

000 001 TOKEN ALIGNMENT HEADS: UNVEILING ATTEN- 002 TION'S ROLE IN LLM MULTILINGUAL TRANSLATION 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Recently, large language models (LLMs) have made remarkable progress, with
012 multilingual capability emerging as a core foundational strengths. However, the
013 internal mechanisms by which these models perform translation remain incom-
014 pletely understood. In this paper, we elucidate the relationship between the at-
015 tention mechanism in LLMs and their translation abilities. We find that certain
016 attention heads, which we term token alignment heads, are specifically responsi-
017 ble for mapping tokens from the source language to the target language during
018 inference. Through a systematic investigation across various models, we confirm
019 that these token alignment heads exhibit several key characteristics: (1) Univer-
020 sality: They are present in all LLMs we studied. (2) Sparsity: They constitute
021 only a small fraction of all attention heads. (3) Consistency: The set of token
022 alignment heads activated by the model shows strong consistency across different
023 language pairs. (4) Causality: Interventionally removing these heads leads to a
024 sharp decline in the model's translation performance, while randomly removing
025 non-token alignment heads has little impact on translation ability. (5) Functional
026 Specificity: Ablating token alignment heads disproportionately harms translation
027 but has a varied impact on other multilingual tasks. We also traced the formation
028 of token alignment heads during pre-training, revealing an evolutionary path of
029 rapid proliferation, stabilization, and eventual pruning. Furthermore we leverage
030 these token alignment heads to filter multilingual training data, and our experi-
031 ments show that these data could enhance translation capabilities of the models.

032 1 INTRODUCTION 033

034 Recently released Large Language Models (LLMs) (Comanici et al., 2025; OpenAI, 2025; An-
035 thropic, 2025; Liu et al., 2024; Yang et al., 2025) have demonstrated remarkable multilingual ca-
036 pabilities, showing significant improvements in both the complexity of multilingual tasks they can
037 handle and the range of languages they support. Multilingual proficiency has now become an essen-
038 tial foundational ability for state-of-the-art LLMs. Among these capabilities, translation is particu-
039 larly crucial, as it not only represents a key application but also underpins the overall multilingual
040 performance of these models. A deeper understanding of the underlying mechanisms of translation
041 in LLMs is therefore vital, not only for a comprehensive view of their inner workings but also for
042 providing valuable insights to guide the development of superior multilingual training strategies, in
043 terms of both model architecture and data selection.

044 A growing body of research has begun to demystify how LLMs process multilingual information
045 (Artetxe et al., 2020; Lindsey et al., 2025; Datta et al., 2020; Chang et al., 2022). Several recent
046 works have explored the internal mechanics of multilingualism in LLMs. For instance, Zhao et al.
047 (2024) found that LLMs often initially process queries by converting multilingual inputs into an
048 English-centric representation before solving tasks. Similarly, Schut et al. (2025) revealed that these
049 models tend to make decisions and reason within an English-dominated semantic space. These
050 findings underscore the hypothesis that an internal translation process is a core component of LLM
051 multilingualism, highlighting the importance of understanding these translation mechanisms to fur-
052 ther elucidate their broader multilingual capabilities.

053 Prior research (Michel et al., 2019; Vig & Belinkov, 2019; Finlayson et al., 2021; Elhage et al.,
2021a) into the attention mechanism has established the functional specialization of individual

054
055
056
057
058
059
060
061
062
063
064
065
066

Response By Llama-3.1-8B:
 Translate this sentence from English to Simplified Chinese
 Sentence: "We now have 4-month-old mice that are non-diabetic that used to be diabetic," he added.
 Answer: "我们现在有4个月大的老鼠, 它们以前是糖尿病的, 现在不再是糖尿病了。" 他补充道

Response By top 30 token alignment heads pruned Llama-3.1-8B:
 Translate this sentence from English to Simplified Chinese
 Sentence: "We now have 4-month-old mice that are non-diabetic that used to be diabetic," he added.
 Answer: "We now have 4-month-old mice that are non-diabetic that used to be diabetic," he added.

Response By 30 random heads pruned Llama-3.1-8B:
 Translate this sentence from English to Simplified Chinese
 Sentence: "We now have 4-month-old mice that are non-diabetic that used to be diabetic," he added.
 Answer: "我们现在有4个月大的老鼠, 它们以前是糖尿病的, 现在不再是糖尿病了。" 他补充道

Figure 1: An example of token alignment head pruning in Llama-3.1-8B. The top panel shows the correct translation generated by the original model. The middle panel demonstrates that after the top 30 token alignment heads are pruned, the model loses its ability to translate and reverts to copying the English source text, while its basic copy-paste functionality is preserved. In contrast, the bottom panel shows that pruning 30 random non-token alignment heads (control) has no impact on the translation output.

heads. Early work on Transformer (Vaswani et al., 2023) models for machine translation showed that many attention heads were redundant and could be pruned with minimal impact on performance (Voita et al., 2019; Kovaleva et al., 2019). Subsequent studies, such as Kim et al. (2021); Ma et al. (2021); Zhang et al. (2025), have investigated the role of attention heads in translation by measuring their impact on translation metrics. These studies identified certain attention heads as critical for translation and noted that the sets of important heads are highly similar across different language pairs. However, these approaches often have limitations. They frequently rely on task-specific evaluation metrics, are typically conducted on smaller or single models, and their methods for identifying important heads can be opaque. Crucially, they often stop at identifying which heads are important, without fully explaining how these heads mechanistically contribute to the translation process.

Unlike previous work that measures head importance based on downstream benchmark performance, our approach shifts the focus to identifying the underlying mechanism. Inspired by the discovery of other functionally specialized circuits in LLMs, such as “induction heads” that implement in-context learning (Olsson et al., 2022) and “retrieval heads” designed for knowledge retrieval (Wu et al., 2025a), we hypothesized that a similar specialization must exist for translation. We posited that beyond attention heads performing generic copy-paste behaviors, there must be a set of heads specifically responsible for the core translation task: mapping tokens from a source language to their corresponding tokens in a target language. We term these attention heads, characterized by their direct cross-lingual token alignment behavior which can be regarded as a form of word alignment(Brown et al., 1993; Och & Ney, 2003), “token alignment heads”. Figure 1 provides an illustration of this functional specialization: when we ablate the top 30 token alignment heads from a Llama-3.1-8B model (Grattafiori et al., 2024), its translation capability collapses. Crucially, the model does not simply fail; it reverts to a more basic copy-paste behavior, reproducing the English input verbatim. This demonstrates that the model’s general ability to copy tokens remains intact, and that the ablated heads perform a specific, non-copying function of cross-lingual mapping.

In this work, we conduct a systematic investigation across a series of LLMs to validate the existence and properties of token alignment heads. We reveal that token alignment heads play a pivotal role in the translation capabilities of LLMs, thereby uncovering the relationship between translation mechanisms and attention. We further identify and validate several key characteristics of these token alignment heads: (1) Universality: All large models we studied possess such token alignment heads; (2) Sparsity: Only a small subset of attention heads function as token alignment heads; (3) Consistency: The token alignment heads activated by the model when translating different language pairs exhibit strong consistency. (4) Causality: Ablating these token alignment heads through causal intervention leads to a significant drop in translation performance, whereas removing a random equivalent number of non-token alignment heads has little impact. (5) Functional Specificity: Ablating token alignment heads disproportionately harms translation but has a varied impact on other

108 multilingual tasks, suggesting the functional specificity of token alignment heads. Furthermore, we
 109 investigate the formation process of token alignment heads by analyzing the entire pre-training life-
 110 cycle of a model. Our analysis reveals a distinct developmental trajectory in three phases: an initial
 111 rapid proliferation of token alignment heads that coincides with the acquisition of translation ability,
 112 followed by a period where the core set of heads stabilizes, and finally a long phase of consolidation
 113 and pruning. This discovery provides insight into how specialized circuits emerge and are optimized
 114 during large-scale training.

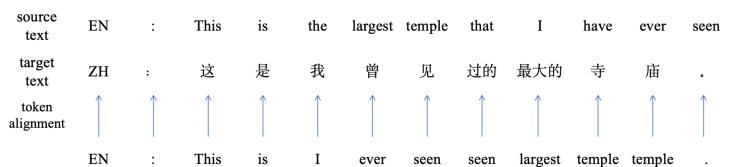
115 We further substantiate these findings through a practical application. We introduce TRater, an al-
 116 gorithm that leverages token alignment heads to score multilingual training data based on its impor-
 117 tance to the translation mechanism. Our experiments reveal that a small fraction of data identified by
 118 TRater is responsible for the model’s final translation proficiency. This result provides evidence for
 119 the causal role of token alignment heads. The discovery that a core capability like translation is gov-
 120 erned by a sparse and functionally specialized circuit provides a concrete target for future research.
 121 Ultimately, these insights pave the way for more efficient and robust multilingual systems, enabling
 122 targeted architectural innovations, data curation strategies guided by mechanistic understanding.

124 2 DETECTING TOKEN ALIGNMENT HEAD

126 In this section, we introduce the algorithm for detecting token alignment heads. Since the model
 127 involves cross-lingual token alignment during the translation process, we first need to identify the
 128 mapping relationship between tokens from the source language to the target language. Then, we
 129 define a metric called the translation score to recognize attention heads that implement the model’s
 130 translation mechanism. The translation score measures the frequency with which an attention head
 131 maps tokens from the source language to the target language. If an attention head exhibits a relatively
 132 high translation score, it indicates that this attention head frequently performs cross-lingual token
 133 alignment when processing different translation texts. Such attention heads are what we refer to as
 134 token alignment heads.

135 2.1 TOKEN ALIGNMENT ANNOTATION

137 Since existing token alignment tools do not cover all languages, we utilize OpenAI’s GPT-4.1 model
 138 to annotate token alignments in translation texts. Specifically, we require the large language model
 139 to identify the corresponding source language token for each target language token and provide a
 140 confidence score for each token alignment. To ensure the accuracy of the annotations, we only retain
 141 token alignments with a confidence score greater than 0.9. If no corresponding source token exists, it
 142 is marked as None. Figure 2 shows an example of token alignment results from English to Chinese.



150 Figure 2: A token alignment example from English text to Chinese text

153 2.2 TRANSLATION SCORE

155 In the decoding process of translation, we define the translation score as the frequency of valid token
 156 alignments by attention heads. Specifically, during the greedy decoding process, let the currently
 157 generated token be t , and let the attention score of the attention head be denoted as $\mathbf{w} \in \mathbb{R}^{|x|}$. If
 158 token t has a corresponding source token s with position idx as s_{idx} , and the attention head assigns
 159 the highest attention score probability to the source token s , then we consider that the attention head
 160 has successfully completed a language-pair token alignment. Formally, we have: $\mathbf{w}_{s_{\text{idx}}} = \max(\mathbf{w})$.

161 Let g_h denote the number of valid language-pair token alignments performed by attention head h ,
 and let m be the total number of target tokens that have a corresponding valid source token. Then,

162 the translation score of attention head h is defined as:
 163

$$164 \quad TS_h = \frac{g_h}{m} \quad (1)$$

165

166 2.3 TOKEN ALIGNMENT HEAD DETECTION

167

168 To empirically identify token alignment heads, we compute the Translation Score for every attention
 169 head in the model using the dev split of the FLORES-101 dataset (Goyal et al., 2021). For each of
 170 the approximately 900 source-target sentence pairs in a given language direction, we calculate the
 171 TS for all heads. The final score for each head is the average TS computed across all examples in that
 172 language pair. An attention head is then classified as a token alignment head if its final Translation
 173 Score exceeds a predefined threshold of 0.1.

174

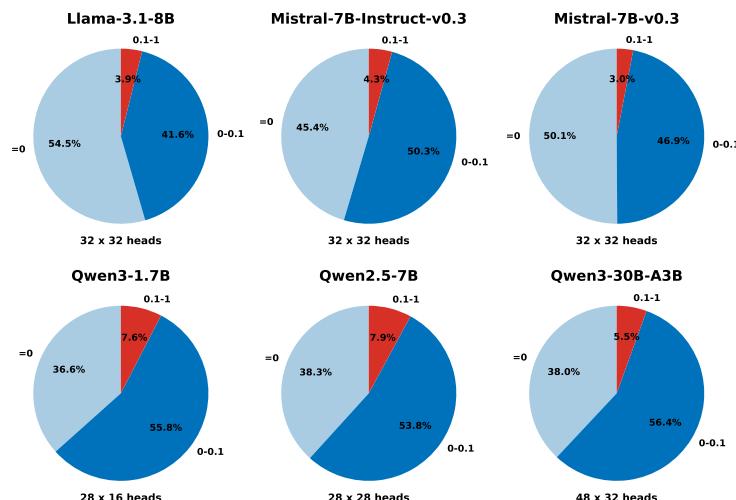
175 3 BASIC PROPERTIES OF TOKEN ALIGNMENT HEADS

176

177 In this section, we characterize the fundamental properties of the identified token alignment heads.
 178 To ensure the robustness and generalizability of our findings, our analysis spans a diverse set of
 179 open-source models. This selection covers a range of parameter sizes (1.7B to 30B), architec-
 180 tures (dense and Mixture-of-Experts (Jacobs et al., 1991)), and training stages (pre-trained and
 181 instruction-tuned). The models include Llama-3.1-8B, Mistral-7B-Instruct-v0.3¹, Mistral-7B-v0.3
 182 (Jiang et al., 2023), Qwen2.5-7B (Qwen et al., 2025), Qwen3-1.7B, and Qwen3-30B (Yang et al.,
 183 2025).

184 3.1 UNIVERSALITY AND SPARSITY

185



203 Figure 3: Translation score distribution for different models. According to the different translation
 204 scores, the model’s attention heads are divided into three categories: token alignment heads (red),
 205 infrequently activated heads (blue), and heads with near-zero activation (light blue). All models
 206 studied have token alignment heads, and the proportion of token alignment heads is relatively small.
 207 Most of the heads are either not activated or are activated at a low frequency.

