

Towards Inference Efficient Deep Ensemble Learning

Ziyue Li*, Kan Ren, Yifan Yang,
Xinyang Jiang, Yuqing Yang, Dongsheng Li

Microsoft Research
litzy0619owned@gmail.com, kan.ren@microsoft.com

Abstract

Ensemble methods can deliver surprising performance gains but also bring significantly higher computational costs, e.g., can be up to 2048X in large-scale ensemble tasks. However, we found that the majority of computations in ensemble methods are redundant. For instance, over 77% of samples in CIFAR-100 dataset can be correctly classified with only a single ResNet-18 model, which indicates that only around 23% of the samples need an ensemble of extra models. To this end, we propose an inference efficient ensemble learning method, to simultaneously optimize for effectiveness and efficiency in ensemble learning. More specifically, we regard ensemble of models as a sequential inference process and learn the optimal halting event for inference on a specific sample. At each timestep of the inference process, a common selector judges if the current ensemble has reached ensemble effectiveness and halt further inference, otherwise filters this challenging sample for the subsequent models to conduct more powerful ensemble. Both the base models and common selector are jointly optimized to dynamically adjust ensemble inference for different samples with various hardness, through the novel optimization goals including sequential ensemble boosting and computation saving. The experiments with different backbones on real-world datasets illustrate our method can bring up to 56% inference cost reduction while maintaining comparable performance to full ensemble, achieving significantly better ensemble utility than other baselines. Code and supplemental materials are available at <https://seqml.github.io/irene>.

Introduction

Recent years have witnessed the great success of deep ensemble learning methods, being applied in practical machine learning applications such as image classification (Lee et al. 2015; Zhang, Liu, and Yan 2020), machine translation (Shazeer et al. 2017; Wen, Tran, and Ba 2020) and reinforcement learning (Yang et al. 2022). The general idea of ensemble method is to conduct the prediction upon aggregating a series of prediction outcomes from several various base models. The benefits of ensemble method mainly lie in two aspects: improved generalization (Zhou, Wu, and Tang

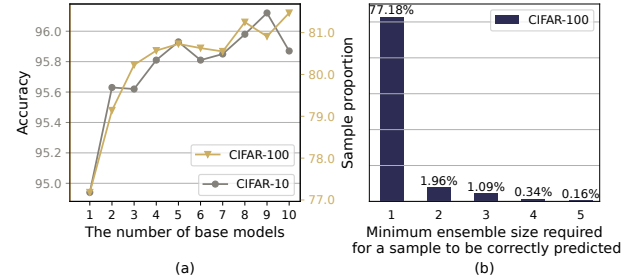


Figure 1: (a) The performance of the average ensemble method (Lakshminarayanan, Pritzel, and Blundell 2016) on CIFAR-10/100 with the different number of base models. (b) The minimum size of the average ensemble for correctly predicting samples in the CIFAR-100 dataset.

2002) and specialization on different samples (Abbasi et al. 2020; Gontijo-Lopes, Dauphin, and Cubuk 2021).

While delivering surprising performance gains, ensembles are typically much more computationally expensive compared to single model inference. The number of the incorporated models can be up to 2048 in large-scale ensemble (Shazeer et al. 2017). However, expanding the capacity of the model pool cannot bring equal benefits. In Figure 1, we illustrate the image classification performance w.r.t. the leveraged model number in average ensemble which has been commonly used as a simple yet effective ensemble method (Garipov et al. 2018; Gontijo-Lopes, Dauphin, and Cubuk 2021). As shown in Figure 1(a), though incorporating more base models increases the overall performance, the marginal benefit is rapidly decreasing when the number of base models is larger than 2. Moreover, we further investigate the minimum model number required in ensemble to conduct correct prediction on CIFAR-100 (Krizhevsky, Hinton et al. 2009) and Figure 1(b) illustrates only a few samples need more than one model to achieve correctness. This is interesting while reasonable since the capability of different models may overlap with each other especially when the model pool becomes larger. We argue that it is of cost and inefficient considering the limited performance gain and the largely increased inference consumption. Therefore, it is crucial to address the *trade-off between ensemble performance and inference efficiency for different samples*, which

*The work was conducted during the internship of Ziyue Li at Microsoft Research Asia. Correspondence to Kan Ren. Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

has not received sufficient attention in previous ensemble works.

Some works aim to reduce computational costs of *single* model inference. One stream is pruning the model architecture (Han, Mao, and Dally 2015; Frankle and Carbin 2018) or quantization the network parameter (Liang et al. 2021; Jacob et al. 2018). The other stream controls the model inference from a dynamic view (Han et al. 2021) by adjusting either the model depth (Teerapittayanon, McDanel, and Kung 2016; Huang et al. 2017a; Figurnov et al. 2017) or model width (Shen et al. 2020). However, these works are mainly for individual model inference and orthogonal to efficient ensemble learning. The base models in ensemble method are individually trained and each of the base models can play an individual role of prediction, upon which ensemble has been conducted. It is non-trivial to directly transfer the techniques of individual model optimization to ensemble structure. The solutions to efficient ensemble learning is required.

However, efficient inferences have rarely been discussed in ensemble learning community. One reason comes from common knowledge that more models derive higher performance gains (Lee et al. 2015; Shazeer et al. 2017). However, as we found in Figure 1, the marginal benefit is lower when conducting ensemble on more models. Some works leverage heuristic selection from the model pool, such as top- K gating (Shazeer et al. 2017), or stop inference based on the obtained prediction confidence in cascading ensemble (Wang et al. 2021). These studies have neither taken inference efficiency into the ensemble learning process, nor explicitly optimized for inference cost reduction. Thus, the adopted heuristic mechanism may derive sub-optimal solutions.

