

# RomanLens: The Role of Latent Romanization in Multilinguality in LLMs

Alan Saji<sup>1†‡</sup>, Jaavid Aktar Husain<sup>2†</sup>, Thanmay Jayakumar<sup>1,3</sup>, Raj Dabre<sup>1,3,4,5§</sup>  
Anoop Kunchukuttan<sup>1,6</sup>, Ratish Puduppully<sup>7‡</sup>

<sup>1</sup>Nilekani Centre at AI4Bharat, <sup>2</sup>Singapore University of Technology and Design,

<sup>3</sup>Indian Institute of Technology Madras, India,

<sup>4</sup>National Institute of Information and Communications Technology, Kyoto, Japan,

<sup>5</sup>Indian Institute of Technology Bombay, India, <sup>6</sup>Microsoft, India, <sup>7</sup>IT University of Copenhagen

## Abstract

Large Language Models (LLMs) exhibit strong multilingual performance despite being predominantly trained on English-centric corpora. This raises a fundamental question: *How do LLMs achieve such multilingual capabilities?* Focusing on languages written in non-Roman scripts, we investigate the role of Romanization—the representation of non-Roman scripts using Roman characters—as a potential bridge in multilingual processing. Using mechanistic interpretability techniques, we analyze next-token generation and find that intermediate layers frequently represent target words in Romanized form before transitioning to native script, a phenomenon we term *Latent Romanization*. Further, through activation patching experiments, we demonstrate that LLMs encode semantic concepts similarly across native and Romanized scripts, suggesting a shared underlying representation. Additionally, for translation into non-Roman script languages, our findings reveal that when the target language is in Romanized form, its representations emerge earlier in the model’s layers compared to native script. These insights contribute to a deeper understanding of multilingual representation in LLMs and highlight the implicit role of Romanization in facilitating language transfer. Code and data are available at <https://github.com/AI4Bharat/Romanlens>.

## 1 Introduction

The majority of modern Large Language Models (LLMs) (Touvron et al., 2023; Dubey et al., 2024; Team et al., 2024) are trained predominantly on English-dominated corpora. Nonetheless, they exhibit strong multilingual performance across diverse languages (Shi et al., 2023; Huang et al., 2023; Zhao et al., 2024a; Zhang et al., 2023). This raises a fundamental question: *How do LLMs develop such robust multilingual capabilities despite their English-centric training?*

To address this, prior work by Wendler et al. (2024) suggests that LLMs encode multilingual information within a shared, language-agnostic latent space, albeit with an inherent bias toward English due to training data composition and architectural choices. Building on this perspective, we investigate a complementary mechanism that may underlie multilingual processing, particularly for languages written in non-Roman scripts.

We hypothesize that LLMs leverage romanized forms of non-Roman script languages as an intermediate bridge between their language-agnostic concept space and language-specific output representations. Romanization—the representation of non-Roman scripts using Roman characters—may facilitate this process by aligning non-English languages more closely with English. Supporting this, Jaavid et al. (2024) demonstrated that explicitly romanizing inputs improves model performance on multilingual tasks, suggesting an inherent alignment between romanized text and English representations. We investigate whether LLMs indeed use romanization as a bridge between language-agnostic concepts and language-specific outputs given its potential implications for understanding multilingual processing in LLMs.

Our primary experiment visualizes next-token generation using the logit lens (Nostalgebraist, 2020), applying the language modeling head to intermediate layers. As illustrated in Figure 1, we prompt the LLaMA-2 7B (Touvron et al., 2023) model with “Francais: porte - हिन्दी:” to translate “door” from French to Hindi. Our results show that in the middle-to-top layers (layers 20–29), romanized Hindi subwords intermittently appear before transitioning to native script, suggesting an internal representation of romanized text as an intermediary.

<sup>†</sup>Work done during employment at AI4Bharat.

<sup>§</sup>Work done during employment at NICT, Japan.

<sup>‡</sup>Correspondence: Alan Saji ([alansaji2001@gmail.com](mailto:alansaji2001@gmail.com)), Ratish Puduppully ([rapu@itu.dk](mailto:rapu@itu.dk))

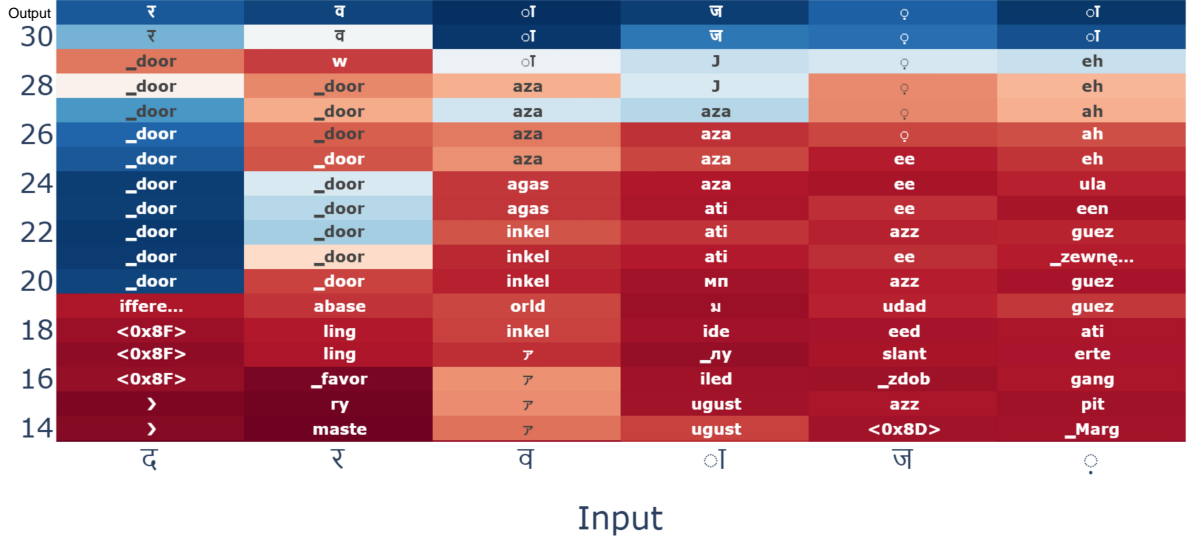


Figure 1: **Logit lens visualization** of Llama-2 7B model translating ‘door’ from French to Hindi. We visualize the output (दरवाजा, ‘Darwaza’ is the romanized form ) taking shape using logit lens producing a next-token distribution for each position (x-axis) and layers 14 and above (y-axis). Interestingly in the middle to top layers (20 - 29) we could observe romanized subwords of the Hindi word (दा - w; अजा - aza ; ज - azz; ज - j) and dependent vowels ( ा -a,eh ; ि - i) before they are represented in their native script. Color represents entropy of next-token generation from low (blue) to high (red). Plotting tool: (Belrose et al., 2023).

Additionally, these romanized representations increase in prominence across timesteps as the target word is generated.

To further probe this phenomenon, we employ activation patching (Ghandeharioun et al., 2024; Variengien and Winsor, 2024; Chen et al., 2024; Dumas et al., 2024), a technique that replaces activations from one forward pass with another to analyze the resulting outputs (c.f. Section 3.3). Dumas et al. (2024) found that LLMs process language and conceptual information as distinct entities. Building on this, we perform layerwise activation patching between romanized and native-script inputs to examine whether LLMs encode conceptual information similarly across scripts.

Based on our experiments, we list down a summary of our contributions below:

- 1. Latent Romanization Across Layers:** During multilingual next-token generation, intermediate layers occasionally represent tokens in Romanized form before resolving to native script (Figure 1). We term this phenomenon *Latent Romanization*.
- 2. Consistent Semantic Encoding Across Scripts:** Activation patching experiments reveal that LLMs encode semantic concepts similarly, regardless of whether the input is in native or Romanized script.
- 3. Earlier Emergence of Target Representations:** When translating into Romanized versus native script, Romanized target representations emerge




Figure 2: **Translation comparison: Romanized vs. Native script.** Next-token generation is visualized using the logit lens for the Gemma-2 9B model translating “flower” from French to Malayalam in romanized (left) and native script (right). The x-axis shows next-token distributions; the y-axis covers layers 30 and above. Target language representations (e.g., “push”, “poo”) appear 1–2 layers earlier in romanized outputs compared to native script outputs. (പൂ, ‘pa’ is its romanized representation). *Push* is a prefix of *pushpam*, the romanized form of the translation of “flower”; also, *poo* is another romanized translation of “flower”. Color represents entropy of next-token generation from low (blue) to high (red). Plotting tool: (Belrose et al., 2023).

earlier in the model’s layers—typically one or two layers prior to native script representations (c.f. Figure 2).

