

# MoE-CT: A Novel Approach For Large Language Models Training With Resistance To Catastrophic Forgetting

Anonymous ACL submission

## Abstract

The advent of large language models (LLMs) has predominantly catered to high-resource languages, leaving a disparity in performance for low-resource languages. Conventional Continual Training (CT) approaches to bridge this gap often undermine a model’s original linguistic proficiency when expanding to multilingual contexts. Addressing this issue, we introduce a novel MoE-CT architecture, a paradigm that innovatively separates the base model’s learning from the multilingual expansion process. Our design freezes the original LLM parameters, thus safeguarding its performance in high-resource languages, while an appended MoE module, trained on diverse language datasets, augments low-resource language proficiency. Our approach significantly outperforms conventional CT methods, as evidenced by our experiments, which show marked improvements in multilingual benchmarks without sacrificing the model’s original language performance. Moreover, our MoE-CT framework demonstrates enhanced resistance to forgetting and superior transfer learning capabilities. By preserving the base model’s integrity and focusing on strategic parameter expansion, our methodology advances multilingual language modeling and represents a significant step forward for low-resource language inclusion in LLMs, indicating a fruitful direction for future research in language technologies.

## 1 Introduction

Large language models (LLMs) (Ouyang et al., 2022; Brown et al., 2020b) have achieved remarkable progress in recent years, particularly in areas such as language generation (Radford et al., 2019; Brown et al., 2020a), machine translation (Vaswani et al., 2017; Wu et al., 2019), text summarization (See et al., 2017; Rush et al., 2015), and language understanding (Devlin et al., 2019a; Peters et al., 2018). However, the majority of these models have




Figure 1: The abscissa represents the number of tokens from the original data incorporated during the Continual Training (CT) process. Although an increased volume of original data may decelerate the model’s forgetting, it can significantly impede the enhancement of multilingual capabilities.

focused on resource-rich languages such as English, leaving substantial potential for performance improvements in low-resource languages. To alleviate above mentioned issues, researchers have aimed to Parameter-Efficient Fine-Tuning (PEFT) methods. Specifically, continual training (CT) has been proposed, which have proven to be effective in enhancing the performance of low-resource languages by training on specific language data. Besides, other researchers have resorted to method of Low-Rank Adaptation (LoRA) (Hu et al., 2021). By introducing a low-rank structure to reduce the model parameters for efficient updating. Thereby, it achieves a successful success between maintaining model performance and saving computation and storage. Despite these developments, research on how to extend the multilingual capabilities of large models are still limited. We conducted an in-depth analysis of existing capability extension technologies.

1. Firstly, **it is difficult to obtain original training data**. In conventional training methods, the original training data accounts for a significant proportion, so that the lack of data poses a significant challenge for the training

067	in use. In addition, due to the different distributions	118
068	between the training data of different	119
069	stages, the absence of original training data	120
070	will significantly exacerbate the catastrophic	121
071	forgetting problems.	122
072	2. Secondly, <b>adding large amounts of raw data</b>	123
073	<b>will limit the improvement of multilingual</b>	124
074	<b>capabilities.</b> As illustrated in Figure 1, with	125
075	the volume of original language data greatly	126
076	alleviated the issue of catastrophic forgetting,	127
077	significantly improving the generation capa-	128
078	bility in original data distribution. When the	129
079	volume of original language data is five times	130
080	greater than oth new one, the catastrophic for-	131
081	getting issue is almost disappeared. However,	132
082	it limits the model’s final performance on mul-	133
083	tilingual tasks. Furthermore, a large amount of	134
084	original data significantly increases the train-	135
085	ing costs of the model.	136
086	In response to the aforementioned issues, we	137
087	propose a method that utilizes a Mixture of Ex-	
088	perts (MoE) (Fedus et al., 2021; Lepikhin et al.,	
089	2020)approach. The Mixture of Experts model inte-	
090	grates multiple experts into the model architecture,	
091	where each expert is tasked with learning a specific	
092	task or feature subspace, thereby enhancing the	
093	model’s learning capacity and generalization per-	
094	formance. Our specific approach involves extend-	
095	ing additional expert networks on top of the pre-	
096	trained model to more efficiently learn new multi-	
097	lingual capabilities. To effectively prevent catas-	
098	trophic forgetting, we froze most of the parameters	
099	of the original LLM. In addition, we employed a	
100	frozen shared feed-forward network (shared-ffn) to	
101	preserve the original knowledge, and implemented	
102	a gating mechanism to dynamically merge the orig-	
103	inal knowledge with the newly acquired knowledge	
104	from the expert networks. The crux of this method	
105	is that it enables us to effectively retain the capa-	
106	bilities of the base model, even with only a limited	
107	amount of original data. Our approach loosens the	
108	restrictions on the proportion of multilingual data,	
109	allowing it to play a more significant role during	
110	training. This not only raises the ceiling for the	
111	model’s performance on multilingual tasks but also	
112	reduces training costs due to a substantial reduction	
113	in the overall volume of data.	
114	In our proposed method, we selected Qwen as	
115	our foundational model and, based on this, ex-	
116	expanded the MoE network with 2-8 experts per layer.	
117	We extracted Chinese and English data, as well as	
	a substantial quantity of multilingual data, from	118
	Wikipedia and mC4 to serve as the continued train-	119
	ing data. Our multilingual capability enhancements	120
	were validated using datasets from the xwinograd	121
	(Tikhonov and Ryabinin, 2021), XCOPA (Ponti	122
	et al., 2020), and pasw-X (Yang et al., 2019) tasks	123
	designed for multilingual understanding, while	124
	the MMLU (Hendrycks et al., 2020)and C-Eval	125
	(Huang et al., 2023)datasets were employed to ver-	126
	ify the model’s retention of original capabilities.	127
	Empirical evidence demonstrates that our MoE ex-	128
	pansion strategy outperforms conventional contin-	129
	ual training (CT) and LoRA in terms of multilin-	130
	gual improvement. Moreover, in tests of original	131
	capabilities, our approach exhibits a notable in-	132
	crease in resistance to forgetting when compared	133
	to conventional CT. Our scheme has achieved com-	134
	mendable results across various sizes of the Qwen	135
	model series.	136
	To sum up, our contributions are as follows:	137
	• Our method introduces a novel MoE training	138
	paradigm for multilingual large model train-	139
	ing, eliminating the need for extensive pre-	140
	training data while achieving compatibility	141
	between original and multilingual capabilities.	142
		143
	• We employed additional expert networks to	144
	learn multilingual competencies, and by in-	145
	creasing the proportion of multilingual data,	146
	we have surpassed the upper limits of multi-	147
	lingual capabilities inherent in conventional	148
	CT methods.	149
	• Our approach leads the way in enhancing mul-	150
	tilingual abilities and resistance to forgetting	151
	when compared to standard CT, demonstrat-	152
	ing generalizability and transferability across	153
	models of varying scales.	154
	<b>2 Related Work</b>	155
	<b>Multilingual Pretraining and Fine-tuning in</b>	156
	<b>Language Models.</b> Developments in deep learning	157
	have significantly improved performance across	158
	a broad spectrum of natural language processing	159
	(NLP) tasks. Early efforts utilizing neural archi-	160
	tectures such as recurrent neural networks (RNNs)	161
	(Zaremba et al., 2014)and long short-term memory	162
	(LSTM) (Sutskever et al., 2011) networks have laid	163
	the foundation for understanding and generating	164