208

209 Our analysis first reveals two fundamental properties: universality and sparsity. As illustrated in
 210 Figure 3, token alignment heads (defined as having a translation score > 0.1) are present in every
 211 model we examined, irrespective of its size, architecture, or training stage. This confirms that they
 212 are a universal, emergent feature of multilingual LLMs.

213 Concurrently, token alignment heads are exceptionally sparse. They constitute less than 8% of the
 214 total attention heads in all models, and as few as 3% in Mistral-7B-v0.3. Additionally, infrequently

215 ¹<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

activated heads during the translation process account for nearly 50%. The remaining attention heads, which are almost never activated, account for between 36% and 55%.

We next analyze the positional distribution of these heads within the model architecture. Figure 4 shows a consistent pattern: token alignment heads are predominantly concentrated in the middle layers of the models. In contrast, the earliest and latest layers contain very few token alignment heads. This aligns with the broader understanding of Transformer architectures, where initial layers are thought to handle surface-level feature extraction and final layers are responsible for structuring the output.

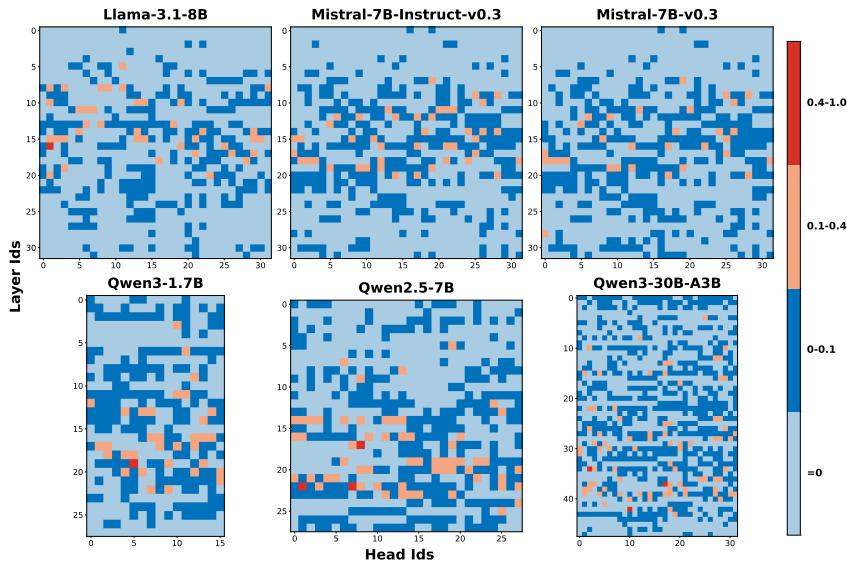


Figure 4: Positional distribution of translation scores in different models. Each heatmap visualizes the translation score distribution. The color intensity corresponds to the translation score, with warmer colors (red/orange) indicating higher scores and cooler colors (blue) indicating lower scores.

3.2 LANGUAGE CONSISTENCY

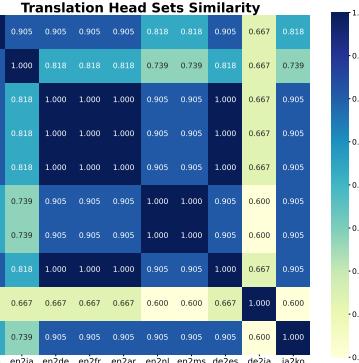


Figure 5: Jaccard similarity matrix of token alignment head sets across various language pairs. The axes list ten different language pairs. The token alignment heads between different language-pairs are very similar, most of the similarity scores are above 0.9.

To investigate the consistency of token alignment heads in the model across different language pairs, we selected the following language pairs: English-Chinese (en2zh), English-Japanese (en2ja),

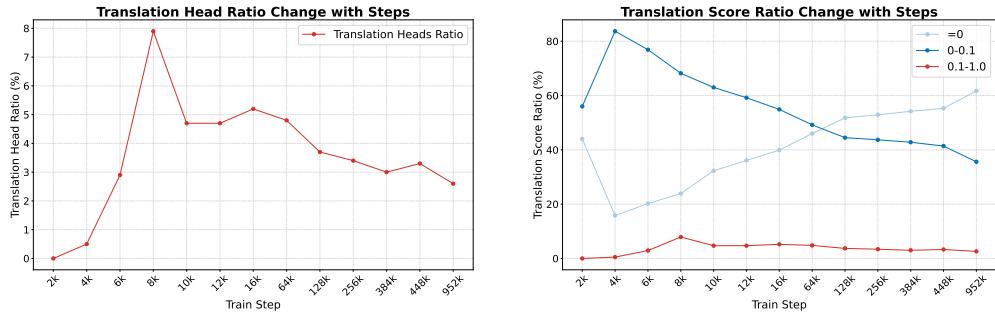
270 English-German (en2de), English-French (en2fr), English-Arabic (en2ar), English-Dutch (en2nl),
 271 English-Malay (en2ms), German-Spanish (de2es), German-Japanese (de2ja), and Japanese-Korean
 272 (ja2ko). These pairs cover a variety of linguistic families. Without loss of generality, we focus this
 273 analysis on the Llama-3.1-8B model. For each language pair, we selected the top 20 token alignment
 274 heads to form its translation set. Then we compute the pairwise similarity between these sets using
 275 the Jaccard index:

$$\text{Sim}_{S,T} = \frac{|S \cap T|}{|S \cup T|} \quad (2)$$

278 As shown in Figure 5, the similarities between all language pairs are relatively high, with Jaccard
 279 similarity scores consistently exceeding 0.8 for most pairs, and never dropping below 0.6. This
 280 indicates that a largely invariant set of attention heads is responsible for translation across diverse
 281 linguistic families, demonstrating the strong cross-lingual generalizability of token alignment heads.

283 4 FORMATION PROCESS OF TOKEN ALIGNMENT HEADS

285 This section investigates the formation process of token alignment heads during the model’s pre-
 286 training lifecycle. To trace their development, we trained an 8B parameter model, architecturally
 287 identical to Llama-2 model (Touvron et al., 2023), from scratch. The model was trained for a total
 288 of 15 trillion tokens. The specific composition of the dataset and the hyperparameters for training
 289 are detailed in Appendix A.1. To understand the formation process of token alignment heads,
 290 we analyzed model checkpoints at multiple intervals throughout training to map the evolutionary
 291 trajectory of token alignment heads.



303 (a) The ratio of token alignment heads rises sharply
 304 in the early training stages, reaching a peak of ap-
 305 proximately 8% at 8k step. Following this peak,
 306 the proportion drops and stabilizes around 5% be-
 307 tween 10k and 64k steps. Subsequently, it enters a
 308 long phase of gradual decline, settling at 2.6% by
 309 the end of training.

310 (b) The distribution of all attention heads by activ-
 311 ity level. The proportion of inactive heads (TS = 0, light blue line) starts at a low point but steadily in-
 312 creases throughout training, reaching over 60% by
 313 the final step. Conversely, the proportion of low-
 314 activity heads (dark blue line) begins as the domi-
 315 nant group but consistently decreases over time.

316 Figure 6: The evolutionary trajectory of token alignment heads during training

317 **Phase 1: Rapid Proliferation** (Early Training Stage, 0-8k steps). In the initial stages of training,
 318 as shown in Figure 6(a), the proportion of token alignment heads experiences a rapid proliferation,
 319 growing from zero to its peak. This period of rapid circuit formation coincides directly with the
 320 steepest gains in the model’s translation performance, where the FLORES chrF++ metric surged
 321 from 12.58 to 45.77. This suggests that the initial acquisition of translation ability is contingent on
 322 the rapid emergence of these specialized heads.

323 **Phase 2: Set Stabilization** (Early-to-Mid Training Stage, 10k-64k steps). From 10k to 64k steps,
 324 the proportion of token alignment heads stabilized around 5%, and the core set of these heads be-
 325 comes remarkably stable. To quantify this, we define a conditional overlap ratio metrics which
 326 measures the overlap between the token alignment head set at any given step (A) and the final refer-
 327 ence set at the end of training (B):

$$\frac{|A \cap B|}{|B|} \quad (3)$$

As shown in Figure 7, from approximately 8k steps onward, conditional overlap ratio remains consistently high. This indicates that the token alignment heads formed rapidly in the early stage of training are largely maintained throughout the subsequent training process.

Phase 3: Consolidation and Pruning (Mid-to-Late Training Stage, 64k-952k steps). In this longest phase of training, we observe a gradual decline in the overall proportion of token alignment heads, which settles at 2.6% (Figure 6(a)). Given that the core set of heads remains stable (Phase 2), this decline implies that heads with weaker or more redundant translation capabilities are being “pruned”—their translation scores fall below the threshold as the network refines its functions.

We hypothesize this pruning is part of a broader network-wide optimization towards increased sparsity and computational efficiency. This is corroborated by the shifting distribution of head activity shown in Figure 6(b). As training progresses, the proportion of completely inactive heads steadily increases, reaching 61.7% by the end. This happens at the expense of low-to-moderately active token alignment heads. In essence, the model learns to solve the translation task not by using more heads, but by relying more heavily on a smaller, more efficient, and highly specialized set of circuits, while deactivating others. This process of over-producing and then refining specialized circuits appears to be a key mechanism in the development of efficient neural networks.

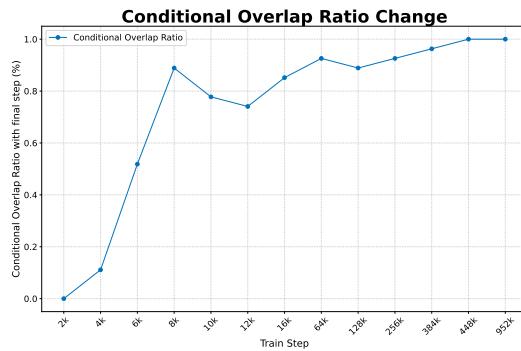


Figure 7: Stability of the token alignment head set over time. The conditional overlap ratio exhibits a steep and rapid increase during the initial training phase, rising from zero to nearly 0.9 at 8k step. From this point onward, the overlap remains consistently high, fluctuating but generally staying above 0.8 and approaching 1.0 by the end of training.

5 INFLUENCE ON DOWNSTREAM TASKS

In this section, we investigate the impact of token alignment heads on downstream benchmarks. First, we analyze the influence of token alignment heads on the model’s translation performance to demonstrate their causality. Here, we select the FLORES101 benchmark to evaluate the model’s translation performance. Next, we examine the impact of token alignment heads on the model’s general multilingual capabilities. We evaluate the following benchmarks: translated Hellaswag (Zellers et al., 2019), ARC-Easy and ARC-Challenge (Clark et al., 2018) which are detailed in Appendix A.2, XWinograd (Tikhonov & Ryabinin, 2021), XStoryCloze (Mostafazadeh et al., 2016), XNLI (Conneau et al., 2018), XCOPA (Ponti et al., 2020), and a localized multilingual variant of the MMMLU² test set denoted as XMMLU which includes JMMLU³, CMMLU (Li et al., 2024), AMMLU⁴, IndoMMLU (Koto et al., 2023), and VMLU⁵. We study the influence of token alignment heads by mask those token alignment heads.

²<https://huggingface.co/datasets/openai/MMMLU>

³<https://huggingface.co/datasets/nlp-waseda/JMMLU>

⁴<https://huggingface.co/datasets/Hennara/ammlu>

⁵<https://vmlu.ai>

378
379

5.1 TRANSLATION CAPACITY

380
381
382
383
384
385
386
387
388
389
390

In this subsection, we investigate the impact of token alignment heads on the translation performance of Llama-3.1-8B, Mistral-7B-v0.3, Qwen2.5-7B, Qwen3-1.7B, and Qwen3-30B using the FLORES-101 benchmark. We compare translation metrics before and after masking token alignment heads, as well as after masking random non-token alignment heads. To clearly illustrate the effects, we report the difference between the metrics of the masked models and those of the baseline models (without masking). Figure 12 presents the changes in BLEU and chrF++ scores. Masking token alignment heads leads to substantial declines in both metrics, with the largest drops exceeding 17 points for BLEU and 25 points for chrF++. In contrast, masking random heads has only a minimal effect. These results demonstrate that token alignment heads have a direct and significant influence on the models’ translation capabilities, a property we refer to as the causality of token alignment heads.

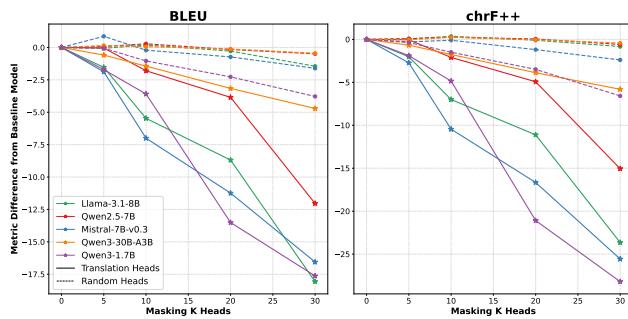
391
392
393
394
395
396
397
398
399
400
401

Figure 8: Impact of masking token alignment heads versus random heads on FLORES benchmark scores. Masking token alignment heads leads to a significantly larger performance drop compared to masking random heads

402
403
404
405
406

5.2 MULTILINGUAL CAPACITY

407
408
409
410
411
412
413
414
415

We then investigate the impact of ablating token alignment heads on a broader suite of multilingual benchmarks. From the results presented in Figure 9, we observe a clear hierarchy of dependency on token alignment heads. Benchmarks such as Hellaswag_ML, ARC_C_ML, and ARC_E_ML exhibit a significant performance drop (up to 10 points), which is consistently larger than the drop from ablating random heads. This suggests that these tasks, while not pure translation, partially rely on the cross-lingual mapping capabilities provided by token alignment heads. This functional overlap may stem from translation artifacts in their data creation process or a genuine need for cross-lingual conceptual alignment to solve the tasks.

