Deep Feature Surgery: Towards Accurate and Efficient Multi-Exit Networks

Cheng Gong¹, Yao Chen^{3*}, Qiuyang Luo¹, Ye Lu^{2,6,7}, Tao Li^{2,5}, Yuzhi Zhang^{1,5}, Yufei Sun^{1*,5}, and Le Zhang⁴

 ¹ College of Software, Nankai University yufei_sun@sina.com
² College of Computer Science, Nankai University
³ National University of Singapore yaochen@nus.edu.sg
⁴ School of Information and Communication Engineering, University of Electronic Science and Technology of China
⁵ HAIHE Lab of ITAI
⁶ Tianjin Key Laboratory of Network and Data Security Technology
⁷ Key Laboratory of Data and Intelligent System Security, Ministry of Education

Abstract. Multi-exit network is a promising architecture for efficient model inference by sharing backbone networks and weights among multiple exits. However, the gradient conflict of the shared weights results in sub-optimal accuracy. This paper introduces Deep Feature Surgery (DFS), which consists of feature partitioning and feature referencing approaches to resolve gradient conflict issues during the training of multiexit networks. The feature partitioning separates shared features along the depth axis among all exits to alleviate gradient conflict while simultaneously promoting joint optimization for each exit. Subsequently, feature referencing enhances multi-scale features for distinct exits across varying depths to improve the model accuracy. Furthermore, DFS reduces the training operations with the reduced complexity of backpropagation. Experimental results on Cifar100 and ImageNet datasets exhibit that DFS provides up to a 50.00% reduction in training time and attains up to a 6.94% enhancement in accuracy when contrasted with baseline methods across diverse models and tasks. Budgeted batch classification evaluation on MSDNet demonstrates that DFS uses about $2 \times$ fewer average FLOPs per image to achieve the same classification accuracy as baseline methods on Cifar100. The code is available at https: //github.com/GongCheng1919/dfs.

1 Introduction

A multi-exit network refers to a neural network architecture designed with multiple points at which the network outputs a result or makes a decision, rather than having a single endpoint. These intermediate exits can be utilized for various purposes such as early exiting for efficiency, where the network can provide a quick

^{*} Corresponding Author



Fig. 1: Solutions to improve the accuracy of the exits for multi-exit networks. (a) Feature enhancement introduces additional features to improve the accuracy of exits. (b) Gradient selection adjusts gradients from conflicting sources to attain consistent update directions. (c) Our method alleviates gradient conflict and ensures end-to-end joint optimization of exits.

prediction without the need for full computation through the entire network depth. Multi-exit networks are widely applied in domains including dynamic computing graph [9], anytime prediction [18], and collaborative inference [19].

In multi-exit network training, it is imperative to formulate multiple losses for distinct exits and concurrently optimize these losses. However, multiple exits manifest inconsistent or diametrically opposing gradient directions, denoted as the gradient conflict issue [1,2,4,12,21,33,34]. This issue hinders the convergence of multi-exit networks, consequently exerting a impact on the model accuracy.

Recent research advancements have introduced enhancements for training multi-exit networks, primarily categorized into feature enhancement [10, 18, 27, 35, 36] and gradient selection [12, 18, 32]. Feature enhancement introduces additional features (more data or wider layer) and training iterations to improve the accuracy of the exits of the model as shown in Fig. 1a. This approach is irrelated to the gradient conflict issue and may lead to elevated training expenses and necessitate model redesign [10, 18, 27, 35, 36]. Gradient selection adjusts gradients from conflicting sources to attain consistent update directions conducive to model optimization as shown in Fig. 1b. This approach is generally heuristic, demonstrating effectiveness in specific models but facing accuracy challenges for universal adoption [12, 18, 32]. In addition, it has to take additional forward-backward phases to merge the gradients to adjust the update directions of the parameters, leading to a long model training time [21, 25, 32, 34].

This paper proposes Deep Feature Surgery (DFS) to tackle both the accuracy and training efficiency challenges in training multi-exit networks. Specifically, DFS utilizes two techniques for feature tensors to address the gradient conflict and enables joint optimization to the different groups of parameters, namely feature partitioning and feature referencing. Feature partitioning divides the features of different exits along the depth axis into exit-specific and shared components as shown in Fig. 1c, and updates the corresponding parameters only during the backpropagation. This partitioning reduces the shared features and the corresponding shared parameters among distinct exits, effectively alleviating gradient conflict while ensuring end-to-end joint optimization of exits with varying depths. Feature referencing allows exits with different depths to reuse the exit-specific features to enhance their accuracy. In addition, DFS reduces

the back-propagation operations and improves the model training efficiency by reducing a substantial portion of gradient computation during model training.

In summary, our contributions are as follows:

- We propose DFS, which consists of feature partitioning and feature referencing to address the gradient conflict issue of training multi-exit networks.
- We reduce the required operations during model training with DFS to improve the training efficiency of multi-exit networks.
- Extensive experiments demonstrate that DFS consistently outperforms baseline methods across various models and tasks with up to 50.00% reduction in training time and a maximum improvement of 6.94% in accuracy.

2 Related work

Training multi-exit networks is challenging since the gradients from multiple exits hold conflicts, which usually lead to sub-optimal model accuracy [34].

2.1 Feature Enhancement

Feature enhancement tries to introduce more robust features to improve the accuracy of multi-exit networks which includes knowledge transfer [18], knowledge distillation [18,27,35,36], weighted sample learning [10], and dense connections [13]. Since feature enhancement is an irrelative technique for solving gradient conflict issues, it was combined with other techniques to further improve the performance of multi-exit architectures. Feature enhancement methods often necessitate extensive modifications to the model [13], loss function [18,35], and even the dataset [10], thereby impacting the model training efficiency.