## 2 Related Work

Recent studies have explored various aspects of LLMs’ multilingual behavior: examining whether

English emerges as a latent language in English-centric LLMs (Wendler et al., 2024), how the composition of training corpus mixtures influences latent representations (Zhong et al., 2024) and how LLMs handle multilingual capabilities (Zhao et al., 2024b). Kojima et al. (2024) describe distinct phases in multilingual information processing: initial layers map language-specific lexical and syntactic representations to a language-independent semantic space, middle layers maintain this semantic abstraction, and final layers transform these representations into language-specific lexical and syntactic forms. Interpretability tools relevant to this work include logit lens (Nostalgebraist, 2020), tuned lens (Belrose et al., 2023) and direct logit attribution (Elhage et al., 2021) which are key tools for decoding intermediate token representations in transformer models. The logit lens applies the language modeling head to earlier layers without additional training, while the tuned lens in addition to this trains an affine mapping to align intermediate states with final token predictions. Direct logit attribution attributes logits to individual attention heads. This work focuses on the logit lens (Section 4.1) to investigate whether English-centric decoder only LLMs when prompted in a non-Roman script language, processes via romanized latent states before producing native language text. Tuned lens is avoided as its training process might obscure the intermediate romanized states by aligning them to final native script outputs, potentially masking the phenomenon under investigation.

Activation patching (Meng et al., 2022) is a key interpretability technique employed in our study. This technique has been used to draw causal interpretations of LLMs representations (Variengien and Winsor, 2024; Geiger et al., 2022; Kramár et al., 2024; Ghandeharioun et al., 2024; Chen et al., 2024). Building on these approaches, we adopt an activation patching-based experimental framework to investigate and compare how concepts are encoded in romanized versus native scripts.

Previous studies have demonstrated that romanization can serve as an effective approach to interact with LLMs (Jaavid et al., 2024). Liu et al. (2024) and Xhelili et al. (2024) employ an approach based on contrastive learning for post-training alignment, contrasting sentences with their transliterations in Roman script to overcome the script barrier and enhance cross-lingual transfer.

However, our work distinguishes itself from prior research by exploring the presence of roman-

ized representations in the latent layers of an LLM during multilingual tasks, an aspect that, to the best of our knowledge, has not yet been investigated.

### 3 Background

We give a quick background of the transformer’s forward pass, romanization and the basics of mechanistic interpretability approaches such as logit lens and activation patching which we leverage in this paper.

**Transformer’s Forward Pass** Decoder-only transformer models (Vaswani, 2017) employ a residual architecture to process input sequences through multiple layers, producing a sequence of hidden states (latents). These latents, whose dimensionality remains the same, are updated iteratively across layers through transformer blocks  $f_j$ , where  $j \in [0, k]$  indicates the layer index and  $k$  is the final layer index. For next-token prediction, the final latent  $h_i^{(k)}$  is transformed by an unembedding matrix  $U \in \mathbb{R}^{v \times d}$  to produce logit scores for vocabulary tokens which are then converted to probabilities via the softmax function (c.f. Appendix A).

#### 3.1 Romanization

Transliteration is the conversion of text written in one script to another. Romanization is a subcategory of transliteration where the target script is English/Latin. Within romanization there are multiple romanization schemes available, each based on different considerations. One key aspect of romanization schema is if it is lossy or lossless. A lossless scheme is required in cases where we have to convert the output back to native script. Typically, deterministic transliterations are lossless, whereas natural transliterations are lossy.

**Example:** The Hindi word for “flower” in Devanagari and its romanization:

- *Devanagari* (native script): फूल
- *Romanization*: phool

#### 3.2 Interpretability Tool: Logit lens

Generally, in a decoder only LLM, the unembedding matrix  $U$  is multiplied with the final hidden state and a softmax is taken on the product to produce the token distributions at that token generation step. Since all hidden states of an LLM are in the same shape, it is possible to apply the unembedding

matrix and softmax on all layers, thereby generating token distributions at all layers. This method of prematurely decoding hidden states is referred to as *logit lens* (Nostalgebraist, 2020). Logit lens reveals how the latent representations evolve across layers to produce the final output, providing insights into the progression of computations within the model.

### 3.3 Interpretability Tool: Activation Patching

Activation patching involves modifying or patching the activations at specific layers during a forward pass and observing the effects on the model’s output. In this work, we adopt the activation patching setup introduced in Dumas et al. (2024).

In the context of activation patching, let  $\ell$  denote the language of a word,  $C$  denote the concept of a word and  $w(C^\ell)$  denote that word. For example, if  $C = \text{cow}$  and  $\ell = \text{'en'}$ , then  $w(C^{\text{en}}) = \text{'cow'}$ . Similarly  $w(C^{\text{fr}}) = \text{'vache'}$ . We use 5-shot translation prompts to create paired source  $S = (C_S, \ell_S^{\text{in}}, \ell_S^{\text{out}})$  and target prompt  $T = (C_T, \ell_T^{\text{in}}, \ell_T^{\text{out}})$ , with different concept, input language, and output language. Unless otherwise specified,  $\ell_S$  and  $\ell_T$  refer to the output languages of  $S$  and  $T$ , respectively.

For each transformer block  $f_j$ , we create two parallel forward passes: one processing the source prompt  $S = (s_1, \dots, s_{n_s}, \dots, s_{n_S})$  and the other processing the target prompt  $T = (t_1, \dots, t_{n_t}, \dots, t_{n_T})$ . It should be noted that  $n_s, n_t$  represents the position of the last token of the object to be translated whereas  $n_S, n_T$  represent the last token position of the source and target prompt to be translated. In Figure 3 in the target prompt translating *sun* from French to Hindi,  $n_t$  would be the position of the subword “*eil*” highlighted in red, whereas  $n_T$  would be the position of “*दी*” the last subword of the prompt. Similarly, for the source prompt in Figure 3 translating *elephant* from German to Italian both  $n_s$  and  $n_S$  would be the position of “*ant*” highlighted in red. After creating the parallel forward passes, we extract the residual stream of the last token of the word to be translated at  $n_s$  after layer  $j$ , denoted as  $h_{n_s}^{(j)}(S)$  and all subsequent layers, and insert it at the corresponding layer at the corresponding position  $n_t$  in the forward pass of the target prompt, i.e. by setting  $h_{n_t}^{(j)}(T) = h_{n_s}^{(j)}(S), h_{n_t}^{(j+1)}(T) = h_{n_s}^{(j+1)}(S), \dots, h_{n_t}^{(k)}(T) = h_{n_s}^{(k)}(S)$ . We then complete the altered forward pass and analyze the next token distribution to evaluate source concept  $C_S$  encoded in the target language. An illustration of

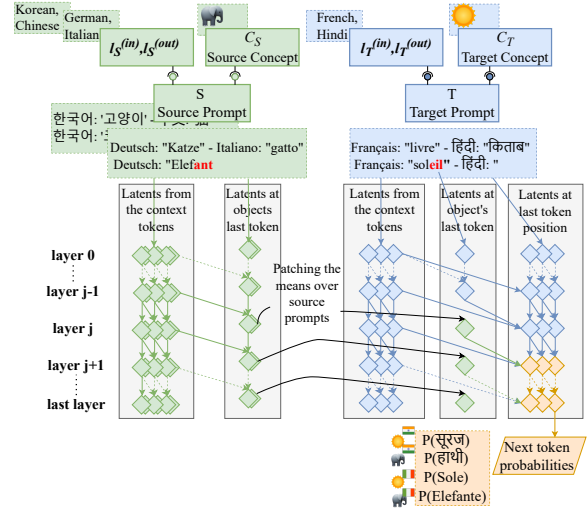


Figure 3: **Activation patching illustration.** For two given concepts, say, *elephant* and *sun*, we generate multiple source prompts which translate *elephant*, and a target prompt for translating *sun* from French to Hindi. We then extract the residual stream associated with the final token of the word to be translated after a specific layer  $j$  and all subsequent layers from the source prompts. The mean residuals at each layer are computed and inserted into the corresponding positions during the forward pass of the target prompt. The resulting next token probabilities will be dominated by the source concept in target language (ELEPHANT<sup>HI</sup>, i.e., हाथी) when patching at layers 0–15, and by the target concept in target language (SUN<sup>HI</sup>, i.e., सूरज) for layers 16–31. Adapted from Dumas et al. (2024).

this setup is shown in Figure 3.

## 4 Methodology

We design our analysis setup with the intention of addressing the following research questions:

**RQ1:** Do LLMs exhibit latent romanization during multilingual text completion tasks? (Section 4.1)

**RQ2:** How does the representation of semantic concepts in LLMs compare between native and romanized scripts of non-Roman script languages? (Section 4.2)

**RQ3:** What are the differences in hidden layer representations when processing the same language in romanized and native scripts? (Section 4.3)

**Prompt design.** We design prompts that facilitate next-token ( $x_{n+1}$ ) prediction from the given context ( $x_1, \dots, x_n$ ). This is adopted across all analysis setups. The prompts are designed around translation, repetition, and cloze tasks, as described below.

**Translation task.** We prompt the model to translate



a word given five in-context examples.

*Repetition task.* We prompt the model to repeat a word in the same language given five in-context examples.

*Cloze task.* We prompt the model to predict the masked word in a sentence given two in-context examples.

These tasks cover a range of multilingual text completion setups. Among these, the repetition task is more syntactic in nature compared to translation and cloze tasks. Appendix B provides a Hindi example of prompts across each of the three tasks, along with their English translations and romanized forms.

