

LORS: Low-rank Residual Structure for Parameter-Efficient Network Stacking

Jialin Li, Qiang Nie, Weifu Fu, Yuhuan Lin, Guangpin Tao, Yong Liu, Chengjie Wang
YouTu Lab, Tencent

{jarenli, ryanwfu, gleelin, guangpintao, choasliu, jasoncjwang}@tencent.com, qnie.cuhk@gmail.com

Abstract

Deep learning models, particularly those based on transformers, often employ numerous stacked structures, which possess identical architectures and perform similar functions. While effective, this stacking paradigm leads to a substantial increase in the number of parameters, posing challenges for practical applications. In today's landscape of increasingly large models, stacking depth can even reach dozens, further exacerbating this issue. To mitigate this problem, we introduce **LORS (LOW-rank Residual Structure)**. LORS allows stacked modules to share the majority of parameters, requiring a much smaller number of unique ones per module to match or even surpass the performance of using entirely distinct ones, thereby significantly reducing parameter usage. We validate our method by applying it to the stacked decoders of a query-based object detector, and conduct extensive experiments on the widely used MS COCO dataset. Experimental results demonstrate the effectiveness of our method, as even with a 70% reduction in the parameters of the decoder, our method still enables the model to achieve comparable or even better performance than its original.

1. Introduction

In the current era of prosperity for large models, a common issue is the significant increase in the number of parameters, which presents challenges for training, inference, and deployment. Various methods have been proposed to reduce the number of parameters in models, such as knowledge distillation [15, 20], which compresses large models into smaller ones while trying to preserve their performance but may still lead to a decrease in model capacity; pruning [16, 55], which removes redundant parameters from the model but can affect the model's stability; quantization [8], which reduces the numerical precision of model parameters to lower storage and computation but may cause model accuracy loss; and parameter sharing [24], which reduces the number of parameters by sharing them across different layers but may limit the model's expressiveness.

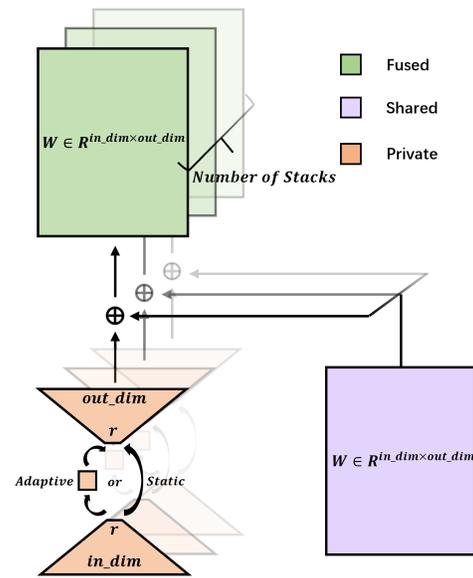


Figure 1. The LORS calculation process, which could be adaptive or static, depending on whether an adaptively generated kernel is used in the matrix manipulation for private parameters.

Different from the aforementioned methods, we have observed an important fact contributing to the large number of parameters: the widespread use of stacking in neural networks. Stacking refers to those modules that have identical architectures and perform the same or similar functions, but possess different parameters due to random initialization as well as training updates. Examples of stacking can be found in numerous prominent neural networks, like the classic ResNet model [18] and Transformers [46]. Particularly, Transformers heavily rely on stacked structures and typically employ completely identical multi-layer stacks in both encoders and decoders. It now serves as an indispensable component of many excellent models in fields such as computer vision and natural language processing.

Although stacking is powerful for enhancing model capacity, as demonstrated by large language models, it also

naturally leads to a sharp increase in the number of parameters. For example, GPT-3 [2] utilizes 175 billion parameters and consists of 96 layers of stacked Transformer [46] layers. How can we enjoy the benefits of stacking while reducing the required number of parameters? We notice that stacked decoders have identical structures and similar functions, indicating that there should be some commonality among their parameters. However, since they handle different input and output distributions, there must also be unique aspects to their parameters. Therefore, a natural idea is: it may be possible to represent the shared aspects with shared parameters while allowing each stacked module to retain only the parameters that capture its unique characteristics, thereby reducing the overall parameter usage.

Based on the above considerations, we propose to decompose the parameters of stacked modules into two parts: shared ones representing the commonality and private ones capturing the specific characteristics. Shared parameters are available for all modules and trained jointly by them, while the private parameters are separately owned by each module. We suppose that the way of sharing parameters related to the commonality may reduce the number of parameters while maintaining the performance of the model. To achieve this goal, we introduce the concept of Low-rank Residual Structure (LORS), inspired by the approach of LoRA [21]. LORS essentially adds the unique parameters to the shared ones, just like residual connections add residual information to the features. While LoRA is originally designed for fine-tuning, we train our LoRA-like operation on parameters from scratch. This approach allows us to effectively reduce the overall parameter usage while maintaining the performance of the model, paving the way for more parameter-efficient network stacking.