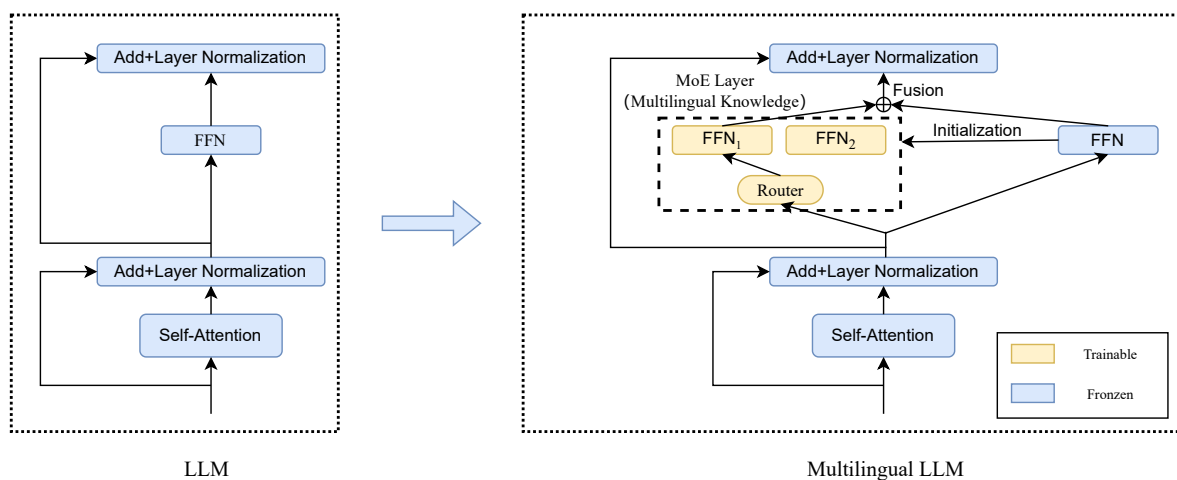


Figure 2: The diagram on the right represents the training structure of MoE-CT, where the blue area indicates that parameters are frozen, and the yellow area indicates that parameters are trainable. The parameters for all experts and the shared feed-forward network (shared-ffn) are initialized from the feed-forward network (ffn) of the original model.

language representations (Cho et al., 2014). The advent of self-attention mechanisms and transformer architectures (Vaswani et al., 2017) has shifted the focus towards models that can be pretrained on large unlabeled corpora for better generalization across NLP tasks. Recent research has expanded the scope of these models to accommodate multiple languages, yielding multilingual models such as mBERT (Devlin et al., 2019b) and XLM (Lample and Conneau, 2019), which are capable of learning cross-lingual representations. By pretraining on text from various languages, these models leverage shared linguistic properties and can be fine-tuned for downstream tasks in different languages.

**Sparsely Gated Networks** Efficiency in neural network scaling has gained paramount importance as the size of state-of-the-art language models has ballooned, resulting in steep computational costs. Sparsely gated networks have emerged as a promising direction for realizing model scaling without a linear increase in resource demands. At the core of this approach lies the concept of conditional computation, where different parts of the network are activated based on the input. The pioneering work of (Shazeer et al., 2017) demonstrated the potential of sparsely gated mixture-of-experts (MoE) layers in language models, where only a subset of experts is chosen for each input, drastically reducing computation during inference. This is achieved through gating mechanisms that learn to distribute the computation across a di-

verse set of experts, each specializing in different aspects of the data. Concurrently, a growing body of work is focusing on extending the Mixture of Experts (MoE) architecture with sparse activations (Hestness et al., 2017; Shazeer et al., 2018; Kudugunta et al., 2021). Switch Transformer (Fedus et al., 2021) and GLaM (Du et al., 2021) showcased a model with orders of magnitude more parameters than traditional models, but with fewer activated parameters per example, leading to improved performance and training efficiency. Similarly, the GShard framework (Lepikhin et al., 2020) demonstrated the scalability of MoE for large-scale multilingual machine translation tasks, reinforcing the viability of sparse gating for extensive language coverage.

**Continual Learning for Large Language models** Advancements in Natural Language Processing (NLP) have led to the exploration of various strategies to support continual learning capabilities. These strategies encompass: i) techniques that leverage the concept of replay to retain knowledge (ROBINS, 1995; Rebuffi et al., 2016; Shin et al., 2017; Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2018); ii) strategies grounded in regularization to prevent overwriting of existing information (Kirkpatrick et al., 2016; Li and Hoiem, 2016); and iii) designs that restructure the neural architecture itself to accommodate new learning (Rusu et al., 2016; Lee et al., 2017; Mallya and Lazebnik, 2017; Wen et al., 2020). The domain of NLP has

witnessed a burgeoning interest in these continual learning paradigms, as evidenced by a series of recent contributions (Wang et al., 2019a; Biesial-ska et al., 2020; Sun et al., 2019a; Huang et al., 2021; Hussain et al., 2021; Ahrens et al., 2021; Jin et al., 2021; Lin et al., 2022). These include embedding alignment with episodic memory (Wang et al., 2019b); improvements to memory-based parameter adaptation via sparse experience replay (de Masson d’Autume et al., 2019); approaches to lifelong language modeling (Sun et al., 2019b); and the adoption of meta-learning frameworks that integrate sparse replay (Holla et al., 2020). (Chen et al., 2023) marks the first application of the Mixture of Experts (MoE) structure to expand NLP tasks, mitigating the issue of catastrophic forgetting encountered in cross-task learning. Traditional approaches have predominantly focused on models of small to medium scale, where recent research has made strides in mitigating catastrophic forgetting to a certain degree. However, these strategies are often not applicable to large models with substantial parameter counts. In the continued training phase, such large models typically confront a more pronounced issue of catastrophic forgetting. Fine-tuning these models to adapt to specific tasks can lead to rapid loss of previously acquired general knowledge, even with minor parameter updates. In light of this, our work is dedicated to exploring how to balance the trade-off between stability (retaining old knowledge) and plasticity (acquiring new knowledge). By investigating and developing strategies that enable large models to maintain previously learned information while also learning new tasks effectively, we aim to contribute to the field’s understanding of how to construct artificial intelligence systems capable of continual learning without sacrificing their pre-trained knowledge.

