

# Task-Based MoE for Multitask Multilingual Machine Translation

**Hai Pham**

Carnegie Mellon University  
htpham@cs.cmu.edu

**Young Jin Kim**

Microsoft  
youki@microsoft.com

**Subhabrata Mukherjee\***

Hippocratic AI

**David P. Woodruff**

Carnegie Mellon University  
dwoodruf@cs.cmu.edu

**Barnabás Póczos**

Carnegie Mellon University  
bapoczos@cs.cmu.edu

**Hany Hassan Awadalla**

Microsoft  
hanyh@microsoft.com

## Abstract

Mixture-of-experts (MoE) architecture has been proven a powerful method for diverse tasks in training deep models in many applications. However, current MoE implementations are task agnostic, treating all tokens from different tasks in the same manner. In this work, we instead design a novel method that incorporates task information into MoE models at different granular levels with shared dynamic task-based adapters. Our experiments and analysis show the advantages of our approaches over the dense and canonical MoE models on multitask multilingual machine translations. With task-specific adapters, our models can additionally generalize to new tasks efficiently.

## 1 Introduction

Mixture-of-Experts (MoE), while not being a novel machine learning algorithm (Yüksel et al., 2012), has revived to combine with deep learning, particularly transformer (Vaswani et al., 2017) and has recently pushed forward various tasks such as natural language processing, computer vision, speech recognition, multimodal and multitask learning due to its advantage in scalability in distributed environments (Fedus et al., 2022). The main advantages of MoE stem from its ensemble design while maintaining the sparsity in computation (Fedus et al., 2021). And with proper design such as using sharded experts (Lepikhin et al., 2020; Fedus et al., 2021), the possibility for enterprise-level scalability is almost boundless. As a result, this method has been more and more widely adopted in many applications that require distributed and intensive workloads.

However, most of the current methods are task-agnostic, only optimizing for performance based on lower levels in the architecture such as at system or communication levels (Rajbhandari et al., 2022). In the case of multitask learning where a

single model is required to learn from heterogeneous tasks, however, the task-specific data could be inherently diverse and vary largely from one to another (Wu et al., 2020). As a result, treating data from such different sources the same makes the learning ineffective, as also evidenced recently by the interference between different task data (Pfeiffer et al., 2022).

As a result, in this work, we design a novel MoE approach where task information is used during training and inference for assigning experts based on individual task information. The intuition is to make the training more task-aware so those similar tasks would be routed to the same group of experts and vice versa. From the architectural perspective, we incorporate high-level application-specific information with the system-level information to make the model become task-aware and hence have a better strategy in allocating experts based on the characteristics of distinct tasks, as also illustrated in Figure 1.

Our proposed architecture allows for grouping experts based on the similarity of tasks, i.e. similar tasks should use a similar group of experts and otherwise for different tasks, by using shared-task adapters. Our design of putting those adapters on top of MoE layers allows for flexibility in future extensions: if we want the model to acquire new tasks while still having similar resources, we only finetune new adapters, and if we want to scale the hardware resources, e.g. for more speed, we simply deal with MoE layers with such new resources.

Our experiments and analysis show the advantages of using task information in MoE architectures in multiple settings including multitask multilingual machine translations, as well as its generalization in few-shot learning. In summary, our contributions are as follows.

- First, we design novel MoE architectures that dynamically allocate experts based on task information in the context of multilingual mul-

\* Work done while at Microsoft (contact email: subhabrata.mukherjee.ju@gmail.com).

