model on D_2 by updating only the selected parameters while keeping the unselected ones frozen. The update package size is counted as the number of selected parameters. We ignore the size of the mask as it can be represented by a random seed.

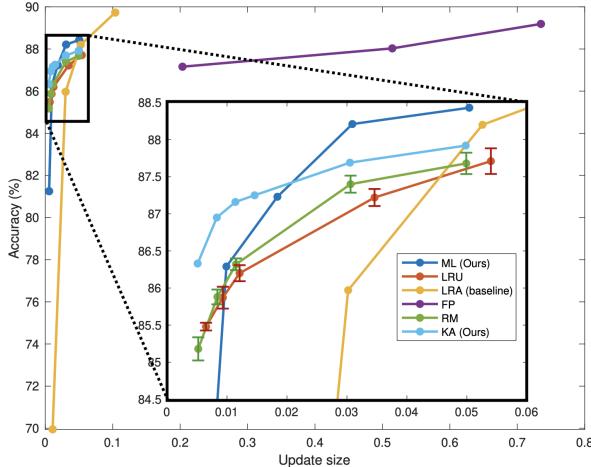
LRU To implement LRU we firstly decompose the weight tensor of each Conv and FC layer to two matrices L and R with rank r in the same way as Eq. (6). We then replace the values of R with random numbers and fix them. During refinement L is initialised with small values and $L \cdot R$ is added to the layer's original weight tensor, *i.e.*, $L \cdot R$ learns the changes of the current weight tensor, as described in [33] Sec.2. The update size of LRU is all the L matrices and other free parameters (such as bias, Batch-Norm layers, etc). Again we ignore the size of the random matrices R as they can be represented by random seeds.

FP The FP method prunes whole filters as well as their connected feature maps in the network. The size of the pruned model is dependent on a compression rate c and the layers selected for the pruning plan. In this experiment we only consider Conv and batch normalisation layers in the pruning plan. We count all remaining parameters after pruning into the update size.

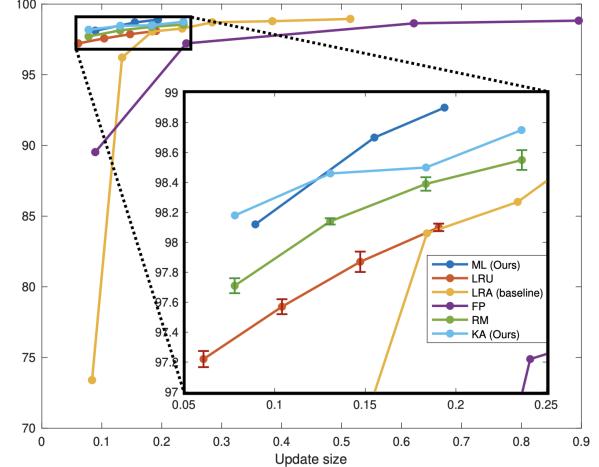
5.2. Update compression for classification models

We evaluate update compression methods for two classification tasks. The first one uses the CIFAR10 [1, 35] dataset on a ResNet18 [18] model and the second one classifies the MNIST [2, 38] dataset with a modified tiny-ResNet model. Since we are mainly concerned with the power- and bandwidth-constrained application, the CIFAR10 and MNIST datasets with low-resolution images are a good choice. The tiny-Resnet architecture is composed of a feature extractor with three layers, each layer consists of a basic residual block and a classifier that include one fully connected layer. Similarly, the VGG-tiny architecture is comprised of a feature extractor that contains a sequence of convolutional layers with 16, 16, 32, and 64 outputs, each is followed by a max-pooling layer. Besides, one fully connected layer is on top of them. For both tasks we use a proportion p of the training set as D_1 and the whole training set as D_2 . The model is initially trained with D_1 and then updated using D_2 with various update methods. We then test the updated model with the testing set and report its accuracy against the update size. The proportion for D_1 is $p = 0.2$ for the CIFAR10 experiment and $p = 0.02$ for MNIST, as MNIST is a much easier dataset and a higher p would make the accuracy gap between the initial model and the refined model too small for the experiment. The initial test accuracy of the model trained with D_1 is 77.5% for ResNet18 on CIFAR10 and 95.5% for tiny-ResNet on MNIST. The results of this experiment is shown in Fig. 4. Note that in this experiment the purpose is not to push the test accuracy as high as possible, but to compare the updating efficiency among different methods, *i.e.*, achieving high accuracy with a small update size.

