

How to Train Your Multi-Exit Model?

Analyzing the Impact of Training Strategies

Piotr Kubaty^{*1} Bartosz Wójcik^{*†1} Bartłomiej Krzepakowski² Monika Michaluk³ Tomasz Trzcinski^{2,4,5}
Jary Pomponi⁶ Kamil Adamczewski^{†7,5}

Abstract

Early exits enable the network’s forward pass to terminate early by attaching trainable internal classifiers to the backbone network. Existing early-exit methods typically adopt either a joint training approach, where the backbone and exit heads are trained simultaneously, or a disjoint approach, where the heads are trained separately. However, the implications of this choice are often overlooked, with studies typically adopting one approach without adequate justification. This choice influences training dynamics and its impact remains largely unexplored. In this paper, we introduce a set of metrics to analyze early-exit training dynamics and guide the choice of training strategy. We demonstrate that conventionally used joint and disjoint regimes yield suboptimal performance. To address these limitations, we propose a mixed training strategy: the backbone is trained first, followed by the training of the entire multi-exit network. Through comprehensive evaluations of training strategies across various architectures, datasets, and early-exit methods, we present the strengths and weaknesses of the early exit training strategies. In particular, we show consistent improvements in performance and efficiency using the proposed mixed strategy.

1. Introduction

Deep neural networks have achieved remarkable results across a variety of machine learning tasks. While the depth

^{*}Equal contribution. [†]Equal supervision. ¹Jagiellonian University ²Warsaw University of Technology ³University of Warsaw ⁴Tooploox ⁵IDEAS Research Institute ⁶Department of Information Engineering, Electronics, and Telecommunications (DIET) at Sapienza, University of Rome, Italy ⁷Wroclaw University of Science and Technology. Correspondence to: Bartosz Wójcik <bartwojc@gmail.com>.

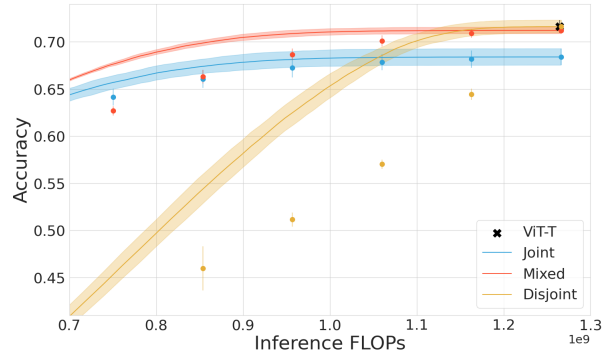


Figure 1. Performance-cost trade-off of the multi-exit ViT-T model trained on the ImageNet-1k dataset by using different training regimes. The choice of the training regime significantly impacts the performance across all computational budgets.

of these networks significantly contributes to their enhanced performance, the necessity of using large models for all inputs, especially in resource-constrained environments like mobile and edge computing devices, is questionable.

Early exit methods for deep neural networks have gained importance due to their potential to significantly improve computational efficiency. By exiting at earlier layers, these methods can decrease the number of operations needed for computation of the forward pass, leading to faster inference times. In doing so they allow the network to adapt its computational cost to the difficulty of the input sample. Simpler inputs can be processed with fewer layers, while more complex inputs can utilize the full capacity of the network.

Early exit methods are implemented through augmentation of the original architecture with internal classifiers (ICs) attached to selected intermediate layers (Kaya et al., 2019). These ICs are designed to perform classification tasks based on the representations available at their respective positions in the network. A common approach for training early-exit models involves training the entire multi-exit network, including the added classifiers, from scratch (Huang et al., 2018; Yang et al., 2020; Meronen et al., 2024) (“joint” regime). Alternatively, some methods train the backbone

network first, then freeze its weights and train the parameters of the newly added ICs in the second, separate phase of training (Teerapittayanon et al., 2016; Liao et al., 2021; Zhou et al., 2020) ("**disjoint**" regime). To the best of our knowledge, no study compares or explores the relationship between these training regimes.

In this study, we perform an extensive assessment of early-exit regimes and notice the choice of training strategy has a significant impact on the final model’s performance. We identify the relationship between computational budget and the choice of the regime. Using the disjoint regime results with a network that is significantly impaired when smaller computational budget is assumed. While the joint regime might initially seem as the appropriate way of training multi-exit networks, we demonstrate that due to its training dynamics it biases the network and produces a model with subpar performance on higher computational budgets.

In order to address the weaknesses of multi-exit networks, we propose a novel "**mixed**" regime: train the backbone network until convergence, then train the entire model jointly, including the internal classifiers, until convergence. This approach ensures that the backbone architecture is adequately trained before optimizing it alongside the internal classifiers for improved performance.

To gain a deeper understanding of learning and optimization in multi-exit architectures, we propose a set of metrics to describe the training dynamics of early-exit models trained under various regimes. Through gradient dominance metric we reveal the set of ICs that have the largest impact on the backbone during training and explain the performance gains of popular technique of loss and gradient scaling depend on the training strategy choices. Furthermore, mode connectivity allows to analyze solution similarities of different regimes. Moreover, numerical rank, and mutual information metrics aid to explain the performance of regimes under different computational budgets.

Finally, we provide a thorough empirical evaluation of early-exit regimes across different network architectures, data modalities, datasets and early-exit methods. Our results show that proposed alternative strategy enables significant improvements in performance in medium and high budgets over the commonly used joint training.

2. Training Regimes

Early exit methods fundamentally alter the organization of neural networks. It is widely believed that neural networks develop a hierarchical representation of features, where earlier layers learn basic shapes and patterns, while later layers progressively capture more complex abstractions (Zeiler & Fergus, 2014). In other words, the earlier layers are characterized by higher frequency features while later layers learn

low frequency elements. This regularity is disrupted in the case of early exit architectures as the backbone network is given additional classifiers that are placed in earlier parts of the network. These changes in architecture require a different approach for training and more nuanced analysis how the training should proceed.

In the early-exit setting, parameters can be divided into *backbone* parameters and *internal classifier (IC)* parameters. Each of these two groups can be trained separately or jointly. In this paper, we frame the training process of any early-exit method as consisting of three following phases:

Phase 1: Train the backbone network parameters θ_b by minimizing the loss at the final output layer (could be the last IC or an added final classifier).

$$\theta_b^* = \arg \min_{\theta_b} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}} \left[\mathcal{L}^{(K)}(\theta_b, \theta_{IC}^{(K)}) \right] \quad (1)$$

During this phase, θ_{IC} are either not trained or not present at all.

Phase 2: Train both the backbone network and the ICs simultaneously.

$$\theta^* = \arg \min_{\theta} \sum_{k=1}^K \alpha_k \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}} \left[\mathcal{L}^{(k)}(\theta_b, \theta_{IC}^{(k)}) \right] \quad (2)$$

Phase 3: Freeze θ_b^* and train only the IC parameters θ_{IC} .

$$\theta_{IC}^* = \arg \min_{\theta_{IC}} \sum_{k=1}^K \alpha_k \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}} \left[\mathcal{L}^{(k)}(\theta_b^*, \theta_{IC}^{(k)}) \right] \quad (3)$$

Correspondingly, we generalize the early exit training regimes into three types based on which of the phases are used:

Disjoint training (Phases 1+3). The model parameters undergo training during the first and third phases. That is, the backbone architecture is trained first, and then the ICs are trained separately while the backbone parameters are frozen.

Joint training (Phase 2). The training consists only of the second phase in which the entire model – including the ICs – is trained. It is currently the most common way of training early-exit models (Matsubara et al., 2022).

Mixed training (Phases 1+2). The training consists of two phases. The backbone is trained in isolation first, and then the entire network, including the ICs, is trained jointly. The regime emphasizes the importance of backbone pre-training as a better way to initialize the architecture for further training. This is our proposed way to improve multi-exit model training.

3. A Framework for Analyzing Multi-Exit Models

This section proposes a framework for analysing and comparing multiple early exit models. This framework comprises multiple evaluations, each focusing on a different aspect of a model.

3.1. Gradient Dominance

The use of internal classifiers during training in early-exit training regime fundamentally alters the training dynamics, as these classifiers contribute to the overall loss. The gradient update now comes from multiple classifiers instead of just the final one, as in a standard neural network.

To study this phenomenon we calculate which gradient is more prominent during training. We define *gradient dominance* (GD) as the cosine similarity between the gradient from an individual internal classifier, \mathbf{g}_i , and the overall gradient of the model $\mathbf{g}_{\text{total}}$:

$$\text{GD}_i = \langle \mathbf{g}_i, \mathbf{g}_{\text{total}} \rangle$$

where $\langle \cdot, \cdot \rangle$ is the cosine similarity. Gradient Dominance measures the consistency of the gradient directions produced by the early-exit classifiers and evaluates how well gradients from separate classifiers align with the overall gradient across the entire model. If the cosine similarity is close to 1, the auxiliary classifier’s gradient is highly aligned with the total gradient, indicating that it potentially dominates other ICs in its impact on the total gradient.