In this paper, we propose an InfeRence EfficieNt Ensemble (IRENE) learning approach that systematically combines ensemble learning procedure and inference cost saving as an organic whole. Specifically, we regard ensemble learning as a sequential process which allows to simultaneously optimize training each base model and learning to select the appropriate ensemble models. At each timestep of the sequential process, we maintain a common selector to judge if the current ensemble has reached ensemble effectiveness and halt further inference, otherwise filters this challenging sample for the subsequent models to conduct more powerful ensemble. As for optimization, the leveraged base models and the selector are jointly optimized to (i) encourage the base models to specialize more on the challenging samples left from the predecessor(s) and (ii) keep the ensemble more effective when leveraging newly subsequent model(s) while (iii) reducing the overall ensemble costs. Through this way, the extensive experiments demonstrate that our proposed method can reduce the average ensemble inference cost by up to 56% while maintaining comparable performance to that on the full ensemble model pool, outperforming the existing heuristic efficient computation methods.

Related Work

Ensemble Learning

Existing ensemble works mainly strive to improve ensemble performance, increasing model diversity to boost gen-

eralization ability. Some methods train models in parallel and encourages the divergence of base models via incorporating additional objectives (Zhou, Wang, and Bilmes 2018; Zhang, Liu, and Yan 2020). While, others propose sequential optimization such as boosting (Freund 1995; Chen and Guestrin 2016; Ke et al. 2017), snapshot ensemble (Huang et al. 2017b), and fast geometric ensembling (Garipov et al. 2018), to encourage model diversity by optimizing the current model on top of the existing model(s). In general, more models yield better ensemble gains (Malinin, Mlodozieniec, and Gales 2019). Thus, previous approaches tend to disregard the computational costs associated with ensemble (Lakshminarayanan, Pritzel, and Blundell 2016). However, the increase in benefits may be less effective compared to increased costs, as we demonstrate in this paper.

Inference Cost Reduction

There have been works on reducing the computational costs of a *single* model. They can be grouped into two main streams: static and dynamic solutions. Within the former stream, numerous efforts are made to design compact network structures (Sandler et al. 2018), and to introduce sparsity (Han et al. 2015), pruning (Sanh, Wolf, and Rush 2020), quantization (Esser et al. 2019) and model distillation (Wang et al. 2018), into existing networks. Static methods treat samples of different hardness equally. In contrast, the latter adaptively activates some components of networks (Bengio et al. 2015), such as dynamic depth (Hu et al. 2020; Zhou et al. 2020), dynamic width (Bengio, Léonard, and Courville 2013) and dynamic routing (Liu and Deng 2018), based on the given input. These efficient inference methods are specifically designed for single-model architectures and are orthogonal to efficient ensemble computation involving multiple independent models.

In ensemble learning, a few recent works consider the computational savings. For example, several works employ heuristic criteria (e.g., confidence threshold) to stop inference (Wang et al. 2021; Inoue 2019). Besides, Shazeer et al. (2017) proposes a top- K gating to determine sparse combinations of models. These heuristics do not explicitly address efficiency as part of the learning objective, and thus may yield sub-optimal solutions.

Methodology

In this section, we first revisit the general ensemble learning framework and devise the overall optimization goal. Then, we formulate ensemble learning as a sequential process for adaptively selecting sufficient models as an ensemble for the given sample; we define the inference efficient ensemble learning as an optimal halting problem and propose a novel framework to model this process. Finally, we detail the inference and optimization of our proposed method.

Preliminaries

Ensemble learning involves training multiple base models that perform diversely and aggregating their predictions. This can lead to significant performance improvements, but

also a corresponding increase in computational cost through the inclusion of various models.

Therefore, considering the computational cost, the general objective of ensemble is to maximize the performance and minimize the inference cost in expectation as

$$\max_{\{\theta_t\}_{t=1}^T} \mathbb{E}_{(x,y) \sim D} \left[V(y, \{\hat{y}_t\}_{t=1}^T) - \alpha C(\{\theta_t\}_{t=1}^T | x) \right], \quad (1)$$

where V measures ensemble performance, C measures the inference cost given a specific sample, and α weighs these two objectives. Here the prediction of t -th model $f_{\theta_t}(\cdot)$ models $\Pr(y|x)$ and $\hat{y}_t = f_{\theta_t}(x)$ where θ_t is the corresponding model parameter. D is the given dataset.

Realization. Here we discuss some specific formulations to realize the performance and cost metrics. To measure ensemble performance, generally task-specific loss \mathcal{L} can be adopted, such as cross-entropy loss for classification and mean squared error for regression (Zhang, Liu, and Yan 2020). While various metrics, e.g., latency or FLOPs (Wang et al. 2021), can be utilized for measuring inference cost.

However, most works only focus on optimizing ensemble performance, by striving to obtain diverse base models for improving the generalization of ensembles (Lee et al. 2015).

As for the inference cost, the existing ensemble works only perform heuristic model selection to save computational effort. For example, some works either choose a fixed number of top-ranked predictions for all the samples (Shazeer et al. 2017), or select a subset of models whose ensemble has prediction confidence above a predefined threshold (Wang et al. 2021). These heuristic solutions may yield sub-optimal solutions in terms of performance and efficiency. Instead, we turn to a joint consideration of ensemble performance and efficiency, modeling the optimization of efficiency as part of ensemble learning.

Sequential Ensemble Framework

In this work we attempt to model ensemble inference as a sequential process in which a learning-based model selection yields a sample-specific efficient ensemble over a subset of models. Specifically, during the process, model inference is sequential and selectively halts at one timestep. In this section, we first discuss two paradigms of ensemble inference and their pros and cons, then we present the model selection problem, introduce our design for sequential model selection, and describe the entire inference process.