416
417
418
419
420
421
422

In contrast, for other benchmarks like XNLI and XCOPA, token alignment heads demonstrate weak causality, as their ablation often results in a smaller performance drop than that of random ablation. This indicates that these tasks depend on different multilingual mechanisms within the model, likely operating at a higher semantic level that does not require the token-level mapping performed by token alignment heads. These findings suggest that token alignment heads provide a foundational cross-lingual alignment capability that various downstream tasks leverage to different degrees.

423
424

5.3 TOKEN ALIGNMENT HEAD AS DATA RATER

425
426
427
428
429
430

To further probe the relationship between token alignment heads and multilingual data, we introduce TRater, a data-filtering algorithm. TRater leverages token alignment heads to score data samples based on their importance to the translation mechanism. We compute the score of sample x as follows:

$$\text{score}(x) = \frac{1}{m} \sum_i (L(\theta_{\text{mask}}, x_i) - L(\theta, x_i)) \quad (4)$$

431

where L denotes the token level cross entropy loss, θ denotes the original model parameters, θ_{mask} represents the model parameters after masking the top 20 token alignment heads, i is the token

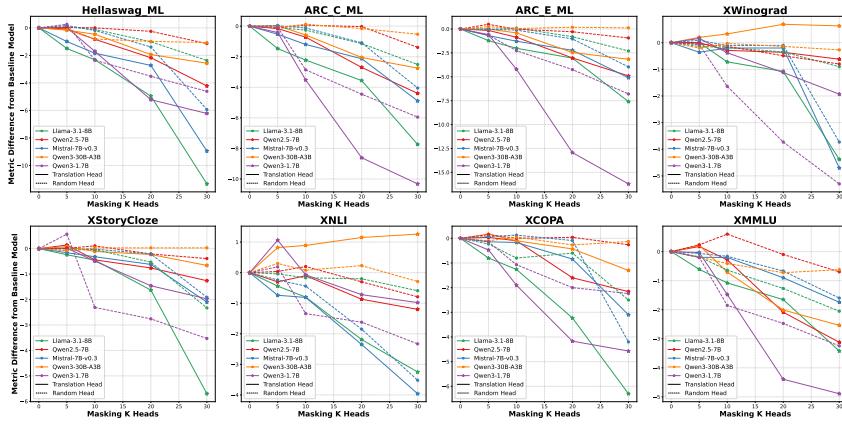


Figure 9: Performance (Accuracy) change across multilingual benchmarks when masking token alignment heads versus random heads. Tasks like Hellaswag_ML and ARC_ML show substantial drops when token alignment heads are ablated, while others such as XNLI and XCOPA are less affected, reflecting varying degrees of reliance on cross-lingual alignment.

index in x , and m is the total number of tokens. This score quantifies the performance degradation on a sample when token alignment heads are removed, with higher scores indicating greater reliance on these heads. We conduct experiments on the 1.5B model, training on a total of 1T tokens. The dataset comprises 700B tokens of English web data and 300B tokens of multilingual web data. Using the TRater algorithm, we score the 300B multilingual data, and we select the top 1.3% for each language. We design the following two experiments to validate the impact of the selected data:

Remove: From the baseline multilingual data, we exclude the selected data. To maintain the data proportions unchanged, we increase the volume of the remaining data per language. And we ensure no additional duplicates compared to the baseline.

Enhance: The selected data is triplicated, while the remaining data is proportionally down-sampled to preserve the overall language distribution.

Table 1: Performance of baseline, remove, and enhance settings across multiple metrics.

Model	flores_chrF++	MMMLU	Hellaswag_ML	ARC_C_ML	ARC_E_ML	XWinograd	XStoryCloze	XNLI	XCOPA	XMMLU
baseline	43.87	26.89	44.88	31.40	53.97	75.70	58.40	41.94	62.88	30.71
remove	41.33	26.58	44.69	31.37	54.54	73.63	58.15	42.10	63.20	30.84
enhance	46.68	26.71	44.95	31.54	54.94	74.33	58.44	41.72	63.70	30.88

The experimental results are presented in Table 1. From the table, we observe that the data filtered by token alignment heads is crucial for the model’s translation capabilities: the **Remove** setup exhibits a noticeable decline in translation performance compared to the baseline, whereas the **Enhance** setup shows a observable improvement.

However, the impact of this selected data on other multilingual benchmarks is less pronounced than its effect on FLORES. This can be attributed to two main factors. Firstly, the performance change observed in FLORES due to the selected data (about 2-3 points) is substantially smaller than the drops seen when ablating token alignment heads (over 10 points). Consequently, any impact on other multilingual benchmarks becomes less perceptible. Secondly, while the function of token alignment heads is partially leveraged by some non-translation tasks, these benchmarks only require the model to possess a foundational translation capability. Once this sufficient baseline is established, further increasing the proportion of translation-centric data yields diminishing returns for the model’s general multilingual performance. Indeed, a qualitative analysis confirms the selected data is highly translation-specific, consisting predominantly of bilingual corpora. Several representative examples are detailed in Appendix A.3.

486 6 RELATED WORK

488 A substantial body of work exists on understanding the internal mechanisms of large language mod-
 489 els. Based on the granularity of analysis, we broadly categorize these studies into three main areas:
 490 semantic space, neuron-level mechanisms, and head-level mechanisms.

492 6.1 SEMANTIC SPACE

494 Prior works (Wendler et al., 2024; Schut et al., 2025; Zhao et al., 2024; Wu et al., 2025b; Harrasse
 495 et al., 2025) study the geometry of multilingual representations and often concludes that models
 496 “think in English” or in a shared latent semantic space in middle layers. These works explain where
 497 multilingual information lives and how information exists (English or Language-Agnostic Space).
 498 Our results are complementary: we identify token alignment heads concentrated in similar middle
 499 layers and show that they implement token-level cross-lingual alignment, routing the aligned source
 500 token’s representation into the target position.

501 6.2 NEURON-LEVEL MECHANISMS

503 Utilizing causal mediation analysis across a diverse range of in-context learning (ICL) tasks, Todd
 504 et al. (2024) identified a key mechanism termed “function vectors”, which trigger the model to exe-
 505 cute specific procedural tasks. Similarly, Wang et al. (2024) employed causal mediation procedures
 506 to locate attention heads pivotal for machine translation, leveraging these heads to construct trans-
 507 lation vectors that mitigate language mismatch errors. In contrast, our work shifts focus from the
 508 task-triggering level to the token-execution level. We find that token alignment heads facilitate the
 509 actual cross-lingual alignment. Analogously, if the function vector acts as the “master switch” ac-
 510 tivating the translation mode, the token alignment heads are the vital machinery carrying out the
 511 translation itself.

512 From the perspective of languages, some studies (Liu et al., 2025; Zhao et al., 2024) pinpoint neurons
 513 that are specialized for encoding language identity and language-specific features and shows that
 514 ablating or fine-tuning them selectively affects particular languages. This explains which subcircuits
 515 are responsible for “being in language X”. In contrast, we operate at the head level and focus on the
 516 cross-lingual alignment step, i.e., how information moves between languages during translation.

517 6.3 HEAD-LEVEL MECHANISMS

519 Works on induction heads, retrieval heads, and circuits (Elhage et al., 2021b; Olsson et al., 2022; Wu
 520 et al., 2025a; Bricken et al., 2024; Zhang et al., 2024) shows that a small number of specialized heads
 521 can explain non-trivial capabilities such as in-context learning or long-context retrieval. Recent work
 522 (Liu et al., 2025; Zhang et al., 2025) identifies language heads or translation-related heads by ranking
 523 heads via their impact on downstream loss, perplexity, or logits on specific benchmarks, sometimes
 524 using path patching. Our approach is closely related but uses a different identification signal: we
 525 define Token Alignment Heads (TAH) using alignment-based translation score—heads are selected
 526 because they consistently link target tokens to their externally aligned source tokens, independent of
 527 any particular evaluation task. Masking experiments are then used only as a causal validation step.
 528 This makes our notion of specialization explicitly lexical and cross-lingual.

530 7 CONCLUSION

532 In this paper, we identified a special class of attention heads responsible for mapping source lan-
 533 guage tokens to target language tokens during translation. We experimentally confirmed that these
 534 heads are universal, consistent, and have a direct causal effect on the model’s translation capabilities.
 535 We also uncovered their evolutionary process during pre-training, which involves rapid formation,
 536 stabilization, and pruning. More importantly, We found that a tiny fraction of critical data filtered by
 537 token alignment heads, proves decisive for translation performance but its impact on other multilin-
 538 gual tasks is less pronounced. This finding suggests that translation operates as a separable module
 539 within LLMs. Our work pave the way for more efficient and robust multilingual systems, enabling
 targeted architectural innovations, data curation strategies guided by mechanistic understanding.

540 REFERENCES

541

542 Anthropic. Claude 4. Large language model, 2025. URL <https://www.anthropic.com/news/clause-4>.

543

544 Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. On the cross-lingual transferability of mono-
545 lingual representations. In *Proceedings of the 58th Annual Meeting of the Association for Com-
546 putational Linguistics*. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.
547 acl-main.421. URL <http://dx.doi.org/10.18653/v1/2020.acl-main.421>.

548

549 Trenton Bricken, Adly Templeton, Alexander G. M. G. Anderson, Liane Lovitt, Jerry Chen,
550 Kamal Ndousse, William Saunders, Christopher Olah, and Tom Henighan. Scaling
551 monosemanticity: The sparse manifold hypothesis. *Transformer Circuits*, May 2024. URL
552 <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>. Accessed: 2025-11-21.

553

554 Peter F Brown, Stephen A Della Pietra, Vincent J Della Pietra, and Robert L Mercer. The mathe-
555 matics of statistical machine translation: Parameter estimation. *Computational linguistics*, 19(2):
556 263–311, 1993.

557

558 Tyler A. Chang, Zhuowen Tu, and Benjamin K. Bergen. The geometry of multilingual language
559 model representations, 2022. URL <https://arxiv.org/abs/2205.10964>.

560

561 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
562 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge,
2018. URL <https://arxiv.org/abs/1803.05457>.

563

564 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
565 Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the
566 frontier with advanced reasoning, multimodality, long context, and next generation agentic capa-
567 bilities. *arXiv preprint arXiv:2507.06261*, 2025.

568

569 Alexis Conneau, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R. Bowman, Holger
570 Schwenk, and Veselin Stoyanov. Xnli: Evaluating cross-lingual sentence representations, 2018.
URL <https://arxiv.org/abs/1809.05053>.

571

572 Arindrima Datta, Bhuvana Ramabhadran, Jesse Emond, Anjuli Kannan, and Brian Roark.
573 Language-agnostic multilingual modeling. In *ICASSP 2020-2020 IEEE International Conference
on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8239–8243. IEEE, 2020.

574

575 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann,
576 Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep
577 Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt,
578 Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and
579 Chris Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*,
2021a. <https://transformer-circuits.pub/2021/framework/index.html>.

580

581 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann,
582 Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Roger
583 Grosse, Zac Hatfield-Dodds, Danny Hernandez, Tristan Jones, Jackson Kernion, Liane Lovitt, Ka-
584 mal Ndousse, Sam Ringer, Sam Johnston, Scott Thompson, Tom B. Brown, Jared Kaplan, Chris
585 Olah, and Dario Amodei. A mathematical framework for transformer circuits. *Transformer Cir-
586 cuits*, Jul 2021b. URL <https://transformer-circuits.pub/2021/framework/index.html>. Accessed: 2025-11-21.

587

588 Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan
589 Belinkov. Causal analysis of syntactic agreement mechanisms in neural language models, 2021.
590 URL <https://arxiv.org/abs/2106.06087>.

591

592 Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju,
593 Sanjana Krishnan, Marc'Aurelio Ranzato, Francisco Guzman, and Angela Fan. The flores-
101 evaluation benchmark for low-resource and multilingual machine translation, 2021. URL
<https://arxiv.org/abs/2106.03193>.

594 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 595 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,
 596 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Ko-
 597 reniev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava
 598 Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux,
 599 Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret,
 600 Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius,
 601 Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary,
 602 Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab
 603 AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco
 604 Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind That-
 605 tai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Kore-
 606 vaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra,
 607 Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Ma-
 608 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu,
 609 Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jong-
 610 soot Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasudevan Alwala,
 611 Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid
 612 El-Arini, Krishika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren
 613 Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin,
 614 Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi,
 615 Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew
 616 Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Ku-
 617 mar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoy-
 618 chev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan
 619 Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajwal Bhargava, Pratik Dubal, Praveen Krishnan,
 620 Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ra-
 621 mon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Ro-
 622 hit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan
 623 Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell,
 624 Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparth, Sheng
 625 Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer
 626 Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman,
 627 Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mi-
 628 haylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor
 629 Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei
 630 Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang
 631 Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Gold-
 632 schlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning
 633 Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh,
 634 Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria,
 635 Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein,
 636 Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, An-
 637 drew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, An-
 638 nie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel,
 639 Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leon-
 640 hardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu
 641 Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Mon-
 642 talvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao
 643 Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia
 644 Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide
 645 Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le,
 646 Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily
 647 Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smo-
 648 thers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni,
 649 Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia
 650 Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan,
 651 Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harri-
 652 son Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj,

648 Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James
 649 Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jen-
 650 nifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang,
 651 Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Jun-
 652 jie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy
 653 Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang,
 654 Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell,
 655 Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa,
 656 Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias
 657 Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L.
 658 Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike
 659 Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari,
 660 Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan
 661 Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong,
 662 Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent,
 663 Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar,
 664 Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Ro-
 665 driguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy,
 666 Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin
 667 Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon,
 668 Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ra-
 669 maswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha,
 670 Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal,
 671 Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satter-
 672 field, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj
 673 Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo
 674 Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook
 675 Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Ku-
 676 mar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov,
 677 Wei Li, Wencheng Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiao-
 678 jian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia,
 679 Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao,
 680 Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhao-
 681 duo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL
 682 <https://arxiv.org/abs/2407.21783>.