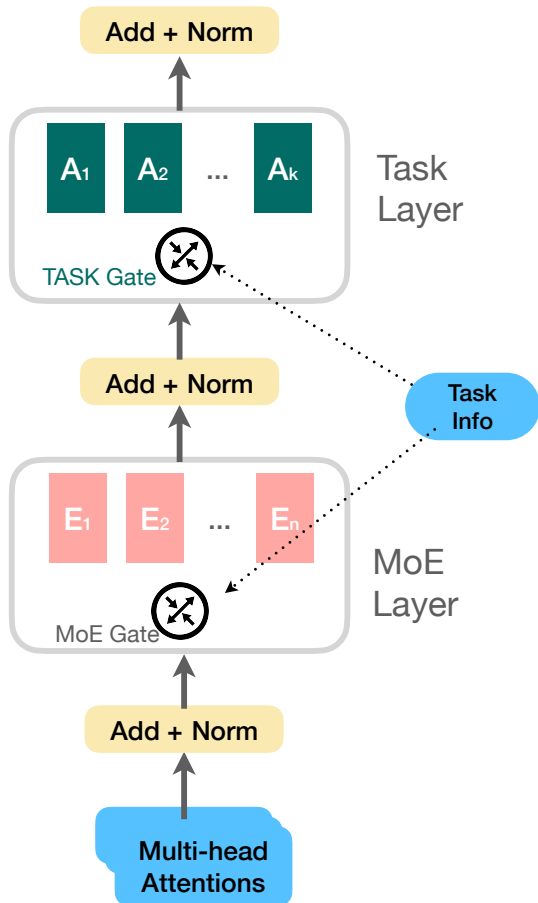


Figure 1: Extended from the typical MoE approaches that do not discriminate tokens from different tasks, we create shared task-related adapters that are trained to route tokens from similar tasks to the same shared adapters, and vice versa.

titask machine translation, with many variations.

- Second, we thoroughly study the pros and cons of our approaches in training from scratch, finetuning as well as transfer learning.
- Third, we implement our models on top of well-proven infrastructures for practicality and scalability including deepspeed (Rasley et al., 2020), fairseq (Ott et al., 2019) and transformer (Vaswani et al., 2017).

## 2 Related Work

**MoE Basic** Transformer-based Mixture-of-Experts (MoE) architecture essentially sparsifies transformer architecture by replacing the heavy feed-forward network (FFN) with a sparse MoE layer with top-1 or top-2 gates (Shazeer et al., 2017). However, increasing the number of experts

does not simply increase the performance (Fedus et al., 2021; Clark et al., 2022), many approaches have been proposed together to tackle the large-scale MoE deployment, such as in (Kim et al., 2021). In large-scale deployment, however, additional techniques should also be employed to battle with memory issues such as “sharding” experts (Lepikhin et al., 2020) or stabilizing the training (Zoph et al., 2022), since the models are often deployed on separate nodes that mainly used GPUs with limited memory. The architecture in this work inherits all of those techniques, and in addition incorporates task information into MoE routing, which in turn directs data into separate task adapters. This kind of routing is, however, hardware-agnostic, unlike some other work such as in (Zheng et al., 2022; Chen et al., 2023; Zeng and Xiong, 2023).

**MoE Routing Techniques** Gating is critical to MoE layer, which works as a weighted sum of the experts and serves for the ultimate purpose of load balancing of all available experts during both training and inference. Unlike the originally proposed top-k experts (Shazeer et al., 2017; Du et al., 2021), it was studied in SwitchTransformer that a single expert can preserve the quality if chosen properly, while significantly reducing the communication and computation cost (Fedus et al., 2021). In more detail, SwitchTransformer first divides evenly amongst all experts with an optional buffer for imbalanced cases and then applies an auxiliary loss to enforce load balancing. Another alternative approach, which is more computationally efficient is to get rid of such extra-heavy complicated loss and instead use a hash function to route every token to its matched expert, which tends to balance the output (Roller et al., 2021). Another interesting approach is to permit each token to appear in the top-k list of multiple experts (Zhou et al., 2022), which has been proven to help, although not applicable for auto-regressive applications. Yet because of the inherent problem of load imbalance, another approach is to replace the gating mechanism with a stochastic selection method, which randomly activates an input during processing (Zuo et al., 2021). The intuition is somewhat similar to the hash approach since it relies on the “fair” randomness to solve the balance problem while keeping the blueprint more lightweight than enforcing an auxiliary loss. Along similar lines, research directions have explored the random dropping of outputs from MoE layers (Liu et al., 2022; Elbayad

et al., 2022). Unlike all of those routing techniques which are application agnostic, our proposed model connects the application level (i.e. task information) with the lower-level MoE layers for better dealing with interference of different tasks in the context of multilingual multitask applications.

**Task-level Routing** Recently task information has been used for improving MoE, e.g. in (Liu et al., 2023). Our model is, however, much simpler and can be trained end-to-end unlike their approach, which requires clustering to be made for off-the-shelf shared representation learning. Probably the most related work to ours is Mod-Squad (Chen et al., 2022) which shares the motivation with us while having several differences. First, their approach has multiple aids to make the task-based MoE work with an additional loss for regularization, while we instead rely mainly on the simple motivation of incorporating task information into MoE. Second, we still stick to a single gate for routing, while they allocate multiple gates, each *per* task. Third, they additionally have MoE attention blocks, which make their architecture more complicated. Finally, our focused application is text-based machine translation, unlike computer vision settings in both works mentioned.