Sequential inference v.s. parallel inference The inference paradigms of ensemble methods with computation saving can be divided into two types: parallel and sequential, as shown in Figure 2(b) and (c). In the parallel paradigm, there is a module that determines which models are activated for inference *prior to* model inference (Shazeer et al. 2017). Therefore, the framework does not utilize actual outcomes or the inference situation of base models to decide (which models would be selected). In contrast, in the sequential paradigm, outcomes can be taken at any step throughout the inference process (Wang et al. 2021) to determine how to balance the effectiveness and efficiency of the ensemble. Thus, we focus on the sequential paradigm, which allows

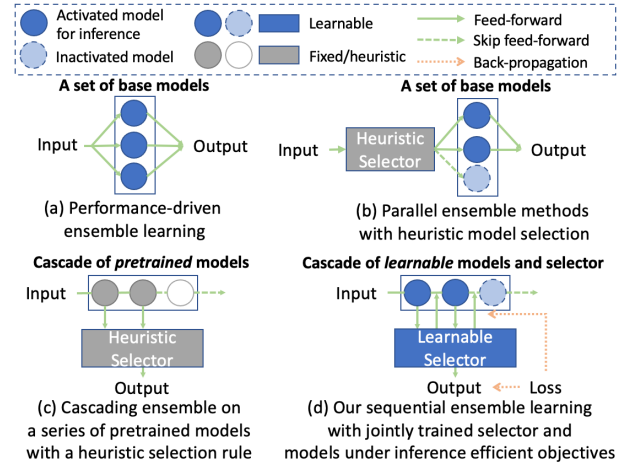


Figure 2: Comparison of (a) the general ensemble, (b) the approach using heuristic model selection under the parallel paradigm and (c) under the sequential paradigm, and (d) our approach. Our approach differs from the existing methods in two distinct ways: (i) we build a learnable selector to explicitly optimize the inference cost, and (ii) model training and selection learning occur as an organic whole.

using model-relevant information to dictate methods for balancing ensemble effectiveness and efficiency, in this paper.

As illustrated in Figure 3, in our framework, all the base models are kept ordered for training (ensemble learning) and inference (ensemble inference). In the inference process for a given sample, each base model will be activated for inference until the optimal halting event occurs, which has been decided by the jointly trained selector. And the predictions of all activated models will be aggregated as output.

Optimal halting Here we define the optimal halting event and describe our inference halting mechanism.

We first illustrate the notations and settings of our sequential inference and optimal halting mechanism. Specifically, we address the problem of selecting an appropriate subset of models for a given sample. Suppose that there are T models cascaded together; and the base model inference is performed sequentially once at each timestep. At each timestep t , the t -th model will be executed and one selector shared by all the timesteps will decide to halt at the current timestep or continue inference for further ensemble. Given one specific sample, halting at a certain timestep implies that, the predictions of the models *at this step and before* will be aggregated as the final ensemble output, and leveraging more following base models will not help improve the prediction performance while only increasing the inference cost.

Formally, given such an ordered sequence of models, we define the optimal halting timestep as z at which it is optimal to halt for some criterion, i.e., the overall optimization objective in Eq. (1) has been achieved. We define the probability distribution p_t on $\{1, \dots, T\}$ to represent the probability that halting at the t -th step is optimal.

With probability distribution p_t , we can define probability

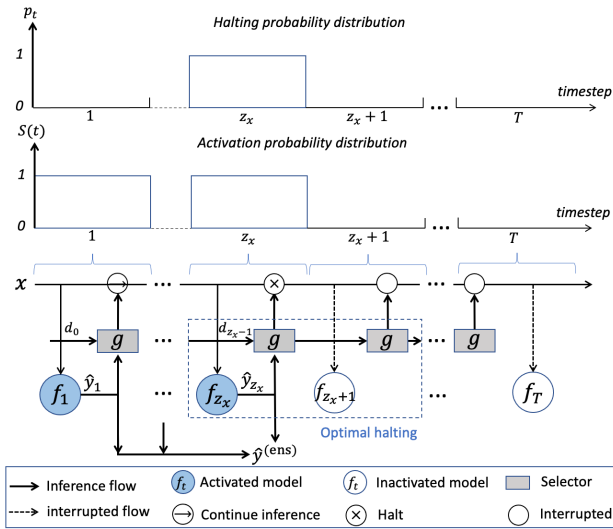


Figure 3: Illustration of our sequential inference and optimal halting mechanism.

that the optimal halting step is *at or after* the t -th step as

$$S(t) = \Pr(z \geq t) = 1 - \Pr(z < t) = 1 - \sum_{i=1}^{t-1} p_i, \quad (2)$$

which is calculated from the cumulative probability of the optimal halting step. Note that, according to our setting, all models *before and at* the optimal halting step are activated for inference. Thus, the t -th model is activated for inference when the optimal halting step is not within the first $(t - 1)$ steps, i.e., $S(t) = 1$. From the definition, we can say that $S(t)$ represents the probability of activating all the models ordered *before and at* t .

We define the conditional probability h_t of halting at the t -th step given that no halting has occurred before as

$$h_t = \Pr(z = t | z > t - 1) = \frac{\Pr(z = t)}{\Pr(z > t - 1)} = \frac{S(t) - S(t + 1)}{S(t)}. \quad (3)$$

To determine whether to halt at each timestep t , our goal is to predict the probability p_t of the optimal halting event. And the probability will be further utilized for efficiency-aware ensemble inference.

