

Focusing on Language: Revealing and Exploiting Language Attention Heads in Multilingual Large Language Models

Xin Liu^{1,2}, Qiyang Song^{1,2}, Qihang Zhou^{1,2}, Haichao Du^{1,2}, Shaowen Xu^{1,2},
Wenbo Jiang³, Weijuan Zhang^{1,2}, Xiaoqi Jia^{1,2*}

¹Institute of Information Engineering, Chinese Academy of Sciences

²School of Cyber Security, University of Chinese Academy of Sciences

³University of Electronic Science and Technology of China

liuxin235@mails.ucas.ac.cn, jiaxiaoqi@iie.ac.cn

Abstract

Large language models (LLMs) increasingly support multilingual understanding and generation. Meanwhile, efforts to interpret their internal mechanisms have emerged, offering insights to enhance multilingual performance. While multi-head self-attention (MHA) has proven critical in many areas, its role in multilingual capabilities remains underexplored. In this work, we study the contribution of MHA in supporting multilingual processing in LLMs. We propose Language Attention Head Importance Scores (LAHIS), an effective and efficient method that identifies attention head importance for multilingual capabilities via a single forward and backward pass through the LLM. Applying LAHIS to Aya-23-8B, Llama-3.2-3B, and Mistral-7B-v0.1, we reveal the existence of both language-specific and language-general heads. Language-specific heads enable cross-lingual attention transfer to guide the model toward target language contexts and mitigate off-target language generation issue, contributing to addressing challenges in multilingual LLMs. We also introduce a lightweight adaptation that learns a soft head mask to modulate attention outputs over language heads, requiring only 20 tunable parameters to improve XQuAD accuracy. Overall, our work enhances both the interpretability and multilingual capabilities of LLMs from the perspective of MHA.

Code — <https://github.com/Linuxin-xxx/LAHIS>

1 Introduction

Large language models (LLMs) have substantially advanced natural language processing (NLP). With their increasing pretraining on multilingual corpora, these models now demonstrate impressive capabilities in understanding and generating text across diverse languages (Huang et al. 2025; Qin et al. 2024). As a result, enhancing and analyzing their multilingual abilities has emerged as a central research objective (Muennighoff et al. 2023; Qi, Fernández, and Bisazza 2023; CONNEAU and Lample 2019).

Alongside performance improvements, there is a growing interest in uncovering the internal mechanisms that enable multilingual processing in LLMs. Gaining a deeper understanding of these mechanisms can facilitate the devel-

opment of more interpretable and controllable models. Recent studies have made progress in this area by identifying language-specific neurons or analyzing the function of each layer (Zhao et al. 2024; Tang et al. 2024; Zhang et al. 2024; Wang et al. 2025; Zhong et al. 2024; Wendler et al. 2024). However, one fundamental component of the Transformer architecture, multi-head self-attention (MHA), remains relatively underexplored in this context.

In other domains, studies have identified attention heads that perform specialized functions, such as induction heads (Olsson et al. 2022), retrieval heads (Wu et al. 2024), successor heads (Gould et al. 2023), and safety heads (Zhou et al. 2025). These roles indicate that attention heads can exhibit functional specialization. Motivated by this, we hypothesize that multilingual LLMs may also contain attention heads that are selectively specialized for processing particular languages.

In this paper, we investigate the contribution of individual attention heads within MHA to the multilingual capabilities of LLMs. To this end, we propose **Language Attention Head Importance Scores (LAHIS)**, an efficient method for attributing head importance with respect to multilingual performance. LAHIS utilizes a soft mask, implemented as a trainable matrix shaped to match the model’s attention structure, to estimate head importance by tracking its updates during training. This process requires only a single forward and backward pass over a target-language corpus.

By applying LAHIS to Aya-23-8B, Llama-3.2-3B, and Mistral-7B-v0.1, we obtain attention head importance matrices across different languages and identify **language heads** that contribute significantly to multilingual capabilities. We define those that are highly important for specific languages as **language-specific heads**, and empirically validate both their effectiveness and specialization. Furthermore, we identify a set of language-general heads that consistently receive high importance scores across all languages.

Language-specific heads can help address challenges faced by multilingual LLMs. In real-world applications such as dialogue systems and retrieval-augmented generation, multilingual input contexts are common. However, when conflicting facts are presented in different languages, the model’s information selection can become unpredictable and biased (Park and Lee 2025). Our experiments demonstrate that by enhancing or suppressing language-specific

*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

heads, we can steer the model’s attention and shift its preference toward the target-language context, thereby enabling more reliable and controllable cross-lingual reasoning.

We further demonstrate that modulating language-specific heads provides a practical mechanism to mitigate off-target language generation (Sennrich, Vamvas, and Mohammadshahi 2024), a common issue where multilingual models are likely to output high-resource languages like English, even when the input is in a different language. By suppressing English-heads during generation, we effectively realign the output language with the input. Our experiments on the XL-Sum dataset show that this targeted intervention restores language fidelity while maintaining high summarization quality, revealing a controllable axis for correcting multilingual generation asymmetries in LLMs.

Moreover, building on the discovery of language heads, we propose a lightweight adaptation method that enhances multilingual performance by training a small matrix aligned with the attention structure of the LLM, without modifying the LLM’s original weights. This trainable matrix updates only a few parameters associated with the identified heads and is inserted into the LLM as a gate to modulate the corresponding attention outputs during training or inference. Experiments on the XQuAD dataset across three models show that this adaptation, tuning only 14–20 parameters, yields an average accuracy improvement of 5 percentage points over the vanilla models, and outperforms the random-heads baseline by 4 points.