2.2 Gradient Selection

Gradient selection is effective in many applications such as computer vision [21. 32] and natural language processing [1, 12] by selecting a proper updating direction. One way is to adjust the gradients to multiple exits. For example, in [18], researchers develop a gradient equilibrium strategy to reduce the gradient variance and stabilize the training process; Meta-GF [32] take account of the importance of the shared weights to each exit, and introduce a meta-learning-based method for weighted fusion of the gradients of each exit; [12] propose Gradient Remedy (GR) to change the angle and magnitude of gradient vectors to solve interference between two gradients in noise-robust speech recognition. Another way is to balance the multiple losses to the multiple exits and obtain better-weighted gradients without complex intervention. [7] design an adaptive loss-weighting policy to prioritize more difficult tasks. GradNorm [2] and [22] introduce the magnitude re-scaling algorithms to balance the gradient magnitudes of different tasks. Other widely used gradient selection methods are multi-objective algorithms, such as the MGDA [29], PCGrad [34], CAGrad [21], GetMTL [1], Nash-MTL [25], and so on. These methods formulate multi-exit network training as 4 Cheng Gong et al.

multi-objective optimization and try to pursue the Pareto optimalities for better trade-offs among multiple losses to multiple exits. Gradient selection inevitably degrades the overall accuracy of networks in the gradient adjusting and compromises among multiple exits, *i.e.*, for one loss to give up updating others. Furthermore, this technique requires computing independent weight gradients for each exit to calculate the merge coefficients, which often entails multiple forward and backward passes for one training iteration, resulting in a manifold increase in training time.

2.3 Feature partitioning

Feature partitioning is a technique proposed for mitigating task interference in multi-task learning [4, 5, 23, 24, 26, 31]. For instance, ETR-NLP [5] partitions the features into task-specific and task-shared ones, to update the corresponding parameters regarding different tasks during model training. It trains the task-specific and shared parameters separately, thus largely increasing training iterations and delaying model training. Mask-based works [26,31] learn the mask module for each task during training to select the proper parameters. The mask modules for different tasks are optimized separately which requires a long training time. [23, 24] employ the attention mechanisms at the filter level as feature selectors allowing each task to select a subset of parameters. They train the attention modules and model parameters simultaneously in one training iteration where attention modules delay the training. GradSplit [4] splits features into multiple groups for multiple tasks and trains each group only using the gradients back-propagated from the task losses. These studies have demonstrated success in multi-task learning, but they do not translate effectively to training multiexit networks due to the training efficiency issue and gap between multi-task learning and multi-exit networks. To the best of our knowledge, DFS is the first study that employs feature partitioning to mitigate the gradient conflict issue in training multi-exit networks, and significantly enhances both model accuracy and training efficiency of multi-exit networks.

3 Preliminary: Gradient Conflict

Gradient conflict is the major cause of the accuracy degradation in multi-exit network training. Let a feature extraction architecture, such as a convolutional neural network, with L layers be the backbone network. Take the classification task as an example, this multi-exit network has L tasks for the exits, denoted as $C = \{c_1, c_2, \dots, c_L\}$ and weights for these layers from the backbone network, denoted as $W_L = \{w_1, w_2, \dots, w_L\}$. Since each exit also has its own exit-specific weights, denoted as $W_C = \{w_{c_1}, w_{c_2}, \dots, w_{c_L}\}$. Let (X, Y) be a dataset, in which X is the sample set and Y is the labels of X, and the *i*-th task c_i of the given multi-exit network is defined as follows.

$$c_i = \pi_i(f_i, w_{c_i}), \quad s.t. \begin{cases} f_i = \varphi_i(X, W_i), \\ i = 1, 2, \cdots, L \end{cases}$$
(1)

Here, $\pi_i(\cdot)$ processes the exit-specific output with the exit-specific weight w_{c_i} and the intermediate feature f_i from the *i*-th layer of backbone network. f_i is extracted by $\varphi_i(\cdot)$ with input X and weights W_i . For weight sets W_i and W_j , they share weights as follows.

$$W_i \cap W_j = \{w_1, w_2, \cdots, w_{\min(i,j)}\}, \quad s.t. \begin{cases} i = 1, 2, \cdots, L\\ j = 1, 2, \cdots, L \end{cases}$$
(2)

As a result, the optimization formulation of this multi-exit network can be derived as follows.

$$W_L^*, W_C^* = \underset{W_L, W_C}{\arg\min} \{ CE(c_1, Y), CE(c_2, Y), \cdots, CE(c_L, Y) \}.$$
(3)

Here, $CE(\cdot)$ computes the cross entropy of prediction $\{c_i | i = 1, 2, \dots, L\}$ and category Y. Obviously, Eq. (3) is a multi-objective optimization [29] problem, which implies that it is hard to achieve the optimal solutions for all classification tasks simultaneously.

The most common methods employ a gradient descent algorithm to optimize Eq. (3), the gradients g_{w_i} of tasks with respect to weight w_i is computed as the summation of $\{g_{w_i}^k | k = i, i + 1, \dots, L\}$ as follows since w_i is the shared weight among these exits [13].

$$g_{w_i} = \sum_{k=i}^{L} g_{w_i}^k, \quad s.t. \quad g_{w_i}^k = \frac{\partial CE(c_k, Y)}{\partial w_i}$$
(4)

Although $g_{w_i}^k$ implies the optimal updating direction of w_i to reduce $CE(c_k, Y)$. The inconsistency in the gradients of $\{g_{w_i}^k | k = 1, 2, \dots, L\}$ is known as gradient conflict and affects the training performance.

4 Deep Feature Surgery

DFS addresses the gradient conflict by creating both individual and shared parameter updates. DFS conducts feature partitioning and feature referencing on the feature tensors to ensure the desirable model accuracy of different exits.