To validate our idea, we choose AdaMixer [12], a strong query-based object detector, as our experimental subject. It contains a large number of both adaptive and static parameters in its stacked decoders, making it an ideal candidate to demonstrate LORS’s effectiveness. The difference between adaptive and static parameters lies in whether they change with different inputs. Our goal is to show that LORS can effectively reduce the overall usage of both types of parameters while maintaining the model’s performance. We conducted extensive experiments on this detector, illustrating that our method succeeded in reducing up to 70% of parameters in AdaMixer’s decoders, while still enabling the model to achieve comparable or even superior performance compared to its vanilla version. In summary, our contributions can be concluded as :

- We propose a novel low-rank residual structure, named LORS, for network stacking, which can greatly reduce the number of parameters while maintaining or even improving the performance compared to the vanilla one.
- We introduce effective methodologies for reducing both

static and adaptive generated parameters in stacked structures, which makes our proposed LORS a more versatile method.

- Our method holds the potential to serve as one of the basic network structures for large models that are affected by the issue of excessive parameters due to stacking, rendering them more parameter-efficient, thereby facilitating an easier implementation in practical applications.

2. Related Work

Models with stacked structures. Many neural networks partially or extensively employ stacked modules. CNN-based models [40], for instance, are widely applied across various computer vision tasks, such as classification [18, 22, 39, 49], detection [19, 25, 36–38, 43], segmentation [7, 13, 32, 54], etc. These models often use stacked smaller modules within larger components. Another kind of models primarily utilize Multi-Layer Perceptrons (MLPs), such as the MLP-Mixer [44], whose main body comprises dozens of completely identical stacked mixer layers. Moreover, since its invention, Transformer [46] has been successfully applied to diverse domains, including computer vision [3, 30, 41, 56], natural language processing [2, 10, 50], multi-modal learning [6, 27, 35, 47], etc. In the thriving arena of large language models, Transformer-based structures are indeed indispensable for many cutting-edge works [10, 26, 34, 42, 48]. Transformers are usually used to form stacked multi-layers encoders or decoders, within each layer being structurally identical.

LoRA and Its Variants. LoRA [21] is a technique proposed for fine-tuning large language models, it introduces low-rank decomposition matrices in the Transformer [46] architecture, significantly reducing trainable parameters and GPU memory requirements. Afterward, a series of works with modifications as well as improvement have been proposed. AdaLoRA [53] adaptively allocates parameter budgets among weight matrices based on their importance scores and uses singular value decomposition for incremental updates, leading to an improved fine-tuning performance, especially in low budget settings. DyLoRA [45] introduces a dynamic low-rank adaptation technique that trains LoRA blocks for a range of ranks instead of a single rank, addressing the issues of fixed block size and rank optimization in LoRA. Delta-LoRA [57] presents a parameter-efficient approach to fine-tune large language models by updating not only the low-rank matrices but also the pre-trained weights using the delta of the product of two low-rank matrices. LoRA-FA [52] introduces a memory-efficient fine-tuning method that reduces activation memory without performance degradation or expensive recomputation, it achieves this by freezing the projection-down weight and updating the projection-up weight in each LoRA layer. ReLoRA [28] is a parameter-efficient method for training

large neural networks using low-rank updates. GLoRA [4], or Generalized LoRA, is an advanced approach for universal parameter-efficient fine-tuning tasks, which enhances the low rank adaptation technique by employing a generalized prompt module to optimize pre-trained model weights and adjust intermediate activations. VeRA [23], or Vector-based Random Matrix Adaptation, further compresses the training parameters compared to the vanilla LoRA [21]. It achieves this by utilizing shared low-rank matrices across all layers and learning small scaling vectors.

Inspired by the aforementioned LoRA series works, we propose our method called LORS (Low Rank Residual Structure) for parameter efficiency. However, our approach is fundamentally distinct from the LoRA series works. While LoRA uses different low-rank weights to fine-tune different pretrained weights, LORS adds different low-rank weights per layer to the same common weights shared by all layers, thus creating parameter-efficient stacked models trained from scratch, requiring no pretrained weights.

3. Approach

In this section, we initially revisit the foundational works upon which our methodology and experiments are primarily based, encompassing the LoRA mechanism [21], query-based object detection [3], and the structure of AdaMixer’s decoders [12]. Then we mathematically formulate our proposed LORS, which comprises two parts: the adaptive one and the static one. Subsequently, we elaborate on how the LORS method is applied to reduce the parameters of AdaMixer’s decoders. Lastly, we provide a quantitative analysis of the parameter reduction percentage attained by our proposed method.

3.1. Preliminary

The mechanism of LoRA. The Low-Rank Adaptation (LoRA) [21] technique is a novel approach designed to adapt large pre-trained language models for specific tasks. The key idea of LoRA is to introduce a low-rank parameter matrix which is able to capture task-specific knowledge while the original pre-trained parameters remain fixed.

Mathematically, given a pre-trained parameter matrix $W \in \mathbb{R}^{d \times h}$, LoRA uses a low-rank matrix $B \in \mathbb{R}^{d \times r}$ and a projection matrix $A \in \mathbb{R}^{r \times h}$ to adapt W , where $r \ll d, h$. The adapted parameter matrix is then given by

$$W + \Delta W = W + BA \quad (1)$$

where BA captures the task-specific knowledge.

