






Tiny Models are the Computational Saver for Large Models

Qingyuan Wang¹, Barry Cardiff¹, Antoine Frappé²,
Benoit Larras², and Deepu John¹

¹ University College Dublin, Ireland

`qingyuan.wang@ucdconnect.ie; {barry.cardiff,deepu.john}@ucd.ie`

² Univ. Lille, CNRS, Centrale Lille, Junia, Univ. Polytechnique Hauts-de-France,
UMR 8520-IEMN, France

`{antoine.frappe,benoit.larras}@junia.com`

Abstract. This paper introduces TinySaver, an early-exit-like dynamic model compression approach which employs tiny models to substitute large models adaptively. Distinct from traditional compression techniques, dynamic methods like TinySaver can leverage the difficulty differences to allow certain inputs to complete their inference processes early, thereby conserving computational resources. Most existing early exit designs are implemented by attaching additional network branches to the model’s backbone. Our study, however, reveals that completely independent tiny models can replace a substantial portion of the larger models’ job with minimal impact on performance. Employing them as the first exit can remarkably enhance computational efficiency. By searching and employing the most appropriate tiny model as the computational saver for a given large model, the proposed approaches work as a novel and generic method to model compression. This finding will help the research community in exploring new compression methods to address the escalating computational demands posed by rapidly evolving AI models. Our evaluation of this approach in ImageNet-1k classification demonstrates its potential to reduce the number of compute operations by up to 90%, with only negligible losses in performance, across various modern vision models.

Keywords: Dynamic Model Architecture · Model Compression · Early Exit

1 Introduction

As we marvel at the exceptional performance of large-scale AI models, their rapidly increasing sizes pose significant challenges in terms of computational demands. In 2023, cutting-edge vision models have attained up to 90% accuracy on the ImageNet-1k dataset, utilizing 362 million parameters [14]. This parameter count is sixfold that of the ResNet-152 model [25], which was deemed large,

² Code available at <https://github.com/QingyuanWang/tinysaver>

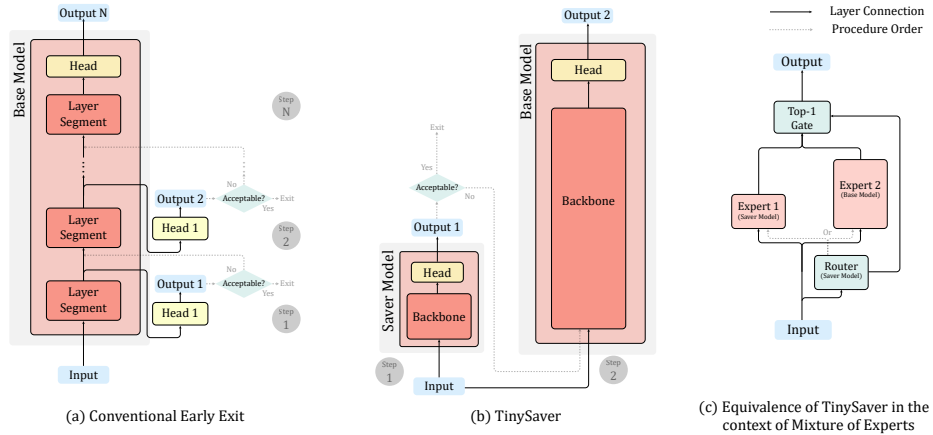


Fig. 1: Illustration of TinySaver with the comparison to related concept of Early Exit and Mixture of Experts.

eight years prior. The AI models in the Natural Language Processing (NLP) domain are expanding even more swiftly, now encompassing tens of billions of parameters [46]. These expansive models necessitate substantial computational resources. Furthermore, the swift escalation in AI services utilizing large models has a notable impact on society’s overall energy consumption [73]. This situation underscores the urgency to develop more efficient methods for computational reduction.

Early Exit (EE) [30, 33, 47, 58] based methods are emerging as a promising approach for computational reduction, offering savings in a dimension distinct from traditional compression methods. It belongs to the category of dynamic model inference techniques [20] that enables input samples to traverse different data paths. This allows for case-by-case treatment of each input, recognizing the inherent variability in sample complexity (e.g., blurred or distorted images present more challenges than clear images). Some samples, being less complex, can be effectively processed by smaller models, while more complex samples necessitate larger models. In contrast to conventional methods that uniformly apply compression across all samples, EE provides the flexibility to tailor compression specifically to each input. As depicted in Fig. 1(a), an EE-enabled system employs multiple exits, each attempting to produce outputs before reaching the final model head. The inference process ceases as soon as the output satisfies predefined exit criteria. Moreover, these early exits can be easily reconfigured, allowing for adjustable compression rates to meet varying needs post-deployment or even in real time. Beyond merely scaling performance, dynamic models offer the potential for integrating multi-task or domain-specific models, sharing a common backbone [43]. The primary drawback of EE models is the requirement for additional structural components, therefore it does not reduce model size.

However, this is less concerning as a large number of model computations will consume more energy, contribute to pollution and strain electricity grids.

The unique benefits of EE models underscore their considerable potential for various applications. However, its development remains in a nascent stage. Predominantly, existing EE approaches involve integrating auxiliary branches into middle layers of a base model’s backbone, with the expectation that these small exits can yield acceptable predictions from intermediate backbone features [32, 53]. Nonetheless, these exit branches from the initial layers often struggle to perform satisfactorily. This is partly because they rely on a backbone that is typically neither designed nor trained to support such exits. Some studies have attempted to incorporate exit strategies into the backbone design, jointly training exits and the base model [7, 36, 60]. However, these approaches are not enough to produce high-performance exits at the level of modern models as joint design will ask shallow backbone layers to provide both high and low-level features at the same time. Additionally, in the vision domain, early-stage features are often limited in their receptive field, lacking the global view and high-level semantic understanding necessary for optimal performance. Consequently, these limitations hinder the advancement and effectiveness of EE based methods.

Moreover, many existing implementation of EEs are application-specific designs [26, 39]. The design and training of these exits are coupled with their main models, posing challenges in applying EE concepts to other models and applications. To facilitate wider employment of EEs, we rethink existing methodologies and recognize the potential of decoupling this concept from specific model architectures, thereby transforming it into a generic dynamic model compression method.

The core concept of EE is to save computation by allowing sub-networks to handle parts of the task with comparable performance. Similarly, efficient-focused independent models [2, 27, 35, 40, 57] are designed to process simpler inputs with significantly lower complexity. By implementing an exiting policy that directs only simpler samples to tiny models, we can establish an EE-like system to minimize computation for straightforward inputs. Additionally, this exiting policy is configurable to balance different efficiency-accuracy trade-offs. As a result, combining tiny models with EE principles forms a dynamic, configurable compression method. It can be applied to larger models within the same task, efficiently process simple inputs while reserving complex tasks for the larger model without making too many mistakes. In essence, tiny models can act as computational savers for larger models by successfully undertaking a portion of the workload with equivalent performance.

To achieve this goal, we introduce TinySaver, a system that leverages pre-trained efficient models as a component of the EE framework, aimed at reducing the overall computation in inference. As depicted in Fig. 1(b), the computational process initiates with a tiny saver model, which acts as the primary early exit point. If the result from the saver model meets the acceptance criteria (step 1), the process halts; otherwise, the larger base model will be invoked (step 2). This approach contrasts with prior early-exit methodologies (illustrated in Fig. 1(a)),

where a complete efficient model guarantees that overall performance is not inferior. Despite their limited size, these tiny models significantly outperform traditional exit points embedded within the backbone, thereby enhancing overall system performance. Importantly, TinySaver is adaptable to model architectures that are difficult to be segmented. Small exits cannot be attached to shallow parts if the model necessitate the feature from deep levels [4, 51]. However, in our case, the employed complete model can still work.

