This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

DeRS: Towards Extremely Efficient Upcycled Mixture-of-Experts Models

Yongqi Huang^{1*} Peng Ye^{2,3*} Chenyu Huang¹ Jianjian Cao¹ Lin Zhang¹ Baopu Li⁴ Gang Yu⁵ Tao Chen^{1†}

¹ Fudan University ² The Chinese University of Hong Kong

³ Shanghai AI Laboratory ⁴ Baidu USA ⁵ StepFun

yqhuang23@m.fudan.edu.cn, eetchen@fudan.edu.cn

Abstract

Upcycled Mixture-of-Experts (MoE) models have shown great potential in various tasks by converting the original Feed-Forward Network (FFN) layers in pre-trained dense models into MoE layers. However, these models still suffer from significant parameter inefficiency due to the introduction of multiple experts. In this work, we propose a novel DeRS (Decompose, Replace, and Synthesis) paradigm to overcome this shortcoming, which is motivated by our observations about the unique redundancy mechanisms of upcycled MoE experts. Specifically, DeRS decomposes the experts into one expert-shared base weight and multiple expert-specific delta weights, and subsequently represents these delta weights in lightweight forms. Our proposed DeRS paradigm can be applied to enhance parameter efficiency in two different scenarios, including: 1) DeRS Compression for inference stage, using sparsification or quantization to compress vanilla upcycled MoE models; and 2) DeRS Upcycling for training stage, employing lightweight sparse or low-rank matrixes to efficiently upcycle dense models into MoE models. Extensive experiments across three different tasks show that the proposed methods can achieve extreme parameter efficiency while maintaining the performance for both training and compression of upcycled MoE models.

1. Introduction

Recently, sparse Mixture-of-Experts [44] (MoE) models have made remarkable progress [1, 8, 19, 23, 31, 49]. These models introduce the MoE layer, which consists of multiple Feed-Forward Network (FFN) experts and a learnable router, to replace a single FFN layer. By dynamically activating only a subset of experts for each input, these MoE models achieve superior performance while maintaining computational efficiency. However, training MoE models from scratch requires substantial computational resources [8, 9, 43] and often en-



Figure 1. Visualization of cosine similarity in the first and last MoE layers of the MoE-LLaVA-Phi [30] model. *FFN* denotes the initial weight while E_i denotes the trained weight of the *i*-th expert.

counters training instabilities [41, 53].

To address these challenges, upcycling pre-trained dense models into MoE models (referred to as upcycling) [21] has emerged as an effective alternative. This approach converts the original FFN layers in a pre-trained dense model into MoE layers with N experts, initializing each expert from the original FFN weight while introducing a randomly initialized router for expert selection. By leveraging the knowledge embedded in pre-trained dense models, upcycled MoE models facilitate more efficient optimization and demonstrate competitive performance under constrained training budgets [14, 30, 47]. Due to its simplicity and effectiveness, upcycling has been widely applied across diverse domains, including natural language processing [6, 14, 21], computer vision [13, 21, 24], and multi-modal learning [20, 27, 30].

While upcycled MoE models have shown significant success, they still suffer from parameter inefficiency due to the large number of extra parameters introduced by MoE experts. For instance, the upcycled MoE-LLaVA-Phi [30] model comprises a total of 5 billion parameters, of which 3.4 billion are occupied by MoE experts. In this paper, we focus on enhancing the parameter efficiency of upcycled MoE models. We first revisit the characteristics of the N experts in upcycled MoE models: 1) We note that N experts share the same initial weight W_{base} , and the weights of trained experts can be uniformly expressed as $W_i = W_{base} + \Delta_i$, where i = 1, ..., N; 2) We observe that, there is an extremely high

[†]Corresponding author. ^{*}Equal contribution.

cosine similarity (higher than 0.999) between the expert's weight and the initial weight, as well as among the weights of different experts, as shown in Fig. 1. This extremely high cosine similarity suggests that Δ_i is a minor adjustment to W_{base} and may exhibit considerable redundancy.

Motivated by these observations, we propose a novel Decompose, Replace and Synthesis (DeRS) paradigm that remodels the weights of upcycled MoE experts $\{W_1,\ldots,W_N\}$ into one expert-shared base weight W_{base} and multiple expert-specific delta weights $\{\Delta_1, \ldots, \Delta_N\}$, enhancing the expert parameter efficiency by applying lightweight representations to these delta weights. Building on the DeRS paradigm, we further design two distinct approaches for upcycled MoE models at different phases: DeRS compression for inference and DeRS upcycling for training. Specifically, DeRS compression employs sparsification or quantization techniques to compress the delta weights of already-trained experts, thereby achieving inference-time parameter efficiency for vanilla upcycled MoE models. On the other hand, DeRS upcycling enables the training of experts in a parameter-efficient way by training only one expertshared FFN weight and multiple sparse or low-rank weights, thus greatly reducing trainable parameters.

We conduct comprehensive experiments to verify the effectiveness of the proposed DeRS compression and DeRS upcycling methods across three upcycled MoE baselines encompassing general multi-modal tasks, medical multi-modal tasks and code generation tasks, as well as six different MoE model architectures. Experimental results demonstrate that our proposed methods bring extreme efficiency to upcycled MoE models. For example, on the general multi-modal task, our DeRS compression effectively reduces the parameter count of a MoE layer by 65% without a performance drop, and our DeRS upcycling method achieves a reduction of up to 2270 times in additional parameters while delivering better performance compared to vanilla upcycling.

Our contributions are as follows:

- We are the first to explore the unique redundancy mechanisms of experts in upcycled MoE models, and propose a novel DeRS paradigm that decomposes multiple experts into one expert-shared weight and multiple expert-specific weights to reduce parameter redundancy.
- Based on our DeRS paradigm, we further propose two application methods to achieve extremely efficient upcycled MoE: 1) DeRS compression to efficiently compress vanilla upcycled MoE models for inference. 2) DeRS upcycling to efficiently upcycle a pre-trained dense model into a MoE model for training and deployment.
- Comprehensive experiments across three tasks and six MoE model architectures consistently verify the efficiency, effectiveness, and generalizability of the proposed DeRS compression and DeRS upcycling techniques.