The key advantage of LoRA is that it can significantly reduce the number of parameters that need to be fine-tuned, thereby reducing the computational cost and lowering down the memory requirement. In some cases, even single-digit values of r are sufficient to fine-tune the model to the desired state, which is often tens of times less expensive than

training the parameters in W . Furthermore, by keeping the original parameters fixed, LoRA avoids catastrophic forgetting, a common issue in fine-tuning large models.

Query-based object detection. In the realm of object detection, query-based detectors have established a new paradigm [11, 12, 30, 41, 51, 56]. Unlike traditional detectors which rely on anchor boxes or sliding windows, query-based models utilize a set of learnable queries to interact with image feature maps. This interaction can be formalized with attention [46] operations as

$$Q_{\text{updated}} = \text{Attention}(Q, K(V)) \quad (2)$$

where Q , K , and V represent the queries, keys, and values. The learnable queries Q are used to predict object classes and bounding boxes ultimately, while K and V typically originate from encoded image features. It is a common practice to continuously refine Q through interactions with K and V using successive decoding layers. These layers are usually composed of structurally identical decoders.

Decoders of AdaMixer. AdaMixer [12] is a query-based detector that features an innovative decoder design. This design includes Adaptive Channel Mixing (ACM) and Adaptive Spatial Mixing (ASM) methods, which greatly enhance its performance [12].

Given a sampled feature $\mathbf{x} \in \mathbb{R}^{P_{\text{in}} \times C}$, where $C = d_{\text{feat}}/g$, and g denotes the number of sampling groups. This sampled feature is obtained through an operation called group sampling. This operation divides the feature space channel d_{feat} into g groups and performs individual sampling for each group. Then, ACM (Adaptive Channel Mixing) utilizes a weight adapted by the object query \mathbf{q} to transform feature \mathbf{x} in the channel dimension, enhancing channel semantics [12]:

$$M_c = \text{Linear}(\mathbf{q}) \in \mathbb{R}^{C \times C} \quad (3)$$

$$\text{ACM}(\mathbf{x}) = \text{ReLU}(\text{LayerNorm}(\mathbf{x}M_c)) \quad (4)$$

where LayerNorm stands for Layer Normalization [1].

Next is the ASM (Adaptive Spatial Mixing) process, which aims to enable the adaptability of the object query \mathbf{q} to spatial structures of sampled features [12] by applying adaptive transformation to the spatial dimension:

$$M_s = \text{Linear}(\mathbf{q}) \in \mathbb{R}^{P_{\text{in}} \times P_{\text{out}}} \quad (5)$$

$$\text{ASM}(\mathbf{x}) = \text{ReLU}(\text{LayerNorm}(\mathbf{x}^T M_s)), \quad (6)$$

Both ACM and ASM train independent parameters for each sampling group, and finally the output with the shape $\mathbb{R}^{g \times C \times P_{\text{out}}}$ is flattened and transformed to the d_q dimension by a linear layer L_{output} to add back to the object query.

ACM, ASM and the output linear transformation L_{output} possess significantly more parameters compared to the decoder’s other operations, making them the main contributors to the number of parameters. Hence, we choose them

as the target components for validating the effectiveness of our LORS method in parameter reduction.

```
def non_dynamic_fused_weight(W_shared, A, B):
    # W_shared: (1, d, h)
    # A: (G, r, h)
    # B: (G, d, r)

    # private weight: (G, d, h)
    W_private = B @ A

    # sum at group dimension to get fused weight: (d, h)
    W = concat([W_shared, W_private], dim=0).sum(axis=0)

    return W
```

Figure 2. Pseudo-code for obtaining a static weight parameter for one layer.

3.2. Formulation of Our Method

The complete computational process of LORS is illustrated in Figure 1. Notably, LORS computation is divided into two types: adaptive and static. In the context of AdaMixer, the term "adaptive" indicates whether the transformation matrix depends on the object query. After completing the formalization of LORS here, we will further elaborate on how it is applied to the tasks in our setting.

Static Low Rank Residual Structure (LORS^T). Suppose there are N stacked layers of modules with identical architecture, and $W_i \in \mathbb{R}^{d \times h}$ be a parameter matrix belonging to the i -th layer. Then, we have:

$$W_i = W^{\text{shared}} + W_i^{\text{private}} \quad (7)$$

Here, $W^{\text{shared}} \in \mathbb{R}^{d \times h}$ represents the shared parameters across all stacked layers, while W_i^{private} denotes the layer-specific parameters for the i -th layer, which is calculated as follows:

$$W_i^{\text{private}} = \sum_{k=1}^K B_{ik} A_{ik} \quad (8)$$

$B_{ik} \in \mathbb{R}^{d \times r}$, $A_{ik} \in \mathbb{R}^{r \times h}$, and $r \ll d, h$. K represents the number of parameter groups used to compute W_i^{private} . Pseudo-code of LORS^T computing W_i^{private} for a specific layer i can be seen in Figure 2.

