

Extracting and Transferring Abilities For Building Multi-lingual Ability-enhanced Large Language Models

Anonymous ACL submission

Abstract

Multi-lingual ability transfer has become increasingly important for the broad application of large language models (LLMs). Existing work highly relies on training with the multi-lingual ability-related data, which may be not available for low-resource languages. To solve it, we propose a **Multi-lingual Ability Extraction and Transfer** approach, named as **MAET**. Our key idea is to decompose and extract language-agnostic ability-related weights from LLMs, and transfer them across different languages by simple addition and subtraction operations without training. Specially, our MAET consists of the extraction and transfer stages. In the extraction stage, we firstly locate key neurons that are highly related to specific abilities, and then employ them to extract the transferable ability-specific weights. In the transfer stage, we further select the ability-related parameter tensors, and design the merging strategy based on the linguistic and ability specific weights, to build the multi-lingual ability-enhanced LLM. To demonstrate the effectiveness of our proposed approach, we conduct extensive experiments on mathematical and scientific tasks in both high-resource lingual and low-resource lingual scenarios. Experiment results have shown that MAET can effectively and efficiently extract and transfer the advanced abilities, and outperform training-based baselines methods. Our code and data will be publicly released.

1 Introduction

Large language models (LLMs) have recently shown remarkable performance on various general tasks, *e.g.*, text generation and question answering (Zhao et al., 2023; OpenAI, 2023; Dubey et al., 2024). Despite the success, LLMs are still struggling to solve complex tasks (*e.g.*, mathematical reasoning), which require LLMs to possess specific advanced abilities (*e.g.*, deductive reasoning) and knowledge (*e.g.*, mathematical theory) (Yue et al.,

2024; Lu et al., 2022). To address it and further improve LLMs, existing work either collects the related data to train LLMs (Du et al., 2024; Chen et al., 2024a), or merges the parameters of existing well-performed LLMs to transfer their advanced abilities into one single model (Ilharco et al., 2023; Yadav et al., 2023; Yu et al., 2024a).

Despite the success, it is not easy to collect sufficient training corpus or well-trained LLMs related to specific abilities, especially in multi-lingual scenarios. Especially, some popular languages (*e.g.*, English) have dominated the linguistic expressions of the open web data, and the amount of available domain-specific data for low-resource languages (*e.g.*, Bengali or Telugu) is highly limited (Magueresse et al., 2020; Patzelt, 2024; Mirashi et al., 2024). Fortunately, existing work (Zhao et al., 2024; Schäfer et al., 2024) has revealed that the learned knowledge from one language by LLMs could be inherited and leveraged by other languages. For example, Llama-series LLMs are trained mainly on English texts, while they can also solve the tasks based on other languages. Such a finding has been widely explored in either improving the overall performance of multi-lingual LLMs (Schäfer et al., 2024) or enhancing fine-grained knowledge (Chen et al., 2024a). However, the related work mostly relies on training with ability-related corpus in the target language, which is not always available for low-resource languages.

To conduct a more effective ability transfer, our idea is to learn and extract the “*ability-specific weights*” that preserves the knowledge about specific abilities for the LLM. If such ability-specific and language-specific weights could be decomposed, it is achievable to transfer the required abilities into target languages by just combining the corresponding weights, even building a multi-lingual ability-enhanced LLM like building blocks. Based on this idea, in this paper, we propose a **Multi-lingual Ability Extraction and Transfer** approach,

named as **MAET**. Specifically, our approach consists of two major stages, *i.e.*, ability extraction and transfer stage. In the extraction stage, we first locate the abilities-related neurons and leverage related corpus in a reference language to continually pre-train the LLM on these identified neurons. Then, based on the LLM trained on the general corpus, we devise the formula to extract the ability-specific weights. In the transfer stage, we utilize the ability-related weights to select related parameter tensors, and design a specific model merging strategy by interpolating linguistic and ability-specific weights. In our approach, we only need ability-specific corpus from any rich-resource language and general multi-lingual corpus, which can effectively mitigate the data scarcity issues in low-resource languages.

To assess the effectiveness of our approach, we conduct the evaluation based on two comprehensive reasoning benchmarks, namely Multi-lingual Grade School Math (MGSM) (Shi et al., 2023) and science tasks from multi-lingual MMLU (Lai et al., 2023) as the evaluation benchmarks. According to the evaluation results, the proposed approach MAET outperforms other competitive baseline methods (*e.g.*, continual pre-training (Gururangan et al., 2020) and model merging methods with task vectors (Ilharco et al., 2023), achieving the 9.1% relative improvement compared to the base LLM. In addition, our approach can work well with relatively fewer training data, demonstrating an improved efficiency in practice.

2 Related Work

We introduce the related work from the following three perspectives:

Continual Pre-training. Although LLMs have shown remarkable performance on various downstream work, they still struggle in several specific tasks, *e.g.*, complex reasoning tasks (Paster et al., 2024; Shao et al., 2024) or low-resource lingual scenarios (Hedderich et al., 2021; Panchbhai and Pankanti, 2021). To adapt LLMs pre-trained on the general corpus to multi-lingual scenarios or specific tasks, existing work (Luo et al., 2022; Taylor et al., 2022; Zhao et al., 2022; Zhang et al., 2024a) has collected the corresponding corpus to continually pre-train (CPT) LLMs. During the continual pre-training process, the mixture strategy between the general corpus and the ability-related corpus should be carefully considered to prevent hurting

the general abilities of LLMs (Ye et al., 2024; Xie et al., 2023; Chen et al., 2024a; Siriwardhana et al., 2024). However, previous study (Chang et al., 2024; Lu et al., 2023) has pointed out that it is difficult to collect the required corpus, especially for low-resource language corpus. Therefore, synthesizing data from powerful LLMs is widely utilized to expand the task-specific training corpus (Chen et al., 2021b; Yu et al., 2024b; Zhou et al., 2024a). Besides, because of the limitation of computation resources, a series of approaches (Hu et al., 2022; Li and Liang, 2021; Dettmers et al., 2023) only train several parameters to reduce the expenses. In this work, we focus on adapting LLMs to multi-lingual complex reasoning scenarios through continually pre-training LLMs on the single-lingual task-specific corpus.

Knowledge Editing. According to lottery ticket hypothesis (Frankle and Carbin, 2019), training a small number of model parameters will achieve comparable or even better performance on downstream tasks. Existing study (Du et al., 2024; Wang et al., 2024b; Gong et al., 2024) has leveraged the inner information of LLMs, *e.g.*, gradient or cosine similarity between different hidden states, to select and train the related sub-network. Besides, the probe (*i.e.*, a newly initialized parameter) can be implemented to detect the knowledge of LLMs and process targeted repair (Wang et al., 2024a; Jiang et al., 2024). However, Since these approaches need to calculate and select a sub-network of each training instance, which might cause the instability of the training process, several study (Chen et al., 2024b; Zhang et al., 2024b) pointed out that the task-related sub-network can be determined before the training process, and only updating the value of the corresponding neurons can achieve better performance. In this work, we focus on editing the task-specific neurons of LLMs to improve the corresponding capacities in multi-lingual tasks.

Model Merging. Given that the CPT process will bring huge computational expenses, previous work leveraged model merging techniques to integrate different abilities (*e.g.*, mathematical reasoning and code synthesizing) into one model (Yang et al., 2024; Xu et al., 2024b; Stoica et al., 2024). During the merging process, the interference between different LLMs might be conflict with each other and affect the final performance. Therefore, researchers proposed the clip (Yadav et al., 2023) or

randomly dropout (Yu et al., 2024a) mechanism to mitigate the performance decrease. Moreover, the selection of the hyper-parameters (*e.g.*, weight of each model) is the challenge of the model merging process, and previous work (Zhou et al., 2024b; Matena and Raffel, 2022) utilized the inner parameters of LLMs or external matrixes to determine the hyper-parameters. Furthermore, a series of work has studied improving the reasoning ability of LLMs in non-English scenarios by merging the reasoning-specialized model and multi-lingual model (Huang et al., 2024; Yoon et al., 2024). Inspired by the above work, we try to locate the task-related sub-networks of LLMs and transfer the advanced abilities.

3 Preliminary

Despite that LLMs exhibit remarkable performance on general tasks, they still have limited advanced abilities, *e.g.*, mathematical and scientific reasoning abilities. A typical approach to enhance these abilities is to continually pre-train (CPT) LLMs with ability-related corpus. However, such training data might not always be available or sufficient, especially for minor domains (*e.g.*, Bengali). In this work, we focus on the task of *ability extraction and transfer* by continual pre-training and merging LLMs. Concretely, LLMs are trained on the collected corpus from a certain domain, and we aim to only transfer its learned advanced capabilities to target domains (Zhuang et al., 2021; Farahani et al., 2021) without further training. In this work, we study the cross-lingual scene where the linguistic-agnostic advanced ability and linguistic abilities should be extracted and transferred, to build a unified multi-lingual ability-enhanced LLM.