2. Related Work

2.1. Training of MoE Models

For MoE models, the straightforward training strategy [7, 8, 23, 48] generally involves designing a MoE architecture, randomly initializing the MoE model weights, and training the model with extensive computational resources and data. Recently, upcycling [21] has been proposed to reduce training costs by initializing MoE models using pre-trained dense models. Due to its simplicity and effectiveness, upcycling has found widespread application across various domains [6, 20, 24, 30]. For instance, XFT [6] employs upcycling in the instruction tuning process of large language models, achieving performance improvement on code generation tasks. In the field of multi-modal learning, MoE-LLaVA [30] successfully trains a vision-language MoE model with only 3B active parameters through upcycling, achieving performance comparable to LLaVA-1.5-7B [35] on various visual understanding benchmarks. In this paper, we first find the unique redundancy of expert parameters in vanilla upcycling, and further propose the DeRS upcycling method, which significantly reduces the parameter count of experts in upcycled MoE models during both training and inference.

2.2. Compression of MoE Models

To date, many studies have explored the compression of experts in MoE models to improve deployment efficiency [15, 28, 33, 39]. Specifically, [39] proposes a plugand-play expert-level sparsification technique by pruning unimportant experts and dynamically skipping specific experts during inference. EEP [33] employs an evolutionary strategy to search the pattern for expert removal and consolidates the knowledge of removed experts into retained experts through weight merging. MC-SMoE [28] first merges infrequently activated experts into those often activated experts, and then compresses the merged experts. In this paper, we also propose a novel expert compression method called DeRS compression, which focuses on compressing MoE models upcycled from pre-trained dense models. Leveraging the characteristics of upcycled MoE models, our DeRS compression innovatively decomposes the weights of multiple experts into an expert-shared important weight and multiple expert-specific redundant weights, and applies sparsification or quantization to those redundant weights.

3. Method

3.1. Preliminary

Mixture-of-Experts. A standard MoE layer consists of a set of N experts (i.e., N FFN layers) $\{E_1, E_2, \ldots, E_N\}$ and a router network R with trainable weight W_R . Given an input x, the output y of the MoE layer can be written as:



Figure 2. Overall procedure of (a) upcycling a dense model into a MoE model through vanilla upcycling and (b) compressing the vanilla upcycled MoE model using the proposed DeRS compression, which first decomposes the trained experts and then applies sparsification or quantization techniques to the expert-specific delta weights. During inference, when an expert is needed, we synthesize its weight online.

$$y = \sum_{i=1}^{N} R(x)_i \cdot E_i(x)$$
(1)
$$R(x) = TopK(softmax(x \cdot W_R), k)$$

where $R(x)_i$ denotes the routing score to the *i*-th expert, $E_i(x)$ stands for the output of *i*-th expert, $TopK(\cdot, k)$ means selecting only the top-*k* experts by setting R(x) of other experts to 0. Usually, $k \ll N$, which means R(x) a sparse *N*-dimensional vector. When $R(x)_i = 0$, $E_i(x)$ does not need to be computed.

Upcycled Mixture-of-Experts. Let $W_{base} \in \mathbb{R}^{d \times d_h}$ denote the weight of original FFN layer in the pre-trained dense model. For the corresponding upcycled MoE layer with N experts, the weight of each expert is uniformly initialized to W_{base} instead of random initialization. After training, the weights of these N experts are updated to $\{W_1, \ldots, W_N\}$.

3.2. Decompose, Replace, and Synthesis

In this paper, we revisit the Upcycled Mixture-of-Experts and remodel it according to the following three aspects: Decompose, Replace, and Synthesize.

Decompose. Since all N experts share the same initial weight W_{base} , the trained weight of a specific expert E_i can be expressed as:

$$W_i = W_{base} + \Delta_i \tag{2}$$

This implies that $\{W_1, \ldots, W_N\}$ can be decomposed into an expert-shared base weight W_{base} and N expert-specific delta weights $\{\Delta_1, \ldots, \Delta_N\}$.

Replace. Since W_{base} is pre-trained knowledge learned from extensive data, the expert-specific delta weight Δ_i is actually minor adjustment to W_{base} . As shown in Fig. 1, the extremely high cosine similarity among the trained experts and the original FFN further supports this opinion, suggesting that these N expert-specific delta weights $\{\Delta_1, \ldots, \Delta_N\}$ may exhibit some degree of redundancy. Based on this hypothesis, it is feasible to adopt a lightweight representation $\mathcal{F}(\Delta_i)$ to reformulate the original redundant Δ_i without a performance drop. In other words, we can replace the original heavyweight set $\{W_1, \ldots, W_N\}$ with a more parameter-efficient set $\{\mathcal{F}(\Delta_1), \ldots, \mathcal{F}(\Delta_N)\} \cup \{W_{base}\}$.

Synthesis. Based on $\{\mathcal{F}(\Delta_1), \ldots, \mathcal{F}(\Delta_N)\}$ and W_{base} , we synthesize the weight of a specific expert E_i by fusing the lightweight expert-specific delta weight $\mathcal{F}(\Delta_i)$ with the expert-shared base weight W_{base} as follows:

$$\hat{W}_i = W_{base} + \mathcal{F}(\Delta_i) \tag{3}$$

Based on the above framework of *Decompose, Replace* and Synthesis (DeRS), we propose two application methods: DeRS compression and DeRS upcycling. DeRS compression focuses on compressing already-trained vanilla upcycled MoE models, while DeRS upcycling is designed to efficiently upcycle existing dense models into MoE architectures for subsequent training and deployment.

