

Enhancing Non-English Capabilities of English-Centric Large Language Models Through Deep Supervision Fine-Tuning

Wenshuai Huo^{1,2}, Xiaocheng Feng^{1,2}*, Yichong Huang¹, Chengpeng Fu^{1,2}, Baohang Li¹, Yangfan Ye¹, Zhirui Zhang³, Dandan Tu³, Duyu Tang³, Yunfei Lu³, Hui Wang², Bing Qin^{1,2}*

¹Harbin Institution of Technology

²Pengcheng Laboratory

³Huawei Technologies Co., Ltd

{wshuo, xcfeng, ychuang, cpf, baohangli, yfy, qinb}@ir.hit.edu.cn, wangh06@pcl.ac.cn
{tudandan, tangduyu, luyunfei6}@huawei.com, zrutc11@gmail.com

Abstract

Large language models (LLMs) have demonstrated significant progress in multilingual language understanding and generation. However, due to the imbalance in training data, their capabilities in non-English languages are limited. Recent studies revealed the English-pivot multilingual mechanism of LLMs, where LLMs implicitly convert non-English queries into English ones at the bottom layers and adopt English for thinking at the middle layers. However, due to the absence of explicit supervision for cross-lingual alignment in the intermediate layers of LLMs, the internal representations during these stages may become inaccurate. In this work, we introduce a deep supervision fine-tuning method (DFT) that incorporates additional supervision over the internal layers of the model to guide its workflow. Specifically, we introduce two training objectives on different layers of LLMs: one at the bottom layers to constrain the conversion of the target language into English, and another at the middle layers to constrain reasoning in English. To effectively achieve the guiding purpose, we designed two types of supervision signals: logits and feature, which represent a stricter constraint and a relatively more relaxed guidance. Our method guides the model to not only consider the final generated result when processing non-English inputs but also ensure the accuracy of internal representations. We conducted extensive experiments on typical English-centric LLMs, LLaMA-2 and Gemma-2. The results on 8 multilingual datasets show that our method significantly outperforms traditional fine-tuning methods.

Introduction

The development of large language models has achieved revolutionary breakthroughs in the field of natural language processing, demonstrating exceptional performance (OpenAI 2023; Touvron et al. 2023; Team et al. 2023). However, performance disparities still exist across different languages due to the imbalance in the training data (Nguyen et al. 2023; Huang et al. 2023; Zhu et al. 2023a). For instance, the well-known LLaMA series models are trained on over 90% English data. This results in a significant gap be-

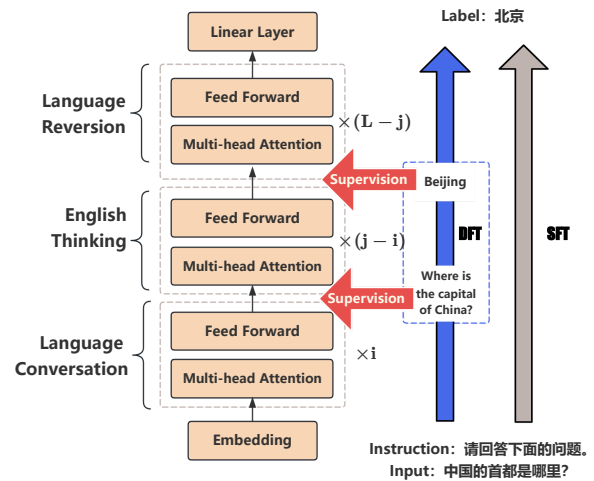


Figure 1: The illustration of Depth Supervision Fine-Tuning (DFT) and Baseline Methods. The left side represents an English-dominated large language model, which can be divided into three parts from shallow to deep layers: Language Conversation, English Thinking, and Language Reversion. The right side shows an sample of Chinese instruction tuning. Blue arrows represent the DFT method, while gray arrows represent the Baseline method. Traditional fine-tuning methods focus only on the model predicting the corresponding target output based on the input instruction. In contrast, our method adds supervision to the process, explicitly guiding the model’s workflow when processing non-English inputs.

tween their multilingual capabilities and their performance in English.

To enhance the non-English processing capabilities of these models, researchers have made numerous attempts. A common and effective approach is to fine-tune pre-trained models on instruction datasets in the target languages (Zhu et al. 2023b; Li et al. 2023; Zhang et al. 2024; Li et al. 2024b; Zhao et al. 2024a). As research advances, some studies have explored the internal mechanisms behind how large models

*Corresponding author

handle multilingualism, revealing an English-centric workflow when processing non-English inputs. Specifically, Zhao et al. (2024b) analyzed activation patterns in specific languages and inferred that the bottom layers of the model convert input from various languages into English, while the top layers perform the reverse conversion. Wendler et al. (2024) found that by decoding early in the model’s intermediate layers, rather than the final layer, LLMs tend to use English as an internal pivot language when processing multilingual inputs. These studies suggest that LLMs dominated by English can be divided into three stages when processing non-English inputs: converting non-English to English (or an English representation space), thinking in English, and converting English back to the target language. However, due to the absence of explicit supervision for cross-lingual alignment in the intermediate layers of LLMs, the internal representations during these stages may become inaccurate. To the best of our knowledge, no studies have yet proposed targeted optimization schemes for this phenomenon.