### 3 Models

Transformer architecture (Vaswani et al., 2017) has been proven to be the core backbone of the pervasive successes in natural language processing, computer vision, and other artificial intelligence fields. The main bottleneck to this architecture is, however, its heavy blueprint leading to intensive resources in training and inference, and is difficult to scale up. MoE is one powerful method to alleviate those problems in transformers.

#### 3.0.1 Sparse Mixture-of-Experts (MoE)

MoE, which was first introduced before the deep learning era (Jacobs et al., 1991), was recently borrowed to address those drawbacks in transformer architecture (Lepikhin et al., 2020). In a nutshell, MoE creates an ensemble of experts in multi-layer transformer blocks in place of a single expert, typically in the form of a feed-forward neural network (FFN) that is dense with many parameters.

In terms of formality, given an original FFN layer called  $\tilde{E}$ , we clone it into another layer containing a set of  $N$  experts with exactly the same architecture  $\{E_i\}_{i=1}^N$ . Likewise, the number of parameters for this particular layer is increased by a

factor of  $N$ .

The typical granular level of applying those experts in the context of natural language processing is the token level. Given a token  $x$ , its learned representation before MoE layer is a vector  $\mathbf{x}$ , its post-MoE output  $\mathbf{y}$  is the weighted average of those experts' output

$$o_i = E_i(\mathbf{x}) \quad (1)$$

$$\mathbf{y} = \sum_{i=1}^N W_i o_i, \quad (2)$$

where  $W_i$  is the weight (importance) of the corresponding expert  $E_i$ .

The key to MoE power and its well-proven successes in tandem with transformer architecture is its sparsity design: only one or few experts are activated (i.e. having non-zero weight) at any point in time in spite of many more parameters just introduced due to the ensemble. Typically the component responsible for this sparsity is a gate that was co-trained with experts to route tokens to their target expert(s), and eventually assigns only a single or few non-zero weights across all experts *per* token to its output  $G(\mathbf{x})$  typically using softmax and top-k method

$$g_{out} = \text{softmax}(W_g \mathbf{x}) \quad (3)$$

$$G(\mathbf{x}) = \text{Top\_K}(g_{out}) \quad (4)$$

With  $G(\mathbf{x})$  being a set of  $K$  chosen experts, equation 2 becomes

$$\mathbf{y} = \sum_{i \in G(\mathbf{x})} W_i o_i \quad (5)$$

The main architectural problem with this design is its scalability: the memory will be quickly used up as we increase experts, given the limitation of current compute resources allocated to a single compute node in any distributed environment. GShard (Lepikhin et al., 2020) was born to fix this issue by trading the memory for communication: allocating each expert to a single node and only aggregating them when needed, e.g. gradient averaging in training or weight averaging when saving a model. This elegant design has unlocked MoE's unlimited scalability and practicality in enterprise-level deployments, especially with the following-up work in optimizing for better architecture in

computation and communication, as mentioned in Section 2.

### 3.1 Task-based Adapters

Yet another problem on which we are focusing is not at the system level but more at the higher application level. As mentioned, in the multitask setting, the interference amongst tasks that are inherently different from each other could lead to the ineffectiveness of training. As a result, we employ *task-based adapters* to separate those different tasks into different adapters. Likewise, data (or tokens) from similar tasks should be routed to a similar group of adapters. There are three modes for those adapters.

First and the simplest is to allocate each adapter for each individual task. Although this setting is straightforward and requires no additional computation for data routing, it has the drawback of acquiring new unseen tasks. The reason is the model has to allocate a new adapter for each new task and fine-tune it with some amount of new data. Another potential problem is memory limitation if we want to extend to many new tasks in the future. This mode is called *static*, as shown in Figure 2a.

To enforce efficient learning of representation of similar task data, as well as alleviating memory problems, we have *dynamic* (Figure 2b) where the number of adapters is less than the number of tasks. As a result, we intentionally guide the model to learn better cross-task representation in terms of similarity and dissimilarity. In other words, data from similar tasks should be routed to the same adapters and vice versa. In practice, we choose the number of adapters to be  $\log_2(n)$  with  $n$  being the number of tasks.

### 3.2 Task-based Adapters with MoE

In this section, we formulate the task-based adapters mentioned in Section 3.1 in combination with MoE, both of which are our core architecture components.

Given  $M$  tasks, we allocate them into  $L$  shared-task adapters ( $L < M$ ). For every single token  $x$ , we have the associated task information  $t$  that makes up an input tuple  $(\mathbf{x}, \mathbf{t})$  per token. As before,  $\mathbf{x}$  is the representation vector from input, and  $\mathbf{t}$  is the task representation vector learned by task embedding.