Our key contributions are as follows:

1. We propose LAHIS, a lightweight and efficient method for estimating attention head importance with respect to multilingual capabilities. Applied to three LLMs across diverse languages, LAHIS reveals the existence of language-specific heads critical to individual language performance, and language-general heads beneficial across languages.
2. We show that language-specific heads offer a mechanism to control model behavior in multilingual contexts. By enhancing or suppressing these heads, we steer the model’s attention toward the target language, enhancing cross-lingual reasoning and reducing bias from conflicting multilingual inputs.
3. We demonstrate that modulating language-specific heads effectively mitigates off-target language generation. Our experiments show that suppressing English-heads realigns the outputs with input languages while maintaining generation quality.
4. We propose a lightweight adaptation that enhances multilingual performance by inserting a simple, attention-aligned trainable matrix to modulate certain language heads’ outputs. Tuning only 14–20 parameters, it improves XQuAD accuracy by 5 points on average and surpasses the random-heads baseline by 4 points.

2 Method

2.1 Preliminary

In this section, we briefly introduce the Transformer architecture and the multi-head self-attention (MHA) mechanism

(Vaswani et al. 2017), as they form the foundation for the rest of the paper. We mainly focus on the decoder-only transformer architecture, which is widely adopted in modern LLMs. This architecture generally consists of multiple layers, each comprising a multi-head self-attention module and a feed-forward network.

The self-attention mechanism computes dependencies between tokens in an input sequence, assigning different weights to enable the LLMs to focus on the most relevant information. Multi-head self-attention further splits internal representations into multiple subspaces, allowing the models to capture diverse semantic features that a single head may struggle to represent effectively (Vaswani et al. 2017).

Each head in the MHA module is implemented as a Scaled Dot-Product Attention unit. The i -th head, $head_i$, is calculated as:

$$head_i = \text{Attention}(Q_i, K_i, V_i) = \text{softmax}\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right) V_i \quad (1)$$

The inputs consist of queries Q_i , keys K_i , and values V_i , each with a dimensionality of d_k . In self-attention, Q_i , K_i , and V_i are obtained from the same internal representation produced by the previous layer, through different linear projections. In this way, each head performs its computation independently. The outputs of all heads \mathcal{H} are then concatenated and projected by a parameter matrix W_O to form the final output of the MHA module:

$$\text{MultiHead} = \text{Concat}(head_1, \dots, head_n) W_O \quad (2)$$

Moreover, several models, such as Llama-3, adopt the Grouped-Query Attention (GQA) mechanism (Ainslie et al. 2023), in which multiple queries within a group share the same keys (K) and values (V). However, this shared use of K and V does not eliminate head-wise differentiation, as each head still employs its own distinct queries (Q).

Specifically, for layer l in an LLM, given the internal representation \mathcal{R}^{l-1} from the previous layer and the trainable parameters W_Q^l , W_K^l , and W_V^l of the current layer, attention head $head_i^l$ is computed as:

$$head_i^l = \text{Attention}(\mathcal{R}^{l-1} W_Q^l, \mathcal{R}^{l-1} W_K^l, \mathcal{R}^{l-1} W_V^l) \quad (3)$$

In this paper, a gate parameter g_i is used to control the magnitude of each attention head $head_i$:

$$h\tilde{e}ad_i = g_i \cdot head_i, \text{ where } g_i \in \mathbb{R}_{\geq 0} \text{ and } g_i \begin{cases} > 1 & (\text{enhance}) \\ \in [0, 1) & (\text{weaken}) \\ = 0 & (\text{disable}) \end{cases} \quad (4)$$

2.2 Language Attention Head Importance Scores

In this paper, we focus on the attention heads within multilingual LLMs. We propose a lightweight and efficient method, Language Attention Head Importance Scores (LAHIS), which estimates the importance of each attention head for multilingual capability. Specifically, it computes an importance matrix that reflects how crucial each attention

head is to a specific language, using only a single forward and backward pass of the LLM on a corpus in that language. This is achieved by a trainable matrix designed to efficiently approximate each head’s contribution.

Following prior works (Zhou et al. 2025; Michel, Levy, and Neubig 2019; Tang et al. 2024; Zhao et al. 2024), we estimate the importance of an attention head h by measuring the change in model loss $\Delta\mathcal{L}$ before and after disabling it via a masking mechanism. Given a language-specific corpus \mathcal{X}_c in language c , the model computes the change in loss $\mathcal{L}(x)$ when an attention head is deactivated. As shown in earlier studies (Voita et al. 2019; Michel, Levy, and Neubig 2019), many attention heads are redundant for specific tasks and can be pruned with minimal degradation in performance. Hence, if disabling a particular head leads to a significant increase in loss, it indicates that the head plays an important role in processing language c .

Formally, for an LLM with n_l layers and n_h heads per layer, we construct the attention head importance matrix $\text{ImpScore}_c \in \mathbb{R}^{n_l \times n_h}$ for language c , where each entry is defined based on the loss difference $\Delta\mathcal{L}$ when deactivating a specific attention head. Specifically, the importance score for attention head h_i is given by:

$$\text{ImpScore}_c(h_i) = \mathbb{E}_{x_c \in \mathcal{X}_c} [\mathcal{L}(h_i=0 | x_c) - \mathcal{L}(x_c)] \quad (5)$$

with $h_i \in \mathcal{H}$ indexing all attention heads.

However, explicitly disabling each attention head one by one is computationally expensive, especially for large-scale models such as Aya-23-8B, which contains 1024 attention heads. To address this, inspired by previous works (Michel, Levy, and Neubig 2019; Molchanov et al. 2017), we approximate the loss difference $\Delta\mathcal{L}$ using a *first-order Taylor expansion*. Specifically, we introduce a trainable matrix $\mathcal{M} \in \mathbb{R}^{n_l \times n_h}$, referred to as the soft attention head mask, to efficiently estimate the head-wise importance through a single forward and backward pass. The approximated loss difference $\Delta\tilde{\mathcal{L}}$ is given by:

$$\Delta\tilde{\mathcal{L}} = \mathbb{E}_{x_c \in \mathcal{X}_c} \left[\left| m_i \cdot \frac{\partial \mathcal{L}(x_c)}{\partial m_i} \right| \right] \quad (6)$$

where $m_i \in \mathcal{M}$ denotes the mask parameter corresponding to attention head h_i .