4.1 Feature Partitioning

Gradient conflict is caused by the shared weights of exits in multi-exit networks, as shown in Eq. (2). In order to alleviate the gradient conflict, we first decouple the weight w_i of *i*-th layer into two distinct parts with a coefficient β . Let $w_i \in \mathbb{R}^{C_{in} \times C_{out} \times K \times K}$, we define:

$$w_{i} = [w_{i}^{+}, w_{i}^{-}]_{\beta}, \ s.t. \begin{cases} w_{i}^{+} = w_{i}[:, : \lfloor \beta C_{out} \rceil, :, :] \\ w_{i}^{-} = w_{i}[:, \lfloor \beta C_{out} \rceil, :, :] \end{cases}$$
(5)

Here, $\lfloor \cdot \rceil$ computes the rounding result of input, and w_i [:] represents a slicing operation on the tensor w_i . $\lfloor \cdot \rceil_{\beta}$ is the concatenate operation of tensors. $\beta \in (0, 1)$



Fig. 2: Deep feature surgery for multi-exit networks. The black arrow represents forward and backward propagation. The blue arrow indicates feature reference. The reddoted line represents the partitioning position. DFS splits the features of each layer into two distinct parts f_i^+ and f_i^- with coefficient β , which reduces the number of shared weights w_i^+ among different exits thus mitigating gradient conflict and reducing backward computation operations. DFS cross-references the shared features and exitspecific features among exits with varying depths in the forward phase while ignoring this in the backward phase, thus using more features for predicting tasks while not introducing more inconsistent gradients.

indicates the partitioning ratio of w_i^+ and $1-\beta$ is that of w_i^- . w_i^+ is shared among exits and w_i^- is exit-specific weights. Based on the partitioned weights, task c_i is computed as follows.

$$c_{i} = \pi_{i}(f_{i}^{-}, w_{c_{i}}), \quad s.t. \begin{cases} f_{i}^{-} = \varphi_{i}(X, \{w_{1}^{+}, w_{2}^{+}, \cdots, w_{i-1}^{+}, w_{i}^{-}\}) \\ i = 1, 2, \cdots, L - 1. \end{cases}$$
(6)

According to the optimization in Eq. (3), the gradients g_{w_i} of w_i can be computed as follows.

$$g_{w_i} = \left[\sum_{k=i+1}^{L} g_{w_i^+}^k, g_{w_i^-}^i\right]_{\beta}, \quad s.t. \ g_{w_i^+}^k = \frac{\partial CE(c_k, Y)}{\partial w_i^+}, \ g_{w_i^-}^i = \frac{\partial CE(c_i, Y)}{\partial w_i^-}. \tag{7}$$

Through decoupling weights into two distinct parts, we alleviate the gradient conflict of $\sum_{k=i}^{L} g_{w_i}^k$ in Eq. (4) to $\sum_{k=i+1}^{L} g_{w_i}^k$ in Eq. (7) as shown in Fig. 2. Notably, $\beta \to 1$ partitions all elements in w_i into w_i^+ , the task c_i is invalid. On the contrary, $\beta \to 0$ partitions all elements in w_i into w_i^- leads to $|w_i^+| \to 0$, the multi-exit network degrades into a cascade of network layers, and the end-to-end joint optimization from input X to tasks is eliminated. We constrain the value of coefficient β to the range of (0, 1) to alleviate the gradient conflict in multi-exit network optimization while ensuring the joint optimization of tasks.

Instead of partitioning the weight tensor, we partition the features for easy deployment in practice. As shown in Fig. 2, taking the first layer as an example, the output feature $f_1 = \varphi_1(X, w_1)$ of the first layer is split into $f_1^+ = \varphi_1(X, w_1^+)$ and $f_1^- = \varphi_1(X, w_1^-)$, which is corresponding to weight partitioning in Eq. (5). The corresponding feature partitions to Eq. (6) is as follows:

$$\begin{cases} f_i^- = \varphi_i(X, \{w_1^+, w_2^+, \cdots, w_{i-1}^+, w_i^-\}) \\ f_i^+ = \varphi_i(X, \{w_1^+, w_2^+, \cdots, w_{i-1}^+, w_i^+\}) \end{cases}, \quad s.t. \; i = 1, 2, \cdots, L-1.$$
(8)

Here, f_i^+ is shared on all the following tasks including $\{c_{i+1}, c_{i+2}, \cdots, c_L\}$, and f_i^- is utilized for predicting c_i . c_i is the only task that contributes to learning w_i^- by supervising the exit-specific feature f_i^- .

4.2 Feature Referencing

Feature partitioning alleviates the gradient conflict in multi-exit network training but reduces the number of features for weight training. For example, f_i^+ is not available for task c_i in Eq. (6). The reduced number of features, e.g., f_i^+ , can easily result in the under-fitting of tasks. We introduce feature referencing to address this issue. Feature referencing reuses these features across tasks to benefit the task performance. Based on feature referencing, task c_i and gradients g_{w_i} of all tasks with respect to weight w_i are defined as follows.

$$\begin{cases} c_i = \pi_i([f_i^+, f_i^-]_\beta, w_{c_i}) \\ f_i^+ = \varphi_i([f_{i-1}^+, f_{i-1}^-]_\beta, w_i^+), \quad s.t. \\ f_i^- = \varphi_i([f_{i-1}^+, f_{i-1}^-]_\beta, w_i^-) \end{cases} \begin{cases} f_0 = X, \\ i = 1, 2, \cdots, L-1 \end{cases}$$
(9)

Here we use both the shared feature f_{i-1}^+ and exit-specific feature f_{i-1}^- for extracting next-layer features f_i^+ and f_i^- in the forward phase with partitioned weights, while both f_i^+ and f_i^- are used for predicting task c_i . This approach increases the involved number of features from the previous output, thus increasing the number of exit-specific weight w_{c_i} for task c_i to improve the accuracy of the task. In the backward phase, to eliminate the gradient conflict, the computation of gradients $\{g_{w_i}^k | k = i+1, \cdots, L\}$ is avoided. As shown in Fig. 2, for the exit-specific weight w_i^- in *i*-th layer, we only compute the gradients $g_{w_i^-}^i$ from the *i*-th loss to ensure the gradient $g_{w_i^-}$ has no conflicts.