Similar to MoE, we use a learnable task gate  $G_t$  that is responsible for this routing with input being the concatenation of the input components

$$G_t(x, t) = \text{Top\_K}(\mathbf{x} \oplus \mathbf{t}) \quad (6)$$

$$\mathbf{y} = \sum_{i \in G_t(\mathbf{x}, \mathbf{t})} W_i o_i \quad (7)$$

And since the number of adapters  $L < M$ , the number of tasks, we call this setting *dynamic*, as demonstrated in Figure 2b, as opposed to *static* (Figure 2a), where each task will go to each individual adapter.

Our main model uses the shared task embedding representation for the task gate as well as MoE gate, which we call *shared-dynamic*, as shown in Figure 2c.

## 4 Experiment Setup

### 4.1 Data

We tackle the problem of multitask multilingual machine translation using the data consisting of 10 different languages ranging from high-resource to low-resource ones including English (En), French (Fr), German (De), Czech (Cs), Finnish (Fi), Latvian (Lv), Estonian (Et), Romanian (Ro), Hindi (Hi), Turkish (Tr), and Gujarati (Gu). In more detail, the data for training, validation, and testing are listed in Table 1 where we can see besides the high-resource ones, we have low-resource languages such as Estonian, Hindi, or Gujarati.

Those data are in the form of Bitext in which there is always English. As a result, we denote EX as the translation from English (E) to another language (X), and similarly for the other way around, XE. Those data are populated from the popular WMT corpus<sup>1</sup>. For the given 1 English and 9 other languages, there are consequently 9 EX and 9XE tasks. More information about data can be found in Table 4 in Appendix A.

### 4.2 Task and Model Training

In this section, we describe the task information, evaluation metrics, and how we deal with data and models for training.

**Task** Our task is multitask multilingual machine translation (MMMT) which uses the EX and XE pairs. Our single model is trained with two main capacities. First, this single model can translate all the training pairs with high accuracy. Second, the model is able to quickly acquire new translation pairs with only zero or a few shots.

<sup>1</sup><https://www.statmt.org/wmt20/index.html>

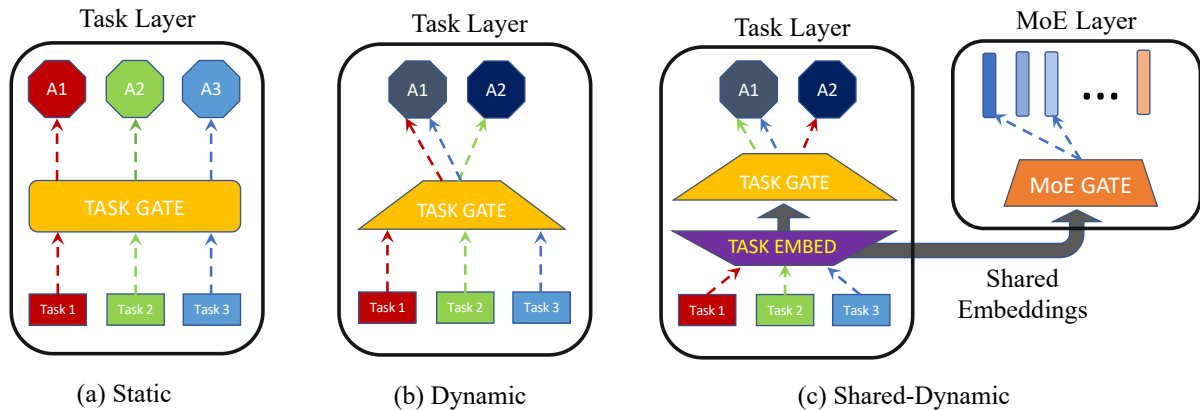


Figure 2: Our MoE models with variants. (a) Static means for each task, there is a separate adapter associated with it. (b) In the Dynamic mode, there is less number of adapters than the number of tasks, in order to learn the shared representation of similar tasks. (c) The last variant is Shared-Dynamic where the gates for task adapters and MoE share the same embedding for their routing decisions.