683 Ping Guo, Yubing Ren, Binbin Liu, Fengze Liu, Haobin Lin, Yifan Zhang, Bingni Zhang, Taifeng
 684 Wang, and Yin Zheng. Exploring polyglot harmony: On multilingual data allocation for large
 685 language models pretraining, 2025. URL <https://arxiv.org/abs/2509.15556>.

686 Wenhan Han, Yifan Zhang, Zhixun Chen, Binbin Liu, Haobin Lin, Bingni Zhang, Taifeng Wang,
 687 Mykola Pechenizkiy, Meng Fang, and Yin Zheng. Mubench: Assessment of multilingual capabil-
 688 ities of large language models across 61 languages, 2025. URL <https://arxiv.org/abs/2506.19468>.

689 Abir Harrasse, Florent Draye, Zhijing Jin, and Bernhard Schölkopf. Tracing multilingual repres-
 690 entations in llms with cross-layer transcoders, 2025. URL <https://arxiv.org/abs/2511.10840>.

691 Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of
 692 local experts. *Neural computation*, 3(1):79–87, 1991.

693 Dongsheng Jiang, Yuchen Liu, Songlin Liu, Jin'e Zhao, Hao Zhang, Zhen Gao, Xiaopeng Zhang, Jin
 694 Li, and Hongkai Xiong. From clip to dino: Visual encoders shout in multi-modal large language
 695 models. *arXiv preprint arXiv:2310.08825*, 2023.

696 Zae Myung Kim, Laurent Besacier, Vassilina Nikoulina, and Didier Schwab. Do multilingual neu-
 697 ral machine translation models contain language pair specific attention heads? *arXiv preprint*
 698 *arXiv:2105.14940*, 2021.

702 Fajri Koto, Nurul Aisyah, Haonan Li, and Timothy Baldwin. Large language models only pass
 703 primary school exams in indonesia: A comprehensive test on indommlu, 2023. URL <https://arxiv.org/abs/2310.04928>.

704

705 Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. Revealing the dark secrets
 706 of bert, 2019. URL <https://arxiv.org/abs/1908.08593>.

707

708 Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy
 709 Baldwin. Cmmlu: Measuring massive multitask language understanding in chinese, 2024. URL
 710 <https://arxiv.org/abs/2306.09212>.

711

712 Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner,
 713 Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly
 714 Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam
 715 Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley
 716 Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language
 717 model. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/attribution-graphs/biology.html>.

718

719 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
 720 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.

721

722 Xin Liu, Qiyang Song, Qihang Zhou, Haichao Du, Shaowen Xu, Wenbo Jiang, Weijuan Zhang,
 723 and Xiaoqi Jia. Focusing on language: Revealing and exploiting language attention heads in
 724 multilingual large language models, 2025. URL <https://arxiv.org/abs/2511.07498>.

725

726 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. URL <https://arxiv.org/abs/1711.05101>.

727

728 Weicheng Ma, Kai Zhang, Renze Lou, Lili Wang, and Soroush Vosoughi. Contributions of trans-
 729 former attention heads in multi-and cross-lingual tasks. *arXiv preprint arXiv:2108.08375*, 2021.

730

731 Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one?, 2019.
 732 URL <https://arxiv.org/abs/1905.10650>.

733

734 Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vander-
 735 wende, Pushmeet Kohli, and James Allen. A corpus and evaluation framework for deeper under-
 736 standing of commonsense stories, 2016. URL <https://arxiv.org/abs/1604.01696>.

737

738 Franz Josef Och and Hermann Ney. A systematic comparison of various statistical alignment mod-
 739 els. *Computational Linguistics*, 29(1):19–51, 2003. doi: 10.1162/089120103321337421. URL
 740 <https://aclanthology.org/J03-1002/>.

741

742 Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan,
 743 Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli,
 744 Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane
 745 Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish,
 746 and Chris Olah. In-context learning and induction heads, 2022. URL <https://arxiv.org/abs/2209.11895>.

747

748 OpenAI. Gpt-5. Large language model, 2025. URL <https://openai.com/index/introducing-gpt-5/>.

749

750 Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korho-
 751 nen. Xcopa: A multilingual dataset for causal commonsense reasoning, 2020. URL <https://arxiv.org/abs/2005.00333>.

752

753 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 754 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 755 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li,
 Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,
 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025.
 URL <https://arxiv.org/abs/2412.15115>.

756 Lisa Schut, Yarin Gal, and Sebastian Farquhar. Do multilingual llms think in english? *arXiv preprint*
 757 *arXiv:2502.15603*, 2025.

758

759 Alexey Tikhonov and Max Ryabinin. It's all in the heads: Using attention heads as a baseline for
 760 cross-lingual transfer in commonsense reasoning, 2021. URL <https://arxiv.org/abs/2106.12066>.

761

762 Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau.
 763 Function vectors in large language models, 2024. URL <https://arxiv.org/abs/2310.15213>.

764

765

766 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
 767 Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher,
 768 Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy
 769 Fu, Wenjin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,
 770 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel
 771 Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee,
 772 Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra,
 773 Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi,
 774 Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh
 775 Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen
 776 Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic,
 777 Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models,
 778 2023. URL <https://arxiv.org/abs/2307.09288>.

779

780 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023. URL <https://arxiv.org/abs/1706.03762>.

781

782 Jesse Vig and Yonatan Belinkov. Analyzing the structure of attention in a transformer language
 783 model, 2019. URL <https://arxiv.org/abs/1906.04284>.

784

785 Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head
 786 self-attention: Specialized heads do the heavy lifting, the rest can be pruned, 2019. URL <https://arxiv.org/abs/1905.09418>.

787

788 Weichuan Wang, Zhaoyi Li, Defu Lian, Chen Ma, Linqi Song, and Ying Wei. Mitigating the lan-
 789 guage mismatch and repetition issues in llm-based machine translation via model editing, 2024.
 790 URL <https://arxiv.org/abs/2410.07054>.

791

792 Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. Do llamas work in En-
 793 glish? on the latent language of multilingual transformers. In Lun-Wei Ku, Andre Martins, and
 794 Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Com-
 795 putational Linguistics (Volume 1: Long Papers)*, pp. 15366–15394, Bangkok, Thailand, August
 796 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.820. URL
 797 <https://aclanthology.org/2024.acl-long.820/>.

798

799 Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. Retrieval head mechanis-
 800 tically explains long-context factuality. In *The Thirteenth International Conference on Learning
 801 Representations*, 2025a. URL <https://openreview.net/forum?id=EytBpUGB1Z>.

802

803 Zhaofeng Wu, Xinyan Velocity Yu, Dani Yogatama, Jiasen Lu, and Yoon Kim. The semantic hub
 804 hypothesis: Language models share semantic representations across languages and modalities,
 805 2025b. URL <https://arxiv.org/abs/2411.04986>.

806

807 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 808 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 809 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger

810 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
811 Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
812

813 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
814 chine really finish your sentence?, 2019. URL <https://arxiv.org/abs/1905.07830>.
815

816 Hongbin Zhang, Kehai Chen, Xuefeng Bai, Xiucheng Li, Yang Xiang, and Min Zhang. Exploring
817 translation mechanism of large language models. *arXiv preprint arXiv:2502.11806*, 2025.
818

819 Ruochen Zhang, Qinan Yu, Matianyu Zang, Carsten Eickhoff, and Ellie Pavlick. The same but
820 different: Structural similarities and differences in multilingual language modeling, 2024. URL
821 <https://arxiv.org/abs/2410.09223>.
822

823 Yiran Zhao, Wenxuan Zhang, Guizhen Chen, Kenji Kawaguchi, and Lidong Bing. How do large
824 language models handle multilingualism? *Advances in Neural Information Processing Systems*,
825 37:15296–15319, 2024.
826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 **A APPENDIX**
865866 **A.1 TRAINING SETTING**
867868 In Section 4, We use AdamW (Loshchilov & Hutter, 2019) optimizer to train a 8B dense model
869 with a structure identical to Llama-2 model Touvron et al. (2023). The training dataset, totaling 15
870 trillion tokens, comprises English and multilingual data sourced from cleaned and filtered ccwarc,
871 along with open-source resources including Wikipedia, books, academic papers, mathematics, code,
872 and parallel corpora. The hyperparameters were set as follows: learning rate (lr) = 3.6×10^{-4} ,
873 global batch size (gbs) = 4096, sequence length = 4096, weight decay = 0.1. We employed a co-
874 sine learning rate scheduler that decayed to 0.1 of the peak learning rate. For our analysis, we se-
875 lected model checkpoints at 2,000, 4,000, 6,000, 8,000, 12,000, 16,000, 64,000, 128,000, 256,000,
876 384,000, 448,000, and 952,000 steps (near the end of training).877 In Section 5.3, we trained the 1.5B model using the AdamW optimizer with the following hyper-
878 parameters: global batch size (gbs) = 4096, sequence length = 4096, weight decay = 0.1, learning
879 rate (lr) = 5.0×10^{-4} , and with lr cosine decay to 5.0×10^{-5} , the multilingual web data includes
880 a total of 17 languages, specifically including German, Spanish, French, Indonesian, Thai, Korean,
881 Vietnamese, Arabic, Turkish, Italian, Malay, Chinese, Portuguese, Japanese, Dutch, Russian and
882 Filipino. The proportion of different languages is determined by the method in Guo et al. (2025).883
884 **A.2 TRANSLATED BENCHMARK**
885886 For evaluating Hellaswag_ML, ARC_C_ML and ARC_E_ML, we use the MuBench dataset (Han
887 et al., 2025). The evaluation covers 18 languages represented in our training data, namely: English,
888 German, Spanish, French, Indonesian, Thai, Korean, Vietnamese, Arabic, Turkish, Italian, Malay,
889 Chinese, Portuguese, Japanese, Dutch, Russian and Filipino.890
891 **A.3 TEXT CASES FOR TRATER**
892893
894 Table 2: Text cases filtered by TRater for German, French and Spanish.
895

896 Language	897 Example
898 899 900 901 902 903 904 905 German	906 Zwischen zwei Seen, die unterschiedlicher kaum sein können, liegt 907 Wandlitz. Der Wandlitzsee, bebaut, kaum zugänglich mit unzähligen 908 Wassergrundstücken, der Liebnitzsee, frei zugänglich, Badestelle, 909 Fähre zur Insel und Naherholungsgebiet im Buchenwald... Wandlitz 910 is located between two lakes that could hardly be more different. The 911 Wandlitzsee, built-up, hardly accessible with countless water proper- 912 ties, the Liebnitzsee, freely accessible, bathing area, ferry to the island 913 and local recreation area in the beech forest...
914 915 916 917 French	918 Bruno Houssin, designer français et professeur à l'école de design de 919 Nantes. Diplômé de l'école Boulle de Paris, en Architecture intérieure 920 et Design en 1986... Bruno Houssin, French designer and teacher at the 921 Nantes School of Design. Graduated from the Boulle school of Paris, 922 in Interior Architecture and Design in 1986...
923 924 925 926 Spanish	927 Las placas tectónicas son como grandes balsas que se reparten por 928 toda la corteza del planeta. Unas son de carácter continental, otras 929 de carácter oceánico, contando las primeras con un espesor mayor que 930 el de las segundas... Le tremblement de terre en Haïti est partie de 931 l'ensemble de la libération des tensions accumulées à l'occasion du 932 mouvement des plaques tectoniques dans les Caraïbes et en Amérique 933 du Nord...

918 Table 3: Text cases filtered by TRater for Italian, Portuguese, Chinese and Dutch.
919

920 Language	921 Example
922	Dorothy Bhawl è un artista autodidatta interessato al mondo contemporaneo, soprattutto alla realtà che appartiene e avvolge questa epoca: quello della comunicazione, social network, spiritualità e grottesco con un sentimento di odi et amo... Dorothy Bhawl is a self-taught artist interested in the contemporary world, especially in the reality that belongs and envelops this era: that of communication, social networks, spirituality and grotesque with a feeling of hatred and love...
923	
924	
925 Italian	
926	
927	
928	
929	André Rigatti Centro Universitário Maria Antonia USP Sempre próximas a suas bordas, as pinturas de André Rigatti possuem pequenas aberturas, por onde se deixa ver o processo que dá origem a trabalhos de textura matérica mais ou menos acentuada, resultados da sobreposição de diversas camadas de tinta, aplicadas cada uma seguindo uma direção diferente do pincel... Always close to its edges, André Rigatti's paintings have small openings, where you can see the process that gives rise to more or less accentuated texture work, results of the overlapping of several layers of paint, applied each following a direction...
930	
931	
932	
933 Portuguese	
934	
935	
936	
937	
938	原文: 版印书籍, 唐人尚未盛为之。自冯瀛王始印五经, 已后典记, 皆为版本。庆历中, 有布衣毕升, 又为活版。其法用胶泥刻字, 薄如钱唇, 每字为一印, 火烧令坚。先设一铁板, 其上以松脂腊和纸灰之类冒之。欲印则以一铁范置铁板上, 乃密布字印。满铁范为一板, 持就火炀之, 药稍熔, 则以一平板按其面, 则字平如砥... 译文: 用刻板印刷书籍, 唐朝人还没有大规模采用它。五代始才开始印刷五经, 以后的各种图书都是雕板印刷本。庆历年间, 有位平民毕升, 又创造了活版。他的方法是用胶泥刻成字, 字薄得像铜钱的边缘, 每个字制成一个字模, 用火来烧使它坚硬。先设置一块铁板, 它的上面用松纸、蜡混合纸灰这一灰东西覆盖它...
939	
940	
941	
942	
943 Chinese	
944	
945	
946	
947	
948	
949	Kees Blom (Apeldoorn, 1968) komt uit een artistieke familie. Zijn vader had al een passie voor de schilderkunst maar pas zoon Kees lukt het om de stap naar zelfstandig kunstschilder te zetten...Kees Blom (Apeldoorn, 1968) comes from an artistic family. While his father already had a passion for painting, son Kees succeeds in taking the step to become an independent painter...
950	
951	
952	
953	
954	
955	
956	
957	
958 A.4 CASE STUDY FOR TOKEN ALIGNMENT HEADS PRUNING	
959	
960	To clarify the role of Token Alignment Heads (TAHs), we conducted a systematic and detailed analysis of the model's performance on translation tasks when TAHs are masked (specifically focusing on cases where the model translated correctly before masking TAHs but failed to do so afterwards). We observed that the failure modes resulting from masking THs can be broadly categorized into three types:
961	
962	
963	
964	
965	1. No Translation (46%) The model fails to generate the target language and merely repeats the source text content. As illustrated in Figure 10, masking the Token Alignment Heads completely disables the model's word alignment capability for certain queries, preventing the generation of target language output.
966	
967	
968	
969	2. Missing Details (36%) The translated output lacks specific details found in the source text. In these instances, the model's word alignment capability is partially compromised. While the model can still generate text in the target language, it fails to correctly map certain source information to the target output. This results in the omission of key details, leading to incomplete or imprecise
970	
971	