Formally, for a certain ability A_i and a set of languages $L = \{L_0, L_1, \dots, L_n\}$, we assume that the general corpus of all languages can be collected, denoted as $C_{\text{general}} = \{C_{L_0}, C_{L_1}, \dots, C_{L_n}\}$, while the ability-related corpus is only available in language L_0 (*i.e.*, English), denoted as C_{L_0, A_i} . Based on the above corpora, our goal is to extract and transfer the advanced ability A_i from language L_0 and linguistic abilities from other languages L_1, \dots, L_n , into a unified LLM.

4 Approach

In this section, we propose the **Multi-lingual Ability Extraction and Transfer** approach, named as **MAET**, which can effectively transfer the ad-

vanced abilities from single-lingual LLMs, to build a multi-lingual ability-enhanced LLM. The key motivation of our approach is to identify and extract ability-related neurons or weights, and transfer the target abilities into a LLM in an efficient way. The framework of MAET is presented in Figure 1.

4.1 Ability-related Weights Extraction

In this part, we aim to locate and learn ability-related parameter weights within an LLM, to enable efficient transferring of the ability into other LLMs. Concretely, it consists of two major steps, *i.e.*, key neurons locating and ability-related parameter weights learning, which are detailed in the following.

Locating the Key Neurons. The gradient of each neuron in LLMs can be utilized to estimate its correlation degree with specific task ability (θ_j), we select those with high gradient values as key neurons. To this end, we first use the ability-related corpus C'_{L_0, A_i} to continually pre-train the LLM, while sampling a small amount to train the model can be also applied to reduce the computation consumption. During training, the LLM learns the language modeling task and each neuron is updated by the gradients associated by the training instances. Due to the high cost of calculating the accumulation of gradient at each training step, we calculate the value changes of the LLM neurons before and after the training process to approximate the importance. Formally, the importance function $I(A_i, \theta_j)$ of neurons can be computed as:

$$I(A_i, \theta_j) = \sum_{d_k \in C'_{L_0, A_i}} \text{Grad}(\theta_j, d_k) \approx \lambda \cdot \|\tilde{\theta}_j - \theta_j\|, \quad (1)$$

where d_k denotes the k -th instance of training corpus C'_{L_0, A_i} and $\tilde{\theta}_j$ denote the value of the j -th neuron of L_{L_0, A_i} after training, respectively. Based on it and inspired by previous work (Yadav et al., 2023), we rank all neurons according to their importance scores, and then select the top $k_1\%$ ones into the set N_{A_i} as the key neurons.

Learning Ability-related Weights. Based on the identified key neurons in N_{A_i} , we further learn the ability-related parameter weights. Our motivation is to decompose the parameter weights according to their changes *before* and *after* the LLM has mastered a specific ability, which is achievable owing to the modularity and composition nature of the LLM parameter matrices (Yu et al., 2024a;




Figure 1: The framework of our approach MAET, including extraction and transfer stages. In the extraction stage, we first locate the key neurons, and utilize the single-lingual ability-related corpus and general corpus to train the LLM on these neurons to obtain the ability-related weight. Then, we remove the parameter tensors related to language knowledge in the ability weight and transfer the remaining to the base LLM. After these stages, we can obtain a powerful LLM with advanced abilities that can solve the corresponding tasks in multi-lingual scenarios.

Shazeer et al., 2017). First, we utilize the key neurons locating method mentioned above to extract the ability-related neuron set \mathcal{N}_{A_i} , and also obtain the language-related neuron set \mathcal{N}_{L_0} via the same way. Then, we train the LLM with the mixture of ability-related corpus and general corpus on the key neuron set $\mathcal{N}_{A_i} \cup \mathcal{N}_{L_0}$ and \mathcal{N}_{L_0} respectively, to obtain two specific models, denoted as LLM_{A_i, L_0} with parameters Θ_{A_i, L_0} and LLM_{L_0} with parameters Θ_{L_0} . Next, we measure the parameter changes between the backbone and the trained models, and obtain the ability-related weights via the parameter decomposition operation as:

$$R(A_i) = \alpha \cdot \underbrace{(\Theta_{A_i, L_0} - \Theta_o)}_{\text{Ability \& language difference}} - \beta \cdot \underbrace{(\Theta_{L_0} - \Theta_o)}_{\text{Language difference}}, \quad (2)$$

where α and β are tunable coefficients to balance the two parts of weight differences, and Θ_o denote the original parameters of the LLM, which serves as the reference for parameter decomposition. As we only train the parameters within the neuron set, its weight difference should preserve the knowledge about the corresponding ability. Thus, it can be regarded as the *ability-related parameter representations*, and is promising to transfer the ability into other LLMs by the addition operation.

4.2 Multi-lingual Ability Transfer

After obtaining the ability-related weights, we utilize them to transfer and integrate the abilities, to build a multi-lingual ability-enhanced LLM.

Ability-related Parameter Tensor Selection. Although we can locate the ability-related key neurons, it is still hard to avoid the involvement of irrelevant ones. Our empirical studies have found that neuron-level features are easy to be affected by the noisy data. Therefore, we consider identifying ability-related parameter tensors, which correspond to the parameter matrices within the LLM. Specifically, we firstly leverage the ability-related weights of languages $R(L_1), \dots, R(L_n)$ to obtain the multi-lingual weight R_{Lang} . Given that large models have varying levels of proficiency in different languages, we use the hyper-parameters μ_1, \dots, μ_n to tune this process as:

$$R_{Lang} = \sum_{i=1}^n \mu_i \cdot R(L_i), \quad (3)$$

where $R(L_i)$ preserves the linguistic ability of language L_i learned based on Equation 2. Therefore, R_{Lang} can be considered as the general language ability of LLMs that spans multiple languages. As we aim to find the parameter tensors that have low

Approaches	MLAR	TPara	AC	AT
CPT	Yes	Full	No	No
MoE	Yes	Full	No	No
LoRA	Yes	Low-Rank	No	No
MoL	Yes	Low-Rank	No	No
TV	Yes	Full	Yes	No
MAET	No	Ability-related	Yes	Yes

Table 1: The difference between our MAET and the methods in previous work (*i.e.*, CPT (Hu et al., 2022), Mixture-of-Expert (MoE) (Shazeer et al., 2017), LoRA (Hu et al., 2022), Mixture-of-LoRA (MoL) (Feng et al., 2024), and Task Vector (TV) (Ilharco et al., 2023). MLAR, TPara, AC, and AT denote the abbreviation of multi-lingual ability-related corpus, parameters for training, ability composition, and ability transfer.

linguistic effects but focus on the desired abilities (*e.g.*, mathematical reasoning), we rank all the tensors according to their similarities with R_{Lang} , and pick up the last $k_2\%$ ones. Formally, for tensor τ_i , we calculate the cosine similarity of this parameter between $R(A_i)$ and R_{Lang} , as follows,

$$S(\tau_i) = \text{sim}(R(A_i)[\tau_i], R_{Lang}[\tau_i]), \quad (4)$$

where we use the cosine similarity to implement the similarity function $\text{sim}(\cdot)$. After obtaining the similarity of all tensors, we rank them in a descending order based on the similarity values, and then select the last $k_2\%$ parameters into the set \mathcal{T} as the ability-related parameters.

Building Multi-lingual Ability-enhanced LLM.

Based on the selected ability-related tensors \mathcal{T} , we design the model merging process by interpolating ability weights and multi-lingual weights, to build the multi-lingual ability-enhanced LLM. Formally, the final parameter tensors of the target LLM are computed as:

$$\tilde{\tau}_i = \tau_i^{(o)} + \begin{cases} \gamma \cdot R(A_i)[\tau_i] + \eta \cdot R_{Lang}[\tau_i], & \tau_i \in \mathcal{T} \\ R_{Lang}[\tau_i], & \tau_i \notin \mathcal{T} \end{cases}, \quad (5)$$

where $\tau_i^{(o)}$ denotes the original value of parameter tensor τ_i , and γ and η are tunable hyper-parameters. This formula can be explained in two different cases. When a parameter tensor serves as the major role for specific abilities, we update it by adding both ability- and linguistic-related weights; otherwise, we simply enhance it with multi-lingual weights. In this way, we can derive a more powerful LLM that is equipped with the multi-lingual abilities and specific advanced abilities.

4.3 The Overall Procedure

To better demonstrate our approach, we present key concepts in Table 4 for further clarifying and provide the complete procedure in Algorithm 1 in the pseudo-code form. The procedure of MAET consists of two main stages, *i.e.*, ability-related weights extraction and multi-lingual ability transferring. For the extraction stage, we first utilize the accumulated gradient to estimate the importance of each neuron by Equation 1. Then, we leverage the model trained on the general corpus to remove the influence of language and obtain the ability-related weight through Equation 2. In the transfer stage, we utilize Equation 3 and Equation 4 to obtain the multi-lingual weight and identify the ability-related parameter tensors in LLM. After it, we leverage Equation 5 to fulfill the multi-lingual ability transfer, to build the multi-lingual ability-enhanced LLM.