3.3. DeRS Compression

As shown in Fig. 2, applying DeRS compression to compress an already-trained vanilla upcycled MoE model follows three steps: 1) decomposing the weights of N trained experts into W_{base} and $\{\Delta_i\}_{i=1}^N$; 2) utilizing a post-training lightweight technique \mathcal{F}_{post} to compress $\{\Delta_i\}_{i=1}^N$ into more efficient $\{\mathcal{F}_{post}(\Delta_i)\}_{i=1}^N$; and 3) when a specific expert E_i is needed,



Figure 3. Comparisons between vanilla upcycling and the proposed DeRS upcycling in the construction of experts. Instead of making N copies of the original FFN, our DeRS upcycling synthesizes experts by combining the shared FFN with expert-specific lightweight parameters (i.e., in the form of sparse matrixes or low-rank matrixes).

synthesizing its weight by summing $\mathcal{F}_{post}(\Delta_i)$ and W_{base} . Two different compression designs are proposed as follows:

Sparsification. A straightforward implementation is to remove unnecessary elements from the weight matrix. Inspired by [50], we randomly drop most elements of the expert-specific delta weight Δ_i and compactly store the obtained sparse matrix $\mathcal{F}_{post}(\Delta_i)$ as vectors. Given a drop rate p, $\mathcal{F}_{post}(\Delta_i)$ can be written as:

$$M_i \sim \text{Bernoulli}(p)$$

$$\mathcal{F}_{post}(\Delta_i) = (\mathbf{1} - M_i) \odot \Delta_i / (1 - p)$$
(4)

Although the theoretical shape of a sparse weight matrix remains $\mathbb{R}^{d \times d_h}$, the actual parameter count of $\mathcal{F}_{post}(\Delta_i)$ is only $d \cdot d_h \cdot (1-p)$ due to its compact storage format. In other words, by applying DeRS-Sparsification with the given drop rate p, the parameter count of N trained experts is reduced from the original $N \cdot d \cdot d_h$ to $(1 + N \cdot (1-p)) \cdot d \cdot d_h$.

Quantization. Additionally, we provide another approach to reduce redundancy by quantizing Δ_i to a lower k-bit representation:

$$\mathcal{F}_{post}(\Delta_i) = Quant(\Delta_i, k) \tag{5}$$

Assuming the original expert weights $\{W_i\}_{i=1}^N$ are represented with K bits, applying DeRS-Quantization can reduce the parameter storage cost from $N \cdot K$ to $K + N \cdot k$.

3.4. DeRS Upcycling

The difference between our proposed DeRS upcycling and vanilla upcycling is presented in Fig. 3. Instead of replicating the original FFN layer N times to form N experts of a MoE layer, we decompose the construction of N experts before training. Specifically, we decompose N expert weights into one trainable expert-shared weight W_{shared} and N trainable expert-specific incremental weights $\{\mathcal{F}_{pre}(\Delta_i)\}_{i=1}^N$, using a parameter-efficient design \mathcal{F}_{pre} . During both training and inference, the weight W_i of expert E_i is synthesized by combining W_{shared} and $\mathcal{F}_{pre}(\Delta_i)$ via addition. Here, W_{shared} is initialized from the weight of the original FFN layer and $\mathcal{F}_{pre}(\Delta_i)$ is zero-initialized. We provide two types of parameter-efficient designs \mathcal{F}_{pre} as follows:

Sparse Matrix. Since sparsification can be used to reduce redundancy in our DeRS compression, it's feasible to adopt sparse matrixes as a parameter-efficient form of the trainable expert-specific incremental weights. However, simply employing a vanilla sparse matrix implementation, where the matrix maintains the shape of $\mathbb{R}^{d \times d_h}$ and most of the element values are zero, can't effectively reduce the number of parameters in practice.

To address this, we reformulate the sparse matrix with a shape of $\mathbb{R}^{d \times d_h}$ and a sparse rate p into two compact row vectors of length $d \cdot d_h \cdot (1-p)$, specifically an index vector I and a value vector V. As shown in Fig. 3(b), the value vector V stores the values of all nonzero elements, while the index vector I stores the positions of these nonzero elements in the original-shaped sparse matrix. These two vectors I and V are mapped back to the sparse matrix using the torch.scatter function. Based on the above efficient implementation of the sparse matrix, we propose the first design of expert-specific incremental weights:

$$\mathcal{F}_{pre}(\Delta_i) = \text{torch.scatter}(I_i, V_i)$$
 (6)

Given a sparse rate p, the index vector I_i is generated by randomly selecting $d \cdot d_h \cdot (1 - p)$ unique values from the range $[0, d \cdot d_h)$ and is fixed thereafter, while V_i serves as the trainable parameter and is initialized to zero.

In terms of the number of trainable expert parameters, vanilla upcycling requires training N expert weights of shape $\mathbb{R}^{d \times d_h}$, while the proposed Sparse-Matrix-based DeRS (DeRS-SM) upcycling only necessitates training one W_{shared} of shape $\mathbb{R}^{d \times d_h}$ along with N row vectors of length $d \cdot d_h \cdot (1-p)$. In other words, our DeRS-SM upcycling decreases the trainable parameter count of experts from $N \cdot d \cdot d_h$ to $(1 + N \cdot (1-p)) \cdot d \cdot d_h$.

Low-rank Matrix. Since representing the expert-specific incremental weight using the high-dimensional space $\mathbb{R}^{d \times d_h}$ is redundant, it's sufficient to utilize two low-rank matrixes $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{r \times d_h}$ for an efficient representation based on the low-rank matrix factorization. Thus, we develop the Low-rank Matrix-based DeRS (DeRS-LM) upcycling . As shown in Fig. 3(c), the expert-specific incremental weight



Figure 4. Performance of applying different DeRS compression methods to compress three vanilla upcycled MoE-LLaVA models respectively.