Adaptive Low Rank Residual Structure (LORS^A). Let $\hat{W}_i \in \mathbb{R}^{d \times h}$ be an adaptive generated parameter in i -th stacked layer, it is similarly calculated as :

$$\hat{W}_i = \hat{W}^{\text{shared}} + \hat{W}_i^{\text{private}} \quad (9)$$

where the cross-layer shared parameter $\hat{W}^{\text{shared}} \in \mathbb{R}^{d \times h}$ and the layer-specific parameter $\hat{W}_i^{\text{private}} \in \mathbb{R}^{d \times h}$ for each layer

```
def dynamic_fused_weight(q, A, B):
    # q: (N, C)
    # A: (1, G, r, h)
    # B: (1, G, d, r)

    # shared weight: (N, 1, d, h)
    W_shared = linear1(pro_feats).reshape(N, 1, d, h)
    # E: (N, G, r, r)
    E = linear2(pro_feats).reshape(N, G, r, r)
    # private weight: (N, G, d, h)
    W_private = B @ E @ A

    # sum at group dimension to get fused weight: (N, d, h)
    W = concat([W_shared, W_private], dim=1).sum(axis=1)

    return W
```

Figure 3. Pseudo-code for obtaining an adaptive weight parameter for one layer.

are both calculated based on q :

$$\hat{W}^{\text{shared}} = \text{Linear}(q) \in \mathbb{R}^{d \times h} \quad (10)$$

$$\hat{W}_i^{\text{private}} = \sum_{k=1}^K \hat{B}_{ik} \hat{E}_{ik} \hat{A}_{ik} \quad (11)$$

$$\hat{E}_{ik} = \text{Linear}(q) \in \mathbb{R}^{r \times r} \quad (12)$$

where $\hat{B}_{ik} \in \mathbb{R}^{d \times r}$ and $\hat{A}_{ik} \in \mathbb{R}^{r \times h}$, $r \ll d, h$. Pseudo-code of LORS^A calculating $\hat{W}_i^{\text{private}}$ for a specific layer i can be seen in Figure 3.

3.3. Applying LORS to AdaMixer's Decoders

We apply LORS to the parameters of linear transformations belonging to ACM, ASM, and L_{output} in each of AdaMixer's decoders. These parameters are explained in 3.1.

The overall pipeline of LORS running in AdaMixer is illustrated in Figure 4. Specifically, for each group of sampling points, LORS^A (see 3.2) is used to reduce parameters in M_c (from \mathbb{R}^{d_q} to $\mathbb{R}^{C \times C}$) of ACM and M_s of ASM (from \mathbb{R}^{d_q} to $\mathbb{R}^{P_{\text{in}} \times P_{\text{out}}}$), and LORS^T (see 3.2) to minimize parameters in L_{output} (from $\mathbb{R}^{C \times P_{\text{out}}}$ to \mathbb{R}^{d_q}).

Thus, the parameter quantities of M_c , M_s , and L_{output} are $d_q \times C \times C$, $d_q \times P_{\text{in}} \times P_{\text{out}}$ and $d_q \times C \times P_{\text{out}}$ respectively. When the sampling strategy consists of 2 groups with 64 points each, which is the default setting in our experiments, the values of the variables are as follows: $d_q = 256$, $C = 64$, $P_{\text{in}} = 64$, and $P_{\text{out}} = 128$. It can be easily calculated that the number of parameters for each of M_c , M_s , and L_{output} exceeds one million. Indeed, these three components collectively account for the most of the total parameter in the AdaMixer model with ResNet-50 as the backbone, whereas they are also the primary drivers of enhanced model performance. This is what motivates us to conduct LORS experiments on them. The effects of LORS applied to them will be discussed in detail in 3.4.

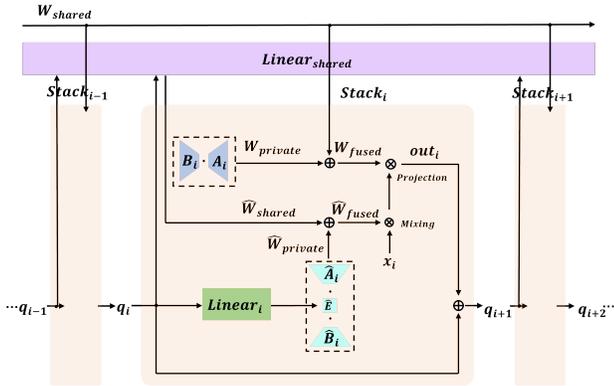


Figure 4. The overall pipeline of our proposed LORS, consisting of both adaptive and static parts, each further composed of shared and private components, works collaboratively. The figure illustrates the entire computation process within one layer of the stacked layers, with an enlarged example of the i -th layer for demonstration.