To highlight the difference between our approach and previous work, we present the comparison of these methods in Table 1. To adapt LLMs to multi-lingual scenarios, most of the existing methods (*e.g.*, CPT and TV) require the multi-lingual ability-related corpus (*i.e.*, ability-related corpus is required for each language) for training the LLM parameters. In comparison, our proposed approach only trains and modifies the ability-related parameters, which can efficiently focus on enhancing the specific ability. A major novelty of our work is that we identify the key units and implement the sparse update in the model training and merging procedure, which can effectively decompose, extract, and transfer the abilities of LLMs. In addition, compared with the LoRA-based methods (*i.e.*, LoRA and MoL) that also sparsely update the LLM parameters, our approach selectively updates the ability-related neurons, while LoRA-based methods utilize the low-rank matrices to approximate the original parameters.

5 Experiment

5.1 Experimental Settings

In this part, we introduce the details of our experimental settings, including the datasets utilized in the training and evaluation process, and the baseline methods. Moreover, we present the implementation details of our approach in Appendix B.

Datasets. In this work, we focus on transferring the advanced abilities (*i.e.*, mathematical and

scientific reasoning abilities) of LLMs from English scenarios to multi-lingual scenarios, including high-resource languages (*i.e.*, Spanish and Chinese) and low-resource languages (*i.e.*, Bengali and Telugu). For the training corpus, we extract the corpus of the corresponding languages from the dataset proposed by previous work (Yang et al., 2023; Scao et al., 2022; Laurençon et al., 2022) as the general training corpus, and utilize OpenWebMath (Paster et al., 2024) and the arXiv papers (Soldaini et al., 2024) as the ability-related corpus for mathematical tasks and scientific tasks respectively. For the evaluation benchmark, we follow the evaluation settings in previous work (OpenAI, 2023), utilizing *Multi-lingual Grade School Math (MGSM)* (Shi et al., 2023) and science tasks from *multi-lingual MMLU* (Lai et al., 2023) (*i.e.*, college biology, college chemistry, college physics, high school biology, high school chemistry, and high school physics) as the downstream tasks for multi-lingual scenarios. The statistical information of the datasets is presented in Table 6.

Baselines. In our evaluation, a baseline can be represented as three parts, *i.e.*, training parameters, training approach, and training data. First, we conduct the full parameters training and the LoRA training (Hu et al., 2022) in our evaluation, denoted as the “*F*” and “*L*” at the prefix of the training approaches, respectively. For the training approach, we employ *continual pre-training (CPT)* (Gururangan et al., 2020), *domain adaption (DA)* (Taylor et al., 2022), and *model merging with task vector (TV)* (Ilharco et al., 2023). Besides, for the training data, “*L*”, “*A*”, and “*T*” refer to the multi-lingual general corpus, the single-lingual ability-related corpus, and the translated multi-lingual ability-related corpus from GPT-4o (Hurst et al., 2024), respectively. Moreover, to conduct a more comprehensive evaluation, we also present the performance of different LLMs, *i.e.*, Baichuan-2 7B (Yang et al., 2023), Mistral 7B (Jiang et al., 2023), LLaMA-2 7B (Touvron et al., 2023), and LLaMA-3 8B (Dubey et al., 2024).

5.2 Main Results

To comprehensively evaluate our proposed MAET, we employ MAET on mathematical and scientific tasks in multi-lingual scenarios and present the results in Table 2.

First, MAET outperforms other baseline methods in the average performance of all downstream

languages. In our experiment, continual pre-training LLMs on the mixture of multi-lingual general corpus and single-lingual ability-related corpus (*i.e.*, F-CPT_{L&A}) can enhance the specific ability of LLMs, achieving the second best performance. However, when adapting LLMs to a new domain or enhancing a new ability of LLM, CPT-based methods should retrain the LLMs on the ability-related and multi-lingual corpus, showing that CPT is leaked of transferability and requires more computational resources. For the issue of new domain adapting, MAET only utilizes a small amount of single-lingual ability-related corpus (*i.e.*, English corpus in practice) to obtain the ability weight, which can be employed to transfer the corresponding advanced ability, achieving both effectiveness and efficiency.

Second, LoRA-based methods (Hu et al., 2022) (*e.g.*, L-CPT_{L&A}, L-CPT_L, L-TV) initialize the low-rank matrices and only update these matrices, performing sparsely optimize on LLM. Since the trainable parameters in LoRA represent the whole parameters of LLM rather than ability-related sub-network, it cannot perform well on the multi-lingual scenarios, indicating the failure of the advanced abilities transferring. In contrast, MAET first identifies the ability-related sub-networks and utilizes the corresponding sub-networks to perform the following operations. Because of the decomposing of the inner abilities of LLMs, MAET can help LLMs improve their specific ability.

Third, translation-based methods are the strong baselines to enhance the LLM performance in low-resource languages. In the experiment, we utilize GPT-4o to translate the ability-related corpus from English to other languages, and present the prompt in Appendix D. According to the experimental results in the above table, we can observe that our MAET outperforms the translation-based method. The translation-based method consumes more computational resources and cannot achieve better performance. The reason might be that the translated corpus shares similar knowledge of the specific domain, which makes LLM overfit the corresponding knowledge and cannot really understand the specific knowledge. In contrast, our approach decomposes the scientific ability and language ability, and transfers the scientific ability from one language to another, preventing overfitting, decreasing the expense, and improving performance.

Last, compared with the model merging based

Methods	Multilingual Mathematical Tasks					Multilingual Scientific Tasks				
	ES	ZH	BN	TE	Avg.	ES	ZH	BN	TE	Avg.
Baichuan-2 7B	17.20	28.00	4.80	2.40	13.10	42.27	46.43	30.17	26.21	36.27
Mistral 7B	38.80	34.40	9.60	2.80	21.40	52.08	45.33	32.91	27.96	39.57
LLaMA-2 7B	7.60	12.00	1.60	0.00	5.30	34.16	31.68	24.56	22.15	28.14
LLaMA-3 8B	<u>48.40</u>	38.80	28.80	20.40	34.10	55.06	47.24	36.63	29.26	42.05
+ F-CPT _{L&A}	46.80	42.00	28.40	27.60	<u>36.20</u>	55.92	<u>48.57</u>	<u>36.84</u>	30.10	<u>42.86</u>
+ L-CPT _{L&A}	44.80	37.60	28.80	23.60	33.70	54.77	46.81	36.41	29.88	41.97
+ F-CPT _{A&T}	-	-	-	-	-	53.73	46.30	35.06	31.73	41.71
+ F-CPT _A	-	-	-	-	-	51.90	45.71	33.35	29.41	40.09
+ F-CPT _T	-	-	-	-	-	50.35	45.36	34.54	34.46	41.18
+ F-CPT _L	38.80	35.60	28.00	23.60	31.50	53.56	47.14	35.89	30.64	41.81
+ F-CPT _L & DA	41.60	39.60	34.40	27.60	35.80	52.71	48.05	35.49	28.62	41.11
+ L-CPT _L	46.40	39.20	28.40	22.80	34.20	55.04	48.09	36.66	30.43	42.56
+ L-CPT _L & DA	46.80	37.60	28.00	<u>27.20</u>	34.90	55.65	49.10	36.48	29.65	42.72
+ F-TV	42.00	32.40	16.00	10.40	25.20	53.36	46.57	36.70	<u>30.73</u>	41.84
+ L-TV	45.60	39.20	30.80	25.60	35.30	55.46	48.27	36.65	<u>30.44</u>	42.71
+ MAET (Ours)	49.60	<u>41.60</u>	<u>32.40</u>	25.20	37.20	56.20	48.00	37.64	30.38	43.06

Table 2: The performance comparison of different approaches on multilingual mathematical tasks and multilingual scientific tasks. Avg. denotes the average accuracy of the multi-lingual tasks. ES, ZH, BN, and TE denote Spanish, Chinese, Bengali, and Telugu, respectively. The best is in bold and the second best is underlined.

approaches (*i.e.*, F-TV and L-TV), experimental results have shown that MAET performs better than these baseline methods, since we decompose the relation between ability and the language of the training corpus. In the previous model merging approaches, they mainly added the parameters of different models to obtain the final model, without considering the the relation between language and abilities. Due to the extraction mechanism of MAET, we mitigate the effect of languages and make the weight more related to ability, which can be transferred in multi-lingual scenarios.

5.3 Detailed Analysis

To comprehensively evaluate our proposed approach MAET and analyze its features, we conduct several experiments and detailed analysis in this part, including the ablation study, analysis of the transfer ratio of LLM parameters, and the generalization of MAET.

Ablation Study. To assess the effectiveness of each component of our proposed MAET, we conduct the ablation study and present the results in Figure 2. We implement MAET on multi-lingual mathematical and scientific tasks without each module of MAET, *i.e.*, key neurons locating (*i.e.*, Eq. 1), ability weight obtaining (*i.e.*, Eq. 2), ability-related parameter tensor identifying (*i.e.*, Eq. 4), and advanced abilities transferring (Eq. 5). First, in most downstream scenarios, removing any module of




Figure 2: The results of ablation study. KNL, AW, API, and AT denote key neurons locating (Eq. 1), ability weights obtaining (Eq. 2), ability-related tensors identifying (Eq. 4), and advanced abilities transferring (Eq. 5).