#### 4.1 Latent Romanization Analysis

Translation, repetition, and cloze tasks are explored by providing the respective prompts as inputs to an LLM to generate the corresponding output word. We romanize the output word, tokenize it, retaining only the tokens present in the model’s vocabulary, and analyze the occurrence of these tokens in the latent layers across timesteps of the output word generation. The analysis is done using logit lens by examining whether the probability of a romanized token in the next token distribution at a given layer exceeds 0.1. We refer to this hereafter as the *latent romanization condition*. The 0.1 threshold is empirically determined to optimize detection accuracy, i.e. minimizing false positives and maximizing true positives (compared to alternative thresholds 0.05 and 0.01). Our analysis focuses on the final 10 layers of an LLM, where coherent romanized representations emerge according to logit lens visualizations (c.f. Figures 1 and 8 - 12).

We track romanized tokens using a timestep-specific tokenization scheme optimized for detection accuracy. In the first output generation timestep, we check for tokens that include the full romanized word and its prefixes. During intermediate timesteps, we check for all possible substrings of the romanized word in the latent layers. In the final output generation timestep, we probe the presence of only the full romanized word and its suffixes as potential tokens (c.f. Appendix C).

Latent romanization is analyzed under three distinct scenarios:

- (a) **Constrained Word Generation:** Using standard prompts and the target word, we guide the model to generate the complete target

word. At each layer, we track how often romanized tokens emerge during the decoding process. We do this by checking each generation timestep for *latent romanization condition*. The ‘*latent fraction*’ for a layer represents how frequently these romanized tokens appear across timesteps, averaged across all samples (c.f. Appendix D).

- (b) **First Subword Only:** We prompt for only the initial subword and compute the *latent fraction*. Despite having a single timestep, we maintain the *latent fraction* terminology for consistency.
- (c) **Last Subword Only:** We augment the standard prompt with all but the final subword of the target, then analyze the generation of the final subword.

We document layerwise *latent fraction* separately for first and last subword generation of the output. Intuitively, there is a distinction between the first and last token generation steps for a given word. In the former, the model faces a greater decision-making burden, while in the latter, the model is typically more confident in its predictions. We hypothesize that, in the latter scenario, the model may reach a decision in the layers just below the final few layers and express the output in a romanized form, as language-specific neurons, responsible for native script processing, are concentrated primarily in the last few layers (Tang et al., 2024). This could lead to romanized tokens appearing more frequently as the model progresses from the first subword to the last.

#### 4.2 Patching With Romanized Representation Versus Native Representation

We intend to compare how concepts are encoded in native script versus romanized script using translation task. In order to do this, we patch representations from source prompt where input language is romanized and compare this with patching from source prompt where input language is in native script. We first perform patching using a single-source prompt. Then, we repeat the process using an averaged multi-source prompt, contrasting multiple romanized source input languages with multiple native script source input languages. Single-source patching might be influenced by language or script-specific characteristics. In contrast, multi-source

patching reduces such biases, leading to more robust and generalizable findings.

In all scenarios from the resulting next token distribution, we compute the probabilities  $P(C_S^{\ell_T})$  i.e. probability of source concept in target language, and  $P(C_T^{\ell_T})$  i.e. probability of target concept in target language. We track  $P(C^\ell)$ , i.e., the probability of the concept  $C$  occurring in language  $\ell$ , by simply summing up the probabilities of all prefixes of  $w(C^\ell)$  and its synonyms in the next-token distribution (c.f. Appendix G). We analyze  $P(C_S^{\ell_T})$  to evaluate how effectively a concept is encoded in a given source input language  $\ell_S^{\text{in}}$ .

### 4.3 Comparing Translations Into Romanized vs. Native Script

This analysis examines translation task with target languages in native script and their romanized equivalents. We focus on first-token generation of the output word, also considering possible synonyms.

In the next token generation step, the probability of target language and latent language (English) (Wendler et al., 2024) at each layer is examined using logit lens. Each probability is computed by summing over probabilities of all possible tokens corresponding to the answer word(s) in that respective language (c.f. Appendix E). Tokens of latent language and target language are derived using tokenization scheme for first token generation timestep mentioned in Section 4.1.

## 5 Experimental Settings

**Languages:** We focus on five Indic languages: Hindi, Gujarati, Tamil, Telugu, and Malayalam, as well as Chinese and Georgian. Among these, Hindi and Gujarati belong to the Indo-Aryan branch of the Indo-European language family and use scripts derived from the Devanagari and Gujarati scripts, respectively. Tamil, Telugu, and Malayalam, on the other hand, are part of the Dravidian language family and each has its own distinct script. Chinese belongs to the Sino-Tibetan language family and is written using logographic characters. Georgian is part of the Kartvelian language family and uses the unique Georgian script. To examine the generality of latent romanization, we perform qualitative analyses on five additional languages that use different writing systems: Greek, Ukrainian, Amharic, Hebrew, and Arabic (c.f. Appendix H).

**Language Models:** In this study, we focus mainly on Gemma-2 9B, Gemma-2 9B instruction-tuned (Team et al., 2024), Llama-2 7B, Llama-2 13B (Touvron et al., 2023) and Mistral-7B (Jiang et al., 2023) (c.f. Appendix K) language models, some of the best performing open weights English-centric LLMs. Although the training data for these models are primarily English, these models have high multilingual capabilities (Huang et al., 2023; Zhao et al., 2024a; Zhang et al., 2023).

**Romanization:** We have used the IndicXlit scheme (Madhani et al., 2023) (c.f. Appendix F) for Indic languages, *pypinyin* (Mozillazg, 2024) for Chinese and *Unidecode* (Šolc, 2025) for Georgian to romanize native scripts.

**Data For Logit Lens Experiments:** We use a curated word-level dataset with synonyms translated from recent works in this field (Wendler et al., 2024) using the Llama 3.3 70B model (Dubey et al., 2024). The quality of translations were manually verified to ensure accuracy and relevance. The datasets are kept simple to facilitate the observation of how latents evolve at each token level.

**Data For Activation Patching Experiments:** We adopt the dataset used in recent studies (Dumas et al., 2024), extending it to include both native script and romanized versions of the languages considered in this study. Translations are performed using the Llama 3.3 70B model (Dubey et al., 2024) and the translations are romanized using IndicXlit (Madhani et al., 2023), *pypinyin* (Mozillazg, 2024) and *Unidecode* (Šolc, 2025). All translations were manually validated to ensure data quality.

## 6 Results

### 6.1 Latent romanization

Our analysis demonstrates that **LLMs do exhibit latent romanization during text completion tasks in six out of seven quantitatively analyzed languages** (c.f. Figure 4a), with Chinese being the only exception where this phenomenon is not observed. In Figure 4a we can see the *latent fraction* of romanized tokens for the last 10 layers of an LLM. This is across all tokens of the output word. The frequency of romanized tokens tends to increase just before the last layers. Qualitative logit lens analysis done for Greek, Ukrainian, Arabic, Amharic and Hebrew reveals similar patterns (c.f. Appendix H).




Figure 4: **Distribution of Romanized Tokens Across Model Layers: Analysis of First, Last, and All Generation Timesteps.** This distribution is plotted across the last 10 layers of Gemma-2 9b IT model for translation task with English as source language and is averaged across 100+ samples. X-axes represents layer index, y-axes represents latent fraction i.e. the fraction of timesteps where romanized tokens occur with a probability  $> 0.1$  averaged over samples for a specific layer. We plot the distributions for Gujarati (gu), Tamil (ta), Telugu (te), Hindi (hi), Malayalam (ml), Georgian (ka) and Chinese (zh).




Figure 5: **Frequency distribution of romanized tokens across translation, repetition and cloze task.** We check if romanized tokens occur with a probability  $> 0.1$  in the last 10 layers of an LLM and compute frequency of this occurrence across 100+ samples. Gemma 2 9B IT is the model used and English is the source language for translation task.

In Figures 4b and 4c it is observed that the range of the *latent fraction* of romanized tokens varies from 0 – 0.01 in the first token generation step to 0 – 0.2 in the last token generation step in most languages. This trend indicates that latent romanization increases progressively from the initial token to the final token of the output across languages. This observation supports our hypothesis that the first token generation involves more intricate decision-making processes compared to the generation of the final token within an output word.

Based on the above information, we quantify romanization across tasks for the last token generation step. Here the criteria would be if a romanized token occur at next token generation step with a probability  $> 0.1$  in the last 10 layers, it is counted as a positive. As depicted in Figure 5, across translation, repetition and cloze task we observe a significant occurrence of romanized tokens in the latent layers.

Among the three text completion tasks consid-

ered in the latent romanization experiment (Section 4.1), a relatively high latent romanization frequency is observed for the repetition task (Figure 5). The translation and cloze tasks function at a semantic level, whereas the repetition task is purely syntactic. This means that the repetition task being less complex, the model may reach a decision of what to predict sooner, potentially in earlier layers and might express its prediction in romanized form in intermediate layers. This behavior could be attributed to language-specific neurons, responsible for native script processing, being predominantly concentrated in the initial and final few layers of LLMs, (Tang et al., 2024) leaving the intermediate layers without them.