Moreover, in the context of TinySaver, all models are pre-trained using their original methods, eliminating the need for additional training. It also effectively decouples early exits from specific applications, facilitating the straightforward incorporation of the concept of dynamic model compression into evolving AI models. This low development cost allows easy application to new models, similar to common compression methods [8]. Additionally, TinySaver’s performance can be upgraded with newly proposed efficient models. This approach not only addresses current challenges but also fosters new perspectives for the development of efficient and scalable AI systems.

The primary objective of introducing TinySaver is to overcome performance limitations encountered when implementing the EE concept. However, this method can also be explained as a specialized Mixture of Experts (MoE) [17, 18, 31, 45, 55] for computation reduction of a given model. In this context, we employ non-uniform pre-trained models as experts and the smaller expert also works as the router as Fig. 1(c). Compared to trained routers which are commonly used in MoE, our routing methods have obvious superiority of performance and system building cost in pure classification tasks and are competitive in other tasks.

We validate our idea by designing dynamic compression systems for vision models. This design demonstrates substantial computational savings for various modern models, achieved with negligible or no loss in performance.

The contribution of this work can be summarized as follows:

1. We propose a simple method to integrate pre-trained tiny models with the EE concept, creating a universal method for computational compression with dynamic model architectures.
2. We analyse the eligibility and efficiency of using tiny models for computational saving purpose in the context of EE and MoE.
3. We introduce a methodology for selecting the most effective saver model to compress a specified base model. Additionally, we discuss the possible extension of the proposed idea.
4. Our research demonstrates a significant reduction in computational requirements for contemporary vision models, thereby validating the practicality and efficiency of dynamic model compression.

2 Related Work

2.1 Efficient Vision models

In recent years, the performance of vision models has advanced rapidly, with numerous studies continuously setting new benchmarks [11, 14, 15, 25, 41, 59, 66].

Alongside these developments, another strand of research focuses on efficiency, striving to achieve optimal performance within a framework of acceptable complexity [2,27,35,40,57]. However, performance and efficiency often are conflicting and difficult to be covered in the same design. Furthermore, leading models typically rely on well-designed training recipe [9,23,66,74,75], access to large private datasets [56], and comprehensive hyper-parameter tuning [16, 57]. While most architectures are scalable [13,23,66], downstream applications usually have limited version choices, restricted to publicly available pre-trained models, which may not be ideally scaled for specific needs.

2.2 Dynamic model inference

Dynamic model inference [1,20,32,62,67,69] introduces the concept of a dynamic computing graph, enabling each input sample to be directed to the most suitable path. This approach uniquely facilitates sample differentiation, assigning each to appropriate layers for computational savings, enhanced performance, or multi-tasking.

Early Exits (EE) Early Exits [30,33,47,58,71], as a leading method in dynamic inference for computational saving, propose the termination of predictions as soon as the necessary computations are completed. While promising as a general model compression technique, research in this domain is still developing. Many implementations of EE are employed in dedicated designs to address efficiency [6, 19, 21, 22, 26, 48, 63, 68]. Some other work explore the feasibility of integrating EE with more cases [30,33,34,64,76]. However, challenges remain in generalizing early exits, such as determining the optimal exit points and constructing efficient exits. Meanwhile, although not enabled by every work, the routing policy in early exits, is usually adjustable post-training and even during runtime, allowing users to finely tune the system’s behavior to meet specific requirements. DyCE [61] was introduced to search the best routing configurations, i.e. which exit to use, for different trade-off preferences in real time. The research presented in this paper, however, tackles the latter challenge, focusing on exploring what are most efficient exits.

Sparsely-Gated Mixture-of-Experts (MoE) Sparsely-Gated MoE [17, 18, 31, 45, 55] is another recent popular dynamic model architecture usually introduced for scaling models to larger sizes. It employs multiple sub-models but only some of them are invoked during the inference according to router’s evaluation for a given sample. Therefore the model capacity can be increased without proportionally increasing the computation. Most successful models with MoE are language models, but there are some related work in vision models such as V-MoE [12, 50] and Swin-MoE [28]. Our work is proposed for saving computation for a pre-trained model which differs from the common MoE. However, it can be considered a special case in the context of MoE where the saver model simultaneously serves as the router and an expert.

3 Tiny Model as the Computational Saver

TinySaver is a simple idea. It only requires two models trained for the same task stacked together. We compare this idea with EE and MoE to demonstrate that using a tiny model as is can be 1) a computational saving replacement of large models for many inputs without a large performance compromise. 2) a competitive router to identify reliable predictions, especially for general classification tasks. We also propose a method to find the most appropriate saver for a given base model and discuss the possible extension of this method.

3.1 Analysis in the context of EEs for performance efficiency

In EE-based methods, whether to exit after a specific layer or not can be controlled by many factors. Confidence-based exiting is one of the easiest to implement. Here, EEs are designed to fulfil two primary functions: (1) Generate a prediction $\tilde{\mathbf{y}}_S$ that aligns with the format of the base model’s output. (2) Providing information for the evaluation of exiting policy. In classification tasks, this information can be a confidence value c to compare against a pre-configured threshold value t . If $c \geq t$, the prediction is deemed acceptable, and the computation process concludes. Otherwise, the computation continues with the remaining layers and finally, the prediction from the base model is accepted $\tilde{\mathbf{y}}_B$. Alternative exiting policies include but not be limited to stopping when output entropy is low enough [58], consecutive exits yield consistent predictions [77], or a trained controller decides to exit [10]. Trained for the same task, a tiny model can perform identically to the base model. Thus, complete tiny models present no significant obstacles to satisfy both requirements and can be integrated into an early exiting system.

Eligibility The maximal computational savings are acquired when EEs can perform tasks equivalently to larger models. As depicted in Fig. 1(a), conventionally attached exits utilise the backbone to minimize overhead, yet they often struggle to achieve high performance. This is primarily due to the fact that early-stage features in these branches are neither designed nor trained for making predictions at such junctures. In contrast, efficiency-focused tiny models are complete, having all necessary stages for output generation do not encounter this issue. Fig. 2 illustrates the comparison between attached exits at different positions in varied scales. However, none of attached exits can outperform independent models with the same level of computation.

Moreover, even some specific co-designs somehow surpass independently developed models in terms of efficiency, the coupled design make them less adaptable to the rapid evolution of modern models. Utilizing widely available efficient model designs can significantly reduce the resources required to build an EE-based system. A straightforward cascading of two models, combined with threshold tuning for early exits, can effectively establish a system with substantial computational savings. These attributes render TinySaver a feasible compression approach for a broad range of applications.

Efficiency We can quantitatively evaluate the efficiency of TinySaver in an EE-based system with two exits. The dataset can be divided into two categories: (I) samples meeting the exit criteria and processed by the saver S , and (II) those processed by the base model B . Cat. I, filtered through exit criteria, is confidently predicted by S , which generally indicates reliable performance. High confidence in predictions typically correlates with reliability, as evidenced in Fig. 3. Here, we examine varying confidence threshold values to chart the performance of both S and base models for Cat. I samples. Initially, when t is large the Cat. I samples are limited to very easy ones, both models exhibit very high performance on them. As the proportion of early exited samples increases, S maintains a relatively high accuracy, while B 's performance declines more rapidly. Although large models generally outperform smaller models on average, this does not imply superiority in every individual sample. Given that these samples are selected based on S 's confidence, it's not surprising that S may perform better on Cat. I samples. S 's performance eventually diminishes as it handles more samples, but there remains a range where results for Cat. I samples are enhanced. Since Cat. II samples continue to be processed by B , the system's performance on this segment remains equivalent to that of the large model. Thus, improvements in Cat. I directly contribute to the overall enhancement of the system.

The cost factor must also be considered. Cat. II samples incur additional overhead, as they require initial processing by S . In contrast, Cat. I samples benefit significantly in cost reduction, bypassing B entirely. For effective compression, the overall complexity should remain below the original complexity of B , as indicated before the crossing point in Fig. 3. Achieving this requires a highly efficient saver model. Fortunately, there are many work focusing on creating efficient models. They offer a variety of options for selecting the most suitable saver model in relation to a specific base model.