Fig. 4a shows that KA has the highest accuracy for update size below 2%, while ML performs best with update size between 2% to 5%. The curves of RM and LRU are much more to the right hand side than KA's, which means in order to achieve the same accuracy, RM and LRU requires 2 to 3 times larger update packages than KA. The low rank approximation nature of ML makes it deteriorates quickly when update size shrinks below 1%. On the other hand, KA is much more robust for small update size thanks to its augmentation mechanism. For the compression method FP, its accuracy at update size 20% is just similar to that of KA at update size 1%. Results in Fig. 4b tell a similar story where KA has the highest accuracy for the smallest update, closely



(a) CIFAR10 classification with ResNet18. Results of RM and LRU are averages of 6 runs and their 90% confidence intervals are plotted.



(b) MNIST classification with tiny-ResNet. Results of RM and LRU are averages of 6 runs and their 90% confidence intervals are plotted.

Figure 4. The update size versus test accuracy after refinement. Update size is measured as a proportion between the number of new parameters ($|\theta'|$) and the number of original parameters ($|\theta|$).

followed by ML. In contrast, it takes RM and LRU at least twice as big an update size as KA’s for the same accuracy.

Note that in both experiments the LRA baseline is capable to achieve the highest accuracy when given enough parameter size but performs poorly at a low update size. This indicates that knowledge recycling has its pros and cons: while it can be super efficient in terms of update size, methods like LRA without preserving any old knowledge can achieve higher accuracy when given enough parameters.

5.3. The effect of parameter redundancy

Besides knowledge recycling, we suspect that another factor contributing to the effectiveness of update compression is parameter redundancy, which allows the model to adapt to the new dataset by updating only a small fraction of parameters. However, for edge devices models deployed on them are usually optimised, distilled, or compressed to perform the target task well enough while minimising resource consumption. Therefore, update compression methods should be robust to the lack of parameter redundancy.

We repeat the CIFAR10 experiment on a customised VGG-tiny [61] model, a much smaller model than ResNet18 with only 26554 parameters, to test different updating methods on compact models that do not have a lot of parameter redundancy. The result in Fig. 5 shows that the gap between KA and the two competing methods RM and LRU is clearly wider than that in ResNet18. This suggests that with a small model the lack of parameter redundancy has restricted the performance of RM and LRU, while the proposed KA is less affected by the small model size, which is attractive for resource limited edge devices.

To further test this hypothesis, we compare KA, RM

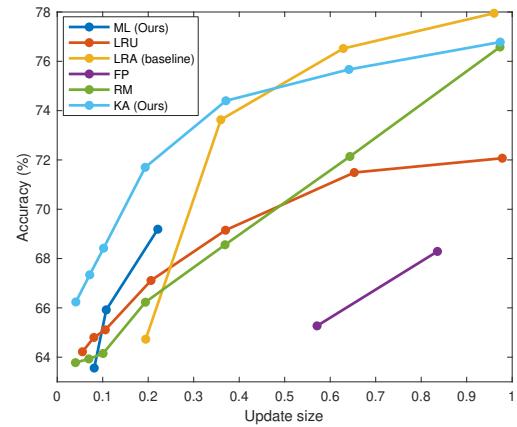


Figure 5. The update size versus test accuracy after refinement on VGG-tiny.

and LRU on different models with parameters ranging from 136K to 23M. For each model, we set the hyper-parameter n of KA to 3 and let the update size of RM and LRU to be similar to but no less than KA’s. We use 20% of CIFAR10 training set as D_1 and 100% as D_2 . We report the accuracy lift as well as other experimental details in Table 1.

For all models in Table 1 regardless of their size, KA outperforms RM and LRU consistently with an equal or smaller update size. Furthermore, as the model becomes more and more compact the accuracy lifts of all methods tend to be smaller as well, supporting our hypothesis about parameter redundancy. This effect is especially obvious for RM, which can hardly lift accuracy for small models. In contrast, KA is less reliant on parameter redundancy and

