



**HAL**  
open science

## **EERO: Early Exit with Reject Option for Efficient Classification with limited budget**

Florian Valade, Mohamed Hebiri, Paul Gay

► **To cite this version:**

Florian Valade, Mohamed Hebiri, Paul Gay. EERO: Early Exit with Reject Option for Efficient Classification with limited budget. 2024. hal-04438166

**HAL Id: hal-04438166**

**<https://hal.science/hal-04438166v1>**

Preprint submitted on 5 Feb 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# EERO: Early Exit with Reject Option for Efficient Classification with limited budget

Florian Valade<sup>1</sup>, Mohamed Hebiri<sup>1</sup>, and Paul Gay<sup>2</sup>

<sup>1</sup>Université Gustave Eiffel

<sup>2</sup>Université de Pau et des pays de l'Addour

February 5, 2024

## Abstract

The increasing complexity of advanced machine learning models requires innovative approaches to manage computational resources effectively. One such method is the Early Exit strategy, which allows for adaptive computation by providing a mechanism to shorten the processing path for simpler data instances. In this paper, we propose EERO, a new methodology to translate the problem of early exiting to a problem of using multiple classifiers with reject option in order to better select the exiting head for each instance. We calibrate the probabilities of exiting at the different heads using aggregation with exponential weights to guarantee a fixed budget. We consider factors such as Bayesian risk, budget constraints, and head-specific budget consumption. Experimental results, conducted using a ResNet-18 model and a ConvNext architecture on Cifar and ImageNet datasets, demonstrate that our method not only effectively manages budget allocation but also enhances accuracy in overthinking scenarios.

## 1 Introduction

Nowadays, vision models are increasing in size rising the issue of their complexity and computation costs. There exist different strategies to train lighter deep learning networks, such as quantization and pruning [19], distillation [25], and dynamic inference where the network adapts its topology on the fly to the input data [11]. Among them, Early Exit [17] is an orthogonal approach which aims at adapting the amount of computation to each input data point, exploiting that most neural networks can be approximated as a stack of layers which process the data sequentially. The idea is to add auxiliary heads at regular intervals along the network (see Figure 1) which are able to produce a prediction with the current state of the features. The intuition is that easy cases can be processed with only the first few layers. Thus, Early Exit provides an interesting flexibility. Moreover, the predictions from the different heads can be used to construct

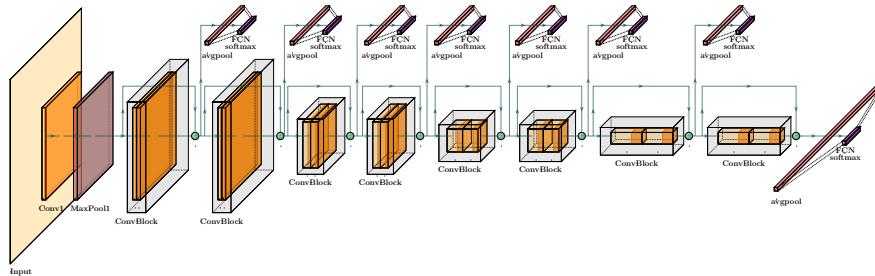


Figure 1: Illustration of the Early Exit principle in a convolutional architecture used in our experiments in Section 4.1.

confidence scores and analyze the importance of the different model blocks [3, 15]. Lastly, it has been observed that training with early exit provides additional and more immediate gradient signal in back propagation, resulting in more discriminant features in lower layers, thus improving accuracy [24].

Despite its advantages, one of the main challenges in applying Early Exit during inference is determining the appropriate moment to exit for a given input. This challenge is intricately connected to assessing confidence in the predictions of neural networks. Specifically, we focus on *Budgeted batch classification* [14], a strategy where a fixed computational budget is allocated for processing a batch of data. The strategy involves calculating thresholds to allocate these resources efficiently across different data points enhancing overall accuracy. This approach theoretically offers improved performance because it allows for the conservation of computational resources on simpler cases, which can then be reallocated to improve accuracy on more complex cases. However, practical investigations into this framework have been limited, with most studies concentrating on architectural design or the reliability of confidence scores. Those few that have addressed Budgeted Batch Classification often make additional assumptions for practicality [14] or utilize less-than-ideal algorithms without fully exploring the associated mathematical challenges [30].

**Contributions** In this work, we introduce a new methodology termed EERO, which stands for Early Exit with Reject Option, designed to optimize the use of Early Exit in the context of Budgeted Batch Classification for inference tasks. We formalize the problem as classification with a reject option and establish the optimal rule for a two-head scenario, deriving a data-driven procedure to define the rejection threshold. We then extend this approach to multiple heads, framing it as a minimization of the joint empirical risk under a global budget constraint, solvable through aggregation with exponential weights. Our work delivers key contributions: (i) we determine the optimal data proportion for a single auxiliary head under a given resources budget expressed in GFlops, and construct an optimal classification rule that balances classification error

and budget adherence and as opposed to previous work, strictly enforces the budget constraint; (ii) we offer a data-driven counterpart of the optimal rule that empirically exhibits a balanced accuracy-budget compromise; (iii) we apply our methodology to ResNet, ConvNext[21] and MSDNet [14] architectures, generalizing it for multiple auxiliary heads; (iv) we validate our approach through real data benchmark experiments on CIFAR 100 and ImageNet, demonstrating its efficacy in adhering to resources constraints while reducing computation and enhancing model accuracy.

The remainder of this paper is structured to provide a comprehensive examination of our research and findings. Following this introduction, Section 2 reviews the existing literature and situates our work within the broader context of related research. Section 3 delves into our proposed methods, starting with a statistical framework and definition in Section 3.1, followed by a detailed exposition of the EERO method in Section 3.2. Section 4 presents our experimental results, showcasing the application of EERO in single-head and multi-head scenarios in Sections 4.1 and 4.2, respectively. Finally, we encapsulate our study with a conclusion in Section 5, where we summarize our key findings, discuss the implications of our work, and offer directions for future research.

## 2 Related work

One of the first works on Early Exit [24] proposed to use auxiliary heads and a weighted sum of losses. The focus was essentially on the design and the positions of the auxiliary heads. MSDNet [14] is an architecture where the heads are carefully designed so that early features which are assumed to be unsuitable for classification are refined. Although this leads to better performances, the complex resulting heads reduce their number, and therefore the flexibility of the model, and makes it more difficult to adapt to new architectures. Overall, the position and the number of the auxiliary heads remains an open discussion for convolutional networks [20] and transformers [1].

Determining the precise timing for an early exit in neural network inference have been addressed by different strategies. A prevalent approach applies a specific threshold to the output probabilities of each auxiliary head, ceasing the inference process when the output for a class surpasses this benchmark. This threshold can be based on a maximum confidence value or on the entropy of the probability distribution, which reflects the certainty of the prediction [14, 2]. Patience-based strategies introduce another layer of decision-making by requiring that multiple consecutive heads agree on the same class prediction before allowing an exit, which utilizes the redundancy in sequential predictions to ensure reliability [33].

More sophisticated supervised methods involve training a cost function that enforces a desired utilization rate of the network’s layers [26]. Modules may also be trained to calculate ‘halting scores’ from each layer’s output, accumulating those scores along the network until a specific exit threshold is reached, thus integrating the model’s confidence into the inference process [9]. In more complex

dynamic inference settings, reinforcement learning techniques have been adopted to finely balance the trade-off between computational efficiency and prediction accuracy [29, 31].

In the context of *Budgeted batch classification*, thresholds are not determined for individual images but for a batch as a whole, facilitating a more strategic deployment of computational resources. To date, the literature presents two notable algorithms addressing this problem in the realm of Early Exit strategies. The authors in [30] employ a genetic algorithm to optimize these thresholds, offering a biological-evolution-inspired approach to the allocation challenge. Conversely, the work [14] operates under the assumption that each exit point in the network has an equal and predetermined likelihood of accurately classifying an image. While this simplifies the optimization problem, it introduces an additional hyper-parameter that is not inherently related to computational complexity. Our work differs by presenting a comprehensive formulation that is grounded on the principles of reject option learning theory.