Feature referencing uses more features for predicting task c_i while not introducing more inconsistent gradients for updating weight w_i . Moreover, the feature segments f_i^+ , and f_i^- are referenced by different exits with varying depth, and the tasks for these exits supervise the w_i^+ and w_i^- to learn varying features with different scales which improve the accuracy of multi-exit networks [13].

5 Training Efficiency Improvement

The training efficiency of multi-exit network models hinders their wide adoption. DFS not only improves the accuracy of the different exits but also improves the training efficiency of multi-exit network models by reducing the total number of training operations. To simplify the representation, we assume the neural operations (NO) in Fig. 2 only contain matrix-multiply operations to describe the reduction of Multiply-Accumulation operation in the training of multi-exit networks with DFS.

Let the *i*-th neural operation of a vanilla multi-exit network be:

$$\begin{cases} f_i = \varphi_i(f_{i-1} \cdot w_i) \\ l_i = CE(\pi_i(f_i \cdot w_{c_i}), Y) \end{cases}, s.t. \begin{cases} f_0 = X, \\ i = 1, 2 \cdots, L. \end{cases}$$
(10)

8 Cheng Gong et al.

Here f_i and f_{i-1} are the features for layer i and w_i is the corresponding layer weights. There are L tasks with loss set $\{l_i | i = 1, 2, \dots, L\}$, and task losses $\{l_j | j >= i\}$ reference the same intermediate feature f_i , recursively. For simplicity, we assume the matrixes f_i , w_i , and w_{c_i} have the same shape [N, N]. Ignoring the computing operations of activation and loss functions ($\varphi_i(\cdot)$ and $CE(\cdot)$), the operations for computing $f_{i-1} \cdot w_i$ and $f_i \cdot w_{c_i}$ are both $2N^3$, and the total forward operations are $4LN^3$.

We can further derivate the gradients of all task losses with respect to the w_i matrix as follows.

$$g_{w_i} = \frac{\partial l_i}{\partial w_i} + \frac{\partial l_{i+1}}{\partial w_i} + \dots + \frac{\partial l_L}{\partial w_i}$$

=
$$\sum_{k=i}^{L} \left(\frac{\partial l_k}{\partial f_k} \prod_{j=i}^{j < k} \left(\frac{\partial f_{j+1}}{\partial f_j}\right)\right) \cdot w_i^T \ s.t. \ i = 1, 2, \dots, L.$$
 (11)

Because of the chain rules and back-propagation algorithms [28], it only requires to compute L-1 times $\partial f_{j+1}/\partial f_j$, L times $\partial l_k/\partial f_k$, and one $\partial f_i/\partial w_i$ to obtain g_{w_i} . The operations for computing gradients g_{w_i} are $4LN^3$. Since there are total 2L trainable weights (w_i and w_{c_i}), and the operations for computing $\partial f_i/\partial w_i$ and $\partial l_i/\partial w_{c_i}$ are both $2N^3$ (one matrix multiplication), the total backward operations are $(8L-2)N^3$, and the total forward-backward operations are $(12L-2)N^3$ for one training step for a vanilla multi-exit network.

Assuming that N can be evenly divided by $1/\beta$, then the *i*-th feature f_i is partitioned into $f_i^+ \in \mathbb{R}^{N*(\beta N)}$ and $f_i^- \in \mathbb{R}^{N*(1-\beta)N}$, and the neural operations are rewritten as follows:

$$f_i = [f_i^+, f_i^-]_\beta = [\varphi_i(f_{i-1} \cdot w_1^+), \varphi_i(f_{i-1} \cdot w_1^-)]_\beta \ s.t. \ f_0 = X, i = 1, 2, \cdots, L.$$
(12)

During the backward process, as described in Eq. (9) in Section 4.2, only the gradients from an exit-specific loss l_i are adopted for weight w_i^- , which have no conflicts, and all gradients from all losses are adopted to update w_i^+ for end-toend optimization. So the computation of g_{w_i} in Eq. (7) is expanded as follows.

$$\begin{cases} g_{w_{i}^{+}} = \sum_{k=i+1}^{L} g_{w_{i}^{+}}^{k} = g_{w_{i}^{+}}^{L} + \sum_{k=i+1}^{L-1} \left(\frac{\partial l_{k}}{\partial f_{k}^{-}} \frac{\partial f_{k}^{-}}{\partial f_{k-1}^{+}} \Pi_{j=i}^{j < k-1} \left(\frac{\partial f_{j+1}^{+}}{\partial f_{j}^{+}}\right)\right) \frac{\partial f_{i}^{+}}{\partial w_{i}^{+}} \\ g_{w_{i}^{-}} = g_{w_{i}^{-}}^{i} = \frac{\partial l_{i}}{\partial f_{i}^{-}} \cdot \frac{\partial f_{i}^{-}}{\partial w_{i}^{-}} = \frac{\partial l_{i}}{\partial f_{i}^{-}} \cdot (w_{i}^{-})^{T}. \end{cases}$$
(13)