We can further derive the calculation of the probability functions $S(t)$ based on the conditional probability h_t at each timestep. The probability that, the optimal halting step is larger than or equal to the current timestep t , can be calculated following a probability chain as

$$\begin{aligned} S(t) &= \Pr(z \neq t - 1, \dots, z \neq 1) \\ &= \Pr(z \neq t - 1 | z > t - 1) \cdots \Pr(z \neq 1 | z > 0) \\ &= \prod_{i=1}^{t-1} (1 - h_i). \end{aligned} \quad (4)$$

Based on Eq. (3) and (4), we can derive the probability p_t as

$$p_t = h_t \prod_{i=1}^{t-1} (1 - h_i). \quad (5)$$

Therefore, we can predict whether to halt at a step based on the predicted conditional probability h_t prior to and at the timestep to be estimated.

Sequential halting decision Based on our derived probabilistic formulation of the optimal halting problem, we propose the selector network that leverages sample and model-related information to predict the conditional probability h_t for the calculation of halting probability p_t at each timestep.

The selector network g_ϕ should be of the form $h_t, d_t = g_\phi(e(x), d_{t-1})$, to generate state-specific information based on the sample-related information encoded by $e(\cdot)$ and the output d_{t-1} of the last timestep. Although various types of neural networks can be applied for g , we choose the recurrent neural network which is commonly used for modeling conditional probabilities over time (Schuster and Paliwal 1997). Besides, we take the output of the t -th model f_{θ_t} , which provides both information related to the sample and the t -th model that has participated in ensemble inference, as $e(x)$. The implementation details refer to Appendix A.1.

Ensemble inference Here we detail the overall inference process. Given a specific sample x , the t -th model and selector infer in turn at each timestep t . Specifically, the base model will be activated to produce the individual prediction \hat{y}_t , and the selector will compute the optimal halting probability p_t to determine if inference of the given sample should be continued ($p_t = 0$) to subsequent base models $\{f\}_{t+1}^T$, or stop the inference ($p_t = 1$). Based on previous definitions, we can derive the halting step, on the basis of which ensemble prediction and ensemble efficiency are calculated. The optimal halting step z_x for the sample x satisfies that

$$z_x = \arg \max_t p_t = \arg \max_t (h_t \prod_{i=1}^{t-1} (1 - h_i)). \quad (6)$$

To sample a unique z_x from p_t , we first sample each h_t using a differentiable sample from the Gumbel-Softmax distribution (Jang, Gu, and Poole 2016). Specifically, after h_t is sampled and binarized using a trainable binary mask in (Mallya, Davis, and Lazebnik 2018), only one timestep z_x has the maximum probability value ($p_{z_x} = 1.0$) when $h_{z_x} = 1$, and $h_t = 0$ for any $t < z_x$. That is, p_t of any timestep t after z_x , being the product of $(1 - h_{z_x})$ with value zero and any number, will be zero in Eq. (5).

As halting at the z_x -th timestep, the ensemble prediction can be calculated as an aggregation from the previous base model outputs as

$$\hat{y}^{(ens)} = \text{ENS}(\{\hat{y}_t\}_{t=1}^{z_x} | z=z_x) = \frac{1}{z_x} \sum_{t=1}^{z_x} \hat{y}_t. \quad (7)$$

Here we simply utilize average ensemble (Huang et al. 2017b; Garipov et al. 2018; Wang et al. 2021) as an example of ensemble aggregation ENS. Note that in our implementation, all base models use the same backbone, thus, the ensemble efficiency can be directly measured as the number of steps taken $z = z_x$.

Ensemble Learning

Note that, we have *no* ground truth of the optimal halting step z for each sample, otherwise we could directly model that through maximum likelihood estimation. Therefore, we propose to optimize our sequential ensemble framework from the perspective of maximizing ensemble performance as well as ensemble efficiency. In the following, we

will first introduce the optimization of base model performance, then the ensemble performance and efficiency, and then introduce inference efficient ensemble optimization by jointly optimizing base models and the selector.

Optimization of base model performance As the cornerstone of the entire framework, the performance of the base models is crucial. We train each base model by minimizing the task-specific loss $\mathcal{L}_t^{(\text{base})}$ over all training samples as

$$\min_{\theta_t} \mathbb{E}_{(x,y) \sim D} \mathcal{L}(y, \hat{y}_t). \quad (8)$$

We will describe its joint optimization with the selector later.

Optimization of ensemble performance Ensemble performance is a critical goal. In particular, we optimize this goal at each timestep to ensure that the halting strategy is optimized toward the final objective.

At each t -th step, the cascade yields its predictions for the current stage, by averaging the predictions of all the activated models for the given sample as

$$\hat{y}_t^{(\text{ens})} = \left(\sum_{i=1}^t (S(i) \hat{y}_i) \right) / \left(\sum_{i=1}^t S(i) \right), \quad (9)$$

where $S(i)$ defined in Eq. (2) indicates the probability of \hat{y}_i will be aggregated into the ensemble. Different from Eq. (7) which computes the final ensemble prediction when the halting step z_x is reached, Eq. (9) computes the (temporary) ensemble predictions obtained at any step, regardless of when the halting step occurs.

We optimize the task-specific loss of ensemble prediction $\mathcal{L}_t^{(\text{ens})}$ at each timestep as

$$\min_{\phi} \mathbb{E}_{(x,y) \sim D} \left[\mathcal{L}(y, \hat{y}_t^{(\text{ens})}) \right]. \quad (10)$$

Note that, this objective is only used to optimize the selector parameter ϕ because it may hinder model training, as proven in Allen-Zhu and Li (2020): optimizing base models by ensemble performance largely degrades performance.