Split	Unit	Task Data								
		de-en	fr-en	cs-en	et-en	fi-en	gu-en	hi-en	lv-en	ro-en
Training	M	4.6	10	10.3	0.7	4.8	0.9	0.3	1.4	0.5
Validation	K	3.0	3.0	3.0	2.0	1.4	2.0	0.5	2.0	2.0
Testing	K	3.0	3.0	3.0	2.0	1.4	2.0	0.5	2.0	2.0

Table 1: Training, Validation, and Testing sizes for all XE tasks (the data for EX are exactly the same). Note that the unit for training is million (M) while that for both validation and testing are thousand (K), and the sizes are the same for validation and testing.

Model	XE Tasks										Avg
	de-en	fr-en	cs-en	et-en	fi-en	gu-en	hi-en	lv-en	ro-en		
1. Dense	29.9	31.2	28	22.4	21.4	22.3	21.4	24.5	36.1	26.4	
2. MoE Token	27.9	29.5	26.3	19.9	19.6	18.9	17.7	22.3	33.8	24.0	
3. MoE Sentence	27.9	29.9	26.2	21.4	19.9	17.9	15.9	23.2	34.4	24.1	
4. MoE Task-Static	<b>32.1</b>	<b>33.3</b>	<b>30.7</b>	<b>24.3</b>	<b>23.4</b>	20.6	<b>22.5</b>	<b>27.2</b>	<b>38.8</b>	<b>28.1</b>	
5. MoE Task-Dynamic	<b>31.4</b>	<b>32.0</b>	<b>29.1</b>	<b>23.4</b>	<b>22.1</b>	18.9	20.5	<b>25.5</b>	<b>37.2</b>	<b>26.7</b>	
Model	EX Tasks										Avg
	en-de	en-fr	en-cs	en-et	en-fi	en-gu	en-hi	en-lv	en-ro		
1. Dense	25.4	28.3	22.4	23.3	20.9	28.4	29.0	26.5	31.5	26.2	
2. MoE Token	22.9	25.1	19.5	20.1	17.9	26.2	26.3	24.0	29.0	23.4	
3. MoE Sentence	23.2	25.7	20.4	22.4	18.7	26.4	27.1	24.2	29.7	24.2	
4. MoE Task-Static	<b>29.5</b>	<b>32.5</b>	<b>27.9</b>	<b>27.4</b>	<b>25.8</b>	<b>28.8</b>	<b>30.8</b>	<b>32.2</b>	<b>34.6</b>	<b>29.9</b>	
5. MoE Task-Dynamic	<b>27.3</b>	<b>29.6</b>	<b>25.0</b>	<b>24.7</b>	<b>22.7</b>	27.7	<b>29.3</b>	<b>28.4</b>	<b>32.7</b>	<b>27.5</b>	

Table 2: Comparison of task-based MoE models (models 4 & 5) to task-agnostic MoE models (models 2 & 3) and the non-MOE (Dense) model (model 1) in BLEU scores. With the help of task information, task-based MoE models show their outperforming BLEU scores over all other types across most of the tasks including both high-resource and low-resource ones.

**Evaluation** While there are many evaluation metrics, we mainly use BLEU score due to its popularity and credibility in evaluating machine translation tasks. This evaluation is implemented by

SacreBLEU<sup>2</sup>. We note that, unlike all available public implementations that we found, our implementation evaluates all BLEU scores on the fly along with the training, so there is no disruption for

<sup>2</sup><https://github.com/mjpost/sacrebleu>

offline evaluation. That also helps in early stopping based on the BLEU scores on the validation sets.

**Pre-Processing and Post-Processing** In terms of preprocessing, we first encode the data using the Byte-Pair encoding (BPE) method and generate shared dictionaries where all the language pairs use the same vocabulary of size 64K, before feeding to the model. To get accurate scores, for post-processing, we again use BPE decoding for reconstructing the whole translated sentences before comparing them with the original sentences before BPE pre-processing. Likewise, we treat all the processing and model manipulation as a black box for calculating the scores.

#### Model Configuration and Implementation

We use transformer architecture (Vaswani et al., 2017) with 12 layers for both encoder and decoder phases, each of which uses a word embedding layer of dimension 1024 and a non-linear layer of dimension 4096. There are 16 attention heads and a dropout rate of 30%. For MoE, all jobs are trained on Azure cloud machines with 8 GPUs, each of which takes around 2 weeks for a model covering 18 aforementioned tasks to reach decent scores. We apply early stopping based on the validation BLEU scores, in which a non-increasing score after 2 epochs is the condition. For task-based information, we have a task embedding dimension of 64 and a task adapter hidden dimension of 256 for every single task adapter. Our implementation inherits the lower-level infrastructure code from Microsoft DeepSpeed and Fairseq.<sup>3</sup>

As for the implementation, an important practical issue with MoE is load balancing among experts for the best utilization of the infrastructure systems. For enforcing the training to have a balanced load, as a result, we employ the auxiliary loss from Lepikhin et al. (2020).