## 6.2 Patching With Romanized Representation vs. Native Representation

In Figure 6, we analyze two patching scenarios: In the single-source setup, we compare patching from Hindi→Italian with Hindi(romanized)→Italian source prompt, to Malayalam→Italian target prompt. In the multi-source setup, we contrast patching from multiple native script prompts (Hindi→Italian, Gujarati→Italian ...) against their romanized counterparts (Hindi(romanized)→Italian, Gujarati(romanized)→Italian ...). We compare the probability distributions of source concept in target language  $P(C_S^{\ell_T})$  across adjacent graphs where native source input language is contrasted with romanized source input language. It is evident that these probability distributions show remarkable similarity whether the source input language is in romanized or native script, consistent across both single-source and multi-source prompt setups. The similarity is quantitatively supported by the KL divergence measurements between adjacent

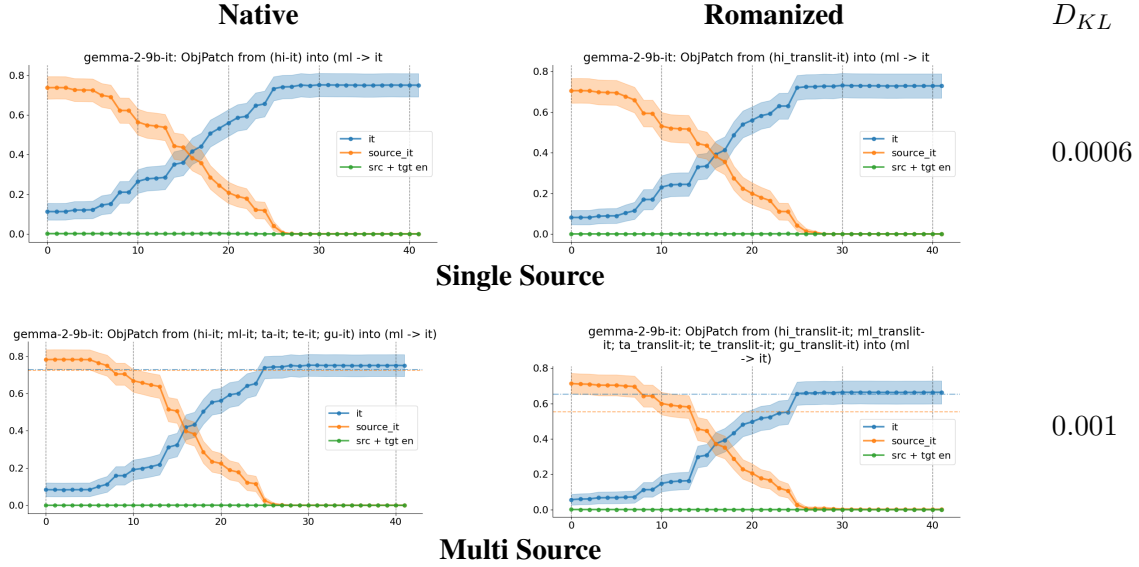


Figure 6: **Comparative Analysis of Patching from Source Prompts: Native Script vs. Romanized Script Inputs.** Concept probabilities across layers for different prompt setups are plotted in each graph. The x-axis represents the patching layer, while the y-axis indicates the probability of correctly predicting the concept in language  $\ell$ . Curves: blue (target concept in Italian), orange (source concept in Italian), and green (source or target concept in English). Results are reported as means with 95% Gaussian confidence intervals, calculated over a dataset of 200 samples. The orange curve is compared across adjacent graphs and KL divergence  $D_{KL}$  quantifies this. Languages involved: Hindi (hi), Tamil (ta), Telugu (te), Malayalam (ml), Gujarati (gu) and Italian (it). Model: Gemma 2 9b it.

graphs, remaining below 0.01 in both setups. KL divergence value close to zero indicates that the two distributions are nearly identical.

This analysis reveals that **LLMs encode semantic concepts similarly regardless of whether the input is in native or romanized script**. Furthermore, this finding demonstrates that the model achieves comparable levels of language understanding when processing non-Roman script languages in their romanized form as in their native script.

### 6.3 Comparing Translations Into Romanized vs. Native Script

We quantify the observations from Figure 2 by analyzing next-token predictions across layers using logit lens. In Figure 7, panels 7c and 7d illustrate that for translations into native scripts, target language tokens begin to emerge from layer 40 onward. Conversely, in panels 7a and 7b, where the target language is in romanized script, target tokens appear 1–2 layers earlier. This pattern indicates that **when processing non-Roman script languages, the model forms internal representations of target tokens in earlier layers for romanized script compared to native script**. This trend is consistent across language pairs and models (c.f. Figures 16–18 in the Appendix). This suggests that

romanization facilitates faster progression toward language-specific embeddings.

In Figure 7, in all four graphs consistent with Wendler et al. (2024), English representations emerge from the middle layers and persist until the final few layers, where the target language representations gradually take shape. It is important to note that native script curves (Figures 7c, 7d) exhibit steeper gradients than their romanized equivalents (Figures 7a, 7b).

**Discussion.** In our investigation of romanized representations in the latent layers, we conclusively identified romanized tokens in the last 6–7 layers of an LLM across various multilingual text completion tasks. Based on previous works in this field (Wendler et al., 2024; Zhao et al., 2024b), in an English-centric decoder only LLM this region corresponds to the transition from an English-centric language-agnostic concept space to a language-specific space where the idea conceived in the concept space is expressed in the target language. Our findings suggest that **romanization serves as a bridge between the concept space and the language-specific region for non-Roman script languages**, an observation strongly supported by our analysis of six diverse writing systems. Romanization acting as a bridge could explain why



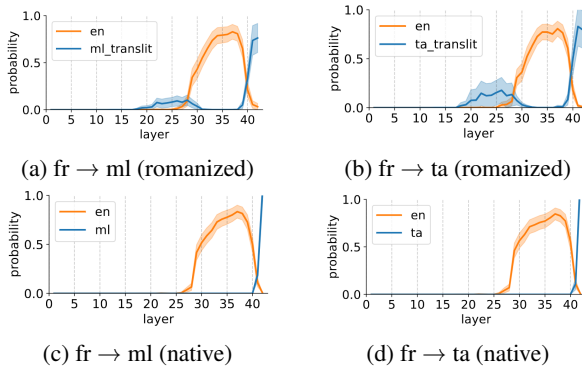


Figure 7: **Language probabilities for latent layers** in translation from French to Malayalam and Tamil in romanized (top row) and native scripts (bottom row) across various samples using Gemma-2 9B IT model. X-axis: layer index; Y-axis: probability of correct next token (via logit lens) in a given language. Error bars: 95% Gaussian confidence intervals. English is the latent language (orange curve). For romanized script, target representations (blue curve) emerge 1-2 layers earlier than native script, appearing before layer 40.

romanization based script barrier breaking methods like Liu et al. (2024) and Xhelili et al. (2024) work. Notably, we do not observe Latent Romanization in Chinese, likely due to its logographic script and relatively high-resource status.

## 7 Conclusion

Our findings show that LLMs implicitly use Romanization as a bridge for non-Roman scripts, exhibiting *Latent Romanization* in intermediate layers before switching to native scripts. Layerwise analyses reveal that semantic concepts are encoded similarly across native and Romanized inputs, indicating a shared internal representation. Moreover, when translating into a Romanized script, target words emerge earlier, highlighting Romanization as a structural link between language-agnostic concepts and language-specific output. While our study reveals initial insights into Latent Romanization, future work could focus on applying these findings to develop training strategies that enhance performance across diverse linguistic communities.

## 8 Limitations

The handling of multilingual text by large language models (LLMs) remains an active area of research. Although evidence suggests that LLMs process English representations within a language-agnostic space, the specific mechanisms by which these models adjust their interactions over different

timesteps during token generation are still not fully understood. In our study, we observe that romanized representations become increasingly prominent in the hidden layers as token generation progresses from the first to the final token. This trend suggests that latent romanization may help the model mitigate differences in token fertility—that is, the average number of tokens required to represent a word—between the output language and its primary latent language, English. This effect appears especially for non-Roman script languages, with high token fertility. However, further research is needed to confirm and generalize these observations.

The interpretability of non-Roman scripts at latent layers is limited when models employ tokenization schemes that split non-Roman characters into multiple bytes, complicating logit lens analysis. Extending this work to models with alternative tokenization methods would offer a more complete understanding of multilingual capabilities and representations.

This work identifies but does not explain the selective occurrence and varying intensity of latent romanization across languages—questions that merit dedicated future investigation.

## 9 Ethics Statement

Through this work, our aim is to democratize access to LLMs and address the issue of limited data availability for low-resource languages. We emphasize that it is not our intention to diminish the value or significance of the native scripts of the languages included in this study.

The code and datasets created in this work will be made available under permissible licenses. Generative AI systems were only used for assistance purely with the language of the paper, e.g., paraphrasing, spell-check, polishing the author’s original content, and for writing boiler-plate code.