Furthermore, since we aim to identify heads whose removal leads to an increase in loss, we compute the proportion of negative gradients with respect to each mask parameter, denoted by W_{neg} :

$$W_{\text{neg}} = \mathbb{E}_{x_c \in \mathcal{X}_c} \left[\mathbb{I} \left(\frac{\partial \mathcal{L}(x_c)}{\partial m_i} < 0 \right) \right] \quad (7)$$

Finally, we define the refined attention head importance matrix $\text{LAHIS}_c \in \mathbb{R}^{n_l \times n_h}$ for language c by weighting the loss difference $\Delta\mathcal{L}$ with the factor W_{neg} that emphasizes heads with beneficial gradients. Formally, each importance score of LAHIS_c is computed as:

$$\text{LAHIS}_c(h_i) = \mathbb{E}_{x_c \in \mathcal{X}_c} \left[\left| m_i \cdot \frac{\partial \mathcal{L}(x_c)}{\partial m_i} \right| \cdot \mathbb{I} \left(\frac{\partial \mathcal{L}(x_c)}{\partial m_i} < 0 \right) \right] \quad (8)$$

In summary, each score in the attention head importance matrix reflects the degree to which a given head contributes to the model’s ability to handle a specific language. Attention heads with higher scores are deemed more critical and can be regarded as language-specific heads that play a specialized role in multilingual processing.

3 Experiments

3.1 Experiment Setting

Models We conduct experiments on three widely-used open-source multilingual models: Aya-23-8B (Aryabumi et al. 2024), Llama-3.2-3B (Grattafiori et al. 2024), and Mistral-7B-v0.1 (Jiang et al. 2023). Aya-23-8B is particularly optimized for multilinguality and supports 23 languages, including both high-resource and low-resource languages. Llama-3.2-3B officially supports eight languages, while it has been trained on a broader collection of languages beyond these supported ones. Mistral-7B-v0.1 is also pretrained on a mixture of multilingual data and exhibits multilingual capabilities.

Datasets For language modeling, we use perplexity (PPL) on the Wikipedia dataset, which contains cleaned articles across multiple languages, with one subset per language. For multilingual summarization, we use XL-Sum (Hasan et al. 2021), a large-scale dataset of article-summary pairs in 44 languages from BBC, covering both low and high-resource languages. Summary quality is measured using BERTScore F1 (Zhang* et al. 2020). For multilingual understanding, we adopt XQuAD (Artetxe, Ruder, and Yogatama 2020), a closed-domain QA benchmark in ten languages that evaluates cross-lingual transfer and comprehension.

Languages Based on the languages supported by the selected models and datasets, our experiments are conducted on the following languages: English (en), Spanish (es), Italian (it), German (de), Portuguese (pt), Greek (el), Chinese (zh), Japanese (ja), Korean (ko), Thai (th), Hindi (hi), Vietnamese (vi), and Indonesian (id). These languages cover both high-resource and low-resource settings to ensure diversity in our evaluation.

3.2 Obtaining Attention Head Importance Matrix

We apply LAHIS to obtain attention head importance matrices for multilingual capabilities in Aya-23-8B, Llama-3.2-3B, and Mistral-7B-v0.1, focusing on the following language sets for each model respectively: (en, zh, hi, vi, es, pt, id, ja, ko, el), (en, pt, it, de, th, id, vi, ja, ko, el) and (en, es, vi, hi, ja, th, el). As described in § 2.2, we perform forward and backward passes over each language’s Wikipedia corpus, which contains hundreds of thousands of tokens, to compute attention head importance scores using Equation 8.

Figure 1 presents the attention head importance matrices as heatmaps, where darker regions represent higher importance scores. Each heatmap is primarily composed of lighter areas with a few dark blocks, indicating that most attention heads contribute little to the model’s performance on the target language, while a small subset has a disproportionately

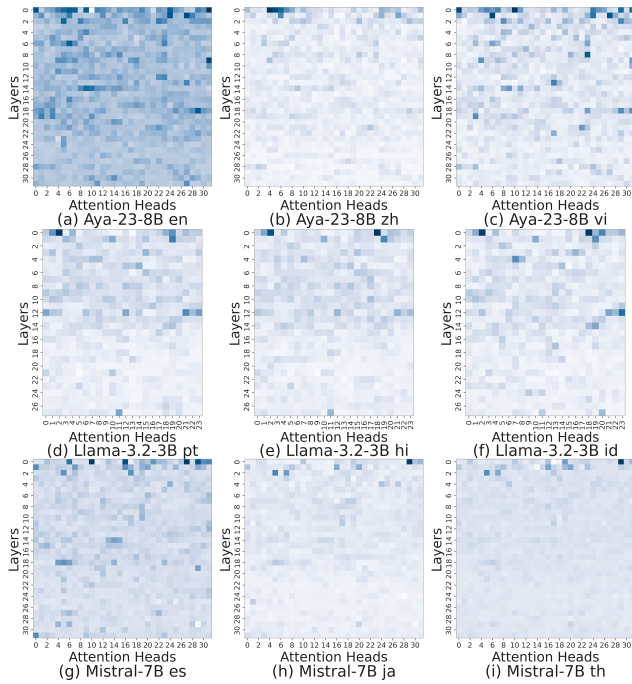


Figure 1: Illustration of attention head importance matrices obtained by LAHIS. Darker cells indicate higher importance scores. Most cells are light, showing that only a small subset of heads are highly important.

high impact. These high-importance heads are likely to capture language-specific patterns.