3.2. Mode Connectivity

Mode connectivity theory suggests that independently trained models often exhibit similar characteristics. Notably, after training two independent models, it is possible to find a continuous path in the parameter space where the loss remains low, enabling the models to be connected without encountering high-loss regions (Garipov et al., 2018).

Building on the observation that independently trained neural networks can be linearly connected in weight space after accounting for permutation symmetries, as described in (Ainsworth et al.), we extend this idea to early-exit architectures trained under different regimes.

Instead of focusing solely on independently trained networks, we investigate early-exit architectures trained in distinct regimes. To align models for mode connectivity, we define an optimal permutation Π^* by solving $\Pi^* = \arg \min_{\Pi} \|\Theta_A - \Pi\Theta_B\|_F$ where $\|\cdot\|_F$ is the Frobenius norm, and Θ_A and Θ_B are the parameters of two early-exit models trained under different strategies as described in Sec. 2. Given the optimal permutation, the permuted interpolation path is:

$$\Theta(\lambda) = (1 - \lambda)\Theta_A + \lambda\Pi^*\Theta_B \quad (4)$$

The core idea is based on the fact that, if two models trained with different strategy can be interpolated such that the resulting model has low loss over the whole interpolation path, then the regime solutions lie in the similar loss basin.

3.3. Numerical Rank

So far, the proposed tools analyse the gradients of the model or their final predictions. Here, we propose to evaluate the expressiveness of a model by analysing its intermediate activations, using the numerical ranks of activation maps (Masarczyk et al., 2024). Mathematically, for a given layer i , the rank is evaluated as:

$$r_i = \text{Rank}(A_i), \quad A_i \in \mathbb{R}^{n \times m} \quad (5)$$

where A_i is the activation matrix of dimensions n (number of samples) and m (number of features).

The rank of the internal representations associated with different layers can provide insight into the “expressiveness”, or capacity of the network. A higher rank (closer to the maximum possible for a given layer’s matrix dimensions) indicates that the layer can capture more complex patterns or features in the data, as it implies a greater degree of linear independence among the feature detectors in that layer. High-rank activations matrices in a network suggest that the network is utilizing its capacity to learn diverse, high-frequency features, whereas a low rank might indicate that the network is not fully exploiting its potential.

3.4. Mutual Information

Adding intermediate early exits to a model could alter the information flow within the model itself. To study how much such exits affect this aspect, we use the concept of mutual information. In the context of neural networks, the mutual information between X and Z represents how much information the input X provides about the internal representation Z after passing through a neural network. For random variables X and Z , the mutual information is defined as: $I(X; Z) = \int_{x \in \mathcal{X}} \int_{z \in \mathcal{Z}} p(x, z) \log \frac{p(x, z)}{p(x)p(z)} dx dz$ where $p(x, z)$ is the joint probability distribution of X and Z , and $p(x)$ and $p(z)$ are the marginal distributions of X and Z , respectively. In practical terms, for neural networks, we use Monte Carlo sampling to estimate $I(X; Z)$ due to the high dimensionality of feature spaces, and as such the results are not always monotonically decreasing (Shwartz-Ziv & Tishby, 2017).

In their work, (Kawaguchi et al., 2023) utilize the concept of mutual information between X and Z to study the information bottleneck (IB) principle. It aims to find a balance




Figure 2. Gradient dominance for different regimes. Each line indicates how well gradients from different ICs align with the total gradient over the course of the training. The last IC dominates the most in the mixed regime, which explains its excellent performance on higher computational budgets (ResNet-50, Tinyimagenet).

between the information carried by Z in predicting the target Y and its complexity in terms of its mutual information with the input X . Specifically, minimizing $I(X; Z)$ reduces the complexity and overfitting by ensuring Z retains only the essential information from X , and maximizing $I(Y; Z)$ ensures that the representation Z is informative enough to predict the target variable Y effectively. To visually show how mutual information is affected by the training regime, we calculate the mutual information for each layer averaged over multiple inputs.

4. Empirical Evaluation of Training Regimes

4.1. Experimental Setup

In this section, we outline the setup for our empirical experiments. We release the source code of our experiments at: <https://github.com/kamadforge/early-exit-benchmark>. A more detailed description can be found in Appendix E.

Architectures and datasets. We conduct experiments across various datasets spanning computer vision (CV) and natural language processing (NLP). For CV, we utilize CIFAR-100 (Krizhevsky, 2009), ImageNet-1k (Russakovsky et al., 2015), TinyImageNet (Le & Yang, 2015), and Imagenette (Howard, 2019). For NLP, we evaluate on 20-Newsgroups (Lang, 1995) and STS-B (Wang et al., 2019) datasets. As far as vision architectures are concerned, we evaluate ResNet (He et al., 2016) and Vision Transformers (ViT) (Dosovitskiy et al.). Additionally, we explore MS-DNet (Huang et al., 2018), an architecture dedicated for multi-exit models. For NLP tasks, we utilize BERT (Devlin et al., 2019). Unless otherwise specified, the Shallow-Deep Network (SDN) (Kaya et al., 2019) early exit model is implemented on top of a chosen backbone.

Training set-up. We train our models, employing the AdamW optimizer (Loshchilov & Hutter, 2019) alongside

the Cosine Annealing scheduler with warm restarts. To ensure fair convergence across different regimes, we incorporate an early stopping mechanism. Training is terminated only when, over n consecutive epochs, none of the exits achieve an improved performance compared to their best scores recorded thus far. These scores – accuracy for classification tasks and loss for regression tasks – are evaluated on a dedicated early-stopping validation set. All results in this section are averaged over three runs, each one having a different initial seed.

Evaluation protocol. The core of our analysis is the framework proposed to understand how adding early exits affects the model’s internal state and hence the predicted values. To this end, in Section 4.2, we use the proposed framework to compare all the training regimes. Additionally, following the core idea that a model having early exits must be used to spare computational power, we evaluate all the resulting models under the lens of FLOPs saved and the obtained accuracy. These results are presented in Section 4.3.

4.2. Framework Evaluation

In this section, we evaluate the training regimes through the framework proposed in Section 3.

Gradient Dominance. In Figure 2, we present the gradient dominance results for the joint and mixed regimes, highlighting their distinct training dynamics. In the joint regime, optimization focuses on subnetworks in the middle of the architecture, where gradients from intermediate classifiers exert the strongest influence. Conversely, in the mixed regime, gradients are dominated by deeper classifiers, causing the early layers to primarily support the learning objectives of later classifiers, potentially at the expense of earlier ones.

The dominance of the final intermediate classifier (IC) in the mixed regime suggests its suitability for scenarios with




Figure 3. Mode connectivity between models trained with different training regimes. Colors represent the values of loss function, with yellow representing high loss (≥ 2.0). Disjoint and mixed regimes produce similar models, while the model trained in joint regime lies in a different basin (ResNet-20, CIFAR-10).

higher computational budgets, which is supported by our empirical results. This behavior mirrors the training of a backbone, where the model is primarily guided by the loss from the final classifier. Consequently, the mixed regime tends to produce models that closely resemble standard neural networks, a fact we confirm in the next section through mode connectivity analysis.

Mode Connectivity. After accounting for permutation symmetries, the loss remains low during linear interpolation of the weights of models trained in the mixed and disjoint regimes, as shown in the mode connectivity plot in Figure 3. This indicates that the resulting weights lie in the same basin, and the two regimes produce similar solutions. However, the disjoint regime is more constrained because it trains only the internal classifiers during the second phase. In contrast, the mixed regime updates the backbone to accommodate the added ICs, resulting in lower overall loss.

Joint training, on the other hand, leads to a model in a different basin of the loss landscape. This suggests that training the entire model at once, without “pre-trained” backbones, produces solutions that differ significantly in structure and performance.

Numerical Rank. In our framework, we analyze the numerical rank of the backbone model under different early-exit regimes. A regular neural network typically exhibits a higher rank in earlier layers and a lower rank in deeper layers, as illustrated in Figure 4. Note that training only the backbone corresponds to the model obtained in the disjoint regime, as this training strategy does not modify the backbone.

