

MoDE: Effective Multi-task Parameter Efficient Fine-Tuning with a Mixture of Dyadic Experts

Lin Ning*, Harsh Lara, Meiqi Guo, Abhinav Rastogi*[†]

Google DeepMind

{linning, harshlara, mqguo, abhirast}@google.com

Abstract

Parameter-efficient fine-tuning techniques like Low-Rank Adaptation (LoRA) have revolutionized the adaptation of large language models (LLMs) to diverse tasks. Recent efforts have explored mixtures of LoRA modules for multi-task settings. However, our analysis reveals redundancy in the down-projection matrices of these architectures. This observation motivates our proposed method, **Mixture of Dyadic Experts (MoDE)**, which introduces a novel design for efficient multi-task adaptation. This is done by sharing the down-projection matrix across tasks and employing atomic rank-one adapters, coupled with routers that allow more sophisticated task-level specialization. Our design allows for more fine-grained mixing, thereby increasing the model’s ability to jointly handle multiple tasks. We evaluate MoDE on the Supernatural Instructions (SNI) benchmark consisting of a diverse set of 700+ tasks and demonstrate that it outperforms state-of-the-art multi-task parameter-efficient fine-tuning (PEFT) methods, without introducing additional parameters. Our findings contribute to a deeper understanding of parameter efficiency in multi-task LLM adaptation and provide a practical solution for deploying high-performing, lightweight models.

1 Introduction

Large language models (LLMs) have demonstrated remarkable capabilities across various natural language processing tasks, ranging from text generation and translation to question-answering and summarization (Brown et al., 2020; Team et al., 2023; GemmaTeam et al., 2024; OpenAI and et al., 2024). Tailoring these models for optimal performance on diverse tasks is crucial for real-world applications, where LLMs must handle various user requests and instructions. However, effectively adapting

¹Equal contribution.

²Corresponding author.

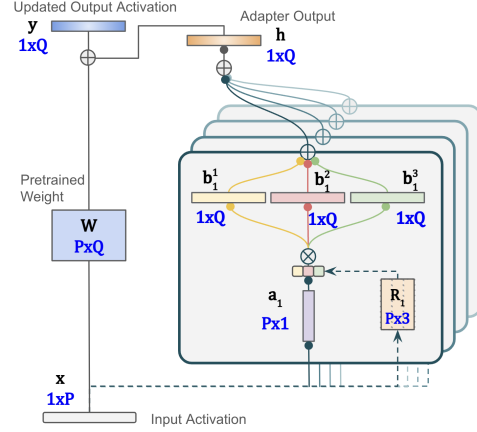


Figure 1: Mixture of Dyadic Experts with 3 experts and a rank of 4, with each slice corresponding to a rank dimension. Our architecture allows independent routing at each rank. When number of mixtures is 1, our architecture is equivalent to traditional LoRA.

these large models to multiple tasks presents significant challenges. Fine-tuning separate models for each task is computationally expensive and requires a vast amount of storage due to the large model sizes. Moreover, independently trained models hinder knowledge transfer between tasks, potentially limiting models’ performance and their generalization capability.

Multi-task learning (MTL) (Caruana, 1997; Ruder, 2017; Frohmann et al., 2024) offers a promising solution to these challenges. By training a single model on multiple tasks simultaneously, MTL aims to improve parameter efficiency, enhance generalization, and potentially boost performance on individual tasks through knowledge transfer. Parameter-efficient fine-tuning (PEFT) techniques, such as Low-Rank Adaptation (LoRA) (Hu et al., 2021), have further enhanced efficiency by introducing only a small number of trainable parameters. LoRA efficiently represents weight changes during fine-tuning using two low-rank pro-

jection matrices: one down-projects input features to a smaller size, and another up-projects the resulting low-dimensional representation back to the original output size.

Mixture-of-Experts (MoE) architectures (Sukhbaatar et al., 2024; Li et al., 2022; Jiang et al., 2024; Fedus et al., 2022) have emerged as a powerful approach to scale model capacity and expertise, enabling LLMs to handle a wider range of tasks. Recent studies (Feng et al., 2024; Zhu et al., 2023; Zadouri et al., 2023; Liu et al., 2023; Li et al., 2024; Lin et al., 2024b) have explored integrating LoRA with Mixture-of-Experts architectures (LoRA-MoE) to extend LLMs’ capabilities to multi-task adaptation. However, our analysis reveals redundancy in down-projection matrices of these architectures. This redundancy leads to an inefficient utilization of parameters, potentially limiting the effectiveness in capturing the unique characteristics of each task.

In this work, we propose Mixture of Dyadic Experts (MoDE) (Figure 1), a novel parameter-efficient framework for multi-task adaptation. MoDE leverages a single, shared down-projection matrix across all experts to reduce parameter redundancy. Furthermore, MoDE introduces atomic rank-one adapters, enabling fine-grained task specialization and knowledge sharing. Crucially, MoDE incorporates a sophisticated routing mechanism that allows for more nuanced and flexible combinations of these rank-one adapters, further enhancing the model’s expressive power while maintaining parameter efficiency.

We rigorously evaluate MoDE on the multi-task Supernatural Instructions benchmark (Wang et al., 2022b). Our results demonstrate that MoDE consistently outperforms state-of-the-art multi-task PEFT methods, including those based on LoRA-MoE, while utilizing comparable number of additional parameters. This underscores MoDE’s efficacy in achieving both strong performance and parameter efficiency, making it a promising approach for deploying multi-task LLMs in real-world applications.

Our **key contributions** are as follows:

- Identify and address the redundancy in down-projection matrices in existing LoRA-based MoE architectures.
- Propose MoDE, a novel architecture leveraging a shared down-projection matrix and atomic rank-one adapters, coupled with a so-

phisticated routing mechanism for efficient and expressive performance.

- Demonstrate the superior performance of MoDE compared to state-of-the-art multi-task LoRA-based MoE methods on the Supernatural Instructions benchmark, while maintaining parameter efficiency.