972 Translate this sentence from English to Simplified Chinese
 973 Sentence: In late 2017, Siminoff appeared on shopping television channel QVC.
 974 [No Heads Pruned] Answer: 2017年末，辛明诺夫出现在购物电视频道QVC上。
 975 [30 Token Alignment Heads Pruned] Answer: In late 2017, Siminoff appeared on shopping television channel QVC.
 976 [30 Random Heads Pruned] Answer: 在2017年晚些时候，Siminoff出现在购物电视频道QVC上。
 977
 978 Sentence: One antibody cocktail, ZMapp, initially showed promise in the field, but formal studies indicated it had less benefit than
 979 sought in preventing death.
 980 [No Heads Pruned] Answer: 一种抗体混合物，ZMapp，最初在野外显示出希望，但正式研究表明它比预期的更少有益于
 981 预防死亡。
 982 [30 Token Alignment Heads Pruned] Answer: One antibody cocktail, ZMapp, showed promise in the field, but formal studies
 983 indicated it had less benefit than sought in preventing death.
 984 [30 Random Heads Pruned] Answer: 一种抗体混合物，ZMapp，最初在野外表现出希望，但正式研究表明，它比预期预防死
 985 亡的作用要小。
 986
 987 Sentence: Scientists say this animal's plumage was chestnut-brown on top with a pale or carotenoid-colored underside.
 988 [No Heads Pruned] Answer: 科学家说，这种动物的羽毛上面是棕色，下面是淡黄色或是橙色。
 989 [30 Token Alignment Heads Pruned] Answer: Scientists say this animal's plumage was chestnut-brown on top with a pale or
 990 carotenoid-colored underside.
 991 [30 Random Heads Pruned] Answer: 科学家们说，这种动物的羽毛是棕色顶部和淡色或橘色底部。
 992
 993
 994

990 Figure 10: No Translation cases for masking token alignment heads. For the queries shown in the
 991 figure, after masking the token alignment heads, the model no longer outputs content related to the
 992 target language, but simply repeats the source text.

995 translations. Figure 11 presents three such cases, where the text highlighted in green represents the
 996 details omitted by the model.

997 3. Hallucination (18%) The translated output contains content that is completely absent from the
 998 source text. This is likely because, without token alignment capabilities, the model cannot directly
 999 associate source tokens with target tokens. Consequently, it resorts to generating content in the target
 1000 language space that is only loosely or tangentially related to the source tokens. Figure 12 displays
 1001 three examples of this category, where the text highlighted in red indicates output generated by the
 1002 model that is entirely unrelated to the source text.

1003
 1004 Translate this sentence from English to Simplified Chinese
 1005 Sentence: Like some other experts, he is skeptical about whether diabetes can be cured, noting that these findings have no relevance
 1006 to people who already have Type 1 diabetes.
 1007 [No Heads Pruned] Answer: 他和其他专家一样，对糖尿病能否治愈持怀疑态度，指出这些发现对已经患有1型糖尿病的人
 1008 没有任何意义。
 1009 [30 Token Alignment Heads Pruned] Answer: 他认为这些发现对那些已经有了糖尿病的人没有什么帮助。
 1010 [30 Random Heads Pruned] Answer: 和其他专家一样，他对糖尿病能否治愈持怀疑态度，指出这些发现对已经患有1型糖
 1011 尿病的人没有任何意义。
 1012
 1013 Sentence: This not only confirms that at least some dinosaurs had feathers, a theory already widespread, but provides details fossils
 1014 generally cannot, such as color and three-dimensional arrangement.
 1015 [No Heads Pruned] Answer: 这不仅证实了至少有一些恐龙有羽毛，一个已经广泛流传的理论，而且提供了化石通常不能提
 1016 供的细节，如颜色和三维排列。
 1017 [30 Token Alignment Heads Pruned] Answer: 这不仅证实了至少有一种恐龙有羽毛，而且还证实了这种羽毛的颜色和羽毛
 1018 的三维结构。
 1019 [30 Random Heads Pruned] Answer: 这不仅证实了至少有些恐龙有羽毛——这一理论已经广为流传——而且还提供了化石
 1020 通常无法提供的细节，比如颜色和三维排列。
 1021
 1022 Sentence: During his trip, Iwasaki ran into trouble on many occasions.
 1023 [No Heads Pruned] Answer: 在他旅行期间，伊瓦斯基遇到了很多麻烦。
 1024 [30 Token Alignment Heads Pruned] Answer: 在他旅行的过程中，他遇到了很多麻烦。
 1025 [30 Random Heads Pruned] Answer: 在旅行期间，伊瓦斯基遇到了很多麻烦。

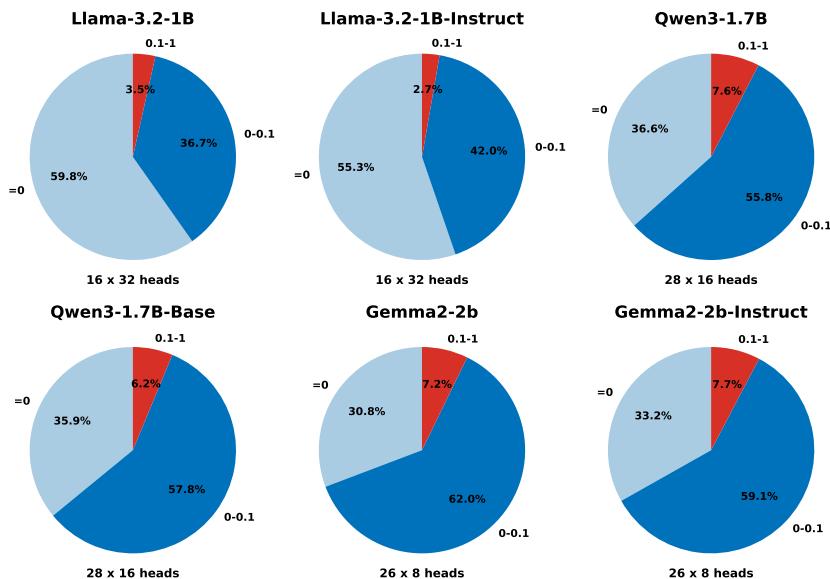
1022 Figure 11: Cases of Missing Details for masking token alignment heads. For the queries shown in
 1023 the figure, after masking the token alignment heads, the model is still able to generate target
 1024 language content. However, because some cross-lingual token alignment capabilities are missing,
 1025 the model loses certain details present in the source text.

1026
1027 Translate this sentence from English to Simplified Chinese
1028 Sentence: During the 1976 elections he advised Carter on foreign policy, then served as National Security Advisor (NSA) from 1977
1029 to 1981, succeeding Henry Kissinger.
1030 [No Heads Pruned] Answer: 1976年选举期间，他为卡特提供建议，1977年至1981年担任国家安全顾问 (NSA)，接替亨利·基辛格。
1031 [30 Token Alignment Heads Pruned] Answer: 1976年，卡特担任美国总统期间，布什担任国务卿，为卡特提供建议。
1032 [30 Random Heads Pruned] Answer: 1976年选举期间，他为卡特总统提供了外交政策建议，1977年至1981年期间，他担任国家安全顾问 (NSA)，接替亨利·基辛格。
1033
1034 Sentence: The vehicle itself was taken away from the scene of the accident at approximately 1200 GMT on the same day.
1035 [No Heads Pruned] Answer: 车辆本身在事故现场被拖走，时间大约是当天下午12点。
1036 [30 Token Alignment Heads Pruned] Answer: 该车辆在事故发生后被送往了1200公里以外的地方。
1037 [30 Random Heads Pruned] Answer: 车辆本身在事故现场被拖走，时间大约是当天下午12点。
1038
1039 Sentence: Several large television screens were installed in various places in Rome to let the people watch the ceremony.
1040 [No Heads Pruned] Answer: 罗马市内各处安装了数块电视屏幕，让人们观看仪式。
1041 [30 Token Alignment Heads Pruned] Answer: 在罗马，人们可以在许多地方看到电视节目。
1042 [30 Random Heads Pruned] Answer: 罗马各地安装了数个大型电视屏幕，供人们观看仪式。

1043 Figure 12: Hallucination cases for masking token alignment heads. For the queries shown in the
1044 figure, after masking the token alignment heads, the model exhibits hallucinations, specifically gen-
1045 erating output that contains information not present in the source text.

1046
1047
1048 **A.5 BASIC PROPERTIES OF TOKEN ALIGNMENT HEADS ACROSS COMPREHENSIVE MODEL**
1049 **FAMILIES**
1050

1051 To provide a more thorough and comprehensive evaluation of Token Alignment Heads, we con-
1052 ducted experiments across four well-established multilingual LLM families: Llama3, Qwen3, Mis-
1053 tral, and Gemma2. Our experimental design ensures broad coverage by including models of varying
1054 scales within each family: small (1B/2B parameters), medium (8B/9B parameters), and large (over
1055 13B parameters). Where applicable, we evaluate the instruction-tuned (Instruct) variants. Addition-
1056 ally, for families that feature Mixture-of-Experts (MoE) architectures, our study includes a specific
1057 analysis of these models.



1078 Figure 13: Translation score distribution for small size group models.
1079

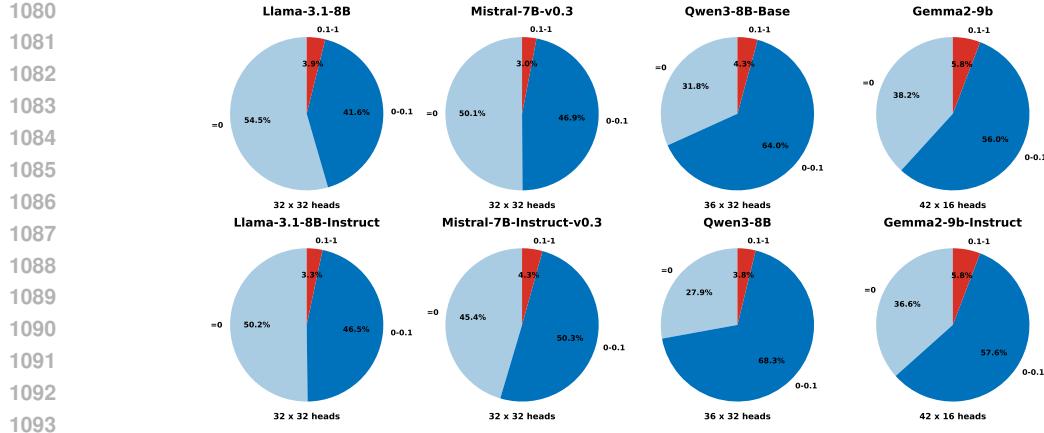


Figure 14: Translation score distribution for medium size group models.

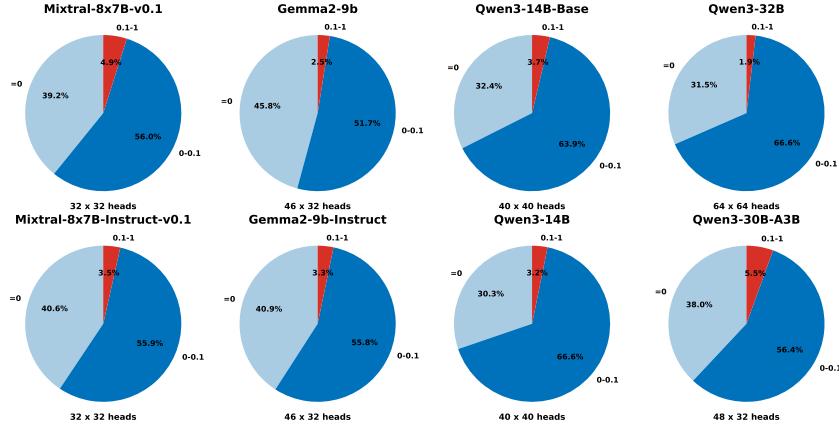


Figure 15: Translation score distribution for large size group models.