MAET will affect the final performance, which has verified the effectiveness of the MAET process. Second, without ability weight obtaining, *i.e.*, directly utilizing the difference between LLM trained on the ability-related corpus and the backbone LLM as the ability weight, we can observe that the performance is seriously hurt in both scenarios, indicating this process can significantly extract the advanced abilities from the single-lingual corpus and decrease the influence of the language




Figure 3: The performance of different proportions for the ability-related parameters identification.

of the training corpus. Third, comparing the results of the models whether adopting the ability transferring process, experimental results show that LLM with the multi-lingual abilities enhanced cannot well solve multi-lingual mathematical and scientific tasks, and leveraging the ability weight provided by MAET can improve the LLM performance on advanced tasks.

Ratio of Key Parameters During Transferring Stage. Identifying and updating the ability-related sub-network of LLMs is the key point of our MAET. We conduct experiments to analyze the influence of the transferring ratio $k_2\%$ and show the results in Figure 3. Observing the results, the performance of LLM has decreased in a lower and higher ratio of the ability-related parameters identifying process. The main reason is that the lower proportion transfers incomplete knowledge to the model and makes LLM unable to possess the corresponding ability, affecting the performance on the downstream tasks. In contrast, the higher proportion cannot extract the ability weight precisely and will transfer too much language-related knowledge to the model, making the conflict with the LLM inner knowledge and hurting the multi-lingual scenarios performance.

Out-of-Domain Performance of MAET. We conduct experiments about adapting mathematical ability on the general LLM through MAET, and assess the performance on out-of-domain (OOD) tasks (*i.e.*, MMLU (Hendrycks et al., 2021), HumanEval (Chen et al., 2021a), MBPP (Austin et al., 2021), and OpenbookQA (Mihaylov et al., 2018)), which can reflect and assess different abilities of LLMs. Results are presented in Table 3. We can

Methods	MMLU	MBPP	OpenbookQA
LLaMA-3 8B	60.85	46.60	65.00
+ CPT	-2.39	-7.00	-3.60
+ MAET	+0.22	+0.80	+0.00

Table 3: The out-of-domain performance of different training methods to train LLaMA-3 8B on OpenWebMath. During the ability-enhancing process, previous methods will hurt the OOD abilities of LLM, while our MAET can maintain the corresponding abilities.

observe that the performance of LLM on all evaluation tasks has decreased through the CPT training process, and the maximum decrease has been achieved 7.32% on the HumanEval task. One of the possible reasons is that LLaMA-3 has been trained on OpenWebMath during pre-training and the CPT process makes it overfit and forget the knowledge of other domains, hurting the performance on OOD tasks. In contrast, our proposed MAET achieves comparable and even better performance with backbone LLM in all downstream scenarios. Since we identify and update the key neurons related to the specific ability, the ability of LLM can be precisely enhanced, and this strategy also helps the OOD tasks needed for mathematical ability, *e.g.*, mathematical sub-tasks in MMLU and code synthesis task MBPP.

6 Conclusion

In this paper, we presented MAET, which extracted the advanced ability-related weights from the LLM and supported simple addition and subtraction operations to transfer the ability across different languages. Concretely, MAET included two main stages, *i.e.*, extraction and transfer. For the extraction stage, we located the key neurons and extracted the ability-related weights. Then, in the transfer stage, we identified the key parameter tensors and leveraged them to transfer the advanced ability into other LLMs. In this process, the multi-lingual ability-related training corpus is not required, and the experimental results have shown that our approach outperformed competitive baselines.

As future work, we will consider better methods to identify the ability-related sub-network to decompose the abilities of LLMs and utilize an automated approach to determine the hyper-parameter. Besides, we will implement MAET on larger-scale models, and scenarios with more languages and requiring more abilities to evaluate its effectiveness.

Limitations

In this section, we discuss the limitations of our work. First, we only implement our approach MAET on 8B LLMs (*i.e.*, LLaMA-3 8B), and do not adopt the LLMs with larger scales (*e.g.*, 13B or 70B LLMs) in the experiment, due to the limitation of computational resources. We will test the effectiveness of our approach on these LLMs in the future. Second, we only evaluate our approach on two downstream tasks (*i.e.*, mathematical and scientific reasoning tasks) in multi-lingual scenarios. Although they are challenging and widely-used testbeds, it is still meaningful to verify our methods on other tasks. Whereas, as we test the performance on diverse high-resource and low-resource languages, it can also provide comprehensive performance estimation for our approach in multi-lingual scenarios. Finally, we do not consider the potential risk and ethics issues that might hurt the alignment of LLMs when using our approach. Actually, our approach is also applicable to transfer the alignment ability across languages. We will investigate to it in the future.

References

Jacob Austin, Augustus Odena, Maxwell I. Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. 2021. Program synthesis with large language models. *CoRR*, abs/2108.07732.

Hsin-Yu Chang, Pei-Yu Chen, Tun-Hsiang Chou, Chang-Sheng Kao, Hsuan-Yun Yu, Yen-Ting Lin, and Yun-Nung Chen. 2024. A survey of data synthesis approaches. *CoRR*, abs/2407.03672.

Jie Chen, Zhipeng Chen, Jiapeng Wang, Kun Zhou, Yutao Zhu, Jinhao Jiang, Yingqian Min, Wayne Xin Zhao, Zhicheng Dou, Jiaxin Mao, Yankai Lin, Ruihua Song, Jun Xu, Xu Chen, Rui Yan, Zhewei Wei, Di Hu, Wenbing Huang, and Ji-Rong Wen. 2024a. Towards effective and efficient continual pre-training of large language models. *CoRR*, abs/2407.18743.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie

Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021a. Evaluating large language models trained on code. *CoRR*, abs/2107.03374.

Richard J Chen, Ming Y Lu, Tiffany Y Chen, Drew FK Williamson, and Faisal Mahmood. 2021b. Synthetic data in machine learning for medicine and healthcare. *Nature Biomedical Engineering*, 5(6):493–497.

Zhipeng Chen, Kun Zhou, Wayne Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2024b. Low-redundant optimization for large language model alignment. *CoRR*, abs/2406.12606.

Pei Cheng, Xiayang Shi, and Yinlin Li. 2024. Enhancing translation ability of large language models by leveraging task-related layers. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pages 6110–6121.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Wenyu Du, Shuang Cheng, Tongxu Luo, Zihan Qiu, Zeyu Huang, Ka Chun Cheung, Reynold Cheng, and Jie Fu. 2024. Unlocking continual learning abilities in language models. *CoRR*, abs/2406.17245.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Alonso, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar,