To tackle the aforementioned problem, we propose the deep supervision fine-tuning method (DFT) that aims to explicitly guide the LLMs’ performance in the three stages when processing non-English inputs. Specifically, DFT incorporates additional supervision over the internal layers of the model. As shown in Figure 1, unlike traditional fine-tuning methods, DFT constrains not only the final output of the model but also the intermediate process. In the bottom layers of the model, DFT guides the conversion from non-English to English. In the middle layers, DFT guides the model to obtain answers in the English space. To effectively constrain the intermediate process, we propose two supervision schemes based on logits and features. Additionally, to accurately identify the critical layers at different stages, we propose an entropy-based selection strategy.

We conducted extensive experiments on 8 commonly used multilingual benchmarks and the results demonstrate the effectiveness of our method. Specifically, for multilingual QA tasks where both input and output are in the target language, our method achieved significant improvements. Our contributions can be summarized as follows.

Background

Instruction Fine-Tuning. The technology of instruction tuning enables pre-trained LLMs to comprehend instructions and handle downstream tasks effectively via the training on an annotated instruction dataset $D = \{(x_i, y_i)\}_{i=1}^N$, where x is the input question and y is the expected output answer. The training objective is to minimize the following negative log-likelihood:

$$\mathcal{L}(\theta) = \sum_{i=1}^N -\log P(y_i|x_i; \theta), \quad (1)$$

where θ denotes the learnable parameters of the model. However, most instruct datasets are in English, which limits the potential of large models to address tasks in non-English languages.

Multilingual Instruction Fine-Tuning. To enhance the multilingual capabilities of large language models, a common approach is to translate English instruction data into the target language (tgt), creating a dataset $D^{tgt} = \{(x_i^{tgt}, y_i^{tgt})\}_{i=1}^N$. The model is then fine-tuned using the translated instruction data. Similar to equation 1, the loss function \mathcal{L}_{TFT} is:

$$\mathcal{L}_{TFT}(\theta) = \sum_{i=1}^N -\log P(y_i^{tgt}|x_i^{tgt}; \theta). \quad (2)$$

In this work, we also fine-tune the model using multilingual instruction datasets constructed through translation. However, unlike traditional fine-tuning methods, we focus not only on the model’s final predictions but also impose constraints to ensure the model achieves accurate intermediate results.

Deep Supervision Networks. Deep Supervision Networks (DSNs) represent a special training strategy that optimizes the training process by incorporating supervision at multiple intermediate layers of the model. Unlike traditional deep learning architectures that compute the loss only at the final layer L , DSNs also calculate loss at several hidden layers k . The total loss function \mathcal{L} can be expressed as:

$$\mathcal{L}(\theta) = \mathcal{L}_L(\theta) + \alpha \mathcal{L}_k(\theta_{1:k}), \quad (3)$$

$\theta_{1:k}$ denotes only the parameters of layer k to layer 1 are learned and updated by minimising the intermediate loss (\mathcal{L}_k); α is a hyper-parameter to control the balance between the intermediate supervision loss (\mathcal{L}_k) and the final output layer loss (\mathcal{L}_L).

In this work, we take inspiration from the training strategy of DSNs by introducing additional supervision over the intermediate layers of the model. Unlike traditional DSNs, where the intermediate supervision is consistent with the final output target (Li et al. 2022), our additional supervision is differ from the final output. The goal is to guide the workflow of LLMs when processing non-English inputs.

Method

In this section, we describe the proposed deep supervision-based fine-tuning approach. Figure 2 illustrates the overall process of our method. The process of handling non-English inputs in an LLM can be roughly divided into three stages from the bottom layers to the top layers: (1) Language Conversion: The model interprets the non-English query and converts the multilingual input into English (or an English representation space); (2) English Thinking: The model employs English for thinking and solving the task; (3) Language Reversion: The model converts the reasoning results back into the target language, consistent with the input. We propose DFT to guide the model’s internal information transformation. Specifically, we introduce additional supervision at the top and middle layers of the model to enhance the model’s ability to convert non-English inputs into English and to improve its logical reasoning in English. We propose two types of supervision: logits-based and feature-based.