2 Related Work

Mixture-of-Experts and LoRA Parameter-efficient fine-tuning (PEFT) methods like LoRA (Low-Rank Adaptation) (Hu et al., 2021) enable effective adaptation of LLMs to downstream tasks with reduced trainable parameters compared to full fine-tuning. Recent research has explored combining LoRA with Mixture-of-Experts (MoE) architectures for multi-task learning. Some approaches, such as Mixture-of-LoRAs (MoA) (Feng et al., 2024) and MOELoRA (Liu et al., 2023), use domain-specific LoRA modules and explicit routing strategies to adapt to diverse tasks. Others, like SiRA (Zhu et al., 2023) and MixLoRA (Li et al., 2024), introduce sparse MoE mechanisms with specialized routing and/or load-balancing techniques to enhance efficiency while maintaining performance. MoLoRA (Zadouri et al., 2023) combines MoE with lightweight LoRA experts to achieve extreme parameter efficiency. AdaMix (Wang et al., 2022a) tunes a mixture of adaptation modules within each Transformer layer to capture multiple views of a single task. It also demonstrates the benefits of sharing up-projection matrices in low-data scenarios. In contrast, our multi-task approach, MoDE, leverages shared project-down matrices across LoRA modules, motivated by their observed similarity across tasks.

Multi-task PEFT Other work has explored extending PEFT to multi-task settings, where a single model needs to adapt to diverse tasks. LoraHub (Huang et al., 2023) investigates LoRA composability for cross-task generalization by dynamically assembling LoRA modules trained on different tasks to adapt to unseen tasks. ZipLoRA (Shah et al., 2023) tackles the challenge of combining independently trained style and subject LoRAs for controllable joint generation. FLix (Lin et al., 2024a) focuses on multi-task multilingual model adaptation by associating each dataset feature with its own low-rank weight update parameters for improved generalization across diverse datasets. MoLE (Wu et al., 2024) implements a




Figure 2: Illustration of a basic LoRA module.

hierarchical weight control approach with learnable gating functions for the optimal composition of trained LoRA layers, treating each layer as a distinct expert. Different from composing over all LoRA adaptors, [Ostapenko et al. \(2024\)](#) explores building a library of trained LoRA adaptors and using a zero-shot routing mechanism (Arrow) to select relevant adaptors for new tasks dynamically.

3 Method

The Mixture of Dyadic Experts (MoDE) architecture presents a novel approach for multi-task learning, building upon and extending the traditional Low-Rank Adaptation (LoRA) and mixture-of-experts (MoE) design.

3.1 Background

Low-Rank Adaption (LoRA) LoRA ([Hu et al., 2021](#)) efficiently adapts LLMs to downstream tasks ([Shah et al., 2023](#)) by freezing pre-trained model weights and injecting trainable rank decomposition matrices into each layer. Given a feed-forward layer with input $\mathbf{x} \in \mathbb{R}^{1 \times P}$ and weight matrix $\mathbf{W}_0 \in \mathbb{R}^{P \times Q}$, LoRA introduces a down-projection matrix $\mathbf{A} \in \mathbb{R}^{P \times r}$ and an up-projection matrix $\mathbf{B} \in \mathbb{R}^{Q \times r}$ (Figure 2). The output of the layer is

$$\mathbf{y} = \mathbf{x}\mathbf{W}_0 + \mathbf{x}\mathbf{A}\mathbf{B}^T.$$

During training, only \mathbf{A} and \mathbf{B} are updated.

Dyadic Product Representation A dyadic product (or outer product) is a matrix multiplication between two vectors. Given vectors $\mathbf{u} \in \mathbb{R}^{p \times 1}$ and $\mathbf{v} \in \mathbb{R}^{q \times 1}$, their dyadic product $\mathbf{u} \otimes \mathbf{v}$ is a matrix of size $p \times q$. The LoRA update $\Delta\mathbf{W} = \mathbf{A}\mathbf{B}^T$ can be expressed as a sum of dyadic products by decomposing \mathbf{A} and \mathbf{B} into their column vectors ([Liu et al., 2024](#)):

$$\begin{aligned} \Delta\mathbf{W} &= [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_r] * [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_r]^T \\ &= \sum_{i=1}^r (\mathbf{a}_i \otimes \mathbf{b}_i) \end{aligned}$$

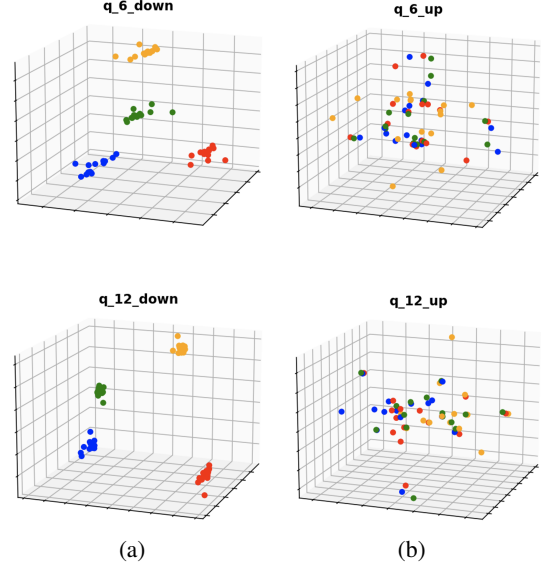


Figure 3: Scatter plots showing the three most prominent principal components of all constituent vectors in the LoRA projection matrices for 15 independently trained single task model with shared initialization. Plots q_6_down and q_12_down (q_6_up and q_12_up) illustrate the down (up) projections of query matrices at layers 6 and 12, respectively. Each point represents a vector of a LoRA matrix, and different colors indicate different positions along the rank dimension in a matrix. The clear clustering of down-projection vectors suggests that the down-projection matrices are task-agnostic, motivating the design of the MoDE architecture.

where \mathbf{a}_i and \mathbf{b}_i are column vectors of \mathbf{A} and \mathbf{B} , respectively. This can be plugged into the output equation to get:

$$\mathbf{y} = \mathbf{x}\mathbf{W}_0 + \mathbf{x} \sum_{i=1}^r (\mathbf{a}_i \otimes \mathbf{b}_i)$$

LoRA-MoE Mixture-of-Experts (MoE) utilizes a combination of sub-models (experts), each specializing in different aspects of the underlying tasks, along with a gating mechanism to dynamically route inputs to the most suitable experts ([Shazeer et al., 2017](#)). Figure 4(a) illustrates a traditional LoRA-MoE approach ([Zadouri et al., 2023](#)), where m LoRA experts ($E^i = \mathbf{A}^i\mathbf{B}^{iT}$, $i \in \{1, \dots, m\}$) are added to each layer. A router \mathcal{R} , parameterized by $\mathbf{W}_R \in \mathbb{R}^{P \times m}$, determines which expert to use, yielding the output:

$$\mathbf{y} = \mathbf{x}\mathbf{W}_0 + \sum_{i=1}^m \mathcal{R}^i(\mathbf{x})(\mathbf{x}\mathbf{A}^i\mathbf{B}^{iT})$$

where $\mathcal{R}^i(\mathbf{x})$ is the routing probability for E^i .