## References

- Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*.
- Haozhe Chen, Carl Vondrick, and Chengzhi Mao. 2024. *Selfie: Self-interpretation of large language model embeddings*. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Clément Dumas, Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. Separating tongue from thought: Activation patching reveals language-agnostic concept representations in transformers. *arXiv preprint arXiv:2411.08745*.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. [A mathematical framework for transformer circuits](#). *Transformer Circuits Thread*, 1(1):12.
- Atticus Geiger, Zhengxuan Wu, Hanson Lu, Josh Rozner, Elisa Kreiss, Thomas Icard, Noah Goodman, and Christopher Potts. 2022. Inducing causal structure for interpretable neural networks. In *International Conference on Machine Learning*, pages 7324–7338. PMLR.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. [Patchscopes: A unifying framework for inspecting hidden representations of language models](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Haoyang Huang, Tianyi Tang, Dongdong Zhang, Xin Zhao, Ting Song, Yan Xia, and Furu Wei. 2023. [Not all languages are created equal in LLMs: Improving multilingual capability by cross-lingual-thought prompting](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12365–12394, Singapore. Association for Computational Linguistics.
- J Jaavid, Raj Dabre, M Aswanth, Jay Gala, Thanmay Jayakumar, Ratish Puduppully, and Anoop Kunchukuttan. 2024. [Romansetu: Efficiently unlocking multilingual capabilities of large language models via romanization](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15593–15615.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Takeshi Kojima, Itsuki Okimura, Yusuke Iwasawa, Hitomi Yanaka, and Yutaka Matsuo. 2024. [On the multilingual ability of decoder-based pre-trained language models: Finding and controlling language-specific neurons](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6919–6971, Mexico City, Mexico. Association for Computational Linguistics.
- János Kramár, Tom Lieberum, Rohin Shah, and Neel Nanda. 2024. Atp\*: An efficient and scalable method for localizing llm behaviour to components. *arXiv preprint arXiv:2403.00745*.
- Anoop Kunchukuttan. 2020. [The indicnlp library](#).
- Yihong Liu, Chunlan Ma, Haotian Ye, and Hinrich Schuetze. 2024. [TransliCo: A contrastive learning framework to address the script barrier in multilingual pretrained language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2476–2499, Bangkok, Thailand. Association for Computational Linguistics.
- Yash Madhani, Sushane Parthan, Priyanka Bedekar, Gokul Nc, Ruchi Khapra, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh Khapra. 2023. [Aksharantar: Open Indic-language transliteration datasets and models for the next billion users](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 40–57, Singapore. Association for Computational Linguistics.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Mozillazg. 2024. [pypinyin: A python library to convert chinese characters to pinyin](#). Accessed: 2025-01-23.
- Nostalgebraist. 2020. [Interpreting GPT: The Logit Lens. LessWrong](#).
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, et al. 2023. [Language models are multilingual chain-of-thought reasoners](#). In *The Eleventh International Conference on Learning Representations*.
- Tianyi Tang, Wenyang Luo, Haoyang Huang, Dongdong Zhang, Xiaolei Wang, Xin Zhao, Furu Wei, and Ji-Rong Wen. 2024. [Language-specific neurons: The key to multilingual capabilities in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5701–5715, Bangkok, Thailand. Association for Computational Linguistics.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrut

Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Alexandre Variengien and Eric Winsor. 2024. [Look before you leap: A universal emergent decomposition of retrieval tasks in language models](#). In *ICML 2024 Workshop on Mechanistic Interpretability*.

Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English? on the latent language of multilingual transformers](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.

Orgest Xhelili, Yihong Liu, and Hinrich Schuetze. 2024. [Breaking the script barrier in multilingual pre-trained language models with transliteration-based post-training alignment](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11283–11296, Miami, Florida, USA. Association for Computational Linguistics.

Wenxuan Zhang, Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2023. [M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Jun Zhao, Zhihao Zhang, Luhui Gao, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024a. Llama beyond english: An empirical study on language capability transfer. *arXiv preprint arXiv:2401.01055*.

Yiran Zhao, Wenxuan Zhang, Guizhen Chen, Kenji Kawaguchi, and Lidong Bing. 2024b. [How do large language models handle multilingualism?](#) In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.

Chengzhi Zhong, Fei Cheng, Qianying Liu, Junfeng Jiang, Zhen Wan, Chenhui Chu, Yugo Murawaki, and Sadao Kurohashi. 2024. Beyond english-centric llms: What language do multilingual language models think in? *arXiv preprint arXiv:2408.10811*.

Tomaž Šolc. 2025. [Unidecode: Ascii transliterations of unicode text](#). Accessed: 2025-01-23.

## A Transformer’s Forward pass: Detailed

For an input sequence  $x_1, \dots, x_n \in V$ , where  $n$  is the sequence length, the initial latents  $h_1^{(0)}, \dots, h_n^{(0)} \in \mathbb{R}^d$  are obtained from a learned embedding matrix. The update rule for the latent at position  $i$  in layer  $j$  is expressed as:

$$h_i^{(j)} = h_i^{(j-1)} + f_j(h_1^{(j-1)}, \dots, h_i^{(j-1)})$$

The logit scores are computed as:

$$z_i = U h_i^{(k)}$$

These are converted to probabilities via the softmax function:

$$P(x_{i+1} = t | x_1, \dots, x_i) \propto \exp(z_{i,t})$$

## B Sample Prompts

A Hindi example, its English translation and transliteration for the translation, repetition and cloze task prompt designs mentioned in Section 4 are provided below.

### Hindi example.

*Translation task.* A translation prompt from French to Hindi.

<b>Français:</b> “poisson”	हिन्दी: “मछली”
<b>Français:</b> “mangue”	हिन्दी: “आम”
<b>Français:</b> “frère”	हिन्दी: “भाई”
<b>Français:</b> “odeur”	हिन्दी: “गंध”
<b>Français:</b> “soleil”	हिन्दी: “सूरज”
<b>Français:</b> “fleur”	हिन्दी:

*Repetition task.*

हिन्दी: “मछली”	हिन्दी: “मछली”
हिन्दी: “आम”	हिन्दी: “आम”
हिन्दी: “भाई”	हिन्दी: “भाई”
हिन्दी: “गंध”	हिन्दी: “गंध”
हिन्दी: “सूरज”	हिन्दी: “सूरज”
हिन्दी: “फूल”	हिन्दी:

*Cloze task..*

एक “\_\_\_” कहानियाँ पढ़ने के लिए उपयोग किया जाता है। उत्तर: “पुस्तक”  
 फुटबॉल और बास्केटबॉल जैसे खेल खेलने के लिए “\_\_\_” का उपयोग किया जाता है। उत्तर: “बॉल”  
 एक “\_\_\_” अक्षर उपहार के रूप में दिया जाता है और यह बगीचों में पाया जा सकता है। उत्तर:

### English Translation.

*Translation task.*

<b>Français:</b> "poisson"	<b>Hindi:</b> "fish"
<b>Français:</b> "mangue"	<b>Hindi:</b> "mango"
<b>Français:</b> "frère"	<b>Hindi:</b> "brother"
<b>Français:</b> "odeur"	<b>Hindi:</b> "smell"
<b>Français:</b> "soleil"	<b>Hindi:</b> "sun"
<b>Français:</b> "fleur"	<b>Hindi:</b>

*Repetition task.*

<b>Hindi:</b> "fish"	<b>Hindi:</b> "fish"
<b>Hindi:</b> "mango"	<b>Hindi:</b> "mango"
<b>Hindi:</b> "brother"	<b>Hindi:</b> "brother"
<b>Hindi:</b> "smell"	<b>Hindi:</b> "smell"
<b>Hindi:</b> "sun"	<b>Hindi:</b> "sun"
<b>Hindi:</b> "flower"	<b>Hindi:</b>

*Cloze task.*

A "\_\_\_" is used to play sports like soccer and basketball. Answer: "ball"  
A "\_\_\_" is used for reading stories. Answer: "book"  
A "\_\_\_" is often given as a gift and can be found in gardens. Answer:

## English Transliteration.

*Translation task.*

<b>Français:</b> "poisson"	<b>Hindi:</b> "machhalee"
<b>Français:</b> "mangue"	<b>Hindi:</b> "aam"
<b>Français:</b> "frère"	<b>Hindi:</b> "bhaee"
<b>Français:</b> "odeur"	<b>Hindi:</b> "gandh"
<b>Français:</b> "soleil"	<b>Hindi:</b> "sooraj"
<b>Français:</b> "fleur"	<b>Hindi:</b>

*Repetition task.*

<b>Hindi:</b> "machhalee"	<b>Hindi:</b> "machhalee"
<b>Hindi:</b> "aam"	<b>Hindi:</b> "aam"
<b>Hindi:</b> "bhaee"	<b>Hindi:</b> "bhaee"
<b>Hindi:</b> "gandh"	<b>Hindi:</b> "gandh"
<b>Hindi:</b> "sooraj"	<b>Hindi:</b> "sooraj"
<b>Hindi:</b> "phool"	<b>Hindi:</b>

*Cloze task.*

Phutball aur baasketball jaise khel khelane ke lie "\_\_\_" ka upayog kiya jaata hai. Uttar: "ball"

Ek "\_\_\_" kahaaniyaan padhane ke lie upayog kiya jaata hai. Uttar: "ball" Ek "\_\_\_" aksar upahaar ke roop mein diya jaata hai aur yah bageechon mein paaya ja sakata hai. Uttar:

## C Latent romanization: Tokenization scheme for the romanized word

Mathematically, we track the following romanized tokens for a given romanized word  $w$  of length  $n$ :

First timestep:  $\{w[0 : i] \mid 1 \leq i \leq n\} \cup \{"_ " + w[0 : i] \mid 1 \leq i \leq n\}$ , " \_ " represents single leading space and  $w[0 : i]$  represents prefixes of  $w$

Intermediate timesteps:  $\{w[i : j] \mid 0 \leq i < j \leq n\}$ ,  $w[i : j]$  represents sub-strings of  $w$