Table 1. Performance comparison between vanilla upcycling and our DeRS upcycling on three MoE-LLaVA models on the general multi-modal task. DeRS-SM and DeRS-LM denote the Sparse-Matrix-based and Low-rank-Matrix-based DeRS upcycling respectively. Added Params represents the number of additional parameters of the upcycled MoE model compared to its corresponding dense model.

| MoF Model | Upcycling | Added | I | mage Q | uestion A | nswering | ç | Ben | chmark ' | Toolkit | Overall |
|-------------------------|----------------|---------|-------------------|--------|-----------|----------|---------|------|----------|---------|---------|
| WIOE WIOUEI | Method | Params. | VQA ^{v2} | GQA | VisWiz | SQAI | VQA^T | POPE | MMB | MM-Vet | Overall |
| | Vanilla | 1.24B | 76.3 | 60.6 | 34.6 | 62.7 | 50.2 | 86.9 | 60.5 | 27.1 | 57.4 |
| MoE-LLaVA-StableLM [30] | DeRS-SM (ours) | 0.26M | 76.5 | 60.7 | 34.9 | 62.9 | 50.2 | 87.3 | 60.4 | 28.4 | 57.7 |
| | DeRS-LM (ours) | 1.20M | 76.6 | 60.6 | 34.4 | 62.3 | 50.0 | 87.1 | 60.1 | 28.6 | 57.5 |
| | Vanilla | 1.22B | 76.2 | 61.2 | 31.8 | 62.4 | 48.1 | 87.5 | 60.4 | 24.4 | 56.5 |
| MoE-LLaVA-Qwen [30] | DeRS-SM (ours) | 0.26M | 76.2 | 61.2 | 31.0 | 62.2 | 47.8 | 87.8 | 60.0 | 23.7 | 56.2 |
| | DeRS-LM (ours) | 1.19M | 76.2 | 61.1 | 31.2 | 62.1 | 47.9 | 87.8 | 60.0 | 24.2 | 56.3 |
| | Vanilla | 2.52B | 77.5 | 61.4 | 42.7 | 68.6 | 50.8 | 86.9 | 66.0 | 32.4 | 60.8 |
| MoE-LLaVA-Phi [30] | DeRS-SM (ours) | 1.11M | 77.7 | 61.3 | 42.4 | 69.0 | 51.7 | 87.3 | 65.5 | 33.8 | 61.1 |
| | DeRS-LM (ours) | 2.42M | 77.6 | 61.3 | 42.0 | 69.1 | 51.8 | 87.4 | 66.3 | 32.6 | 61.0 |

can be expressed as $\mathcal{F}_{pre}(\Delta_i) = A_i \cdot B_i$, where A_i is randomly initialized and B_i is zero-initialized.

Given the rank $r, r \ll min(d, d_h)$, our DeRS-LM upcycling reduces the number of trainable expert parameters from $N \cdot d \cdot d_h$ to $d \cdot d_h + N \cdot r \cdot (d + d_h)$.

4. Experiments

In this section, we conduct a series of experiments to evaluate our DeRS compression and DeRS upcycling methods across a range of tasks, including general multi-modal tasks, medical multi-modal tasks, and code generation tasks.

4.1. General Multi-modal Task

Model architecture. We adopt the MoE-LLaVA [30] framework for the experiments of the general multi-modal task. We employ CLIP-Large [42] as the visual encoder. Three upcycled MoE models, respectively initialized from StableLM-2-1.6B [4], Qwen-1.8B [3] and Phi-2-2.7B [18], serve as the language backbone. Before upcycling, all the dense models have been previously fine-tuned on datasets collected from MIMIC-IT [25], LRV [34], SVIT [52] and LVIS [46]. In each dense model, every other block's FFN layer is upcycled into a MoE layer with 4 experts, where the top-2 experts are dynamically activated when processing each input.

Datasets. The upcycled MoE models undergo fine-tuning on the LLaVA-mix-665k [35] dataset, followed by evaluation across five image question-answering benchmarks: VQAv2 [10], GQA [17], VisWiz [12], ScienceQA-IMG [38] and TextVQA [45]. Besides, three benchmark toolkits, including POPE [29], MMBench [37] and MM-Vet [51], are adopted to evaluate the multi-modal understanding capabilities.

Results of DeRS Compression. The experimental results of DeRS compression are shown in Fig. 4. As we can see, the proposed two compression techniques, sparsification and quantization, can effectively reduce the parameter redundancy while not affecting the model's performance. Specifically, as shown in Fig. 4a, applying DeRS-Sparsification to remove 90% of the elements in the delta weights hardly affects the model's performance, but the parameter count of a MoE layer can be reduced from the original 4 experts'



Figure 5. Performance of applying different DeRS compression methods to compress two vanilla upcycled Med-MoE models respectively.

Table 2. Performance comparison between vanilla upcycling and our DeRS upcycling on two Med-MoE models on the medical multi-modal task. DeRS-SM and DeRS-LM denote the Sparse-Matrix-based and Low-rank-Matrix-based DeRS upcycling respectively. The light-gray Added Params denotes the additional parameters introduced by the universal FFN layers that are not considered as experts of MoE layers.