### 3 Method

As discussed above, it is evident that the multilingual capabilities of large-scale models remain constrained. Our goal is to maintain a robust performance in the generation of high-resource languages, effectively mitigating the risk of catastrophic forgetting, while simultaneously enhancing the multilingual capabilities of the model.

To achieve this, we introduce a novel architectural paradigm named MoE-CT, which substantially integrates diverse knowledge domains into our model, by employing an MoE (Mixture of

Experts) approach tailored specifically for the multilingual context.

#### 3.1 Mixture of Experts

In the described model, we have a collection of  $N$  feed-forward neural networks, which are all structurally equivalent and operate independently, denoted as a set of experts  $E_i, i = 1^N$ . These experts are integrated with a dedicated gating mechanism, denoted as  $G(\cdot)$ , which functions as a decision-maker. This gating mechanism is responsible for determining the contribution of each expert’s output in the final response of the system. Specifically, if we consider  $h$  to be the resultant vector from an attention mechanism in any block of the model, then the final output  $y$  from the mixture of experts (MoE) layer is derived by taking a weighted sum of each expert’s output. Mathematically, this can be described by the equation:

$$y = \sum_i G(h)_i \cdot E_i(h)$$

In this expression,  $E_i(h)$  represents the output of the  $i$ -th expert network, while  $G(h)_i$  signifies the weight allocated to that output by the gating function. The gating function itself is defined using a softmax operation applied to the dot product of the input  $h$  and a set of learned parameters encapsulated in the matrix  $W_g$ , which is expressed as:

$$G(\cdot) = \text{Softmax}(h \cdot W_g)$$

In essence, the matrix  $W_g$  holds the parameters that the gating function adapts during training to optimize the routing of information through the various experts in the MoE architecture.

#### 3.2 Routing mechanism design

Our goal is to store multilingual knowledge in the MoE layer and store English/Chinese knowledge in the old FFN layer of pretrained LLM. We hope that the sparse extended multilingual LLM can combine the outputs of the two layers to achieve the best world of multilingual and English/Chinese abilities. Therefore, we have added a fusion module with different structures to explore the best combination of MoE and FFN layers. The output of the fusion module can be described as follows:

$$\text{Fusion}(x, y) = w \cdot x + (1 - w) \cdot y \quad (1)$$

In the equation,  $w$  represents the weight of the MoE layer, which, when combined with the weight of the shared-ffn, sums to 1.

	en	zh	ar	id	vi	de	fr	ja	th	pt	es	it
Tokens	2B	2B	2B	2B	2B	2B	2B	2B	2B	2B	2B	2B

Table 1: The composition of the continual training dataset, there are 24B tokens in total, and each language is allocated an equal proportion of the data.

Type	Expert	Params	Act-params	Layers	Hidden size	Heads
Dense	-	1.8B	1.8B	24	2048	16
MoE	2	3.4B	2.6B	24	2048	16
Dense	-	7B	7B	32	4096	32
MoE	2~8	16B~40B	11B	32	4096	32

Table 2: The size and architecture of our MoE model and dense model

$$output = Fusion\left\{\sum_{i=1}^K w_i FFN_i(h_{in}), FFN(h_{in})\right\} \quad (2)$$

where  $w_i$  is the weight of  $i$ -th expert determined by router, and  $K$  is the Top- $K$  experts selected by the router.

To prevent the problem of *catastrophic forgetting* during multilingual training, most of the parameters in our sparse-extended multilingual LLM will be frozen, and only the MoE layer, embedding layer and fusion module are trainable, as shown in Figure 2, the blue modules are frozen, and the yellow modules are trainable.

### 3.3 An innovative model training methodology

As shown in Figure 2, our method extends large language models from dense structures to sparse structures by adding MoE layer parallel to the Feed Forward Layer. The MoE layer includes multiple FFN layers, but only a small portion will be activated, and the activation strategy is determined by the router module. We use the parameters of FFN layer from pretrained LLM to initialize multiple FFN layers in the MoE layer, which can significantly accelerate the training process and provide valuable knowledge transfer from old one to new one..

Throughout the continuation training process, we exclusively train the expert networks and the embedding layer within the expanded MoE model architecture, while all other structural parameters and shared-FFN parameters are kept fixed. Through experimental validation, we have ascertained that such a training strategy optimally retains the model’s original capabilities while effectively expanding its multilingual capacity.

## 4 Experiment

### 4.1 Training Datasets

The composition of our data and the proportion of each language variety are reflected in Table 1. Our continued training data set comprises a total of 24 billion tokens, of which 20 billion are multilingual data encompassing ten languages (Arabic, Indonesian, Vietnamese, German, French, Japanese, Thai, Portuguese, Spanish, and Italian), with each language accounting for 2 billion tokens. Additionally, there are 2 billion tokens of Chinese data and 2 billion tokens of English data. All of the data were extracted from mC4 (Xue et al., 2020) and Wikipedia.

In addition, we prepared an extra 50 billion tokens for Chinese and 50 billion tokens for English to ensure that the proficiency in these languages does not diminish during the standard continual training (CT) process. Empirical evidence suggests that for the expanded multilingual dataset of 20 billion tokens, we need to incorporate at least five times more Chinese and English data to maintain their resistance to forgetting. However, within our Mixture of Experts (MoE) expanded architecture, we found that only 4 billion tokens of Chinese and English data are required to achieve resistance to forgetting in these languages, which substantially reduces the cost of training data.

### 4.2 Architecture Setting

We leverage Qwen (Bai et al., 2023), a large language model pre-trained using a large amount of Chinese and English data as our base model, exhibits industry-leading performance in both Chinese and English. Due to limitations in training resources, we employ the 1b8 and 7b versions of Qwen’s models as our foundational architectures.