The feature number of f_i^+ and f_i^- is β and $(1 - \beta)$ times that of f_i . Therefore, the operations for computing $\partial l_i / \partial f_i^+$ and $\partial l_i / \partial f_i^-$ are also β and $(1 - \beta)$ times that for $\partial l_i / \partial f_i$, and they are $\beta \cdot 2N^3$ and $(1 - \beta) \cdot 2N^3$, respectively. The operations to compute $\partial f_k^- / \partial f_{k-1}^+$ and $\partial f_{j+1}^+ / \partial f_j^+$ are $2\beta(1 - \beta)N^3$ and $2\beta^2N^3$, respectively. Employing the back-propagation algorithm [28], we only need to compute L times of $\partial l_k / \partial f_k^-$, L - 2 times of $\partial f_k^- / \partial f_{k-1}^+$, L - 2 times of $\partial f_{j+1}^+ / \partial f_j^+$ and one time of $\partial l_L / \partial f_L$ and $\partial f_L / \partial f_{L-1}^+$ to compute $g_{w_i^+}^L$), and one time of $\partial f_i^+ / \partial w_i^+$ and $\partial f_i^- / \partial w_i^-$ in Eq. (13). The operations for computing $g_{w_i} = [g_{w_i^+}, g_{w_i^-}]_{\beta}$ is $2(L+1)N^3$, which is irrelated to coefficient β . There are a Deep Feature Surgery: Towards Accurate and Efficient Multi-Exit Networks

total number of 2L trainable weights (w_i and w_{c_i}), and the operations for computing $\left[\frac{\partial f_i^+}{\partial w_i^+}, \frac{\partial f_i^-}{\partial w_i^-}\right]_{\beta}$ and $\frac{\partial l_i}{\partial w_{c_i}}$ are both $2N^3$, the total backward operations are $6LN^3$. The total forward-backward operations are $10LN^3$ for one training step with DFS. Compared with one vanilla training step, applying DFS reduces the operations by up to 16.67%, which is formulated as follows and can be observed that it is irrelated to the value of β .

$$1 - \frac{10LN^3}{(12L-2)N^3} \in [0\%, 16.67\%), \ s.t. \ L \in [1, +\infty).$$
(14)

When $\beta \to 0$, the multi-exit networks degrade to a cascade of shallow networks with total training operations of $10LN^3$. When $\beta \to 1$, the multi-exit networks degrade to backbone networks with multiple independent tasks, and the total training operations are $10LN^3$. In conclusion, for a typical multi-exit network, employing DFS can save up to 16.67% of operations in model training.

6 Experiments

In this section, we conduct experiments to evaluate the accuracy and training efficiency improvement of multi-exit networks with DFS.

6.1 Model Accuracy Evaluation

We use Deeply-Supervised Nets (DSN) [17], BYOT [35], Gradient Equilibrium (GE) [18], MSDNet [13], PCGrad [34], CAGrad [21], Meta-GF [32], Block-Dependent Loss (BDL) [8], WPN [10], Nash-MTL [25], and DR-MGF [33] as the baseline methods. All accuracy results of baselines are cited from the original papers. We use the same datasets and models used in the baselines and adopt DFS to the model training framework. Specifically, the datasets are Cifar100 [15] and ImageNet [3]. The data augmentation policies for Cifar100 and ImageNet refer to [35] and [6], respectively. The models on these datasets are MSDNet [13] models and ResNet18 [11] and their implementations in [18] and [35] are utilized for reproducibility and fair comparison. $\beta = 0.5$ is used as the default setting for DFS. For Cifar100 dataset, we use the stochastic gradient descent (SGD) optimizer with a momentum of 0.9 for all models. We train all models from scratch for a total of 300 epochs with an initial learning rate of 0.1 and drops to 0.01, 1e-3, and 1e-4, after 250, 280, and 295 epochs, a weight decay of $5e^{-4}$, and a batch size of 500. For ImageNet dataset, we use the SGD optimizer with a batch size of 256, and train MSDNet models from scratch for a total of 90 epochs, with an initial learning rate of 0.1 and drops to 0.01, 1e-3, and 1e-4, after 70, 80, and 85 epochs. Besides, we do not employ extra training tricks, such as knowledge distillation strategies [13, 35], gradient selection strategies [25, 32, 34], and weighted losses or samples [10, 33] in this evaluation. The top-1 accuracy is used as the metric for accuracy evaluation for each of the exits.

The results of MSDNet are cited from [18], and the results of PCGrad and CAGrad for multi-exit network training are cited from [32].

Tuble 1. comparison results of restretto on charito.							
Methods	Exit1	Exit2	Exit3	Exit4	Ensemble		
Params (M)	0.43	0.96	3.11	11.17	-		
FLOPs (M)	169.62	300.22	431.05	559.75	-		
DSN [17]	67.23	73.80	77.75	78.38	79.27		
BYOT [35]	67.85	74.57	78.23	78.64	79.67		
DFS	74.17	77.43	78.82	79.93	80.89		

Table 1: Comparison results of ResNet18 on Cifar100.

*Ensemble is the accuracy of ensembled logits from all tasks [35].

Table 2: Comparison results of MSDNet on Cifar100.								
Methods	Exit1	Exit2	Exit3	Exit4	Exit5	Exit6	Exit7	
Params (M)	0.90	1.84	2.80	3.76	4.92	6.10	7.36	
FLOPs (M)	56.43	101.00	155.31	198.10	249.53	298.05	340.64	
MSDNet [13]	64.10	67.46	70.34	72.38	73.06	73.81	73.89	
GE [18]	64.00	68.41	71.86	73.50	74.46	75.39	75.96	
CAGrad [21]	68.78	72.55	74.23	74.97	75.35	75.82	76.08	
PCGrad [34]	67.06	71.37	74.86	75.78	76.25	76.95	76.71	
Meta-GF [32]	67.97	72.27	75.06	75.77	76.38	77.11	77.47	
BDL-1 [8]	66.29	69.31	71.15	72.05	72.61	73.23	73.59	
BDL-2 [8]	65.81	68.93	71.01	72.45	72.98	73.42	74.27	
WPN [10]	62.26	67.18	70.53	73.10	74.80	76.05	76.31	
Nash-MTL $[25]$	64.14	69.23	72.64	74.89	75.32	75.88	76.75	
DR-MGF [33]	67.82	72.45	74.77	75.77	76.56	76.90	77.01	
DFS	72.46	75.64	77.50	78.54	79.2	79.45	79.56	

Results of ResNet18 on Cifar100. The experimental results for ResNet18 on Cifar100 dataset are shown in Tab. 1. Compared with DSN [17] and BYOT [35], DFS achieves higher accuracy across all exits by mitigating gradient conflict and reusing the multi-scale features of different exits. For early exits, such as exit 1 and exit 2, the accuracy improvements are up to **6.94%** and **3.63%**, respectively. The ensemble accuracy results of DFS also outperform the same counterparts from DSN and BYOT by **1.62%** and **1.22%**, respectively. The results show that applying DFS largely improves the accuracy of the exits, especially the accuracy of early exits. This is because the early exits share all parameters with all exits, which are more susceptible to the impact of gradient conflict, thus their accuracy results often decrease more when compared to single-task models and deep tasks. DFS mitigates the gradient conflict in training, leading to a significant improvement in the performance of early exits.