Optimization of ensemble inference efficiency The selector should be aware of the inference cost imposed by its own strategy to discourage trading full usage for higher performance. In this work, we directly use the number of steps taken before halting (i.e., the ensemble size for a sample) as a measure of inference cost $\mathcal{L}_t^{(\text{cost})}$ and optimize as

$$\min_{\phi} \mathbb{E}_{\left\{ \begin{array}{c} (x,y) \sim D, \\ \{h_i \sim q_{\phi}(\cdot, i)\}_{i=1}^{t-1} \end{array} \right\}} \left[\sum_{t=1}^T \underbrace{(t \cdot (h_t \prod_{i=1}^{t-1} (1 - h_i)))}_{p_t} \right]. \quad (11)$$

Since all the models in our framework use the same backbone, as a common practice in ensemble learning, the number of activated models for ensemble is a direct measurement of the inference cost. If the base models use different network backbones, the measurement could be other metrics such as the activated model parameters or inference FLOPs, which does not influence our proposed approach.

Joint optimization for effective and efficient ensemble

During the sequential inference process in our ensemble method, the selector relying on the inference situation tries to decide whether or not to halt (stop inference for ensemble) at the current timestep. Thus, the whole system (both selector and the base models) should balance the performance gain when incorporating more base models with the increased inference cost.

We expect that adding a model to the ensemble leads to an improvement in performance, which means that the newly added model is more specialized on these samples. To this end, one way is, for the selector, to adjust its selection of models and the other way is to encourage each base model to be more focusing on the samples to which they are assigned (by the sequential inference process) accordingly.

We propose an objective $\mathcal{L}_t^{(\text{rank})}$ to optimize from both perspectives simultaneously as

$$\min_{\theta_t, \phi} \mathbb{E}_{(x,y) \sim D} \left[\max(0, S(t)(\mathcal{L}(y, \hat{y}_t) - \hat{\mathcal{L}}(y, \hat{y}_{t-1}^{(\text{ens})})) \right], \quad (12)$$

where $S(t)$ in Eq. (2) denotes the probability that inference does not halt at reaching the t -th step. Note that, with $S(t)$ equal to 1, the right-hand element of the maximum function will be the difference between the loss $\mathcal{L}(y, \hat{y}_t)$ of the t -th model and the referential loss $\hat{\mathcal{L}}(y, \hat{y}_{t-1}^{(\text{ens})})$ of the ensemble of the first $(t-1)$ models. $\hat{\mathcal{L}}$ plays a role as the target value for bootstrapping *without* backpropagation at timestep t . Otherwise when the difference is less than 0 or $S(t) = 0$, the objective output is 0, as it implies the loss of the t -th model being less than that of the other one; or inference has been halted before the t -th step.

On one hand, $\mathcal{L}_t^{(\text{rank})}$ encourages the potentially incorporated base models to obtain lower task-specific losses, i.e., better performance, on the sample that is assigned to them than *the ensemble of the previous models*. On the other hand, it incents the selector to let subsequent models get samples they are better at than *the ensemble of the previous models*.

Sequential training paradigm Aligned with the sequential inference process, the training process is also sequentially conducted. Specifically, we initialize the sequentially ordered base models and optimize them step by step. Based on our proposed optimization objectives, we minimize the expected task-specific loss and inference cost by jointly optimizing base models and the selector at each timestep t as

$$\mathcal{L}_t^{(\text{total})} = \underbrace{\mathcal{L}_t^{(\text{base})}}_{\text{model optimization}} + \underbrace{\omega_1 \mathcal{L}_t^{(\text{ens})} + \omega_2 \mathcal{L}_t^{(\text{cost})}}_{\text{selector optimization}} + \underbrace{\omega_3 \mathcal{L}_t^{(\text{rank})}}_{\text{joint optimization}}, \quad (13)$$

where ω_1 , ω_2 , and ω_3 are the loss weights. The overall training algorithm has been illustrated in Appendix B.

Together, these objectives serve to learn an inference efficient ensemble. The optimization of a single base model for $\mathcal{L}_t^{(\text{base})}$ is relatively independent from the other objectives, but is also the basis of our entire framework. For the selector optimization, the two objectives applied are adversarial, as optimizing $\mathcal{L}_t^{(\text{ens})}$ readily employs all models in exchange for advanced performance, while optimizing $\mathcal{L}_t^{(\text{cost})}$ may decrease the ensemble performance while saving inference

costs. Though these two objectives can in essence optimize both ensemble performance and ensemble efficiency, it is difficult to expect the selector to work reasonably well, e.g., to perform further inference only when subsequent models can bring improvements. Therefore, we further propose $\mathcal{L}_t^{(\text{rank})}$ to jointly optimize an effective and efficient ensemble, ensuring that adding more models is beneficial and models are more focused on samples assigned to them.

Discussion: Paradigm differences from prior methods Our ensemble paradigm varies to existing paradigms in the optimal halting strategy learned from the interaction of the selector and base models via their joint optimization, as shown in Figure 2. From a paradigm perspective, we introduce a learning-based selector rather than heuristics. In addition, our selector interacts with base models to facilitate their attention to samples assigned to them, instead of receiving model outcomes to independently determine the halting step. From the optimization perspective, we propose novel objectives for inference efficient ensemble learning as discussed above. Note that, by optimizing Eq. (12), we realize the interaction between the base models and the selector, rationalizing the strategies made by the selector and aligning the performance of the base models with the selection.

Experiment

Experimental Setup

Here we present the details of experimental setup, including datasets, backbones used, and baselines for comparison.

Datasets and backbones. We conduct experiments on two image classification datasets, CIFAR-10 and CIFAR-100, the primary focus of neural ensemble methods (Zhang, Liu, and Yan 2020; Rame and Cord 2021). CIFAR (Krizhevsky, Hinton et al. 2009) contains 50,000 training samples and 10,000 test samples, which are labeled as 10 and 100 classes in CIFAR-10 and CIFAR-100, respectively. Following the previous setup of ensemble methods, we adopt ResNet-32 and ResNet-18 (He et al. 2016) as backbones and all the ensemble methods to be compared use three base models, i.e., $T = 3$. We also provide ablation studies on T .