Model	Method	Update size	Update %	Model size	Acc. lift	Refined acc.	Init acc.
VGG-small	KA	7413	5.44	136266	2.28	78.68	76.4
VGG-small	RM	7413	5.44	136266	0.02	76.42	76.4
VGG-small	LRU	7878	5.78	136266	1.45	77.85	76.4
VGG-medium	KA	14672	2.72	539786	2.83	80.97	78.14
VGG-medium	RM	14681	2.72	539786	0.07	78.21	78.14
VGG-medium	LRU	15638	2.90	539786	1.44	79.58	78.14
MobileNetV2	KA	122632	5.34	2296922	3.67	87.92	84.25
MobileNetV2	RM	122660	5.34	2296922	3.18	87.43	84.25
MobileNetV2	LRU	128734	5.60	2296922	3.54	87.79	84.25
ResNet18	KA	127939	1.14	11173962	9.67	87.16	77.49
ResNet18	RM	128378	1.15	11173962	8.97	86.46	77.49
ResNet18	LRU	135670	1.21	11173962	9.02	86.51	77.49
ResNet50	KA	310656	1.32	23520842	6.57	84.34	77.77
ResNet50	RM	312837	1.33	23520842	4.61	82.38	77.77
ResNet50	LRU	327185	1.39	23520842	6.16	83.93	77.77

Table 1. Accuracy lift on different models. Update size and model size are both measured in number of parameters.

remains effective for models of all sizes.

5.4. Effectiveness of knowledge renewal

All three methods of KA, RM, and LRU reuses significant old knowledge to achieve update compression. RM keeps all unselected parameters and only learns selected ones to update the model, while LRU renews the old model by learning to approximate the residuals of the old parameters. We evaluate the effectiveness of knowledge renewal amongst the three different techniques in the circumstance that the model is compact and there is little parameter redundancy. We train a VGG-tiny model with different initial training set size, *i.e.*, $|D_1|$ ranges from 0.1 to 0.7 of $|D_2|$. We then perform update compression to the model with different techniques and report their accuracy lifts after refinement with D_2 . For KA the hyper-parameter n is set to 1. For RM and LRU their update size are set to be as close as but not less than KA's.

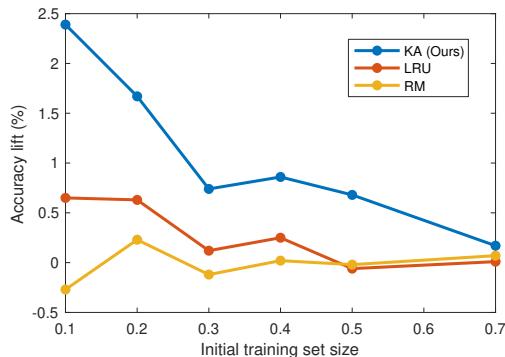


Figure 6. Initial training set size vs accuracy lift on VGG-tiny. The initial training set size is computed by $|D_1|/|D_2|$.

Fig. 6 shows that while the gap between the initial set D_1 and the target set D_2 enlarges, KA is more effective in learning new knowledge to update the model than RM and LRU, even when parameter redundancy is lacking. This implies the KA update strategy would benefit applications where D_2 grows ever larger over time.

6. Limitations

The proposed update compression methods only targets Conv and FC layers. It is unable to handle models that largely consist of other types of layers, such as RNN [13], LSTM [23] and Transformer [70]. Due to limited resources and paper space, we have also limited our experiments to classification tasks and small to medium size models, the effect of practical application of our methods outside of these settings requires further testings.

7. Conclusion

In this paper we address the problem of server-to-edge device model update with a limited communication bandwidth. The proposed update compression framework achieves superior communication efficiency via novel compact refinement and reconstitution techniques. Experiments further show that our method is robust to the lack of parameter redundancy.

Acknowledgements

We gratefully acknowledge generous funding from SmartSat CRC, Research Program 2: Advanced Satellite Systems, Sensors and Intelligence.

Tat-Jun Chin is SmartSat CRC Profesorial Chair of Sentient Satellites.

References

[1] Cifar10 dataset. <https://www.cs.toronto.edu/~kriz/cifar.html>. 6

[2] Mnist dataset. <http://yann.lecun.com/exdb/mnist/>. License: Creative Commons Attribution-Share Alike 3.0 license. 6

[3] Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In *NeurIPS*, 2017. 3

[4] Milad Alizadeh, Arash Behboodi, Mart van Baalen, Christos Louizos, Tijmen Blankevoort, and Max Welling. Gradient l_1 regularization for quantization robustness. *arXiv preprint arXiv:2002.07520*, 2020. 1, 2

[5] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *NeurIPS*, 2020. 1

[6] Jiasi Chen and Xukan Ran. Deep learning with edge computing: A review. *Proceedings of the IEEE*, 107(8):1655–1674, 2019. 1, 2