Task category	Task	Test	Lang.	Metric	Prompt
NLU	XNLI	5,010	4	Acc.	[Premise], right? {Yes/Also/No}, [Hypothesis]
	XCOPA	500	7	Acc.	[Prefix] {because/therefore} {choice1/choice2} [Suffix]
	PAWS-X	2,000	7	Acc.	[Sentence1], right? {Yes/No}, [Sentence2]
Knowledge	C-Eval	1590	1	Acc.	[Question]{Choices}[Answer]
	MMLU	14213	1	Acc.	[Question]{Choices}[Answer]
MT	WMT/IWSLT	991-3,002	9	BLEU	[INPUT] Translate this sentence from [SRC] to [TGT].

Table 3: Multilingual benchmark

Model	Multilingual Ability							Chinese Ability	English Ability
	XCOPA	PAWS-X	XNLI	MT (en2xx)	MT (xx2en)	MT (en2zh)	MT (zh2en)	C-eval	MMLU
Qwen-1b8	54.7	53.0	38	15.1	21.0	35.6	22.1	48.5	37.0
Qwen-1b8-CT	62.3	51.0	44.7	19.1	20.9	27.2	15.2	29.7	29.5
Qwen-1b8-MoE-CT (Experts=1)	62.4	51.7	44.2	19.0	21.1	29.7	17.8	40.5	32.2
Qwen-1b8-MoE-CT (Experts=2)	62.5	52.9	44.4	19.5	21.9	30.8	18.5	41.5	33.4

Table 4: Main results on Qwen-1b8. When compared to Qwen-CT, the Qwen-MoE-CT exhibits a more robust resistance to forgetting and also demonstrates a greater enhancement in multilingual capabilities.

Table 2 shows the parameter structure of Qwen as a dense model, as well as the parameter structure after extending Qwen based on MoE. "Expert" represents the number of experts utilized in the MoE architecture. "Params" refers to the total number of parameters in the model. "Act-params" denotes the quantity of parameters activated during model inference. "Layers" indicates the number of layers in the model. "Hidden size" is the hidden dimension of the feed-forward layers, and "Heads" is the number of self-attention heads.

### 4.3 Multilingual Benchmark

We used six test datasets to evaluate the multilingual and Chinese/English capabilities of our models, and we divided them into three task categories: **NLU**, **Knowledge task** and **Machine Translation**. Please refer to Table 3 for detailed information.

**NLU** For Natural Language Understanding (NLU) task category, we choose three multilingual test task: XNLI, XCOPA and PAWS-X.

**Knowledge** The Knowledge task category includes two test dataset, C-Eval and MMLU, which are used to evaluate the Chinese and English capabilities of our models respectively.

**MT** For Machine Translation tasks, we selected 9 language test datasets from WMT and IWSLT to evaluate the translation ability from these languages to English and English to these languages.

### 4.4 Main Results

The main results are shown in Table 4 and Table 5, we conducted experiments on Qwen-1b8 and Qwen-7b models to verify the general effectiveness of our method. As shown in Table 4, Qwen-1b8-CT

is a multilingual continue-training version of Qwen-1b8, which has a significant improvements of multilingual abilities over Qwen-1b8 model, improves the accuracy of XCOPA from 54.7% to 62.3%, XNLI from 38.0% to 44.7%. For machine translation task, multilingual continue-training brings in an average of 4.0 BLEU on En-XX directions, but no improvement is found on XX-EN directions.

Although continue-training strategy can improve multilingual ability of LLM, it always causes problem of *catastrophic forgetting*, which damages the original abilities of the model. We can find that Qwen-1b8-CT model has a significant drop on English-to-Chinese and Chinese-to-English translation performance, from 35.6 to 27.2 and 22.1 to 15.2, respectively. For Chinese and English ability evaluation, Qwen-1b8-CT achieves a performance decline from 48.5 to 29.7 and 37.0 to 29.5 on C-Eval and MMLU test set, respectively. What's more, simply continue train the LLM on multilingual corpus may not improve multilingual abilities, as we find a performance decline on PAWS-X test dataset (53.0 vs 51.0). *Catastrophic forgetting* problem can also be found on Qwen-7b model in Table 5. We have also experimented with the use of LoRA-CT, setting the LoRA dimension to 8, as shown in Table 5, although LoRA-CT can effectively alleviate the issue of catastrophic forgetting, its limited number of changed parameters results in a significant performance gap in extended multilingual capabilities when compared to conventional CT.

By using the sparse MoE architecture for continual training of LLM, the best results can be achieved in both multilingual and Chinese English ability. In Table 5, we can find that Qwen-7b-

Model	Multilingual Ability							Chinese Ability	English Ability
	XCOPA	PAWS-X	XNLI	MT (en2xx)	MT (xx2en)	MT (en2zh)	MT (zh2en)	C-eval	MMLU
Qwen-7b	61.6	60.1	41.6	24.0	30.2	40.1	29.0	57.4	42.7
Qwen-7b-CT	69.8	60.1	46.8	27.5	30.3	37.7	26.8	50.1	39.2
Qwen-7b-LoRA-CT	62.3	59.6	44.0	24.4	29.1	40.0	27.8	55.4	43.3
Qwen-7b-MoE-CT (Experts=2)	69.5	60.6	46.0	27.8	30.3	39.3	27.5	53.6	41.7
Qwen-7b-MoE-CT (Experts=4)	69.3	60.5	46.0	28.3	30.5	39.8	27.7	55.9	42.1
Qwen-7b-MoE-CT (Experts=8)	69.6	60.6	46.5	28.4	30.7	39.8	28.1	55.5	42.3

Table 5: Main results on Qwen-7b. Our MoE architecture is equally applicable to various model sizes, and we have achieved substantial resistance to forgetting on Qwen-7b, alongside a comprehensive enhancement of multilingual capabilities that surpasses those of conventional CT and LoRA-CT.

MoE-CT performs equivalently to Qwen-7b-CT in multilingual tasks, and significantly prevents the performance degradation of Chinese and English capabilities, where performance on C-eval only decreases from 57.4 to 55.5 and MMLU decreases from 42.7 to 42.3. Our method also alleviates the decline in the performance of English-to-Chinese and Chinese-to-English translation, and has a significant improvement in the translation ability on English-XX directions from 24.0 to 28.4. The details of translation results on Qwen-1b8 and Qwen-7b can be found in Table 9 and Table 10. In our experiments, we have identified a significant data ratio problem when employing conventional continual training (CT) methods to combat catastrophic forgetting. Specifically, to preserve the proficiency in Chinese and English, the volume of data for these languages must be at least five times greater than that of the multilingual dataset. However, as indicated in Table 6, such an excessive reliance on Chinese and English data significantly hampers the enhancement of multilingual capabilities. Moreover, it substantially increases the training cost, posing a major obstacle to the expansion of large language models.