3.4. Analysis on Parameter Reduction

Let $W \in \mathbb{R}^{d \times h}$ be a weight parameter that exists in every layer of stacked structures, N is the number of stacked layers. If static, it originally has $d \times h$ parameters, while using LORS^T requires $\frac{1}{N} \times d \times h + K \times (d \times r + r \times h)$ parameters on average per layer; if it is adaptive, generating it by q with linear transformation requires $d_q \times d \times h$ parameters, where d_q is the dimension of q , and using LORS^A requires $\frac{1}{N} \times d_q \times d \times h + K \times (d_q \times r^2 + d \times r + r \times h)$ parameters on average per layer. To more intuitively display the parameter reduction effect of LORS, we set $d_q = 256$, $d = 64$, $h = 128$, $K = 2$ for the ASM process and $d = 2 \times 128 \times 128$, $h = 256$, $K = 1$ for the L_{output} , which is actually the case of AdaMixer’s decoders in our experiment, and show the parameter reduction for different r values using LORS^T and LORS^A in Table 1.

	r=4	r=8	r=16	r=32
W/. LORS ^T	1.53M	1.66M	1.93M	2.46M
W/O. LORS ^T	8.39M	8.39M	8.39M	8.39M
Percentage	18.3%	19.8%	23.0%	29.3%
W/. LORS ^A	0.36M	0.39M	0.49M	0.89M
W/O. LORS ^A	2.10M	2.10M	2.10M	2.10M
Percentage	17.2%	18.6%	23.3%	42.4%

Table 1. Analysis of the parameter reduction effect of LORS^T and LORS^A with varying rank r values. Our default selections are colored gray. Here the average amount of shared parameters over each layer has already been taken into account.

4. Experiments

In this section, we first explain the implementation details of our experiments, including the dataset, training specifics, and different parameter initialization methods for static and adaptive LORS. Next, we present the main experimental results. Finally, we perform ablation studies to examine LORS’s design and its individual components.

4.1. Implementation Details

Dataset. Our experiments were conducted on the widely-used MS COCO [29] dataset in mmdetection codebase [5], using the standard metrics for object detection. All models were trained on the *train2017* split ($\sim 118k$ images) and then evaluated on the *val2017* split (5k images).

Training. We use the AdamW optimizer [33] with a weight decay of 0.0001 and train all models on 8 Nvidia V100 GPUs with a batch size of 16 and a learning rate of $2.5 \times e^{-5}$. Models are trained for either 12 or 36 epochs, with learning rate decreasing by a factor of 10 at epochs 8 and 11 for 12-epoch training, and at epochs 24 and 33 for 36-epoch training. The low-rank values are set as $r = 16$ for LORS^A and $r = 8$ for LORS^T. The number of parameter groups is set as $K = [1, 1, 2, 2, 3, 3]$ for LORS^A in all experiments, applied to ACM and ASM in AdaMixer’s decoders, and $K = [1, 1, 1, 1, 1, 1]$ for L_{output} in all experiments. We divide the feature channels into 2 groups with 64 sampling points each instead of AdaMixer’s default 4 groups with 32 sampling points each, aiming to increase the parameter compressible space for LORS, which does not improve performance according to both AdaMixer paper [12] and our experiments. Backbones are initialized with ImageNet-1k [9] pre-trained models, and LORS-related parameter initialization is detailed below, remaining parameters are Xavier-initialized [14]. Finally, all other aspects about model training, like the data augmentation, the loss function, etc., just follow AdaMixer’s settings [12].

Initialization Strategies We tried various initialization methods for each component in LORS and determined the overall initialization method as follows:

- **LORS^T:** For static LORS, we employ Kaiming initialization [17] for W^{shared} and each B , and zero initialization for each A .
- **LORS^A:** For adaptive LORS, we apply Kaiming initialization [17] to the linear transformation weights forming each \hat{W}^{shared} , as well as each \hat{B} and \hat{A} . Additionally, we use zero initialization for the linear transformation weights forming each \hat{E} .

4.2. Main Results

To reliably demonstrate that LORS can reduce the number of parameters in stacked structures while maintaining model performance, which is measured by AP metric, we follow

Method	Queries	Epochs	Decoder Params(M)	Params(M)	Decoder GFLOPs	GFLOPs	Training Hours	AP	AP_{50}	AP_{75}	AP_s	AP_m	AP_l
AdaMixer	100	12	110	135	12	104	8.5h	42.7	61.5	45.9	24.7	45.4	59.2
AdaMixer + LORS	100	12	35	60	18	110	9.5h	42.6	61.4	46.0	25.0	45.6	58.5
AdaMixer	300	12	113	139	39	132	10h	44.1	63.4	47.4	27.0	46.9	59.5
AdaMixer + LORS	300	12	35	60	56	149	12h	44.1	63.0	47.7	27.8	47.0	59.5

Table 2. **1× training scheme** performance on COCO 2017 val set with ResNet-50 as backbone. AdaMixer with LORS can achieve competitive results while employing a notably reduced number of parameters. FPS is obtained with a single Nvidia V100 GPU.