[7] Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. In *ICLR*, 2016. 1, 2

[8] Ting Chen, Ji Lin, Tian Lin, Song Han, Chong Wang, and Denny Zhou. Adaptive mixture of low-rank factorizations for compact neural modeling. 2018. 1, 2

[9] Yunpeng Chen, Haoqi Fan, Bing Xu, Zhicheng Yan, Yannis Kalantidis, Marcus Rohrbach, Shuicheng Yan, and Jiashi Feng. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In *ICCV*, 2019. 1

[10] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*, 2017. 2

[11] Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *NeurIPS*, 2014. 5

[12] Nikoli Dryden, Tim Moon, Sam Ade Jacobs, and Brian Van Essen. Communication quantization for data-parallel training of deep neural networks. In *2nd Workshop on Machine Learning in HPC Environments (MLHPC)*, pages 1–8. IEEE, 2016. 3

[13] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. 8

[14] Timur Garipov, Dmitry Podoprikhin, Alexander Novikov, and Dmitry Vetrov. Ultimate tensorization: compressing convolutional and fc layers alike. *arXiv preprint arXiv:1611.03214*, 2016. 1, 2

[15] Graham Gobieski, Brandon Lucia, and Nathan Beckmann. Intelligence beyond the edge: Inference on intermittent embedded systems. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 199–213, 2019. 1

[16] Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, and Edward Choi. Morphnet: Fast & simple resource-constrained structure learning of deep networks. In *CVPR*, 2018. 1, 2

[17] Corentin Hardy, Erwan Le Merrer, and Bruno Sericola. Distributed deep learning on edge-devices: feasibility via adaptive compression. In *International Symposium on Network Computing and Applications (NCA)*, 2017. 3

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 6

[19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*. Springer, 2016. 1

[20] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *ECCV*, 2018. 1, 2

[21] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017. 1, 2

[22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2

[23] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 8

[24] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *ICCV*, 2019. 1, 2

[25] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 1

[26] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456. PMLR, 2015. 3

[27] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *Proceedings of the British Machine Vision Conference*, 2014. 5

[28] Hyuk-Jin Jeong, InChang Jeong, Hyeon-Jae Lee, and Soo-Mook Moon. Computation offloading for machine learning web apps in the edge server environment. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pages 1492–1499. IEEE, 2018. 3

[29] Kumara Kahatapitiya and Ranga Rodrigo. Exploiting the redundancy in convolutional filters for parameter reduction. In *WACV*, 2021. 1

[30] Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. *ACM SIGARCH Computer Architecture News*, 45(1):615–629, 2017. 3

[31] Hyeji Kim, Muhammad Umar Karim Khan, and Chong-Min Kyung. Efficient neural network compression. In *CVPR*, 2019. 1, 2

[32] Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. In *ICLR*, 2016. 1, 2, 5

[33] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016. 2, 3, 6

[34] Jean Kossaifi, Adrian Bulat, Georgios Tzimiropoulos, and Maja Pantic. T-net: Parametrizing fully convolutional nets with a single high-order tensor. In *CVPR*, 2019. 1, 2

[35] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images, 2009. 6

[36] Nicholas D Lane, Sourav Bhattacharya, Petko Georgiev, Claudio Forlivesi, Lei Jiao, Lorena Qendro, and Fahim Kawsar. Deepx: A software accelerator for low-power deep learning inference on mobile devices. In *15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12. IEEE, 2016. 3

[37] V Lebedev, Y Ganin, M Rakhuba, I Oseledets, and V Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. In *ICLR*, 2015. 5

[38] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 6

[39] Hao Li, Soham De, Zheng Xu, Christoph Studer, Hanan Samet, and Tom Goldstein. Training quantized nets: A deeper understanding. In *NeurIPS*, 2017. 1, 2

[40] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *ICLR*, 2017. 6

[41] He Li, Kaoru Ota, and Mianxiong Dong. Learning iot in edge: Deep learning for the internet of things with edge computing. *IEEE network*, 32(1):96–101, 2018. 1

[42] Shaohui Lin, Rongrong Ji, Chao Chen, Dacheng Tao, and Jiebo Luo. Holistic cnn compression via low-rank decomposition with knowledge transfer. *TPAMI*, 41(12):2889–2905, 2018. 5

[43] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 1

[44] Xiangzhong Luo, Di Liu, Hao Kong, and Weichen Liu. Edgenas: Discovering efficient neural architectures for edge systems. In *International Conference on Computer Design (ICCD)*, pages 288–295. IEEE, 2020. 1