Final timestep:  $\{w[i : n] \mid 0 \leq i < n\}$ ,  $w[i : n]$  represents suffixes of  $w$

Similarly we construct token sets for the native script ( $T_{\text{native}}$ ) and English ( $T_{\text{English}}$ ) by including prefixes of the corresponding word, both with and without leading spaces, for all timesteps except the last. For the final timestep, we use suffixes of the corresponding word. We discard the sample if there is any overlap between romanized tokens ( $T_{\text{romanized}}$ ) and either the native ( $T_{\text{native}}$ ) or English tokens ( $T_{\text{English}}$ ) as shown by the following condition:

$$T_{\text{romanized}} \cap (T_{\text{native}} \cup T_{\text{English}}) = \emptyset$$

Lets take an example with a prompt translating "rope" from French to Hindi and derive romanized, English and native tokens for its first token generation timestep. The Hindi translation of "rope" is रस्सी and its romanized form is "rassi". So the romanized word tokens would be  $T_{\text{romanized}} = \{"r", "ra", "ras", "rass", "rassi", "_r", "_ra", "_ras", "_rass" \text{ and } "_rassi"\}$ . The corresponding English word tokens would be  $T_{\text{English}} = \{"r", "ro", "rop", "rope", "_r", "_ro", "_rop", \text{ and } "_rope"\}$ . The corresponding Hindi word tokens would be  $T_{\text{native}} = \{र, रस, रस्स, रस्सी, र, रस, रस्स, रस्सी\}$ . Here  $T_{\text{romanized}} \cap (T_{\text{native}} \cup T_{\text{English}}) = \{r, _r\}$  which is not null. As such we will exclude



this example translating rope to Hindi from the dataset to analyze *Latent Romanization*.

## D Latent Fraction

Formally, we compute the latent fraction as follows:

For layer  $l$ , timestep  $t$ , sample  $i$  and set of corresponding romanized tokens  $R$ :

### 1. Latent romanization condition:

$$r_{l,t}^{(i)} = \begin{cases} 1, & \text{if } \max_{r \in R} P(x_t = r \mid l, t) > 0.1 \\ 0, & \text{otherwise} \end{cases}$$

### 2. Latent fraction for a layer $\ell$ :

$$\text{L.F}(\ell) = \frac{1}{N} \sum_{i=1}^N \frac{1}{T} \sum_{t=1}^T r_{l,t}^{(i)}$$

where  $N$  is the number of samples,  $T$  is the number of generation timesteps and  $P(x_t = r \mid l, t)$  is the probability of generating token  $r$  at timestep  $t$  and layer  $\ell$ .

## E Computing Language probabilities - For translation towards native script vs translation towards romanized script task

To compute language probabilities, we search the LLM’s vocabulary for all tokens that could be the first token of the correct output word(s) in the respective language. We search the models vocabulary for all prefixes of the word(s) with and without leading space. For a language  $\ell$  with corresponding output word  $w_1$  and its synonyms  $w_2, w_3, \dots$ , we define:

$$P(\text{lang} = \ell) = \sum_{t_\ell \in W_{\text{prefix}}} P(x_{n+1} = t_\ell)$$

where  $W_{\text{prefix}}$  is the set of all prefixes of output word  $w_1$  and its synonyms  $w_2, w_3, \dots$ , including versions with and without leading spaces. For example to get probability for english when the output word is “fast” and its synonym is “swift”, then  $P(\text{lang} = \text{EN}) = P(x_{n+1} = \text{“f”}) + P(x_{n+1} = \text{“fa”}) + P(x_{n+1} = \text{“fas”}) + P(x_{n+1} = \text{“fast”}) + P(x_{n+1} = \text{“_f”}) + P(x_{n+1} = \text{“_fa”}) + P(x_{n+1} = \text{“_fas”}) + P(x_{n+1} = \text{“_fast”}) + P(x_{n+1} = \text{“s”}) + P(x_{n+1} = \text{“sw”}) + P(x_{n+1} = \text{“swi”}) + P(x_{n+1} = \text{“swif”}) + P(x_{n+1} =$

“swift”) +  $P(x_{n+1} = \text{“_s”}) + P(x_{n+1} = \text{“_sw”}) + P(x_{n+1} = \text{“_swi”}) + P(x_{n+1} = \text{“_swif”}) + P(x_{n+1} = \text{“_swift”})$  (all the token-level prefixes of “fast”, “\_fast”, “swift” and “\_swift”). “\_” represents a single leading space.

## F Romanization scheme: Indic languages

We have taken into consideration two romanization schemes for Indic languages: (a) ITRANS scheme from IndicNLP library (Kunchukuttan, 2020) and (b) IndicXlit scheme (Madhani et al., 2023). Based on our initial experiments, we observed that the IndicXlit scheme produces better romanization than ITRANS scheme. Thus for romanization we have used the IndicXlit scheme (Madhani et al., 2023). It generates romanization as is commonly used by native speakers and is trained on parallel transliteration corpora.

## G Computing Probabilities : Activation Patching Experiment

Probability for a concept  $C$  in language  $\ell$  can be formulated as :

$$P(C^\ell) = \sum_{t_\ell \in W_{\text{prefix}}} P(x_{n+1} = t_\ell)$$

where  $W_{\text{prefix}}$  is the set of all prefixes of output word  $w(C^\ell)$  and its synonyms (note that a word’s tokens are its prefixes).

We keep source concept  $C_S$  and target concept  $C_T$  distinct to avoid ambiguity when both are expressed in the same target language  $l_{\text{out}}^T$ .

Cases of token overlap between  $w(C_S^{\ell_T})$  i.e. word representing source concept in target language and  $w(C_T^{\ell_T})$  i.e. word representing target concept in target language and their synonyms are excluded. Token overlap would cause ambiguity. Therefore in the final dataset,

$$T(w(C_S^{\ell_T})) \cap T(w(C_T^{\ell_T})) = \emptyset$$

Where  $T(w)$  represents all the prefixes of  $w$  and its synonyms.

## H Latent Romanization Qualitative Analysis

We list qualitative logit lens analysis for Greek, Ukrainian, Hebrew, Arabic and Amharic (see Figures 8 to 12).

**Languages.** Greek is part of the Hellenic branch of the Indo-European language family and is written using the Greek alphabet. Ukrainian belongs to the East Slavic group of the Indo-European family and employs the Ukrainian alphabet, a variant of the Cyrillic script. Amharic is a South Ethiopian Semitic language within the Afroasiatic family and is written using the Ge'ez script, an abugida where each character represents a consonant-vowel combination. Hebrew is a Northwest Semitic language within the Afroasiatic family and is written using the Hebrew alphabet, an abjad script originating from the Aramaic alphabet. Arabic is a Central Semitic language, also part of the Afroasiatic family, and utilizes the Arabic script, another abjad that evolved from the Nabataean alphabet. Notably, both Hebrew and Arabic scripts are written from right to left.

## **I Latent Romanization Quantitative Analysis: Additional examples**

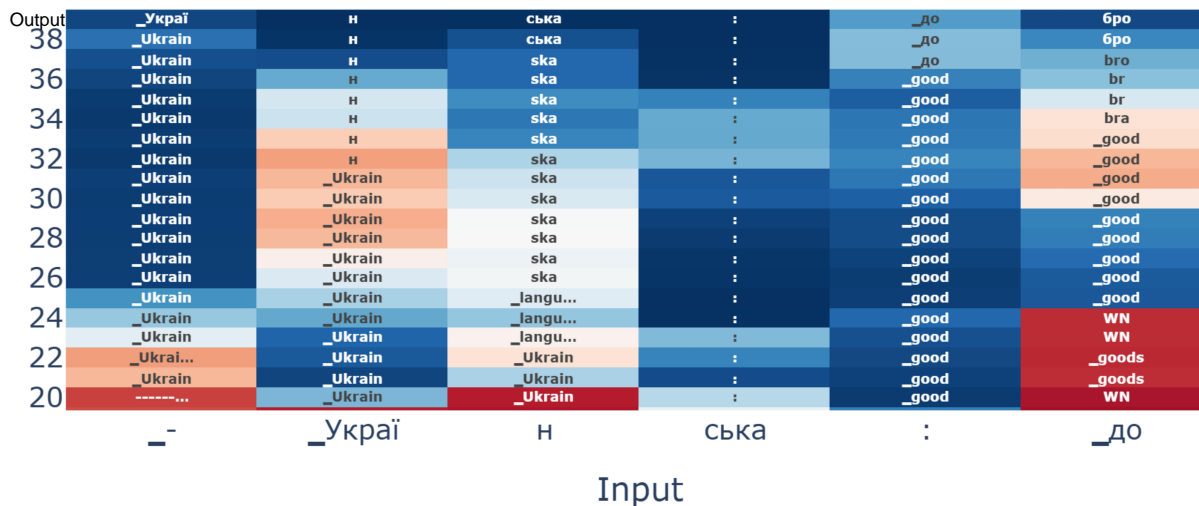

Quantitative Analysis of latent romanization for repetition task and cloze task with gemma 2 9b it model can be seen in Figures 13 and 14 respectively. Layerwise fractional distribution of romanized tokens across output token generation timesteps for translation, repetition and cloze task with Gemma 2 9b, Llama 2 7b, and Llama 2 13b models are present in Figure 15.