Furthermore, the identified high-importance heads are predominantly located in the lower layers of the LLMs, where attention heads exhibit greater feature differentiation (Voita et al. 2019).

3.3 Language-General Heads

Certain attention heads show consistently high importance across all languages, suggesting a key role in general language processing. To validate this, we disable these heads and evaluate LLMs on the XL-Sum dataset (500 samples per language) using BERTScore F1, which reflects both language understanding and generation abilities. Mistral-7B is excluded due to its tendency to generate English summaries for non-English inputs, resulting in artificially low evaluation scores. For Aya-23-8B and Llama-3.2-3B, we select the top 4% and 19% most important heads per language, respectively, and identify the ones shared across all languages, which account for approximately 1% and 5% of total heads.

As shown in Table 1, disabling these heads consistently degrades summarization quality across languages on two models (-12.3 and -19.2 F1 score on average), while randomly removing the same number of heads has minimal impact. In some cases, disabling these heads causes the models to produce repetitive or nonsensical outputs, indicating severe damage to their language capabilities. Therefore, we argue that these heads are universally important across languages and refer to them as language-general heads.

		Aya-23-8B							
		zh	hi	vi	es	pt	id	ko	Avg
VM		89.1	85.7	79.3	69.6	72.7	68.3	84.6	78.5
RH		88.5	86.5	77.8	66.8	72.9	67.0	84.8	77.7
GH		72.0	84.0	69.0	58.4	63.5	47.9	69.0	66.2
		Llama-3.2-3B							
		en	pt	th	id	vi	ja	ko	Avg
VM		56.2	59.1	56.3	66.5	71.4	72.2	82.1	66.3
RH		54.8	61.2	55.2	63.7	75.5	74.2	82.8	66.8
GH		48.0	53.9	39.0	46.9	38.7	61.4	59.0	49.6

Table 1: Impact of deactivating language-general heads on XL-Sum. Disabling these heads significantly reduces the average F1 scores (\uparrow), whereas randomly removing heads has negligible impact. *VM*: Vanilla Model; *RH*: Random Heads Deactivated; *GH*: Language-General Heads Deactivated.

3.4 Language-Specific Heads

§ 3.2 reveals that a small subset of attention heads plays a dominant role in supporting individual language capabilities in multilingual LLMs. We refer to these as language-specific heads. Specifically, based on the attention head importance matrices obtained by LAHIS, we select the top 2% highest-scoring heads per language, excluding language-general heads to ensure specificity. We then evaluate both the effectiveness and the language specificity of these heads.

To evaluate effectiveness, we deactivate the identified language-specific heads and measure the impact on perplexity (PPL) using the multilingual Wikipedia corpus with several hundred thousand tokens per language. An increase in PPL indicates performance degradation in language ability, reflecting the importance of the deactivated heads. As a baseline, we randomly deactivate the same number of heads. As shown in Table 2, disabling language-specific heads consistently leads to larger PPL increases than both the vanilla model and the random baseline, demonstrating the effectiveness of the language-specific heads identified by LAHIS.

To assess specificity, we evaluate whether each set of language-specific heads primarily affects its corresponding language by examining the cross-lingual impact of deactivating them. Figure 2 presents the resulting PPL increases, where each cell (i, j) denotes the PPL change on language j when the heads specific to language i are disabled. The consistently darker diagonal entries indicate that almost every language is affected the most by the deactivation of its own heads, confirming their specificity.

3.5 Transferring Cross-Lingual Attention

In interactive dialogue or retrieval-augmented generation (RAG) scenarios, input contexts are often multilingual (Qi, Fernández, and Bisazza 2025; Chirkova et al. 2024). However, when cross-lingual information conflicts, they exhibit a preference for the language of the context during information extraction, driven by prompt design or input-language factors (Park and Lee 2025; Chirkova et al. 2024; Sharma,

Aya-23-8B					
	en	zh	hi	vi	es
Vanilla Model	8.18	8.66	2.89	6.38	4.22
Random Heads	8.26	8.94	2.92	6.42	4.36
Language Heads	9.57	14.1	4.08	8.90	6.60
Llama-3.2-3B					
	en	pt	it	de	th
Vanilla Model	8.36	9.09	8.76	6.26	6.35
Random Heads	8.65	9.28	8.84	6.47	6.63
Language Heads	10.01	12.35	12.16	9.04	8.71
	id	vi	ja	ko	el
Vanilla Model	2.98	8.93	9.04	10.92	4.78
Random Heads	3.00	9.13	9.41	11.15	4.89
Language Heads	6.47	11.75	12.32	15.15	6.81
Mistral-7B					
	en	es	vi	hi	ja
Vanilla Model	5.08	3.26	5.14	3.43	6.96
Random Heads	5.17	3.32	5.20	3.46	6.99
Language Heads	5.73	5.16	8.91	4.80	9.81
	th	el	-	-	-
Vanilla Model	4.35	2.73	-	-	-
Random Heads	4.36	2.76	-	-	-
Language Heads	7.70	3.41	-	-	-

Table 2: Impact of deactivating language-specific heads on PPL (\downarrow). Disabling language-specific heads significantly increases PPL compared to both the vanilla models and random heads removal, highlighting the effectiveness of the LAHIS method in identifying language-specific heads.

Murray, and Xiao 2025), potentially misaligning with users’ task-specific language requirements.

In this section, we investigate whether language-specific heads can steer the model’s attention across languages. Specifically, we test whether enhancing or weakening them can guide the model’s attention toward a target language.

As illustrated in Figure 3, the input provides conflicting facts about Alice’s occupation: “scientist” in Chinese and “teacher” in Hindi. When queried in English, the model initially answers “teacher”, reflecting the Hindi input. Enhancing Chinese-heads or suppressing Hindi-heads shifts the answer to “scientist”, aligning with the Chinese context. Figure 4 further visualizes vocabulary-projected activations across Aya-23-8B’s layers using LogitLens (Wang 2025), showing that manipulating language-specific heads effectively redirects attention to the intended language, altering the final output.