Method	backbone	Queries	Epochs	Params(M)	GFLOPs	AP	AP_{50}	AP_{75}	AP_s	AP_m	AP_l
AdaMixer	R-50	100	36	135	104	43.2	61.8	46.7	25.0	46.1	58.8
AdaMixer + LORS	R-50	100	36	60	110	43.7	62.3	47.3	25.5	46.4	60.0
AdaMixer	R-50	300	36	139	132	47.0	66.0	51.1	30.1	50.2	61.8
AdaMixer + LORS	R-50	300	36	60	149	47.6	66.6	52.0	31.1	50.2	62.5
AdaMixer	R-101	300	36	158	208	48.0	67.0	52.4	30.0	51.2	63.7
AdaMixer + LORS	R-101	300	36	79	225	48.2	67.5	52.6	31.7	51.3	63.8
AdaMixer	Swin-S	300	36	164	234	51.3	71.2	55.7	34.2	54.6	67.3
AdaMixer + LORS	Swin-S	300	36	85	250	51.8	71.6	56.4	35.4	55.0	68.4

Table 3. **3× training scheme** performance on COCO 2017 val set, considering different combinations of backbone and query numbers. Longer training time allows LORS to be fully trained and perform better. The comprehensive improvement in performance and the significant reduction in parameters demonstrate the effectiveness of LORS. The best results in each category are highlighted in bold.

common practices in the object detection field and present experimental comparisons in two categories based on training schemes. The first involves training for 12 epochs using a weaker data augmentation, and the second involves training for 36 epochs using a stronger data augmentation, both of which are identical to those in AdaMixer [12].

Table 2 presents the comparison of performance with and without the LORS technique, under a 1× training scheme. Metrics such as the number of parameters, GFLOPs, and Average Precision (AP) at various scales are evaluated. When the AdaMixer model is combined with LORS, it demonstrates a remarkable reduction in the number of parameters. This reduction is consistently observed across various query quantities and different training epochs, encompassing both the decoder and overall model. For example, when trained with 100 queries for 12 epochs, the AdaMixer+LORS model uses only 35M decoder parameters and 60M total parameters, compared to 110M and 135M respectively in the AdaMixer model. This shows reduction of approximately 70% in the decoder’s parameters, which is substantial. However, the AdaMixer+LORS model shows a slight increase in GFLOPs. In terms of model performance, the AdaMixer+LORS model achieves competitive results, even with a significantly reduced number of parameters. When trained with 100 queries, although the AP of 42.6 is slightly lower than the 42.7 AP of the vanilla AdaMixer, it still slightly outperforms in AP_s and AP_m . When trained with 300 queries, the use of LORS does not

affect the performance and allows an advantage in all AP_s , AP_m and AP_l . These findings suggest that even under conditions of limited model training and data augmentation, LORS does not impact training result, the model can still converge quickly to its original limit. Additionally, we also demonstrate the slight impact of the current LORS method on inference speed, which might be improved by reducing the serial and redundant computations within LORS but not the focus of this paper.

Table 3 showcases the remarkable performance of the AdaMixer + LORS method under the 3× training scheme with different backbones and query numbers. It can be observed that the proposed method consistently outperforms the vanilla AdaMixer across all backbones, query numbers, and evaluation metrics. This result is somewhat surprising to us, as the LORS enables the model to use significantly fewer parameters during both training and inference.

Specifically, when adopting the ResNet-50 backbone and 100 queries, LORS improves the AP value of AdaMixer by 0.5 points (43.7 vs 43.2). It is worth noting that in Table 2, for experiments with the same configuration except for training 12 epochs, the AP with LORS is 0.1 lower than without LORS. However, this situation is reversed when training for 36 epochs. The longer training time under the 3× training scheme allows LORS to be fully trained and perform better. This emphasizes the importance of sufficient training time for the LORS technique to fully exploit its potential in enhancing the performance. Furthermore,

when using more powerful backbones and a larger number of queries, the performance metrics of models utilizing LORS consistently and comprehensively surpass those of their counterparts without LORS. Even when employing a strong backbone such as Swin-S [31] and using 300 queries, LORS still manages to further improve the model’s performance while reducing nearly half of the overall parameters. Much fewer parameters, yet improved performance.

In conclusion, the results in Table 2 and Table 3 show the effectiveness of the LORS method in reducing the number of parameters while achieving even better performance. The method’s adaptability to various backbones and query numbers further demonstrates its versatility and potential for broader applications. As for why LORS can achieve better performance with fewer parameters, a plausible explanation is that after the model has been sufficiently trained and stabilized, the parameter matrix within the stacked structure is inherently sparse and low-rank. Without using LORS, the model needs to learn such a low-rank structure through training. In contrast, LORS directly and explicitly makes the structure sparse, which, to some extent, can be considered as adding a form of regularization to the model training, thereby facilitating the model’s learning.

4.3. Ablation Study

To save computational resources, all ablation studies are conducted using a ResNet-50 backbone and a 1× training scheme.

LORS ^A	LORS ^T	Decoder Params(M)	AP
		110	42.5 [†]
✓		79	42.6
	✓	66	42.6
✓	✓	35	42.6

Table 4. Effect of LORS^A and LORS^T. “†” denotes this AP result was reproduced by ourselves. Both LORS^T and LORS^A can reduce parameters without compromising performance.