Fig. 2: Exits attached on ConvNeXtv2_{tiny} vs independent models in the similar range of complexity. Exits are attached at various backbone positions, and have different sizes. However, efficient-focused models can easily outperform them.

3.2 Analysis in the context of MoE for routing efficiency

As illustrated in Fig. 1(c), this work can be viewed as a variant of the top-1 sparse gated Mixture of Experts (MoE) architecture, featuring two heterogeneous experts, S and B with a top-1 router. The router assess the raw input x to determine which expert to engage. Actually, if the output of S yields a confidence related to per sample performance, it could effectively function as the router in a MoE framework. Our subsequent analysis shows that this model is a competitive alternative to a trained router in complexity reduction scenario.

Eligibility To reduce computation, the router should activate the smaller expert, S , when sufficient. Therefore the router must predict the reliability of the result of S for a given input \mathbf{x} . In classification, that is $P(\text{argmax}_i(\hat{\mathbf{y}}_S^{(m,i)}) = \mathbf{y}^{(m)} \mid \mathbf{x}^{(m)})$. However, this formula has the same meaning with $\text{max}_i(\hat{\mathbf{y}}_S^{(m,i)})$, which is also the probability of the prediction matches the ground truth. That implies a classification model S inherently performs the routing function, as its confidence level can serve as a prediction of accuracy. This concept extends to tasks where confidence can be derived from outputs, rendering routers redundant as the two models can inherently form an MoE-like system for computational efficiency. However, an pre or post evaluator is needed for outputs lacking inherent confidence, such as in regression tasks.

Efficiency In TinySaver, computational complexity is reduced to S only, for successfully early exited samples. However, samples failing to meet early exit criteria incur the combined complexities of both S and B , as they progress through S before B . Thus, the expectation of TinySaver’s computational complexity, $E[C_{Ts}]$ can be quantified by Eq. (1), where C_S and C_B represent computational costs, e.g. FLOPs or MACs. r_S is the ratio of samples exited via S

$$E[C_{Ts}] = r_S C_S + (1 - r_S)(C_B + C_S) \quad (1)$$

In a conventional MoE setup with a prior router, there is potential to bypass S for difficult inputs, theoretically saving computations. However, this router must handle raw inputs, necessitating a basic encoder. It makes the cost of the router, C_{Rt} , include the costs of an encoder besides the linear and softmax layer. Then the expectation of systematic complexity can be formulated as:

$$E[C_{MoE}] = r_S (C_{Rt} + C_S) + (1 - r_S)(C_{Rt} + C_B) \quad (2)$$

Compare Eq. (2) with the proposed method in Eq. (1):

$$E[C_{MoE}] < E[C_{Ts}] \quad (3)$$

$$r_{S2}(C_{Rt} + C_S) + (1 - r_{S2})(C_{Rt} + C_B) < r_{S1}C_S + (1 - r_{S1})(C_B + C_S) \quad (4)$$

$$C_{Rt} < (r_{S2} - r_{S1})C_B + (1 - r_{S2})C_S \quad (5)$$

Where r_{S1} and r_{S2} are the exiting ratio of S for each system when they are achieving the equivalent performance. Eq. (5) suggests that, the trained router must exhibit significantly lower complexity for the identical routing performance, i.e. $r_{S1} = r_{S2}$, to obtain advantages. In pure classification tasks, the router’s learning objective closely align with the classifier, making it improbable for the trained router to offer superior routing with reduced complexity. However, for complex tasks, such as object detection involving bounding box regression, a trained router may provide more efficiency by concentrating on the specific routing task. This is a known limitation of using simple tiny model as-is. Nevertheless, placing an evaluator on the top of S ’s final layer might be a solution while preserving most features of TinySaver. This post-evaluator can share the backbone with S and introduce limited overhead to plain TinySaver approach.

In summary, employing S as the router offers clear benefits over trained routers in classification tasks. For other types of tasks, a case-by-case analysis is

necessary, guided by the conditions in Eq. (5). Moreover, utilizing a pre-trained model as a router significantly reducing deployment and upgrade costs. Consequently, a simple tiny model emerges as a competitive option for computational savings in the context of MoE.

3.3 Saver model selection

We propose a method that estimates the compatibility between B and S based on a computation reduction metric ΔC_{tr} on the training dataset. As illustrated in Fig. 3, the performance of B tends to decline more rapidly as the exit threshold is progressively lowered to contain more samples in Cat. I (Reflecting as the increase of the early exit ratio in Fig. 3). At a certain point, the performance curves of S and B intersect, indicating equivalent performance on the early exited samples. This intersection represents the optimal job ratio r_{match} that S model can undertake without compromising performance. This methodology is encapsulated in Eq. (6), where $m \in \{0, 1, \dots, M\}$ denoting the index of each sample in a dataset consisting of M samples. r_{match} is maximized through minimizing t , such that the performance of S is no less than B .

$$r_{match} = \max_t \frac{\#\{m : Conf_S^{(m)} \geq t\}}{M} \quad (6)$$

s.t. $Perf_S \geq Perf_B$

Specifically, for the classification task, this equation can be formulated in Eq. (7) taking into account the predictions made by both S and B for each sample, denoted as $\tilde{\mathbf{y}}_S^{(m,i)}$ and $\tilde{\mathbf{y}}_B^{(m,i)}$, respectively. Where i is the index of N_c classes, $i \in \{0, 1, \dots, N_c\}$. The confidence level associated with each prediction is derived from the maximum value in the probability vector, which can be expressed as $\max_i(\tilde{\mathbf{y}}_S^{(m,i)})$. Here, r_{match} is the maximized ratio of samples that S can take without incurring systematic performance loss.

$$r_{match} = \max_t \frac{\#\{m : \max_i(\tilde{\mathbf{y}}_S^{(m,i)}) \geq t\}}{M} \quad (7)$$

s.t. $\#\{m : \underset{i}{\operatorname{argmax}}(\tilde{\mathbf{y}}_S^{(m,i)}) = \mathbf{y}^{(m)}\}$
 $\geq \#\{m : \underset{i}{\operatorname{argmax}}(\tilde{\mathbf{y}}_B^{(m,i)}) = \mathbf{y}^{(m)}\}$

Then the change of complexity, normalized by the complexity of B , ΔC is computed by comparing Eq. (1) with the standalone complexity C_B . This ΔC is




Fig. 3: Performance vs early exit ratio on the ImageNet-1k training set. The crossing point (marked in red circle) denotes the estimation of how much base model’s job that can be replaced by S model without influencing the systematic performance.

also the computation saving ratio for exiting system.

$$\Delta C = \frac{E[C_{Ts}] - C_B}{C_B} \quad (8)$$

$$= \frac{1}{C_B} (r_{match}C_S + (1 - r_{match})(C_B + C_S) - C_B) \quad (9)$$

$$= \frac{C_S}{C_B} - r_{match} \quad (10)$$

We can apply the aforementioned process on the training set to accurately determine the extent of computational savings, i.e. ΔC_{tr} . Then the metric ΔC_{tr} serves as an indicator of the efficacy with which saver model collaborates with the base model. $\Delta C_{tr} < 0$ signifies that S can compress B without any performance loss on the training set. This calculation method provides a evaluation of the computational change resulting from the integration of the saver and base models within an EE framework.

However, it is important to note the gap between training and testing environments. To ensure reliable compression in practical applications, ΔC_{tr} should be much lower than 0 to guarantee the compression in practice. Additionally, data augmentation should also be enabled while collecting predictions on the training set to prevent saver from being over-confident.