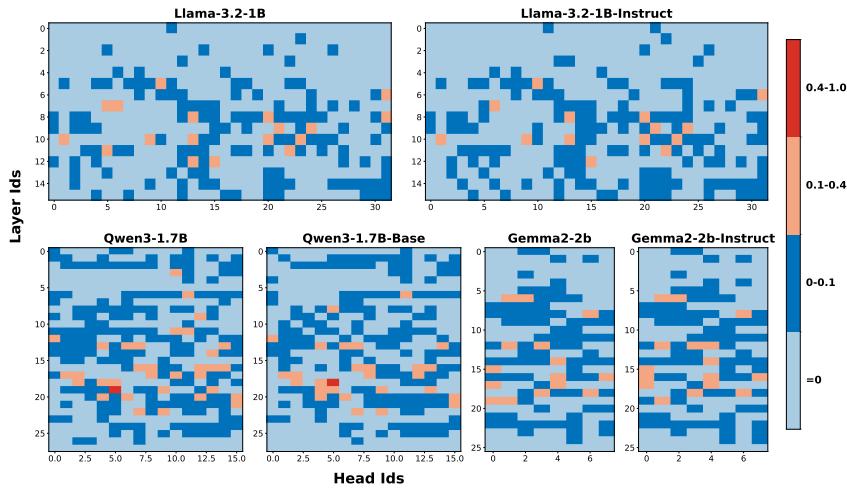
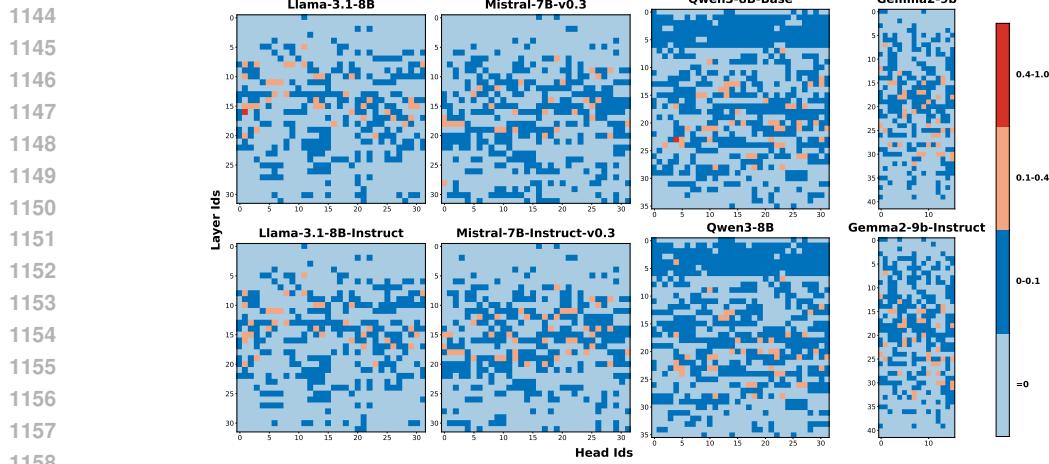


Figure 16: Positional distribution of translation scores in small size group models.

For clarity in presenting our experimental findings, we have organized the models into three distinct groups according to their scale: small, medium, and large. (The Mixtral-8x7B model is classified

1134 within the medium-scale group, as its quantity of attention heads is analogous to that of other models
 1135 in this tier.) Our initial analysis focuses on demonstrating the basic characteristics of Token Align-
 1136 ment Heads. As illustrated in Figure 13, Figure 14 and Figure 15, Token Alignment Heads are
 1137 a pervasive phenomenon, consistently identified across all models under investigation—irrespective
 1138 of model scale, architecture (dense versus MoE), or training paradigm (Base versus Instruct). Cru-
 1139 cially, these heads universally demonstrate the property of sparsity. Additionally, Figure 16, Figure
 1140 17 and Figure 18 illustrates the positional distribution of token alignment heads across these mod-
 1141 els. The general distribution pattern observed aligns with the description provided in the main body
 1142 of the paper.

1143



1158

1159

Figure 17: Positional distribution of translation scores in medium size group models.

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

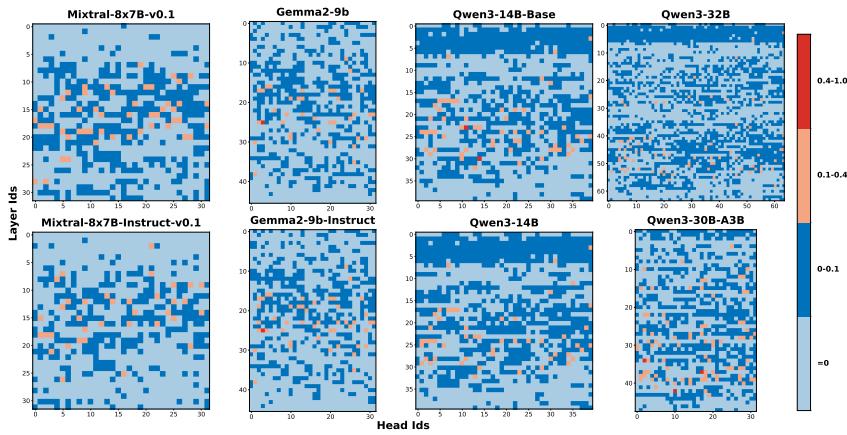


Figure 18: Positional distribution of translation scores in large size group models.

1188
1189

A.6 DOWNSTREAM TASKS INFLUENCE ACROSS COMPREHENSIVE MODEL FAMILIES

1190
1191
1192

In this section, we present the impact of token alignment heads on downstream task performance, analyzed on a group-by-group basis. To better evaluate translation performance, we supplement the BLEU and chrF++ metrics from the main text with two additional metrics: BLEURT and COMET.

1193
1194
1195
1196
1197
1198
1199
1200
1201
1202

As can be seen, the experimental conclusions are consistent with those in the main text. As illustrated in Figure 19, Figure 20 and Figure 21, for BLEURT and COMET, masking token alignment heads leads to a significant drop in scores, whereas masking random heads results in only a minimal decrease. This demonstrates the causal role of token alignment heads in translation capability. For Hellaswag_ML, ARC_C_ML, and ARC_E_ML, masking token alignment heads causes a larger performance drop than masking random heads, but the overall magnitude of the decrease is far less pronounced than that for the translation metrics. This suggests that these metrics rely to some extent on token alignment capabilities. In contrast, for metrics like XWinograd and XNLI, performance after masking token alignment heads can be better than after masking random heads, indicating that these metrics prioritize other model abilities, such as reasoning, over translation capability.

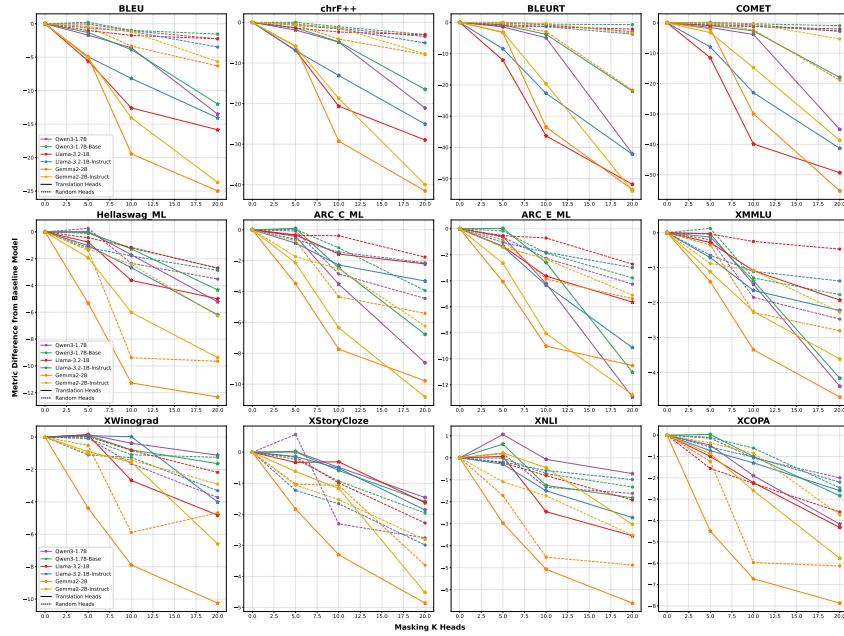
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223

Figure 19: Performance change across downstream benchmarks for small size group models when masking token alignment heads versus random heads.

1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

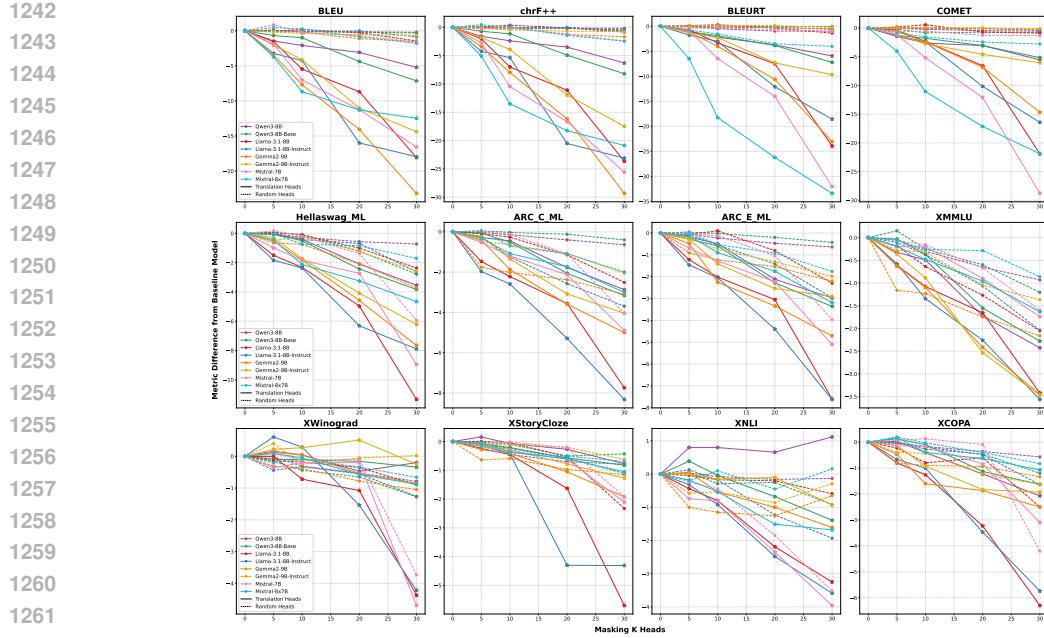


Figure 20: Performance change across downstream benchmarks for medium size group models when masking token alignment heads versus random heads.

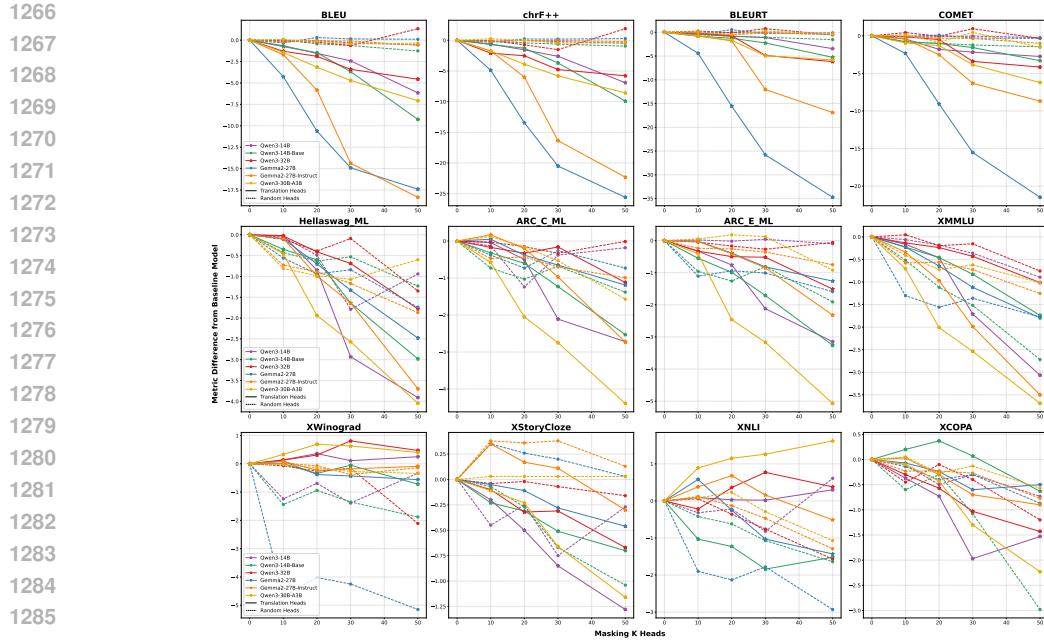


Figure 21: Performance change across downstream benchmarks for large size group models when masking token alignment heads versus random heads.

A.7 THE ROLE OF TOKEN ALIGNMENT HEADS ON MULTILINGUAL TASKS

In this section, we showcase the utility of Token Alignment Heads on multilingual tasks through a series of case studies. To elucidate their performance characteristics under varying conditions, we analyze two distinct scenarios: the Hellaswag_ML task, which is moderately dependent on Token Alignment Heads, and the XNLI/XWinograd tasks, where the dependency is substantially weaker.

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

The Role of Token Alignment Heads in Hellaswag_FR

We need to choose the best ending for the context text.

context_text: Ceci est un **tutoriel** sur comment allumer un feu de camp. Il montre le **feu de camp** qui **brûle** sur le sol. il . endings:

- A: zeigt den Schornstein des Hauses.
- B: zeigt, wie man ein paar drumherum legen kann, um etwas Licht zu haben.
- C: beginnt mit den Dingen, die man für den Anfang braucht.
- D: zeigt, wie man das Feuer hält und den Topf ein wenig darauf stellt, um das Propangas aufzufangen.

