

# Finding the Translation Switch: Discovering and Exploiting the Task-Initiation Features in LLMs

Xinwei Wu<sup>1\*</sup>, Heng Liu<sup>2\*</sup>, Xiaohu Zhao<sup>2</sup>, Yuqi Ren<sup>1</sup>, Linlong Xu<sup>2</sup>,  
Longyue Wang<sup>2</sup>, Deyi Xiong<sup>1†</sup>, Weihua Luo<sup>2</sup>, Kaifu Zhang<sup>2</sup>

<sup>1</sup> TJUNLP, Tianjin University, China

<sup>2</sup> Alibaba International Digital Commerce, China  
{wuxw2021, ryq20, dyxiong}@tju.edu.cn

## Abstract

Large Language Models (LLMs) frequently exhibit strong translation abilities, even without task-specific fine-tuning. However, the internal mechanisms governing this innate capability remain largely opaque. To demystify this process, we leverage Sparse Autoencoders (SAEs) and introduce a novel framework for identifying task-specific features. Our method first recalls features that are frequently co-activated on translation inputs and then filters them for functional coherence using a PCA-based consistency metric. This framework successfully isolates a small set of “translation initiation” features. Causal interventions demonstrate that amplifying these features steers the model towards correct translation, while ablating them induces hallucinations and off-task outputs, confirming they represent a core component of the model’s innate translation competency. Moving from analysis to application, we leverage this mechanistic insight to propose a new data selection strategy for efficient fine-tuning. Specifically, we prioritize training on “mechanistically hard” samples—those that fail to naturally activate the translation initiation features. Experiments show this approach significantly improves data efficiency and suppresses hallucinations. Furthermore, we find these mechanisms are transferable to larger models of the same family. Our work not only decodes a core component of the translation mechanism in LLMs but also provides a blueprint for using internal model mechanism to create more robust and efficient models.

## Introduction

Large Language Models (LLMs), particularly after instruction fine-tuning, often exhibit remarkable zero-shot translation capabilities, even without dedicated training on translation tasks (Hendy et al. 2023; Zhang, Haddow, and Birch 2023). The origin of this emergent behavior remains a subject of debate. Early investigations posited that this ability stems from inadvertently included parallel corpora in the vast pre-training datasets (Li and Flanigan 2024). However, subsequent work has shown that translation capabilities persist even when such data is meticulously removed (Briakou, Cherry, and Foster 2023; Fu et al. 2023), suggesting

\*These authors contributed equally.

†Corresponding author.

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

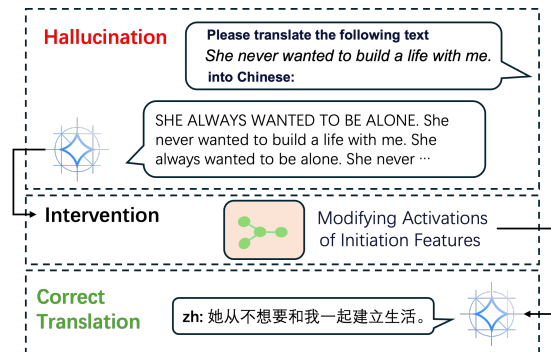


Figure 1: An illustration of the effect of translation initiation features. Initially, the model hallucinates by failing to switch to the target language (top). By amplifying the identified “translation initiation” features (middle), we causally steer the model to produce the correct translation (bottom).

a more fundamental mechanism is at play. This data-centric approach to explanation has become increasingly untenable as pre-training corpora scale to trillions of tokens, making exhaustive analysis computationally prohibitive (Shen et al. 2023; Chang et al. 2024). This leaves a fundamental question unanswered: what internal mechanism within instruction-tuned models is responsible for this innate translation ability?

In this work, we shift the focus from the pre-training corpus to the model’s internal representations, employing Sparse Autoencoders (SAEs) to demystify the translation mechanism (Cunningham et al. 2023; Gao et al. 2024; Templeton et al. 2024). SAEs allow us to decompose a model’s dense hidden states into a high-dimensional, yet sparsely activated, feature space, enabling more interpretable analysis of neural computations (Cunningham et al. 2023; Lieberum et al. 2024). To this end, we propose a novel, three-stage framework for identifying task-specific features. First, we recall features that are frequently co-activated across critical prompt positions during translation tasks. Second, we derive a canonical, directionally stable feature influence vector for each candidate feature directly from the sparse representation of SAEs. Finally, we employ a PCA-based consistency metric to filter for a functionally coherent group of features

whose vectors are highly aligned, representing a unified neural mechanism.