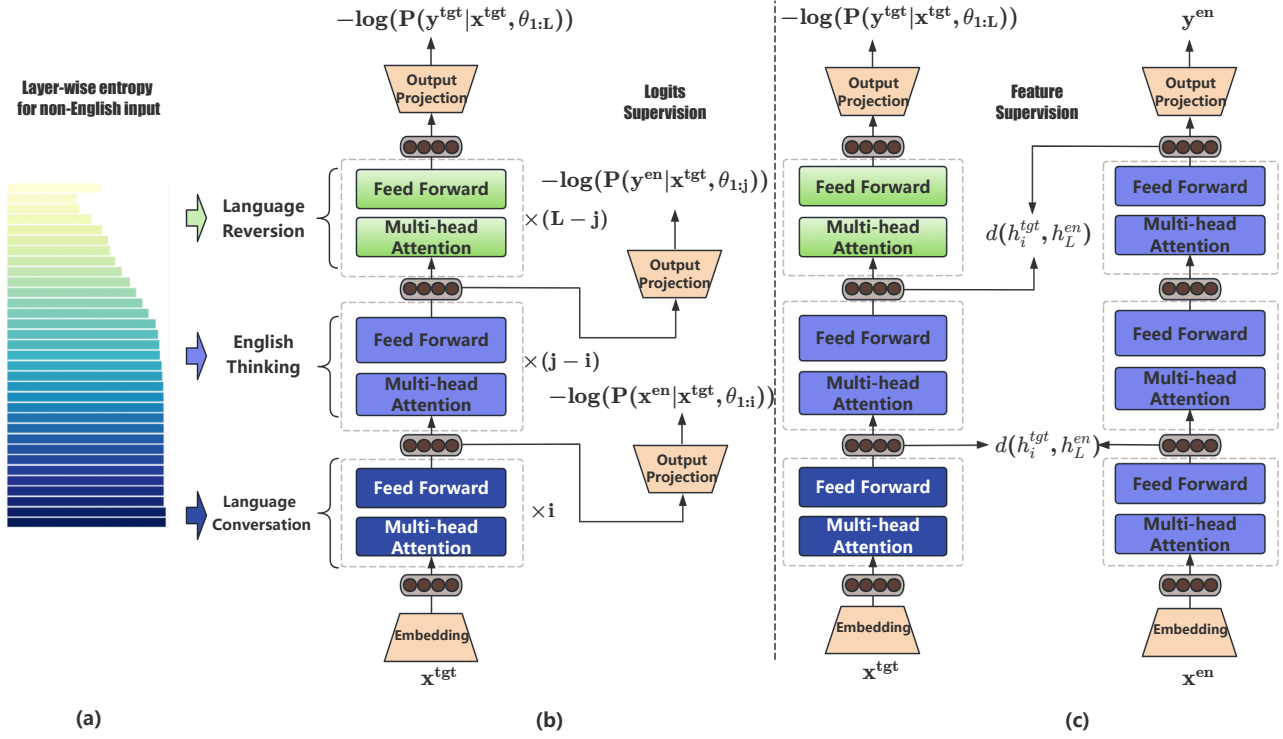


Figure 2: The illustration of the proposed methods DFT-logits (b) and DFT-feature (c). The heatmap (a) represents the entropy values of each layer in English-dominated Large Language Model when processing non-English inputs. The process of handling non-English inputs in an LLM can be roughly divided into three stages from the bottom layers to the top layers: Language Conversion, English Thinking and Language Reversion.

Language Conversion Constraints

When processing non-English inputs, the initial layers of large language models are primarily tasked with converting the input query into an internal English representation that the model can effectively process. Due to the significant differences in grammar, vocabulary, and linguistic structures across languages, this conversion process is inherently challenging. Failure to accurately perform this conversion can lead to semantic biases or errors in subsequent reasoning and generation stages. Thus, it is critical that the model excels in language conversion within its early layers.

To enhance the model’s language conversion capabilities, we propose two constraint methods: logits supervision and feature supervision.

Logits-Based: For transformer-based LLMs, each layer outputs a hidden representation of shape $(batch_size \times sequence_length \times hidden_size)$. The typical training strategy feeds the hidden representation from the final layer into a linear matrix W^{out} of shape $(hidden_size \times vocab_size)$, projecting the final hidden representation into the vocabulary space and predicting the probability distribution over the vocabulary using softmax. This distribution is then compared with the ground truth to calculate the loss. However, since the hidden states have the same shape across all layers, we can apply the W^{out} and softmax operations at

any layer to make predictions. To enforce language conversion, we apply W^{out} at an internal layer i in the model (as shown in Figure 2b) and compute the probability distribution. Then calculate the loss with a parallel English query, thereby forcing the model to convert the target language query into its English version in the early layers. Where the logits-based language conversion loss \mathcal{L}_{LC} is computed as:

$$\mathcal{L}_{LC} = -\log P(x^{en}|x^{tgt}; \theta_{1:i}, W^{out}), \quad (4)$$

$\theta_{1:i}$ represents the parameters of the first i layers of the model. We use the LLM’s head as W^{out} and restrict W^{out} to only participate in the forward pass during the language conversion stage, without being updated.

Feature-Based: Another approach to enhancing the language conversion process in the first stage is through feature alignment. During instruction tuning in the target language, semantically equivalent English queries are also fed into the model. The hidden representations of these inputs are then extracted, and a similarity loss is calculated to ensure alignment between the English and non-English representations. The feature-based language conversion loss function \mathcal{L}_{LC} can be defined as:

$$\mathcal{L}_{LC} = 1 - \left(\frac{h_i^{en} \cdot h_i^{tgt}}{\|h_i^{en}\| \|h_i^{tgt}\|} \right), \quad (5)$$

where i represents the critical layer in the first stage, h^{en} and h^{tgt} represent the hidden states of the same semantic input in English and the target language, respectively. Compared to the logits-based supervision strategy, feature supervision is relatively more relaxed. It does not require the target language input to be strictly converted into English but instead aligns the representations across languages.