We observe a distinct change in network expressiveness once intermediate classifiers are attached and permitted to influence the backbone. The numerical rank rises across

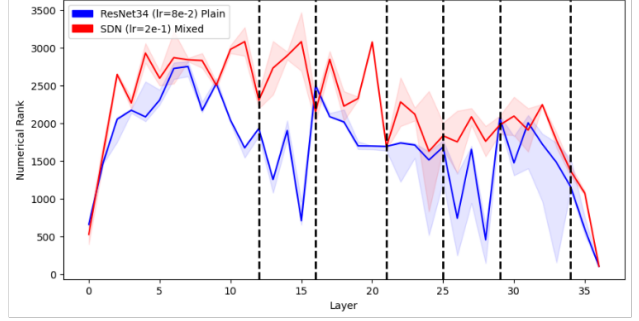


Figure 4. Numerical ranks as measured in the backbone network of the multi-exit model which was trained with different regimes. The rank increases when the ICs are allowed to affect the backbone network.

the layers, becoming comparable to the numerical rank of a multi-exit network trained entirely from scratch. This suggests that well-functioning multi-exit models require layers with greater expressiveness, which underpins the effective operation of all subsequent internal classifiers. Such an increase in expressiveness is impeded when employing the disjoint training strategy.

Mutual Information. Early-exit architectures attach intermediate classifiers (ICs) to internal layers, altering the distribution of information flow as seen in Figure 5. The effect is two-fold and differs between earlier and later layers. *Earlier layers:* The mutual information between X and Z is *larger* compared to a network trained without additional classifiers. *Deeper layers:* The mutual information for early-exit architecture is *lower* in the final layers.

The above effect is seen in both regimes but is more pronounced in the joint regime. The information flow in the joint regime is more skewed and different from backbone-only training. Backbone training in the mixed regime makes the information flow fall between backbone-only and joint training. This is due to the fact that the representation of easy samples is not complex (that is, it is processed with just a few layers before exiting through an early IC). As




Figure 5. Mutual information $I(X; Z)$ between the input X and the internal representation Z of the backbone at different depths for the three multi-exit model training regimes.

Table 1. Testset accuracy of multi-exit models trained with different regimes on the CV datasets. Each column represents the maximum averaged computational budget assumed for the model, indicated as a percentage of the computational cost of the backbone. The mixed regime achieves significant improvements over the commonly used joint and disjoint regimes.

Setting	Regime	25%	50%	75%	100%
ResNet-34 C-100	Disjoint	52.57	67.56	73.49	73.79
	Joint	62.84	72.80	74.32	74.17
	Mixed	62.24	73.81	75.92	75.88
MSDNet C-100	Disjoint	56.74	63.96	68.59	70.36
	Joint	65.93	72.02	74.73	75.86
	Mixed	65.94	72.01	75.46	76.51
ViT-T C-100	Disjoint	28.33	47.32	61.50	63.99
	Joint	40.87	60.48	66.22	66.49
	Mixed	41.65	64.07	70.09	70.25
ResNet-50 Tiny-IN	Disjoint	38.49	49.25	60.50	65.71
	Joint	52.98	62.16	65.14	65.01
	Mixed	52.89	63.28	67.20	67.24
ViT-T IN-1k	Disjoint	10.22	35.15	61.23	71.61
	Joint	36.39	61.89	68.08	68.39
	Mixed	35.91	62.96	70.79	71.20
ViT-S IN-1k	Disjoint	10.23	33.91	68.02	78.38
	Joint	50.10	73.99	76.38	76.44
	Mixed	49.50	75.17	78.28	78.33

the sample is easy, it is clearly and distinctly located within the boundaries of a single class. To describe it in terms of mutual information, the network does not need to reduce the complexity of X to fit the internal representation Z , as X has little irrelevant details. Consequently, the input X is not compressed and the internal representation Z has similar complexity to the representation of X , hence $I(X; Z)$ is higher.

Following this observation, we note that with higher $I(X; Z)$ in earlier layers, the joint strategy is more suitable for easy datasets where more samples exit at earlier layers. Similarly, the mixed regime learns more uniform representation of the $I(X; Z)$ across the network (one may observe an analogy to the numerical rank results in Section 3.3) and may be preferred for more difficult datasets that exit at later internal classifiers.

4.3. Performance-Cost Evaluation

Table 1 presents the performance of the training regimes on various dataset and architectures. The **disjoint regime shows weaker performance under lower computational budgets**, as freezing the backbone limits the ability to learn effective early-layer representations for classification. The results for ImageNet-1k, which are also reported in Fig-

Table 2. Testset accuracy (classification) and MSE (regression) of multi-exit BERT-B models trained using different regimes on the NLP datasets. The results align with those observed for classification tasks.

Dataset	Regime	25%	50%	75%	100%
Newsgr. (class.)	Disjoint	68.53	83.75	85.75	85.69
	Joint	84.24	84.41	84.41	84.41
	Mixed	84.99	85.25	85.25	85.25
STS-B (regr.)	Disjoint	2.37	1.55	0.54	0.51
	Joint	1.70	0.61	0.53	0.52
	Mixed	2.43	0.59	0.53	0.51

ure 1, demonstrate a similar trend, but also highlighting that the **joint regime experiences a performance gap at higher computational budgets**. The mixed regime performs consistently well across all scenarios. Due to space constraints in the main paper we present the result averaged over three runs, but in Appendix C we present these results with standard deviations. In Appendix A we also examine four alternative training regimes.

Generalization to Natural Language Processing & Regression Tasks. To explore the generalizability of these results, we first conduct similar experiments on natural language processing data and on a regression task, with the results presented in Table 2. These results are consistent with those observed before, further highlighting the weaknesses of both the joint and disjoint regimes.

Generalization to Other Early-Exit Methods. To show that our conclusions extend to other early-exit approaches, we conduct three additional experiments for different early-exit methods: GPF (Liao et al., 2021), PBEE (Zhou et al., 2020), and one that uses entropy instead of max-softmax probability as a proxy for confidence (Teerapittayanon et al., 2016). The findings from these experiments are consistent with our main results, again showing the effectiveness of the mixed regime.

Generalization to the Pre-trained Setup. Nowadays, most applications begin with a model pre-trained on a large dataset, which is then fine-tuned for a specific target task. We emphasize that during the fine-tuning step, any of the multi-exit training regimes can be applied. To assess whether our findings extend to this scenario, we use the weights of a pre-trained ViT-B model from the *torchvision* library and fine-tune it on CIFAR-100 using the three training regimes. The results, presented in Table 4, demonstrate patterns similar to those observed in the from-scratch training setup. Specifically: (1) the disjoint regime continues to show significantly lower performance under lower computational budgets, and (2) the joint regime still exhibits

Table 3. CIFAR-100 testset accuracy of multi-exit models trained with different regimes on the following early-exit methods: GPF (Liao et al., 2021), PBEE (Zhou et al., 2020), and entropy-based exit criterion (Teerapittayanon et al., 2016). The results are consistent with the main experiments.

EE Method	Regime	25%	50%	75%	100%
PBEE	Disjoint	19.28	44.11	57.85	63.91
	Joint	28.66	56.49	64.31	66.52
	Mixed	28.87	59.42	68.33	70.33
GPF	Disjoint	33.16	53.19	62.54	64.01
	Joint	46.21	62.23	66.84	67.19
	Mixed	47.21	64.51	68.92	69.10
Entropy	Disjoint	27.14	46.09	60.30	63.99
	Joint	41.16	58.60	65.69	66.49
	Mixed	41.54	62.20	69.63	70.25

Table 4. Test accuracy of a ViT-B model pre-trained on ImageNet-1k and fine-tuned on CIFAR-100. The disjoint regime struggles at low budgets, while the joint regime shows a slight but notable performance gap at higher budgets. The mixed regime performs consistently well, aligning with the patterns observed in the main experiments.

Regime	35%	50%	75%	100%
Disjoint	28.50	55.64	87.49	89.95
Joint	73.18	85.41	87.82	87.85
Mixed	73.38	85.70	88.06	88.23

a performance gap at higher computational budgets.

4.4. Impact of Gradient and Loss Scaling

The gradient dominance results demonstrated how the mixed regime prioritizes deeper intermediate classifiers. A similar effect can also be achieved by adjusting the loss coefficients for deeper ICs, as proposed in prior work (Kaya et al., 2019; Han et al., 2022), or by scaling the gradients of each IC, as proposed by Li et al. (Li et al., 2019). We revisit these techniques within both the joint and mixed regimes to evaluate their impact, with the results presented in Table 5.

We first evaluate the loss scaling method proposed by Kaya et al. (Kaya et al., 2019) (SDN) by training multi-exit ResNet-50 models on the TinyImageNet dataset. We also test (Han et al., 2022) constant loss weighting scheme, which adjusts the loss coefficients by either increasing (Inc.) or decreasing (Dec.) them along the model’s depth. These schemes allow the model to prioritize different computational budgets. However, **the improvements observed in the mixed regime are notably smaller compared to those seen in the joint regime.**

Finally, we evaluate the gradient equilibrium method (Li

Table 5. Accuracy improvement on the Tinyimagenet dataset when using various IC loss and gradient scaling methods. While gradient equilibrium enhances performance in the joint regime, it does not benefit the mixed regime, highlighting the effectiveness of the mixed regime for training multi-exit models.