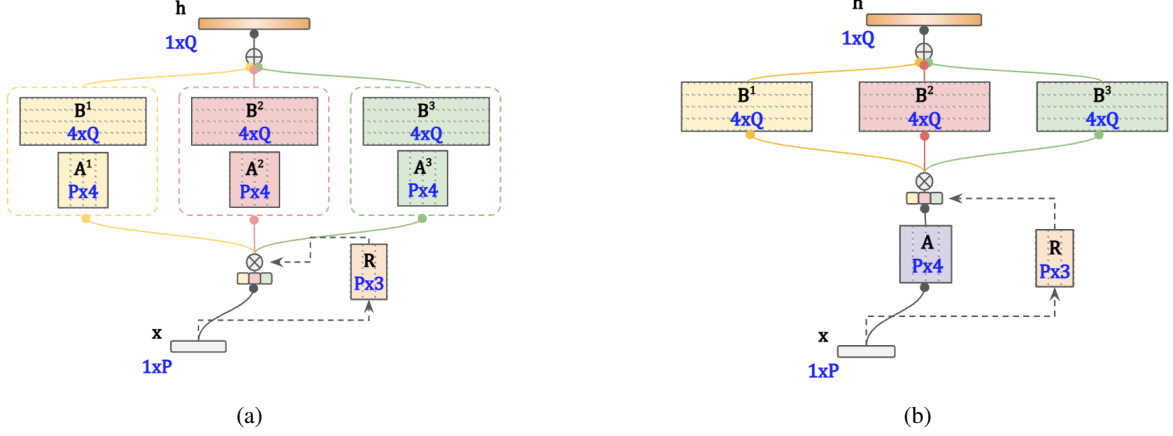


Figure 4: Illustration of (a) traditional LoRA Mixture-of-Experts, (b) traditional LoRA Mixture-of-Experts with shared down-projection matrix.

3.2 Motivating Observation

To motivate our proposed MoDE architecture, we first present an empirical analysis of projection matrices of LoRA modules that are independently trained on a set of tasks from the same initialization (B_i to 0 and A_i from a random normal distribution with a standard deviation of 0.01 and mean 0). We selected 15 diverse tasks from the Supernatural Instructions benchmark (see Section 4.1 for details) and trained 15 LoRA modules for this study. We visualized the learned LoRA parameters using Principal Component Analysis (PCA), focusing on the distribution of vectors obtained by slicing the up-projection and down-projection matrices along their rank dimension, i.e., the vectors featuring in the dyadic product representation in Section 3.1.

Figure 3 shows the resulting scatter plots. Notably, we observe that down-projection matrix vectors from different LoRA modules tend to cluster into distinct groups, with vectors corresponding to the same position along the rank dimension forming tight clusters. In contrast, up-projection matrices exhibit no such clustering. This suggests down-projection matrices are task-agnostic, while up-projection matrices are more task-specific.

This empirical finding motivates the MoDE architecture, which leverages a shared down-projection matrix to reduce parameter redundancy. We further improve this design by leveraging the dyadic formulation to introduce a more sophisticated routing strategy which enables a more fine-grained task-specific adaptation. Note that vectors from the up-projection matrix of each task (Figure 3(b)) are very different from each other. There-

fore, we don't consider sharing an up-projection matrix in our design. Subsequent sections will detail MoDE's architecture and demonstrate its effectiveness in achieving both parameter efficiency and strong multi-task performance.

3.3 Mixture of Dyadic Experts (MoDE)

Inspired by the observations in Section 3.2, we introduce Mixture of Dyadic Experts (MoDE), a novel framework for efficient multi-task adaptation that incorporates two key innovations: (i) shared down-projection matrices for more efficient parameter utilization, and (ii) a sophisticated routing strategy to promote better task-level specialization.

3.3.1 Shared Down-Projection Matrix

Before introducing MoDE, we first present a simple modification of the traditional LoRA-MoE, termed LoRA-MoE-SD (Figure 4(b)), where all experts share a single down-projection matrix A . The output of a layer with LoRA-MoE-SD is:

$$y = xW_0 + \sum_{i=1}^m \mathcal{R}^i(x)(xA B^i T)$$

where A is the shared down-projection matrix, and B^i is the up-projection matrix for expert E_i .

Parameter Efficiency By sharing a single down-projection matrix A across all experts, LoRA-MoE-SD reduces the number of trainable parameters for these matrices from $m \cdot P \cdot r$ to $P \cdot r$.

3.3.2 Fine-Grained Routing

LoRA-MoE-SD, albeit with smaller parameters, offers only m choices for up projection. This is because the router inherently introduces a constraint

that all the r dimensions of each expert must route together. In other words, the dyadic representation of LoRA-MoE update contains $m \times r$ terms, but the router only provides m weights. MoDE addresses this by introducing atomic rank-one adapters. This design choice allows MoDE to leverage the dyadic product representation of LoRA, where each rank-one update captures a specific direction of change in the original weight matrix.

MoDE employs m rank-one experts for each column vector \mathbf{b}_j in \mathbf{B} (Figure 1), resulting in $m \times r$ experts. Each expert E_j^i (where $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, r\}$) specializes in a specific component of the up-projection, represented as a dyadic product $\mathbf{a}_j \otimes \mathbf{b}_j^{iT}$, where \mathbf{a}_j is the j -th column vector of the shared down-projection matrix \mathbf{A} and \mathbf{b}_j^i is the vector representing the i -th rank-one expert for the j -th component of \mathbf{B} .