Baselines. We compare IRENE with various ensemble methods, the implementation details of which are given in Appendix A.2. Traditional ensemble methods that do not consider costs, includes **MoE** (Shazeer et al. 2017), **average ensemble**, **Snapshot ensemble** (Huang et al. 2017b) and **fast geometric ensembling** (Garipov et al. 2018) (FGE). Other than them, we also compare with **Sparse-MoE** (Shazeer et al. 2017) and **WoC** (Wang et al. 2021), both of which adopt heuristic computation saving methods. For Sparse-MoE, using the same setup as theirs, two models are activated for each sample in the inference. For WoC, which uses a confidence threshold-based halting on a cascade of pretrained models, we follow them and implement it on the trained base models of average ensemble.

Evaluation metrics. We report, for each method, its utility value (the trade-off between ensemble performance and efficiency), top-1 accuracy, and the corresponding inference cost, i.e., average number of utilized models in ensemble. Their detailed description is in Appendix C.

	Methods	Top-1 (%) \uparrow	Cost \downarrow	Utility \uparrow
ResNet-32	Single model	93.10 ± 0.12	1.00	1.00
	Average ensemble	94.46 ± 0.13	3.00	1.00
	Snapshot ensemble	93.73 ± 0.36	3.00	0.56
	FGE	93.19 ± 0.11	3.00	0.36
	MoE	93.26 ± 0.35	3.00	0.38
	Sparse-MoE	81.31 ± 2.50	2.00	0.00
	WoC	94.29 ± 0.00	2.48 ± 0.09	1.05
	IRENE	94.46 ± 0.07	2.21 ± 0.12	1.37
ResNet-18	Single model	94.94 ± 0.07	1.00	1.00
	Average ensemble	95.62 ± 0.04	3.00	1.00
	Snapshot ensemble	95.61 ± 0.14	3.00	0.98
	FGE	94.67 ± 0.06	3.00	0.22
	MoE	94.41 ± 0.06	3.00	0.14
	Sparse-MoE	84.07 ± 1.42	2.00	0.00
	WoC	95.48 ± 0.01	1.15 ± 0.02	2.10
	IRENE	95.81 ± 0.08	1.32 ± 0.13	3.10

Table 1: Experiment results on CIFAR-10.

	Methods	Top-1 (%) \uparrow	Cost \downarrow	Utility \uparrow
ResNet-32	Single model	69.58 ± 0.52	1.00	1.00
	Average ensemble	74.94 ± 0.30	3.00	1.00
	Snapshot ensemble	74.26 ± 0.18	3.00	0.97
	FGE	71.19 ± 0.27	3.00	0.86
	MoE	70.64 ± 1.00	3.00	0.84
	Sparse-MoE	49.48 ± 1.09	2.00	0.00
	WoC	73.90 ± 0.24	2.31 ± 0.21	1.05
	IRENE	74.84 ± 0.06	2.53 ± 0.08	1.16
ResNet-18	Single model	77.18 ± 0.16	1.00	1.00
	Average ensemble	80.28 ± 0.25	3.00	1.00
	Snapshot ensemble	79.17 ± 0.19	3.00	0.68
	FGE	77.84 ± 0.37	3.00	0.42
	Div ²	79.12	3.00	0.67
	MoE	77.49 ± 0.37	3.00	0.37
	Sparse-MoE	59.04 ± 0.82	2.00	0.00
	WoC	79.86 ± 0.01	1.93 ± 0.06	1.35
	IRENE	80.10 ± 0.14	1.88 ± 0.05	1.50

Table 2: Experiment results on CIFAR-100.

Evaluation Results

We demonstrate the effectiveness of IRENE on two benchmark datasets, CIFAR-10 and CIFAR-100, using two different backbones, with results shown in Table 1 and 2. We mark the best results in bold with arrows (\uparrow/\downarrow) indicating the direction of better outcomes for the metrics.

Performance Comparison

IRENE achieves better trade-offs than traditional ensemble methods and heuristic cost-constrained methods. While comparing with traditional methods with fixed costs, we are interested in the inference cost saved by IRENE. As shown in tables, the inference costs of IRENE are 73.67%, 44.00%, 84.33%, and 62.67% of those in traditional ensemble methods, resp. That is, a large fraction of the inference costs can be saved while the performance penalty turns out to be small or even negligible. Accordingly, the utility of IRENE is higher than average ensemble by an average of 78.25%. Compared with Sparse-MoE and WoC, two methods that use heuristic computation saving, IRENE also scores higher in

utility, indicating that it is superior to heuristic solutions.

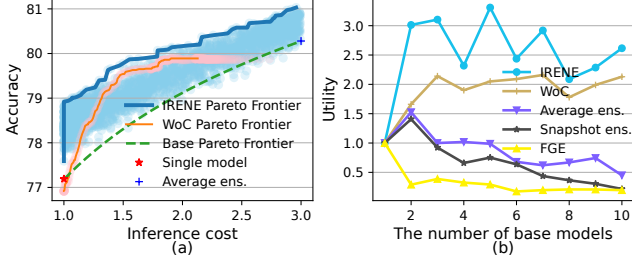


Figure 4: (a) Pareto frontier on CIFAR-100 using ResNet-18. The curve of single model (average ensemble) is denoted as base Pareto frontier. (b) Utility with varied number of ResNet-18 models on CIFAR-10.