Regime	Scaling	25%	50%	75%	100%
Joint	Inc.	−1.7	+0.4	+0.98	+1.06
Mixed	Inc.	−1.08	−0.27	+0.22	+0.25
Joint	Dec.	+0.6	−0.29	−0.08	+0.0
Mixed	Dec.	+0.69	+0.32	−0.44	−0.46
Joint	SDN	−3.75	−0.33	+0.89	+0.94
Mixed	SDN	−3.39	−1.1	+0.07	+0.23
Joint	GE	+0.31	+0.4	+0.73	+0.71
Mixed	GE	−1.16	−0.11	+0.02	−0.02

et al., 2019). For the joint regime, gradient scaling improves performance at higher computational budgets. However, for the mixed regime, there is no additional benefit, further confirming that they achieve a similar effect. **The mixed regime achieve superior results and it obviates the need for application of gradient equilibrium.**

4.5. Impact of IC Size

The size of an internal classifier (IC) in early exit architectures refers to the number of layers and neurons within the internal classifier. Smaller ICs are computationally efficient and enable faster early exits with minimal overhead. Conversely, larger ICs offer greater capacity, potentially enhancing accuracy, but may offset the computational benefits of multi-exit models.

A recent study by (Wójcik et al., 2023) investigated the impact of IC size and found that larger heads significantly improve performance. However, this analysis focused solely on the disjoint training regime. In our study, we examine the effect of varying IC sizes using a ViT-T model trained on the TinyImagenet dataset. Each IC consists of either one or two fully connected layers with hidden dimensions of 1024 or 2048, followed by a softmax layer. As summarized in Table 6, under the disjoint regime larger ICs indeed yield clear gains: moving from a single-layer IC with 1024 units to a two-layer IC with 2048 units can increase top-1 accuracy by up to 7.6 percentage points, and this effect is especially visible for lower budgets. However, in the mixed regime this trend reverses, and the change can result in an actual performance degradation. These results demonstrate that optimal IC architecture depends critically on the chosen training strategy, and that conclusions drawn under a disjoint setup may not transfer when the backbone is also trained.

Table 6. The effect of varying head size on the ViT-T model. Larger heads translate to inferior performance for the mixed regime. In contrast, the disjoint model clearly benefits from larger ICs when small budgets are considered.

Regime	IC arch.	25%	50%	75%	100%
Disjoint	1L	26.35	41.69	53.06	56.72
	2L-1024	33.85	<u>47.21</u>	<u>54.97</u>	56.72
	2L-2048	<u>33.96</u>	45.98	54.74	56.72
Joint	1L	42.47	53.24	56.02	56.03
	2L-1024	<u>45.59</u>	<u>55.11</u>	<u>57.73</u>	<u>57.66</u>
	2L-2048	43.60	53.89	56.94	57.00
Mixed	1L	44.03	57.91	60.42	60.28
	2L-1024	44.92	57.11	59.32	59.22
	2L-2048	44.61	56.94	60.18	60.15

4.6. Impact of IC Placement Scheme

The scheme of placing the internal classifiers in early exit architectures refers to where these classifiers are inserted at different layers within the model. This can range from being placed at every layer to being placed at strategically selected layers. To explore the impact of the placement scheme on the final performance, we train a ResNet-50 model with different regimes for every placement scheme. An IC can be placed at any block from index 0 to 14 for this architecture; however, we do not place it at index 0, as this would entirely bypass the backbone. In the Every- n placement scheme, an IC is inserted at every n -th block. In the Dense-Sparse configuration, ICs are placed at blocks: [1, 2, 3, 4, 5, 6, 7, 11], while in the Sparse-Dense scheme, they are placed at blocks: [1, 4, 8, 9, 10, 11, 12, 13, 14].

The placement scheme plays a significant role in shaping network performance across different training regimes, as shown in Table 7. While densely placed ICs generally ensure strong performance across a range of computational budgets, the model can be tailored to prioritize a particular budget by using the Dense-Sparse or Sparse-Dense schemes. In most cases, models trained with the mixed regime outperform those trained under other regimes.

4.7. Training time

The ultimate aim of a majority of multi-exit models is achieving the Pareto frontier in model performance and computational cost of the model during inference (Scardapane et al., 2020; Han et al., 2021). As such, to prevent under-training, we apply an early-stopping criterion in every experiment in this work. However, this raises an interesting question: how long does it take for the training in each regime to converge?

We explore the training time of the ResNet-50 model trained on the TinyImagenet dataset (see Table 1). Disjoint training

Table 7. The effect of varying head placement schemes on the SDN early-exit architecture with ResNet-50 as a backbone, trained on the TinyImagenet dataset.

Scheme	Regime	25%	50%	75%	100%
Every-1	Disjoint	38.92	49.25	60.10	65.76
	Joint	52.20	62.49	65.52	65.59
	Mixed	<u>52.22</u>	<u>63.03</u>	<u>67.21</u>	<u>67.35</u>
Every-2	Disjoint	37.34	48.03	60.34	65.65
	Joint	51.81	62.60	65.55	65.38
	Mixed	<u>52.41</u>	<u>63.19</u>	<u>67.14</u>	<u>67.33</u>
Every-3	Disjoint	—	47.91	60.95	65.77
	Joint	—	63.33	<u>67.22</u>	<u>67.21</u>
	Mixed	—	62.52	66.72	66.71
Every-4	Disjoint	—	41.32	57.78	65.78
	Joint	—	<u>62.54</u>	66.30	66.27
	Mixed	—	62.07	<u>67.08</u>	<u>67.14</u>
Every-5	Disjoint	—	39.85	56.20	65.72
	Joint	—	61.10	65.61	65.79
	Mixed	—	<u>61.95</u>	67.33	<u>67.40</u>
Den.-Spa.	Disjoint	38.47	50.64	62.04	65.74
	Joint	53.14	62.23	64.76	64.93
	Mixed	53.48	<u>63.17</u>	<u>66.24</u>	<u>66.27</u>
Spa.-Den.	Disjoint	37.12	47.03	59.79	65.68
	Joint	50.47	61.19	65.36	65.42
	Mixed	<u>51.19</u>	<u>62.27</u>	<u>67.26</u>	67.47

was the fastest and took 523 (± 156 , averaged over multiple seeds) epochs on average. Joint training converged after 1610 (± 395), while mixed training required 1166 (± 136) epochs. These results indicate that while disjoint training converges most rapidly, the mixed regime offers a substantial speed-up over joint training by alleviating the interference effect between intermediate classifiers.

5. Related Work

Early exiting is a notable application of the conditional computation paradigm (Bengio et al., 2013). While conceptually similar to earlier classifier cascades (Xu et al., 2014; Wang et al., 2018), it differs in that all classifiers are integrated within a single model, enabling end-to-end training. The first multi-exit model was introduced by (Teerapittayanon et al., 2016), and the field has expanded considerably since its inception.

Joint training is the most widely used and well-established strategy for early-exit models (Matsubara et al., 2022). This approach has been successfully applied to dynamic inference under various constraints, such as energy or time limitations (Wang et al., 2020), and extended to diverse early-exit

applications, including low-resolution classification (Xing et al., 2020), quality enhancement (Yang et al., 2020), and Question-Answering systems (Soldaini & Moschitti, 2020). While joint training has proven effective, several studies have demonstrated significant improvements through modifications of the training process. For instance, knowledge distillation from the final classifier to earlier internal classifiers has been shown to enhance their performance (Phuong & Lampert, 2019; Li et al., 2019; Liu et al., 2020). Similarly, ensembling multiple intermediate classifiers can improve the prediction accuracy (Qendro et al., 2021; Sarti et al., 2023). The Global Past-Future (GPF) method (Liao et al., 2021) goes a step further and incorporates information from both earlier predictions and surrogate later predictions to improve inference. Additionally, recent works (Han et al., 2022; Yu et al., 2023; Chataoui et al., 2023) identify a train-test mismatch in conventional multi-exit approaches and propose strategies that address this issue, further enhancing the robustness of early-exit models.

SDN (Kaya et al., 2019) was one of the first to explore the training of early-exit models through the pre-training of the architecture’s backbone followed by separate training of the classifiers. Multiple subsequent works have focused on optimizing early-exit models based only on this setup (Wójcik et al., 2023; Lahiany & Aperstein, 2022; Liu et al., 2020), potentially limiting the general applicability of their findings. For instance, (Wołczyk et al., 2021) employ an ensembling technique that combines predictions from earlier internal classifiers, weights of which are trained in a separate, third training phase. (Lahiany & Aperstein, 2022) propose PTEENet, which augments pre-trained networks with confidence heads that dynamically adjust based on available resources and unlabeled data.