733	Jeet Shah, Jelmer van der Linde, Jennifer Billock,	Zixian Huang, Wenhao Zhu, Gong Cheng, Lei Li, and	789
734	Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi,	Fei Yuan. 2024. Mindmerger: Efficient boosting	790
735	Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu,	LLM reasoning in non-english languages. <i>CoRR</i> ,	791
736	Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph	abs/2405.17386.	792
737	Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia,		
738	Kalyan Vasuden Alwala, Kartikeya Upasani, Kate	Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Day-	793
739	Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and	iheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang,	794
740	et al. 2024. The llama 3 herd of models. <i>CoRR</i> ,	Bowen Yu, Kai Dang, An Yang, Rui Men, Fei Huang,	795
741	abs/2407.21783.	Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and	796
		Junyang Lin. 2024. Qwen2.5-coder technical report.	797
742	Abolfazl Farahani, Behrouz Pourshojae, Khaled	<i>CoRR</i> , abs/2409.12186.	798
743	Rasheed, and Hamid R. Arabnia. 2021. A concise		
744	review of transfer learning. <i>CoRR</i> , abs/2104.02144.	Aaron Hurst, Adam Lerer, Adam P Goucher, Adam	799
		Perelman, Aditya Ramesh, Aidan Clark, AJ Os-	800
745	Wenfeng Feng, Chuzhan Hao, Yuewei Zhang, Yu Han,	trow, Akila Welihinda, Alan Hayes, Alec Radford,	801
746	and Hao Wang. 2024. Mixture-of-loras: An efficient	et al. 2024. Gpt-4o system card. <i>arXiv preprint</i>	802
747	multitask tuning method for large language models.	<i>arXiv:2410.21276</i> .	803
748	In <i>Proceedings of the 2024 Joint International Con-</i>		
749	<i>ference on Computational Linguistics, Language Re-</i>	Gabriel Ilharco, Marco Túlio Ribeiro, Mitchell Worts-	804
750	<i>sources and Evaluation, LREC/COLING 2024, 20-25</i>	man, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali	805
751	<i>May, 2024, Torino, Italy</i> , pages 11371–11380.	Farhadi. 2023. Editing models with task arithmetic.	806
		In <i>The Eleventh International Conference on Learn-</i>	807
752	Jonathan Frankle and Michael Carbin. 2019. The lottery	<i>ing Representations, ICLR 2023, Kigali, Rwanda,</i>	808
753	ticket hypothesis: Finding sparse, trainable neural	<i>May 1-5, 2023</i> .	809
754	networks. In <i>7th International Conference on Learn-</i>		
755	<i>ing Representations, ICLR 2019, New Orleans, LA,</i>	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Men-	810
756	<i>USA, May 6-9, 2019</i> .	sch, Chris Bamford, Devendra Singh Chaplot, Diego	811
		de Las Casas, Florian Bressand, Gianna Lengyel,	812
757	Zhuocheng Gong, Ang Lv, Jian Guan, Junxi Yan, Wei	Guillaume Lample, Lucile Saulnier, L��lio Ren-	813
758	Wu, Huishuai Zhang, Minlie Huang, Dongyan Zhao,	nard Lavaud, Marie-Anne Lachaux, Pierre Stock,	814
759	and Rui Yan. 2024. Mixture-of-modules: Reinvent-	Teven Le Scao, Thibaut Lavril, Thomas Wang, Timo-	815
760	ing transformers as dynamic assemblies of modules.	th��e Lacroix, and William El Sayed. 2023. Mistral	816
761	<i>CoRR</i> , abs/2407.06677.	7b. <i>CoRR</i> , abs/2310.06825.	817
762	Suchin Gururangan, Ana Marasovic, Swabha	Yuxin Jiang, Yufei Wang, Chuhan Wu, Wanjun Zhong,	818
763	Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey,	Xingshan Zeng, Jiahui Gao, Liangyou Li, Xin Jiang,	819
764	and Noah A. Smith. 2020. Don’t stop pretraining:	Lifeng Shang, Ruiming Tang, Qun Liu, and Wei	820
765	Adapt language models to domains and tasks. In	Wang. 2024. Learning to edit: Aligning llms with	821
766	<i>Proceedings of the 58th Annual Meeting of the</i>	knowledge editing. In <i>Proceedings of the 62nd An-</i>	822
767	<i>Association for Computational Linguistics, ACL</i>	<i>nual Meeting of the Association for Computational</i>	823
768	<i>2020, Online, July 5-10, 2020</i> , pages 8342–8360.	<i>Linguistics (Volume 1: Long Papers), ACL 2024,</i>	824
		<i>Bangkok, Thailand, August 11-16, 2024</i> , pages 4689–	825
769	Michael A. Hedderich, Lukas Lange, Heike Adel, Jan-	4705.	826
770	nik Str��tgen, and Dietrich Klakow. 2021. A survey		
771	on recent approaches for natural language process-	Viet Dac Lai, Chien Van Nguyen, Nghia Trung Ngo,	827
772	ing in low-resource scenarios. In <i>Proceedings of</i>	Thuat Nguyen, Franck Dernoncourt, Ryan A. Rossi,	828
773	<i>the 2021 Conference of the North American Chap-</i>	and Thien Huu Nguyen. 2023. Okapi: Instruction-	829
774	<i>ter of the Association for Computational Linguistics:</i>	tuned large language models in multiple languages	830
775	<i>Human Language Technologies, NAACL-HLT 2021,</i>	with reinforcement learning from human feedback.	831
776	<i>Online, June 6-11, 2021</i> , pages 2545–2568.	In <i>Proceedings of the 2023 Conference on Empirical</i>	832
		<i>Methods in Natural Language Processing, EMNLP</i>	833
777	Dan Hendrycks, Collin Burns, Steven Basart, Andy	<i>2023 - System Demonstrations, Singapore, December</i>	834
778	Zou, Mantas Mazeika, Dawn Song, and Jacob Stein-	<i>6-10, 2023</i> , pages 318–327.	835
779	hardt. 2021. Measuring massive multitask language		
780	understanding. In <i>9th International Conference on</i>	Hugo Lauren��on, Lucile Saulnier, Thomas Wang,	836
781	<i>Learning Representations, ICLR 2021, Virtual Event,</i>	Christopher Akiki, Albert Villanova del Moral,	837
782	<i>Austria, May 3-7, 2021</i> .	Teven Le Scao, Leandro von Werra, Chenghao Mou,	838
		Eduardo Gonz��lez Ponferrada, Huu Nguyen, J��rg	839
783	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan	Frohberg, Mario Sasko, Quentin Lhoest, Angelina	840
784	Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and	McMillan-Major, G��rard Dupont, Stella Biderman,	841
785	Weizhu Chen. 2022. Lora: Low-rank adaptation of	Anna Rogers, Loubna Ben Allal, Francesco De Toni,	842
786	large language models. In <i>The Tenth International</i>	Giada Pistilli, Olivier Nguyen, Somaieh Nikpoor,	843
787	<i>Conference on Learning Representations, ICLR 2022,</i>	Maram Masoud, Pierre Colombo, Javier de la Rosa,	844
788	<i>Virtual Event, April 25-29, 2022</i> . OpenReview.net.	Paulo Villegas, Tristan Thrush, Shayne Longpre, Se-	845
		bastian Nagel, Leon Weber, Manuel Mu��oz, Jian	846

847	Zhu, Daniel van Strien, Zaid Alyafeai, Khalid Al-	Kale. 2024. On importance of pruning and distilla-	904
848	mubarak, Minh Chien Vu, Itziar Gonzalez-Dios,	tion for efficient low resource nlp. <i>arXiv preprint</i>	905
849	Aitor Soroa, Kyle Lo, Manan Dey, Pedro Ortiz	<i>arXiv:2409.14162</i> .	906
850	Suarez, Aaron Gokaslan, Shamik Bose, David Ife-		
851	oluwa Adelani, Long Phan, Hieu Tran, Ian Yu, Suhas	OpenAI. 2023. GPT-4 technical report. <i>CoRR</i> ,	907
852	Pai, Jenny Chim, Violette Lepercq, Suzana Ilic,	abs/2303.08774.	908
853	Margaret Mitchell, Alexandra Sasha Luccioni, and		
854	Yacine Jernite. 2022. The bigscience ROOTS corpus:	Anand Panchbhai and Smarana Pankanti. 2021. Explor-	909
855	A 1.6tb composite multilingual dataset. In <i>Advances</i>	ing large language models in a limited resource sce-	910
856	<i>in Neural Information Processing Systems 35: Annual</i>	nario. <i>2021 11th International Conference on Cloud</i>	911
857	<i>Conference on Neural Information Processing</i>	<i>Computing, Data Science & Engineering (Conflu-</i>	912
858	<i>Systems 2022, NeurIPS 2022, New Orleans, LA, USA,</i>	<i>ence)</i> , pages 147–152.	913
859	<i>November 28 - December 9, 2022</i> .		
860	Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning:	Keiran Paster, Marco Dos Santos, Zhangir Azerbayev,	914
861	Optimizing continuous prompts for generation. In	and Jimmy Ba. 2024. Openwebmath: An open	915
862	<i>Proceedings of the 59th Annual Meeting of the Asso-</i>	dataset of high-quality mathematical web text. In	916
863	<i>ciation for Computational Linguistics and the 11th</i>	<i>The Twelfth International Conference on Learning</i>	917
864	<i>International Joint Conference on Natural Language</i>	<i>Representations, ICLR 2024, Vienna, Austria, May</i>	918
865	<i>Processing, ACL/IJCNLP 2021, (Volume 1: Long</i>	<i>7-11, 2024</i> .	919
866	<i>Papers)</i> , Virtual Event, August 1-6, 2021, pages 4582–		
867	4597. Association for Computational Linguistics.	Tim Patzelt. 2024. Medical concept normaliza-	920
		tion in a low-resource setting. <i>arXiv preprint</i>	921
		<i>arXiv:2409.14579</i> .	922
868	Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-	Morgane Rivière, Shreya Pathak, Pier Giuseppe	923
869	Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Pe-	Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard	924
870	ter Clark, and Ashwin Kalyan. 2022. Learn to ex-	Hussenot, Thomas Mesnard, Bobak Shahriari,	925
871	plain: Multimodal reasoning via thought chains for	Alexandre Ramé, Johan Ferret, Peter Liu, Pouya	926
872	science question answering. In <i>Advances in Neural</i>	Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos,	927
873	<i>Information Processing Systems 35: Annual Confer-</i>	Ravin Kumar, Charline Le Lan, Sammy Jerome, An-	928
874	<i>ence on Neural Information Processing Systems 2022,</i>	ton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan	929
875	<i>NeurIPS 2022, New Orleans, LA, USA, November 28</i>	Girgin, Nikola Momchev, Matt Hoffman, Shantanu	930
876	<i>- December 9, 2022</i> .	Thakoor, Jean-Bastien Grill, Behnam Neyshabur,	931
		Olivier Bachem, Alanna Walton, Aliaksei Severyn,	932
877	Yingzhou Lu, Huazheng Wang, and Wenqi Wei. 2023.	Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin	933
878	Machine learning for synthetic data generation: a	Abdagic, Amanda Carl, Amy Shen, Andy Brock,	934
879	review. <i>CoRR</i> , abs/2302.04062.	Andy Coenen, Anthony Laforge, Antonia Pater-	935
		son, Ben Bastian, Bilal Piot, Bo Wu, Brandon	936
880	Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng	Royal, Charlie Chen, Chintu Kumar, Chris Perry,	937
881	Zhang, Hoifung Poon, and Tie-Yan Liu. 2022.	Chris Welty, Christopher A. Choquette-Choo, Danila	938
882	Biogpt: generative pre-trained transformer for	Sinopalnikov, David Weinberger, Dimple Vijayku-	939
883	biomedical text generation and mining. <i>Briefings</i>	mar, Dominika Rogozinska, Dustin Herbison, Elisa	940
884	<i>Bioinform.</i> , 23(6).	Bandy, Emma Wang, Eric Noland, Erica Moreira,	941
		Evan Senter, Evgenii Eltyshhev, Francesco Visin,	942
885	Alexandre Magueresse, Vincent Carles, and Evan	Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus	943
886	Heetderks. 2020. Low-resource languages: A re-	Martins, Hadi Hashemi, Hanna Klimczak-Plucinska,	944
887	view of past work and future challenges. <i>CoRR</i> ,	Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda	945
888	abs/2006.07264.	Mein, Jack Zhou, James Svensson, Jeff Stanway,	946
		Jetha Chan, Jin Peng Zhou, Joana Carrasqueira,	947
889	Michael Matena and Colin Raffel. 2022. Merging mod-	Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost	948
890	els with fisher-weighted averaging. In <i>Advances in</i>	van Amersfoort, Josh Gordon, Josh Lipschultz,	949
891	<i>Neural Information Processing Systems 35: Annual</i>	Josh Newlan, Ju-yeong Ji, Kareem Mohamed, Kar-	950
892	<i>Conference on Neural Information Processing Sys-</i>	tikeya Badola, Kat Black, Katie Millican, Keelin	951
893	<i>tems 2022, NeurIPS 2022, New Orleans, LA, USA,</i>	McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish	952
894	<i>November 28 - December 9, 2022</i> .	Greene, Lars Lowe Sjöstrand, Lauren Usui, Laurent	953
		Sifre, Lena Heuermann, Leticia Lago, and Lilly Mc-	954
895	Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish	Nealus. 2024. Gemma 2: Improving open language	955
896	Sabharwal. 2018. Can a suit of armor conduct elec-	models at a practical size. <i>CoRR</i> , abs/2408.00118.	956
897	tricity? A new dataset for open book question an-		
898	swering. In <i>Proceedings of the 2018 Conference on</i>	Teven Le Scao, Angela Fan, Christopher Akiki, El-	957
899	<i>Empirical Methods in Natural Language Processing,</i>	lie Pavlick, Suzana Ilic, Daniel Hesslow, Roman	958
900	<i>Brussels, Belgium, October 31 - November 4, 2018,</i>	Castagné, Alexandra Sasha Luccioni, François Yvon,	959
901	pages 2381–2391.	Matthias Galle, Jonathan Tow, Alexander M. Rush,	960
		Stella Biderman, Albert Webson, Pawan Sasanka Am-	961
902	Aishwarya Mirashi, Purva Lingayat, Srushti Sona-	manamanchi, Thomas Wang, Benoît Sagot, Niklas	962
903	vane, Tejas Padhiyar, Raviraj Joshi, and Geetanjali		