Learning with reject option is a related line of work where the main focus is to abstain from predicting when there is a doubt in the predicted value. Reject option has first been exploited in [4] and has been considered by the statistical learning community in the early 2000 with the development of *conformal prediction* in [27, 28]. The papers [13, 23, 10, 32, 18, 5, 6] (and references therein) contributed to a good comprehension of the reject option rule. In particular, there are three main approaches for learning with reject option (without any clear advantage of one approach on the others). Either, we focus on ensuring a predefined level of coverage, or on a predefined rejection rate, or on a tradeoff of the two. In our context, rejection will be considered at the level of heads to determine whether an instance  $\mathbf{x}$  should be classified by a given head or considered by the following heads. While this approach is completely original in its application to deep learning and energy saving, it is worth noticing that reject option have already been successfully used to handle other types of learning issues [28, 8, 7].

### 3 Method

In this section, we introduce the Early Exit with Reject Option (EERO) framework, applying reject option learning theory to enable efficient early exits in neural networks while adhering to computational budgets. We detail the statistical underpinnings and describe how EERO strategically manages resources to balance accuracy with computational expenditure. This approach not only optimizes performance but also ensures strict compliance with predefined computational constraints.

#### 3.1 Statistical framework

This section describes the mathematical framework for Early Exit. Let  $(\mathbf{X}, Y)$  be a tuple on  $\mathcal{X} \times [K]$ , where  $[K] := \{1, \dots, K\}$ , having a joint distribution  $\mathbb{P}$ . Here,

$\mathcal{X} \subset \mathbb{R}^d$  is the feature space, and  $Y$  is the label corresponding to the feature  $\mathbf{X}$ . We focus on the problem of  $K$ -class classification with  $K \geq 2$  and one of our goals is to build a prediction rule  $g : \mathcal{X} \rightarrow [K]$  that reduces the misclassification risk  $\mathbb{P}(g(\mathbf{X}) \neq Y)$ . This risk is minimized by the Bayes rule  $g^*$  that is given, in the multi-class setting, for all  $\mathbf{x} \in \mathcal{X}$  by

$$g^*(\mathbf{x}) = \operatorname{argmax}_{k=1, \dots, K} p_k(\mathbf{x}) \ , \quad (1)$$

where  $p_k(\mathbf{x}) = \mathbb{P}(Y = k | \mathbf{X} = \mathbf{x})$  are the conditional probabilities. Because the distribution  $\mathbb{P}$  is unknown, the Bayes rule itself is unknown, and we need to approximate it. In general, this approximation seeks to maximize the classification accuracy. In our case, the goal is also to reduce the energy consumption of the model, and thus, we assume a constraint on the computation budget available. Notably, we will show that the latter is strongly connected to classification with reject option (classification with abstention) that we describe in Section 3.2.

The framework of classification with reject option assumes that classifiers are allowed to abstain from classifying (on the empirical side, this means that the classifier abstains on a given proportion of the data). Let  $\varepsilon \in (0, 1)$  be a parameter that denotes the probability of classifying an image. We define the optimal rule that abstains with probability  $1 - \varepsilon$  as follows:

**Definition 1** *Let  $\varepsilon \in (0, 1)$ . The optimal classifier with  $1 - \varepsilon$  rejection rate is defined as*

$$h_\varepsilon^* \in \operatorname{argmin}_h \{ \mathbb{P}(\{h(\mathbf{X}) \neq Y\} \cap \{h(\mathbf{X}) \neq \mathfrak{R}\}) \text{ s.t. } \mathbb{P}(h(\mathbf{X}) = \mathfrak{R}) = 1 - \varepsilon \} \ ,$$

where  $h$  is a classifier that is allowed to reject, that is,  $h : \mathcal{X} \rightarrow [K] \cup \{\mathfrak{R}\}$  and  $\mathfrak{R}$  is the output when the classifier rejects all elements from  $[K]$ .

**Remark 1** *It might be suitable to replace the equality constraint in Definition 1 by an inequality, that is,  $\mathbb{P}(h(\mathbf{X}) = \mathfrak{R}) \leq 1 - \varepsilon$ . However, it is easy to observe that, thanks to the continuity assumption, such modification does not change the optimal rule since the risk of  $h^*$  is a non-increasing function of  $1 - \varepsilon$ . Indeed, when  $\varepsilon$  decreases, the risk of  $h^*$  is evaluated on smaller region (w.r.t  $\mathbb{P}$ ), corresponding to  $\{\mathbf{x} : h^*(\mathbf{x}) \neq \mathfrak{R}\}$ .*

The opportunity of using the reject option is important in applications where ambiguity occurs between classes, which is often the case when the total amount of classes  $K$  is large. There are several ways to handle this reject option. In this paper, we constrain the rejection rate as it is in accordance with the resource limitation.

Now, we aim at providing the explicit expression of the optimal rule given by Definition 1. Let us denote by  $s$  the score function defined for each  $\mathbf{x} \in \mathcal{X}$  by

$$s(\mathbf{x}) = \max_{k \in [K]} \{p_1(\mathbf{x}), \dots, p_K(\mathbf{x})\} \ , \quad (2)$$

From the definition of the Bayes rule (1), we can establish the following characterization of the optimal rule.

**Proposition 1** Assume the cumulative distribution function (CDF)  $F_s$  of  $s(\mathbf{X})$  is continuous, Then for all  $\mathbf{x} \in \mathcal{X}$

$$h_\varepsilon^*(\mathbf{x}) = \begin{cases} g^*(\mathbf{x}) & \text{if } F_s(s(\mathbf{x})) \geq 1 - \varepsilon \\ \mathfrak{R} & \text{otherwise.} \end{cases}$$

In particular, we have  $\mathbb{P}(h_\varepsilon^*(\mathbf{X}) = \mathfrak{R}) = 1 - \varepsilon$ .

Proposition 1 extends the result established in [6] (c.f. Proposition 1) to multi-class setting. Its proofs can be found in the Section B of the Appendix. The main assumption used to build this result is the continuity assumption on the CDF of  $s(\mathbf{X})$  and requires that the random variable has no atoms. It is fundamental to ensure that the classifier  $h_\varepsilon^*$  has a rejection rate exactly  $1 - \varepsilon$ . In the next section, we will see that from an empirical point of view, this condition can always be satisfied by randomizing the estimation of the conditional probabilities.

## 3.2 Data-driven procedure

Before considering the reject option arguments and since we deal with an Early Exit strategy, we train a neural network with  $M$  possible exits that correspond to  $M - 1$  auxiliary heads and the classical output of the last layer, as illustrated in Figure 1. To this end, we collect a *labeled* dataset  $\mathcal{D}_n$ , that consists of  $n \in \mathbb{N}$  *i.i.d.* copies of  $(\mathbf{X}, Y)$ , and train for each head  $\ell$  an estimator  $(\hat{p}_1^\ell, \dots, \hat{p}_K^\ell)$  of the conditional probability vector  $(p_1, \dots, p_K)$ . Since larger  $\ell$  means that we go deeper in the network, it is reasonable to assume that for all layers indices  $\ell, \ell' \in [M - 1]$  with  $\ell < \ell'$ , the auxiliary head  $\ell'$  consumes more resources than the auxiliary head  $\ell$  and in general, this increase in consumption might come with a better accuracy.

### 3.2.1 Classification rule based on reject option

Our methodology highlights that early exiting can be efficiently performed, borrowing tools from learning with reject option. Let us then define, for each auxiliary head  $\ell$ , a rejection rate  $1 - \varepsilon^\ell \in (0, 1)$  – or a non abstention rate  $\varepsilon^\ell$ . In order to specify the prediction rule for each head, we will mimic the optimal rule provided by Proposition (1) using the plug-in principle. This step requires collecting a sample of *unlabeled* instances  $\mathcal{D}_N$ , that consists in  $N$  *i.i.d.* copies of  $\mathbf{X}$ . Since the process is the same at each auxiliary head, let us develop our methodology for a specific head  $\ell$ .

The goal is to understand whether the head  $\ell$  is a good early exit for a given instance  $\mathbf{x} \in \mathcal{X}$ . If we translate this into the classification with reject option vocabulary, the question becomes whether the classifier  $\hat{g}^\ell(\cdot) = \operatorname{argmax}_{k \in [K]} \{\hat{p}_k^\ell(\cdot) + u_k\}$ , should classify the instance  $\mathbf{x}$  or reject it. The  $u_k$  variable is introduced to randomize the  $\hat{p}_k^\ell$  estimations for a technical reason that we now explain.