In contrast to the aforesaid conventional CT, our proposed MoE expansion technique significantly reduces the dependence on Chinese and English data, with these languages’ data requiring only a one-fifth proportion of the multilingual dataset. As demonstrated in Table 6, due to the reduced incorporation of Chinese and English data in our MoE framework, the model can more effectively assimilate multilingual knowledge. Consequently, the amplification of multilingual abilities is more pronounced compared to the conventional CT methods. This approach highlights the efficiency of our MoE expansion method in achieving a better and more efficient balance between preserving original language capabilities and enhancing expanded multilingual proficiencies.

## 4.5 Ablation Study

In this study, we conducted ablation experiments on the freezing strategy, routing mechanism, and the number of experts used in the MoE expanded architecture to verify their effects on multilingual improvement and resistance to forgetting of Chinese-English bilingual capabilities.

**Training strategy.** In our experiments on the Qwen-1.8B scale model, we attempted to freeze different parts of the MoE model to determine which components could retain the model’s original capabilities. As indicated in Table 7, continual training all parameters or attention layers will result in a significant decrease in the model’s original ability, while training expert layers and embedding layers can balance the original ability and expanded ability. Therefore, in the final model, we chose to train the parameters of the experts and the embedding layer.

**Routing mechanism.** In our experiments, we attempted various combination methods for the outputs of the shared feed-forward network and the MoE layer. Initially, we utilized a weighted sum approach for integrating the two, assigning MoE output weights from 0.1 to 0.9. Our experiments revealed that a lower MoE weight corresponded to a more pronounced resistance to catastrophic forgetting, yet resulted in a smaller improvement in multilingual capabilities. Conversely, a higher MoE weight weakened the resistance to catastrophic forgetting while yielding a greater enhancement in multilingual capabilities. As shown in table 8, when the weights for both the MoE and the shared-ffn are set to 0.5, an optimal balance is achieved between the enhancement of multilingual capabilities and the resistance to forgetting.

**The number of experts.** As shown in Table 4 and Table 5, two experts already exhibits considerable resistance to forgetting of Chinese-English bilingual capabilities. To verify whether a greater

Model	Continue training Tokens			Chinese Ability	English Ability	Multilingual Ability
	EN	ZH	Multilingual	C-eval	MMLU	Avg
Qwen-1b8	-	-	-	48.5	37	48.6
Qwen-1b8-CT	50B	50B	20B	39.5	31.3	50.2
Qwen-1b8-CT	2B	2B	20B	29.7	29.5	52.6
Qwen-1b8-MoE-CT	2B	2B	20B	<b>41.5</b>	<b>33.4</b>	<b>53.3</b>
Qwen-7b	-	-	-	57.4	42.7	54.4
Qwen-7b-CT	50B	50B	20B	53.3	41.1	57.6
Qwen-7b-CT	2B	2B	20B	50.1	39.2	58.9
Qwen-7b-MoE-CT	2B	2B	20B	<b>55.5</b>	<b>42.3</b>	<b>58.9</b>

Table 6: The data ratios required for conventional CT and those necessitated by the MoE expansion indicate that the MoE architecture can achieve better resistance to forgetting in Chinese and English, as well as enhanced multilingual capabilities, with only a minimal amount of Chinese and English data.

Trainable parameters	original	expanded
Qwen-1b8	42.7	48.6
all	30.8	52.8
attention	31.3	51.5
embedding	35.3	52.4
experts	35.9	52.7
embedding&experts	37.5	53.3

Table 7: The Impact of Training Strategies on Original and Expanded Capabilities in the MoE architecture, all means that all parameters participate in the training

MoE weights	original	expanded
0.1	37.8	52.1
Weighted 0.9	34.4	53.3
0.5	37.5	53.3

Table 8: The impact of different combinations of MoE layers and shared FFN on original and expanded capabilities in the Qwen-1B8 model, where MoE weight represents the proportion of the MoE layer’s output in the total output.

number of experts would yield further enhancements in multilingual capabilities, we expanded the model to include 4 and 8 experts at the Qwen-7B scale. However, according to the experimental results, increasing the number of experts does not significantly improve multilingual abilities. In light of this, we hypothesize that the quantity of continued training data may be insufficient, preventing the multiple experts from converging to their optimal performance. Consequently, we plan to utilize a larger corpus of continued training data, provided sufficient training resources, to support an increased number of experts in future experiments.

## 5 Conclusion

In conclusion, our research presents a significant advancement in the field of multilingual language modeling, particularly addressing the challenges posed by catastrophic forgetting in large language models (LLMs). Through the introduction of the MoE-CT structure, we have demonstrated a novel approach that not only enhances the extension of LLMs to low-resource languages but also preserves the original linguistic competencies in high-resource languages. Our experiments on models such as Qwen-1b8 and Qwen-7b have validated the effectiveness of MoE-CT, marking clear improvements in multilingual benchmarks while maintaining or even improving performance in the original languages.

The MoE-CT structure showcases a delicate balance between stability and plasticity, ensuring that the base model’s parameters remain undisturbed and that the newly introduced MoE layers absorb the additional linguistic knowledge. This balance is crucial for achieving a harmonious integration of multilingual capabilities without the detriment of pre-existing language proficiencies. Our findings indicate that MoE-CT can achieve substantial resistance to forgetting with a minimal amount of pre-training data, which is a considerable stride towards reducing training costs and resources.

The implications of our work are manifold. Primarily, it facilitates the creation of more inclusive language technologies that do not favor solely high-resource languages. Furthermore, it paves the way for future research into continual learning for LLMs, emphasizing the importance of models that can continually evolve and adapt to new language without losing previously established knowledge.

## 568 Limitations

569 In summary, we have proposed the MoE-CT archi-  
570 tecture to address the issue of catastrophic forget-  
571 ting encountered by LLMs during the expansion  
572 of multilingual capabilities. Due to the limitations  
573 of computational resources, we have not attempted  
574 to extend the MoE-CT architecture to other open-  
575 source models beyond Qwen, which may not fully  
576 demonstrate the catastrophic forgetting challenges  
577 faced by all LLMs. Therefore, our future work will  
578 explore whether this structure can be adapted to a  
579 wider range of open-source models.