To verify cross-lingual attention transfer via language-specific heads, we design prompts with conflicting descriptions of Alice’s profession in two languages (Context1

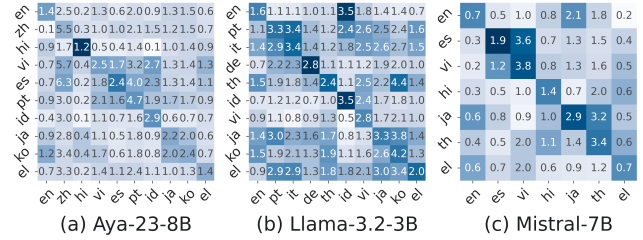


Figure 2: PPL impact from deactivating language-specific heads. Cell (i, j) shows the PPL increase on language j when disabling language i ’s heads. The dark diagonals indicate the specificity of heads identified by LAHIS.

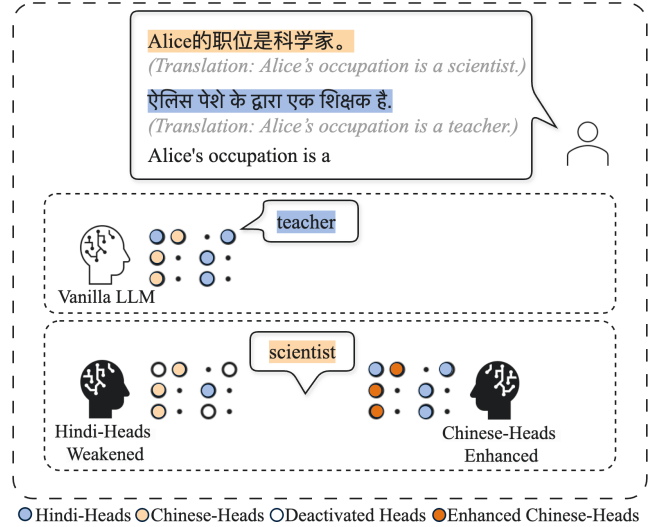


Figure 3: An illustration of shifting LLM’s multilingual attention via head intervention. The vanilla model chooses “teacher” from the Hindi context, but switches to “scientist” from the Chinese context after weakening Hindi-heads or enhancing Chinese-heads.

in language A and Context2 in language B) and an English question. The model’s predicted profession indicates its preference for one language’s context. We then enhance language A’s heads or deactivate language B’s heads to examine whether the output shifts toward Context1.

Figure 5 shows experimental results on Aya-23-8B, Llama-3.2-3B, and Mistral-7B with 450–720 samples per language. Models initially prefer Context2, likely due to its proximity to the question. After intervention, preference for Context1 rises by about 10 percentage points, while reliance on Context2 drops by 12 points. These results confirm that language-specific heads enable effective cross-lingual attention transfer.

3.6 Mitigating Off-Target Language Generation

Multilingual LLMs exhibit *off-target language generation* (Tang et al. 2024; Sennrich, Vamvas, and Mohammadshahi 2024; Chen et al. 2023; Gu et al. 2019), where outputs do not match the input language. This misalignment ad-

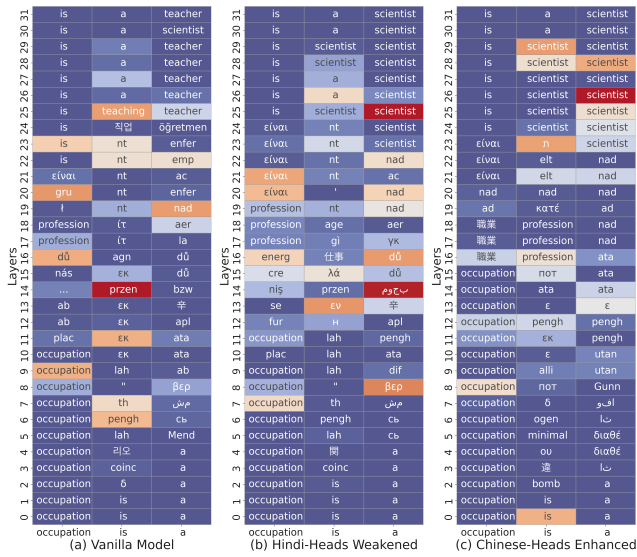


Figure 4: Layer-wise vocabulary projections of Aya-23-8B illustrating the multilingual attention shift.

versely affects tasks requiring language consistency. This phenomenon is also observed in our experiments on XL-Sum dataset with Mistral-7B, where the model generates English summaries despite non-English inputs and prompts, resulting in reduced evaluation scores. The bias toward English is primarily attributed to its dominance in the pre-training corpus, leading the model to develop stronger proficiency and preference for English (Hasan et al. 2024).

To address this issue, we reduce the model’s attention to English to encourage generation in the target non-English language. Experimenting on XL-Sum with Mistral-7B (100 samples per language), we deactivate 5% of English-heads for *es* and *vi*, and 1% for *hi*, *ja*, and *th*. The language of the generated summaries is classified using FastText (Joulin et al. 2016). Figure 6 illustrates a Thai summarization example where weakening English-heads shifts the output from English to Thai. Table 3 shows that deactivating a small subset of English-heads increases correct-language outputs to 100%. Additionally, improved BERTScore F1 confirms that summary quality remains high after this language shift.

These results suggest that language-specific heads can be leveraged to modulate the model’s language bias, providing a simple yet effective solution to the off-target language generation problem. Additionally, these findings indicate an asymmetry in multilingual LLMs, showing that while their understanding capability spans many languages, their generation ability remains disproportionately influenced by high-resource languages.