Let us denote by  $(u_k)_{k \in [K]}$  *i.i.d.* variables which follow a uniform distribution on  $[0, u]$  with  $u$  being a non-negative real number which is usually chosen very

small. (In practice, we set  $u = 10^{-5}$ .) It ensures that the random variables  $p_k^\ell(\mathbf{X}) + u_k$  has no atoms and is actually the key argument (see proof in appendix B) to have a good control on the rejection rate of the produced classifier with reject option. Formally, this classifier is the empirical counterpart of the classifier with reject option given by Proposition 1 and is based on the empirical score function given for all  $\mathbf{x} \in \mathcal{X}$  by

$$\hat{s}^\ell(\mathbf{x}) = \max_{k \in [K]} \{ \hat{p}_1^\ell(\mathbf{x}) + u_1, \dots, \hat{p}_K^\ell(\mathbf{x}) + u_K \} . \quad (3)$$

we can derive the plug-in estimator of the classifier with reject option  $h_\varepsilon^*$  at the level of the  $\ell$ -th auxiliary head.

**Definition 2** For all  $\mathbf{x} \in \mathcal{X}$

$$\hat{h}_{\varepsilon^\ell}^\ell(\mathbf{x}) = \begin{cases} \hat{g}^\ell(\mathbf{x}) & \text{if } \hat{F}_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{x})) \geq 1 - \varepsilon^\ell \\ \mathfrak{R} & \text{otherwise,} \end{cases}$$

where conditional on the dataset  $\mathcal{D}_n$ , we denote by  $\hat{F}_{\hat{s}^\ell}(t) = \frac{1}{N} \sum_{\mathbf{x} \in \mathcal{D}_N} \mathbf{1}_{\{\hat{s}^\ell(\mathbf{x}) \leq t\}}$  the empirical CDF of  $\hat{s}^\ell(\mathbf{X})$  on the sample  $\mathcal{D}_N$ .

**Remark 2** Notably, the score function  $\hat{s}^\ell$  in the above definition can be replaced by any suitable metrics providing an estimation of the prediction confidence. We tried entropy-based confidence and breaking ties [22] (which is the difference between the two highest scores) in our experiments and select the latter which performed slightly better empirically.

According to the above definition, the classifier  $\hat{h}_{\varepsilon^\ell}^\ell$  abstains when it is not confident about the classification. In our case, having  $\hat{h}_{\varepsilon^\ell}^\ell$  assigns the output  $\mathfrak{R}$  to a given instance  $\mathbf{x}$  means that the observation should not be treated by  $\hat{h}_{\varepsilon^\ell}^\ell$ , but rather should be delayed to next auxiliary head treatment. In contrast, when  $\hat{h}_{\varepsilon^\ell}^\ell(\mathbf{x}) = \hat{g}^\ell(\mathbf{x})$ , then we use the early exit at the head  $\ell$  and the observation  $\mathbf{x}$  does not go through the rest of the network, thus reducing the computation. We can establish the following result whose proof can be found in the Section B of the Appendix.

**Proposition 2** For all  $\ell \in [M]$  and all  $\varepsilon^\ell \in (0, 1)$ , there exists a constant  $C > 0$  such that, whatever the distribution  $\mathbb{P}$  of the data and whatever the estimators  $\hat{p}_k^\ell$  of the conditional probabilities we consider, we have

$$\left| \mathbb{P} \left( \hat{h}_{\varepsilon^\ell}^\ell(\mathbf{X}) = \mathfrak{R} \right) - (1 - \varepsilon^\ell) \right| \leq \frac{C}{\sqrt{N}} .$$

The above result confirms that the rejection rate of the classifier  $\hat{h}_{\varepsilon^\ell}^\ell$  is indeed of the right order. However, this result suggests that the head  $\ell$  might reject more data than it should (by a proportion of order  $C/\sqrt{N}$  which is not suitable from the budget perspective. In order to solve this issue, it is sufficient to impose a smaller rate of rejection. More precisely, if we replace  $\varepsilon^\ell$  by  $\tilde{\varepsilon}^\ell = \varepsilon^\ell + C/\sqrt{N}$  when we run our algorithm, we force the rejection rate to be less than  $1 - \varepsilon^\ell$ . From our numerical study, we observed that  $C = \varepsilon^\ell$  leads to good results in order to ensure the budget limitation.



**Remark 3** *Our methodology, applicable for semi-supervised learning, capitalizes on both labeled and unlabeled data, making it ideal when acquiring labels is costly. If only labeled data is available, we advise splitting the dataset, facilitating budget control theoretically.*

**Remark 4** *The result in Proposition 2 is valid when  $M > 2$  with the assumption that previous heads did not reject data. In a dynamic inference process, we need to adjust the rejection rate at the level of the head  $\ell$  based on the rejection rates of the earlier heads. We will elaborate on this in Section 3.2.3.*

### 3.2.2 Calibration of the rejection rates

To fully exploit the layered complexity of deep learning networks, our EERO methodology must handle multiple exit. We incorporate an aggregation with exponential weights to optimize the decision-making process at various network depths. This adaptation is crucial for leveraging the diverse representational capabilities of neural networks, ensuring computational efficiency, and maintaining high accuracy across multiple exit points. We recall that our main constraint here is the maximum allowed budget  $B$ , in GFlops.

Let us then develop the process to build the vector  $\hat{\varepsilon} = (\hat{\varepsilon}^1, \dots, \hat{\varepsilon}^M)$  of discrete probabilities that provides us the rates of classifying at each head (in other words, the  $1 - \hat{\varepsilon}^\ell$  are the rejection rates). As inputs, we assume

- we have already trained all heads classifiers that are called  $\hat{g}^1 \dots \hat{g}^M$ ;
- we have already computed for each of the classifiers  $\hat{g}^\ell$  an evaluation of its risk  $\hat{R}^\ell = \hat{R}_n(\hat{g}^\ell) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{g}^\ell(\mathbf{x}_i) \neq Y_i\}}$  based on the training set  $\mathcal{D}_n$  – notice that these error rates are already computed during any classic deep learning training;
- we have a prior distribution  $\boldsymbol{\pi} = (\pi^1, \dots, \pi^M)$  on the simplex  $\Lambda_{M-1}$  defined as  $\pi^\ell = \frac{(\hat{B}^\ell)^{-1}}{\sum_{j=1}^M (\hat{B}^j)^{-1}}$ , where  $\hat{B}^\ell$  is the budget required by the head classifier  $\hat{g}^\ell$  to provide an inference for one instance,
- we have fixed the overall budget  $B$  we are allowed to use and the size  $T$  of the batch of new data points we need to predict.

Our proposal is based on aggregation with exponential weights. The vector  $\hat{\varepsilon} = (\hat{\varepsilon}^1, \dots, \hat{\varepsilon}^M) \in \Lambda_{M-1}$  is solution in  $\boldsymbol{\varepsilon} = (\varepsilon^1, \dots, \varepsilon^M)$  of the following minimization problem:

$$\min_{\boldsymbol{\varepsilon} \in \mathbb{R}^M} \sum_{\ell=1}^M \varepsilon^\ell \hat{R}^\ell + \beta \sum_{\ell=1}^M \varepsilon^\ell \log \left( \frac{\varepsilon^\ell}{\pi^\ell} \right), \quad (4)$$

such that

$$\varepsilon^\ell \geq 0, \quad \sum_{\ell=1}^M \varepsilon^\ell = 1, \quad \sum_{\ell=1}^M \varepsilon^\ell \hat{B}^\ell \leq \bar{B}, \quad (5)$$

where  $\beta \geq 0$  is a tuning parameter that controls the strength of the Kullback-Leibler divergence between  $\varepsilon$  and  $\pi$  and  $\bar{B} = B/T$  is the average budget we can spend to infer one instance. Notice that the constraint  $\sum_{\ell=1}^M \varepsilon^\ell \hat{B}^\ell \leq \bar{B}$  reads as  $\sum_{\ell=1}^M (T\varepsilon^\ell) \hat{B}^\ell \leq B$ . In particular,  $T\varepsilon^\ell$  interprets as the number of data points that should be classified by the head  $\ell$  so that the total budget remains less (or equal) than the allocated budget  $B$ . Considering the Lagrangian of this problem, we can exhibit the following form of the probability vector  $\hat{\varepsilon}$ .