Improvement of IRENE is related to dataset difficulty and neural network capability. As shown in Table 1, the relatively advanced backbone ResNet-18 shows a surprising result over the easier CIFAR-10 dataset. Specifically, the inference cost of IRENE is reduced by 56.0% versus average ensemble, while its performance is improved by 0.19%. This supports that the existence of redundancy between models and IRENE seizes it for efficient and effective inference. Additionally, both backbones use less inference cost in CIFAR-10 than in CIFAR-100, which is reasonable since the task is relatively more difficult in CIFAR-100 and it also illustrates that our proposed method of inference efficient ensemble learning can dynamically adjust the ensemble efficiency according to the task difficulty.

Sensitivity analysis A: The Pareto frontier We compare IRENE, WoC, and average ensemble for their performance-cost trade-off through the Pareto frontier, and their results using ResNet-18 on CIFAR-100 are shown in Figure 4(a). From the figure, *for a small increase in cost (around 1.0)*, IRENE achieves a significant improvement in accuracy. In addition, IRENE competitively obtains Pareto optimal values under all cost regimes. Moreover, IRENE yields significantly better performance than traditional ensemble methods under various cost constraints, indicating IRENE benefits performance-cost trade-off as well as model training. In contrast, WoC can only approach the optimum in a limited cost region and leaves a clear cutoff point in the curve, suggesting that its heuristic solution is sub-optimal since adding more inference cost does not bring promising performance gains which also illustrates that our proposed method is superior for effective and efficient ensemble learning.

Sensitivity analysis B: Utility of IRENE with different number of base models IRENE aiming to learn inference efficient ensemble is expected to achieve better trade-off between performance and cost thus preventing introducing worthless inference cost. To verify this, we compare the utility of various ensemble methods when the number of base models varies, with CIFAR-10 results using ResNet-18 shown in Fig. 4(b). The other compared methods include average ensemble, snapshot ensemble, FGE, and WoC. An insight drawn from the results is that, as the number of models

	Top-1 (%) \uparrow	Cost \downarrow	Utility \uparrow
w/o $\mathcal{L}_t^{(rank)}, \mathcal{L}_t^{(ens)}, \mathcal{L}_t^{(cost)}$ on selector	95.78 ± 0.04	1.82 ± 0.05	2.15
w/o $\mathcal{L}_t^{(rank)}$ on base models	95.74 ± 0.08	1.47 ± 0.18	2.52
w/o $\mathcal{L}_t^{(rank)}$	96.14 ± 0.05	2.69 ± 0.19	2.57
w/o $\mathcal{L}_t^{(ens)}$	95.75 ± 0.10	1.41 ± 0.23	2.67
w/o $\mathcal{L}_t^{(cost)}$	95.98 ± 0.24	2.00 ± 0.64	2.88
IRENE	95.81 ± 0.08	1.32 ± 0.13	3.10

①: $\mathcal{L}_t^{(cost)}$ helps to control the **ensemble efficiency**, as cost increases w/o it.
 ②: $\mathcal{L}_t^{(ens)}$ benefits **performance and efficiency**, both of which degrade w/o it.
 ③: $\mathcal{L}_t^{(rank)}$ proves its validity in **effectiveness and efficiency trade-offs**, given a sharp cost increase and a significant utility drop from ablation results.
 ④: Not optimizing base models by $\mathcal{L}_t^{(rank)}$ leads to a worse utility, indicating that it **facilitates model training** and thus improves the final performance.
 ⑤: Not optimizing the selector by $\mathcal{L}_t^{(rank)}, \mathcal{L}_t^{(ens)},$ and $\mathcal{L}_t^{(cost)}$ leads the most distinct utility drop, showing their critical role in **efficient ensemble learning**.
 ④ v.s. ⑤: Optimizing selector benefits utility more than optimizing models.

Figure 5: Ablation study of the optimization objectives with ResNet-18 on CIFAR-10.

increases, the utility of traditional ensemble methods generally first increases and then continues to decline. This further suggests that existing ensemble methods may be trading unnecessary computation for performance gains. IRENE and WoC, in contrast, can maintain decent utility despite adding more base models and IRENE outperforms WoC by a large margin, proving the robustness and effectiveness of IRENE.

Sensitivity analysis C: Ablation studies We perform ablation studies to analyze the effects of our propose objectives in Eq. (13), except for the indispensable base model optimization objective $\mathcal{L}_t^{(base)}$ in ensemble learning, on performance. Furthermore, we want to verify the effectiveness of IRENE for (i) optimal halting over base models trained independently of the selector and (ii) sequential model training only, resp. Thus, we perform ablation experiments for the optimization of the selector and for the optimization of the base model. To be specific, the selector is optimized by $\mathcal{L}_t^{(rank)}, \mathcal{L}_t^{(ens)},$ and $\mathcal{L}_t^{(cost)}$ together, while base models is optimized by $\mathcal{L}_t^{(rank)}$. Figure 5 presents the ablation results with detailed descriptions. As shown in the figure, all the learning objectives play an indispensable role and the optimization of selector benefits the final performance more than that of base models. Note that in all ablation settings, IRENE still achieves higher utility than all the baselines in Table 1.

Conclusion

In this paper, we focus on balancing the trade-off between ensemble effectiveness and efficiency, which is largely overlooked by existing ensemble approaches. By modeling efficient ensemble inference as an optimal halting problem, we propose an *effectiveness and efficiency-aware* selector network that is optimized jointly with base models via novel optimization objectives, to determine the halting strategies. We demonstrate that IRENE can significantly reduce the inference cost while maintaining comparable performance to full ensembles, and beats existing computation-saving methods. Optimal halting modeling also offers the possibility of sequential model selection with skipping to further boost computational savings, which we leave for further work.