3.7 Enhancing Multilingual Performance via Language-Specific Head Mask

Language heads can be leveraged to enhance the multilingual capabilities of LLMs. For a given LLM with n_l layers and n_h attention heads per layer, we construct a trainable matrix of shape (n_l, n_h) . Only the parameters corresponding

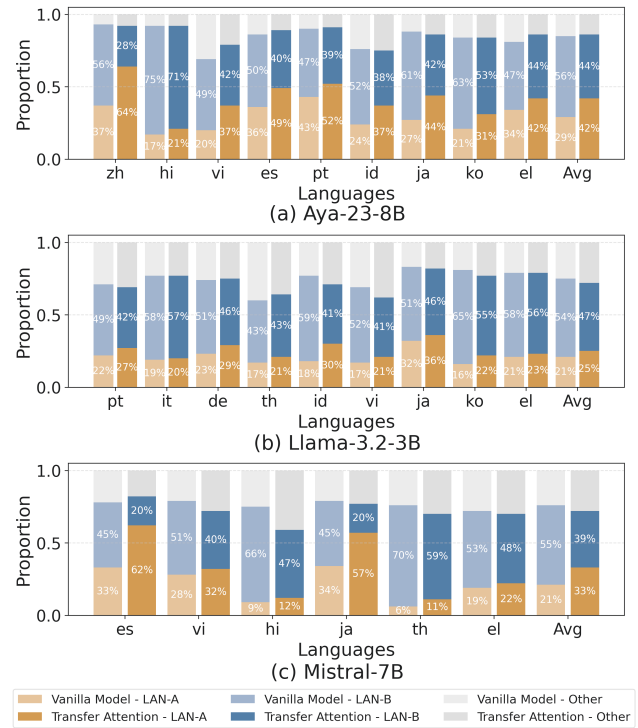


Figure 5: Influence of language-specific heads on multilingual prediction shifting. Given prompts with conflicting facts in two languages, enhancing heads for language A or disabling heads for B increases the model’s reliance on language A’s context, indicating that language-specific heads can steer cross-lingual attention and influence predictions.

to the identified language heads are set as trainable, while the rest remain frozen, which we refer to as the **language-specific head mask**. During training or inference, the mask parameters are multiplied with the attention outputs (prior to projection through W_O), thereby influencing the model’s final predictions.

Based on each language’s attention head importance matrix, we select the top 2% scoring heads and set the corresponding positions in the language-specific head mask as trainable parameters. To support language task performance, language-general heads are not excluded from this set. To assess its multilingual effectiveness, we evaluate the language-specific head mask on three models using the XQuAD across languages. We train the language-specific head masks for two epochs on 200 training samples and evaluate on 800 test samples. As a baseline, we train head masks by randomly selecting an equal number of parameters. As shown in Table 4, the language-specific head masks yield an average accuracy improvement of 5 percentage points over the vanilla models, and outperform the random baseline by 4 points. These results demonstrate the effectiveness of the language heads identified using LAHIS, and show that updating only a small number of parameters within the language-specific head mask (e.g., 14–20 for these three LLMs), requiring only 30 seconds, can improve multilingual task performance.

Prompt in Thai: สรุปข้อความที่กำหนดโดยย่อ [ข้อความ] Text in Thai [สรุป] (Translation: Summarize the given text briefly. [Text]Text in Thai [Summary])
Output from Vanilla Model: African twin girls, joined at the head, separated after 18-hour surgery
Output after English-heads Weakened: เพิ่งถูกแยกร่างออกจากกัน หลังเข้ารับการผ่าตัดใหญ่ นาน 18 ชั่วโมง เออร์วินา และเพ (Translation: They had recently been separated after having had a major surgery of 18 hours)

Figure 6: An example of off-target language generation and mitigation. The LLM summarizes Thai input in English but outputs correctly after weakening English-heads.

		es	vi	hi	ja	th
Lan-Acc	VM	0.67	0.35	0.74	0.99	0.78
	EH	1.00	1.00	1.00	1.00	1.00
Sum-Qua	VM	57.41	50.21	70.19	81.48	58.99
	EH	71.70	80.27	85.59	81.54	69.07

Table 3: Language-accuracy (\uparrow) and summarization-quality (\uparrow) of Mistral-7B on XL-Sum. Deactivating English-heads yields 100% target language outputs and significantly improves F1 scores versus the vanilla model. *Lan-Acc*: Language Accuracy; *Sum-Qua*: Summarization Quality; *VM*: Vanilla Model; *EH*: English-Heads Deactivated.

4 Related Work

Multilingual Mechanism Understanding how LLMs internally process different languages remains an open challenge. Wendler et al. (2024) applies the LogitLens method (Wang 2025) and finds that in English-centric models like Llama, token representations gradually shift from input space to an English-biased concept space before transitioning to the target language’s output space. Tang et al. (2024) proposes a method to identify language-specific neurons, suggesting that a small subset of neurons governs language capability. Similarly, Zhao et al. (2024) introduces the PLND method to detect such neurons and verifies a multilingual workflow in LLMs: understanding in multiple languages, reasoning in English, and generating output in the target language. Wang et al. (2025) uses mechanistic interpretability to explore cross-lingual inconsistencies, revealing that most layers encode language-agnostic knowledge, with language-specific transformations occurring mainly in the final layers, where failures often cause output errors. While these works offer valuable insights, they largely focus on entire layers or neurons. In contrast, our work investigates the multilingual roles of attention heads, offering a fresh perspective on cross-lingual behavior in LLMs.