**Proposition 3** *For all  $\ell \in [M]$ , the  $\ell$ -th coordinates of the rejection rates vector  $\hat{\varepsilon}$  is given by*

$$\hat{\varepsilon}^\ell = \frac{\pi^\ell \exp \left\{ -\frac{\hat{R}^\ell + \hat{\mu} \hat{B}^\ell}{\beta} \right\}}{\sum_{j=1}^M \pi^j \exp \left\{ -\frac{\hat{R}^j + \hat{\mu} \hat{B}^j}{\beta} \right\}},$$

where  $\hat{\mu} = \max\{0, \bar{\mu}\}$ , with  $\bar{\mu}$  being solution of the equation

$$\sum_{\ell=1}^M (\bar{B} - \hat{B}^\ell) \pi^\ell \exp \left\{ -\frac{\hat{R}^\ell + \bar{\mu} \hat{B}^\ell}{\beta} \right\} = 0 .$$

The only tuning parameter in the above procedure is the temperature  $\beta$ . Higher values force the probability vector  $\hat{\varepsilon}$  to get closer to the prior distribution  $\pi = (\pi^1, \dots, \pi^M)$  that is created to take into account the budget required by each head to produce a prediction.

### 3.2.3 Heads with reject option

In the case of multiple heads, it is difficult to ensure the validity of Proposition 2 for all heads (*c.f.* Remark 4). However, we can guarantee a weaker result that is sufficient to ensure the good control on the allocated budget. Based on the previous section, we built  $M$  head classifiers on one hand and a probability vector  $\hat{\varepsilon}$  whose  $\ell$ -th components gives the rates of classification at the  $\ell$ -th auxiliary head on the other hand. The probability vector  $\hat{\varepsilon}$  takes into account both the accuracy and the resources of each head and allows achieving the suitable control on the budget. However, in the case of multiple auxiliary heads, a careful analysis of our methodology imposes some modifications according to the sequential calibration of the probabilities of classification.

The calibration of the classification rates of all heads is based on the estimation of the CDF  $F_s$  of the score function. Since there are  $M - 1$  auxiliary heads, we estimate  $M - 1$  times the function  $F_s$ . Importantly, all these estimates are built on the same dataset  $\mathcal{D}_N$ . While the calibration at the level of the first head is perfectly valid, the calibration of the classification rate starting from the second head needs to be adjusted to get the  $\hat{\varepsilon}^\ell$  classification rate. In particular, for all  $\ell \in \{2, \dots, M - 1\}$ , we enforce a higher classification rate as

$$\hat{\varepsilon}_{\text{seq}}^\ell = \sum_{j=1}^{\ell} \hat{\varepsilon}^j . \quad (6)$$

---

**Algorithm 1:** EERO method : calibration phase and classification of a batch of data points

---

**Input:** Batch of data points:  $\mathbf{X}_1, \dots, \mathbf{X}_T \in \mathcal{X}$ ; Calibration sets  $\mathcal{D}_N$ ;  
Model  $p$  with  $M$  heads, Risk of heads:  $R$ , Allowed budget:  $B$

**Output:** Prediction:  $\mathcal{P} \in [K]^T$

$F_s \leftarrow \text{ComputeCDF}(\mathcal{D}_N)$ ; // Compute empirical CDF of score  
 $s \{(3)\}$

$\tilde{\varepsilon}_{\text{seq}} \leftarrow \text{AggregateEpsilons}(M, R, B)$ ; // {Prop: 3, (6), (7) }  
 $\mathcal{P} \leftarrow$  empty list

**for**  $i = 1$  **to**  $T$  **do**

**for**  $\ell = 1$  **to**  $M$  **do**

$s^\ell, g^\ell \leftarrow \text{ComputeOutput}(\mathbf{X}_i, p_\ell)$ ; // Get score and class at  
 $\ell \{(3), (1)\}$

**if**  $F_{s^\ell}(s^\ell(\mathbf{x})) \geq 1 - \tilde{\varepsilon}_{\text{seq}}^\ell$  **then**

$\mathcal{P} \leftarrow [\mathcal{P}, g^\ell]$ ; // save the classification from head  $\ell$

**break**

**end**

**end**

**end**

**return**  $\mathcal{P}$

---

This choice is motivated by the fact that we need, for each head  $\ell$ , to calibrate the threshold for the rejection rule so that at most a proportion  $1 - \sum_{j=1}^{\ell} \hat{\varepsilon}^j$  of the data in  $\mathcal{D}_N$  is rejected.

If we consider this adjustment together with the correction we have detailed right after Proposition 2, we can show that for all  $\ell \in [M]$ , if we replace the probability of classification  $\hat{\varepsilon}^\ell$  in Definition 2 by

$$\tilde{\varepsilon}_{\text{seq}}^\ell = \hat{\varepsilon}_{\text{seq}}^\ell + C/\sqrt{N} \text{ ,} \quad (7)$$

we show that

$$\mathbb{P} \left( \hat{h}_{\tilde{\varepsilon}_{\text{seq}}^\ell}^\ell(\mathbf{X}) = \mathfrak{R} \right) \leq 1 - \sum_{j=1}^{\ell} \hat{\varepsilon}^j = \sum_{j=\ell+1}^M \hat{\varepsilon}^j \text{ ,}$$

meaning that after using the head  $\ell$ , there is at most a proportion  $\sum_{j=\ell+1}^M \hat{\varepsilon}^j$  of the data that remains to classify. This result will be illustrated in the next section through numerical experiments.

## 4 Experiments

Building upon the methodological foundations established for EERO, this section seeks to empirically substantiate the approach. We initiate with an exploration of a single-head EERO within the ResNet-18 architecture, employing the CIFAR dataset. The scope then broadens to encapsulate the multi-head EERO variant,

tested against the more recent ConvNext and MSDNet architectures on the expansive ImageNet dataset. Those experiments validate our theory and demonstrate the scalability of EERO across diverse neural network configurations.

#### 4.1 EERO with only 1 auxiliary head

In the case of one auxiliary head, computations of the rate of classifying  $\hat{\varepsilon}^1$  (of this single head) can be simplified. In particular, we do not require aggregation with exponential weights to find it and even more, we can express it analytically. Let  $B > 0$  be the amount of *GFlops* we were allocated for the task of labeling  $T$  inputs vectors  $\mathbf{X}_1, \dots, \mathbf{X}_T \in \mathcal{X}$ . Assume that the classifiers  $\hat{h}_{\hat{\varepsilon}^1}^1$  and  $\hat{g}^2$  burn off respectively  $\hat{B}^1$  and  $\hat{B}^2$  *GFlops* at each call with  $\hat{B}^1 < \hat{B}^2$ . (We assume that both  $\hat{B}^1$  and  $\hat{B}^2$  are small as compared to the total budget  $B$  and that  $B \geq T\hat{B}^1$  so that we are guaranteed to label all  $T$  points). Let  $\hat{\varepsilon}^1$  be such that  $\hat{\varepsilon}^1 T \hat{B}^1 + (1 - \hat{\varepsilon}^1) T \hat{B}^2 = B$ . We can state the following result.

**Proposition 4** *Let  $\hat{\varepsilon}^1 = \frac{\frac{B}{T} - \hat{B}^2}{\hat{B}^1 - \hat{B}^2}$ . Then the total amount of GFlops used to label  $T$  instances is not larger than  $B$ .*

This result highlights that our strategy succeeds to comply with the constraint of budget we imposed. The proof of the result lies in the fact that we use the early exit on a proportion  $\hat{\varepsilon}^1$  of the data. Then we consume  $T\hat{\varepsilon}^1\hat{B}^1$  *GFlops* (up to a  $1/\sqrt{N}$  additive term – see the comment after Proposition 2 for a correction). The rest of the data is treated by the classifier  $\hat{g}^2$  and then uses  $T(1 - \hat{\varepsilon}^1)\hat{B}^2$  *GFlops*. In total, we then get

$$T(\hat{\varepsilon}^1 \hat{B}^1 + (1 - \hat{\varepsilon}^1) \hat{B}^2) = T(\hat{B}^2 + \hat{\varepsilon}^1 (\hat{B}^1 - \hat{B}^2)) = T \left( \hat{B}^2 + \frac{\frac{B}{T} - \hat{B}^2}{\hat{B}^1 - \hat{B}^2} (\hat{B}^1 - \hat{B}^2) \right) = B . \quad (8)$$

**Remark 5** *It is important to observe that in this case where we have only one head, we enforce the rejection rate to exactly burn off the whole budget  $B$ . This is highlighted in Equation (8). We make this choice to explain some interesting broader effects that illustrate the importance of using auxiliary heads – see below and Figure 2. On the other hand, we notify that it is extremely simple to modify the algorithm so that it consumes less than or equal to the total budget, that is,  $(\hat{\varepsilon}^1 T) \hat{B}^1 + ((1 - \hat{\varepsilon}^1) T) \hat{B}^2 \leq B$ . In this case, the accuracy curve would always be increasing w.r.t. the budget.*

We then implement <sup>1</sup> our procedure with one Early Exit on a ResNet [12] model with Cifar-100 dataset [16] (see Figure 1). We test 7 different positions for the early exit by adding auxiliary heads at regular intervals along the network.