| MaE Madal | Upcycling | Added VQA-RAD | | SLAKE | | PathVQA | | Overall | |
|-----------------------|----------------|---------------|------|--------|------|---------|------|---------|---------|
| MOE MOUE | Method | Params. | Open | Closed | Open | Closed | Open | Closed | Overall |
| Med-MoE-StableLM [20] | Vanilla | 0.42B+1.24B | 51.0 | 82.3 | 82.4 | 85.3 | 33.4 | 91.4 | 71.0 |
| | DeRS-SM (ours) | 0.42B+0.29M | 49.5 | 84.2 | 83.8 | 84.9 | 33.1 | 91.6 | 71.2 |
| (EMINLP 24) | DeRS-LM (ours) | 0.42B+1.20M | 53.5 | 81.6 | 83.7 | 84.1 | 33.6 | 91.5 | 71.3 |
| Mad MaE Db: [20] | Vanilla | 0.84B+2.52B | 55.1 | 85.3 | 84.6 | 85.8 | 35.1 | 91.5 | 72.9 |
| (EMNI D 24) | DeRS-SM (ours) | 0.84B+1.11M | 55.5 | 85.3 | 84.3 | 86.3 | 35.0 | 91.6 | 73.0 |
| (EIVIINEF 24) | DeRS-LM (ours) | 0.84B+2.42M | 55.7 | 84.9 | 84.3 | 85.6 | 35.2 | 91.6 | 72.9 |

parameters to the equivalent of $(1 + 4 \times 0.1)$ experts' parameters. Similarly, we can quantize the expert-specific delta weights from the original 16 bits down to 2 bits using DeRS-Quantization, reducing the parameter storage cost of a MoE layer's 4 experts from original 4×16 to $16 + 4 \times 2$, while demonstrating the same or even better performance.

Moreover, we can observe that even under extreme DeRS compression settings, such as sparsification with a 0.99 drop rate or quantization with the 1-bit width, there is no significant degradation in the model's performance. This may be attributed to the fact that the three dense models used in the MoE-LLaVA framework have been previously fine-tuned on some multi-modal data. In other words, the expert-shared base weight obtained through decomposition during DeRS compression already has a certain degree of multi-modal understanding capability. As a result, even with the extreme compression of the redundant expert-specific delta weights, the model's performance remains robust.

Results of DeRS Upcycling. Tab. 1 shows the performance of different upcycling methods on three MoE-LLaVA architectures. As we can see, vanilla upcycling results in billions of additional parameters by duplicating the original FFN layer to construct experts, while our DeRS upcycling introduces only millions of extra parameters to achieve the same or superior performance. For example, when upcycling Phi-2-2.7B into the MoE-LLaVA-Phi architecture, vanilla upcycling adds 2.52 billion parameters. In comparison, our DeRS upcycling based on Sparse Matrix (DeRS-SM) and Low-rank Matrix (DeRS-LM) improve overall performance by 0.3% and 0.2% with just 1.11 and 2.42 million extra parameters (2270× and 1041× reduction), respectively.

These experimental results highlight the efficiency and effectiveness of our DeRS upcycling, which reduces the number of trainable expert parameters by sharing a base FFN across experts and employing multiple expert-specific lightweight parameters.

4.2. Medical Multi-Modal Task

Model architecture. We adopt the settings in Med-MoE [20] to conduct our medical multi-modal experiments. Similarly, the CLIP-Large model is used as the vision encoder. Two dense models, StableLM-2-1.6B and Phi-2-2.7B, which have been fine-tuned on the medical language-image instructionfollowing data of [26], are upcycled into MoE architectures to serve as the language backbones. Specifically, the original FFN layer of every other block is converted into a parallel structure, which consists of a universal FFN layer and a MoE layer with four experts. The MoE layer activates only the top-1 expert for a given input, while the universal FFN processes all inputs. The outputs from the MoE layer and the universal FFN are summed to produce the final output. Datasets. We use three medical image question-answering datasets with open- and closed-end question-answering pairs, including VQA-RAD [22], SLAKE [32], and PathVQA [16], to fine-tune and evaluate upcycled MoE models.

Results of DeRS Compression. Fig. 5 shows the performance of applying DeRS compression to two vanilla upcycled Med-MoE models. It can be observed that extreme compression of expert-specific delta weights (removing 99% of their elements or quantizing them to 1bit) has negligible impact on model performance. This may be attributed to two factors. First, the two dense models utilized within the Med-MoE framework have been previously fine-tuned on relevant

Table 3. Performance comparison between vanilla upcycling and our DeRS upcycling on the code generation task. DeRS-SM and DeRS-LM denote the Sparse-Matrix-based and Low-rank-Matrix-based DeRS upcycling respectively. The light-gray Added Params denotes the additional parameters introduced by the universal FFN layers that are not considered as experts of MoE layers.

| MoE Model | Upcycling Method | Added Params. | HumanEval | HumanEval+ | MBPP | MBPP+ | Overall |
|---------------|---------------------|---------------------|-----------|------------|------|-------|---------|
| Coder MoE [6] | Vanilla | 0.81B+2.43B | 64.6 | 61.0 | 63.9 | 51.4 | 60.2 |
| (ACL 24) | DeRS-SM (ours) | 0.81B+325M | 66.5 | 62.8 | 63.4 | 51.4 | 61.0 |
| (ACL 24) | DeRS-LM (ours) | 0.81B +9.09M | 65.9 | 62.8 | 62.9 | 51.9 | 60.9 |

yet non-overlapping medical multi-modal datasets. Second, the Med-MoE framework involves replacing the original FFN layer in the dense models with a parallel structure comprising a MoE layer and a universal FFN layer. And in our main experiments, the universal FFN layer is not considered as an expert for compression, as it's not sparsely activated by the router. Consequently, despite the decomposition and delta weight compression applied to the MoE layer's experts, the overall model performance is sustained by the universal FFN. In the Appendix, we further provide experimental results of applying DeRS compression to both experts and the universal FFN. Even with additional compression of the universal FFN, DeRS-Sparsification with a 0.8 drop rate or DeRS-Quantization with a 4-bit width can still significantly reduce redundancy while maintaining performance.