#### 4.2.1 Baselines

In order to show the performance of the task-based MoE models, the following baselines are selected:

**Dense** This is the traditional transformer model without any MoE layer, i.e., no change to the fully connected (FFN) layer in each layer of encoders or decoders.

**MoE - Token** This is the MoE model that is usually considered the default option where each FFN layer is replaced by an MoE layer. In our experiments, each MoE layer comprises 8 experts

(each has the same size as the original FFN being replaced) and a gate for routing purposes.

**MoE - Sentence** This is yet another MoE architecture with exactly the same architecture configuration as the MoE - Token baseline. The difference is in the routing layer, which functions at a different granularity: sentences instead of tokens. In more detail, while the gate decides which expert for each token separately in MoE - Token model, it instead routes all tokens belonging to a single sentence to the same chosen expert.

## 5 Results and Discussions

### 5.1 Multitask Multilingual Machine Translation

We first present the main results for models capable of translating 18 tasks (see Section 4.2) concurrently. As shown in Table 2, our models that incorporate MoE layers and are enhanced with task information show great advantages over all the baseline models on most tasks, in both directions EX and XE, in accordance with our hypothesis that using task adapters in conjunction with MoE is helpful in multilingual multitask translation.

An outstanding drawback with which the task-based MoE models are facing, however, is for the low-resource translation pairs, e.g. Gu-En, Hi-En, or En-Gu. As we can see from the results in Table 2, training those pairs with Dense models seems to benefit more than with MoE models. We hypothesize the problem is due to the undersampling of the training data for those languages, which have much less data than their high-resource counterparts. In more detail, our training routine concatenates all the tasks' data in a single big dataset before drawing batches. However, without adjusting the sampling process, high-resource language pairs are being trained significantly more given their much larger data in place. In particular, for the case of Gujarati where the Task-Dynamic MoE model underperforms in comparison to the baselines, our hypothesis is that linguistically, this language is the most different from all other languages, which makes the models very hard to learn effective shared representation with any other pairs.

In the future, we plan to explore ideas such as custom sampling or contrastive representation learning to tackle with such issues with the low-resource language pairs, in order to make MoE work as well for those languages as in high-resource pairs.

<sup>3</sup><https://github.com/facebookresearch/fairseq>

Model	Design		Routing		Tasks				Average
	MoE	Task	MoE	Task	de-en	fr-en	et-en	fi-en	
MoE	Y	N	Token	-	32.4	33.7	24.2	23.6	28.5
Dense + Task Static	N	Y	Task	Static	32.2	33.7	21.0	22.8	27.4
Dense + Task Dynamic				Dynamic	31.9	33.0	22.0	22.5	27.4
MoE + Task Static	Y	Y	Task	Static	30.7	32.0	19.9	20.8	25.9
MoE + Task Dynamic				Dynamic	<b>32.6</b>	<b>33.9</b>	24.0	<b>23.9</b>	<b>28.6</b>
MoE + Task Shared-Dynamic				Shared-Dynamic	32.2	33.3	<b>24.3</b>	<b>24.5</b>	<b>28.6</b>

Table 3: Performance of different models with changes on whether MoE layers exist, whether Task Adapters exist, and how routing for those components is undertaken. The scores better than the baseline are highlighted. Task-based MoE shows advantages, especially with shared-dynamic adapters between MoE and Task Adapters on the low-resource language pair.




Figure 3: Ablation study about merging 2 checkpointed models of different capabilities. Model 1 is trained with 4 tasks: de-en, fr-en, et-en and fi-en. Model 2 is trained with the other 4 tasks: cs-en, gu-en, en-et, and en-fi. Although those 2 models are under-trained with only a few thousand steps, in the merged model that has the capabilities of those two combined, many pairs have quickly picked up to a similar levels as in the previous single models.

## 5.2 Ablation Study

### 5.2.1 Implications of Different Task Layers and MoE Layers

In this study, we limit the number of tasks to four (De-En, Fr-En, Et-En, and Fi-En), which can be divided into 2 groups of similar tasks: (De-En, Fr-En) is the first group and (Et-En, Fi-En) is the second one, to study the performance implications of different model variants when there is a task layer and/or MoE layer.