<sup>1</sup>All computations are run on a server with an Intel(R) Xeon(R) Gold 5120 CPU and a Tesla V100 GPU with 32GB of Vram and 64GB of RAM. Code associated with paper can be found in repository here : <https://anonymous.4open.science/r/Early-Exit-With-Reject-Option-9896>

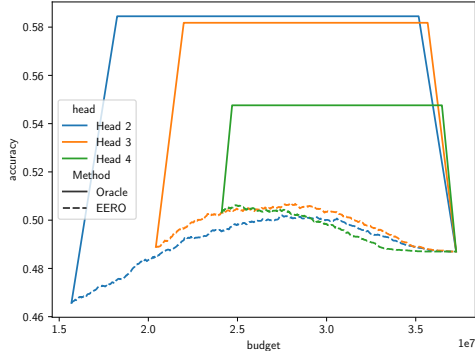


Figure 2: Accuracy versus computation budget for our multi-headed ResNet-18 (given in Figure 1). In this illustration, we consider only a two heads procedure by selecting a specific auxiliary head – 2, 3, or 4 – among the initial 7. Each curve shows how the accuracy evolves according to the probability of use of the auxiliary head. For each color, we display our algorithm as well as the oracle.

In more details, we adopt a simple approach and flatten the outputs of the convolution layers and append a Multilayer Perceptron with one hidden layer which has the same size as the output dimension. We then train the model over the sum of the cross entropy losses from the different exits. We use a simple gradient descent algorithm with a learning rate of 0.003 and a batch size of 128. In the CIFAR100 dataset, we partition the original 50,000 training examples into a training set  $\mathcal{D}_n$  with 49,000 examples and a calibration set  $\mathcal{D}_N$  of 1,000 examples used to compute the  $\tilde{\varepsilon}_{seq}^1$  variables – see Equation (7) – leaving  $\mathcal{D}_T$ , the test set, with 5,000 examples. We also implement an oracle which will select the best head to maximize the accuracy for a given computation budget. To build this oracle, we solve this problem with Integer Linear Programming, where the auxiliary head selection is modeled as binary variables.

We plot the accuracy for the different exits in Figure 2 for different budgets.

First, we can see a general trend where our procedure gradually improves the accuracy as the budget increases since images wrongly classified by the auxiliary head and correctly classified by the final layer are forwarded from the Early Exit to the latter. Then there is an accuracy decrease since augmenting the budget only affects images which are wrongly classified by both Early Exit and final model exit. Secondly, we notice that the oracle quickly outperforms the accuracy of the main model for only a fraction of the total budget required to process all images with the full model.

This fact shows that Early Exits can have valid predictions, whereas the last model output is wrong. This phenomenon is known as over-thinking and motivates the potential of Early Exit methods, as a way of better understanding the flaws of deep learning models. Here, this over-thinking situation seems

important, as some of the Early Exits have a better accuracy than the last model output, *c.f.*, Remark 5. A second reason explaining the oracle improvements is its capacity to save computation budget by choosing an Early Exit when both early and last exit predictions are wrong.

This advantage, only achievable in our specific context of Budgeted Batch Classification, aligns well with our primary objective of reducing computational power for our algorithms. Furthermore, it indirectly enhances the overall accuracy of our model. The next sections validate our finding with the more realistic case of multiple exits with the ImageNet dataset.

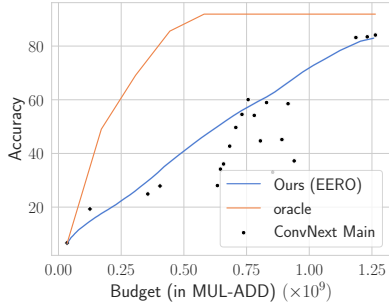
## 4.2 EERO with multiple heads

We now consider the case of multiple heads. We recall that the procedure consists of two steps. The first one focuses on determining the rejection rates at each head. The second one consists in specifying the rejection rule at the level of each given head, *c.f.*, Definition 2. While in Section 4.1, we consider two heads ( $M = 2$ ) and then there is only one rejection rate to compute (*c.f.* Proposition 4), in this case of multiple heads, the computation is more involved. We therefore exploit the methodology based on aggregation with exponential weights developed in Section 3.2.2.

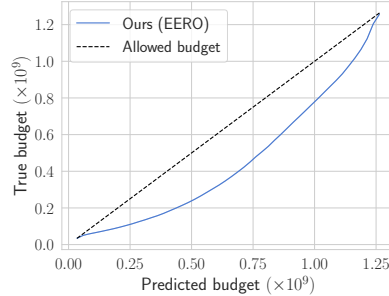
We follow a similar protocol as in Section 4.1 with the more recent ConvNext Architecture and the ImageNet dataset where we extract 5,000 images from the 50,000 images in the validation set as our calibration set. More precisely, we flattened the outputs of the convolution layers and append a Multilayer Perceptron with one hidden layer which has the same size as the input dimension. To avoid an increase of the computational cost, we apply Adaptive Average pooling so that the input dimension remains below 3,000. To train this model, we used pretrained weights on ImageNet and fine tune only different heads for 300 epochs with a 0.9 momentum and a weight decay  $10^{-4}$ . This corresponds to a classical use-case where an Early Exit application needs to be applied to a carefully trained model which would be difficult to modify. We use the same oracle as in Section 4.1 – with the exception that we do not enforce the allowed budget to be totally consumed – and follow the algorithm described in 1 to do the inference on the test set.

As in Section 4.1, we plot the results for the different heads and for different computation budgets in Figure 3a. We can notice several similar trends already observed on the Cifar-100 dataset. We can observe that our proposed multiple heads approach (blue curve) improves on the trade-off of the different exits used independently (black dots). Regarding the accuracy of the auxiliary heads alone, placing an exit further in the model does not necessarily result in a better accuracy. Lastly, the oracle shows a great potential in combining these different exits, reaching a final accuracy of 89% which is comparable to the state-of-the-art on ImageNet and with only 60% of the budget used by the original ConvNext.

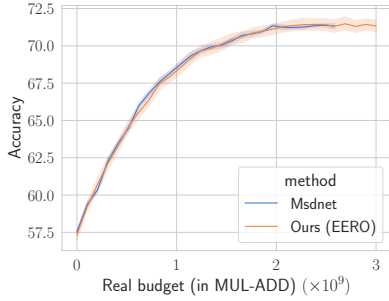
We further check the budget used by our model in Figure 3b and show that our methodology respects the allowed budget by always being inferior or equal to it, validating our theoretical result from Proposition 2. On some budget, The



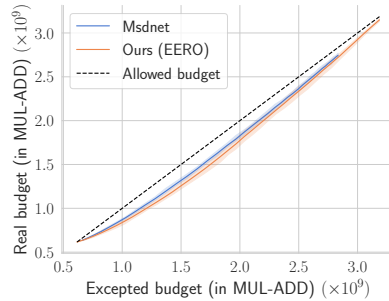
(a) Accuracy *w.r.t.* the budget for our EERO methodology based on Convnext – comparison with the oracle.



(b) Measured budget of our algorithm with Convnext architecture.



(c) Accuracy *w.r.t.* the real budget based on MSDNet architecture – comparison of EERO methodology to the original MSDNet algorithm.



(d) Measured budget on MSDNet architecture – comparison of EERO methodology to the original MSDNet algorithm.

Figure 3: Evaluation of EERO on Convnext and MSDNet architectures: Subfigures (a) and (b) show EERO’s accuracy and budget on Convnext, while subfigures (c) and (d) compare EERO with MSDNet’s original method, focusing on accuracy and budget metrics.

gap between the allowed budget and the actual consumption suggests that a better accuracy could be obtained.

To illustrate how  $\varepsilon$  is evolving we provide Table 1 in Appendix with different budget which range from a tight budget, where we have to focus more on the first heads to a more permissive one which allow using the latest heads. Lastly, by inspecting the  $\hat{\varepsilon}$  values, we noted that our procedure effectively exploits the information from the Risks  $\hat{R}^\ell$  to better calibrate the probabilities of classifying *i.e.*, heads which do offer the best accuracy will be used more often as Early Exits.