3.4 Extension with early exit sequence

Pre-trained tiny models are more efficient than attached exits. However, traditional EE methods possess notable attributes, such as the ability to reuse the backbone, access to early features, be trained specifically, and the potential for multitasking. These flexibilities are not retained with a single tiny model. To address this, we propose the concept of Exit Sequence Network (ESN). This network is designed to collect features provided by the saver model and establish them as a foundation for exits attached to the backbone of the base model. This approach preserves all the benefits of the early exit strategy and further integrates the saver model into the existing network. It enables intermediate exits to merge progressively refined features from the base model with high-level features from the saver model, offering trade-offs at different levels in predictions. Similar to other EE-based models, the heads built on the ESN is not necessarily having the same task as the original model, providing rich flexibility. However, ESN uniquely includes the tiny model, ensuring a lower bound of performance. The detailed design considerations for ESN are discussed in the Appendix ???. Additionally, as ESN integrates the concept of TinySaver into traditional EE contexts, it allows us evaluate the impact of incorporating an independent tiny model within an EE-based system.

4 Experiment

In this section, we present the comprehensive analysis of TinySaver with its compression impact on a diverse set of vision models, mainly focusing on the ImageNet-1k [52] classification task. Additionally, we explore the saver selection

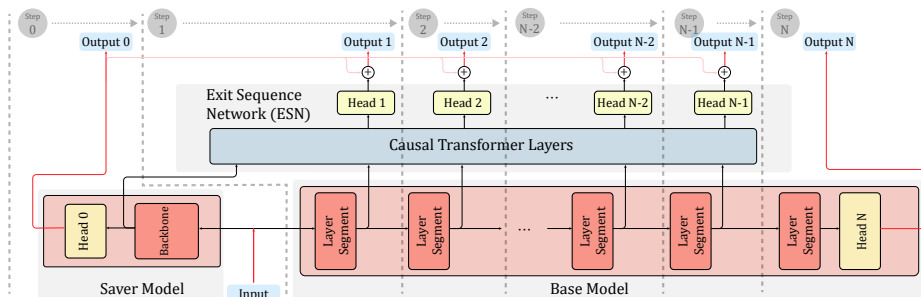


Fig. 4: TinySaver with the Exit Sequence Network

methodology and the ESN framework. Our experiments encompass a pool of 44 models, each capable of utilizing another model as its saver. Detailed descriptions of these models, experimental configurations and the complete results are documented in the appendix ???. Moreover, we extend our validate to the object detection task on the COCO dataset [38].

4.1 Compression of general models

We conduct simulations to assess TinySaver’s efficiency in compressing general models. Inference is performed on the training set, incorporating data augmentation techniques [9], and we log all predictions and associated confidences from a variety of pre-trained models. Subsequently, for each model, we identify a smaller model with the lowest ΔC_{tr} to act as its saver. Following this selection, we test the system on the validation set across varying confidence thresholds $t \in [0, 1]$.

The simulation results for eight models, split between four large and four small, are illustrated in Tab. 1. They demonstrate TinySaver’s effectiveness in significantly reducing overall FLOPs, particularly in larger models. Most models experienced reductions in complexity without sacrificing performance. These models also receive considerable savings while incurring negligible accuracy losses, especially for four large models. The TinySaver even makes some models surpass their original performance with certain threshold settings. Notably, all selected savers are efficiency-focused designs, rather than simply being scaled-down versions of high-performance architectures. This suggests that despite many SOTA works claiming scalability, efficiency-focused models may offer even superior performance at reduced scales.

4.2 ΔC_{tr} and validation performance

To validate our saver selection approach, we conducted simulations by applying various savers to the same base model, with outcomes presented in Tab. 2. Due to the difference between the training and validation set, the saver with the smallest ΔC_{tr} does not invariably equate to the best validation performance. Nonetheless, a high correlation exists between lower ΔC_{tr} and greater computational savings,

Table 1: Compression effect while TinySaver applied to models at various architectures and scales. Evaluation is performed on IN-1K classification.

Base Model Name	GFLOPs.	Acc.	Saver Model Name	GFLOPs.	Acc.	Required GFLOPs(%) when Acc. Drop \leq				Max Performance	
						1%	0.5%	0.01%	0%	Avg. GFLOPs	Acc.
ConvNeXtV2 _{huge} [66]	114.92	86.25%	EfficientFormerV2 _t [35]	2.51	83.53%	13.21	19.41	31.17	67.34	70.70	86.25%
						(-88.5%)	(-83.1%)	(-72.9%)	(-41.4%)	(-38.5%)	(+0.00%)
MaxViT _{large} [59]	42.80	84.83%	EfficientFormerV2 _t [35]	2.51	83.53%	3.46	5.48	7.81	9.04	11.24	84.92%
						(-91.9%)	(-87.2%)	(-81.8%)	(-78.9%)	(-73.7%)	(+0.09%)
DaViT _{base} [14]	15.22	84.49%	EfficientFormerV2 _{s2} [35]	1.22	82.16%	2.40	3.16	4.40	5.19	6.50	84.53%
						(-84.3%)	(-79.2%)	(-71.1%)	(-65.9%)	(-57.3%)	(+0.04%)
Swin _{base} [41]	15.13	85.15%	EfficientFormerV2 _{s2} [35]	1.22	82.16%	2.81	3.70	4.79	9.06	9.40	85.15%
						(-81.4%)	(-75.6%)	(-68.3%)	(-40.1%)	(-37.9%)	(+0.00%)
EfficientNet _{b2} [57]	0.66	77.90%	MobileNetV3 _{large100} [27]	0.22	75.78%	0.27	0.30	0.34	0.36	0.46	78.29%
						(-59.4%)	(-53.9%)	(-47.8%)	(-45.5%)	(-30.9%)	(+0.39%)
EfficientViT _{m5} [40]	0.52	77.08%	MobileNetV3 _{large100} [27]	0.22	75.78%	0.23	0.25	0.28	0.29	0.38	77.55%
						(-55.7%)	(-51.2%)	(-44.9%)	(-43.0%)	(-27.1%)	(+0.47%)
EfficientFormerV2 _{so} [35]	0.38	76.25%	EfficientViT _{so} [2]	0.10	71.35%	0.18	0.20	0.22	0.23	0.26	76.38%
						(-51.0%)	(-46.0%)	(-41.3%)	(-38.8%)	(-30.6%)	(+0.13%)
MobileNetV3 _{large100} [27]	0.22	75.78%	MobileNetV3 _{small100} [27]	0.06	67.65%	0.14	0.16	0.19	0.22	0.22	75.78%
						(-33.6%)	(-26.5%)	(-12.2%)	(+2.8%)	(+3.4%)	(+0.00%)

indicating that this metric can effectively and quickly identify promising saver options in practice.

Table 2: Compression when different savers are selected. Lower ΔC_{tr} correlates with better results. Best cases are marked in bold.

Base Model	Saver Model	Required GFLOPs(%) when Acc. Drop \leq				ΔC		
		1%	0.5%	0.01%	0%			
MaxViT _{large} [59]	EfficientFormerV2 _t [35]	3.47	(-91.9%)	7.81	(-81.8%)	9.31	(-78.2%)	-0.81
	EfficientViT _{so} [2]	6.11	(-85.7%)	10.92	(-74.5%)	13.50	(-68.5%)	-0.79
	DaViT _{large} [14]	7.71	(-82.0%)	13.78	(-67.8%)	16.41	(-61.7%)	-0.77
	EfficientFormerV2 _{so} [35]	4.89	(-88.0%)	11.22	(-73.8%)	12.80	(-70.1%)	-0.77
	ConvNeXtV2 _{large} [66]	7.18	(-83.2%)	12.18	(-71.5%)	15.06	(-64.8%)	-0.76
DaViT _{base} [14]	EfficientFormerV2 _{so} [35]	2.40	(-84.3%)	4.40	(-71.1%)	5.19	(-65.9%)	-0.78
	EfficientFormerV2 _t [35]	-	-	3.96	(-74.0%)	4.19	(-72.4%)	-0.76
	EfficientNet _{b2} [57]	4.86	(-68.1%)	8.71	(-42.8%)	15.18	(-0.5%)	-0.73
	EfficientNet _{so} [57]	5.06	(-66.7%)	9.35	(-58.6%)	15.22	(-0.0%)	-0.71
	ConvNeXtV2 _{base} [66]	4.00	(-73.7%)	6.46	(-57.5%)	7.98	(-47.6%)	-0.7
EfficientNet _{b2} [57]	MobileNetV3 _{large100} [27]	0.27	(-59.2%)	0.35	(-47.5%)	0.36	(-45.5%)	-0.58
	EfficientViT _{so} [2]	0.28	(-57.3%)	0.38	(-42.6%)	0.39	(-41.1%)	-0.47
	MobileNetV3 _{small100} [27]	0.34	(-48.5%)	0.48	(-27.2%)	0.53	(-20.0%)	-0.42
	EfficientViT _{m4} [40]	0.41	(-37.5%)	0.49	(-26.0%)	0.51	(-23.0%)	-0.4
	EfficientViT _{m3} [40]	0.41	(-38.2%)	0.50	(-23.9%)	0.54	(-17.5%)	-0.39