Let $\mathcal{R}_j^i(\mathbf{x})$ is the routing probability for expert E_j^i given input \mathbf{x} , the MoDE module outputs:

$$\mathbf{y} = \mathbf{x}\mathbf{W}_0 + \sum_{i=1}^m \sum_{j=1}^r \mathcal{R}_j^i(\mathbf{x})(\mathbf{x}(\mathbf{a}_j \otimes \mathbf{b}_j^{iT})).$$

Model Expressivity The router in MoDE independently selects the expert used for each vector \mathbf{b}_j . This fine-grained control allows for flexible combination of these dyadic product experts, enabling MoDE to dynamically compose a specialized up-projection matrix tailored to the input and task. For example, if \mathbf{B} has rank 4, the router might select E_1^1 for \mathbf{b}_1 , E_2^3 for \mathbf{b}_2 , E_3^2 for \mathbf{b}_3 , and E_4^1 for \mathbf{b}_4 .

With m rank-one experts per vector \mathbf{b}_j of a rank- r up-projection matrix, MoDE can model m^r different expert compositions, allowing for a wide range of task-specific adaptations compared to the m experts in a traditional LoRA-MoE, given a similar number of parameters. This increased expressivity, derived from the flexible combination of dyadic products, allows MoDE to better capture the nuances of individual tasks while maintaining scalability for a large number of tasks.

MoDE Routing MoDE utilizes a token-level soft routing strategy, where the router \mathcal{R} assigns a weight to each rank-one expert for a given input token. The weighted sum of experts' outputs determines the final output. This approach enables dynamic utilization of the most relevant experts for each input, facilitating nuanced and context-aware adaptation.

The router network is denoted as $\mathbf{W}_{\mathcal{R}} \in \mathbb{R}^{r \times P \times m}$, where $\mathbf{W}_{\mathcal{R};j} \in \mathbb{R}^{P \times m}$ represents the

network for vector \mathbf{b}_j in the up-projection matrix. For an input \mathbf{x} , the routing weights $\mathcal{R}_j \in \mathbb{R}^{1 \times m}$ for the experts corresponding to \mathbf{b}_j are

$$\mathcal{R}_j(\mathbf{x}) = \text{softmax}(\mathbf{x} \cdot \mathbf{W}_{\mathcal{R};j}).$$

This mechanism allows MoDE to adaptively combine the expertise of multiple rank-one adapters, leading to improved multi-task performance.

3.4 Generalization

The rank-1 adapters in MoDE can be generalized to rank- p , where the router selects a composition of rank- p adapters for each input. This requires the LoRA rank r to be divisible by adapter rank p . The generalized output calculation becomes:

$$\mathbf{y} = \mathbf{x}\mathbf{W}_0 + \sum_{i=1}^m \sum_{k=1}^{r/p} \mathcal{R}_k^i(\mathbf{x}) \cdot \mathbf{x}\mathbf{A}_k \mathbf{B}_k^{iT}$$

where

$$\mathbf{A}_k \mathbf{B}_k^{iT} = \sum_{j=1}^p (\mathbf{a}_{j+p(k-1)} \otimes \mathbf{b}_{j+p(k-1)}^{iT}).$$

We denoted this generalized module as MoDE $m \times r \times p$. Note that MoDE $1 \times r \times r$ is functionally equivalent to a LoRA module of rank r , and MoDE $m \times r \times r$ is functionally equivalent to LoRA-MoE-SD with rank r and m experts.

4 Experiments

We comprehensively evaluate MoDE's performance and analyze its design choices through three sets of experiments on the Supernatural Instructions (SNI) benchmark (Wang et al., 2022b): (1) multi-task evaluation on the full dataset, (2) an ablation study on generalized MoDE architecture (Section 3.4) on the full dataset, and (3) a case study with a fixed number of tasks and parameter budgets. This section details the experimental setup and presents the results.

4.1 Datasets and Metrics

We leverage the Supernatural Instructions (SNI) dataset (with 1,616 diverse instruction-following tasks covering 76 distinct task types) for our experiments, focusing on the 756 English-only tasks from the default train split for both training and evaluation. For each task, we split the examples into a 90% training set and 10% evaluation set. For multi-task experiments, we create mixed training and evaluation datasets by combining examples from all 756 tasks.

We also curate a diverse subset of 15 individual tasks from the 756 tasks for the fixed parameter

Category (Task ID)	Instruct	Input	Output
QuestionAnswering(24)	104	87/90.9	8/8.8
WrongCandidateGeneration(25)	127	100/104.1	8/8.8
QuestionGeneration(74)	97	143/155	12/12.5
GrammarErrorDetection(89)	89	10/10.6	6/6
LinguisticProbing(114)	66	25/25.7	1/1
PosTagging(155)	19	23/23.7	1/1
Explanation(192)	43	123/131.4	30/33.7
StoryComposition(269)	184	82/82.2	34/4.7
StereotypeDetection(279)	89	15/15.5	2/2
CommonsenseClassification(291)	54	17/17.8	1/1
ProgramExecution(622)	62	94/93.6	99/98.5
FillInTheBlank(672)	24	13/3.2	1/1
PoemGeneration(1711)	83	3/3.6	44/59
DialogueGeneration(1729)	58	155/156.9	13/12.7

Table 1: Sequence lengths of instruction, input (median/mean) and output (median/mean) in the selected SNI datasets for the fixed parameter budget case study.

budgets case study. These tasks are selected from 15 different categories to ensure a comprehensive evaluation across different domains. Each selected task has over 5k training instances and 500-650 evaluation instances. See Table 1 for details.

Evaluation Metric We report ROUGE-L, the default metric for SNI dataset, for all experiments.

	Eval	Add. Params.
LoRA 64	56.11	6.31%
MoLORA 16×4	57.77	7.62%
MoLORA-SD 16×4	58.28	2.71%
MoDE 16×4	60.00	6.64%

Table 2: Multitask performance comparison between LoRA, MoLORA, MoLORA-SD and MoDE. The evaluation metric used is ROUGE-L. The last column represents the total number of adapter parameters as a percentage of the total number of non-embedding parameters in Gemma-2B.

4.2 Implementation Details

Model The Gemma 2B language model (GemmaTeam et al., 2024) serves as the foundational LLM for all experiments due to its state-of-the-art performance on a variety of natural language processing tasks and its efficient size.