Key Word For Task: **tutoriel**, **feu de camp**, **brûle**

Qwen3-30B-A3B Attention Heads Translation Score: L40-H14 = 0.26 L16-H28 = 0.22 L26-H2 = 0.126 L28-H14 = 0.124

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

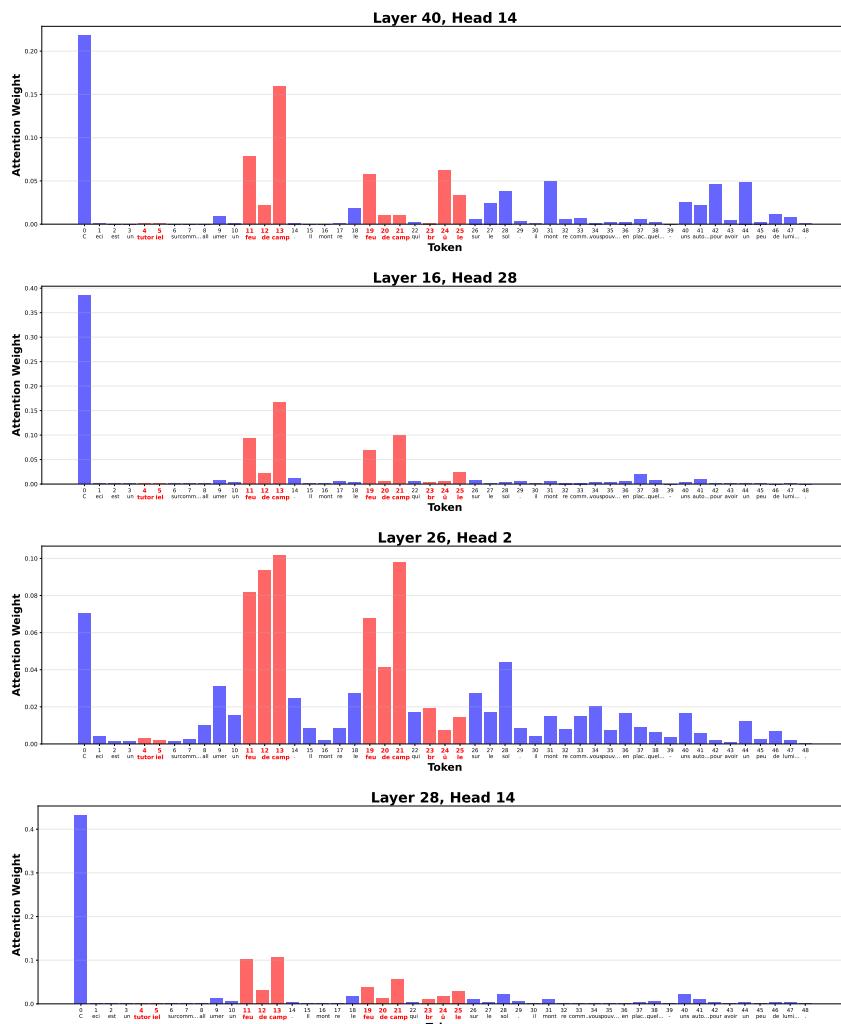


Figure 22: French Hellaswag case. The notation L40-H14 indicates Layer 40, Head 14, where "L" stands for Layer and "H" stands for Head. In this French Hellaswag case, L40-H14, L16-H28, L26-H2 and L28-H14 have translation scores greater than 0.1, identifying them as token alignment heads. We observe that all token alignment heads here can attend to the key tokens.

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

The Role of Token Alignment Heads in Hellaswag DE

We need to choose the best ending for the context text.

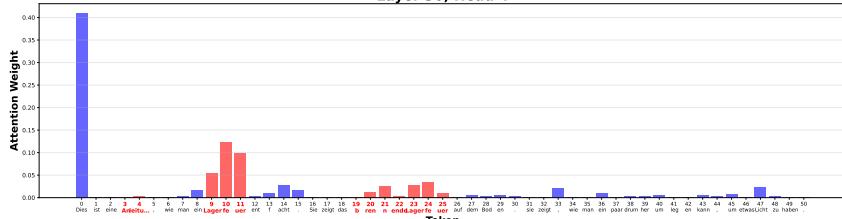
context_text: Dies ist eine Anleitung, wie man ein Lagerfeuer entfacht. Sie zeigt das brennende Lagerfeuer auf dem Boden. sie. endings:

- A: zeigt den Schornstein des Hauses.
- B: zeigt, wie man ein paar drumherum legen kann, um etwas Licht zu haben.
- C: beginnt mit den Dingen, die man für den Anfang braucht.
- D: zeigt, wie man das Feuer hält und den Topf ein wenig darauf stellt, um das Propangas aufzufangen.

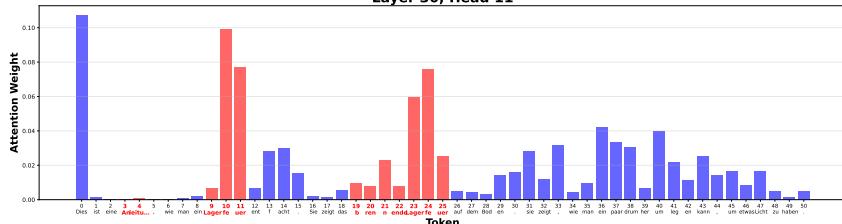
Key Word For Task: **tutoriel, feu de camp, brûle**

Qwen3-30B-A3B Attention Heads Translation Score: L30-H4 = 0.175 L30-H11 = 0.127 L40-H14 = 0.262 L26-H2 = 0.126

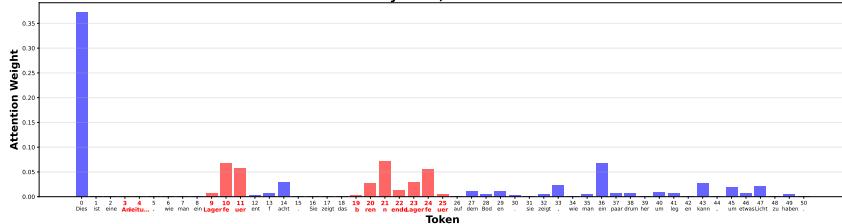
Layer 30, Head 4



Layer 30, Head 11



Layer 40, Head 14



Layer 26, Head 2

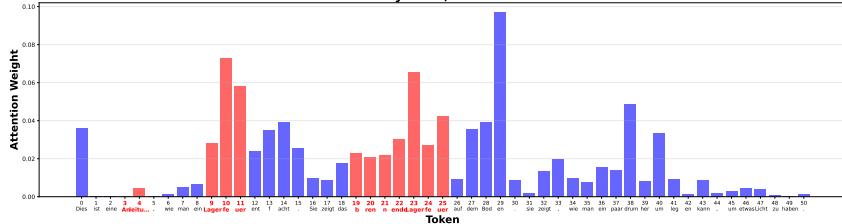
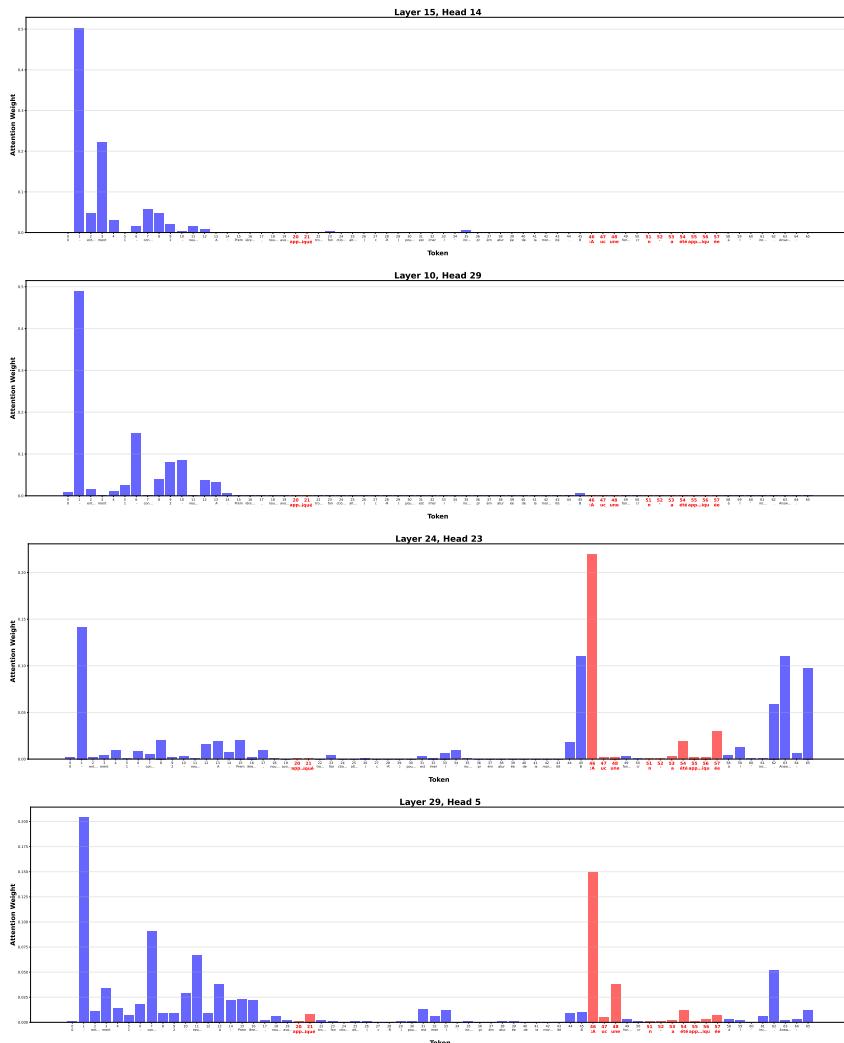


Figure 23: German Hellaswag case. The notation L30-H4 indicates Layer 30, Head 4, where "L" stands for Layer and "H" stands for Head. In this German Hellaswag case, L30-H4, L30-H1, L40-H14 and L26-H2 have translation scores greater than 0.1, identifying them as token alignment heads. We observe that all token alignment heads here can attend to the key tokens.

1404
1405
1406
14071408 **The Role of Token Alignment Heads in XNLI_FR**1409 We need to decide whether A and B are entailment, contradiction or neutral.
1410 A: Premièrement , nous avons **appliqu ** trois fonctions alternatives (c-R) pour estimer l' incidence pr  matur  e de la mortalit   .
1411 B: **Aucune** fonction cr n' a   t   **appliqu **    l' incidence .1412 Key Word For Task: **appliqu **, **Aucune** , **appliqu **

1413 Qwen-3-30B-A3B Attention Heads Translation Score: L15-H14 = 0.155 L10-H29 = 0.127 L24-H23 = 0 L29-H5 = 0

1414

1448 Figure 24: French XNLI case. The notation L15-H14 indicates Layer 15, Head 14, where "L" stands
1449 for Layer and "H" stands for Head. In this French XNLI case, Both L15-H14 and L10-H29 have
1450 translation scores greater than 0.1, identifying them as token alignment heads, whereas L24-H23
1451 and L29-H5 are not. We observe that the token alignment heads tend to have attention weights
1452 close to zero on key tokens, as seen with L15-H14 and L10-H29 in the figure. In contrast, heads
1453 with relatively high attention weights on key tokens, such as L24-H23 and L29-H5, have very low
1454 translation scores.1455
1456
1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

The Role of Token Alignment Heads in XNLI_DE

We need to decide whether A and B are entailment, contradiction or neutral.

A: Sieh mal , das ist nicht so **schlamm** ein paar **Stunden** .

B: Ich **mag es nicht** , zu warten , nicht mal für 10 **Minuten** .

Key Word For Task: **schlamm, Stunden, mag es nicht, Minuten**

Qwen3-30B-A3B Attention Heads Translation Score: L13-H17 = 0.115 L10-H29 = 0.127 L0-H29 = 0.01 L1-H13 = 0.005

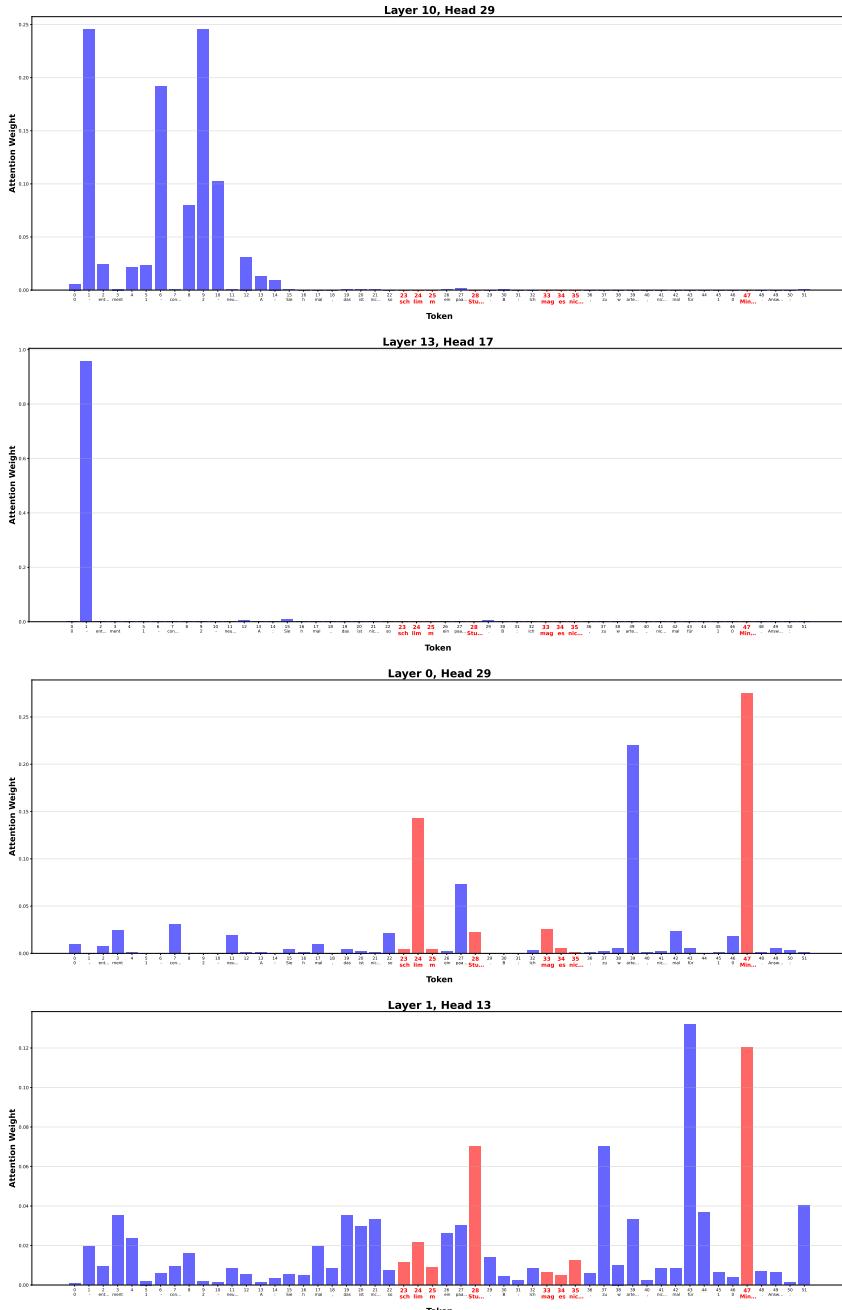


Figure 25: German XNLI case. The notation L13-H17 indicates Layer 13, Head 17, where "L" stands for Layer and "H" stands for Head. In this German XNLI case, Both L13-H17 and L10-H29 have translation scores greater than 0.1, identifying them as token alignment heads, whereas L0-H29 and L1-H13 are not. We observe that the token alignment heads tend to have attention weights close to zero on key tokens, as seen with L13-H17 and L10-H29 in the figure. In contrast, heads with relatively high attention weights on key tokens, such as L0-H29 and L1-H13, have very low translation scores.