By applying constraints at this stage, non-English inputs can be more accurately transformed into English, thereby laying a strong foundation for subsequent reasoning and output generation.

English Thinking Constraints

After language conversion, the model needs to further process the input information and reason how to provide an appropriate response. Therefore, we add supervisory signals to constrain the model’s reasoning ability in English. Similar to the previous section, we implement with both logits constraint and feature constraint methods.

Logits-Based: For the logits constraint, we output probabilities from the hidden vectors at the end of the second stage and compute the loss, aiming for the model to provide corresponding English answers based on the input target question. The logits-based english thinking loss function \mathcal{L}_{ET} is formulated as follows:

$$\mathcal{L}_{ET} = -\log P(y^{en}|x^{tgt}; \theta_{1:j}, W^{out}), \quad (6)$$

where j represents the critical layer in the second stage responsible for reasoning in English. This constraint forces the model to focus on producing the correct English reasoning path, ultimately leading to accurate final outputs.

Feature-Based: For the feature constraint, as shown in the Figure 2, this is implemented by feeding semantically equivalent English and non-English questions into the model. For the English input, we extract the features from the top layers of the model, while for the non-English input, we extract the features from the critical layer of the English Thinking stage, where the reasoning in English is assumed to occur. The alignment of these features is crucial for ensuring that the model’s internal representations remain consistent across language.

$$\mathcal{L}_{ET} = 1 - \left(\frac{h_L^{en} \cdot h_j^{tgt}}{\|h_L^{en}\| \|h_j^{tgt}\|} \right), \quad (7)$$

where L represents the final layer, j represents the critical layer in the second stage. This constraint helps the model to align its reasoning process in English with the semantic content derived from the target language input, ensuring that the reasoning remains accurate and consistent in the model.

Training Objective

The total loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{TFT} + \mathcal{L}_{LC} + \mathcal{L}_{ET}, \quad (8)$$

where \mathcal{L}_{LC} is either logits-based or feature-based, and \mathcal{L}_{ET} is also either logits-based or feature-based.

Experiment

Setup

We use LLaMA-2-7B (Touvron et al. 2023) and Gemma-2-2B (Team 2024) as the base models. The training data consists of Stanford Alpaca instruction data (Taori et al. 2023) and its translations in the target languages, which include Chinese (zh), Vietnamese (vi), and Arabic (ar). For the translated data, we directly used publicly available datasets from (Zhu et al. 2023b). Our code implementation is based on stanford_alpaca¹. All experiments were conducted on 8 × A100 GPUs with a batch size of 128. The models were trained for 3 epochs with a learning rate of 2e-5. To accelerate training, we utilized the FSDP training strategy (Zhao et al. 2023).

Comparison of Methods

- **SFT** (Ouyang et al. 2022), which is instruction-tuned with English instruction datasets.
- **TFT** (Zhu et al. 2023b), which is instruction-tuned using the original English instruction datasets translated into the target languages.
- **SDRRL** (Zhang et al. 2024), which is a method based on Self-Distillation. Besides using English instruction-tuning data and its multilingual code-switching extensions, it also incorporates partially translated data and completion data for fine-tuning.
- **DFT-logits**, our method that applies logits-based supervision to guide the model’s intermediate layers.
- **DFT-feature**, our method that uses feature alignment to maintain consistent internal representations between English and tgt language.

Evaluation Dataset

- **XQUAD (Cross-lingual Question Answering Dataset):** XQUAD (Artetxe, Ruder, and Yogatama 2019) is a high-quality cross-lingual question answering dataset containing 240 paragraphs and 1,190 question-answer pairs, which have been manually translated into 10 languages.
- **MLQA (Multilingual Question Answering):** MLQA (Lewis et al. 2019) is a multilingual question answering dataset covering 7 languages. Each question in the dataset is accompanied by a paragraph and an answer in the corresponding language.
- **MKQA (Multilingual Knowledge Questions and Answers):** The MKQA (Longpre, Lu, and Daiber 2021) dataset contains 2,600 common sense question-answer pairs across 26 languages.
- **TruthfulQA:** TruthfulQA (Lin, Hilton, and Evans 2021) includes questions from various domains, specifically designed to test the truthfulness and accuracy of models when answering complex questions.
- **XNLI (Cross-lingual Natural Language Inference):** XNLI (Conneau et al. 2018) is a widely used language

¹https://github.com/tatsu-lab/stanford_alpaca.