## 580 Ethics Statement

581 Our work on the MoE-CT for LLMs considers sev-  
582 eral ethical concerns. Primarily, we aim to address  
583 linguistic biases by enhancing LLMs performance  
584 in low-resource languages, promoting inclusivity  
585 and cultural diversity. We recognize the risk of po-  
586 tential biases in model training and commit to their  
587 mitigation. We also acknowledge the responsibility  
588 to prevent the misuse of our model for deceptive  
589 purposes and advocate for its transparent and re-  
590 sponsible use. Environmental impacts due to the  
591 high computational requirements of LLMs are also  
592 considered. Our model aims to reduce training re-  
593 sources, and we encourage sustainable practices in  
594 AI research.

## 595 Acknowledgements

## 596 References

597 Kyra Ahrens, Fares Abawi, and Stefan Wermter. 2021.  
598 [DRILL: dynamic representations for imbalanced life-  
599 long learning](#). *CoRR*, abs/2105.08445.

600 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang,  
601 Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei  
602 Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin,  
603 Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu,  
604 Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren,  
605 Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong  
606 Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-  
607 guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang,  
608 Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu,  
609 Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingx-  
610 uan Zhang, Yichang Zhang, Zhenru Zhang, Chang  
611 Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang  
612 Zhu. 2023. [Qwen technical report](#).

613 Magdalena Biesialska, Katarzyna Biesialska, and  
614 Marta R. Costa-jussà. 2020. [Continual lifelong learn-  
615 ing in natural language processing: A survey](#). In  
616 *Proceedings of the 28th International Conference  
617 on Computational Linguistics*, pages 6523–6541,

Barcelona, Spain (Online). International Committee  
on Computational Linguistics. 618  
619

Tom Brown, Benjamin Mann, Nick Ryder, Melanie  
Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind  
Neelakantan, Pranav Shyam, Girish Sastry, Amanda  
Askell, Sandhini Agarwal, Ariel Herbert-Voss,  
Gretchen Krueger, Tom Henighan, Rewon Child,  
Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens  
Winter, Chris Hesse, Mark Chen, Eric Sigler, Ma-  
teusz Litwin, Scott Gray, Benjamin Chess, Jack  
Clark, Christopher Berner, Sam McCandlish, Alec  
Radford, Ilya Sutskever, and Dario Amodei. 2020a.  
[Language models are few-shot learners](#). In *Ad-  
vances in Neural Information Processing Systems*,  
volume 33, pages 1877–1901. Curran Associates,  
Inc. 620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie  
Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind  
Neelakantan, Pranav Shyam, Girish Sastry, Amanda  
Askell, Sandhini Agarwal, Ariel Herbert-Voss,  
Gretchen Krueger, Tom Henighan, Rewon Child,  
Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,  
Clemens Winter, Christopher Hesse, Mark Chen,  
Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin  
Chess, Jack Clark, Christopher Berner, Sam Mc-  
Candlish, Alec Radford, Ilya Sutskever, and Dario  
Amodei. 2020b. [Language models are few-shot learn-  
ers](#). *CoRR*, abs/2005.14165. 634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645

Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus  
Rohrbach, and Mohamed Elhoseiny. 2018. [Ef-  
ficient lifelong learning with A-GEM](#). *CoRR*,  
abs/1812.00420. 646  
647  
648  
649

Wuyang Chen, Yanqi Zhou, Nan Du, Yanping Huang,  
James Laudon, Zhifeng Chen, and Claire Cui. 2023.  
[Lifelong language pretraining with distribution-  
specialized experts](#). In *International Conference on  
Machine Learning, ICML 2023, 23-29 July 2023,  
Honolulu, Hawaii, USA*, volume 202 of *Proceedings  
of Machine Learning Research*, pages 5383–5395.  
PMLR. 650  
651  
652  
653  
654  
655  
656  
657

Kyunghyun Cho, Bart van Merriënboer, Çağlar  
Gülçehre, Fethi Bougares, Holger Schwenk, and  
Yoshua Bengio. 2014. [Learning phrase represen-  
tations using RNN encoder-decoder for statistical  
machine translation](#). *CoRR*, abs/1406.1078. 658  
659  
660  
661  
662

Cyprien de Masson d’Autume, Sebastian Ruder, Ling-  
peng Kong, and Dani Yogatama. 2019. [Episodic  
memory in lifelong language learning](#). In *Ad-  
vances in Neural Information Processing Systems*,  
volume 32. Curran Associates, Inc. 663  
664  
665  
666  
667

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and  
Kristina Toutanova. 2019a. [Bert: Pre-training of  
deep bidirectional transformers for language under-  
standing](#). In *North American Chapter of the Associa-  
tion for Computational Linguistics*. 668  
669  
670  
671  
672

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and  
Kristina Toutanova. 2019b. [BERT: Pre-training of](#) 673  
674