## **J Comparing Translations Into Romanized vs. Native Script: Additional examples**

Translation towards native script is compared with translation towards romanized script for gemma 2 9b it, gemma 2 9b and llama 2 13b models (see Figures 16 to 18).

## **K Other Models: Mistral**

We also perform our experiments on Mistral-7B (Jiang et al., 2023), a popular LLM known for its performance and efficiency. Layerwise distribution of romanized tokens for initial, final and all token generation steps are presented in Figure 19.






Figure 10: **Logit lens illustration.** We input Llama-2 13b model with a prompt translating ‘time’ from French to Hebrew. We visualize the output (זמן, ‘zman’ is the romanized form ) taking shape using logit lens producing a next-token distribution for each position (x-axis) and layers 20 and above (y-axis). Interestingly in the middle to top layers (20 - 29) we could observe romanized subwords of the Hebrew word (מן - man ; אן - an ). Color represents entropy of next-token generation from low (blue) to high (red). Plotting tool: (Belrose et al., 2023).






Figure 11: **Logit lens illustration.** We input Llama-2 13b model with a prompt translating ‘door’ from French to Arabic. We visualize the output (باب , ‘bab’ is the romanized form ) taking shape using logit lens producing a next-token distribution for each position (x-axis) and layers 20 and above (y-axis). Interestingly in the middle to top layers (20 - 29) we could observe romanized subwords of the Arabic word ( باب - bab; با - ab) before they are represented in their native script. Color represents entropy of next-token generation from low (blue) to high (red). Plotting tool: (Belrose et al., 2023).








(a) Layerwise distribution of romanized tokens averaged across output token generation steps and samples

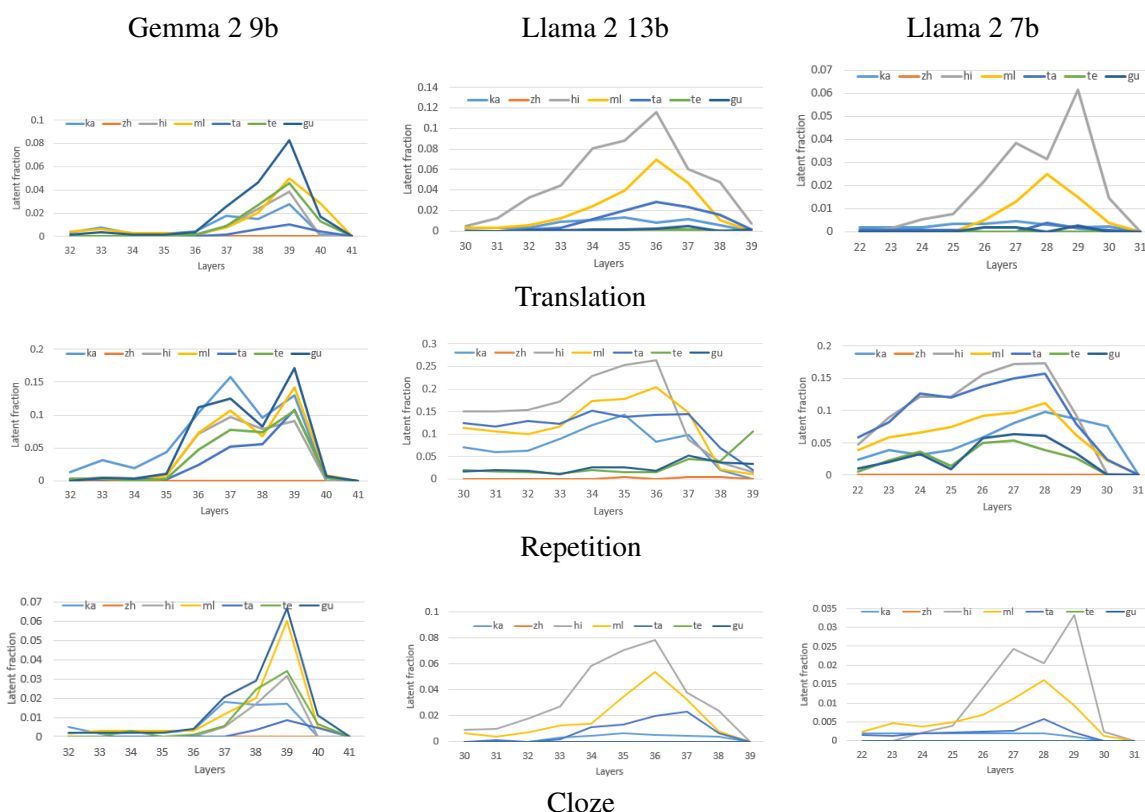


(b) Layerwise distribution of romanized tokens in the first token generation step averaged across samples



(c) Layerwise distribution of romanized tokens in the last token generation step averaged across samples

**Figure 14: Distribution of Romanized Tokens Across Model Layers: Analysis of First, Last, and All Generation Timesteps.** This distribution is plotted across the last 10 layers of Gemma-2 9b IT model for **cloze task** and is averaged across 100+ samples. X-axes represents layer index, y-axes represents latent fraction i.e. the instances where romanized tokens occur with a probability  $> 0.1$  averaged over samples for a specific layer. We plot the distributions for Gujarati (gu), Tamil (ta), Telugu (te), Hindi (hi) and Malayalam (ml), Georgian (ka) and Chinese (zh).



**Figure 15: Layerwise fractional distribution of romanized tokens across output token generation timesteps.** This distribution is plotted across the last 10 layers of Gemma 2 9b, Llama 2 13b and Llama 2 7b models (columns) for (a) translation task from French, (b) Repetition task and is averaged across 100+ samples. X-axes represents layer index, y-axes represents latent fraction i.e. the fraction of timesteps where romanized tokens occur with a probability  $> 0.1$  averaged over samples for a specific layer. We plot the distributions for Gujarati (gu), Tamil (ta), Telugu (te), Hindi (hi) and Malayalam (ml).

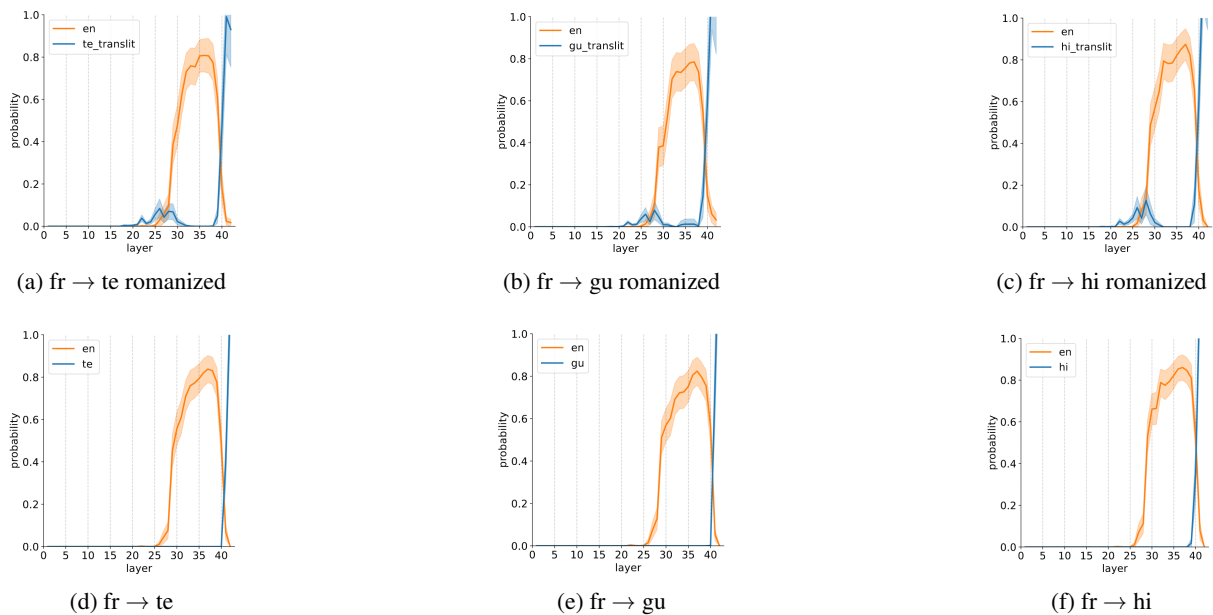


Figure 16: **Language probabilities for latent layers** in translation from French to Telugu, Gujarati and Hindi in romanized (top row) and native scripts (bottom row) across various samples using the Gemma-2 9B IT model. On  $x$ -axes, layer index; on  $y$ -axes, probability (according to logit lens) of correct next token in a given language. Error bars represent 95% Gaussian confidence intervals over input. In translations to non-English languages in romanized scripts (top row), target representations emerge slightly earlier—approximately one to two layers before layer 40—compared to their native script counterparts (bottom row), which only begin to appear from layer 40 onwards.