Multi-Head Attention Multi-head attention in LLMs has been shown to be sparse, with only a small subset of heads capturing task-relevant features while others can be pruned without performance degradation (Michel, Levy, and Neu-

Aya-23-8B						
	en	zh	vi	es	el	Avg
VM	76.00	57.63	46.75	52.13	43.88	55.28
RH	75.25	60.75	48.00	52.88	43.88	56.15
LH	77.38	63.63	53.88	60.38	50.25	61.10
Llama-3.2-3B						
	en	de	th	vi	el	Avg
VM	56.13	29.13	27.88	37.38	14.38	32.98
RH	56.50	32.00	27.00	38.25	14.63	33.68
LH	59.25	35.88	31.25	39.00	18.50	36.78
Mistral-7B						
	en	es	vi	hi	el	Avg
VM	44.88	30.25	21.38	9.63	6.50	22.53
RH	52.75	32.13	22.25	11.38	7.25	25.15
LH	60.13	37.13	24.13	14.88	8.88	29.03

Table 4: Accuracy (\uparrow) (%) on XQuAD. Training language-specific head masks (14–20 parameters) improves accuracy by 5 points over vanilla models and outperforms random masks by 4. *VM*: Vanilla Model; *RH*: Random Head Mask; *LH*: Language-Specific Head Mask.

big 2019; Voita et al. 2019). However, certain heads have also been demonstrated to play critical roles in many core capabilities of LLMs (Su, Kempe, and Ullrich 2025; Vig 2019). Gould et al. (2023) identifies successor heads that increment tokens in natural order, capturing abstract representations shared across architectures. Wu et al. (2024) discovers retrieval heads responsible for extracting relevant information from long contexts, which are essential for chain-of-thought (CoT) reasoning. Zhou et al. (2025) examines the role of attention heads in safety-related behaviors, introducing attribution methods to highlight their influence. Motivated by these findings, our work investigates how certain attention heads contribute to multilingual processing, highlighting a new facet of their functionality.

5 Conclusion

This paper provides new insights into the role of MHA in enabling multilingual capabilities in LLMs. We introduce LAHIS, an effective and efficient method for computing attention head importance with respect to multilingual processing. Through comprehensive experiments on three LLMs, we confirm the existence of both language-specific and language-general heads. Our analysis demonstrates that language-specific heads are instrumental in facilitating cross-lingual attention transfer and mitigating off-target language generation, thereby addressing the challenges in multilingual LLMs. Leveraging language heads, we further propose a lightweight adaptation strategy that improves multilingual performance by updating only a small subset of parameters. Collectively, our work not only enhances the interpretability of multilingual behaviors in LLMs from the perspective of MHA, but also presents practical approaches to strengthen the multilingual capabilities of LLMs.

Acknowledgments

We are grateful to all the anonymous reviewers for their insightful comments, which greatly improved our paper. The research is supported in part by the National Natural Science Foundation of China (No. 62202465) and the National Key Research and Development Program of China (No.2021YFB2910109), and the Outstanding Talent Scheme (Category B) - Qihang Zhou (E3YY141116).

References

- Ainslie, J.; Lee-Thorp, J.; de Jong, M.; Zemlyanskiy, Y.; Lebron, F.; and Sanghvi, S. 2023. GQA: Training Generalized Multi-Query Transformer Models from Multi-Head Checkpoints. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 4895–4901. Singapore: Association for Computational Linguistics.
- Artetxe, M.; Ruder, S.; and Yogatama, D. 2020. On the Cross-lingual Transferability of Monolingual Representations. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4623–4637. Online: Association for Computational Linguistics.
- Aryabumi, V.; Dang, J.; Talupuru, D.; and et al. 2024. Aya 23: Open Weight Releases to Further Multilingual Progress. arXiv:2405.15032.
- Chen, L.; Ma, S.; Zhang, D.; Wei, F.; and Chang, B. 2023. On the Off-Target Problem of Zero-Shot Multilingual Neural Machine Translation. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023*, 9542–9558. Toronto, Canada: Association for Computational Linguistics.
- Chirkova, N.; Rau, D.; Déjean, H.; Formal, T.; Clinchant, S.; and Nikoulina, V. 2024. Retrieval-augmented generation in multilingual settings. In Li, S.; Li, M.; Zhang, M. J.; Choi, E.; Geva, M.; Hase, P.; and Ji, H., eds., *Proceedings of the 1st Workshop on Towards Knowledgeable Language Models (KnowLLM 2024)*, 177–188. Bangkok, Thailand: Association for Computational Linguistics.
- CONNEAU, A.; and Lample, G. 2019. Cross-lingual Language Model Pretraining. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Gould, R.; Ong, E.; Ogden, G.; and Conmy, A. 2023. Successor Heads: Recurring, Interpretable Attention Heads In The Wild. arXiv:2312.09230.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; and et al. 2024. The Llama 3 Herd of Models. arXiv:2407.21783.
- Gu, J.; Wang, Y.; Cho, K.; and Li, V. O. 2019. Improved Zero-shot Neural Machine Translation via Ignoring Spurious Correlations. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1258–1268. Florence, Italy: Association for Computational Linguistics.
- Hasan, M. A.; Tarannum, P.; Dey, K.; Razzak, I.; and Naseem, U. 2024. Do Large Language Models Speak All Languages Equally? A Comparative Study in Low-Resource Settings. arXiv:2408.02237.
- Hasan, T.; Bhattacharjee, A.; Islam, M. S.; Mubasshir, K.; Li, Y.-F.; Kang, Y.-B.; Rahman, M. S.; and Shahriyar, R. 2021. XL-Sum: Large-Scale Multilingual Abstractive Summarization for 44 Languages. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 4693–4703. Online: Association for Computational Linguistics.
- Huang, K.; Mo, F.; Zhang, X.; Li, H.; Li, Y.; Zhang, Y.; Yi, W.; Mao, Y.; Liu, J.; Xu, Y.; Xu, J.; Nie, J.-Y.; and Liu, Y. 2025. A Survey on Large Language Models with Multilingualism: Recent Advances and New Frontiers. arXiv:2405.10936.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; de las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; Lavaud, L. R.; Lachaux, M.-A.; Stock, P.; Scao, T. L.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2023. Mistral 7B. arXiv:2310.06825.
- Joulin, A.; Grave, E.; Bojanowski, P.; Douze, M.; Jégou, H.; and Mikolov, T. 2016. FastText.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Michel, P.; Levy, O.; and Neubig, G. 2019. Are Sixteen Heads Really Better than One? In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Molchanov, P.; Tyree, S.; Karras, T.; Aila, T.; and Kautz, J. 2017. Pruning Convolutional Neural Networks for Resource Efficient Inference. In *International Conference on Learning Representations*.
- Muennighoff, N.; Wang, T.; Sutawika, L.; Roberts, A.; Biderman, S.; Le Scao, T.; Bari, M. S.; Shen, S.; Yong, Z. X.; Schoelkopf, H.; Tang, X.; Radev, D.; Aji, A. F.; Almubarak, K.; Albanie, S.; Alyafeai, Z.; Webson, A.; Raff, E.; and Raffel, C. 2023. Crosslingual Generalization through Multitask Finetuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15991–16111. Toronto, Canada: Association for Computational Linguistics.
- Olsson, C.; Elhage, N.; Nanda, N.; Joseph, N.; DasSarma, N.; Henighan, T.; Mann, B.; Askell, A.; Bai, Y.; Chen, A.; Conerly, T.; Drain, D.; Ganguli, D.; Hatfield-Dodds, Z.; Hernandez, D.; Johnston, S.; Jones, A.; Kernion, J.; Lovitt, L.; Ndousse, K.; Amodei, D.; Brown, T.; Clark, J.; Kaplan, J.; McCandlish, S.; and Olah, C. 2022. In-context Learning and Induction Heads. arXiv:2209.11895.
- Park, J.; and Lee, H. 2025. Investigating Language Preference of Multilingual RAG Systems. arXiv:2502.11175.
- Qi, J.; Fernández, R.; and Bisazza, A. 2023. Cross-Lingual Consistency of Factual Knowledge in Multilingual Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 10650–10666. Singapore: Association for Computational Linguistics.