Applying this framework, we successfully isolate a small cluster of “translation initiation” features. Our analysis reveals two key properties of these features: (1) **Causal Generalization**: Causal interventions show that amplifying these features improves translation quality and reduces hallucinations (Wu et al. 2025) across diverse language pairs. (2) **Mechanistic Function**: They steer the model to generate explicit translation-framing tokens, locking it into the correct task context. These findings provide compelling evidence for a dedicated, general-purpose circuit that controls the initiation of the translation task.

To demonstrate the practical utility of these mechanistic insights, we move beyond direct intervention, which incurs prohibitive computational overhead (Li et al. 2024; Zhang et al. 2025). We propose a novel data selection strategy that uses the features’ activation as an inner metric for “instructional difficulty”, selectively fine-tuning on “mechanistically hard” samples where the model fails to naturally activate the initiation features. Our experiments demonstrate that this method significantly improves data efficiency and suppresses hallucinations. Furthermore, we show that these mechanistic insights are highly transferable to larger models *within the same model family* (e.g., Gemma-2B to Gemma-9B), but do not generalize across different model architectures (e.g., Gemma to LLaMA).

Our contributions are threefold:

1. We propose a novel, multi-stage framework for identifying functionally coherent, task-specific features in LLMs.
2. We discover and analyze a “translation initiation” features, providing a mechanistic explanation for the model’s innate translation abilities through causal and statistical evidence.
3. We introduce a new mechanistic data selection method for fine-tuning, demonstrating that insights into internal model workings can lead to more efficient and robust model specialization.

## Methodology

In this section, we detail our framework for identifying and filtering task-level features within LLMs. Specifically, our goal is to pinpoint the features responsible for initiating the translation task from the vast feature space decomposed by SAEs. Our methodology comprises three main stages: (1) **High-Frequency Feature Recall** to identify candidate features correlated with the task, (2) **Feature influence vector Characterization** to represent their directional influence, and (3) **Consistency-Based Filtering** to isolate functionally coherent and causally significant features.

### Preliminaries: Sparse Feature Extraction via SAEs

Our analysis focuses on the Gemma-2 model family (Team et al. 2024). We leverage the publicly available, pre-trained Sparse Autoencoders (SAEs) released by Google (Lieberum et al. 2024), which were trained on over 4 billion tokens. These SAEs are designed to decompose the model’s dense hidden states into a more interpretable, sparse set of features.

Given a translation input to the LLM, we perform a forward pass to obtain the hidden states from a specific layer  $l$ . Let  $\mathbf{h}^{(l)} \in \mathbb{R}^{d_{\text{model}}}$  be the hidden state vector at a particular token position from layer  $l$ , where  $d_{\text{model}}$  is the dimension of the hidden state.

An SAE is trained to reconstruct this hidden state  $\mathbf{h}^{(l)}$  from a sparse, high-dimensional feature activation vector  $\mathbf{a}(\mathbf{h}^{(l)}) \in \mathbb{R}^{d_{\text{sae}}}$ , where the feature dimension  $d_{\text{sae}}$  is typically much larger than  $d_{\text{model}}$ . The SAE consists of an encoder and a decoder. The encoder, with weight matrix  $\mathbf{W}_{\text{enc}} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{sae}}}$ , maps the hidden state to feature activations using the JumpReLU activation function, which is known to improve feature stability (Lieberum et al. 2024):

$$\mathbf{a}(\mathbf{h}^{(l)}) = \text{JumpReLU}(\mathbf{W}_{\text{enc}}^T \mathbf{h}^{(l)} + \mathbf{b}_{\text{enc}}) \quad (1)$$

The decoder, with weight matrix  $\mathbf{W}_{\text{dec}} \in \mathbb{R}^{d_{\text{sae}} \times d_{\text{model}}}$ , then attempts to reconstruct the original hidden state from these sparse activations:

$$\hat{\mathbf{h}}^{(l)} = \mathbf{W}_{\text{dec}} \mathbf{a}(\mathbf{h}^{(l)}) + \mathbf{b}_{\text{dec}} \quad (2)$$

In this framework, we define a “feature”  $f_{l,j}$  as a specific dimension  $j$  of the SAE dictionary at layer  $l$ . Its activation for a given hidden state  $\mathbf{h}^{(l)}$  is the  $j$ -th component of the sparse representation  $\mathbf{a}(\mathbf{h}^{(l)})$ . Each feature represents an abstract concept learned by the SAE. Our goal is to search through the tens of thousands of such features to find those causally linked to the translation task.