675	<a href="#">deep bidirectional transformers for language understanding</a> . In <i>Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)</i> , pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.	
676		
677		
678		
679		
680		
681		
682	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, Kathy Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc V. Le, Yonghui Wu, Zhifeng Chen, and Claire Cui. 2021. <a href="#">Glam: Efficient scaling of language models with mixture-of-experts</a> . <i>CoRR</i> , abs/2112.06905.	
683		
684		
685		
686		
687		
688		
689		
690		
691		
692	William Fedus, Barret Zoph, and Noam Shazeer. 2021. <a href="#">Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity</a> . <i>CoRR</i> , abs/2101.03961.	
693		
694		
695		
696	Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. <a href="#">Measuring massive multitask language understanding</a> . <i>CoRR</i> , abs/2009.03300.	
697		
698		
699		
700	Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory F. Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. 2017. <a href="#">Deep learning scaling is predictable, empirically</a> . <i>CoRR</i> , abs/1712.00409.	
701		
702		
703		
704		
705	Nithin Holla, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. <a href="#">Meta-learning with sparse experience replay for lifelong language learning</a> . <i>ArXiv</i> , abs/2009.04891.	
706		
707		
708		
709	Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. <a href="#">Lora: Low-rank adaptation of large language models</a> . <i>arXiv preprint arXiv:2106.09685</i> .	
710		
711		
712		
713		
714	Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang, and Diyi Yang. 2021. <a href="#">Continual learning for text classification with information disentanglement based regularization</a> . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 2736–2746, Online. Association for Computational Linguistics.	
715		
716		
717		
718		
719		
720		
721		
722	Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. <a href="#">C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models</a> .	
723		
724		
725		
726		
727		
728	Aman Hussain, Nithin Holla, and Pushkar Mishra. 2021. <a href="#">Towards a robust experimental framework and benchmark for lifelong language learning</a> . In <i>NeurIPS Datasets and Benchmarks</i> .	
729		
730		
731		
	Xisen Jin, Mohammad Rostami, and Xiang Ren. 2021. <a href="#">Lifelong learning of few-shot learners across nlp tasks</a> . <i>ArXiv</i> , abs/2104.08808.	732
		733
		734
	James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2016. <a href="#">Overcoming catastrophic forgetting in neural networks</a> . <i>CoRR</i> , abs/1612.00796.	735
		736
		737
		738
		739
		740
		741
	Sneha Kudugunta, Yanping Huang, Ankur Bapna, Maxim Krikun, Dmitry Lepikhin, Minh-Thang Luong, and Orhan Firat. 2021. <a href="#">Beyond distillation: Task-level mixture-of-experts for efficient inference</a> . <i>arXiv preprint arXiv:2110.03742</i> .	742
		743
		744
		745
		746
	Guillaume Lample and Alexis Conneau. 2019. <a href="#">Cross-lingual language model pretraining</a> . <i>CoRR</i> , abs/1901.07291.	747
		748
		749
	Jeongtae Lee, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. 2017. <a href="#">Lifelong learning with dynamically expandable networks</a> . <i>CoRR</i> , abs/1708.01547.	750
		751
		752
	Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. <a href="#">Gshard: Scaling giant models with conditional computation and automatic sharding</a> . <i>CoRR</i> , abs/2006.16668.	753
		754
		755
		756
		757
		758
	Zhizhong Li and Derek Hoiem. 2016. <a href="#">Learning without forgetting</a> . <i>CoRR</i> , abs/1606.09282.	759
		760
	Bill Yuchen Lin, Sida I. Wang, Xi Victoria Lin, Robin Jia, Lin Xiao, Xiang Ren, and Wen tau Yih. 2022. <a href="#">On continual model refinement in out-of-distribution data streams</a> . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	761
		762
		763
		764
		765
	David Lopez-Paz and Marc’Aurelio Ranzato. 2017. <a href="#">Gradient episodic memory for continuum learning</a> . <i>CoRR</i> , abs/1706.08840.	766
		767
		768
	Arun Mallya and Svetlana Lazebnik. 2017. <a href="#">Packnet: Adding multiple tasks to a single network by iterative pruning</a> . <i>CoRR</i> , abs/1711.05769.	769
		770
		771
	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. <a href="#">Training language models to follow instructions with human feedback</a> .	772
		773
		774
		775
		776
		777
		778
		779
	Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. <a href="#">Deep contextualized word representations</a> . In <i>Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)</i> , pages 2227–2237,	780
		781
		782
		783
		784
		785
		786