Results of MSDNet on Cifar100. As shown in Tab. 2, the results demonstrate that DFS consistently outperforms all baselines across all exits of MS-DNet, and significantly improves the overall performance. Compared with the baselines, the average accuracy improvements of DFS on the 7 exits are 6.64%, 5.72%, 4.86%, 4.47%, 4.42%, 3.99%, and 3.76%, respectively. It can be observed that the accuracy improvements of early exits are more significant than those of deep ones, which is because of the mitigation of the gradient conflict that

Table 5. comparison results of hisbride on imageree.							
Methods	Exit1	Exit2	Exit3	Exit4	Exit5	Average	
Params (M)	4.24	8.77	13.07	16.75	23.96	-	
FLOPs (G)	0.34	0.69	1.01	1.25	1.36	-	
MSDNet [13]	56.64	65.14	68.42	69.77	71.34	66.26	
GE [18]	57.28	66.22	70.24	71.71	72.43	67.58	
CAGrad [21]	58.37	64.21	66.88	68.22	69.42	65.42	
PCGrad [34]	57.62	64.87	68.93	71.05	72.45	66.98	
Mate-GF $[32]$	57.43	64.82	69.08	71.67	73.27	67.25	
DFS	61.80	68.03	70.75	71.79	72.88	69.05	

Table 3: Comparison results of MSDNet on ImageNet.

Average is the average accuracy of 5 tasks [32].

improves the performance of early tasks, as described in Sec. 6.1. The average accuracy improvement of DFS across all tasks and all baselines is 4.84%.

Results of MSDNet on ImageNet. The results are shown in Tab. 3. Results show that DFS outperforms all the baselines at the early exits even without extra training tricks, and the average accuracy improvements of DFS on the 5 tasks are 4.33%, 2.98%, 2.04%, 1.31%, 1.10%, respectively. Similar to the experiments on Cifar100, the accuracy improvements on early exits are more significant than those of deep ones, which further demonstrates the effectiveness of DFS on conflict mitigation. The average accuracy achieved by DFS is 69.05%, which significantly exceeds baselines, and the average accuracy improvement achieves 2.35%. One exception happens at the fifth exit, the accuracy of DFS is slightly lower than that of Mate-GF [32]. The main reason is that DFS relies on the over-parameterization characteristics of models for higher accuracy while the model size of MSDNet is relatively small, especially when training on the ImageNet dataset. It hinders the accuracy improvement of DFS for the later tasks that have greater differences in model capacity compared to early tasks. Besides, Mate-GF employed a cautious while time-consuming gradient merge method to achieve high accuracy, while DFS do not employ any additional methods for high training efficiency, which is another cause for the results.

Budgeted Batch Classification. We evaluate the achieved classification accuracy of the model within a given computational budget [18, 32]. We conduct experiments on Cifar100 dataset using MSDNet [13] model. The baseline methods include MSD-Net [13], GE [18], and Meta-GF [32]. The results are shown in Fig. 3. The results show that DFS consistently outperforms all baselines across all budgets. With an average budget of 170MFLOPs, DFS achieves an accu-



Fig. 3: Accuracy of budgeted batch classification. The X-axis is the average computational budget per image for MSDNet model on the Cifar100, and Y-axis top-1 accuracy.



Fig. 4: Training efficiency comparison. Y1-axis is the average latency (divided into forward, backward, and other parts) for one training step, Y2-axis is the average FPS, and X-axis is the results of DSN, BYOT, and DFS under different batch sizes.

racy of $\sim 79.5\%$, which is $\sim 2.5\%$ higher than that of Meta-GF with the same budget. DFS uses $\sim 2 \times$ fewer FLOPs to achieve the same classification accuracy compared to Meta-GF, which is the best among all baselines.

6.2 Training Efficiency Validation

In this experiment, we evaluate the training efficiency of DFS. Specifically, the computation reduction during the backpropagation when adopting DFS contributes to efficient training of multi-exit networks. We select $\beta = 0.5$ as the default setting for DFS.

We adopt the VGG7-64 in [6] with 6 convolution layers and train it on Cifar100 [15] dataset, and VGG16 [30] with 13 convolution layers and train it on ImageNet [16] dataset to evaluate the training efficiency improvement from DFS. Both models employ a Batch Normalization (BN) [14] layer after each convolution layer and the early classifiers are added after the BNs. Each classifier only contains a Global Average Pooling (GAP) [20] layer and a Fully-Connected (FC) layer for classification. The baseline methods include DSN [17] and BYOT [35]. DSN applies the vanilla training in Sec. 5 which optimizes all losses with the same weight and updates the trainable parameters. BYOT introduces self-distillation supervision based on DSN, which employs the outputs and logits of the final exit to teach the early tasks. We use the latency for training one batch of samples and Frames Per Second (FPS) on the same devices to evaluate the training efficiency. The overall training latency is divided into forward latency, backward latency, and other latency for detailed comparison. The other latency includes data loading, loss computation, parameter update, etc. The lower latency and higher FPS stand for a higher training efficiency. For a fair comparison, the results are collected with different batch sizes. We train VGG7-64 on CPU $(2 \times$ 20-core Intel Xeon Gold 6248) and VGG16 on GPU $(1 \times \text{NVIDIA RTX A6000})$, both are trained with 100 batches of data inputs.