As illustrated in Table 3, we again see that combining MoE and Task Adapters yields the best models, the same trend as shown in Table 2, particularly when the dynamic adapters are used to enforce similar tasks to share the same representations.

However, when task adapters are not used in conjunction with MoE, the performance is worse than MoE alone. This also means MoE should be the foundational infrastructure, and on top of that, task adapters should be used. It concurs with the motivation that the interference of different tasks or languages makes the training of experts difficult.

In other words, there is not so much help when there is only one expert for all the tasks (i.e. Dense models).

### 5.2.2 Flexibility of Task-based MoE in Merging Two Trained Models

One of the important capabilities in multitask learning and in general learning problems is how to quickly acquire new capabilities given current models with minimal resources and effort. Aligned with this goal, this ablation explores how quickly our task-based MoE models can be merged with each other from 2 different models to newly establish only 1 model that has the combination of their capabilities.

In merging those two models, we restore two respective checkpoints and merge layer-by-layer as follows. First, task-based adapters are kept and combined with each other: each model has 2 adapters (for 4 tasks in the model) and the combined model has 4 adapters (for 8 tasks in combination). Second, the task routers will also be merged and changed so that the routing of each

data will now have 4 selections instead of 2 outputs as in the previous models. Finally, the rest of the transformer and MoE layers will have their weights averaged.

The tasks in the original two models are hand-picked as in Section 5.2.1 to have 2 different groups, each of which has 2 similar tasks. Model 1 has de-en, fr-en, et-en, and fi-en, while Model 2 has cs-en, gu-en, en-et and en-fi.

As shown in Figure 3, while two original models have been trained with just a few thousand steps (a couple of hours), the combined model shows that it can quickly pick up their original capabilities with just a few hundred steps after merging. Although there are a few uncommon pairs that seem to fail, such as gu-en or en-et, the chart shows the optimistic result of combining trained models with our flexible task-based MoE architectures.

## 6 Conclusion

In the era of large language models, more efficient and effective modeling techniques are essential to, where MoE in combination with transformer-based models has proven its great advantages. It is, however, complicated to enable that implementation in practice due to the difficulties of training a single model serving diverse tasks. The proposed task-based MoE, which employs both task adapters in tandem with MoE has shown its promising advantages in the application of multitask multilingual machine translations. This novel design enforces shared representation of similar tasks and separates different task data to counter the interference effects. In addition, it also offers the flexibility of changing adapters based on new tasks or changing the MoE infrastructure without affecting the application level. Besides outperforming the traditional approaches using Dense models, however, our MoE models still need to improve on low-resource language pairs. To tackle that issue, in the future, exploring custom sampling for those pairs, and enforcing the shared representation learning explicitly using such additional techniques as contrastive learning or mutual information are worth exploring.

## 7 Acknowledgements

The authors would like to thank the great feedback and help from Yiren Wang, Muhammad ElNokrashy, Alex Muzio, Akiko Eriguchi and other members of Microsoft’s Machine Translation Group.

## References

- Chang-Qin Chen, Min Li, Zhihua Wu, Dianhai Yu, and Chao Yang. 2023. [Ta-moe: Topology-aware large scale mixture-of-expert training](#). *ArXiv*, abs/2302.09915.
- Z. Chen, Yikang Shen, Mingyu Ding, Zhenfang Chen, Hengshuang Zhao, Erik G. Learned-Miller, and Chuang Gan. 2022. [Mod-squad: Designing mixture of experts as modular multi-task learners](#). *ArXiv*, abs/2212.08066.
- Aidan Clark, Diego de Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake A. Hechtman, Trevor Cai, Sebastian Borgeaud, George van den Driessche, Eliza Rutherford, T. W. Hennigan, Matthew G. Johnson, Katie Millican, Albin Cassirer, Chris Jones, Elena Buchatskaya, David Budden, L. Sifre, Simon Osindero, Oriol Vinyals, Jack W. Rae, Erich Elsen, Koray Kavukcuoglu, and Karen Simonyan. 2022. [Unified scaling laws for routed language models](#). In *International Conference on Machine Learning*.
- Nan Du, Yanping Huang, Andrew M. Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, Barret Zoph, Liam Fedus, Maarten Bosma, Zongwei Zhou, Tao Wang, Yu Emma Wang, Kellie Webster, Marie Pellat, Kevin Robinson, Kathleen S. Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc V. Le, Yonghui Wu, Z. Chen, and Claire Cui. 2021. [Glam: Efficient scaling of language models with mixture-of-experts](#). *ArXiv*, abs/2112.06905.
- Maha Elbayad, Anna Sun, and Shruti Bhosale. 2022. [Fixing moe over-fitting on low-resource languages in multilingual machine translation](#). *arXiv preprint arXiv:2212.07571*.
- William Fedus, Jeff Dean, and Barret Zoph. 2022. [A review of sparse expert models in deep learning](#). *ArXiv*, abs/2209.01667.
- William Fedus, Barret Zoph, and Noam M. Shazeer. 2021. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *J. Mach. Learn. Res.*, 23:120:1–120:39.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. [Adaptive mixtures of local experts](#). *Neural Computation*, 3:79–87.
- Young Jin Kim, Ammar Ahmad Awan, Alexandre Muzio, Andrés Felipe Cruz-Salinas, Liyang Lu, Amr Hendy, Samyam Rajbhandari, Yuxiong He, and Hany Hassan Awadalla. 2021. [Scalable and efficient moe training for multitask multilingual models](#). *ArXiv*, abs/2109.10465.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam M. Shazeer, and Z. Chen. 2020. [Gshard: Scaling giant models with conditional computation and automatic sharding](#). *ArXiv*, abs/2006.16668.