Fine-tuning Setup For all experiments, we fine-tune the parameter-efficient adaptors using Adafactor optimizer(Shazeer and Stern, 2018) with learning rate 1e-3, total sequence length 1024, and batch size 128 for 20,000 steps. We reported the averaged result of 5 runs for each experiments.

Model	LoRA	MoLORA	MoLORA-SD
LoRA	\	\	\
MoLORA	69%	\	\
MoLORA-SD	69%	60%	\
MoDE	78%	73%	68%

Table 3: Task-level win rate against baseline models.

4.3 Multi-task performance

We assess model performance on the full dataset comprising 756 tasks. We compare MoDE with vanilla LoRA and MoLORA (Zadouri et al., 2023), a strong baseline approach of LoRA-MoE for multi-task adaptation. To better understand the benefit of removing redundancy in down-projection matrices, we also apply down-projection sharing to MoLORA, referred to as MoLORA-SD.

The specific models and configurations are:

- *LoRA 64*: A base LoRA module with rank 64, having a comparable number of parameters as other MoE models.
- *MoLORA 16×4*: A LoRA-MoE model with 4 experts, each using a rank-4 LoRA module.
- *MoLORA-SD 16×4*: A LoRA-MoE model with a shared down-projection matrix and 4 experts with rank-4 up-projection matrices.
- *MoDE 16×4*: A MoDE model with 4 experts per vector of a rank-4 up-projection matrix.

Model Comparison Table 6 reveal several key findings. First, all LoRA-MoE methods outperforms a single LoRA, with improvements ranging from 2.96% (MoLORA 16x4) to 6.93% (MoDE 16x4) in ROUGE-L scores. This demonstrates the effectiveness of MoE-based architectures in multi-task settings, allowing the model to leverage specialized experts for different tasks.

The advantage of parameter sharing is evident in the comparison between MoLORA 16x4 and MoLORA-SD 16x4. By sharing the down-projection matrix, MoLORA-SD achieves an 0.88% improvement over MoLORA while using only 36% of the additional parameters, highlighting the benefit of reducing parameter redundancy.

Among all the models, MoDE 16x4 achieves the highest overall performance by leveraging both shared down-projection and rank-one adapters. This showcases the effectiveness of MoDE in balancing parameter efficiency with models’ expressivity and adaptability in multi-task scenarios.

We further conduct a task-level analysis of all the 756 evaluation tasks and report the win rate between each pair of models in Table 3. Each row of the table shows the win rate of target model against models in each column. We find that all LoRA-MoE methods outperform the single LoRA baseline in around 70%-80% of tasks. Importantly, MoDE significantly outperforms all three baselines, passing the significance test over 50% win rate at 0.99 confidence, demonstrating its consistent superiority across a wide range of tasks.

4.4 Generalized MoDE Architecture

To gain a deeper understanding of the impact expert rank (p) on model performance, we conduct two sets of experiments with the generalized MoDE architecture (Section 3.4). Following the notations m and r , the number of experts becomes $m \times r/p$, where each expert is a rank p adaptor.

Varying Expert Rank (p) In the first set of experiments, we vary the expert rank p while keeping m and r fixed on two m and r combinations (4×16 and 16×4). The results are presented in Table 4. We observe that, with fixed m and r , increasing the expert rank generally leads to improved performance, as indicated by higher ROUGE-L scores. This suggests that increasing the expressiveness of individual experts contributes to better overall multi-task performance.

Model Config.			Eval	Add. Params.
m	r	p		
4	16	16	58.51	2.71%
4	16	8	59.15	3.04%
4	16	4	59.56	3.69%
4	16	2	59.76	5.00%
4	16	1	59.93	7.62%
16	4	4	58.97	2.71%
16	4	2	59.30	4.02%
16	4	1	59.91	6.64%

Table 4: Experiments on generalized MoDE architecture. r : the rank of the LoRA matrices. p : the rank of each experts. Number of experts: $m \times r/p$. The last column represents the total number of adapter parameters as a percentage of the total number of non-embedding parameters in Gemma-2B.

Iso-parametric Configurations In the second set of experiments, we explore iso-parametric configurations of MoDE, where the total number of added parameters remains approximately constant across different model configurations. We vary the

LoRA rank r (4, 8, or 16) and expert rank p (from 1 to r), adjusting m to maintain a consistent parameter budget. Table 5 presents the results of these experiments, providing insights into the trade-offs between different hyperparameter choices under a fixed resource constraint.

Model Config.			Eval	Add. Params.
m	r	p		
42	4	4	59.89	6.58%
27	4	2	60.06	6.55%
16	4	1	59.91	6.64%
27	8	8	60.52	6.48%
20	8	4	60.74	6.61%
12	8	2	60.23	6.19%
7	8	1	59.73	6.18%
15	16	16	60.77	6.55%
12	16	8	60.94	6.49%
8	16	4	60.77	6.07%
5	16	2	60.28	5.92%
3	16	1	59.42	6.04%

Table 5: Experiments on generalized MoDE architecture with iso-parametric constraint.

Impact of Expert Rank (p): For a fixed LoRA rank r , increasing the expert rank p generally improve performance, as seen when comparing configurations with the same LoRA rank but different expert ranks (e.g., 16×1 v.s. 16×2 v.s. 16×4 v.s. 16×8). This indicates that enhancing the expressiveness of individual experts with higher expert ranks contributes to better multi-task performance under the iso-parametric setting. However, gains diminish as the expert rank p approaches to the LoRA rank r , suggesting that using $p < r$ is beneficial.

The best overall performance is achieved by the $12 \times 16 \times 8$ configuration, which balances a moderate number of experts with a reasonably high LoRA rank and expert rank. This emphasizes finding the optimal balance between these hyperparameters for strong multi-task performance.

4.5 Case Study with Fixed Parameter Budget

In real-world scenarios, there are often constraints on the number of additional parameters that can be introduced during model adaptation. To assess the effectiveness of MoDE under such constraints, we conduct a case study with a fixed parameter budget determined by the baseline models.

We leverage the diverse subset of 15 individual tasks from the SNI dataset, each belonging to a distinct category, as described in Section 4.1. Our baseline model consists of 15 individual rank-4 LoRA adapters (denoted as *LoRA* 15×4), one for




Figure 5: Performance comparison among various model configurations on 15 tasks with a fixed parameter budget.

each task, resulting in approximately 6 million additional trainable parameters compared to the frozen LLM backbone. This baseline establishes our parameter budget for further experimentation. We also compare the performance with 15 full finetuned models (denoted as *Full*), each dedicated to a single task.