Adaptive & Static LORS. We first conduct ablation studies on the impact of LORS^A and LORS^T on model parameters and performance, as shown in Table 4. In this table, a ✓ indicates that we use the LORS method to replace the corresponding component in the AdaMixer model mentioned earlier, while a blank cell means we use the original AdaMixer module. If neither of the two is present, the model is identical to the original AdaMixer. As can be seen from the table, both adaptive and static LORS can effectively reduce the parameter size of the decoder in the model (31M and 44M, respectively) without affecting the model performance. Note that the symbol † here represents that the 42.5 AP value is reproduced by us, which is slightly dif-

ferent from the 42.7 AP recorded in the original paper. We believe this discrepancy can be attributed to the randomness of 1x training and does not affect our conclusions.

Shared and private weights. The second point of interest is to determine which of the shared weight and private weight has a greater impact on the decoder’s performance. Experiment in Table 5 is designed to address this question, which is performed on ACM and ASM of AdaMixer’s decoder. Recall that the formula we use to generate the final weight parameters is $W = W^{\text{shared}} + W^{\text{private}}$. In the table, a ✓ indicates that the corresponding term is present on the right side of the formula, otherwise, it is removed. It can be seen that the absence of either term will cause a nonnegligible decrease in model performance, and the impact of missing W^{private} is even more severe. This suggests that although the number of parameters in each W^{private} is smaller than that of W^{shared} , it could be more important, possibly because the most informative part of each layer in the stacked structure is contained within it.

W^{shared}	W^{private}	Decoder Params(M)	AP
✓		58	40.9
	✓	43	41.4
✓	✓	35	42.6

Table 5. Effect of shared parameters W^{shared} and private parameters W^{private} . Default choice for our model is colored gray.

Hyperparameters of LORS^A. Another aspect we would like to explore is the optimal number of parameter groups and the value of rank r for adaptive LORS. Here, we feel it is necessary to explain why we use grouped parameters. Recalling formula 11, the adaptive private parameters are obtained by continuous multiplication of three matrices, so the matrix with the smallest rank among them determines the rank of multiplication, which is the upper limit for the private parameters. If the task-specific requirement for the rank of the private parameters matrix unfortunately exceeds the rank of the smallest component, there will be a bottleneck when training the model. Therefore, we designed this parameter group mechanism, using the sum of the results of multiple matrix multiplication to break the aforementioned problem. Another advantage is that the rank of the final matrix possibly increases linearly with the number of parameters, while simply using a larger r would lead to a quadratic increase in the number of parameters. In the end, looking at Table 6, we observed that the experimental result for the 6-layer decoder of the AdaMixer object detector is relatively better when using [1,1,2,2,3,3] groups of parameters per layer and $r = 16$ for each group. This slightly strange parameter group numbers is inspired by AdaLoRA [53]. It suggests that higher-level neural network parameters may

need to allocate more ranks when performing LoRA fine-tuning, which we suppose may also hold true for LORS, and the experiment results confirmed that.

Parameter Groups	Rank r	Decoder Params(M)	AP
1 1 1 1 1 1	8	32	41.8
1 1 1 1 1 1	16	34	42.1
2 2 2 2 2 2	16	35	42.4
1 1 2 2 3 3	16	35	42.6
1 1 1 1 1 1	32	38	41.2

Table 6. Validation accuracy with different groups of parameters and rank r in LORS^A. Default choice for our model is colored gray. Each row’s first six numbers indicate the number of parameter groups sequentially used in each of the six decoder layers.

Hyperparameters of LORS^T. We further investigate the optimal configuration settings for LORS^T to achieve satisfactory performance, as shown in Table 7. We observed that the best option is to use one group of low-rank parameters with $r = 8$ for the decomposition of each layer L_{output} in the stacked decoder of AdaMixer, which can achieve a model performance comparable to the original detector. Since using one group of parameters already yields reasonably good results, we did not explore the use of more parameter groups. Moreover, we should also note that the value of rank r is not necessarily better when larger or smaller, which may indicate that the low-rank structure required by the model for a specific task has its own particularity.

Parameter Groups	Rank r	Decoder Params(M)	AP
1 1 1 1 1 1	2	34	42.3
1 1 1 1 1 1	4	34	41.9
1 1 1 1 1 1	8	35	42.6
1 1 1 1 1 1	16	37	42.1
1 1 1 1 1 1	32	40	41.7

Table 7. Validation accuracy with different rank r in LORS^T. Default choice for our model is colored gray. Each row’s first six numbers indicate the number of parameter groups sequentially used in each of the six decoder layers.

Number of decoders. Finally, we conclude our ablation studies with experiments on the optimal number of decoder layers incorporating the LORS structure for better performance. Since our method significantly reduces the total number of model parameters and allows each decoder layer to train and use shared parameters that capture common features, it is natural to consider whether the model’s inherent characteristics are affected and whether using more or fewer decoding layers would further improve the model’s performance. With this idea in mind, we conducted sev-

eral experiments as shown in Table 8 to explore this issue. However, similar to the original AdaMixer without LORS, 6 layers still yield the best performance for the AdaMixer with LORS, having too many or too few layers would degrade the model’s performance. This may also suggest that LORS merely captures the intrinsic structure of the model without altering its original properties.