Methods	Exit1	Exit2	Exit3	Exit4	Exit5	Exit6
Params (K)	8.32	45.31	125.70	273.41	581.63	1171.97
FLOPs (M)	2.04	40.05	59.06	96.94	115.89	153.71
DFS $\beta = 0.1$	22.81	41.56	53.93	58.58	63.87	65.46
DFS $\beta = 0.3$	21.23	40.34	55.53	61.31	67.11	69.44
DFS $\beta = 0.5$	20.77	39.21	55.31	61.83	68.61	70.56
DFS $\beta = 0.7$	19.62	36.42	52.96	60.95	69.04	72.17
DFS $\beta = 0.9$	17.88	32.24	49.69	57.12	68.13	73.06

Table 4: The accuracy comparison of multiple tasks in the VGG7-64 model with different coefficient values.

The experimental results are shown in Fig. 4. For different batch sizes and devices, all the DSN, BYOT, and DFS have similar forward latency, since their forward operations are the same. For the backward latency, DFS has shorter latency compared with DSN and BYOT across all batch sizes and devices mainly because DFS reduces the required number of operations for backward computation. BYOT uses knowledge distillation and intermediate feature distillation to improve model accuracy, resulting in much higher backward latency and other latency than DSN and DFS.

The FPS results of VGG7-64 with DFS outperforms DSN and BYOT by up to **14.43%** (76.67*fps*) and **50.00%** (202.70*fps*), and the average FPS improvements across batch sizes are **11.69%** (62.43*fps*) and **40.77%** (172.77*fps*), respectively. On the GPU device with VGG16 as backbone, the training FPS results of DFS exceed DSN and BYOT by up to **14.99%** (25.48*fps*) and **22.23%** (35.55*fps*), and the average FPS improvements across batch sizes are **13.51%** (22.28*fps*) and **21.27%** (32.84*fps*), respectively. The results demonstrate that DFS reduces the backward latency significantly during model training and improves FPS across different models, batch sizes, and devices. The actual average FPS improvements obtained through experiments are **11.69%** and **13.51%**, both of them are consistent with the theoretical FPS improvements of VGG7-64 and VGG16 according to Eq. (14), which are **14.29%** with L = 6 (convolutional layers of VGG7-64) and **15.58%** with L = 13 (convolutional layers of VGG16).

6.3 Impact of Coefficient

We evaluate the impact of the coefficient on the accuracy and training efficiency of DFS in this experiment. We select VGG7-64 on Cifar100 with batch size 500 to evaluate the impact of β on the model accuracy and training latency (FPS).

The accuracy of different exits with different coefficients is shown in Tab. 4 and the normalized accuracy with training latency (FPS) is shown in Fig. 5. The normalized accuracy is normalized to the maximum accuracy of different methods on each task to show the impact of β . The results show that a lower coefficient value results in a higher accuracy for early exits while a lower accuracy for later exits. On the contrary, the higher β improves the accuracy of later exits while decreasing that of early exits. In summary, **lower** β **contributes to early exits while higher** β **benefits later exits**. This is because the



Fig. 5: The impact of the coefficient on task accuracy and training efficiency. The left side shows the normalized accuracy of DFS with different β . It focuses on the sort instead of the absolute value of accuracy results on each task. The right side is the training efficiency of VGG7-64 using DFS with different β . Y1-axis is the average latency and y2-axis is the average FPS metric.

lower β assigns more exit-specific features and decreases the shared features. When decreasing β , the number of trainable weights for later exits decreases while the number of trainable weights for early tasks increases. For exits with different depths, the best coefficient values are different. As shown in Fig. 5, an optimal coefficient value improves the training efficiency of multi-exit networks. According to Eq. (13), increasing β will increase the portion of shared weights, thus increasing the computational load of $g_{w_i^+}$, while decreasing β increases the computational load for $g_{w_i^-}$. In our experiments, $\beta = 0.5$ offers a balanced computational load among multiple tasks and leads to higher training efficiency.

7 Conclusion

In this paper, we proposed a novel method DFS to mitigate the gradient conflict in the training of multi-exit networks. Specifically, DFS consists of feature partitioning and feature referencing techniques. The former decouples the shared weights among multiple exits and mitigates the gradient conflict in model training and ensures an end-to-end joint optimization for each exit. The latter reuses the multi-scale features among exits and enhances the overall performance of multi-exit networks. DFS improves the accuracy by up to 6.94% of the individual exits and up to 3.63% for the overall ensembled model accuracy. In addition, we provide a detailed analysis of the operation reduction for backpropagation by applying DFS and the experimental results show that DFS saves up to 50.00% training time in the training of multi-exit networks and still provides improved model accuracy. Deep Feature Surgery: Towards Accurate and Efficient Multi-Exit Networks

15

8 Acknowledgments

This work is partially supported by the China Postdoctoral Science Foundation (2022M721707), the National Natural Science Foundation (62272248), the Natural Science Foundation of Tianjin (23JCZDJC01010, 23JCQNJC00010), the National Natural Science Foundation (62372253), the Natural Science Foundation of Tianjin Fund (23JCYBJC00010), the CCF-Baidu Open Fund (CCF-Baidu202310), the Open Project Fund of State Key Laboratory of Computer Architecture, ICT, CAS (CARCHB202016), the Innovation Fund of Qiyuan Lab (2022-JCJ0-LA-001-068), Open Fund of Civil Aviation Smart Airport Theory and System Key laboratory, Civil Aviation University of China (SATS202303), the Key Program for International Cooperation of Ministry of Science and Technology, China (2024YFE0100700).