	Method	xquad			mlqa			truthfulQA			mkqa		
		vi	ar	zh	vi	ar	zh	vi	ar	zh	vi	ar	zh
<i>Performance on LLaMA-2-7b</i>													
Baselines	SFT	23.11	19.42	3.34	22.79	19.18	6.91	29.37	29.11	30.49	33.16	32.84	33.51
	TFT	27.36	21.23	23.41	28.80	21.62	23.86	28.54	27.04	26.40	38.59	35.02	40.70
	SDRRL	28.79	25.75	22.01	29.59	27.26	24.67	28.03	28.11	28.30	39.20	35.81	40.63
Ours	DFT-logits	29.65	27.57	24.77	31.54	28.26	25.68	28.64	28.81	28.93	39.64	35.84	40.87
	DFT-feature	28.27	26.35	24.34	32.19	29.51	25.29	28.61	28.34	28.00	39.36	35.65	40.90
<i>Performance on Gemma-2-2b</i>													
Baselines	SFT	19.02	18.82	9.74	20.21	20.45	11.79	30.45	28.98	27.79	36.04	33.69	32.46
	TFT	17.58	24.73	15.95	20.26	25.31	18.26	26.44	26.26	24.87	36.84	34.49	40.20
	SDRRL	17.55	18.16	12.03	19.89	23.91	14.24	26.62	26.52	26.39	37.69	34.51	40.75
Ours	DFT-logits	21.23	25.79	17.37	24.36	27.51	20.13	27.34	26.78	27.16	38.59	34.81	40.84
	DFT-feature	22.78	27.38	16.94	22.19	30.18	19.51	27.88	27.64	26.90	38.04	34.73	40.87

Table 1: Results of baselines and our method on multilingual question and answer benchmark. Bold indicates the best result of all methods. Our method outperforms the baselines in almost all languages.

	Method	xnli			xcopa		xstory_cloze		mmlu		
		vi	ar	zh	vi	zh	ar	zh	vi	ar	zh
<i>Performance on LLaMA-2-7b</i>											
Baselines	SFT	36.59	35.38	36.39	63.00	65.00	49.70	59.56	29.37	27.6	30.49
	TFT	44.14	33.73	36.87	66.00	65.00	57.51	64.06	33.23	28.05	32.34
	SDRRL	43.41	33.98	37.43	65.40	65.20	58.17	64.53	31.91	27.47	32.74
Ours	DFT-logits	43.37	34.06	36.94	66.20	66.40	58.78	64.93	32.60	29.06	33.19
	DFT-feature	43.50	34.38	37.64	66.60	68.60	58.78	64.40	32.19	28.86	32.82
<i>Performance on Gemma-2-2b</i>											
Baselines	SFT	41.33	35.62	40.36	68.60	70.40	61.02	68.23	36.07	31.16	36.22
	TFT	43.05	33.94	36.63	67.00	70.40	61.48	64.88	36.95	32.71	36.37
	SDRRL	42.25	36.31	39.20	65.40	66.60	61.09	63.20	34.47	31.31	34.23
Ours	DFT-logits	43.29	35.14	38.46	68.40	70.60	62.08	66.91	36.89	31.84	35.30
	DFT-feature	43.25	36.73	39.74	67.20	70.20	61.88	66.25	36.70	31.63	34.43

Table 2: Results of baselines and our method on multilingual understanding benchmark. Our method outperforms the baselines in almost all languages.

understanding dataset to evaluate models’ performance in cross-lingual inference tasks.

- **XCOPA (Cross-lingual Choice of Plausible Alternatives):** XCOPA (Ponti et al. 2020) is a benchmark designed to evaluate the ability of models to apply commonsense reasoning, requiring both world knowledge and the ability to generalize it to new languages.
- **XStoryCloze (Cross-lingual Story Cloze Test)** XStoryCloze (Lin et al. 2022) is a cross-lingual dataset for evaluating models’ ability to understand stories and generate plausible endings.
- **MMLU (Massive Multitask Language Understanding)** MMLU (Hendrycks et al. 2020) is a large-scale multitask language understanding dataset covering multiple domains (such as history, geography, science, law, etc.) and various languages.

For all evaluation datasets, we conducted tests using a zero-shot setting. We used the F1 score for XQuAD, MLQA, and MKQA, and the MC1 metric for TruthfulQA. For other NLU datasets, accuracy was used as the evaluation metric.

Main Results

Table 1 and Table 2 present the results on multilingual QA and NLU tasks, respectively. From the experimental results, we can observe that: (1) Our method outperforms the baselines in almost all languages for both QA and NLU tasks. This indicates that our approach successfully enhances the model’s capabilities in the target language by guiding the internal workflow. (2) The improvement is more pronounced in QA tasks, as our method is specifically designed for tasks where both the input and output are in the target language, making it better suited for QA scenarios in the target language. (3) Fine-tuning the model solely on English instruc-

	TFT	+LC	+ET
XQUAD	23.41	24.75 / 23.99	24.47 / 24.26
MLQA	23.86	25.62 / 24.27	25.68 / 24.80
TruthfulQA	26.40	27.11 / 27.51	28.93 / 28.00
MKQA	40.70	40.73 / 40.81	40.83 / 40.70
XNLI	36.87	36.27 / 36.59	37.61 / 37.75
XCOPA	65.00	65.60 / 66.80	66.40 / 68.40
Xstory_Cloze	64.06	63.60 / 64.92	64.95 / 64.26
MMLU	32.34	32.75 / 33.73	33.29 / 32.85

Table 3: Results on 8 Chinese evaluation datasets with separately added language conversion supervision (LC) and english thinking supervision (ET). The supervision types are represented in “logits-based / feature-based” form.

tion data (SFT) outperformed all other results fine-tuned on target language instruction datasets for TruthfulQA. This suggests that for the MC1 metric in TruthfulQA, generation capabilities in the target language are less important. (4) We observed that DFT (feature) performs better than DFT (logits) on understanding tasks, possibly because the stricter logits-based supervision is more suitable for generation tasks, whereas the feature-based supervision offers better generalization across different types of tasks.