4.3 Testing with ESN

In Tab. 3, we detail TinySaver’s performance when enhanced with ESN, testing it across the same eight models and savers as in Tab. 1. We conducted a basic search for optimal hyper-parameters, including the number of attention layers and dimensions of hidden neurons. The listed results reflect the best configurations for different performance goals. Additionally, we employed DyCE [61] to tune exit thresholds, with full details in ???. Our findings reveal no substantial improvements using ESN compared to the plain TinySaver, suggesting that a tiny model is already highly efficient and effective, in classification, making traditional EE architectures unable to provide further significant benefits.

Table 3: TinySaver with ESN enabled. Each base model is employing the same saver as in Tab. 1, figures have improvement are marked in bold.

Base Model Name	Saver Model Name	Required GFLOPs(%) when Acc. Drop \leq			
		1%	0.5%	0.01%	0%
ConvNeXtV2 _{large} [66]	EfficientFormerV2 _t [35]	13.07	19.67	-	-
		(-88.4%)	(-82.0%)	-	-
MaxViT _{large} [59]	EfficientFormerV2 _t [35]	3.46	5.48	7.79	9.03
		(-91.9%)	(-87.2%)	(-81.8%)	(-78.9%)
DaViT _{base} [14]	EfficientFormerV2 _{so} [35]	2.40	3.16	4.38	5.12
		(-84.2%)	(-79.2%)	(-71.2%)	(-66.4%)
Swin _{base} [41]	EfficientFormerV2 _{so} [35]	2.80	3.69	4.72	-
		(-81.5%)	(-75.6%)	(-68.8%)	-
EfficientNet _{b2} [57]	MobileNetV3 _{large100} [27]	0.27	0.30	0.34	0.36
		(-59.3%)	(-54.0%)	(-47.7%)	(-45.5%)
EfficientViT _{m5} [40]	MobileNetV3 _{large100} [27]	0.23	0.25	0.28	0.29
		(-55.8%)	(-51.3%)	(-45.3%)	(-43.2%)
EfficientFormerV2 _{so} [35]	EfficientViT _{so} [2]	0.19	0.20	0.22	0.23
		(-51.0%)	(-45.8%)	(-41.1%)	(-38.4%)
MobileNetV3 _{large100} [27]	MobileNetV3 _{small100} [27]	0.14	0.16	0.19	0.21
		(-33.6%)	(-26.5%)	(-12.2%)	(-2.9%)

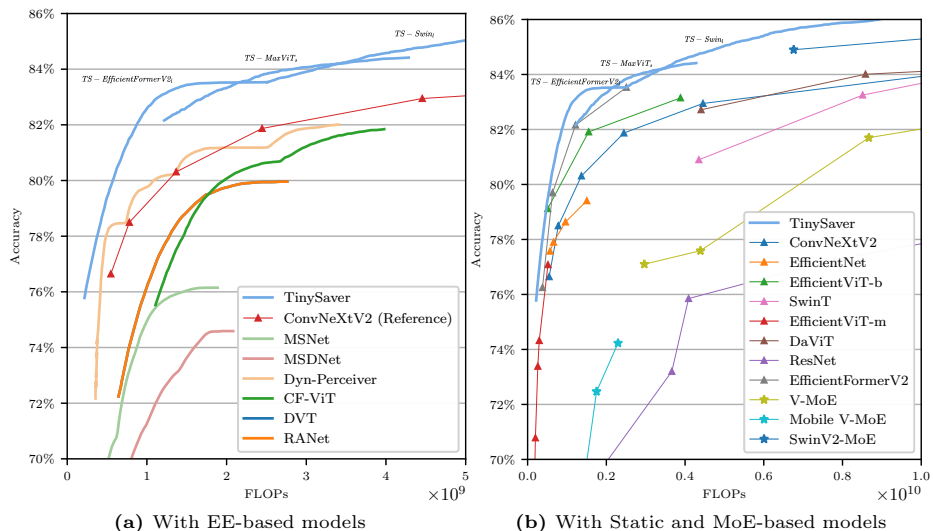


Fig. 5: Comparison of the top-1 accuracy on ImageNet-1k between TinySaver enabled Models and related work

4.4 Comparison with related work

In Fig. 5, we compare TinySaver against various existing models from EE-based [6, 19, 22, 26, 63], MoE-based [12, 28, 50], and static categories [2, 14, 15, 25, 35, 40, 57, 59, 66]. As presented in Fig. 5(a), TinySaver outperforms EE-based methods significantly due to its model-agnostic feature, which allows it to seamlessly leverage improvement across models, datasets, and training recipes, while the progress of other EE-based work relies on refining design gradually. Compared to static models in Fig. 5(b), TinySaver surpasses scaling down model architectures directly. It enhances the original model with efficiencies from other models, outperforming smaller versions of the original. The TinySaver enabled model also usually has higher accuracy than all static models in a similar complexity range. Moreover, as a dynamic method which has a continuous performance curve, it is configurable in fine-grain, during runtime to adapt to specific needs. It is also orthogonal to most existing compression methods like quantization. Combining these approaches could yield further efficiency gains. We also include the comparison with MoE-based models in Fig. 5(b). Although MoE-based architectures are efficient in scaling models up, they are not as efficient as TinySaver in reducing complexity. All these observations demonstrate the effectiveness of integrating tiny models as computational savers.

4.5 Test on COCO detection

We extended our experiments to the COCO dataset [38] for object detection, utilizing the scores of detected boxes as the confidence metric. Specifically, we

filter out boxes with scores below 0.05 and use the average score of the remaining boxes as the model’s confidence for each sample. We employ a real-time YOLOv8 [29] detectors as savers. Images not confidently predicted by YOLOv8 [29] are passed to a larger model. Results compared with other detectors are presented in Tab. 4. The proposed method demonstrates superiority in several scenarios. While Co-DINO-Swin-L [78] excels in high complexity scenarios, YOLOv8 [29] offers greater efficiency though with constrained performance. TinySaver manages to extend their advantages to other complexity ranges. However, in object detection, it does not outperform all static models as it does in classification, likely due to bounding box quality being less correlated with score. This outcome aligns with our analysis in Sec. 3.2 and suggests further exploration of external evaluators for such tasks.

5 Conclusion

This paper explores the feasibility and effectiveness of using tiny models to decrease the computation of large models. We introduce TinySaver, a novel methodology that employs the most appropriate tiny model to reduce the workload of large models on simpler samples, thus reducing overall complexity. Furthermore, we analyze our approach within early exit (EE) and Mixture of Experts (MoE) frameworks and discuss its broader application in complexity reduction. Our methods have undergone thorough testing with various models which are designed and trained in different ways, showing significant computational reductions across most of them, highlighting the potential of integrating tiny models into extensive AI systems to save on computation.

Table 4: Results of objection detection on COCO 2017val [38].