(Kaya et al., 2019) were the first to explore both joint and disjoint training approaches for early-exit models. These approaches are also briefly reviewed in surveys such as (Scardapane et al., 2020; Matsubara et al., 2022). Furthermore, techniques like weighting the losses at each exit head (Zhou et al., 2020; Kaya et al., 2019; Han et al., 2022) or scaling the individual gradients (Li et al., 2019) can be regarded as variations of the joint training paradigm. In the context of LLMs, a concurrent work of Bae et al. (2024) includes an empirical comparison of joint and disjoint training strategies, with similar conclusions about the weaknesses of both approaches. To the best of our knowledge, our work is the first to present a broad and systematic analysis of the early-exit training strategies.

6. Conclusion

This study emphasizes the critical importance of training regimes for early-exit models, an aspect that has been largely neglected in prior research. Addressing this gap, we pro-

vide a detailed analysis and evaluation of various training approaches for early-exit architectures in deep neural networks. Our findings reveal that the manner in which the backbone and internal classifiers are trained has a significant impact on the performance and efficiency of these models. Furthermore, we show that when using the mixed training regime, gradient scaling is unnecessary. Below we summarize some practical takeaways for the selection of the training regime.

Mixed. Mixed regime demonstrates substantial robustness across various factors, including different data modalities and early-exit approaches with varying exit criteria. Therefore, the mixed regime is generally preferred, combining the benefits of both disjoint and joint training. The mixed regime ensures that the backbone network is well-optimized before integrating internal classifiers, leading to improved computational efficiency and accuracy. It is particularly recommended for cases where the performance on medium and high computational budgets is the main requirement.

Joint. The joint regime is simple to implement and can perform well for small computational budgets. However, it generally underperforms, especially for medium to higher budgets, which are more relevant in practical applications.

Disjoint. This regime is generally inferior compared to the others in most setups, but performs well for the highest budgets. It may be preferred when the backbone is shared, or the lack of resources prevents us from training the backbone network.

Acknowledgements

The research was funded by the program Excellence Initiative – Research University at the Jagiellonian University in Kraków.

This research has been supported by a grant from the Faculty of Mathematics and Computer Science under the Strategic Programme Excellence Initiative at Jagiellonian University.

This research was partially funded by National Science Centre, Poland, grant no 2022/45/B/ST6/02817.

We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grants no. PLG/2023/016270, PLG/2024/017185, and PLG/2024/017385, and the LUMI consortium through PLL/2024/07/017641.

The contribution of Bartosz Wójcik to this research was conducted at the Faculty of Mathematics and Computer Science, and the Doctoral School of Exact and Natural Sciences of the Jagiellonian University.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- Ainsworth, S., Hayase, J., and Srinivasa, S. Git re-basin: Merging models modulo permutation symmetries. In *The Eleventh International Conference on Learning Representations*. 3
- Bae, S., Fisch, A., Harutyunyan, H., Ji, Z., Kim, S., and Schuster, T. Relaxed recursive transformers: Effective parameter sharing with layer-wise lora. *ArXiv preprint*, abs/2410.20672, 2024. URL <https://arxiv.org/abs/2410.20672>. 9
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *ArXiv preprint*, abs/1308.3432, 2013. URL <https://arxiv.org/abs/1308.3432>. 8
- Chataoui, J., Coates, M., et al. Jointly-learned exit and inference for a dynamic neural network. In *The Twelfth International Conference on Learning Representations*, 2023. 9
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>. 4
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*. 4
- Elhoushi, M., Shrivastava, A., Liskovich, D., Hosmer, B., Wasti, B., Lai, L., Mahmoud, A., Acun, B., Agarwal, S., Roman, A., et al. Layerskip: Enabling early exit inference and self-speculative decoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12622–12642, 2024. 13
- Garipov, T., Izmailov, P., Podoprikin, D., Vetrov, D. P., and Wilson, A. G. Loss surfaces, mode connectivity, and fast ensembling of dnns. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 8803–8812, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/be3087e74e9100d4bc4c6268cdbe8456-Abstract.html>. 3
- Han, Y., Huang, G., Song, S., Yang, L., Wang, H., and Wang, Y. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44 (11):7436–7456, 2021. 8
- Han, Y., Pu, Y., Lai, Z., Wang, C., Song, S., Cao, J., Huang, W., Deng, C., and Huang, G. Learning to weight samples for dynamic early-exiting networks. In *European Conference on Computer Vision*, pp. 362–378. Springer, 2022. 7, 9
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90.4>
- Howard, J. Imagenette: A smaller subset of 10 easily classified classes from imagenet, 2019. URL <https://github.com/fastai/imagenette>. Accessed: 2025-01-28. 4
- Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., and Weinberger, K. Q. Multi-scale dense networks for resource efficient image classification. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=Hk2aImxAb>. 1, 4
- Kawaguchi, K., Deng, Z., Ji, X., and Huang, J. How does information bottleneck help deep learning? In *International Conference on Machine Learning*, pp. 16049–16096. PMLR, 2023. 3
- Kaya, Y., Hong, S., and Dumitras, T. Shallow-deep networks: Understanding and mitigating network overthinking. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3301–3310. PMLR,

2019. URL <http://proceedings.mlr.press/v97/kaya19a.html>. 1, 4, 7, 9, 17
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009. 4
- Lahiany, A. and Aperstein, Y. Pteenet: Post-trained early-exit neural networks augmentation for inference cost optimization. *IEEE Access*, 10:69680–69687, 2022. 9
- Lang, K. 20 newsgroups dataset, 1995. URL <http://qwone.com/~jason/20Newsgroups/>. Accessed: 2025-01-28. 4
- Le, Y. and Yang, X. S. Tiny imagenet visual recognition challenge. 2015. URL <https://api.semanticscholar.org/CorpusID:16664790>. 4
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. Visualizing the loss landscape of neural nets. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 6391–6401, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/a41b3bb3e6b050b6c9067c67f663b915-Abstract.html>. 13
- Li, H., Zhang, H., Qi, X., Yang, R., and Huang, G. Improved techniques for training adaptive deep networks. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 1891–1900. IEEE, 2019. doi: 10.1109/ICCV.2019.00198. URL <https://doi.org/10.1109/ICCV.2019.00198>. 7, 9
- Liao, K., Zhang, Y., Ren, X., Su, Q., Sun, X., and He, B. A global past-future early exit method for accelerating inference of pre-trained language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2013–2023, Online, 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.162. URL <https://aclanthology.org/2021.naacl-main.162>. 2, 6, 7, 9
- Liu, W., Zhou, P., Wang, Z., Zhao, Z., Deng, H., and Ju, Q. FastBERT: a self-distilling BERT with adaptive inference time. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6035–6044, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.537. URL <https://aclanthology.org/2020.acl-main.537>. 9
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>. 4
- maintainers, T. and contributors. Torchvision: Pytorch’s computer vision library. <https://github.com/pytorch/vision>, 2016. 20
- Masarczyk, W., Ostaszewski, M., Imani, E., Pascanu, R., Miłoś, P., and Trzcinski, T. The tunnel effect: Building data representations in deep neural networks. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- Matsubara, Y., Levorato, M., and Restuccia, F. Split computing and early exiting for deep learning applications: Survey and research challenges. *ACM Computing Surveys*, 55(5):1–30, 2022. 2, 8, 9
- Meronen, L., Trapp, M., Pilzer, A., Yang, L., and Solin, A. Fixing overconfidence in dynamic neural networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2680–2690, 2024. 1
- Phuong, M. and Lampert, C. Distillation-based training for multi-exit architectures. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 1355–1364. IEEE, 2019. doi: 10.1109/ICCV.2019.00144. URL <https://doi.org/10.1109/ICCV.2019.00144>. 9
- Qendro, L., Campbell, A., Lio, P., and Mascolo, C. Early exit ensembles for uncertainty quantification. In *Machine Learning for Health*, pp. 181–195. PMLR, 2021. 9
- Rahmath P, H., Srivastava, V., Chaurasia, K., Pacheco, R. G., and Couto, R. S. Early-exit deep neural network-a comprehensive survey. *ACM Computing Surveys*, 57(3):1–37, 2024. 13
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y. 4
- Sarti, S., Lomurno, E., and Matteucci, M. Anticipate, ensemble and prune: Improving convolutional neural networks via aggregated early exits. *Procedia Computer Science*, 222:519–528, 2023. 9
- Scardapane, S., Scarpiniti, M., Baccarelli, E., and Uncini, A. Why should we add early exits to neural networks? *Cognitive Computation*, 12(5):954–966, 2020. 8, 9, 13