References

- Abbasi, M.; Rajabi, A.; Gagné, C.; and Bobba, R. B. 2020. Toward adversarial robustness by diversity in an ensemble of specialized deep neural networks. In *Canadian Conference on Artificial Intelligence*, 1–14. Springer.
- Allen-Zhu, Z.; and Li, Y. 2020. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816*.
- Bengio, E.; Bacon, P.-L.; Pineau, J.; and Precup, D. 2015. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Esser, S. K.; McKinstry, J. L.; Bablani, D.; Appuswamy, R.; and Modha, D. S. 2019. Learned step size quantization. *arXiv preprint arXiv:1902.08153*.
- Figurnov, M.; Collins, M. D.; Zhu, Y.; Zhang, L.; Huang, J.; Vetrov, D.; and Salakhutdinov, R. 2017. Spatially adaptive computation time for residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1039–1048.
- Frankle, J.; and Carbin, M. 2018. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*.
- Freund, Y. 1995. Boosting a weak learning algorithm by majority. *Information and computation*, 121(2): 256–285.
- Garipov, T.; Izmailov, P.; Podoprikin, D.; Vetrov, D.; and Wilson, A. G. 2018. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 8803–8812.
- Gontijo-Lopes, R.; Dauphin, Y.; and Cubuk, E. D. 2021. No One Representation to Rule Them All: Overlapping Features of Training Methods. *arXiv:2110.12899*.
- Han, S.; Mao, H.; and Dally, W. J. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*.
- Han, S.; Pool, J.; Tran, J.; and Dally, W. 2015. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28.
- Han, Y.; Huang, G.; Song, S.; Yang, L.; Wang, H.; and Wang, Y. 2021. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, T.-K.; Chen, T.; Wang, H.; and Wang, Z. 2020. Triple wins: Boosting accuracy, robustness and efficiency together by enabling input-adaptive inference. *arXiv preprint arXiv:2002.10025*.
- Huang, G.; Chen, D.; Li, T.; Wu, F.; Van Der Maaten, L.; and Weinberger, K. Q. 2017a. Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844*.
- Huang, G.; Li, Y.; Pleiss, G.; Liu, Z.; Hopcroft, J. E.; and Weinberger, K. Q. 2017b. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*.
- Inoue, H. 2019. Adaptive ensemble prediction for deep neural networks based on confidence level. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 1284–1293. PMLR.
- Jacob, B.; Kligys, S.; Chen, B.; Zhu, M.; Tang, M.; Howard, A.; Adam, H.; and Kalenichenko, D. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2704–2713.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; and Liu, T.-Y. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30: 3146–3154.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. Master thesis, University of Toronto.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2016. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*.
- Lee, S.; Purushwalkam, S.; Cogswell, M.; Crandall, D.; and Batra, D. 2015. Why M heads are better than one: Training a diverse ensemble of deep networks. *arXiv preprint arXiv:1511.06314*.
- Liang, T.; Glossner, J.; Wang, L.; Shi, S.; and Zhang, X. 2021. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing*, 461: 370–403.
- Liu, L.; and Deng, J. 2018. Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Malinin, A.; Mlodozieniec, B.; and Gales, M. 2019. Ensemble distribution distillation. *arXiv preprint arXiv:1905.00076*.
- Mallya, A.; Davis, D.; and Lazebnik, S. 2018. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 67–82.
- Rame, A.; and Cord, M. 2021. Dice: Diversity in deep ensembles via conditional redundancy adversarial estimation. *arXiv preprint arXiv:2101.05544*.

Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4510–4520.

Sanh, V.; Wolf, T.; and Rush, A. 2020. Movement pruning: Adaptive sparsity by fine-tuning. *Advances in Neural Information Processing Systems*, 33: 20378–20389.

Schuster, M.; and Paliwal, K. K. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11): 2673–2681.

Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.

Shen, J.; Wang, Y.; Xu, P.; Fu, Y.; Wang, Z.; and Lin, Y. 2020. Fractional skipping: Towards finer-grained dynamic cnn inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5700–5708.

Teerapittayanon, S.; McDanel, B.; and Kung, H.-T. 2016. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, 2464–2469. IEEE.

Wang, X.; Kondratyuk, D.; Christiansen, E.; Kitani, K. M.; Movshovitz-Attias, Y.; and Eban, E. 2021. Wisdom of Committees: An Overlooked Approach To Faster and More Accurate Models. In *International Conference on Learning Representations*.

Wang, Y.; Xu, C.; Xu, C.; and Tao, D. 2018. Adversarial learning of portable student networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Wen, Y.; Tran, D.; and Ba, J. 2020. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. *arXiv preprint arXiv:2002.06715*.

Yang, Z.; Ren, K.; Luo, X.; Liu, M.; Liu, W.; Bian, J.; Zhang, W.; and Li, D. 2022. Towards Applicable Reinforcement Learning: Improving the Generalization and Sample Efficiency with Policy Ensemble. *arXiv preprint arXiv:2205.09284*.

Zhang, S.; Liu, M.; and Yan, J. 2020. The diversified ensemble neural network. *Advances in Neural Information Processing Systems*, 33.

Zhou, T.; Wang, S.; and Bilmes, J. A. 2018. Diverse ensemble evolution: Curriculum data-model marriage. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 5909–5920.

Zhou, W.; Xu, C.; Ge, T.; McAuley, J.; Xu, K.; and Wei, F. 2020. Bert loses patience: Fast and robust inference with early exit. *Advances in Neural Information Processing Systems*, 33: 18330–18341.

Zhou, Z.-H.; Wu, J.; and Tang, W. 2002. Ensembling neural networks: many could be better than all. *Artificial intelligence*, 137(1-2): 239–263.