963	Muennighoff, Albert Villanova del Moral, Olatunji	15725–15788. Association for Computational Lin-	1021
964	Ruwase, Rachel Bawden, Stas Bekman, Angelina	guistics.	1022
965	McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile	George Stoica, Daniel Bolya, Jakob Bjorner, Pratik	1023
966	Saulnier, Samson Tan, Pedro Ortiz Suarez, Vic-	Ramesh, Taylor Hearn, and Judy Hoffman. 2024.	1024
967	tor Sanh, Hugo Laurençon, Yacine Jernite, Julien	Zipit! merging models from different tasks without	1025
968	Launay, Margaret Mitchell, Colin Raffel, Aaron	training. In <i>The Twelfth International Conference</i>	1026
969	Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri	<i>on Learning Representations, ICLR 2024, Vienna,</i>	1027
970	Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg	<i>Austria, May 7-11, 2024.</i>	1028
971	Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue,		
972	Christopher Klammer, Colin Leong, Daniel van Strien,	Tianyi Tang, Wenyang Luo, Haoyang Huang, Dong-	1029
973	David Ifeoluwa Adelani, and et al. 2022. BLOOM:	dong Zhang, Xiaolei Wang, Xin Zhao, Furu Wei,	1030
974	A 176b-parameter open-access multilingual language	and Ji-Rong Wen. 2024. Language-specific neurons:	1031
975	model. <i>CoRR</i> , abs/2211.05100.	The key to multilingual capabilities in large language	1032
976	Anton Schäfer, Shauli Ravfogel, Thomas Hofmann,	models. In <i>Proceedings of the 62nd Annual Meeting</i>	1033
977	Tiago Pimentel, and Imanol Schlag. 2024. Lan-	<i>of the Association for Computational Linguistics (Vol-</i>	1034
978	guage imbalance can boost cross-lingual generali-	<i>ume 1: Long Papers), ACL 2024, Bangkok, Thailand,</i>	1035
979	sation. <i>CoRR</i> , abs/2404.07982.	<i>August 11-16, 2024</i> , pages 5701–5715.	1036
980	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu,	Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas	1037
981	Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu,	Scialom, Anthony Hartshorn, Elvis Saravia, An-	1038
982	and Daya Guo. 2024. Deepseekmath: Pushing the	drew Poulton, Viktor Kerkez, and Robert Stojnic.	1039
983	limits of mathematical reasoning in open language	2022. Galactica: A large language model for science.	1040
984	models. <i>CoRR</i> , abs/2402.03300.	<i>CoRR</i> , abs/2211.09085.	1041
985	Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarczyk,	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	1042
986	Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	1043
987	Dean. 2017. Outrageously large neural networks:	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	1044
988	The sparsely-gated mixture-of-experts layer. <i>arXiv</i>	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-	1045
989	<i>preprint arXiv:1701.06538.</i>	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	1046
990	Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang,	Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller,	1047
991	Suraj Srivats, Soroush Vosoughi, Hyung Won Chung,	Cynthia Gao, Vedanuj Goswami, Naman Goyal, An-	1048
992	Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das,	thony Hartshorn, Saghar Hosseini, Rui Hou, Hakan	1049
993	and Jason Wei. 2023. Language models are multi-	Inan, Marcin Kardas, Viktor Kerkez, Madihan Khabsa,	1050
994	lingual chain-of-thought reasoners. In <i>The Eleventh</i>	Isabel Kloumann, Artem Korenev, Punit Singh Koura,	1051
995	<i>International Conference on Learning Representa-</i>	Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Di-	1052
996	<i>tions, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.</i>	ana Liskovich, Yinghai Lu, Yuning Mao, Xavier Mar-	1053
997	Shamane Siriwardhana, Mark McQuade, Thomas Gau-	tinet, Todor Mihaylov, Pushkar Mishra, Igor Moly-	1054
998	thier, Lucas Atkins, Fernando Fernandes Neto, Luke	bog, Yixin Nie, Andrew Poulton, Jeremy Reizen-	1055
999	Meyers, Annekeith Vij, Tyler Odenhall, Charles God-	stein, Rishi Rungta, Kalyan Saladi, Alan Schelten,	1056
1000	dard, Mary MacCarthy, and Jacob Solawetz. 2024.	Ruan Silva, Eric Michael Smith, Ranjan Subrama-	1057
1001	Domain adaptation of llama3-70b-instruct through	nian, Xiaoqing Ellen Tan, Binh Tang, Ross Tay-	1058
1002	continual pre-training and model merging: A com-	lor, Adina Williams, Jian Xiang Kuan, Puxin Xu,	1059
1003	prehensive evaluation. <i>CoRR</i> , abs/2406.14971.	Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan,	1060
1004	Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin	Melanie Kambadur, Sharan Narang, Aurélien Ro-	1061
1005	Schwenk, David Atkinson, Russell Authur, Ben Bo-	driguez, Robert Stojnic, Sergey Edunov, and Thomas	1062
1006	gin, Khyathi Raghavi Chandu, Jennifer Dumas, Yanai	Scialom. 2023. Llama 2: Open foundation and fine-	1063
1007	Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin	tuned chat models. <i>CoRR</i> , abs/2307.09288.	1064
1008	Kumar, Li Lucy, Xinxin Lyu, Nathan Lambert, Ian	Huanqian Wang, Yang Yue, Rui Lu, Jingxin Shi, An-	1065
1009	Magnusson, Jacob Morrison, Niklas Muennighoff,	drew Zhao, Shenzhi Wang, Shiji Song, and Gao	1066
1010	Aakanksha Naik, Crystal Nam, Matthew E. Peters,	Huang. 2024a. Model surgery: Modulating llm’s	1067
1011	Abhilasha Ravichander, Kyle Richardson, Zejiang	behavior via simple parameter editing. <i>CoRR</i> ,	1068
1012	Shen, Emma Strubell, Nishant Subramani, Oyvind	abs/2407.08770.	1069
1013	Tafjord, Evan Pete Walsh, Luke Zettlemoyer, Noah A.	Mengru Wang, Ningyu Zhang, Ziwen Xu, Zekun Xi,	1070
1014	Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groen-	Shumin Deng, Yunzhi Yao, Qishen Zhang, Linyi	1071
1015	evelt, Jesse Dodge, and Kyle Lo. 2024. Dolma:	Yang, Jindong Wang, and Huajun Chen. 2024b.	1072
1016	an open corpus of three trillion tokens for language	Detoxifying large language models via knowledge	1073
1017	model pretraining research. In <i>Proceedings of the</i>	editing. In <i>Proceedings of the 62nd Annual Meeting</i>	1074
1018	<i>62nd Annual Meeting of the Association for Compu-</i>	<i>of the Association for Computational Linguistics (Vol-</i>	1075
1019	<i>tational Linguistics (Volume 1: Long Papers), ACL</i>	<i>ume 1: Long Papers), ACL 2024, Bangkok, Thailand,</i>	1076
1020	<i>2024, Bangkok, Thailand, August 11-16, 2024</i> , pages	<i>August 11-16, 2024</i> , pages 3093–3118.	1077