- Shwartz-Ziv, R. and Tishby, N. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017. 3
- Soldaini, L. and Moschitti, A. The cascade transformer: an application for efficient answer sentence selection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5697–5708, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.504. URL <https://aclanthology.org/2020.acl-main.504>. 9
- Teerapittayanon, S., McDanel, B., and Kung, H.-T. Branchynet: Fast inference via early exiting from deep neural networks. In *Proceedings of the International Conference on Pattern Recognition, ICPR*, pp. 2464–2469, 2016. 2, 6, 7, 8
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJ4km2R5t7>. 4
- Wang, X., Luo, Y., Crankshaw, D., Tumanov, A., Yu, F., and Gonzalez, J. E. IDK cascades: Fast deep learning by learning not to overthink. In Globerson, A. and Silva, R. (eds.), *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pp. 580–590. AUAI Press, 2018. URL <http://auai.org/uai2018/proceedings/papers/212.pdf>. 8
- Wang, Y., Shen, J., Hu, T.-K., Xu, P., Nguyen, T., Baraniuk, R., Wang, Z., and Lin, Y. Dual dynamic inference: Enabling more efficient, adaptive, and controllable deep inference. *IEEE Journal of Selected Topics in Signal Processing*, 14(4):623–633, 2020. 8
- Wójcik, B., Przewieźlikowski, M., Szatkowski, F., Wołczyk, M., Bałazy, K., Krzepkowski, B., Podolak, I., Tabor, J., Śmieja, M., and Trzciński, T. Zero time waste in pre-trained early exit neural networks. *Neural Networks*, 168: 580–601, 2023. 7, 9
- Wołczyk, M., Wójcik, B., Bałazy, K., Podolak, I. T., Tabor, J., Śmieja, M., and Trzciński, T. Zero time waste: Recycling predictions in early exit neural networks. *Advances in Neural Information Processing Systems*, 34: 2516–2528, 2021. 9
- Xin, J., Tang, R., Yu, Y., and Lin, J. BERxiT: Early exiting for BERT with better fine-tuning and extension to regression. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 91–104, Online, 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.eacl-main.8>. 13
- Xing, Q., Xu, M., Li, T., and Guan, Z. Early exit or not: Resource-efficient blind quality enhancement for compressed images. In *European Conference on Computer Vision*, pp. 275–292. Springer, 2020. 9
- Xu, Z., Kusner, M. J., Weinberger, K. Q., Chen, M., and Chapelle, O. Classifier cascades and trees for minimizing feature evaluation cost. *The Journal of Machine Learning Research*, 15(1):2113–2144, 2014. 8
- Yang, L., Han, Y., Chen, X., Song, S., Dai, J., and Huang, G. Resolution adaptive networks for efficient inference. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 2366–2375. IEEE, 2020. doi: 10.1109/CVPR42600.2020.00244. URL <https://doi.org/10.1109/CVPR42600.2020.00244>. 1, 9
- Yu, H., Li, H., Hua, G., Huang, G., and Shi, H. Boosted dynamic neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 10989–10997, 2023. 9
- Yun, S., Han, D., Chun, S., Oh, S. J., Yoo, Y., and Choe, J. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 6022–6031. IEEE, 2019. doi: 10.1109/ICCV.2019.00612. URL <https://doi.org/10.1109/ICCV.2019.00612>. 19
- Zeiler, M. D. and Fergus, R. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I 13*, pp. 818–833. Springer, 2014. 2
- Zhang, H., Cissé, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=r1Ddp1-Rb>. 19
- Zhou, W., Xu, C., Ge, T., McAuley, J., Xu, K., and Wei, F. Bert loses patience: Fast and robust inference with early exit. *Advances in Neural Information Processing Systems*, 33:18330–18341, 2020. 2, 6, 7, 9

A. Alternative Regimes

In addition to the training regimes discussed in the main paper, early-exit surveys describe alternative training strategies (Scardapane et al., 2020; Rahmath P et al., 2024). Moreover, (Xin et al., 2021) have proposed the "Alternating" approach, which also attempts to combine the advantages of joint and disjoint regimes. Additionally, we propose a refinement of the mixed regime, termed the mixed-gradual training regime.

We describe these alternative regimes below:

- **Branch-wise training.** In this regime, ICs are trained sequentially, with only the part of the backbone corresponding to the current IC being unfrozen.
- **Separate training.** Similar to branch-wise training, ICs are optimized sequentially. However, at each step, the preceding ICs are jointly trained along with the current one. The entire model is unfrozen during the training process.
- **Alternating training.** Proposed in (Xin et al., 2021), this regime alternates between training the backbone independently (Equation (1)) and jointly training the backbone with all ICs (Equation (2)) in each training step.
- **Mixed-gradual training.** Building on the strengths of the mixed regime, we propose a more gradual optimization approach. Training proceeds in m phases, where m is the number of ICs. In the i -th phase, the last i ICs are jointly optimized, and no weights are frozen at any stage. This ensures flexibility throughout the training process while allowing the model to adapt more effectively to the optimization landscape. This approach is similar to *gradual early exit curriculum* proposed by the concurrent work of Elhoushi et al. (2024), with the difference being that we utilize early-stopping, while they enable the next IC at constant intervals of training steps.

In Appendix C, we present the results for these alternative training regimes. Both the branch-wise and separate training approaches demonstrate subpar performance when evaluated under higher computational budgets. Although the alternating training regime resembles the mixed regime in its motivation, its performance consistently falls short of both the mixed and joint regimes. Interestingly, the mixed-gradual regime outperforms the mixed regime, providing strong evidence in favor of our approach to enhancing early-exit optimization.

B. Loss Landscape

The concept of a loss landscape in the context of neural networks is crucial for understanding the training dynamics and generalization properties of models. The loss landscape provides a visual and analytical representation of how the loss function changes with respect to the model's parameters. By visualizing the loss landscapes of different neural network architectures, we can understand how design choices affect the shape of the loss function.

For a trained model with parameters θ^* , one can evaluate the loss function for the numbers x, y

$$f(x, y) = L(\theta^* + x\delta + y\eta) \quad (6)$$

such that δ, η are random directions sampled from a probability distribution, usually a Gaussian distribution, filter-normalized (Li et al., 2018), obtaining a 3D plot. In contrast to a typical neural network architecture, in early-exit set-up, both the final and internal classifiers are considered. We consider total training loss and separate losses for each IC.

When evaluating head losses, we use common random directions(δ, η) for each IC. The δ, η directions contain both backbone and head parameters.

As shown in Figure 6 there is a significant difference in loss landscapes between the Disjoint regime and the Joint one. The Joint and Mixed regimes are similar in this regard.

C. Full Results

Due to space constraints in the main paper we have presented only the averaged results. In this section, we present the main results with the standard deviation reported for each experiment. Furthermore, we include the results for additional training regimes, which are described in Appendix A.




Figure 6. Training loss landscapes: comparison for Joint, Mixed, and Disjoint regimes (left to right), head 1. Landscapes for SDN architecture with Resnet20 as backbone trained on the CIFAR-10 dataset.

Table 8. ResNet-34 CIFAR-100

Regime	25%	50%	75%	100%	Unlimited
Disjoint	52.57 \pm 0.07	67.56 \pm 0.22	73.49 \pm 0.93	73.79 \pm 1.00	73.79 \pm 1.00
Joint	62.84 \pm 0.67	72.80 \pm 0.34	74.32 \pm 0.08	74.17 \pm 0.10	74.19 \pm 0.09
Mixed	62.24 \pm 0.78	73.81 \pm 0.51	75.92 \pm 0.44	75.88 \pm 0.40	75.88 \pm 0.40
Branch-wise	64.19 \pm 0.47	67.03 \pm 0.76	66.88 \pm 0.60	66.75 \pm 0.68	66.75 \pm 0.68
Separate	66.25 \pm 0.54	72.57 \pm 0.38	72.87 \pm 0.41	72.82 \pm 0.46	72.82 \pm 0.46
Alternating	60.35 \pm 1.15	71.99 \pm 0.45	74.57 \pm 0.17	74.37 \pm 0.02	74.36 \pm 0.09
Mixed-gradual	60.11 \pm 0.68	73.16 \pm 0.70	76.08 \pm 0.54	75.94 \pm 0.75	75.94 \pm 0.75