Results of DeRS Upcycling. As shown in Tab. 2, our DeRS upcycling remains efficient and effective across two Med-MoE models on the medical multi-modal task. For the Med-MoE-StableLM, our DeRS-SM and DeRS-LM upcycling strategies significantly reduce the additional parameter count by 1.24 billion compared to vanilla upcycling, while achieving performance improvement of 0.2% and 0.3%, respectively. Furthermore, for the Med-MoE-Phi architecture, the two DeRS upcycling methods reduce the number of additional parameters by 2.52 billion while maintaining comparable performance. Due to the presence of the universal FFN layers, the Med-MoE model obtained through DeRS upcycling still incurs an increase of 0.42 billion or 0.84 billion parameters compared to its corresponding dense model. In the Appendix, we further extend DeRS upcycling to the universal FFN, treating the construction of the universal FFN and the MoE layer's experts as a unified entity. This allows us to achieve performance comparable to that of vanilla upcycling with only a million-level increase in additional parameters over the dense model.

4.3. Code Generation Task

Model architecture. Following the settings in [6], we further evaluate our DeRS compression and DeRS upcycling on the code generation task. We directly upcycle the open-source language model, DeepSeek-Coder-Base-1.3B [11] into the Coder-MoE architecture for fine-tuning. Specifically, each block's FFN layer is replaced with a combination of a MoE layer and a universal FFN layer. The MoE layer contains



Figure 6. Performance of applying different DeRS compression methods to compress the vanilla upcycled Coder-MoE model.

four experts, with the top-1 expert activated for each input. **Datasets.** We utilize *evol-codealpaca-v1*, an open-source Evol-Instruct [40] dataset with 110K instruction-output pairs, to fine-tune the upcycled MoE model. After fine-tuning, we evaluate the model on widely used Python code generation benchmarks, including HumanEval [5] and MBPP [2], as well as on extended HumanEval+ and MBPP+ that contain additional tests generated by EvalPlus [36].

Results of DeRS Compression. As shown in Fig. 6, for the expert-specific delta weights in the Coder-MoE model, removing 60% of their elements or quantizing them to 2 bits can effectively eliminate redundancy without degrading performance. In comparison to the compression results of Med-MoE models shown in Fig. 5, the delta weight redundancy in the Coder-MoE model is notably lower, despite both experimental setups involving a parallel structure comprising a universal FFN and a MoE layer. The lower redundancy may be linked to the fact that the dense model utilized for Coder-MoE has not undergone any prior fine-tuning. Further experiments of compressing both the universal FFN and MoE experts are also provided in the Appendix.

Results of DeRS Upcycling. As shown in Tab 3, our DeRS-SM and DeRS-LM upcycling methods yield significant overall performance improvement of 0.7%-0.8% on the code generation task for the Coder-MoE model, while reducing over 2 billion extra parameters compared to vanilla upcycling. This substantial performance improvement may be attributed to the sparse activation mechanism and the lowbudget training setting. In vanilla upcycling, each expert is an independent FFN with extensive parameters, struggling to learn effectively due to the limited training data. In contrast, our DeRS upcycling allows multiple experts to share a base FFN while retaining their specific lightweight incremental parameters. This design enables the base FFN to leverage all training data for robust learning, while partial training data Table 4. Ablation studies on the hyper-parameters of two DeRS upcycling methods. DeRS-SM and DeRS-LM denote the Sparse-Matrix-based and Low-rank-Matrix-based DeRS upcycling respectively. The light-gray Added Params represents the additional parameters introduced by the universal FFN layers that are not considered as experts of MoE layers.

| DeRS-SM Rate | Added Params. | HumanEval (+) | MBPP (+) | Overall |
|-----------------|------------------|---------------|-------------|---------|
| 0.9999 | 0.81B+0.72M | 62.8 (60.4) | 63.7 (51.6) | 59.6 |
| 0.999 | 0.81B+3.64M | 64.0 (60.4) | 63.4 (51.6) | 59.9 |
| 0.99 | 0.81B+32.9M | 64.6 (61.6) | 62.7 (51.4) | 60.1 |
| 0.9 | 0.81B+325M | 66.5 (62.8) | 63.4 (51.4) | 61.0 |

(a) Different sparse rates for DeRS-SM upcycling

Table 5. Ablation studies on whether freezing the expert-shared base FFN for our two DeRS upcycling methods.

| Upcycling Method | Freeze Shared | HumanEval (+) | MBPP (+) | Overall |
|---------------------|------------------|----------------------------|----------------------------|--------------|
| DeRS-SM | × ✓ | 66.5 (62.8) 65.2 (61.6) | 62.4 (51.4) 61.4 (50.4) | 61.0 59.7 |
| DeRS-LM | × ✓ | 65.9 (62.8) 64.0 (61.0) | 62.9 (51.9) 61.7 (50.6) | 60.9 59.3 |

suffices to effectively train the sparsely activated lightweight incremental parameters. In the Appendix, we also present the experimental results of extending DeRS upcycling to the universal FFN within the Coder-MoE architecture, highlighting the ability of DeRS upcycling to achieve extremely efficient upcycled MoE models.

4.4. Ablation and Analysis

In this section, we conduct ablation studies and cost analysis of our DeRS upcycling on the code generation task.

Effect of freezing the expert-shared base FFN. As shown in Tab. 5, freezing the expert-shared base FFN results in a significant overall performance drop for both DeRS-SM and DeRS-LM upcycling strategies, with declines of 1.3% and 1.6%, respectively. These results validate the necessity of setting the expert-shared base FFN as trainable, as it represents foundational knowledge for the MoE experts and plays a crucial role in enabling effective learning.

Effect of the hyper-parameter. In Tab. 4, we explore the effect of the hyper-parameter of DeRS upcycling, which determines the parameter count of expert-specific incremental weights, specifically the sparse rate for DeRS-SM and the rank for DeRS-LM. As shown in Tab. 4a, for DeRS-SM, where the incremental weights take the form of sparse matrixes, a lower sparsity generally leads to better performance, albeit with an increase in the number of additional parameters. This can be attributed to the inherent limitation of the sparse matrix, which can only modify a subset of elements in another matrix of the same shape, as shown in Fig 3(b). Therefore, to enhance each expert's distinctiveness for better performance, a lower sparsity is necessary to allow the sparse-matrix-based incremental weights to make more substantial adjustments to the expert-shared base weight.