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

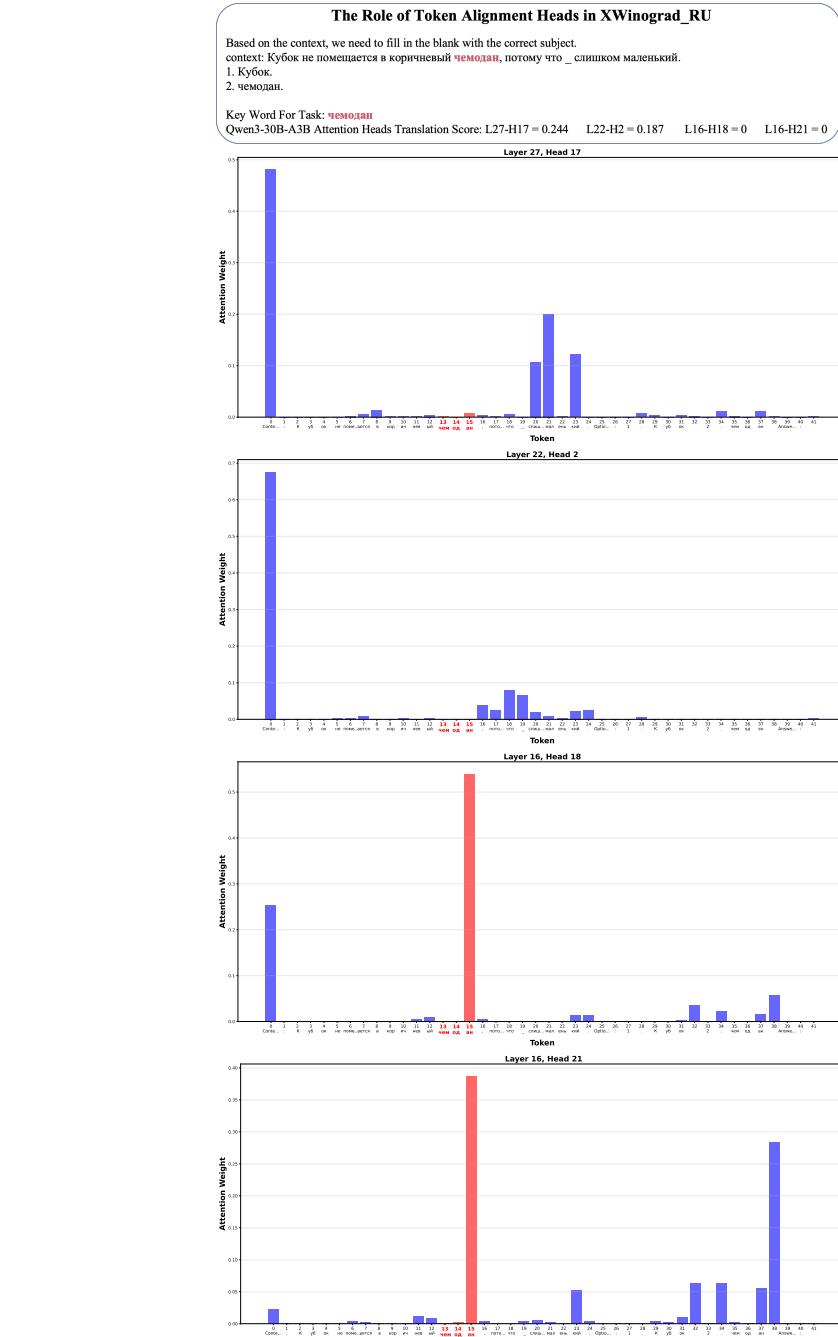


Figure 26: Russian XWinograd case. The notation L27-H17 indicates Layer 27, Head 17, where "L" stands for Layer and "H" stands for Head. In this Russian XWinograd case, Both L27-H17 and L22-H2 have translation scores greater than 0.1, identifying them as token alignment heads, whereas L16-H18 and L16-H21 are not. We observe that the token alignment heads tend to have attention weights close to zero on key tokens, as seen with L27-H17 and L22-H2 in the figure. In contrast, heads with relatively high attention weights on key tokens, such as L16-H18 and L16-H21, have very low translation scores.

1564

1565

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585

1586

1587

1588

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

The Role of Token Alignment Heads in XWinograd_FR

Based on the context, we need to fill in the blank with the correct subject.

context: Claire a frappé à la porte de **Sylvie**, mais _ n'a pas répondu.

1. Claire.

2. Sylvie.

Key Word For Task: **Sylvie**

Qwen3-30B-A3B Attention Heads Translation Score: L13-H18 = 0.24 L22-H2 = 0.187 L45-H29 = 0

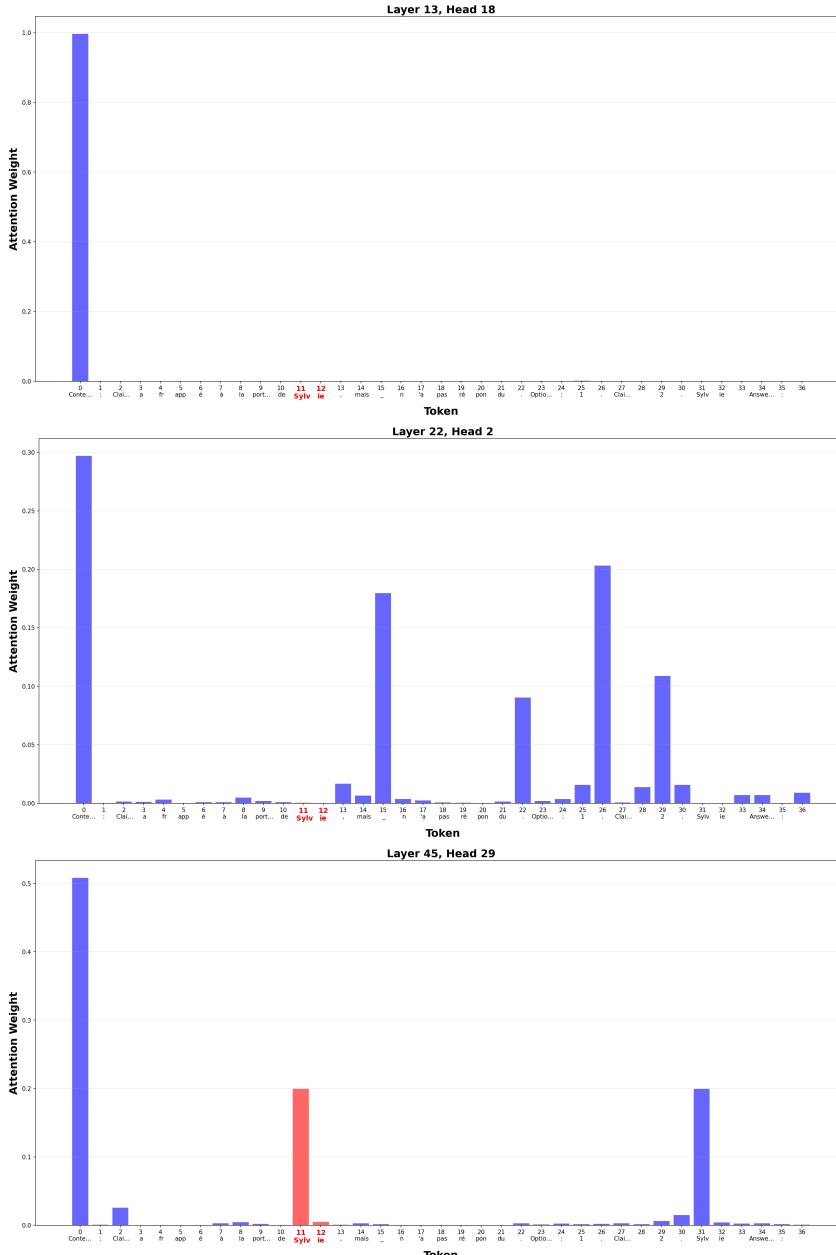


Figure 27: French XWinograd case. The notation L13-H18 indicates Layer 13, Head 18, where "L" stands for Layer and "H" stands for Head. In this French XWinograd case, Both L13-H18 and L22-H2 have translation scores greater than 0.1, identifying them as token alignment heads, whereas L45-H29 is not. We observe that the token alignment heads tend to have attention weights close to zero on key tokens, as seen with L13-H18 and L22-H2 in the figure. In contrast, heads with relatively high attention weights on key tokens, such as L40-H14, have very low translation scores.

1620
1621

A.8 LANGUAGE CONSISTENCY ACROSS 20 LANGUAGE PAIRS

1622
1623
1624

1625

1626

1627

1628

1629

1630

1631

1632

1633

1634

1635

1636

1637

1638

1639

1640

1641

1642

1643

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

1654

1655

1656

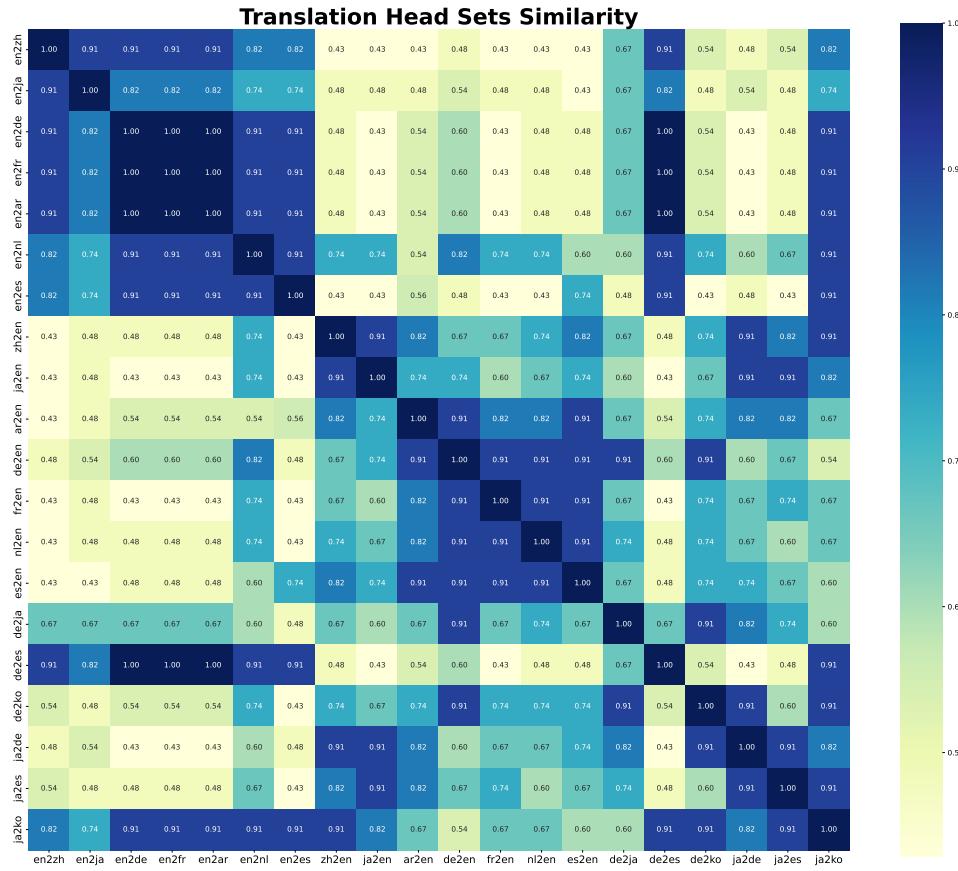


Figure 28: Jaccard similarity matrix of token alignment head sets across 20 language pairs. Here, we broadly categorize the selected language pairs into four groups: en2others (en2zh, en2ja, en2de, en2fr, en2ar, en2nl, en2es), others2en (zh2en, ja2en, ar2en, de2en, fr2en, nl2en, es2en), de2ja/es/ko, and ja2de/es/ko. It is observable that the similarity among the en2others pairs is high, whereas the others2en group and the other language pairs generally exhibit lower similarity.

1661

1662

1663

1664

1665

1666

1667

1668

1669

1670

1671

1672

1673