Decoder Number	Decoder Params(M)	AP
3	28	38.9
6	35	42.6
9	42	42.0
12	50	41.0

Table 8. Model performance using different numbers of stacked decoders when employing LORS. The default choice for our model is colored gray. This experiment with varying numbers of decoders demonstrates that six layers still yield the best performance for the model, which is not changed by LORS.

5. Conclusion

This research introduced a novel approach to the reduction of parameters within deep learning models utilizing stacked structures, *i.e.*, the Low-rank Residual Structure (LORS). It is potentially an effective methodology in reducing both static and adaptive parameters. Broadly speaking, LORS allows the parameters of stacked modules to be largely shared, while maintaining only a small number of unique parameters for each module, thereby significantly reducing the total number of parameters whilst not compromising performance. We validated our method via object detection task across various extensive experimental cases on the widely-used COCO dataset, and the results demonstrates a superior model performance even with a 50% – 70% reduction in decoder parameters, surpassing our expectations.

There are some noticeable limitations in our study. While effective in reducing model parameters without compromising performance, enhancing model capabilities requires a comparatively long training process for the method. Additionally, our method slightly increases the inference time, as even without a detailed analysis, it can be roughly seen that LORS has unparalleled and repetitive computations needing more optimization. Finally, we have only tested our approach on object detection tasks, specifically with the AdaMixer model and its decoder structure. It’s clear that our method can be applied to more tasks (e.g., NLP), different models (e.g., language models), and other neural network components (e.g., backbones). These aspects will be the focus of our future work.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [3](#)
- [2] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. [2](#)
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. [2](#), [3](#)
- [4] Arnav Chavan, Zhuang Liu, Deepak Gupta, Eric Xing, and Zhiqiang Shen. One-for-all: Generalized lora for parameter-efficient fine-tuning, 2023. [3](#)
- [5] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. [5](#)
- [6] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multi-modal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. [2](#)
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. [2](#)
- [8] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems*, 28, 2015. [1](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [5](#)
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [2](#)
- [11] Peng Gao, Minghang Zheng, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fast convergence of detr with spatially modulated co-attention. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3621–3630, 2021. [3](#)
- [12] Ziteng Gao, Limin Wang, Bing Han, and Sheng Guo. Adamixer: A fast-converging query-based object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5364–5373, 2022. [2](#), [3](#), [5](#), [6](#)
- [13] Yanqi Ge, Qiang Nie, Ye Huang, Yong Liu, Chengjie Wang, Feng Zheng, Wen Li, and Lixin Duan. Beyond prototypes: Semantic anchor regularization for better representation learning. *arXiv preprint arXiv:2312.11872*, 2023. [2](#)
- [14] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010. [5](#)
- [15] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129:1789–1819, 2021. [1](#)
- [16] Song Han, Jeff Pool, John Tran, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *International Conference on Learning Representations*, 2016. [1](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. [5](#)
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#), [2](#)
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [2](#)
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [1](#)
- [21] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [2](#), [3](#)
- [22] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. [2](#)
- [23] Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki Markus Asano. Vera: Vector-based random matrix adaptation, 2023. [3](#)
- [24] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019. [1](#)
- [25] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. [2](#)
- [26] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019. [2](#)
- [27] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. [2](#)

- [28] Vladislav Lialin, Namrata Shivagunde, Sherin Muckatira, and Anna Rumshisky. Stack more layers differently: High-rank training through low-rank updates, 2023. [2](#)
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [5](#)
- [30] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. [2](#), [3](#)
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. [7](#)
- [32] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [2](#)
- [33] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [5](#)
- [34] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. [2](#)
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#)
- [36] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017. [2](#)
- [37] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [2](#)
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [2](#)
- [40] Farhana Sultana, Abu Sufian, and Paramartha Dutta. A review of object detection models based on convolutional neural network. *Intelligent computing: image processing based applications*, pages 1–16, 2020. [2](#)
- [41] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14454–14463, 2021. [2](#), [3](#)
- [42] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022. [2](#)
- [43] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. [2](#)
- [44] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems*, 34:24261–24272, 2021. [2](#)
- [45] Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan Kobzyev, and Ali Ghodsi. Dylora: Parameter efficient tuning of pre-trained models using dynamic search-free low-rank adaptation, 2023. [2](#)
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#), [3](#)
- [47] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *arXiv preprint arXiv:2305.11175*, 2023. [2](#)
- [48] BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022. [2](#)
- [49] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. [2](#)
- [50] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019. [2](#)
- [51] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. [3](#)
- [52] Longteng Zhang, Lin Zhang, Shaohuai Shi, Xiaowen Chu, and Bo Li. Lora-fa: Memory-efficient low-rank adaptation for large language models fine-tuning, 2023. [2](#)
- [53] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive budget allocation for parameter-efficient fine-tuning, 2023. [2](#), [7](#)
- [54] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [2](#)

- [55] Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv preprint arXiv:1710.01878*, 2017. [1](#)
- [56] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [2](#), [3](#)
- [57] Bojia Zi, Xianbiao Qi, Lingzhi Wang, Jianan Wang, Kam-Fai Wong, and Lei Zhang. Delta-lora: Fine-tuning high-rank parameters with the delta of low-rank matrices, 2023. [2](#)