1078	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien	Model merging in llms, mllms, and beyond: Meth-	1135
1079	Chaumond, Clement Delangue, Anthony Moi, Pier-	ods, theories, applications and opportunities. <i>CoRR</i> ,	1136
1080	ric Cistac, Tim Rault, Remi Louf, Morgan Funtow-	abs/2408.07666.	1137
1081	icz, Joe Davison, Sam Shleifer, Patrick von Platen,		
1082	Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu,	1138	
1083	Teven Le Scao, Sylvain Gugger, Mariama Drame,	1139	
1084	Quentin Lhoest, and Alexander Rush. 2020. Trans-	Optimizing data mixtures by predicting language	1140
1085	formers: State-of-the-art natural language processing.	modeling performance. <i>CoRR</i> , abs/2403.16952.	1141
1086	In <i>Proceedings of the 2020 Conference on Empirical</i>		
1087	<i>Methods in Natural Language Processing: System</i>	1142	
1088	<i>Demonstrations</i> , pages 38–45, Online. Association	1143	
1089	for Computational Linguistics.	1144	
		1145	
1090	Chaojun Xiao, Zhengyan Zhang, Chenyang Song, Dazhi	1146	
1091	Jiang, Feng Yao, Xu Han, Xiaozhi Wang, Shuo Wang,	1147	
1092	Yufei Huang, Guanyu Lin, et al. 2024. Configurable	1148	
1093	foundation models: Building llms from a modular	7522. Association for Computational Linguistics.	1149
1094	perspective. <i>arXiv preprint arXiv:2409.02877</i> .		
1095	Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du,	1150	
1096	Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V. Le,	1151	
1097	Tengyu Ma, and Adams Wei Yu. 2023. Doremi: Op-	1152	
1098	timizing data mixtures speeds up language model	1153	
1099	pretraining. In <i>Advances in Neural Information Pro-</i>	1154	
1100	<i>cessing Systems 36: Annual Conference on Neural</i>	1155	
1101	<i>Information Processing Systems 2023, NeurIPS 2023,</i>		
1102	<i>New Orleans, LA, USA, December 10 - 16, 2023</i> .		
1103	Haoyun Xu, Runzhe Zhan, Derek F. Wong, and Lidia S.	1156	
1104	Chao. 2024a. Let’s focus on neuron: Neuron-level su-	1157	
1105	pervised fine-tuning for large language model. <i>CoRR</i> ,	1158	
1106	abs/2403.11621.	1159	
		1160	
1107	Zhengqi Xu, Ke Yuan, Huiqiong Wang, Yong Wang,	1161	
1108	Mingli Song, and Jie Song. 2024b. Training-free	1162	
1109	pretrained model merging. <i>CoRR</i> , abs/2403.01753.		
1110	Prateek Yadav, Derek Tam, Leshem Choshen, Colin A.	1163	
1111	Raffel, and Mohit Bansal. 2023. Ties-merging: Re-	1164	
1112	solving interference when merging models. In <i>Ad-</i>	1165	
1113	<i>vances in Neural Information Processing Systems 36:</i>		
1114	<i>Annual Conference on Neural Information Process-</i>	1166	
1115	<i>ing Systems 2023, NeurIPS 2023, New Orleans, LA,</i>	1167	
1116	<i>USA, December 10 - 16, 2023</i> .	1168	
		1169	
1117	Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang,	1170	
1118	Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang,	1171	
1119	Dong Yan, Fan Yang, Fei Deng, Feng Wang, Feng		
1120	Liu, Guangwei Ai, Guosheng Dong, Haizhou Zhao,	1172	
1121	Hang Xu, Haoze Sun, Hongda Zhang, Hui Liu,	1173	
1122	Jiaming Ji, Jian Xie, Juntao Dai, Kun Fang, Lei	1174	
1123	Su, Liang Song, Lifeng Liu, Liyun Ru, Luyao Ma,	1175	
1124	Mang Wang, Mickel Liu, MingAn Lin, Nuolan Nie,	1176	
1125	Peidong Guo, Ruiyang Sun, Tao Zhang, Tianpeng	1177	
1126	Li, Tianyu Li, Wei Cheng, Weipeng Chen, Xian-	1178	
1127	grong Zeng, Xiaochuan Wang, Xiaoxi Chen, Xin		
1128	Men, Xin Yu, Xuehai Pan, Yanjun Shen, Yiding	1179	
1129	Wang, Yiyu Li, Youxin Jiang, Yuchen Gao, Yu-	1180	
1130	peng Zhang, Zenan Zhou, and Zhiying Wu. 2023.	1181	
1131	Baichuan 2: Open large-scale language models.	1182	
1132	<i>CoRR</i> , abs/2309.10305.		
1133	Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang,	1183	
1134	Xiaochun Cao, Jie Zhang, and Dacheng Tao. 2024.	1184	
		1185	
		1186	
		1187	
		1188	
		1189	
		1190	

1191	Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang,	
1192	Xiaolei Wang, Yupeng Hou, Yingqian Min, Be-	
1193	ichen Zhang, Junjie Zhang, Zican Dong, Yifan Du,	
1194	Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao	
1195	Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang	
1196	Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen.	
1197	2023. A survey of large language models. <i>CoRR</i> ,	
1198	abs/2303.18223.	
1199	Kun Zhou, Beichen Zhang, Jiapeng Wang, Zhipeng	
1200	Chen, Wayne Xin Zhao, Jing Sha, Zhichao Sheng,	
1201	Shijin Wang, and Ji-Rong Wen. 2024a. Jiuzhang3.0:	
1202	Efficiently improving mathematical reasoning by	
1203	training small data synthesis models. <i>CoRR</i> ,	
1204	abs/2405.14365.	
1205	Yuyan Zhou, Liang Song, Bingning Wang, and Weipeng	
1206	Chen. 2024b. Metagpt: Merging large language mod-	
1207	els using model exclusive task arithmetic. <i>CoRR</i> ,	
1208	abs/2406.11385.	
1209	Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi,	
1210	Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing	
1211	He. 2021. A comprehensive survey on transfer learn-	
1212	ing. <i>Proc. IEEE</i> , 109(1):43–76.	

A Empirical Study

A surge of work (Zhang et al., 2024b; Xiao et al., 2024; Tang et al., 2024) has pointed out that LLMs sparsely activate the specific sub-modules to perform corresponding tasks. Based on these findings, we conduct empirical experiments to explore whether the specific sub-module, which is related to advanced abilities, can be extracted and transferred. We utilize the forum corpus (*i.e.*, Zhihu for Chinese forum corpus and Reddit for English forum corpus) to continually pre-train LLMs, and then assess the training performance (*i.e.*, the value of loss function) and similarity of LLM neurons.

The forum corpus can be considered as containing the question-answering (QA) ability, which is necessary and important for LLMs. The results from Figure 4a have shown that only training the top 5% relevant neurons of LLMs can achieve the lower training loss and fit into the training set more quickly, indicating that LLMs contain the sub-module corresponding to the QA ability. Moreover, from Figure 4b and Figure 4c, we can observe that the LLM trained on Zhihu has shown higher similarity with the LLM trained on Reddit than the LLM trained on Github (*i.e.*, lower L1 Norm and higher cosine similarity), and the cosine similarity of different layers in LLM are largely different.

According to the above results, we have found that the different sub-networks of LLMs control the different abilities, and precisely selecting the correct sub-module of LLMs will help the extraction of advanced abilities from the single-lingual corpus and the transfer of these abilities to multi-lingual scenarios. Concretely, although Zhihu and Reddit are in different languages, they will influence the similar sub-modules of LLM and make these sub-networks show high similarity with each other. These sub-networks can be referred to the ability-related sub-networks, which are slightly influenced by languages.