To ensure a fair comparison, we identify configurations for LoRA (trained on mixture of tasks), MoLoRA, MoLoRA-SD, and MoLoRA-SD that introduce a similar number of parameters (approximately 6 million). We systematically explore combinations of experts (m) and ranks (r) for each method, aiming to keep the total number of additional parameters as close as possible to the baseline budget. Specifically, we experiment with ranks of 4, 8, 16, and 32, adjusting the number of experts accordingly to maintain the desired parameter count. Notation " $m \times r$ " indicates a model with m experts, each using adapters associated with a rank r LoRA. This leads to the following configurations:

- *LoRA* 1×60
- *MoLoRA* 14×4 , 6×8 , 3×16
- *MoLoRA* 12×4 , 6×8 , 3×16
- *MoLoRA-SD* 36×4 , 24×8 , 12×16 , 5×32

Note that MoLoRA-SD, benefiting from its parameter efficiency due to the shared down-projection matrix, can support a higher-rank (5×32) configuration within the budget.

Overall Performance Figure 5 displays the model performance for each configuration on each task. The reported performance is the normalized ROUGE-L score, designed to highlight performance differences between models. Let \bar{R}_r be

the average ROUGE-L score of a single model on a task across different runs and $\bar{R}_{r,m}$ be the mean of \bar{R}_r scores across all models on a task. The normalized ROUGE-L score is calculated as $(\bar{R}_r - \bar{R}_{r,m}) / \bar{R}_{r,m}$. This metric indicates how a model performs relative to the average performance of all models on a given task. A higher normalized ROUGE-L score indicates better performance (higher \bar{R}_r).

MoLoRA consistently achieves comparable or superior performance to the 15 individual LoRA adapters baseline, the full finetuned models, and other MoE approaches (MoLoRA and MoLoRA-SD). This demonstrates MoLoRA’s ability to effectively leverage its shared down-projection matrix and dyadic experts to achieve strong multi-task performance while maintaining parameter efficiency. Notably, MoLoRA substantially improves the average ROUGE-L score across all tasks compared to all other models.

While the full finetuned models are better than the LoRA 15×4 or LoRA 1×60 configurations (as expected), they are potentially susceptible to overfitting due to the limited training data available for each individual task. MoLoRA’s strong performance suggests its ability to mitigate this overfitting.

Benefit of down-projection sharing MoLoRA-SD, which shares the down-projection matrix like MoLoRA, generally outperforms MoLoRA with independent down-projection matrices. This highlights the benefits of reduced parameter redundancy through a shared down-projection matrix.

Impact of Experts and Rank While individual task performance varies across different MoLoRA configurations, we observe that the overall average

performance across all tasks and examples remains relatively stable despite changes in the number of experts (m) and the rank (r). This suggests that MoDE’s performance is robust to these hyperparameter choices, and there may not be a single "best" configuration for all scenarios. The optimal choice of experts and rank depends on specific task characteristics, resource constraints, or desired trade-offs between model size and performance.

Benefit of MoDE routing We observed that MoDE performs better than MoLORA-SD. It suggests that the dyadic experts in MoDE contribute to its superior expressivity and adaptability compared to simply down-projection matrix sharing.

These findings demonstrate that MoDE can effectively leverage a fixed parameter budget to achieve strong multi-task performance. Its shared down-projection matrix and mixtures of dyadic experts enable a balance between parameter efficiency and expressive power, making it a promising approach for deploying multi-task LLMs in resource-constrained environments.

5 Conclusion and Future Work

This paper introduces MoDE (Mixture of Dyadic Experts), a novel parameter-efficient fine-tuning method for multi-task LLM adaptation. MoDE leverages a shared down-projection matrix and rank-one adapters with a sophisticated routing mechanism to improve parameter efficiency and expressive power. Experiments on the Supernatural Instructions benchmark show that MoDE consistently outperforms state-of-the-art multi-task PEFT methods, achieving superior performance with comparable parameter efficiency. This highlights MoDE’s potential as an efficient and effective solution for multi-task LLM adaptation, particularly in resource-constrained environments. Future work will explore more routing strategies, analyze task-specific patterns, and evaluate MoDE on larger models and alternative PEFT techniques.

6 Limitations

While our proposed MoDE architecture demonstrates promising results in multi-task LLM adaptation, there are several limitations that warrant further investigation.

Routing Strategy The current MoDE implementation utilizes a relatively simple routing mechanism based on a softmax function. While effective

in our experiments, exploring more sophisticated routing strategies that incorporate task relationships or input-specific features could potentially further improve performance.

Hyperparameter Sensitivity The optimal number of experts and rank of the LoRA matrices can vary depending on the specific task distribution and available resources. While our ablation study provides some insights, a more comprehensive exploration of hyperparameter sensitivity could help identify optimal configurations for different scenarios.

Computational Overhead While MoDE significantly reduces parameter count compared to traditional LoRA-MoE, the routing mechanism introduces additional computational overhead during inference. This overhead could become a bottleneck in real-time applications with strict latency requirements. Investigating ways to optimize the routing process or reduce its computational cost would be beneficial.

Evaluation Benchmark Our evaluation primarily focuses on the Supernatural Instructions benchmark. While this dataset covers a wide range of tasks, it may not fully represent the diversity of real-world applications. Evaluating MoDE on other multi-task benchmarks or in specific domains could further validate its effectiveness and generalizability. Addressing these limitations could lead to even more efficient and adaptable multi-task LLM architectures, further expanding the potential of parameter-efficient fine-tuning for a wider range of applications.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28:41–75.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Preprint*, arXiv:2101.03961.
- Wenfeng Feng, Chuzhan Hao, Yuewei Zhang, Yu Han, and Hao Wang. 2024. Mixture-of-loras: An efficient