- Rui Liu, Young Jin Kim, Alexandre Muzio, and Hany Hassan. 2022. [Gating dropout: Communication-efficient regularization for sparsely activated transformers](#). In *International Conference on Machine Learning*, pages 13782–13792. PMLR.
- Zhili Liu, Kai Chen, Jianhua Han, Lanqing Hong, Hang Xu, Zhenguo Li, and James Tin-Yau Kwok. 2023. [Task-customized masked autoencoder via mixture of cluster-conditional experts](#). In *International Conference on Learning Representations*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *North American Chapter of the Association for Computational Linguistics*.
- Jonas Pfeiffer, Naman Goyal, Xi Victoria Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe. 2022. [Lifting the curse of multilinguality by pre-training modular transformers](#). In *North American Chapter of the Association for Computational Linguistics*.
- Samyam Rajbhandari, Conglong Li, Zhewei Yao, Minjia Zhang, Reza Yazdani Aminabadi, Ammar Ahmad Awan, Jeff Rasley, and Yuxiong He. 2022. [Deepspeed-moe: Advancing mixture-of-experts inference and training to power next-generation ai scale](#). In *International Conference on Machine Learning*.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. [Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters](#). *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Stephen Roller, Sainbayar Sukhbaatar, Arthur Szlam, and Jason Weston. 2021. [Hash layers for large sparse models](#). In *Neural Information Processing Systems*.
- 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*.
- Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NIPS*.
- Sen Wu, Hongyang Zhang, and Christopher Ré. 2020. [Understanding and improving information transfer in multi-task learning](#). *ArXiv*, abs/2005.00944.
- Seniha Esen Yüksel, Joseph N. Wilson, and Paul D. Gader. 2012. [Twenty years of mixture of experts](#). *IEEE Transactions on Neural Networks and Learning Systems*, 23:1177–1193.
- Zhiyuan Zeng and Deyi Xiong. 2023. [Scomoe: Efficient mixtures of experts with structured communication](#). In *International Conference on Learning Representations*.
- Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yuanzhong Xu, Danyang Zhuo, Joseph Gonzalez, and Ion Stoica. 2022. [Alpa: Automating inter- and intra-operator parallelism for distributed deep learning](#). *ArXiv*, abs/2201.12023.
- Yan-Quan Zhou, Tao Lei, Han-Chu Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew M. Dai, Zhifeng Chen, Quoc V. Le, and James Laudon. 2022. [Mixture-of-experts with expert choice routing](#). *ArXiv*, abs/2202.09368.
- Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam M. Shazeer, and William Fedus. 2022. [St-moe: Designing stable and transferable sparse expert models](#).
- Simiao Zuo, Xiaodong Liu, Jian Jiao, Young Jin Kim, Hany Hassan, Ruofei Zhang, Tuo Zhao, and Jianfeng Gao. 2021. [Taming sparsely activated transformer with stochastic experts](#). *arXiv preprint arXiv:2110.04260*.

## A WMT Data Information

Code	Language	Test Split
de	German	wmt2013
fr	French	wmt2013
cs	Czech	wmt2013
et	Estonian	wmt2018dev
fi	Finish	wmt2015
gu	Gujarati	wmt2019dev
hi	Hindi	wmt2014dev
lv	Latvian	wmt2017dev
ro	Romanian	wmt2016dev

Table 4: More details about our datasets for comparison and reproducibility.