### 4.3 MSDNet Comparison

Building on these insights, we extended our investigation to include a comparative analysis with the Multi-Scale Dense Network (MSDNet) [14]. This comparison not only benchmarks the accuracy of our EERO method against MSDNet, but also evaluates how effectively each approach adheres to the computational budget. The results, as depicted in Figures 3c and 3d, reinforce the findings observed with other architectures. It’s noteworthy that in MSDNet, the input parameter is the distribution of a control variable  $\epsilon$ , rather than the computational budget directly. To achieve a specific budget target using MSDNet, one needs to iteratively run the method with various  $\epsilon$  values until finding the one that aligns with the desired budget. In contrast, EERO directly takes the computational budget as its input parameter, thereby offering a more efficient and straightforward approach to budget-constrained optimization.

Our experiments demonstrate that, while MSDNet presents a robust framework, our EERO approach offers distinct advantages. Firstly, it exhibits a commendable level of accuracy, closely paralleling the performance of MSDNet. This is particularly notable in Figure 3c, where the accuracy of EERO consistently aligns with or exceeds that of MSDNet across various budget thresholds.

EERO demonstrates a significant capability in adhering to the computational budget. Figure 3d illustrates that EERO not only meets but often remains below the allocated budget. This efficiency, reinforced by our theoretical results from Proposition 2, highlights its effectiveness. It is important to note that both EERO and MSDNet exhibit comparable levels of efficiency and adherence to budget constraints. However, a key distinction lies in the theoretical grounding of EERO’s budget compliance. Our framework provides a formal proof of budget respect, which adds a layer of reliability and predictability to its performance. This theoretical assurance of budget adherence sets EERO apart from MSDNet, enhancing its practical applicability in scenarios where strict budget compliance is critical.

An additional key advantage of our EERO methodology is its generalizability. Unlike MSDNet, which is tailored to specific architectural constructs, EERO can be applied broadly across a variety of models equipped with multiple output heads. This versatility enhances EERO’s applicability in diverse real-world scenarios, where the ability to adapt to different architectures without compromising on budget adherence or accuracy is crucial.

In summary, our comparative analysis with MSDNet validates the efficacy of the EERO approach. Not only does it maintain competitive accuracy, but it also demonstrates adherence to computational budgets and boasts a level of generalizability that positions it as a versatile tool in the realm of efficient neural network inference.



## 5 Conclusion

In this paper, we presented EERO, a novel mathematical framework devised to construct classification rules for Early Exit in deep learning networks. This framework targets the pivotal challenge of optimizing computational efficiency while enhancing model performance. EERO approaches the problem by modeling it as a classification with a reject option, providing a feasible solution for scenarios involving a single Early Exit. For multi-exit scenarios, our innovative approach hinges on an aggregation procedure employing exponential weights. A key strength of EERO lies in its flexibility and generalizability, enabling it to be applied across various model architectures. This adaptability makes EERO a versatile tool, well-suited for diverse applications in the realm of efficient deep learning.

Our extensive experiments on ImageNet and Cifar-100 using the Resnet-18, MSDNet and the ConvNext architectures respectively demonstrated the efficacy of our proposed framework. These results not only shed light on the intricate workings of multi-head deep neural networks but also offer practical strategies to enhance model accuracy while reducing computational budget. This has important implications for edge computing and making the field of deep learning more sustainable.

Despite our promising results, the study presents opportunities for further refinement. For instance, the disparity between our model’s performance and the oracles’ suggests potential areas of improvement in our methodology. Future work could, therefore, focus on developing more sophisticated strategies for optimizing Early Exits and further closing this performance gap.

In conclusion, our study charts a new course for exploiting multi-head deep learning models, delivering promising results and setting clear targets for future work. We anticipate that our findings will stimulate further research in this direction, paving the way for more efficient and sustainable deep learning practices.

## References

- [1] A. Bakhtiarnia, Q. Zhang, and A. Iosifidis. Multi-exit vision transformer for dynamic inference. *arXiv preprint arXiv:2106.15183*, 2021.
- [2] T. Bolukbasi, J. Wang, O. Dekel, and V. Saligrama. Adaptive neural networks for fast test-time prediction. *arXiv preprint arXiv:1702.07811*, 1(3), 2017.
- [3] X. Chen, H. Dai, Y. Li, X. Gao, and L. Song. Learning to stop while learning to predict. In *International Conference on Machine Learning*, pages 1520–1530. PMLR, 2020.
- [4] C. Chow. An optimum character recognition system using decision functions. *IRE Transactions on Electronic Computers*, EC-6(4):247–254, 1957.
- [5] C. Cortes, G. DeSalvo, and M. Mohri. Learning with rejection. In *International Conference on Algorithmic Learning Theory*, pages 67–82. Springer, 2016.
- [6] C. Denis and M. Hebiri. Consistency of plug-in confidence sets for classification in semi-supervised learning. *Journal of Nonparametric Statistics*, 2019.

- [7] C. Denis, M. Hebiri, B. Njike, and X. Siebert. Active learning algorithm through the lens of rejection arguments. *arXiv preprint arXiv: arXiv:2208.14682*, 2022.
- [8] C. Denis, M. Hebiri, and A. Zaoui. Regression with reject option and application to knn. *arXiv preprint arXiv:2006.16597*, 2020.
- [9] M. Figurnov, M. Collins, Y. Zhu, L. Zhang, J. Huang, D. Vetrov, and R. Salakhutdinov. Spatially adaptive computation time for residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1039–1048, 2017.
- [10] Y. Grandvalet, A. Rakotomamonjy, J. Keshet, and S. Canu. Support vector machines with a reject option. In *NIPS*, pages 537–544, 2009.
- [11] Y. Han, G. Huang, S. Song, L. Yang, H. Wang, and Y. Wang. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [13] R. Herbei and M. Wegkamp. Classification with reject option. *Canad. J. Statist.*, 34(4):709–721, 2006.
- [14] G. Huang, D. Chen, T. Li, F. Wu, L. Van Der Maaten, and K. Weinberger. Multi-scale dense networks for resource efficient image classification. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net, 2018.
- [15] Y. Kaya, S. Hong, and T. Dumitras. Shallow-deep networks: Understanding and mitigating network overthinking. In *International conference on machine learning*, pages 3301–3310. PMLR, 2019.
- [16] A. Krizhevsky, V. Nair, and G. Hinton. Cifar-100 (canadian institute for advanced research).
- [17] S. Laskaridis, A. Kouris, and N. Lane. Adaptive inference through early-exit networks: Design, challenges and directions. In *Proceedings of the 5th International Workshop on Embedded and Mobile Deep Learning*, pages 1–6, 2021.
- [18] J. Lei. Classification with confidence. *Biometrika*, 101(4):755–769, 2014.
- [19] T. Liang, J. Glossner, L. Wang, S. Shi, and X. Zhang. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing*, 461:370–403, 2021.
- [20] S. Lin, B. Ji, R. Ji, and A. Yao. A closer look at branch classifiers of multi-exit architectures. *arXiv preprint arXiv:2204.13347*, 2022.
- [21] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022.
- [22] T. Luo, K. Kramer, D. Goldgof, L. Hall, S. Samson, A. Remsen, T. Hopkins, and D. Cohn. Active learning to recognize multiple types of plankton. *Journal of Machine Learning Research*, 6(4), 2005.
- [23] M. Naadeem, J.D. Zucker, and B. Hanczar. Accuracy-rejection curves (ARCs) for comparing classification methods with a reject option. In *MLSB*, pages 65–81, 2010.
- [24] S. Teerapittayanon, B. McDanel, and H.-T. Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2464–2469. IEEE, 2016.
- [25] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021.

- [26] A. Veit and S. Belongie. Convolutional networks with adaptive inference graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [27] V. Vovk, A. Gammerman, and C. Saunders. Machine-learning applications of algorithmic randomness. In *In Proceedings of the Sixteenth International Conference on Machine Learning*, pages 444–453. Morgan Kaufmann, 1999.
- [28] V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic learning in a random world*. Springer, New York, 2005.
- [29] X. Wang, F. Yu, Z.-Y. Dou, T. Darrell, and J. Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 409–424, 2018.
- [30] Y. Wang, R. Huang, S. Song, Z. Huang, and G. Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. *Advances in Neural Information Processing Systems*, 34:11960–11973, 2021.
- [31] Z. Wu, T. Nagarajan, A. Kumar, S. Rennie, L. Davis, K. Grauman, and R. Feris. Blockdrop: Dynamic inference paths in residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8817–8826, 2018.
- [32] M. Yuan and M. Wegkamp. Classification methods with reject option based on convex risk minimization. *J. Mach. Learn. Res.*, 11:111–130, 2010.
- [33] W. Zhou, C. Xu, T. Ge, J. McAuley, K. Xu, and F. Wei. Bert loses patience: Fast and robust inference with early exit. In *Advances in Neural Information Processing Systems*, volume 33, pages 18330–18341, 2020.