Table 9. MSDNet CIFAR-100

Regime	25%	50%	75%	100%	Unlimited
Disjoint	56.74 \pm 0.41	63.96 \pm 0.74	68.59 \pm 0.72	70.36 \pm 0.84	70.54 \pm 0.93
Joint	65.93 \pm 0.20	72.02 \pm 0.27	74.73 \pm 0.18	75.86 \pm 0.26	75.93 \pm 0.21
Mixed	65.94 \pm 0.55	72.01 \pm 0.47	75.46 \pm 0.15	76.51 \pm 0.08	76.71 \pm 0.17
Branch-wise	58.24 \pm 0.45	59.73 \pm 0.41	60.16 \pm 0.47	59.97 \pm 0.38	59.73 \pm 0.29
Separate	64.90 \pm 0.56	67.84 \pm 0.17	69.02 \pm 0.25	69.21 \pm 0.23	69.15 \pm 0.23
Alternating	64.69 \pm 0.45	71.67 \pm 0.27	74.93 \pm 0.16	76.06 \pm 0.25	76.34 \pm 0.31
Mixed-gradual	64.73 \pm 0.66	72.10 \pm 0.47	75.66 \pm 0.51	77.21 \pm 0.61	77.49 \pm 0.64

Table 10. ViT-T CIFAR-100

Regime	25%	50%	75%	100%	Unlimited
Disjoint	28.33 \pm 0.47	47.32 \pm 0.81	61.50 \pm 1.82	63.99 \pm 1.68	63.99 \pm 1.68
Joint	40.87 \pm 0.84	60.48 \pm 0.94	66.22 \pm 1.07	66.49 \pm 1.08	66.49 \pm 1.08
Mixed	41.65 \pm 0.19	64.07 \pm 0.38	70.09 \pm 0.64	70.25 \pm 0.56	70.25 \pm 0.56
Branch-wise	32.41 \pm 1.32	33.46 \pm 0.76	34.33 \pm 0.49	34.51 \pm 0.45	34.51 \pm 0.45
Separate	42.67 \pm 0.89	51.48 \pm 3.58	54.29 \pm 4.68	54.42 \pm 4.83	54.42 \pm 4.83
Alternating	37.85 \pm 0.40	59.11 \pm 0.20	67.81 \pm 0.69	68.39 \pm 0.78	68.39 \pm 0.78
Mixed-gradual	40.61 \pm 0.08	64.83 \pm 0.46	71.42 \pm 0.34	71.60 \pm 0.38	71.60 \pm 0.38

Table 11. ResNet-50 Tinyimagenet

Regime	35%	50%	75%	100%	Unlimited
Disjoint	38.49 \pm 0.46	49.25 \pm 1.33	60.50 \pm 1.27	65.71 \pm 0.90	65.80 \pm 0.90
Joint	52.98 \pm 0.77	62.16 \pm 0.61	65.14 \pm 0.57	65.01 \pm 0.75	65.00 \pm 0.75
Mixed	52.89 \pm 0.23	63.28 \pm 0.59	67.20 \pm 0.65	67.24 \pm 0.71	67.22 \pm 0.71

Table 12. ViT-T ImageNet-1k

Regime	25%	50%	75%	100%	Unlimited
Disjoint	10.22 \pm 0.15	35.15 \pm 0.97	61.23 \pm 1.37	71.61 \pm 0.68	71.61 \pm 0.68
Joint	36.39 \pm 0.12	61.89 \pm 0.71	68.08 \pm 0.84	68.39 \pm 0.87	68.39 \pm 0.87
Mixed	35.91 \pm 0.19	62.96 \pm 0.12	70.79 \pm 0.36	71.20 \pm 0.34	71.20 \pm 0.34

Table 13. BERT-B 20-Newsgroups

Regime	25%	50%	75%	100%	Unlimited
Disjoint	68.53 \pm 0.97	83.75 \pm 0.34	85.75 \pm 0.16	85.69 \pm 0.28	85.54 \pm 0.34
Joint	84.24 \pm 0.48	84.41 \pm 0.41	84.41 \pm 0.41	84.41 \pm 0.41	84.51 \pm 0.60
Mixed	84.99 \pm 0.84	85.25 \pm 0.64	85.25 \pm 0.64	85.25 \pm 0.64	85.46 \pm 0.44
Branch-wise	80.20 \pm 0.19	80.05 \pm 0.12	80.05 \pm 0.12	80.05 \pm 0.12	79.83 \pm 0.06
Separate	79.87 \pm 0.36	79.87 \pm 0.36	79.87 \pm 0.36	79.87 \pm 0.36	79.78 \pm 0.30
Alternating	83.51 \pm 0.33	83.92 \pm 0.06	83.92 \pm 0.06	83.92 \pm 0.06	84.18 \pm 0.34
Mixed-gradual	81.35 \pm 1.24	81.35 \pm 1.24	81.35 \pm 1.24	81.35 \pm 1.24	81.24 \pm 0.93

Table 14. BERT-B STS-B

Regime	25%	50%	75%	100%	Unlimited
Disjoint	2.37 \pm 0.01	1.55 \pm 0.11	0.54 \pm 0.00	0.51 \pm 0.02	0.51 \pm 0.02
Joint	1.70 \pm 0.23	0.61 \pm 0.02	0.53 \pm 0.01	0.52 \pm 0.02	0.52 \pm 0.01
Mixed	2.43 \pm 0.03	0.59 \pm 0.01	0.53 \pm 0.01	0.51 \pm 0.00	0.51 \pm 0.00
Branch-wise	2.14 \pm 0.24	0.82 \pm 0.09	0.84 \pm 0.10	0.83 \pm 0.10	0.84 \pm 0.10
Separate	2.69 \pm 0.14	1.41 \pm 0.09	1.41 \pm 0.09	1.40 \pm 0.08	1.39 \pm 0.0
Alternating	1.76 \pm 0.33	0.60 \pm 0.00	0.51 \pm 0.01	0.50 \pm 0.01	0.50 \pm 0.01
Mixed-gradual	2.37 \pm 0.03	0.51 \pm 0.00	0.48 \pm 0.00	0.46 \pm 0.01	0.46 \pm 0.01

Table 15. ViT-T PBEE CIFAR-100

Regime	25%	50%	75%	100%	Unlimited
Disjoint	19.28 \pm 0.32	44.11 \pm 1.73	57.85 \pm 1.61	63.91 \pm 1.68	63.99 \pm 1.68
Joint	28.66 \pm 0.50	56.49 \pm 1.48	64.31 \pm 0.87	66.52 \pm 1.08	66.49 \pm 1.08
Mixed	28.87 \pm 0.44	59.42 \pm 0.50	68.33 \pm 0.60	70.33 \pm 0.68	70.25 \pm 0.56

Table 16. ViT-T GPF CIFAR-100

Regime	25%	50%	75%	100%	Unlimited
Disjoint	33.16 \pm 0.64	53.19 \pm 0.66	62.54 \pm 1.59	64.01 \pm 1.67	63.99 \pm 1.68
Joint	46.21 \pm 0.30	62.23 \pm 0.81	66.84 \pm 0.68	67.19 \pm 0.78	67.19 \pm 0.80
Mixed	47.21 \pm 1.02	64.51 \pm 0.84	68.92 \pm 0.30	69.10 \pm 0.38	69.10 \pm 0.38

Table 17. ViT-T Entropy CIFAR-100

Regime	25%	50%	75%	100%	Unlimited
Disjoint	27.14 \pm 0.36	46.09 \pm 0.59	60.30 \pm 1.41	63.99 \pm 1.68	63.99 \pm 1.68
Joint	41.16 \pm 0.82	58.60 \pm 0.89	65.69 \pm 1.23	66.49 \pm 1.08	66.49 \pm 1.08
Mixed	41.54 \pm 0.52	62.20 \pm 0.35	69.63 \pm 0.47	70.25 \pm 0.56	70.25 \pm 0.56

Table 18. ViT-B pretrained on ImageNet-1k to CIFAR-100

Regime	35%	50%	75%	100%	Unlimited
Disjoint	28.50 \pm 2.02	55.64 \pm 3.71	87.49 \pm 0.40	89.95 \pm 0.18	89.95 \pm 0.18
Joint	73.18 \pm 1.36	85.41 \pm 0.23	87.82 \pm 0.22	87.85 \pm 0.29	87.84 \pm 0.29
Mixed	73.38 \pm 0.35	85.70 \pm 0.63	88.06 \pm 0.31	88.23 \pm 0.29	88.23 \pm 0.29