References

- Chai, H., Cui, J., Wang, Y., Zhang, M., Fang, B., Liao, Q.: Improving gradient trade-offs between tasks in multi-task text classification. In: ACL. pp. 2565–2579 (2023) 2, 3
- Chen, Z., Badrinarayanan, V., Lee, C.Y., Rabinovich, A.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In: ICML. pp. 794–803. PMLR (2018) 2, 3
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE CVPR. pp. 248–255 (2009) 9
- Deng, W., Suh, Y., Yu, X., Faraki, M., Zheng, L., Chandraker, M.: Split to learn: Gradient split for multi-task human image analysis. In: WCACV. pp. 4351–4360 (2023) 2, 4
- Ding, C., Lu, Z., Wang, S., Cheng, R., Boddeti, V.N.: Mitigating task interference in multi-task learning via explicit task routing with non-learnable primitives. In: IEEE CVPR. pp. 7756–7765 (2023) 4
- Gong, C., Lu, Y., Xie, K., Jin, Z., Li, T., Wang, Y.: Elastic significant bit quantization and acceleration for deep neural networks. IEEE Transactions on Parallel and Distributed Systems (2021). https://doi.org/10.1109/TPDS.2021.3129615 9, 12
- Guo, M., Haque, A., Huang, D.A., Yeung, S., Fei-Fei, L.: Dynamic task prioritization for multitask learning. In: ECCV. pp. 270–287 (2018) 3
- Han, D.J., Park, J., Ham, S., Lee, N., Moon, J.: Improving low-latency predictions in multi-exit neural networks via block-dependent losses. TNNLS (2023) 9, 10
- Han, Y., Huang, G., Song, S., Yang, L., Wang, H., Wang, Y.: Dynamic neural networks: A survey. IEEE TPAMI 44(11), 7436–7456 (2022) 2
- Han, Y., Pu, Y., Lai, Z., Wang, C., Song, S., Cao, J., Huang, W., Deng, C., Huang, G.: Learning to weight samples for dynamic early-exiting networks. In: ECCV. pp. 362–378. Springer (2022) 2, 3, 9, 10
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE CVPR (2016) 9
- Hu, Y., Chen, C., Li, R., Zhu, Q., Chng, E.S.: Gradient remedy for multi-task learning in end-to-end noise-robust speech recognition. In: ICASSP. pp. 1–5. IEEE (2023) 2, 3

- 16 Cheng Gong et al.
- 13. Huang, G., Chen, D.: Multi-scale dense networks for resource efficient image classification. ICLR (2018) 3, 5, 7, 9, 10, 11
- Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML. pp. 448–456 (2015) 12
- Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009) 9, 12
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. pp. 1097–1105 (2012) 12
- Lee, C., Xie, S., Gallagher, P.W., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: AISTATS. vol. 38. JMLR.org (2015) 9, 10, 12
- Li, H., Zhang, H., Qi, X., Yang, R., Huang, G.: Improved techniques for training adaptive deep networks. In: IEEE ICCV. pp. 1891–1900 (2019) 2, 3, 9, 10, 11
- Lin, L., Liao, X., Jin, H., Li, P.: Computation offloading toward edge computing. Proceedings of the IEEE 107(8), 1584–1607 (2019) 2
- Lin, M., Chen, Q., Yan, S.: Network in network. arXiv preprint arXiv:1312.4400 (2013) 12
- Liu, B., Liu, X., Jin, X., Stone, P., Liu, Q.: Conflict-averse gradient descent for multi-task learning. NIPS 34, 18878–18890 (2021) 2, 3, 9, 10, 11
- Liu, L., Li, Y., Kuang, Z., Xue, J., Chen, Y., Yang, W., Liao, Q., Zhang, W.: Towards impartial multi-task learning. In: ICLR (2021) 3
- Liu, S., Johns, E., Davison, A.J.: End-to-end multi-task learning with attention. In: IEEE CVPR. pp. 1871–1880 (2019) 4
- Maninis, K.K., Radosavovic, I., Kokkinos, I.: Attentive single-tasking of multiple tasks. In: IEEE CVPR. pp. 1851–1860 (2019) 4
- Navon, A., Shamsian, A., Achituve, I., Maron, H., Kawaguchi, K., Chechik, G., Fetaya, E.: Multi-task learning as a bargaining game. In: ICML. pp. 16428–16446. PMLR (2022) 2, 3, 9, 10
- Pascal, L., Michiardi, P., Bost, X., Huet, B., Zuluaga, M.A.: Maximum roaming multi-task learning. In: AAAI. vol. 35, pp. 9331–9341 (2021) 4
- Phuong, M., Lampert, C.H.: Distillation-based training for multi-exit architectures. In: IEEE ICCV. pp. 1355–1364 (2019) 2, 3
- Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by backpropagating errors. nature 323(6088), 533–536 (1986)
- Sener, O., Koltun, V.: Multi-task learning as multi-objective optimization. NIPS 31 (2018) 3, 5
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv (2014) 12
- Strezoski, G., Noord, N.v., Worring, M.: Many task learning with task routing. In: IEEE ICCV. pp. 1375–1384 (2019) 4
- Sun, Y., Li, J., Xu, X.: Meta-gf: Training dynamic-depth neural networks harmoniously. In: ECCV. pp. 691–708. Springer (2022) 2, 3, 9, 10, 11
- 33. Sun, Y., Xu, X., Li, J., Hu, X., Shi, Y., Zeng, L.L.: Learning task-preferred inference routes for gradient de-conflict in multi-output dnns. arXiv preprint arXiv:2305.19844 (2023) 2, 9, 10
- Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., Finn, C.: Gradient surgery for multi-task learning. NIPS 33, 5824–5836 (2020) 2, 3, 9, 10, 11
- 35. Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., Ma, K.: Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In: IEEE ICCV. pp. 3713–3722 (2019) 2, 3, 9, 10, 12
- Zhang, L., Tan, Z., Song, J., Chen, J., Bao, C., Ma, K.: Scan: A scalable neural networks framework towards compact and efficient models. NIPS 32 (2019) 2, 3