## Supplementary material

The supplementary material is organized as follows. In Appendix A we provide an additional numerical description of our algorithm. The proof of our main results are provided in Section B.

### A Figures

			6.55 * 10 <sup>12</sup> Flops		3.21 * 10 <sup>13</sup> Flops		5.64 * 10 <sup>13</sup> Flops	
$\ell$	$\hat{R}^\ell$	$\pi^\ell$	$\hat{\varepsilon}^\ell$	$\hat{\varepsilon}_{\text{out}}^\ell$	$\hat{\varepsilon}^\ell$	$\hat{\varepsilon}_{\text{out}}^\ell$	$\hat{\varepsilon}^\ell$	$\hat{\varepsilon}_{\text{out}}^\ell$
1	0.93	0.46	0.69	0.68	0.18	0.17	0.00	0.00
2	0.81	0.13	0.17	0.28	0.07	0.19	0.00	0.01
3	0.75	0.04	0.03	0.04	0.03	0.16	0.00	0.01
4	0.72	0.04	0.03	0.00	0.03	0.12	0.00	0.01
5	0.72	0.02	0.01	0.00	0.03	0.08	0.00	0.02
6	0.66	0.02	0.01	0.00	0.03	0.07	0.00	0.01
7	0.64	0.02	0.01	0.00	0.03	0.05	0.01	0.02
8	0.57	0.02	0.01	0.00	0.03	0.04	0.01	0.02
9	0.51	0.02	0.01	0.00	0.04	0.03	0.01	0.02
10	0.46	0.02	0.01	0.00	0.04	0.02	0.01	0.01
11	0.40	0.02	0.01	0.00	0.04	0.02	0.01	0.02
12	0.46	0.02	0.01	0.00	0.04	0.03	0.01	0.04
13	0.55	0.02	0.00	0.00	0.03	0.02	0.01	0.03
14	0.41	0.02	0.01	0.00	0.04	0.01	0.02	0.03
15	0.67	0.02	0.00	0.00	0.03	0.01	0.02	0.08
16	0.55	0.02	0.00	0.00	0.04	0.00	0.02	0.15
17	0.42	0.02	0.00	0.00	0.04	0.00	0.03	0.09
18	0.62	0.02	0.00	0.00	0.03	0.00	0.03	0.13
19	0.17	0.01	0.00	0.00	0.06	0.00	0.20	0.05
20	0.17	0.01	0.00	0.00	0.06	0.00	0.27	0.06
21	0.16	0.01	0.00	0.00	0.06	0.00	0.33	0.20

Table 1: Illustration of the evolution of the probability vector of classification for the different heads on ImageNET dataset with Convnext architecture and for three different budgets:  $\pi$  is the prior distribution,  $\hat{\varepsilon}$  are the weights produced by the aggregation, and  $\hat{\varepsilon}_{\text{out}}$  are the proportion of the data that was actually classified by each head.

The above Table 1 highlights that: i) when higher budget are allocated, the weights produced by the aggregation put higher weights on heads with lower risks; ii) heads with high risks are assigned smaller weights by the aggregation, see for instance the weights corresponding to the heads 13 and 18 that have a higher risk and so a lower  $\varepsilon^\ell$ .

## B Proofs

In this section, we provide the proofs of our main results.

**Proof 1 (Proposition 1)** *Let us write*

$$\mathcal{R}(h) = \mathbb{P}(h(\mathbf{X}) \neq Y, h(\mathbf{X}) \neq \mathfrak{A}),$$

for short. We then need to solve the problem

$$h_\varepsilon^* = \operatorname{argmin}_h \{ \mathcal{R}(h) : \mathbb{P}(h(\mathbf{X}) = \mathfrak{A}) = 1 - \varepsilon \},$$

where the minimum is taken over all measurable functions. Considering the Lagrangian of the above problem, we solve

$$\min_h \max_{\lambda \in \mathbb{R}^+} \underbrace{\{ \mathcal{R}(h) + \lambda (\mathbb{P}(h(\mathbf{X}) = \mathfrak{A}) - (1 - \varepsilon)) \}}_{:= \mathcal{L}(h, \lambda)},$$

where  $\lambda \in \mathbb{R}$  is the dual variable. Observe that by weak duality we have

$$\min_h \max_{\lambda \in \mathbb{R}^+} \mathcal{L}(h, \lambda) \geq \max_{\lambda \in \mathbb{R}^+} \min_h \mathcal{L}(h, \lambda),$$

we then consider first the minimization problem over  $h$  of  $\mathcal{R}$ . We have

$$\begin{aligned} \mathcal{R}(h) &= \mathbb{P}(h(\mathbf{X}) \neq Y, h(\mathbf{X}) \neq \mathfrak{A}) \\ &= \mathbb{E} \left[ \sum_{k \in [K]} \mathbb{1}_{\{Y=k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq \mathfrak{A}\}} \right] = \mathbb{E} \left[ \sum_{k \in [K]} p_k(\mathbf{X}) \mathbb{1}_{\{h(\mathbf{X}) \neq k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq \mathfrak{A}\}} \right] \cdot \\ &= \mathbb{P}(h(\mathbf{X}) \neq \mathfrak{A}) - \mathbb{E} \left[ \sum_{k \in [K]} p_k(\mathbf{X}) \mathbb{1}_{\{h(\mathbf{X})=k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq \mathfrak{A}\}} \right]. \end{aligned}$$

Moreover,

$$\mathbb{P}(h(\mathbf{X}) \neq \mathfrak{A}) = \mathbb{E} \left[ \sum_{k \in [K]} \mathbb{1}_{\{h(\mathbf{X})=k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq \mathfrak{A}\}} \right].$$

Therefore, we can write

$$\begin{aligned} \mathcal{L}(h, \lambda) &= \lambda - \lambda(1 - \varepsilon) + (1 - \lambda) \mathbb{P}(h(\mathbf{X}) \neq \mathfrak{A}) - \mathbb{E} \left[ \sum_{k \in [K]} p_k(\mathbf{X}) \mathbb{1}_{\{h(\mathbf{X})=k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq \mathfrak{A}\}} \right] \\ &= \lambda \varepsilon - \mathbb{E} \left[ \sum_{k \in [K]} (p_k(\mathbf{X}) - (1 - \lambda)) \mathbb{1}_{\{h(\mathbf{X})=k\}} \mathbb{1}_{\{h(\mathbf{X}) \neq \mathfrak{A}\}} \right]. \quad (9) \end{aligned}$$

We need to maximize the expectation w.r.t.  $h$  that first leads to the optimum  $h_\lambda^*$  is such that

$$h_\lambda^*(\mathbf{x}) \neq \mathfrak{R} \quad \iff \quad \sum_{k \in [K]} (p_k(\mathbf{x}) - (1 - \lambda)) \mathbb{1}_{\{h_\lambda^*(\mathbf{x})=k\}} > 0 ,$$

for all  $\mathbf{x} \in \mathcal{X}$ . Moreover we have that on the event  $\{h_\lambda^*(\mathbf{x}) \neq \mathfrak{R}\}$ , the mapping  $h$  that maximizes  $h \mapsto \sum_{k \in [K]} (p_k(\mathbf{x}) - (1 - \lambda)) \mathbb{1}_{\{h(\mathbf{x})=k\}}$  is simply  $h_\lambda^*(\mathbf{x}) = \operatorname{argmax}_{k \in [K]} p_k(\mathbf{x})$ . At this level of the proof, we have shown that the problem  $\min_h \mathcal{L}(h, \lambda)$  leads to the rule

$$h_\lambda^*(\mathbf{x}) = \begin{cases} \operatorname{argmax}_{k \in [K]} p_k(\mathbf{x}) & \text{if } \max_{k \in [K]} p_k(\mathbf{x}) \geq 1 - \lambda \\ \mathfrak{R} & \text{otherwise .} \end{cases} \quad (10)$$