### Stage 1: Recalling High-Frequency Feature

A feature fundamental to translation should activate consistently when the model is prompted for the task (Geiger et al. 2024; Deng et al. 2025). In our first stage, we filter for such features by identifying those that are frequently active across a corpus of  $N$  translation samples. Recognizing that the task-related signal may appear at multiple locations, we monitor feature activations at three critical token positions: the final token of the source text (`src_last`), the target language token (`tgt_lang`), and the final query token (`input_last`).

For each sample, a feature  $f_{l,j}$  is considered “present” if its activation is non-zero at the **union** of these three positions—that is, if it is active at `src_last`, `tgt_lang`, or `input_last`. We then recall all features that are present in at least 60% of the samples ( $\tau_{\text{freq}} = 0.6$ ). This multi-position filtering strategy effectively narrows our search space to a manageable set of candidate features strongly correlated with the translation task.

### Stage 2: Characterizing Feature Influence Vectors

Stage 1 identifies features that are correlated with the translation task, but this does not guarantee their causal specificity or functional coherence (Deng et al. 2025). A truly task-specific feature should exert a consistent influence on the model’s residual stream across different contexts.

To operationalize this, we define a **feature influence vector**,  $\mathbf{v}_{l,j}$ , which captures the causal effect of a feature  $f_{l,j}$  within the context of a specific hidden state  $\mathbf{h}^{(l)}$ . This vector

is computed by measuring the change in the SAE’s reconstructed output when we force the activation of feature  $f_{l,j}$  to a high value,  $\alpha_{\text{act}}$ , while keeping all other feature activations as determined by  $\mathbf{h}^{(l)}$ .

Formally, let  $\hat{\mathbf{h}}_{\text{base}}$  be the standard reconstructed output from the SAE, and let  $\hat{\mathbf{h}}_{\text{intervene}}$  be the output after the aforementioned intervention on feature  $f_{l,j}$ . The feature influence vector is defined as:

$$\mathbf{v}_{l,j} \triangleq \hat{\mathbf{h}}_{\text{intervene}} - \hat{\mathbf{h}}_{\text{base}} \quad (3)$$

By this definition, each vector  $\mathbf{v}_{l,j}$  is context-dependent. This allows us to proceed to our next stage: analyzing the directional consistency of these vectors across a diverse set of inputs to identify features with a stable, monosemantic function.

### Step 3: Consistency-Based Filtering

The final stage validates whether our candidate features form a functionally coherent group. If a set of features truly represents a single underlying function, their corresponding feature influence vectors should be highly aligned, pointing in a near-identical direction in the residual stream. To quantify this collective alignment, we introduce a **PCA Consistency Score**.

Given a set of  $n$  candidate feature influence vectors  $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ , our method first normalizes each vector to unit length, focusing purely on directional information. The PCA Consistency Score is then defined as the proportion of variance explained by the first principal component (PC1) of these unit vectors. This score measures how well a single direction can represent the entire set.

As implemented in our methodology, this score can be computed efficiently. Let  $\mathbf{U} \in \mathbb{R}^{n \times d_{\text{model}}}$  be the matrix of the  $n$  unit-normalized feature influence vectors. The score,  $\rho$ , is the largest eigenvalue of the covariance matrix  $\mathbf{S} = \frac{1}{n} \mathbf{U}^T \mathbf{U}$ .

$$\rho = \lambda_{\max} \left( \frac{1}{n} \mathbf{U}^T \mathbf{U} \right) \quad (4)$$

A score close to 1 indicates that the vectors are almost perfectly aligned, while a low score suggests functional divergence.

We apply this metric to the set of candidate features recalled in Stage 1. A set is deemed to represent a unified neural circuit and is selected only if its PCA Consistency Score exceeds a stringent threshold of  $\tau_{\text{cons}} = 0.95$ . This final step ensures we isolate a small, highly-curated group of features that are not only correlated with the task but are also functionally synonymous.