- multitask tuning for large language models. *arXiv preprint arXiv:2403.03432*.
- Markus Frohmann, Carolin Holtermann, Shahed Masoudian, Anne Lauscher, and Navid Rekabsaz. 2024. ScaLearn: Simple and highly parameter-efficient task transfer by learning to scale. In *Findings of the Association for Computational Linguistics: ACL 2024*.
- GemmaTeam, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. 2023. Lorahub: Efficient cross-task generalization via dynamic lora composition. *arXiv preprint arXiv:2307.13269*.
- 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, Giana Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2024. *Mixtral of experts*. *Preprint*, arXiv:2401.04088.
- Dengchun Li, Yingzi Ma, Naizheng Wang, Zhiyuan Cheng, Lei Duan, Jie Zuo, Cal Yang, and Mingjie Tang. 2024. Mixlora: Enhancing large language models fine-tuning with lora based mixture of experts. *arXiv preprint arXiv:2404.15159*.
- Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A. Smith, and Luke Zettlemoyer. 2022. *Branch-train-merge: Embarrassingly parallel training of expert language models*. *Preprint*, arXiv:2208.03306.
- Chu-Cheng Lin, Xinyi Wang, Jonathan H Clark, Han Lu, Yun Zhu, Chenxi Whitehouse, and Hongkun Yu. 2024a. Multitask multilingual model adaptation with featurized low-rank mixtures. *arXiv preprint arXiv:2402.17934*.
- Zhisheng Lin, Han Fu, Chenghao Liu, Zhuo Li, and Jianling Sun. 2024b. PEMT: Multi-task correlation guided mixture-of-experts enables parameter-efficient transfer learning. In *Findings of the Association for Computational Linguistics: ACL 2024*.
- Qidong Liu, Xian Wu, Xiangyu Zhao, Yuanshao Zhu, Derong Xu, Feng Tian, and Yefeng Zheng. 2023. Moelora: An moe-based parameter efficient fine-tuning method for multi-task medical applications. *arXiv preprint arXiv:2310.18339*.
- Yijiang Liu, Rongyu Zhang, Huanrui Yang, Kurt Keutzer, Yuan Du, Li Du, and Shanghang Zhang. 2024. *Intuition-aware mixture-of-rank-1-experts for parameter efficient finetuning*. *Preprint*, arXiv:2404.08985.
- OpenAI and Josh Achiam et al. 2024. *Gpt-4 technical report*. *Preprint*, arXiv:2303.08774.
- Oleksiy Ostapenko, Zhan Su, Edoardo Maria Ponti, Laurent Charlin, Nicolas Le Roux, Matheus Pereira, Lucas Caccia, and Alessandro Sordoni. 2024. Towards modular llms by building and reusing a library of lorae. *arXiv preprint arXiv:2405.11157*.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Viraj Shah, Nataniel Ruiz, Forrester Cole, Erika Lu, Svetlana Lazebnik, Yanzhen Li, and Varun Jampani. 2023. Ziplora: Any subject in any style by effectively merging lorae. *arXiv preprint arXiv:2311.13600*.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- Sainbayar Sukhbaatar, Olga Golovneva, Vasu Sharma, Hu Xu, Xi Victoria Lin, Baptiste Rozi  re, Jacob Kahn, Daniel Li, Wen tau Yih, Jason Weston, and Xian Li. 2024. *Branch-train-mix: Mixing expert llms into a mixture-of-experts llm*. *Preprint*, arXiv:2403.07816.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Yaqing Wang, Sahaj Agarwal, Subhabrata Mukherjee, Xiaodong Liu, Jing Gao, Ahmed Hassan Awadallah, and Jianfeng Gao. 2022a. Adamix: Mixture-of-adaptations for parameter-efficient model tuning. *arXiv preprint arXiv:2205.12410*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoor-molabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, et al. 2022b. Super-naturalinstructions: generalization via declarative instructions on 1600+ tasks. In *EMNLP*.

Xun Wu, Shaohan Huang, and Furu Wei. 2024. [Mixture of lora experts](#).

Ted Zadouri, Ahmet Üstün, Arash Ahmadian, Beyza Ermiş, Acyr Locatelli, and Sara Hooker. 2023. Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning. *arXiv preprint arXiv:2309.05444*.

Yun Zhu, Nevan Wichers, Chu-Cheng Lin, Xinyi Wang, Tianlong Chen, Lei Shu, Han Lu, Canoe Liu, Liangchen Luo, Jindong Chen, et al. 2023. Sira: Sparse mixture of low rank adaptation. *arXiv preprint arXiv:2311.09179*.

A Mutli-task performance - Ablation Study

We conducted an ablation study to investigate the impact of varying the number of experts (m) and the rank (r) of the LoRA matrices on MoDE’s performances and present the result in Table 6. The model configuration *MoDE* $m \times r$ represent MoDE models with different number of experts (m) and ranks (r). We explore the combinations 8×4 , 6×4 , 4×4 , 4×6 , 4×8 , and 4×16 .

	Eval	Add. Params.
MoDE 8×4	59.00	3.48%
MoDE 6×4	60.91	2.69%
MoDE 4×4	60.18	1.90%
MoDE 4×6	60.53	2.86%
MoDE 4×8	58.92	3.81%
MoDE 4×16	60.04	7.62%

Table 6: Ablation study on MoDE with varying number of ranks and experts. The evaluation metric used is ROUGE-L. The last column represents the total number of adapter parameters as a percentage of the total number of non-embedding parameters in Gemma-2B.

Number of Experts (m): Increasing the number of experts initially improves performance (MoDE 4×4 v.s. MoDE 6×4), suggesting that having more experts allows for better specialization. However, further increasing the number of experts to 8 or 16 (MoDE 8×4 or MoDE 16×4) does not lead to any improvement in performance, suggesting diminishing returns beyond a certain point.

Rank r : For a fixed number of experts (4), increasing the rank of the LoRA matrices from 4 to 6 (MoDE 4×4 vs. MoDE 4×6) results in a slight performance improvement (0.6018 vs. 0.6053 ROUGE-L). This suggests that higher rank matrices can capture more nuanced information, leading to better adaptation to different tasks. Further increasing the rank of both down-projection and up-projection matrices to 8 or 16 (MoDE 4×8 or MoDE 4×16) leads to a decrease in performance.

B Detailed Results for Case Study with Fixed Parameter Budget

Table 7 presents the detailed model performance with 12 model configures on the 15 tasks.