Table 19. ResNet-50 Tinyimagenet loss and gradient scaling experiment

Regime	25%	50%	75%	100%	Unlimited
Joint	52.98 \pm 0.77	62.16 \pm 0.61	65.14 \pm 0.57	65.01 \pm 0.75	65.00 \pm 0.75
Joint GE	53.29 \pm 0.90	62.56 \pm 0.15	65.87 \pm 0.36	65.72 \pm 0.56	65.69 \pm 0.57
Joint Inc.	51.28 \pm 0.19	62.56 \pm 0.12	66.12 \pm 0.39	66.07 \pm 0.42	66.06 \pm 0.41
Joint Dec.	53.58 \pm 0.44	61.87 \pm 0.45	65.06 \pm 0.35	65.01 \pm 0.29	65.00 \pm 0.30
Joint SDN	49.23 \pm 0.64	61.83 \pm 0.32	66.03 \pm 0.18	65.95 \pm 0.25	65.93 \pm 0.24
Mixed	52.89 \pm 0.23	63.28 \pm 0.59	67.20 \pm 0.65	67.24 \pm 0.71	67.22 \pm 0.71
Mixed GE	51.73 \pm 0.12	63.17 \pm 0.33	67.22 \pm 0.39	67.22 \pm 0.10	67.21 \pm 0.12
Mixed Inc.	51.81 \pm 0.24	63.01 \pm 0.49	67.42 \pm 0.39	67.49 \pm 0.32	67.47 \pm 0.34
Mixed Dec.	53.58 \pm 0.69	63.60 \pm 0.25	66.76 \pm 0.50	66.78 \pm 0.47	66.75 \pm 0.47
Mixed SDN	49.50 \pm 0.74	62.18 \pm 0.76	67.27 \pm 0.68	67.47 \pm 0.55	67.47 \pm 0.56

Table 20. ViT-T Imagenette head size experiment

Regime, Head	25%	50%	75%	100%	Unlimited
Disjoint 1L	75.61 \pm 0.10	78.72 \pm 1.24	78.32 \pm 1.62	78.29 \pm 1.63	78.28 \pm 1.65
Disjoint 2L-1024	77.41 \pm 0.39	79.04 \pm 1.26	78.36 \pm 1.63	78.29 \pm 1.63	78.28 \pm 1.65
Disjoint 2L-2048	77.10 \pm 0.40	79.02 \pm 1.26	78.39 \pm 1.69	78.30 \pm 1.65	78.28 \pm 1.65
Joint 1L	81.44 \pm 1.54	82.51 \pm 1.40	82.46 \pm 1.41	82.46 \pm 1.41	82.45 \pm 1.40
Joint 2L-1024	80.00 \pm 0.69	81.06 \pm 0.57	80.99 \pm 0.69	80.99 \pm 0.69	81.00 \pm 0.68
Joint 2L-2048	79.68 \pm 0.28	80.67 \pm 0.08	80.64 \pm 0.15	80.64 \pm 0.15	80.61 \pm 0.13
Mixed 1L	82.21 \pm 0.08	83.20 \pm 0.39	83.19 \pm 0.41	83.19 \pm 0.41	83.18 \pm 0.45
Mixed 2L-1024	80.25 \pm 0.59	81.28 \pm 0.44	81.13 \pm 0.40	81.13 \pm 0.40	81.11 \pm 0.37
Mixed 2L-2048	79.63 \pm 0.67	80.77 \pm 0.58	80.70 \pm 0.51	80.70 \pm 0.51	80.67 \pm 0.52

D. Results as Performance-Cost plots

Following [Kaya et al. \(2019\)](#), we presented most of the results in this paper in the form of tables. In this section, we provide the same results as model performance vs. computational costs plots.




Figure 7. Performance–cost trade-off of the multi-exit ResNet-34 on the CIFAR-100 dataset




Figure 8. Performance–cost trade-off of the MSDNet architecture on the CIFAR-100 dataset




Figure 9. Performance–cost trade-off of the multi-exit ViT-T on the CIFAR-100 dataset




Figure 10. Performance–cost trade-off of the ViT-T architecture on the ImageNet-1k dataset




Figure 11. Performance–cost trade-off of the multi-exit BERT-B on the 20NewsGroups dataset




Figure 12. Performance–cost trade-off of the BERT-B architecture on the STS-B dataset




Figure 13. Performance–cost trade-off of the multi-exit PBEE with ViT-T as the backbone on the CIFAR-100 dataset

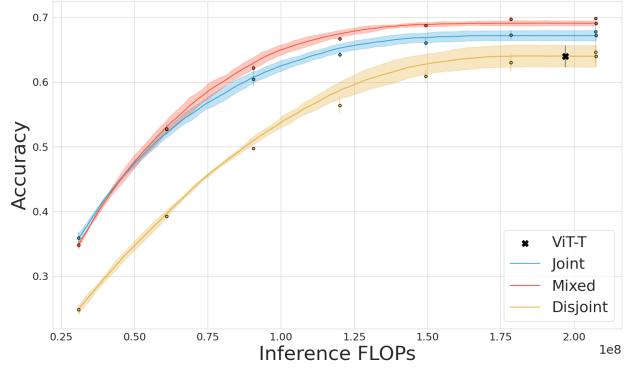


Figure 14. Performance–cost trade-off of the multi-exit GPF with ViT-T as the backbone on the CIFAR-100 dataset




Figure 15. Performance–cost trade-off of the multi-exit ViT-T on the CIFAR-100 dataset, using normalized entropy as the exit criterion




Figure 16. Performance–cost trade-off of the multi-exit ViT-B on the CIFAR-100 dataset, pre-trained on ImageNet-1k and fine-tuned on CIFAR-100




Figure 17. Performance–cost trade-off of the multi-exit ViT-S on the ImageNet-1k dataset

E. Training Details

Here we provide more information training setup, in addition to what is described in Section 4.1

E.1. ResNet-34, CIFAR-100

Model set-up. Exit heads are placed at positions $\{2, 4, \dots, 14\}$.

Training set-up. We train each model with batch size of 128. We use a learning rate of $5e-4$ and no weight decay. We set the early stopping patience to 50 epochs. CutMix and Mixup are used as augmentations.

E.2. MSDNet, CIFAR-100

Model set-up. We use CIFAR variant of MSDNet with 7 blocks. Exit heads are placed at positions $\{0, 1, \dots, 17\}$.

Training set-up. We train each model with batch size of 512. We use a learning rate of $1e-3$ and no weight decay. We set the early stopping patience to 50 epochs. CutMix and Mixup are used as augmentations.

E.3. ViT-T, CIFAR-100

Model set-up. ViT-T hyperparameters are as follows: a patch size of 4, an embedding size of 256, an MLP dimension of 256, 7 layers and 8 attention heads. Exit heads are placed at positions $\{0, 1, \dots, 6\}$. For GPF we set head embedding size to 512.

Training set-up. We train each model with batch size of 256. We use a learning rate of $5e-4$ and no weight decay. We set the early stopping patience to 30 epochs. Following augmentations are used: random resizing, cropping, rotation, contrast adjustment, random erasing, CutMix and Mixup.

E.4. ResNet-50, Tinyimagenet

Model set-up. Exit heads are placed at block positions $\{1..14\}$.

Training set-up. We train each model with batch size of 256. We use a learning rate of $1e-3$ and no weight decay. We set the early stopping patience to 50 epochs. CutMix and Mixup are used as augmentations.

E.5. ViT-T, ImageNet-1k

Model set-up. ViT-T hyperparameters are as follows: a patch size of 16, an embedding size of 192, an MLP dimension of 768, 12 layers, and 3 attention heads. Exit heads are placed at positions $\{1, 2, \dots, 10\}$.

Training set-up. We train each model using 4 A-100 GPUs with an effective batch size of 2048. We use a learning rate of $5e-4$ and no weight decay. We set the early stopping patience to 25 epochs. Following augmentations are used: random resizing, cropping, rotation, contrast adjustment, random erasing, CutMix (Yun et al., 2019) and Mixup (Zhang et al., 2018).

E.6. 20 Newsgroups, Bert-B

Model set-up. We set maximum sequence length to 512. Exit heads are placed at positions $\{1, 3, \dots, 21\}$

Training set-up. We train each model with batch size of 32. We use a learning rate of $5e-5$ and a weight decay of $1e-2$. We set the early stopping patience to 3 epochs.

E.7. STS-B, Bert-B

Model set-up. We set maximum sequence length to 128. Exit heads are placed at positions $\{1, 3, \dots, 21\}$

Training set-up. We train each model with batch size of 16. We use a learning rate of $1e-5$ and a weight decay of $1e-4$. We set the early stopping patience to 3 epochs.

E.8. ViT-B, CIFAR-100

Model set-up. ViT-B hyperparameters are as follows: a patch size of 16, an embedding size of 768, an MLP dimension of 3072, 12 layers and 12 attention heads. Exit heads are placed at positions $\{3, 5, \dots, 9\}$.

Training set-up We use pretrained weights from `torchvision` ([maintainers & contributors, 2016](#)), trained on ImageNet-1k. To fine-tune, we use a learning rate of $2e-5$ and a weight decay of $3e-2$. We set the early stopping patience to 30 epochs. We use Mixup as an augmentation.