[45] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018. 1, 2

[46] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Training binary neural networks with real-to-binary convolutions. In *ICLR*, 2020. 1, 2

[47] Yoshitomo Matsubara, Sabur Baidya, Davide Callegaro, Marco Levorato, and Sameer Singh. Distilled split deep neural networks for edge-assisted real-time systems. In *Proceedings of the 2019 Workshop on Hot Topics in Video Analytics and Intelligent Edges*, pages 21–26, 2019. 3

[48] Yoshitomo Matsubara, Marco Levorato, and Francesco Restuccia. Split computing and early exiting for deep learning applications: Survey and research challenges. *arXiv preprint arXiv:2103.04505*, 2021. 3

[49] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017. 2

[50] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *ICML*. PMLR, 2020. 1, 2

[51] Zhaolong Ning, Peiran Dong, Xiaojie Wang, Xiping Hu, Lei Guo, Bin Hu, Yi Guo, Tie Qiu, and Ricky YK Kwok. Mobile edge computing enabled 5g health monitoring for internet of medical things: A decentralized game theoretic approach. *IEEE Journal on Selected Areas in Communications*, 39(2):463–478, 2020. 1

[52] Bo Peng, Wenming Tan, Zheyang Li, Shun Zhang, Di Xie, and Shiliang Pu. Extreme network compression via filter group approximation. In *ECCV*, 2018. 1, 2

[53] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. In *ICLR*, 2018. 1, 2

[54] Jiaxiong Qiu, Cai Chen, Shuaicheng Liu, Heng-Yu Zhang, and Bing Zeng. Slimconv: Reducing channel redundancy in convolutional neural networks by features recombining. *IEEE Transactions on Image Processing*, 30:6434–6445, 2021. 1

[55] Santhosh K Ramakrishnan, Dinesh Jayaraman, and Kristen Grauman. An exploration of embodied visual exploration. *International Journal of Computer Vision*, 129(5):1616–1649, 2021. 1

[56] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, 2016. 1, 2

[57] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 1, 2

[58] Khalid Sayood. *Lossless compression handbook*. Elsevier, 2002. 2

[59] Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth Annual Conference of the International Speech Communication Association*. Citeseer, 2014. 3

[60] Wenling Shang, Kihyuk Sohn, Diogo Almeida, and Honglak Lee. Understanding and improving convolutional neural networks via concatenated rectified linear units. In *ICML*. PMLR, 2016. 1

[61] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 1, 7

[62] Nikko Strom. Scalable distributed dnn training using commodity gpu cloud computing. In *Sixteenth Annual Conference of the International Speech Communication Association*, 2015. 3

[63] Ke Sun, Mingjie Li, Dong Liu, and Jingdong Wang. Igcv3: Interleaved low-rank group convolutions for efficient deep neural networks. *arXiv preprint arXiv:1806.00178*, 2018. 1, 2

- [64] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. [1](#)
- [65] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. [1](#)
- [66] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*, 2019. [1, 2](#)
- [67] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*. PMLR, 2019. [1, 2](#)
- [68] Shitao Tang, Litong Feng, Wenqi Shao, Zhanghui Kuang, Wei Zhang, and Yimin Chen. Learning efficient detector with semi-supervised adaptive distillation. *arXiv preprint arXiv:1901.00366*, 2019. [1, 2](#)
- [69] Frederick Tung and Greg Mori. Clip-q: Deep network compression learning by in-parallel pruning-quantization. In *CVPR*, 2018. [1, 2](#)
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. [8](#)
- [71] Jianyu Wang and Gauri Joshi. Cooperative sgd: A unified framework for the design and analysis of communication-efficient sgd algorithms. *arXiv preprint arXiv:1808.07576*, 2018. [2](#)
- [72] Guotian Xie, Jingdong Wang, Ting Zhang, Jianhuang Lai, Richang Hong, and Guo-Jun Qi. Interleaved structured sparse convolutional neural networks. In *CVPR*, 2018. [1, 2](#)
- [73] Zheng Xu, Yen-Chang Hsu, and Jiawei Huang. Training shallow and thin networks for acceleration via knowledge distillation with conditional adversarial networks. *arXiv preprint arXiv:1709.00513*, 2017. [1, 2](#)
- [74] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [1](#)
- [75] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018. [1, 2](#)
- [76] Zhuangwei Zhuang, Mingkui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu. Discrimination-aware channel pruning for deep neural networks. In *NeurIPS*, 2018. [1, 2](#)