Task ID	LoRA		MoLORA			MoLORA-SD				MoDE			Full
	15×4	1×60	12×4	6×8	3×16	36×4	24×8	12×6	5×32	14×4	6×8	3×16	
24	25.91	25.30	26.67	26.16	26.29	26.81	27.26	27.12	27.71	27.40	27.31	27.55	27.36
25	41.27	41.35	41.7	41.43	41.99	41.91	41.73	41.92	42.36	42.29	41.92	42.97	42.04
74	42.73	42.53	43.12	42.86	42.64	42.67	42.88	43.53	43.30	42.34	43.35	43.03	43.08
89	33.90	38.21	64.41	63.64	57.24	71.42	71.19	64.71	61.48	77.73	76.43	78.20	70.31
114	84.15	87.85	91.54	89.85	90.15	91.08	92.00	90.92	91.38	90.77	90.62	90.62	91.14
141	75.79	82.87	94.00	91.08	91.85	94.15	94.31	93.95	93.49	96.00	95.33	94.92	94.62
155	51.09	53.88	62.42	59.16	56.99	61.18	61.80	54.35	60.71	63.82	65.37	63.66	61.21
192	72.47	77.21	83.35	80.41	80.86	82.99	82.68	84.47	83.61	84.78	83.42	85.69	83.79
269	75.40	75.51	75.46	75.56	75.50	75.61	75.52	75.33	75.72	75.59	75.71	75.52	75.57
279	69.90	82.59	89.52	88.39	88.44	90.60	89.68	90.45	91.42	91.53	91.73	91.58	90.96
291	79.97	85.14	86.64	86.31	85.31	86.31	86.31	85.81	85.81	87.48	87.65	86.81	86.61
622	99.95	99.70	99.97	99.93	99.96	99.96	99.93	99.93	99.94	99.96	99.94	99.93	99.94
672	39.69	45.08	48.92	46.46	45.85	48.31	49.69	51.69	48.62	50.46	51.85	50.31	50.46
1711	4.19	8.33	10.40	7.66	8.43	9.81	10.16	9.85	10.19	9.97	8.37	9.16	9.71
1729	17.21	17.11	17.43	16.77	17.34	17.25	17.35	17.60	17.68	17.31	17.75	17.61	17.54
Avg	53.83	57.07	61.93	60.63	60.16	62.25	62.42	61.66	61.79	63.39	63.37	63.39	62.53

Table 7: Performance comparison among various model configurations on 15 tasks with a fixed parameter budget. The scores in blue and green correspond to the highest and second-highest scores for the corresponding task.

C PCA Clustering of LoRA Matrices

Scatter plots after applying Principal Component Analysis (PCA) on all of the LoRA projection matrices at layer 6 and 12 are shown in Figure 6. The distinct grouping of down-projection vectors indicates common representations across tasks, providing the inspiration for the MoDE architecture.

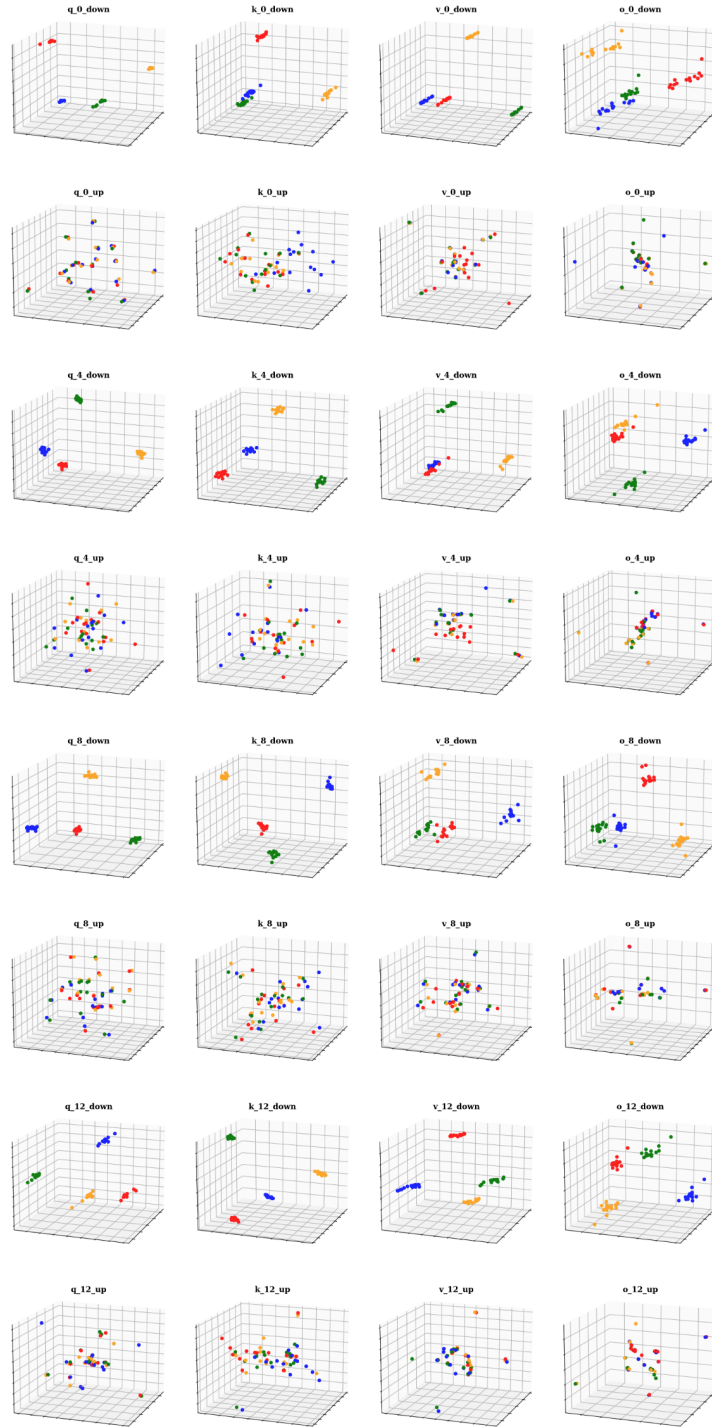


Figure 6: Scatter plots after applying Principal Component Analysis (PCA) on all LoRA projection matrices, i.e. query, key, value, and output, sliced along the rank dimension. The clear clustering of down-projection vectors suggests the presence of shared representations across tasks, motivating the design of the MoDE architecture.