787	New Orleans, Louisiana. Association for Computational Linguistics.	
788		
789	Edoardo Maria Ponti, Goran Glavaš, Olga Majewska,	
790	Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020.	
791	<a href="#">XCOPA: A multilingual dataset for causal common-</a>	
792	<a href="#">sense reasoning</a> . In <i>Proceedings of the 2020 Con-</i>	
793	<i>ference on Empirical Methods in Natural Language</i>	
794	<i>Processing (EMNLP)</i> , pages 2362–2376, Online. As-	
795	sociation for Computational Linguistics.	
796	Alec Radford, Jeff Wu, Rewon Child, David Luan,	
797	Dario Amodei, and Ilya Sutskever. 2019. <a href="#">Language</a>	
798	<a href="#">models are unsupervised multitask learners</a> .	
799	Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, and	
800	Christoph H. Lampert. 2016. <a href="#">icarl: Incremen-</a>	
801	<a href="#">tal classifier and representation learning</a> . <i>CoRR</i> ,	
802	abs/1611.07725.	
803	ANTHONY ROBINS. 1995. <a href="#">Catastrophic forgetting,</a>	
804	<a href="#">rehearsal and pseudorehearsal</a> . <i>Connection Science</i> ,	
805	7(2):123–146.	
806	Alexander M. Rush, Sumit Chopra, and Jason Weston.	
807	2015. <a href="#">A neural attention model for abstractive sen-</a>	
808	<a href="#">tence summarization</a> . In <i>Proceedings of the 2015</i>	
809	<i>Conference on Empirical Methods in Natural Lan-</i>	
810	<i>guage Processing</i> , pages 379–389, Lisbon, Portugal.	
811	Association for Computational Linguistics.	
812	Andrei A. Rusu, Neil C. Rabinowitz, Guillaume	
813	Desjardins, Hubert Soyer, James Kirkpatrick, Ko-	
814	ray Kavukcuoglu, Razvan Pascanu, and Raia Had-	
815	sell. 2016. <a href="#">Progressive neural networks</a> . <i>CoRR</i> ,	
816	abs/1606.04671.	
817	Abigail See, Peter Liu, and Christopher Manning.	
818	2017. <a href="#">Get to the point: Summarization with pointer-</a>	
819	<a href="#">generator networks</a> . In <i>Association for Computa-</i>	
820	<i>tional Linguistics</i> .	
821	Noam Shazeer, Youlong Cheng, Niki Parmar, Dustin	
822	Tran, Ashish Vaswani, Penporn Koanantakool, Peter	
823	Hawkins, Hyoungho Lee, Mingsheng Hong, Cliff	
824	Young, Ryan Sepassi, and Blake A. Hechtman. 2018.	
825	<a href="#">Mesh-tensorflow: Deep learning for supercomputers</a> .	
826	<i>CoRR</i> , abs/1811.02084.	
827	Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz,	
828	Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and	
829	Jeff Dean. 2017. <a href="#">Outrageously large neural networks:</a>	
830	<a href="#">The sparsely-gated mixture-of-experts layer</a> . <i>CoRR</i> ,	
831	abs/1701.06538.	
832	Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon	
833	Kim. 2017. <a href="#">Continual learning with deep generative</a>	
834	<a href="#">replay</a> . <i>CoRR</i> , abs/1705.08690.	
835	Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee.	
836	2019a. <a href="#">Lamol: Language modeling for lifelong lan-</a>	
837	<a href="#">guage learning</a> . <i>arXiv preprint arXiv:1909.03329</i> .	
838	Fan-Keng Sun, Cheng-Hao Ho, and Hung yi Lee. 2019b.	
839	<a href="#">Lamol: Language modeling for lifelong language</a>	
840	<a href="#">learning</a> . In <i>International Conference on Learning</i>	
841	<i>Representations</i> .	
	Ilya Sutskever, James Martens, and Geoffrey E Hinton.	842
	2011. <a href="#">Generating text with recurrent neural networks</a> .	843
	In <i>Proceedings of the 28th international conference</i>	844
	<i>on machine learning (ICML-11)</i> , pages 1017–1024.	845
	Alexey Tikhonov and Max Ryabinin. 2021. <a href="#">It’s All in</a>	846
	<a href="#">the Heads: Using Attention Heads as a Baseline for</a>	847
	<a href="#">Cross-Lingual Transfer in Commonsense Reasoning</a> .	848
	In <i>Findings of the Association for Computational</i>	849
	<i>Linguistics: ACL-IJCNLP 2021</i> , pages 3534–3546,	850
	Online. Association for Computational Linguistics.	851
	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob	852
	Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz	853
	Kaiser, and Illia Polosukhin. 2017. <a href="#">Attention is all</a>	854
	<a href="#">you need</a> . <i>CoRR</i> , abs/1706.03762.	855
	Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo,	856
	Shiyu Chang, and William Yang Wang. 2019a. <a href="#">Sen-</a>	857
	<a href="#">tence embedding alignment for lifelong relation ex-</a>	858
	<a href="#">traction</a> . In <i>Proceedings of the 2019 Conference of</i>	859
	<i>the North American Chapter of the Association for</i>	860
	<i>Computational Linguistics: Human Language Tech-</i>	861
	<i>nologies, Volume 1 (Long and Short Papers)</i> , pages	862
	796–806, Minneapolis, Minnesota. Association for	863
	Computational Linguistics.	864
	Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo,	865
	Shiyu Chang, and William Yang Wang. 2019b. <a href="#">Sen-</a>	866
	<a href="#">tence embedding alignment for lifelong relation ex-</a>	867
	<a href="#">traction</a> . In <i>Proceedings of the 2019 Conference of</i>	868
	<i>the North American Chapter of the Association for</i>	869
	<i>Computational Linguistics: Human Language Tech-</i>	870
	<i>nologies, Volume 1 (Long and Short Papers)</i> , pages	871
	796–806, Minneapolis, Minnesota. Association for	872
	Computational Linguistics.	873
	Yeming Wen, Dustin Tran, and Jimmy Ba. 2020.	874
	<a href="#">Batchensemble: An alternative approach to ef-</a>	875
	<a href="#">ficient ensemble and lifelong learning</a> . <i>CoRR</i> ,	876
	abs/2002.06715.	877
	Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le,	878
	Mohammad Norouzi, Wolfgang Macherey, Maxim	879
	Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al.	880
	2019. <a href="#">Google’s neural machine translation system:</a>	881
	<a href="#">Bridging the gap between human and machine trans-</a>	882
	<a href="#">lation</a> . arxiv 2016. <i>arXiv preprint arXiv:1609.08144</i> ,	883
	2.	884
	Linting Xue, Noah Constant, Adam Roberts, Mihir Kale,	885
	Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and	886
	Colin Raffel. 2020. <a href="#">mt5: A massively multilingual</a>	887
	<a href="#">pre-trained text-to-text transformer</a> . In <i>North Amer-</i>	888
	<i>ican Chapter of the Association for Computational</i>	889
	<i>Linguistics</i> .	890
	Yinfei Yang, Yuan Zhang, Chris Tar, and Jason	891
	Baldrige. 2019. <a href="#">PAWS-X: A cross-lingual adversar-</a>	892
	<a href="#">ial dataset for paraphrase identification</a> . In <i>Proceed-</i>	893
	<i>ings of the 2019 Conference on Empirical Methods</i>	894
	<i>in Natural Language Processing and the 9th Inter-</i>	895
	<i>national Joint Conference on Natural Language Pro-</i>	896
	<i>cessing (EMNLP-IJCNLP)</i> , pages 3687–3692, Hong	897
	Kong, China. Association for Computational Linguis-	898
	tics.	899

900 Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals.  
901 2014. [Recurrent neural network regularization](#).  
902 *CoRR*, abs/1409.2329.

## 903 **A Appendix A**

904 In this section, we provide the detailed information  
905 about translation results on Qwen-1b8 and Qwen-  
906 7b models. The results are shown in Table 9 and  
907 Table 10.

Model	En-X								X-EN							
	Ar	De	Es	Fr	Ru	Ko	Zh	Ja	Ar	De	Es	Fr	Ru	Ko	Zh	Ja
Qwen-1b8	1.4	16.9	22.2	22.4	9.4	4.0	35.6	8.8	15.9	30.7	28.2	30.8	25.8	6.2	22.1	8.0
Qwen-1b8-CT	7.3	21.3	25.1	26.7	12.4	11.9	27.2	20.5	24.4	31.0	28.9	31.0	23.2	6.5	15.2	7.2
Qwen-1b8-MoE-CT (Experts=1)	7.1	20.7	24.9	26.0	11.8	12.6	29.7	18.8	22.8	31.3	28.5	30.8	24.0	6.0	17.8	7.9
Qwen-1b8-MoE-CT (Experts=2)	7.5	21.0	25.3	26.5	12.3	12.1	30.8	20.2	24.8	31.7	29.2	31.2	26.1	6.5	18.5	7.4

Table 9: Translation results on Qwen-1b8 in details.

Model	En-X								X-EN							
	Ar	De	Es	Fr	Ru	Ko	Zh	Ja	Ar	De	Es	Fr	Ru	Ko	Zh	Ja
Qwen-7b	7.1	27.7	30.6	35.9	16.9	13.4	40.1	20.0	31.5	41.3	34.7	39.6	34.3	14.0	29.0	17.2
Qwen-7b-CT	11.7	30.0	32.2	36.3	19.1	21.0	37.7	31.6	35.2	40.9	34.4	39.0	33.6	15.8	26.8	16.3
Qwen-7b-MoE-CT (Experts=2)	11.6	30.0	32.3	36.8	18.9	21.7	39.3	31.9	35.4	41.1	34.5	39.3	34.0	14.3	27.5	16.6
Qwen-7b-MoE-CT (Experts=4)	11.6	30.6	32.7	37	19.4	23.0	39.8	32.2	34.7	41.2	34.4	39.5	34.0	14.9	27.7	17.6
Qwen-7b-MoE-CT (Experts=8)	11.8	30.5	32.8	37.2	19.3	23.5	39.8	32.5	35.4	41.1	34.4	39.2	34.1	15.5	28.1	18.0

Table 10: Translation results on Qwen-7b in details.