B Implementation Details

In the experiment, we adapt LLaMA-3 8B as the backbone LLM, and employ Transformers (Wolf et al., 2020) and DeepSpeed () framework to perform the training process. And we also present the evaluation results of different backbone LLM (*i.e.*, Qwen2.5 0.5B (Hui et al., 2024) and Gemma2 2B (Rivière et al., 2024)) in Appendix E. For the training process, the learning rate, batch size, and training step are set as 5×10^{-5} , 1M tokens, and 2B

Concepts	Meaning
Key Neurons	Neuron refers to one of the trainable values of the tensors in LLMs. As previous work pointed out (Xu et al., 2024a), different neurons might control the different abilities of LLMs. Following this finding, in our work, we define the neurons that control the specific ability as the "Key Neurons". Key neurons can be regarded as a set without duplication, and a neuron belonging to the set means that this neuron can control the specific ability (Chen et al., 2024b). During the following training process, only the neurons belonging to the key neurons will be trained and optimized.
Ability-related Weights	Ability-related weights refer to the value of the whole neuron in LLM, which can represent the corresponding ability of LLM (Yu et al., 2024a; Ilharco et al., 2023). In MAET, we obtain the ability-related weights through equation 2. The ability-related weights contain the value of all neurons. Since only the key neurons will be trained during the training process, the value of the neurons not belonging to key neurons is zero in the ability-related weights.
Ability-related Tensors	Ability-related tensors can be regarded as a set of LLM tensors, which is related to the corresponding ability. Previous work has studied how the LLM layers influence the ability (Cheng et al., 2024). Different from key neurons, ability-related tensors focus on higher-level information, integrating the sparse neurons into a coarser-grained element (Xiao et al., 2024). A tensor belonging to the ability-related tensors denotes that this tensor is highly related to the corresponding ability and can control this ability.
Language-specific Weights	Similar to the ability-related weights, language-specific weights also refer to the value of the whole neurons in LLMs (Zhang et al., 2024b). However, language-specific weights represent the language abilities of LLM that include multiple abilities (i.e., one language can be regarded as one ability) (Tang et al., 2024), and the method of obtaining them is also different from ability-specific weights. In MAET, we first calculate the ability-related weights of each language and then Integrating these weights together to obtain the language-specific.

Table 4: The key concepts of our approach.




Figure 4: The results of empirical experiments. We present the loss of different training methods during the training process, the cosine similarity of LLM layers after being trained on Zhihu and Reddit, and the similarity of LLMs being trained on different training corpus.

tokens, respectively. Besides, for the key neurons locating, we select the top 5% relevant neurons as the key neuron set \mathcal{N} for both stages and identify the last 80% and 60% similar tensor as the key sub-network \mathcal{T} for mathematical reasoning tasks and scientific reasoning tasks respectively. Besides, we released all of the hyper-parameters during our experiment to reproduce our proposed approach better. Table 5 shows the details of hyper-parameters of different stages.

C Details of Dataset

We present the statistical information of the datasets in Table 6. We mainly consider English, Spanish, Chinese, Bengali, and Telugu in our experiment, and utilized English as the in-domain

language while others as the out-of-domain languages. For the evaluation datasets, we select MGSM and multi-lingual MMLU as the evaluation benchmarks, which contain the parallel data in different languages and are useful for multi-lingual complex tasks evaluation.

D Prompt for Translation

You should translate the following text from English to {TARGET LANGUAGE} and should not modify the latex code or website code. You should not add any details that are not mentioned in the original text.

English

Algorithm 1: The complete procedure of our proposed approach.

Input : Single-lingual ability-related corpus C_{L_0, A_i} , multi-lingual general corpus $C_{L_0}, C_{L_1}, \dots, C_{L_n}$, and the parameters of the backbone model Θ_o .
Output : A well-trained multi-lingual ability-enhanced LLM.

```

// Ability-related Weights Extraction
1  $\theta' \leftarrow \text{CPT}(C_{L_0, A_i}, \Theta_o)$ ;
2 for  $j$ -th neuron in  $\Theta_o$  do
3   | Calculate the importance score of the corresponding neuron using Eq. 1;
4   Identify the key neuron set  $\mathcal{N}_{A_i}$ ;
5    $\text{LLM}_{A_i, L_0} \leftarrow \text{CPT}(C_{L_0, A}, \Theta_o, \mathcal{N}_{A_i} \cup \mathcal{N}_{L_0})$ ;
6    $\text{LLM}_{L_0} \leftarrow \text{CPT}(C_{L_0}, \Theta_o, \mathcal{N}_{L_0})$ ;
7   Learning the ability-related weight  $R(A_i)$  using Eq. 2;

// Multi-lingual Ability Transfer
8   Obtaining the multi-lingual weight  $R_{Lang}$  using Eq. 3;
9   for  $j$ -th parameter tensor in LLM do
10    | Calculate the correlation using Eq. 4;
11   Identify the ability-related parameters  $\mathcal{T}$ ;
12   Transfer the ability to multi-lingual scenarios using Eq. 5;
13   Obtain the well-trained multi-lingual ability-enhanced LLM.

```

{ENGLISH TEXT}

{TARGET LANGUAGE}

E Performance of Small Scale LLMs

We conduct the different LLMs with different sizes (*i.e.*, Qwen2.5-0.5B and Gemma2-2B) in our experiment to valid the practicality of our approach. We assess MAET and baselines on multi-lingual scientific reasoning tasks and present the evaluation results in Table 7. Comparing the performance of MAET and the baseline methods, we can observe that MAET can also enhance the performance of small scale models and outperform competitive baselines. Therefore, the evaluation results have shown the effectiveness of MAET and verified that MAET is a general LLM enhancement technology.

F Ability-related Sub-networks of LLM

To assess and probe the ability-related sub-networks of LLMs, we only transfer the specific tensors (*i.e.*, tensors in self-attention and MLP mechanism) from the ability weight to the final models through Eq. 5, to analyze the LLM inner abilities. The experimental results are presented in Table 8. From the experiment, we can observe that although the proportion of MLP layers (41.38%) is lower than the attention layers (45.26%), only trans-

ferring the MLP layers outperforms transferring the attention layers, indicating that the MLP layers are more related to the advanced abilities and stores the corresponding knowledge. In the MLP layers of LLM, the gate mechanism (*i.e.*, MLP Gate) will control the transmission of information and the down project mechanism (*i.e.*, MLP Down) will integrate the knowledge from previous layers, so that transferring the MLP layers can achieve better performance on the downstream tasks.

Stage	Hyper-Parameter	Mathematical Tasks	Scientific Tasks
Extracting	Learning Rate	5×10^{-5}	5×10^{-5}
	Batch Size	1M Tokens	1M Tokens
	Training Steps	2B Tokens	2B Tokens
	α in Extraction	0.8	0.8
	β in Extraction	0.2	0.2
	Ratio of Key Neurons	5%	5%
Transferring	Learning Rate	5×10^{-5}	5×10^{-5}
	Batch Size	1M Tokens	1M Tokens
	Training Steps	2B Tokens	2B Tokens
	γ in Transferring	0.2	0.2
	η in Transferring	1.0	1.0
	Ratio of Key Neurons	80%	60%
	μ for Spanish	1.5	1.5
	μ for Chinese	2.0	2.0
	μ for Bengali	1.2	1.2
	μ for Telugu	1.2	1.2

Table 5: The details of hyper-parameters in the training and evaluation process.

Language	Training Dataset (Tokens)		Evaluation Dataset (Instances)	
	General Corpus	Ability-related Corpus	Mathematical Tasks	Scientific Tasks
English	1.81B	1.30B (Math) / 1.82B (Sci)	250	1,245
Spanish	1.81B	-	250	1,232
Chinese	1.80B	-	250	1,229
Bengali	1.81B	-	250	1,137
Telugu	1.81B	-	250	1,036

Table 6: The statistical information of the training and evaluation datasets.

Methods	Qwen2.5 0.5B			Gemma2 2B		
	ES	TE	Avg.	ES	TE	Avg.
Backbone LLM	36.64	25.69	31.17	43.41	30.01	36.71
+ F-CPT _{L&A}	32.90	22.43	27.67	38.48	30.39	34.62
+ F-CPT _A	32.62	25.26	28.94	37.83	25.39	31.61
+ MAET w/o API	36.72	28.91	32.82	43.23	29.59	36.41
+ MAET (Ours)	36.91	29.62	33.27	43.62	30.37	37.00

Table 7: The performance comparison of different LLMs on multilingual scientific tasks.

LLM Tensors	Proportion of \mathcal{T}	ES	ZH	BN	TE	Avg.
All Tensors	100.00%	49.60	41.60	32.40	25.20	37.20
Attention All	45.26%	48.80	41.60	28.80	26.40	36.40
Attention Q	12.07%	47.60	40.80	30.80	26.40	36.40
Attention K	10.34%	47.20	42.40	29.60	24.40	35.90
Attention V	9.48%	47.60	42.40	28.80	25.20	36.00
Attention O	13.36%	48.00	40.40	30.80	27.20	36.60
MLP All	41.38%	48.80	39.60	31.60	27.60	36.90
MLP Up	13.79%	50.00	40.00	28.80	25.20	36.00
MLP Gate	13.79%	46.00	41.20	30.00	24.00	35.30
MLP Down	13.79%	49.60	41.60	30.40	26.00	36.90

Table 8: The effect of only merging the specific LLM tensors during the transferring process (*i.e.*, Eq.5) on multilingual mathematical tasks.