Now, we deal with the maximization of  $\mathcal{L}(h_\lambda^*, \lambda)$  w.r.t.  $\lambda$ . Substituting the above value of  $h = h_\lambda^*$  in (9), we can show that

$$\mathcal{L}(h_\lambda^*, \lambda) = \lambda \varepsilon - \mathbb{E} \left[ \left( \max_{k \in [K]} \{p_k(\mathbf{X}) - (1 - \lambda)\} \right)_+ \right] ,$$

where for all  $a \in \mathbb{R}$ , we write  $(a)_+ = \max\{a, 0\}$ . The above function is then concave in  $\lambda$ , therefore we can write the first order optimality condition as  $0 \in \partial \mathcal{L}(h_{\lambda^*}^*, \lambda^*)$ . Observe that because we assumed that the r.v.  $p_k(X)$  has no atom for all  $k \in [K]$ , we have that  $\mathbb{P}(\exists j \in [K] : p_k(X) = p_j(X)) = 0$  and then the subgradient reduces to the gradient. As a consequence, we obtain the following condition on  $\lambda^*$

$$\mathbb{P} \left( \exists k \in [K] : p_k(\mathbf{X}) > \max \left\{ \max_{j \in [K]} \{p_j(\mathbf{X})\} ; 1 - \lambda^* \right\} \right) = \varepsilon .$$

Notice that this last condition can rewrite as

$$\varepsilon = \mathbb{P}(s(\mathbf{X}) > 1 - \lambda^*) = \mathbb{P}(h_{\lambda^*}^*(\mathbf{X}) \neq \mathfrak{R}) ,$$

where  $s(\mathbf{X}) = \max_{k \in [K]} p_k(\mathbf{X})$ , which guarantee that the optimal rule has indeed the correct rejection rate. Moreover, from the above relation, we can exhibit the value of  $\lambda$ . Indeed, using the continuity condition on the CDF  $F_s$  of  $s(\mathbf{X})$ , we have

$$\mathbb{P}(s(\mathbf{X}) > 1 - \lambda^*) = \varepsilon \iff 1 - F_s(1 - \lambda^*) = \varepsilon \iff \lambda^* = 1 - F_s^{-1}(1 - \varepsilon) ,$$

where  $F_s^{-1}$  is the generalized inverse of  $F_s$ . We conclude the proof substituting this value into the expression of the optimal rule given by (10).

**Remark 6** If we replace the equality by an inequality, everything is the same expect the part of  $\lambda^*$ . We need to consider the case where  $\lambda^* = 0$  separately, and in this case, we would get  $\mathbb{P}(h_{\lambda^*}^*(\mathbf{X}) \neq \mathfrak{R}) \geq \varepsilon$ .

**Proof 2 (Proposition 2)** First, observe that conditionally to the training dataset  $\mathcal{D}_n$  and due to the continuity of the CDF of  $\hat{s}^\ell(\mathbf{X})$ , the random variable  $F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X}))$  is uniformly distributed. Therefore, for any  $u \in [0, 1]$ , we have  $\mathbb{P}(F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X})) \leq u) = u$ . We then can write

$$\begin{aligned} \left| \mathbb{P}\left(\hat{h}_{\varepsilon^\ell}^\ell(\mathbf{X}) = \mathfrak{R}\right) - (1 - \varepsilon^\ell) \right| &= \left| \mathbb{E}\left[\mathbb{1}_{\{\hat{F}_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X})) \leq 1 - \varepsilon^\ell\}} - \mathbb{1}_{\{F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X})) \leq 1 - \varepsilon^\ell\}}\right] \right| \\ &\leq \left| \mathbb{E}\left[\mathbb{1}_{\{|F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X})) - (1 - \varepsilon^\ell)| \leq |\hat{F}_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X})) - F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X}))|\}}\right] \right| \\ &\leq \left| \mathbb{E}\left[\mathbb{1}_{\{|F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X})) - (1 - \varepsilon^\ell)| \leq \|\hat{F}_{\hat{s}^\ell} - F_{\hat{s}^\ell}\|_\infty\}}\right] \right| \\ &\leq 2\mathbb{E}\left\|\hat{F}_{\hat{s}^\ell} - F_{\hat{s}^\ell}\right\|_\infty, \end{aligned}$$

where we used again in the last line the fact that  $F_{\hat{s}^\ell}(\hat{s}^\ell(\mathbf{X}))$  is uniformly distributed. Moreover, we wrote  $\|\hat{F}_{\hat{s}^\ell} - F_{\hat{s}^\ell}\|_\infty = \sup_{t \in \mathbb{R}} |\hat{F}_{\hat{s}^\ell}(t) - F_{\hat{s}^\ell}(t)|$ . We conclude the proof using the Dvoretzky–Kiefer–Wolfowitz inequality, that states that

$$\mathbb{E}\left\|\hat{F}_{\hat{s}^\ell} - F_{\hat{s}^\ell}\right\|_\infty \leq \sqrt{\frac{\pi}{2n}}.$$

**Proof 3 (Proposition 3)** The Lagrangian of the minimization problem (4)-(5) is given by

$$\mathcal{L}(\varepsilon, \lambda, \mu) = \sum_{j=1}^M \varepsilon^j \hat{R}^j + \beta \sum_{j=1}^M \varepsilon^j \log\left(\frac{\varepsilon^j}{\pi^j}\right) + \lambda \left(\sum_{j=1}^M \varepsilon^j - 1\right) + \mu \left(\sum_{j=1}^M \varepsilon^j \hat{B}^j - \bar{B}\right),$$

with  $(\lambda, \mu) \in \mathbb{R} \times \mathbb{R}^+$ . Considering the KKT condition of the problem we get for all  $j \in [M]$

$$\hat{R}^j + \beta \left(\log\left(\frac{\hat{\varepsilon}^j}{\pi^j}\right) + 1\right) + \hat{\lambda} + \hat{\mu} \hat{B}^j = 0,$$

that leads in turns to

$$\hat{\varepsilon}_{\hat{\lambda}, \hat{\mu}}^j = \pi^j \exp\left(-\frac{\hat{R}^j + \hat{\lambda} + \hat{\mu} \hat{B}^j}{\beta} - 1\right). \quad (11)$$

Plug-in these values into the equality constraint leads to

$$\begin{aligned} \sum_{j=1}^M \hat{\varepsilon}_{\hat{\lambda}, \hat{\mu}}^j = 1 &\iff \sum_{j=1}^M \pi^j \exp\left(-\frac{\hat{R}^j + \hat{\mu} \hat{B}^j}{\beta} - 1\right) = \exp\left(\frac{\hat{\lambda}}{\beta}\right) \\ &\iff \hat{\lambda} = \beta \log\left(\sum_{j=1}^M \pi^j \exp\left(-\frac{\hat{R}^j + \hat{\mu} \hat{B}^j}{\beta} - 1\right)\right). \end{aligned}$$

Substituting back this value into (11), we get

$$\hat{\varepsilon}_{\hat{\mu}}^j = \frac{\pi^j \exp\left(-\frac{\hat{R}^j + \hat{\mu} \hat{B}^j}{\beta}\right)}{\sum_{k=1}^M \pi^k \exp\left(-\frac{\hat{R}^k + \hat{\mu} \hat{B}^k}{\beta}\right)}.$$

According to the parameter  $\hat{\mu}$ , we need to consider the constraints  $\hat{\mu} \geq 0$  and  $\sum_{j=1}^M \hat{\varepsilon}_{\hat{\mu}}^j \hat{B}^j \leq \bar{B}$  together with the complementary condition  $\hat{\mu} \left( \sum_{j=1}^M \hat{\varepsilon}_{\hat{\mu}}^j (\hat{B}^j - \bar{B}) \right) = 0$ . Therefore, when  $\hat{\mu} \neq 0$ , this parameter should be taken such that

$$\sum_{j=1}^M \hat{\varepsilon}_{\hat{\mu}}^j (\hat{B}^j - \bar{B}) = 0 \quad \iff \quad \sum_{j=1}^M \pi^j (\hat{B}^j - \bar{B}) \exp \left( -\frac{\hat{R}^j + \hat{\mu} \hat{B}^j}{\beta} \right) = 0 .$$

Otherwise  $\hat{\mu} = 0$  and in this case,  $\hat{\varepsilon}_{\hat{\mu}}^j$  becomes

$$\hat{\varepsilon}_{\hat{\mu}}^j = \frac{\pi^j \exp \left( -\hat{R}^j / \beta \right)}{\sum_{k=1}^M \pi^k \exp \left( -\hat{R}^k / \beta \right)} .$$