Method	Model	Pretrained Dataset	Box AP	GFLOPs
-	YOLOv8n [29]	IN1k	37.3	9
-	YOLOv8s [29]	IN1k	44.9	29
Efficient-DETR [72]	R50 [25]	IN1k	45.1	210
TinySaver($t = 0.25$)	YOLOv8x [29]	IN1k	45.3	71
-	+YOLOv8n [29]			
-	YOLOv8m [29]	IN1k	50.2	79
TinySaver($t = 0.25$)	YOLOv8l [29]	IN1k	50.7	114
-	+YOLOv8s [29]			
Cascade Mask R-CNN [3]	Swin-T [41]	IN1k	50.4	745
TinySaver($t = 0.25$)	Co-DINO-Swin-L [78]	Obj365 [54], IN1k	51.5	589
-	+YOLOv8n [29]			
Cascade Mask R-CNN [3]	Swin-S [41]	IN1k	51.9	838
-	YOLOv8l [29]	IN1k	52.9	165
Mask-RCNN [24]	ConvNeXt V2-B [66]	IN1k	52.9	486
TinySaver($t = 0.26$)	Co-DINO-Swin-L [78]	Obj365 [54], IN1k	53.1	645
-	+YOLOv8n [29]			
RT-DETR [42]	R50 [25]	IN1k	53.1	136
-	YOLOv8x [29]	IN1k	53.9	258
HTC++ [5]	Swin-B [41]	IN22k	56.4	1043
TinySaver($t = 0.34$)	Co-DINO-Swin-L [78]	Obj365 [54], IN1k	56.5	671
-	+YOLOv8x [29]			
HTC++ [5]	Swin-L [41]	IN22k	57.1	1470
TinySaver($t = 0.37$)	Co-DINO-Swin-L [78]	Obj365 [54], IN1k	60.3	1076
-	+YOLOv8n [29]			
Soft-teacher [70], HTC++ [5]	Swin-L [41]	Obj365 [54]	60.1	1470
TinySaver($t = 0.69$)	Co-DINO-Swin-L [78]	Obj365 [54], IN1k	63.0	1497
-	+YOLOv8n [29]			
Co-DINO [78]	Swin-L [41]	Obj365 [54]	64.1	1527 [†]

[†]The FLOPs is measured by us, otherwise are collected from corresponding work.

Acknowledgment

This publication has emanated from research conducted with the financial support of 1) Science Foundation Ireland under Grant number 18/CRT/6183 2) Horizon 2020 ERA-NET Chist-Era grant, CHIST-ERA-18-ACAI-003 and 3) Microelectronic Circuits Centre Ireland. For the purpose of Open Access, the

author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

1. Cai, H., Gan, C., Wang, T., Zhang, Z., Han, S.: Once-for-all: Train one network and specialize it for efficient deployment. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net (2020), <https://openreview.net/forum?id=HylxE1HKwS>
2. Cai, H., Li, J., Hu, M., Gan, C., Han, S.: EfficientViT: Multi-Scale Linear Attention for High-Resolution Dense Prediction (2022), <https://arxiv.org/abs/2205.14756>
3. Cai, Z., Vasconcelos, N.: Cascade R-CNN: delving into high quality object detection. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 6154–6162. IEEE Computer Society (2018). <https://doi.org/10.1109/CVPR.2018.00644>, http://openaccess.thecvf.com/content_cvpr_2018/html/Cai_Cascade_R-CNN_Delving_CVPR_2018_paper.html
4. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-End Object Detection with Transformers (2020), <https://arxiv.org/abs/2005.12872>
5. Chen, K., Pang, J., Wang, J., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: Hybrid task cascade for instance segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 4974–4983. Computer Vision Foundation / IEEE (2019). <https://doi.org/10.1109/CVPR.2019.00511>, http://openaccess.thecvf.com/content_cvpr_2019/html/Chen_Hybrid_Task_Cascade_for_Instance_Segmentation_CVPR_2019_paper.html
6. Chen, M., Lin, M., Li, K., Shen, Y., Wu, Y., Chao, F., Ji, R.: CF-ViT: A General Coarse-to-Fine Method for Vision Transformer (2022), <https://arxiv.org/abs/2203.03821>
7. Chen, X., Dai, H., Li, Y., Gao, X., Song, L.: Learning to stop while learning to predict. In: Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event. Proceedings of Machine Learning Research, vol. 119, pp. 1520–1530. PMLR (2020), <http://proceedings.mlr.press/v119/chen20c.html>
8. Cheng, Y., Wang, D., Zhou, P., Zhang, T.: A Survey of Model Compression and Acceleration for Deep Neural Networks (2017), <https://arxiv.org/abs/1710.09282>
9. Cubuk, E.D., Zoph, B., Mané, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 113–123. Computer Vision Foundation / IEEE (2019). <https://doi.org/10.1109/CVPR.2019.00020>, http://openaccess.thecvf.com/content_cvpr_2019/html/Cubuk_AutoAugment_Learning_Augmentation_Strategies_From_Data_CVPR_2019_paper.html
10. Dai, X., Kong, X., Guo, T.: Epnet: Learning to exit with flexible multi-branch network. In: d’Aquin, M., Dietze, S., Hauff, C., Curry, E., Cudré-Mauroux, P. (eds.) CIKM ’20: The 29th ACM International Conference on Information and

- Knowledge Management, Virtual Event, Ireland, October 19-23, 2020. pp. 235–244. ACM (2020). <https://doi.org/10.1145/3340531.3411973>, <https://doi.org/10.1145/3340531.3411973>
11. Dai, Z., Liu, H., Le, Q.V., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. In: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*. pp. 3965–3977 (2021), <https://proceedings.neurips.cc/paper/2021/hash/20568692db622456cc42a2e853ca21f8-Abstract.html>
 12. Daxberger, E., Weers, F., Zhang, B., Gunter, T., Pang, R., Eichner, M., Emersberger, M., Yang, Y., Toshev, A., Du, X.: *Mobile V-MoEs: Scaling Down Vision Transformers via Sparse Mixture-of-Experts (2023)*, <https://arxiv.org/abs/2309.04354>
 13. Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A.P., Caron, M., Geirhos, R., Alabdulmohsin, I., Jenatton, R., Beyer, L., Tschannen, M., Arnab, A., Wang, X., Ruiz, C.R., Minderer, M., Puigcerver, J., Evci, U., Kumar, M., Steenkiste, S.V., Elsayed, G.F., Mahendran, A., Yu, F., Oliver, A., Huot, F., Bastings, J., Collier, M., Gritsenko, A.A., Birodkar, V., Vasconcelos, C.N., Tay, Y., Mensink, T., Kolesnikov, A., Pavetic, F., Tran, D., Kipf, T., Lucic, M., Zhai, X., Keysers, D., Harmsen, J.J., Houlsby, N.: *Scaling Vision Transformers to 22 Billion Parameters*. In: *Proceedings of the 40th International Conference on Machine Learning*. pp. 7480–7512. PMLR, <https://proceedings.mlr.press/v202/dehghani23a.html>
 14. Ding, M., Xiao, B., Codella, N., Luo, P., Wang, J., Yuan, L.: *DaViT: Dual Attention Vision Transformers*. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 74–92. *Lecture Notes in Computer Science*, Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-20053-3_5
 15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: *An image is worth 16x16 words: Transformers for image recognition at scale*. In: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net (2021), <https://openreview.net/forum?id=YicbFdNTTy>
 16. Elskens, T., Metzen, J.H., Hutter, F.: *Neural Architecture Search: A Survey*
 17. Fedus, W.e.a.: *Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity (2021)*, <https://arxiv.org/abs/2101.03961>
 18. Gururangan, S.e.a.: *Scaling Expert Language Models with Unsupervised Domain Discovery (2023)*, <https://arxiv.org/abs/2303.14177>
 19. Han, Y., Han, D., Liu, Z., Wang, Y., Pan, X., Pu, Y., Deng, C., Feng, J., Song, S., Huang, G.: *Dynamic Perceiver for Efficient Visual Recognition (2023)*, <https://arxiv.org/abs/2306.11248>
 20. Han, Y., Huang, G., Song, S., Yang, L., Wang, H., Wang, Y.: *Dynamic Neural Networks: A Survey (2021)*, <https://arxiv.org/abs/2102.04906>
 21. Han, Y., Pu, Y., Lai, Z., Wang, C., Song, S., Cao, J., Huang, W., Deng, C., Huang, G.: *Learning to Weight Samples for Dynamic Early-exiting Networks (2022)*, <https://arxiv.org/abs/2209.08310>
 22. Hang, R., Qian, X., Liu, Q.: *MSNet: Multi-Resolution Synergistic Networks for Adaptive Inference* **33**(5), 2009–2018. <https://doi.org/10.1109/TCSVT.2022.3218891>

23. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.B.: Masked autoencoders are scalable vision learners. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 15979–15988. IEEE (2022). <https://doi.org/10.1109/CVPR52688.2022.01553>, <https://doi.org/10.1109/CVPR52688.2022.01553>
24. He, K., Gkioxari, G., Dollár, P., Girshick, R.B.: Mask R-CNN. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. pp. 2980–2988. IEEE Computer Society (2017). <https://doi.org/10.1109/ICCV.2017.322>, <https://doi.org/10.1109/ICCV.2017.322>
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. pp. 770–778. IEEE Computer Society (2016). <https://doi.org/10.1109/CVPR.2016.90>, <https://doi.org/10.1109/CVPR.2016.90>
26. Hong, S., Kaya, Y., Modoranu, I., Dumitras, T.: A panda? no, it’s a sloth: Slow-down attacks on adaptive multi-exit neural network inference. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021), <https://openreview.net/forum?id=9xC2tWEwBD>
27. Howard, A., Pang, R., Adam, H., Le, Q.V., Sandler, M., Chen, B., Wang, W., Chen, L., Tan, M., Chu, G., Vasudevan, V., Zhu, Y.: Searching for mobilenetv3. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 1314–1324. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00140>, <https://doi.org/10.1109/ICCV.2019.00140>
28. Hwang, C., Cui, W., Xiong, Y., Yang, Z., Liu, Z., Hu, H., Wang, Z., Salas, R., Jose, J., Ram, P., Chau, J., Cheng, P., Yang, F., Yang, M., Xiong, Y.: Tutel: Adaptive Mixture-of-Experts at Scale (2022), <https://arxiv.org/abs/2206.03382>
29. Jocher, G., Chaurasia, A., Qiu, J.: Ultralytics yolov8 (2023), <https://github.com/ultralytics/ultralytics>
30. Kaya, Y., Hong, S., Dumitras, T.: Shallow-deep networks: Understanding and mitigating network overthinking. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA. Proceedings of Machine Learning Research, vol. 97, pp. 3301–3310. PMLR (2019), <http://proceedings.mlr.press/v97/kaya19a.html>
31. Komatsuzaki, A.e.a.: Sparse Upcycling: Training Mixture-of-Experts from Dense Checkpoints (2022), <https://arxiv.org/abs/2212.05055>
32. Laskaridis, S., Kouris, A., Lane, N.D.: Adaptive Inference through Early-Exit Networks: Design, Challenges and Directions. In: Proceedings of the 5th International Workshop on Embedded and Mobile Deep Learning. pp. 1–6. ACM. <https://doi.org/10.1145/3469116.3470012>, <https://dl.acm.org/doi/10.1145/3469116.3470012>
33. Laskaridis, S., Venieris, S.I., Almeida, M., Leontiadis, I., Lane, N.D.: SPINN: Synergistic progressive inference of neural networks over device and cloud. In: Proceedings of the 26th Annual International Conference on Mobile Computing and Networking. pp. 1–15. MobiCom ’20, Association for Computing Machinery. <https://doi.org/10.1145/3372224.3419194>, <https://dl.acm.org/doi/10.1145/3372224.3419194>
34. Laskaridis, S., Venieris, S.I., Kim, H., Lane, N.D.: HAPI: Hardware-Aware Progressive Inference. vol. abs/2008.03997 (2020), <https://arxiv.org/abs/2008.03997>

35. Li, Y., Hu, J., Wen, Y., Evangelidis, G., Salah, K., Wang, Y., Tulyakov, S., Ren, J.: Rethinking Vision Transformers for MobileNet Size and Speed. pp. 16889–16900. https://openaccess.thecvf.com/content/ICCV2023/html/Li_Rethinking_Vision_Transformers_for_MobileNet_Size_and_Speed_ICCV_2023_paper.html
36. Liao, K., Zhang, Y., Ren, X., Su, Q., Sun, X., He, B.: A global past-future early exit method for accelerating inference of pre-trained language models. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 2013–2023. Association for Computational Linguistics, Online (2021). <https://doi.org/10.18653/v1/2021.naacl-main.162>, <https://aclanthology.org/2021.naacl-main.162>
37. Lin, T., Dollár, P., Girshick, R.B., He, K., Hariharan, B., Belongie, S.J.: Feature pyramid networks for object detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017. pp. 936–944. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.106>, <https://doi.org/10.1109/CVPR.2017.106>
38. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft COCO: Common Objects in Context. <https://doi.org/10.48550/arXiv.1405.0312>, <http://arxiv.org/abs/1405.0312>
39. Liu, W., Zhou, P., Wang, Z., Zhao, Z., Deng, H., Ju, Q.: FastBERT: a self-distilling BERT with adaptive inference time. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 6035–6044. Association for Computational Linguistics, Online (2020). <https://doi.org/10.18653/v1/2020.acl-main.537>, <https://aclanthology.org/2020.acl-main.537>
40. Liu, X., Peng, H., Zheng, N., Yang, Y., Hu, H., Yuan, Y.: EfficientViT: Memory Efficient Vision Transformer With Cascaded Group Attention. pp. 14420–14430. https://openaccess.thecvf.com/content/CVPR2023/html/Liu_EfficientViT_Memory_Efficient_Vision_Transformer_With_Cascaded_Group_Attention_CVPR_2023_paper.html
41. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021. pp. 9992–10002. IEEE (2021). <https://doi.org/10.1109/ICCV48922.2021.00986>, <https://doi.org/10.1109/ICCV48922.2021.00986>
42. Lv, W., Zhao, Y., Xu, S., Wei, J., Wang, G., Cui, C., Du, Y., Dang, Q., Liu, Y.: DETRs Beat YOLOs on Real-time Object Detection (2023), <https://arxiv.org/abs/2304.08069>
43. Ma, J., Zhao, Z., Yi, X., Chen, J., Hong, L., Chi, E.H.: Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In: Guo, Y., Farooq, F. (eds.) Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19–23, 2018. pp. 1930–1939. ACM (2018). <https://doi.org/10.1145/3219819.3220007>, <https://doi.org/10.1145/3219819.3220007>
44. maintainers, T., contributors: TorchVision: PyTorch’s Computer Vision library (2016), <https://github.com/pytorch/vision>
45. Mustafa, B., Riquelme, C., Puigcerver, J., Jenatton, R., Houlsby, N.: Multi-modal Contrastive Learning with LIMoE: The Language-Image Mixture of Experts (2022), <https://arxiv.org/abs/2206.02770>
46. OpenAI: GPT-4 Technical Report (2023), <https://arxiv.org/abs/2303.08774>
47. Panda, P., Sengupta, A., Roy, K.: Conditional Deep Learning for Energy-Efficient and Enhanced Pattern Recognition (2015), <https://arxiv.org/abs/1509.08971>