In contrast, as shown in Tab. 4b, our DeRS-LM upcycling

(b) Different ranks for DeRS-LM upcycling

| | - | | | |
|-----------------|------------------|---------------|-------------|---------|
| DeRS-LM Rank | Added Params. | HumanEval (+) | MBPP (+) | Overall |
| 1 | 0.81B+2.57M | 64.6 (61.6) | 63.4 (51.1) | 60.2 |
| 4 | 0.81B+9.09M | 65.9 (62.8) | 62.9 (51.9) | 60.9 |
| 16 | 0.81B+35.2M | 65.9 (62.2) | 63.2 (51.4) | 60.7 |
| 64 | 0.81B+140M | 63.4 (59.1) | 62.7 (50.9) | 59.0 |

Table 6. Cost efficiency comparison between vanilla upcycling and our two DeRS upcycling methods.

| Upcycling Method | Model Size | Training Memory | Inference Memory | Overall Performance |
|---------------------|---------------|--------------------|---------------------|------------------------|
| Vanilla | 4.59B | 25.9G | 10.5G | 60.2 |
| DeRS-SM (0.9rate) | 2.48B | 24.9G | 9.0G | 61.0 |
| DeRS-SM (0.99rate) | 2.19B | 21.0G | 6.1G | 60.1 |
| DeRS-LM (4rank) | 2.17B | 20.4G | 5.9G | 60.9 |

method can achieve strong performance with minimal extra parameters by using only a low rank (1, 4 or 16). This is due to the fact that the low-rank-matrix-based incremental weights are always able to make global adjustments to the expert-shared base weight, regardless of the rank, as shown in Fig 3(c). Moreover, if the rank is set too high (64), it can lead to performance degradation. This is likely because a high rank introduces significant redundancy in the incremental weights, making them less efficient for training.

Cost Analysis. Tab. 6 demonstrates the cost efficiency of our DeRS upcycling compared to vanilla upcycling. As shown, both DeRS upcycling methods achieve comparable or even superior performance while reducing model size and GPU memory consumption during both training and inference. For instance, our DeRS-LM upcycling significantly reduces model size by 52.7%, training memory by 21.2%, and inference memory by 43.8%, while achieving an overall performance improvement of 0.7%. These results showcase that DeRS upcycling propels upcycled MoE models towards a new level of efficiency.

5. Conclusion

In this work, we investigate the distinctive redundancy mechanisms of upcycled MoE models, and introduce an innovative DeRS paradigm that remodels MoE experts into one shared base weight and multiple exclusive compact weights. Further, we propose DeRS compression and DeRS upcycling to enhance the efficiency of expert parameters in upcycled MoE models during inference and training, respectively. We conduct comprehensive experiments to support the effectiveness of our proposals. Future works could focus on improving the parameter-efficient techniques in DeRS compression and DeRS upcycling or extending the DeRS paradigm to scenarios with higher training budgets.

6. Acknowledgments

This work is supported by National Key Research and Development Program of China (No. 2022ZD0160101), Shanghai Natural Science Foundation (No. 23ZR1402900), Shanghai Municipal Science and Technology Major Project (No. 2021SHZDZX0103). The computations in this research were performed using the CFFF platform of Fudan University.

References

- [1] Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. arXiv preprint arXiv:2404.14219, 2024. 1
- [2] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021. 7
- [3] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. arXiv preprint arXiv:2309.16609, 2023. 5
- [4] Marco Bellagente, Jonathan Tow, Dakota Mahan, Duy Phung, Maksym Zhuravinskyi, Reshinth Adithyan, James Baicoianu, Ben Brooks, Nathan Cooper, Ashish Datta, et al. Stable Im 2 1.6 b technical report. arXiv preprint arXiv:2402.17834, 2024. 5
- [5] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. 7
- [6] Yifeng Ding, Jiawei Liu, Yuxiang Wei, Terry Yue Zhuo, and Lingming Zhang. Xft: Unlocking the power of code instruction tuning by simply merging upcycled mixture-of-experts. arXiv preprint arXiv:2404.15247, 2024. 1, 2, 7
- [7] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR, 2022. 2
- [8] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022. 1, 2
- [9] Trevor Gale, Deepak Narayanan, Cliff Young, and Matei Zaharia. Megablocks: Efficient sparse training with mixtureof-experts. *Proceedings of Machine Learning and Systems*, 5: 288–304, 2023. 1
- [10] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. 5

- [11] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, YK Li, et al. Deepseek-coder: When the large language model meets programming-the rise of code intelligence. arXiv preprint arXiv:2401.14196, 2024. 7
- [12] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham.
 Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617, 2018.
 5
- [13] Xumeng Han, Longhui Wei, Zhiyang Dou, Zipeng Wang, Chenhui Qiang, Xin He, Yingfei Sun, Zhenjun Han, and Qi Tian. Vimoe: An empirical study of designing vision mixtureof-experts. arXiv preprint arXiv:2410.15732, 2024. 1
- [14] Ethan He, Abhinav Khattar, Ryan Prenger, Vijay Korthikanti, Zijie Yan, Tong Liu, Shiqing Fan, Ashwath Aithal, Mohammad Shoeybi, and Bryan Catanzaro. Upcycling large language models into mixture of experts. arXiv preprint arXiv:2410.07524, 2024. 1
- [15] Shwai He, Daize Dong, Liang Ding, and Ang Li. Demystifying the compression of mixture-of-experts through a unified framework. *arXiv preprint arXiv:2406.02500*, 2024. 2
- [16] Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. Pathvqa: 30000+ questions for medical visual question answering. arXiv preprint arXiv:2003.10286, 2020.
- [17] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 5
- [18] Mojan Javaheripi, Sébastien Bubeck, Marah Abdin, Jyoti Aneja, Sebastien Bubeck, Caio César Teodoro Mendes, Weizhu Chen, Allie Del Giorno, Ronen Eldan, Sivakanth Gopi, et al. Phi-2: The surprising power of small language models. *Microsoft Research Blog*, 2023. 5
- [19] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. arXiv preprint arXiv:2401.04088, 2024. 1
- [20] Songtao Jiang, Tuo Zheng, Yan Zhang, Yeying Jin, Li Yuan, and Zuozhu Liu. Med-moe: Mixture of domain-specific experts for lightweight medical vision-language models, 2024. 1, 2, 6
- [21] Aran Komatsuzaki, Joan Puigcerver, James Lee-Thorp, Carlos Riquelme Ruiz, Basil Mustafa, Joshua Ainslie, Yi Tay, Mostafa Dehghani, and Neil Houlsby. Sparse upcycling: Training mixture-of-experts from dense checkpoints. arXiv preprint arXiv:2212.05055, 2022. 1, 2
- [22] Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10, 2018. 6
- [23] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam

Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv* preprint arXiv:2006.16668, 2020. 1, 2

- [24] Bo Li, Yifei Shen, Jingkang Yang, Yezhen Wang, Jiawei Ren, Tong Che, Jun Zhang, and Ziwei Liu. Sparse mixture-ofexperts are domain generalizable learners. arXiv preprint arXiv:2206.04046, 2022. 1, 2
- [25] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Fanyi Pu, Jingkang Yang, Chunyuan Li, and Ziwei Liu. Mimic-it: Multi-modal in-context instruction tuning. *arXiv preprint arXiv:2306.05425*, 2023. 5
- [26] Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large languageand-vision assistant for biomedicine in one day. Advances in Neural Information Processing Systems, 36, 2024. 6
- [27] Jiachen Li, Xinyao Wang, Sijie Zhu, Chia-Wen Kuo, Lu Xu, Fan Chen, Jitesh Jain, Humphrey Shi, and Longyin Wen. Cumo: Scaling multimodal llm with co-upcycled mixture-ofexperts. arXiv preprint arXiv:2405.05949, 2024. 1
- [28] Pingzhi Li, Zhenyu Zhang, Prateek Yadav, Yi-Lin Sung, Yu Cheng, Mohit Bansal, and Tianlong Chen. Merge, then compress: Demystify efficient smoe with hints from its routing policy. arXiv preprint arXiv:2310.01334, 2023. 2
- [29] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. arXiv preprint arXiv:2305.10355, 2023. 5
- [30] Bin Lin, Zhenyu Tang, Yang Ye, Jiaxi Cui, Bin Zhu, Peng Jin, Junwu Zhang, Munan Ning, and Li Yuan. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*, 2024. 1, 2, 5
- [31] Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Dengr, Chong Ruan, Damai Dai, Daya Guo, et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. arXiv preprint arXiv:2405.04434, 2024. 1
- [32] Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A semantically-labeled knowledgeenhanced dataset for medical visual question answering. In 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), pages 1650–1654. IEEE, 2021. 6
- [33] Enshu Liu, Junyi Zhu, Zinan Lin, Xuefei Ning, Matthew B Blaschko, Shengen Yan, Guohao Dai, Huazhong Yang, and Yu Wang. Efficient expert pruning for sparse mixture-of-experts language models: Enhancing performance and reducing inference costs. arXiv preprint arXiv:2407.00945, 2024. 2
- [34] Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Aligning large multi-modal model with robust instruction tuning. *arXiv preprint arXiv:2306.14565*, 2023. 5
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 26296–26306, 2024. 2, 5
- [36] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chatgpt really correct?

rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems*, 36, 2024. 7

- [37] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European Conference on Computer Vision*, pages 216–233. Springer, 2025. 5
- [38] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022.
- [39] Xudong Lu, Qi Liu, Yuhui Xu, Aojun Zhou, Siyuan Huang, Bo Zhang, Junchi Yan, and Hongsheng Li. Not all experts are equal: Efficient expert pruning and skipping for mixture-of-experts large language models. arXiv preprint arXiv:2402.14800, 2024. 2
- [40] Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. Wizardcoder: Empowering code large language models with evol-instruct. arXiv preprint arXiv:2306.08568, 2023. 7
- [41] Joan Puigcerver, Carlos Riquelme, Basil Mustafa, and Neil Houlsby. From sparse to soft mixtures of experts. arXiv preprint arXiv:2308.00951, 2023. 1
- [42] 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. 5
- [43] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. Advances in Neural Information Processing Systems, 34:8583–8595, 2021. 1
- [44] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixtureof-experts layer. arXiv preprint arXiv:1701.06538, 2017. 1
- [45] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. 5
- [46] Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. To see is to believe: Prompting gpt-4v for better visual instruction tuning. *arXiv preprint arXiv:2311.07574*, 2023. 5
- [47] Tianwen Wei, Bo Zhu, Liang Zhao, Cheng Cheng, Biye Li, Weiwei Lü, Peng Cheng, Jianhao Zhang, Xiaoyu Zhang, Liang Zeng, et al. Skywork-moe: A deep dive into training techniques for mixture-of-experts language models. arXiv preprint arXiv:2406.06563, 2024. 1
- [48] Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. Openmoe: An early effort

on open mixture-of-experts language models. *arXiv preprint arXiv:2402.01739*, 2024. 2

- [49] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. arXiv preprint arXiv:2407.10671, 2024. 1
- [50] Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*, 2024. 4
- [51] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. arXiv preprint arXiv:2308.02490, 2023. 5
- [52] Bo Zhao, Boya Wu, Muyang He, and Tiejun Huang. Svit: Scaling up visual instruction tuning. *arXiv preprint arXiv:2307.04087*, 2023. 5
- [53] Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam Shazeer, and William Fedus. Stmoe: Designing stable and transferable sparse expert models. *arXiv preprint arXiv:2202.08906*, 2022. 1