- Qi, J.; Fernández, R.; and Bisazza, A. 2025. On the Consistency of Multilingual Context Utilization in Retrieval-Augmented Generation. arXiv:2504.00597.
- Qin, L.; Chen, Q.; Zhou, Y.; Chen, Z.; Li, Y.; Liao, L.; Li, M.; Che, W.; and Yu, P. S. 2024. Multilingual Large Language Model: A Survey of Resources, Taxonomy and Frontiers. arXiv:2404.04925.
- Sennrich, R.; Vamvas, J.; and Mohammadshahi, A. 2024. Mitigating Hallucinations and Off-target Machine Translation with Source-Contrastive and Language-Contrastive Decoding. In Graham, Y.; and Purver, M., eds., *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, 21–33. St. Julian’s, Malta: Association for Computational Linguistics.
- Sharma, N.; Murray, K.; and Xiao, Z. 2025. Faux Polyglot: A Study on Information Disparity in Multilingual Large Language Models. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 8090–8107. Albuquerque, New Mexico: Association for Computational Linguistics. ISBN 979-8-89176-189-6.
- Su, J.; Kempe, J.; and Ullrich, K. 2025. From Concepts to Components: Concept-Agnostic Attention Module Discovery in Transformers. arXiv:2506.17052.
- Tang, T.; Luo, W.; Huang, H.; Zhang, D.; Wang, X.; Zhao, X.; Wei, F.; and Wen, J.-R. 2024. Language-Specific Neurons: The Key to Multilingual Capabilities in Large Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5701–5715. Bangkok, Thailand: Association for Computational Linguistics.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Vig, J. 2019. A Multiscale Visualization of Attention in the Transformer Model. In Costa-jussà, M. R.; and Alfonseca, E., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 37–42. Florence, Italy: Association for Computational Linguistics.
- Voita, E.; Talbot, D.; Moiseev, F.; Sennrich, R.; and Titov, I. 2019. Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5797–5808. Florence, Italy: Association for Computational Linguistics.
- Wang, M.; Adel, H.; Lange, L.; Liu, Y.; Nie, E.; Strötgen, J.; and Schütze, H. 2025. Lost in Multilinguality: Dissecting Cross-lingual Factual Inconsistency in Transformer Language Models. arXiv:2504.04264.
- Wang, Z. 2025. LogitLens4LLMs: Extending Logit Lens Analysis to Modern Large Language Models. arXiv:2503.11667.
- Wendler, C.; Veselovsky, V.; Monea, G.; and West, R. 2024. Do Llamas Work in English? On the Latent Language of Multilingual Transformers. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15366–15394. Bangkok, Thailand: Association for Computational Linguistics.
- Wu, W.; Wang, Y.; Xiao, G.; Peng, H.; and Fu, Y. 2024. Retrieval Head Mechanistically Explains Long-Context Factuality. arXiv:2404.15574.
- Zhang*, T.; Kishore*, V.; Wu*, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*.
- Zhang, Z.; Zhao, J.; Zhang, Q.; Gui, T.; and Huang, X. 2024. Unveiling Linguistic Regions in Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6228–6247. Bangkok, Thailand: Association for Computational Linguistics.
- Zhao, Y.; Zhang, W.; Chen, G.; Kawaguchi, K.; and Bing, L. 2024. How do Large Language Models Handle Multilingualism? In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Zhong, C.; Cheng, F.; Liu, Q.; Jiang, J.; Wan, Z.; Chu, C.; Murawaki, Y.; and Kurohashi, S. 2024. Beyond English-Centric LLMs: What Language Do Multilingual Language Models Think in? arXiv:2408.10811.
- Zhou, Z.; Yu, H.; Zhang, X.; Xu, R.; Huang, F.; Wang, K.; Liu, Y.; Fang, J.; and Li, Y. 2025. On the Role of Attention Heads in Large Language Model Safety. In *The Thirteenth International Conference on Learning Representations*.