Ablation Study

We further analyzed the effects of applying supervision over either the first or the second stage. We compared the performance of separately adding Language Conversion supervision (LC) and English Thinking supervision (ET) on 8 Chinese datasets. The results, based on LLaMA2, are shown in the Table 3, where the supervision types are represented as “logits-based / feature-based.”

From the experimental results, we can observe that: (1) Adding supervision during the English Thinking stage yielded more significant improvements, indicating that aligning responses has a greater impact. (2) Logits-based supervision led to greater improvements in QA tasks, while feature-based supervision was more beneficial for understanding tasks. This may be because generation tasks require stricter supervision signals. (3) Logits-based language conversion supervision caused performance drops in some tasks, suggesting that strong supervision over the earlier layers may harm the model’s original capabilities.

Entropy-Based Critical Layer Selection

To implement our approach, it is crucial to accurately identify the critical layers that separate the different stages, particularly when determining the $i - th$ and $j - th$ layers as shown in the Figure 2. Although previous works (Wendler et al. 2024; Zhao et al. 2024b) have discovered the mechanisms of information transformation within models when processing non-English inputs, identifying the critical layers at different stages remains challenging.

Entropy, in the context of information theory, is a measure of uncertainty or randomness in the information being processed. In neural networks, entropy can help identify where

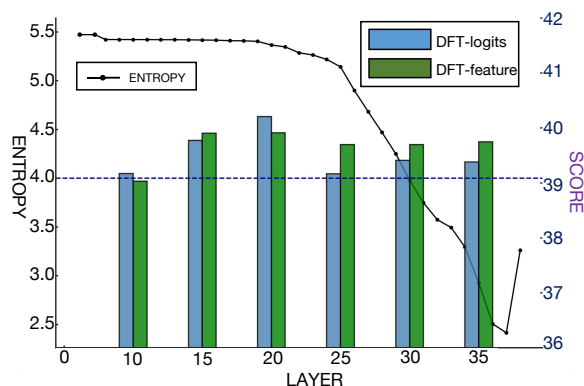


Figure 3: The bars in the figure represent the results of DFT-logits and DFT-feature on all evaluation datasets, with English Thinking supervision applied at different layers. The target language is Chinese, and the base model used is LLaMA-2-7b. The dashed line indicates the results of the TFT method. The average scores across various datasets are reported. The broken line represents the change in entropy as the layer depth increases.

significant transformations or reductions in uncertainty occur. We observed that when LLMs processes non-English inputs, there are two significant drops in entropy (as shown in Figure 2 (a)) . These drops indicate key points where the model undergoes substantial information transformation. Based on these observations, we hypothesize that the initial entropy drop corresponds to the model processing the non-English input into an English representation that it can handle. After this, the model transitions to the English reasoning and processing stage. The subsequent entropy drop marks the model gradually completing the reasoning process and progressively forming the final output.

To validate the effectiveness of our hypothesis, we applied English Thinking supervision at layers 5, 10, 15, 20, 25, 30 of LLaMA-2-7b and fine-tuned the model on the Chinese instruction dataset. Figure 3 shows the average results across 8 Chinese datasets using the DFT-logits and DFT-feature methods. We observed that: (1) Entropy first drops at layer 2, stabilizes for a period, and then begins to drop again around layer 15. (2) Both DFT-logits and DFT-feature achieve better performance when applied at layer 15, indicating that our hypothesis is valid. (3) For the DFT-logits method, performance declines when supervision is applied at later layers, possibly because the model has already started the language reversion process (converting English back to the target language). Adding constraints to predict English results at this stage may interfere with the model’s generation of the target language. In contrast, the relatively more relaxed DFT-feature method performs better at later layers.

Although this hypothesis is rough (as different tokens may exhibit different behaviors), it provides us with guidance for selecting the critical layer at each stage.




Figure 4: t-SNE visualizations of sentence representations from FLORES-200 dataset by LLaMA-2 before and after applying DFT.

Analysis of Representation Alignment

We used the t-SNE (Van der Maaten and Hinton 2008) method to visualize the representations of input sentences to analyze the impact of DFT on aligning cross-lingual representations.

Specifically, we encoded parallel English and Chinese sentences from the FLORES-200 dataset and obtain sentence representations by the mean pooling method using the representation for each token.

The results are shown in Figure 4. In the vanilla model, the representations of the two languages are far apart. After applying the DFT method, they become more aligned. This indicates that our method can help bring the target language representations closer to the English representations.

Analysis of Translation Task

The model’s performance on translation tasks can reflect its overall language conversion and generation capabilities. Therefore, although cross-lingual generation tasks do not align with the workflow of our method, we still analyzed our approach on translation tasks. We compared the performance of separately adding Language Conversion supervision and English Thinking supervision, and the results are shown in Table 3.