48. Phuong, M., Lampert, C.: Distillation-based training for multi-exit architectures. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 1355–1364. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00144>, <https://doi.org/10.1109/ICCV.2019.00144>
49. Rasley, J., Rajbhandari, S., Ruwase, O., He, Y.: Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In: Gupta, R., Liu, Y., Tang, J., Prakash, B.A. (eds.) KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020. pp. 3505–3506. ACM (2020), <https://dl.acm.org/doi/10.1145/3394486.3406703>
50. Riquelme, C., Puigcerver, J., Mustafa, B., Neumann, M., Jenatton, R., Pinto, A.S., Keyzers, D., Houlsby, N.: Scaling vision with sparse mixture of experts. In: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual. pp. 8583–8595 (2021), <https://proceedings.neurips.cc/paper/2021/hash/48237d9f2dea8c74c2a72126cf63d933-Abstract.html>
51. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation (2015), <https://arxiv.org/abs/1505.04597>
52. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge **115**(3), 211–252
53. Scardapane, S., Comminiello, D., Scarpiniti, M., Baccarelli, E., Uncini, A.: Differentiable branching in deep networks for fast inference. In: 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020. pp. 4167–4171. IEEE (2020). <https://doi.org/10.1109/ICASSP40776.2020.9054209>, <https://doi.org/10.1109/ICASSP40776.2020.9054209>
54. Shao, S., Li, Z., Zhang, T., Peng, C., Yu, G., Zhang, X., Li, J., Sun, J.: Objects365: A large-scale, high-quality dataset for object detection. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 8429–8438. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00852>, <https://doi.org/10.1109/ICCV.2019.00852>
55. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q.V., Hinton, G.E., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net (2017), <https://openreview.net/forum?id=B1ckMDq1g>
56. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. pp. 843–852. IEEE Computer Society (2017). <https://doi.org/10.1109/ICCV.2017.97>, <https://doi.org/10.1109/ICCV.2017.97>
57. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA. Proceedings of Machine Learning Research, vol. 97, pp. 6105–6114. PMLR (2019), <http://proceedings.mlr.press/v97/tan19a.html>

58. Teerapittayanon, S., McDanel, B., Kung, H.T.: BranchyNet: Fast Inference via Early Exiting from Deep Neural Networks (2017), <https://arxiv.org/abs/1709.01686>
59. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: MaxViT: Multi-axis Vision Transformer. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 459–479. *Lecture Notes in Computer Science*, Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-20053-3_27
60. Wang, J., Chen, K., Chen, G., Shou, L., McAuley, J.: SkipBERT: Efficient inference with shallow layer skipping. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 7287–7301. Association for Computational Linguistics, Dublin, Ireland (2022). <https://doi.org/10.18653/v1/2022.acl-long.503>, <https://aclanthology.org/2022.acl-long.503>
61. Wang, Q., Cardiff, B., Frappé, A., Larras, B., John, D.: DyCE: Dynamic Configurable Exiting for Deep Learning Compression and Scaling (2024), <https://arxiv.org/abs/2403.01695>
62. Wang, X., Yu, F., Dou, Z.Y., Darrell, T., Gonzalez, J.E.: SkipNet: Learning Dynamic Routing in Convolutional Networks (2017), <https://arxiv.org/abs/1711.09485>
63. Wang, Y., Huang, R., Song, S., Huang, Z., Huang, G.: Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. In: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*. pp. 11960–11973 (2021), <https://proceedings.neurips.cc/paper/2021/hash/64517d8435994992e682b3e4aa0a0661-Abstract.html>
64. Wang, Y., Lv, K., Huang, R., Song, S., Yang, L., Huang, G.: Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual (2020)*, <https://proceedings.neurips.cc/paper/2020/hash/1963bd5135521d623f6c29e6b1174975-Abstract.html>
65. Wightman, R., Raw, N., Soare, A., Arora, A., Ha, C., Reich, C., Guan, F., Kaczmazzyk, J., MrT23, Mike, SeeFun, Contrastive, Rizin, M., Hyeongchan Kim, Kertész, C., Dushyant Mehta, Cucurull, G., Kushajveer Singh, Hankyul, Tatsunami, Y., Lavin, A., Juntang Zhuang, Hollemans, M., Rashad, M., Sameni, S., Shults, V., Lucain, Wang, X., Yonghye Kwon, Uchida, Y.: Rwrightman/pytorch-image-models: V0.8.10dev0 Release. <https://doi.org/10.5281/ZENODO.4414861>, <https://zenodo.org/record/4414861>
66. Woo, S., Debnath, S., Hu, R., Chen, X., Liu, Z., Kweon, I.S., Xie, S.: ConvNeXt V2: Co-Designing and Scaling ConvNets With Masked Autoencoders. pp. 16133–16142. https://openaccess.thecvf.com/content/CVPR2023/html/Woo_ConvNeXt_V2_Co-Designing_and_Scaling_ConvNets_With_Masked_Autoencoders_CVPR_2023_paper.html
67. Wu, Z., Nagarajan, T., Kumar, A., Rennie, S., Davis, L.S., Grauman, K., Feris, R.S.: Blockdrop: Dynamic inference paths in residual networks. In: *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. pp. 8817–8826. IEEE Computer Society (2018).

- <https://doi.org/10.1109/CVPR.2018.00919>, http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_BlockDrop_Dynamic_Inference_CVPR_2018_paper.html
68. Xing, Q., Xu, M., Li, T., Guan, Z.: Early Exit or Not: Resource-Efficient Blind Quality Enhancement for Compressed Images. vol. abs/2006.16581 (2020), <https://arxiv.org/abs/2006.16581>
 69. Xu, C., McAuley, J.: A survey on dynamic neural networks for natural language processing. In: Findings of the Association for Computational Linguistics: EACL 2023. pp. 2370–2381. Association for Computational Linguistics, Dubrovnik, Croatia (2023), <https://aclanthology.org/2023.findings-eacl.180>
 70. Xu, M., Zhang, Z., Hu, H., Wang, J., Wang, L., Wei, F., Bai, X., Liu, Z.: End-to-end semi-supervised object detection with soft teacher. In: 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021. pp. 3040–3049. IEEE (2021). <https://doi.org/10.1109/ICCV48922.2021.00305>, <https://doi.org/10.1109/ICCV48922.2021.00305>
 71. Yang, L., Han, Y., Chen, X., Song, S., Dai, J., Huang, G.: Resolution adaptive networks for efficient inference. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020. pp. 2366–2375. IEEE (2020). <https://doi.org/10.1109/CVPR42600.2020.00244>, <https://doi.org/10.1109/CVPR42600.2020.00244>
 72. Yao, Z., Ai, J., Li, B., Zhang, C.: Efficient DETR: Improving End-to-End Object Detector with Dense Prior (2021), <https://arxiv.org/abs/2104.01318>
 73. Yokoyama, A.M., Ferro, M., de Paula, F.B., Vieira, V.G., Schulze, B.: Investigating hardware and software aspects in the energy consumption of machine learning: A green AI-centric analysis **35**(24), e7825. <https://doi.org/10.1002/cpe.7825>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.7825>
 74. Yun, S., Han, D., Chun, S., Oh, S.J., Yoo, Y., Choe, J.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 6022–6031. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00612>, <https://doi.org/10.1109/ICCV.2019.00612>
 75. Zhang, H., Cissé, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net (2018), <https://openreview.net/forum?id=r1Ddp1-Rb>
 76. Zhang, M., Chu, C., Zhmoginov, A., Howard, A., Jou, B., Zhu, Y., Zhang, L., Hwa, R., Kovashka, A.: Basisnet: Two-stage model synthesis for efficient inference. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2021, virtual, June 19–25, 2021. pp. 3081–3090. Computer Vision Foundation / IEEE (2021). <https://doi.org/10.1109/CVPRW53098.2021.00344>, https://openaccess.thecvf.com/content/CVPR2021W/ECV/html/Zhang_BasisNet_Two-Stage_Model_Synthesis_for_Efficient_Inference_CVPRW_2021_paper.html
 77. Zhu, W.: LeeBERT: Learned early exit for BERT with cross-level optimization. In: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). pp. 2968–2980. Association for Computational Linguistics, Online (2021). <https://doi.org/10.18653/v1/2021.acl-long.231>, <https://aclanthology.org/2021.acl-long.231>
 78. Zong, Z., Song, G., Liu, Y.: DETRs with Collaborative Hybrid Assignments Training (2022), <https://arxiv.org/abs/2211.12860>