## Experiments

### Experimental Setup

**Models and Interpretability Tools.** Our interpretability analysis primarily targeted **Gemma-2-2B-IT** (Team et al. 2024). This model was selected due to the public availability of high-quality, pre-trained SAEs provided by Google (Lieberum et al. 2024). These SAEs were trained on the

intermediate activations of the model’s residual streams, which extended hidden states size from 2304 to 16384.

To assess the generalizability of our findings, we extended our causal analysis and fine-tuning experiments to two additional instruction-tuned models: **Gemma-2-9B-IT**, **LLaMA-3.1-1B-IT** and **LLaMA-3.2-8B-IT** (Dubey et al. 2024). These models, varying in architecture and scale, allow us to test the robustness of our conclusions across different model families.

**Datasets and Partitioning.** Our primary analysis utilized the **WMT24++** dataset (Deutsch et al. 2025). To assess performance across diverse linguistic typologies, we select four translation directions: English-to-Chinese (en-zh), English-to-Arabic (en-ar), English-to-Russian (en-ru), and English-to-Japanese (en-ja). For each language pair, we used a subset of approximately 1,000 parallel sentences.

We used random samples of 98 sentence pairs as feature identification set. This set was used exclusively within our methodology (Section ) to discover task-relevant features. The remaining  $\approx 900$  sentence pairs were used for testing. This larger, held-out partition was used for all subsequent evaluations of translation performance and causal interventions.

Additionally, for the fine-tuning experiments presented in Section , we constructed a larger data pool by randomly sampling 100,000 English-to-Chinese sentence pairs from the general **WMT24** benchmark (Kocmi et al. 2024). This dataset was used exclusively for training and evaluating our data selection strategies.

**Evaluation Metrics.** We employ two core metrics to capture both the quality and faithfulness of the translations. We use the **COMET** score (Rei et al. 2020). COMET is a widely used, reference-based metric that has demonstrated a high correlation with human judgments. Following the taxonomy proposed by Huang et al. (2025), we classify any translation that is unfaithful to the source text as a *hallucination*. This includes instances of unexpected languages, empty output, and refusal to translate (e.g., repeating the source language). To quantify this, we adopt an LLM-as-a-Judge approach, closely following the evaluation setup proposed by Gogoulou et al. (2025). The **Hallucination Rate** is defined as the percentage of translations judged as unfaithful by our judge model across the test set.

### R1: Results of Identifying Translation Features

**Stage 1: High-Frequency Feature Recall.** Following our methodology, we began by identifying an initial candidate set of features. We analyzed the activation patterns on our *Feature Identification Set* within the MLP layers of Gemma-2-2B-IT and Gemma-2-9B-IT. Each layer’s activations were decomposed into 16,384 sparse features by its pre-trained SAE. We flagged a feature as a candidate if it activates at the **union** of three critical prompt locations (`src_last`, `tgt_lang`, or `input_last`) in over 60% of the samples.

The results of this recall stage are presented in Figure 2. The figure plots the proportion of recalled features per layer and reveals a strong, shared pattern: the density of task-correlated features consistently increases with model depth.

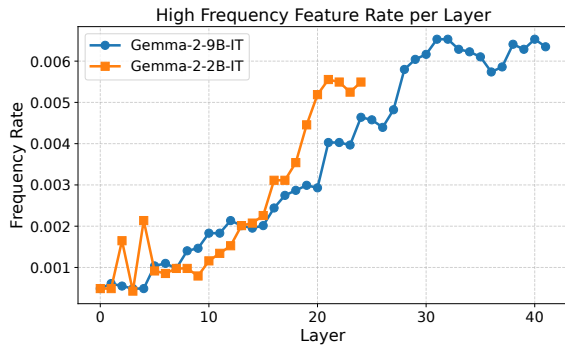


Figure 2: Comparison of high-frequency feature rates per layer for Gemma-2-2B-IT and Gemma-2-9B-IT.

Despite the significant difference in scale (26 vs. 42 layers), the distributional shape is strikingly similar, pointing to a robust architectural principle for representing abstract tasks. This stage yields a total of 1,004 candidate features for Gemma-2-2B-IT and 2,485 for Gemma-2-9B-IT.