From the experimental results, we can observe that: (1) Adding logits-based supervision led to a catastrophic drop in the en-zh direction. This suggests that strict supervision of target language conversion into English within the model’s internal layers leads to a decline in the model’s overall ability to convert English to the target language. (2) Adding feature-based supervision significantly improved translation results for both en-zh and zh-en directions, indicating that aligning representations between languages is beneficial for cross-lingual tasks.

In fact, our method is better suited for scenarios where both the input and output are in the target language, rather than for cross-lingual tasks. Nevertheless, DFT-feature still achieved strong performance on translation tasks, indicating that our approach has broad potential for application.

Related Work

Aligning Non-English Capabilities of Large Language Models To enhance the non-English capabilities of LLMs,

	TFT	+LC	+ET
en-zh	68.37	36.79 / 75.62	46.50 / 74.49
zh-en	63.12	64.10 / 82.58	63.68 / 83.98

Table 4: Translation performance on FLORES-200 with separately added language conversion supervision (LC) and english thinking supervision (ET), evaluated using COMET score.

researchers have explored several approaches. Pre-training LLMs on diverse multilingual datasets has proven effective in improving multilingual performance. However, this approach requires the collection of large amounts of data and significant computational resources (Le Scao et al. 2023; Cui, Yang, and Yao 2023; HIT-SCIR 2024). Instruction fine-tuning on translation datasets has also been successful in enhancing non-English performance (Li et al. 2023; Zhu et al. 2023b; Zhang et al. 2024; Xu, Li, and Xiong 2023; Li et al. 2024b). In the inference stage, cross-lingual transfer methods, such as leveraging knowledge from resource-rich languages and using self-translation prompts, have been effective (Qin et al. 2023; Xu, Li, and Xiong 2023; Huang et al. 2023). Additionally, Li et al. (2024a) aligned English representations to enable the model to fully leverage its English capabilities when processing non-English inputs.

Deep Supervision Networks Lee et al. (2015) first proposed the deeply supervised network, where auxiliary classifiers are added on various intermediate layers, and each classifier contributes to the overall loss during training. Huang et al. (2022) enhanced machine translation performance by predicting outputs layer by layer in a non-autoregressive manner. Elbayad et al. (2019) improved decoding efficiency by making predictions at different layers based on token prediction differences.

Our method draws on the implementation of deeply supervised networks by introducing supervision over the middle part of the model. Our goal is to guide the internal information transformation process within the model.

Conclusion and Future Work

In this work, we propose Deep supervision Fine-Tuning, effectively enhancing the multilingual capabilities of English-dominated LLMs. Our method guides the workflow of LLMs when processing non-English inputs by adding cross-lingual supervision over intermediate layers, constraining models to achieve more accurate language conversion and obtain more precise intermediate results. Experimental results demonstrate that our method significantly improves performance on various multilingual tasks.

We devise an entropy-based method for critical layer selection and have preliminarily validated its effectiveness. However, variations among tokens within samples suggest that this guidance is still imprecise. We will further explore this issue in the future.

Acknowledgments

Bing Qin and Xiaocheng Feng are the co-corresponding authors of this work. We thank the anonymous reviewers for their insightful comments. This work was supported by the National Key R&D Program of China via grant No. 2021ZD0112905, National Natural Science Foundation of China (NSFC) via grant 62276078, the Key R&D Program of Heilongjiang via grant 2022ZX01A32, the Major Key Project of PCL via grant No. PCL2023A09, the International Cooperation Project of PCL, PCL2022D01 and the Fundamental Research Funds for the Central Universities (Grant No.HIT.OCEF.2023018).