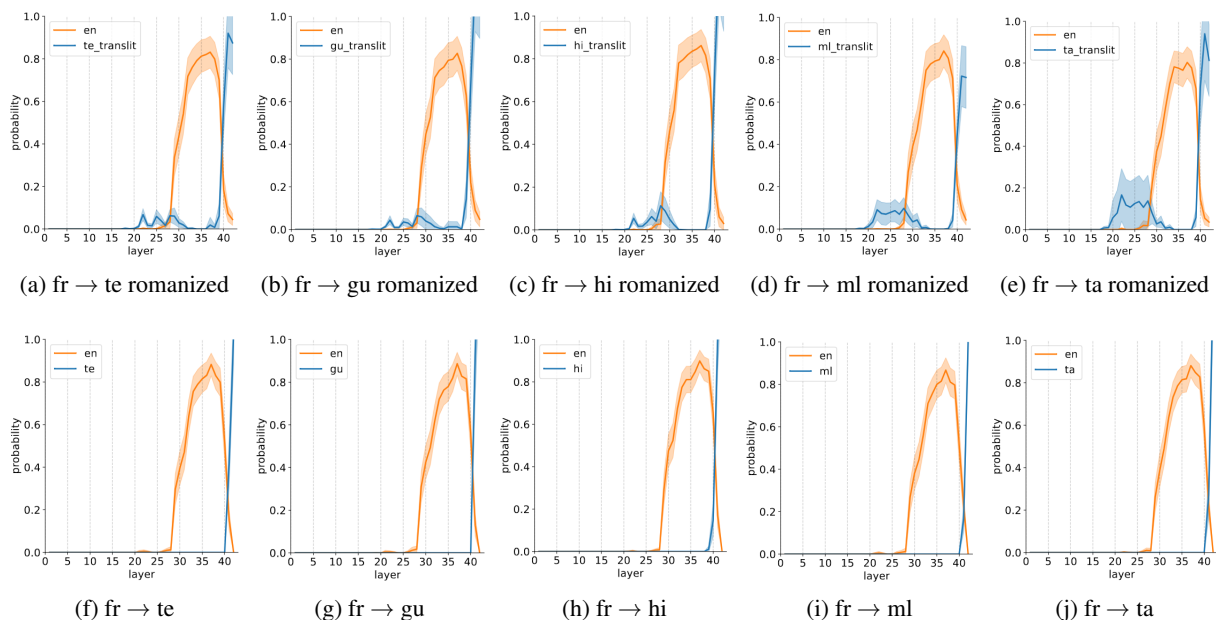


Figure 17: **Language probabilities for latent layers** in translation from French to five Indic languages (Telugu, Gujarati, Hindi, Malayalam, and Tamil) in romanized (top row) and native scripts (bottom row) using the Gemma-2 9B model. On  $x$ -axes, layer index; on  $y$ -axes, probability of correct next token in a given language. Error bars represent 95% Gaussian confidence intervals over input. In translations using romanized scripts (top row), target representations emerge approximately 1-2 layers earlier than their native script counterparts (bottom row).

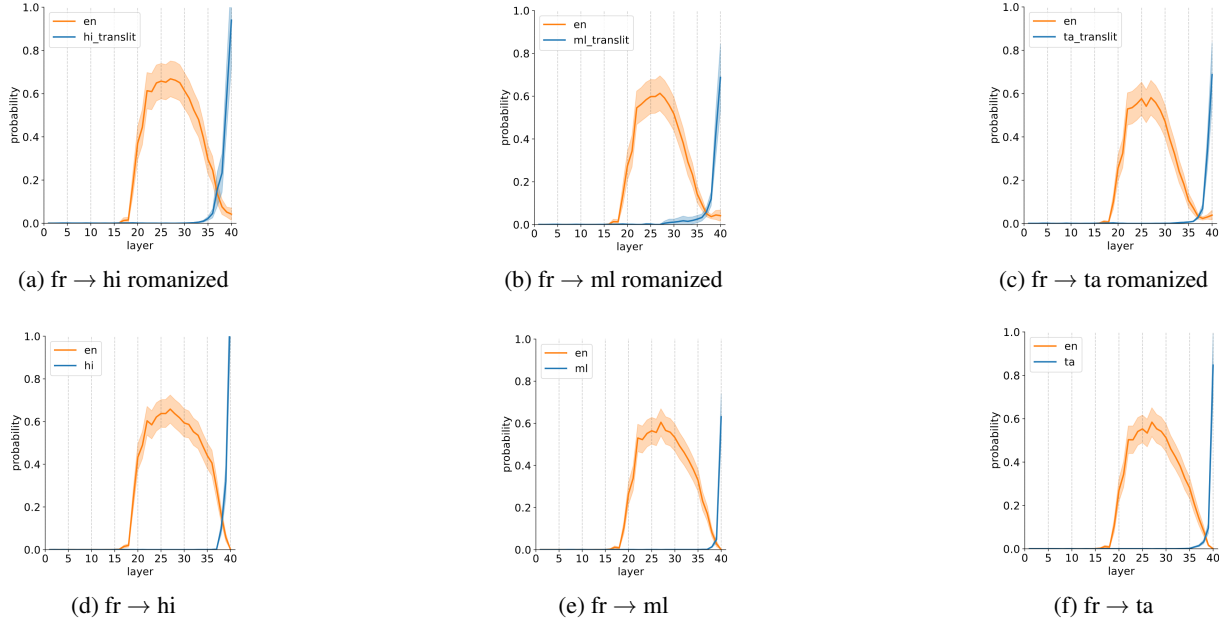


Figure 18: **Language probabilities for latent layers** in translation from French to Indic languages (Hindi, Malayalam, and Tamil) in romanized (top row) and native scripts (bottom row) using the Llama-2 13b model. On  $x$ -axes, layer index; on  $y$ -axes, probability of correct next token in a given language. Error bars represent 95% Gaussian confidence intervals over input. In translations using romanized scripts (top row), target representations emerge approximately 1-2 layers earlier than their native script counterparts (bottom row).

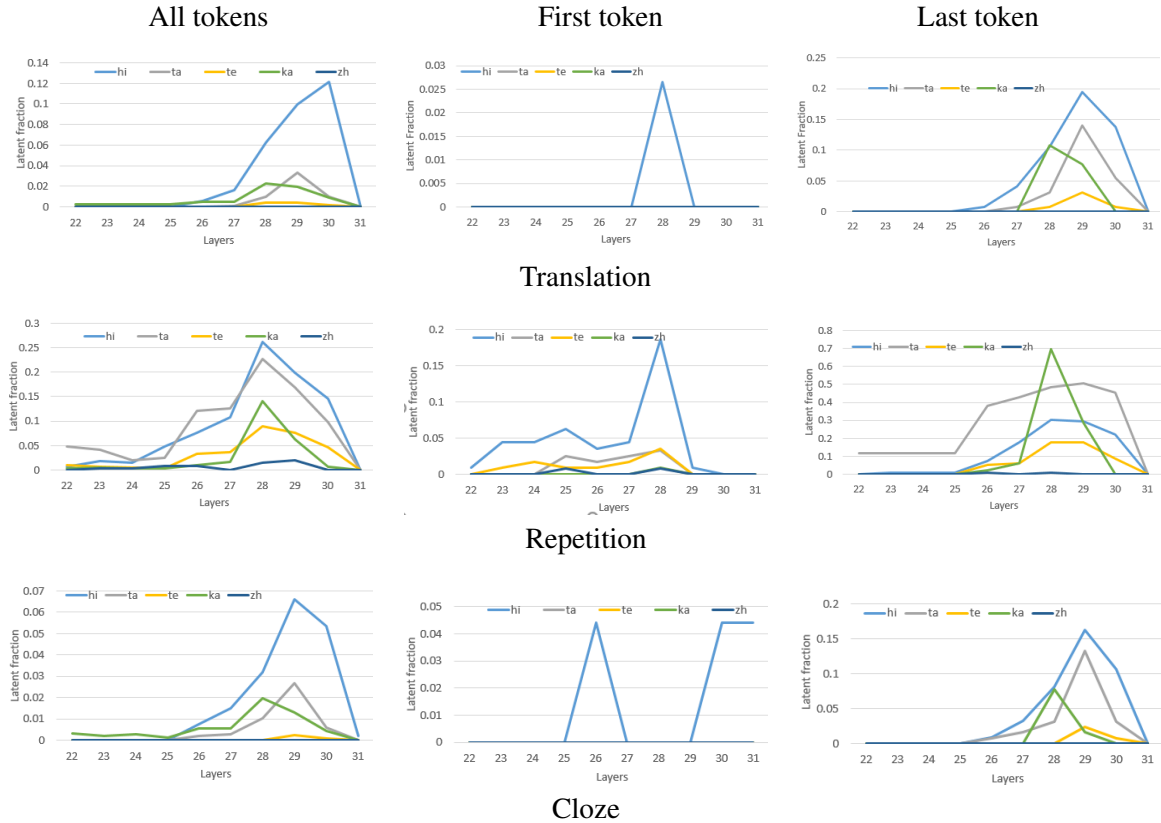


Figure 19: **Distribution of Romanized Tokens Across Model Layers: Analysis of First, Last, and All Generation Timesteps.** This distribution is plotted across the last 10 layers of **Mistral-7B** model for initial, final and all token generation steps (columns) for (a) translation task from English, (b) Repetition task, (c) Cloze task (rows) and is averaged across 100+ samples.  $X$ -axes represents layer index,  $y$ -axes represents latent fraction i.e. the fraction of timesteps where romanized tokens occur with a probability  $> 0.1$  averaged over samples for a specific layer. We plot the distributions for Tamil (ta), Telugu (te), Hindi (hi), Georgian (ka) and Chinese (zh).