References

- Artetxe, M.; Ruder, S.; and Yogatama, D. 2019. On the cross-lingual transferability of monolingual representations. *arXiv preprint arXiv:1910.11856*.
- Conneau, A.; Lample, G.; Rinott, R.; Williams, A.; Bowman, S. R.; Schwenk, H.; and Stoyanov, V. 2018. XNLI: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053*.
- Cui, Y.; Yang, Z.; and Yao, X. 2023. Efficient and effective text encoding for chinese llama and alpaca. *arXiv preprint arXiv:2304.08177*.
- Elbayad, M.; Gu, J.; Grave, E.; and Auli, M. 2019. Depth-adaptive transformer. *arXiv preprint arXiv:1910.10073*.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- HIT-SCIR. 2024. Chinese-Mixtral-8x7B: An Open-Source Mixture-of-Experts LLM. <https://github.com/HIT-SCIR/Chinese-Mixtral-8x7B>.
- Huang, C.; Zhou, H.; Zaïane, O. R.; Mou, L.; and Li, L. 2022. Non-autoregressive translation with layer-wise prediction and deep supervision. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 10776–10784.
- Huang, H.; Tang, T.; Zhang, D.; Zhao, W. X.; Song, T.; Xia, Y.; and Wei, F. 2023. Not all languages are created equal in llms: Improving multilingual capability by cross-lingual-thought prompting. *arXiv preprint arXiv:2305.07004*.
- Le Scao, T.; Fan, A.; Akiki, C.; Pavlick, E.; Ilić, S.; Hesslow, D.; Castagné, R.; Luccioni, A. S.; Yvon, F.; Gallé, M.; et al. 2023. Bloom: A 176b-parameter open-access multilingual language model.
- Lee, C.-Y.; Xie, S.; Gallagher, P.; Zhang, Z.; and Tu, Z. 2015. Deeply-supervised nets. In *Artificial intelligence and statistics*, 562–570. Pmlr.
- Lewis, P.; Oğuz, B.; Rinott, R.; Riedel, S.; and Schwenk, H. 2019. MLQA: Evaluating cross-lingual extractive question answering. *arXiv preprint arXiv:1910.07475*.
- Li, C.; Wang, S.; Zhang, J.; and Zong, C. 2024a. Improving In-context Learning of Multilingual Generative Language Models with Cross-lingual Alignment. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 8051–8069.
- Li, C.; Yang, W.; Zhang, J.; Lu, J.; Wang, S.; and Zong, C. 2024b. X-Instruction: Aligning Language Model in Low-resource Languages with Self-curated Cross-lingual Instructions. *arXiv preprint arXiv:2405.19744*.
- Li, H.; Koto, F.; Wu, M.; Aji, A. F.; and Baldwin, T. 2023. Bactrian-x: Multilingual replicable instruction-following models with low-rank adaptation. *arXiv preprint arXiv:2305.15011*.
- Li, R.; Wang, X.; Huang, G.; Yang, W.; Zhang, K.; Gu, X.; Tran, S. N.; Garg, S.; Alty, J.; and Bai, Q. 2022. A comprehensive review on deep supervision: Theories and applications. *arXiv preprint arXiv:2207.02376*.
- Lin, S.; Hilton, J.; and Evans, O. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Lin, X. V.; Mihaylov, T.; Artetxe, M.; Wang, T.; Chen, S.; Simig, D.; Ott, M.; Goyal, N.; Bhosale, S.; Du, J.; et al. 2022. Few-shot learning with multilingual generative language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 9019–9052.
- Longpre, S.; Lu, Y.; and Daiber, J. 2021. MKQA: A linguistically diverse benchmark for multilingual open domain question answering. *Transactions of the Association for Computational Linguistics*, 9: 1389–1406.
- Nguyen, X.-P.; Aljunied, S. M.; Joty, S.; and Bing, L. 2023. Democratizing LLMs for low-resource languages by leveraging their english dominant abilities with linguistically-diverse prompts. *arXiv preprint arXiv:2306.11372*.
- OpenAI. 2023. ChatGPT: August 12 Version. <https://chat.openai.com/>. Large language model.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Ponti, E. M.; Glavaš, G.; Majewska, O.; Liu, Q.; Vulić, I.; and Korhonen, A. 2020. XCOPA: A multilingual dataset for causal commonsense reasoning. *arXiv preprint arXiv:2005.00333*.
- Qin, L.; Chen, Q.; Wei, F.; Huang, S.; and Che, W. 2023. Cross-lingual prompting: Improving zero-shot chain-of-thought reasoning across languages. *arXiv preprint arXiv:2310.14799*.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford alpaca: An instruction-following llama model.
- Team, G. 2024. Gemma.
- Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).

Wendler, C.; Veselovsky, V.; Monea, G.; and West, R. 2024. Do llamas work in english? on the latent language of multilingual transformers. *arXiv preprint arXiv:2402.10588*.

Xu, S.; Li, J.; and Xiong, D. 2023. Language Representation Projection: Can We Transfer Factual Knowledge across Languages in Multilingual Language Models? *arXiv preprint arXiv:2311.03788*.

Zhang, Y.; Wang, Y.; Liu, Z.; Wang, S.; Wang, X.; Li, P.; Sun, M.; and Liu, Y. 2024. Enhancing Multilingual Capabilities of Large Language Models through Self-Distillation from Resource-Rich Languages. *arXiv preprint arXiv:2402.12204*.

Zhao, J.; Zhang, Z.; Zhang, Q.; Gui, T.; and Huang, X. 2024a. Llama beyond english: An empirical study on language capability transfer. *arXiv preprint arXiv:2401.01055*.

Zhao, Y.; Gu, A.; Varma, R.; Luo, L.; Huang, C.-C.; Xu, M.; Wright, L.; Shojanazeri, H.; Ott, M.; Shleifer, S.; et al. 2023. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*.

Zhao, Y.; Zhang, W.; Chen, G.; Kawaguchi, K.; and Bing, L. 2024b. How do Large Language Models Handle Multilingualism? *arXiv preprint arXiv:2402.18815*.

Zhu, W.; Liu, H.; Dong, Q.; Xu, J.; Huang, S.; Kong, L.; Chen, J.; and Li, L. 2023a. Multilingual machine translation with large language models: Empirical results and analysis. *arXiv preprint arXiv:2304.04675*.

Zhu, W.; Lv, Y.; Dong, Q.; Yuan, F.; Xu, J.; Huang, S.; Kong, L.; Chen, J.; and Li, L. 2023b. Extrapolating large language models to non-english by aligning languages. *arXiv preprint arXiv:2308.04948*.