**Stage 2 & 3: Isolating a Coherent Feature set.** The initial recall stage provides a broad set of candidates. To prune this set and isolate a functionally coherent set, we applied our PCA-based consistency filtering. We hypothesize that a true translation-initiation function is represented by a cluster of features whose canonical feature influence vectors are highly aligned. We tested this by grouping candidate features by layer and calculating a single **PCA Consistency Score** for each group, as defined in Equation 4. A group is considered a coherent circuit only if its score exceeds  $\tau_{\text{cons}} = 0.95$ .

While our filtering operates at the group level, for finer-grained analysis, we can measure the alignment of each individual feature’s vector with its group’s principal component (PC1). Figure 3 visualizes the distribution of these individual alignment scores. The boxplots reveal a critical insight: the majority of high-frequency features exhibit poor alignment, with median scores often below 0.4. This validates that co-activation alone is an insufficient criterion. However, a small number of outlier features achieve near-perfect alignment scores ( $> 0.95$ ). It is these features, residing within a handful of highly consistent layer-groups, that form our final, validated set of translation initiation features. For Gemma-2-2B-IT, this rigorous process distills the initial 1,004 candidates down to just **45** highly-consistent features.

**Validating Causal Impact via Intervention.** To confirm that our consistency metric successfully isolates causally significant features, we conducted targeted intervention experiments on Gemma-2-2B-IT. On the en-zh test set, we intervened on features grouped by their individual alignment score: low ( $< 0.5$ ), medium ( $[0.5, 0.95)$ ), and high ( $> 0.95$ ). We performed two types of **multiplicative interventions**:

- **Ablation:** We multiplied a feature’s activation by a coefficient of 0.
- **Amplification:** We multiplied a feature’s activation by a coefficient of 2.0.




Figure 3: Distribution of feature influence vector consistency scores for high-frequency features in (a) Gemma-2-2B-IT and (b) Gemma-2-9B-IT.

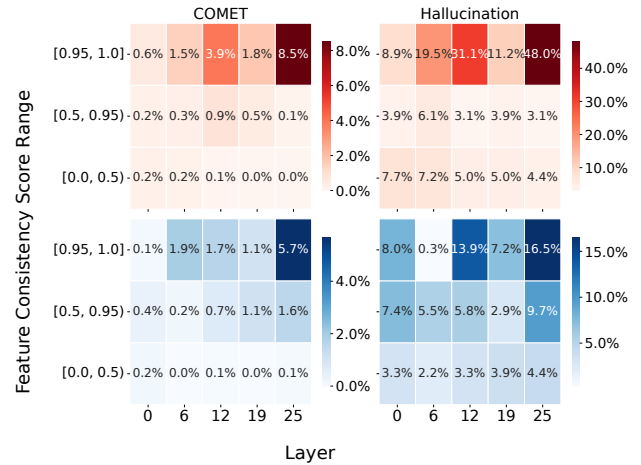


Figure 4: Absolute change in COMET and Hallucination Rate after feature intervention, grouped by layer and feature consistency score. Top row (red): ablation (coeff.=0). Bottom row (blue): amplification (coeff.=2.0).

We then measured the resulting absolute change in COMET and Hallucination Rate to quantify each feature group’s causal impact.

The results, summarized in Figure 4, are decisive. Across all layers, a clear gradient emerges: **a feature’s causal impact is directly proportional to its consistency score**. High-consistency features induce dramatic changes in model behavior, while low-consistency features have a negligible effect. For instance, ablating high-consistency features in layer 25 increases the hallucination rate by a staggering 47.99% and degrades the COMET score by 8.49%, confirming their necessity for faithful translation. This demonstrates that our framework successfully distills a large, noisy pool of candidates into a small set of features that act as the causal levers for the model’s translation capabilities.

## R2: Analysis of Translation-task Features

**Cross-Lingual Generalization of Task Features.** Our framework isolates a small set of high-impact features. We selected two prime candidates from Gemma-2-2B-IT, 112-f2291 and 113-f3517<sup>1</sup>, to test for functional gen-

<sup>1</sup>112-f2291 denotes the feature at position 2291 of layer 12.

Model	en-zh		en-ja		en-ru		en-ar	
	COMET $\uparrow$	Halluc. $\downarrow$	COMET $\uparrow$	Halluc. $\downarrow$	COMET $\uparrow$	Halluc. $\downarrow$	COMET $\uparrow$	Halluc. $\downarrow$
Gemma-2-2B-IT	73.62	19.15%	44.80	30.76%	54.36	29.46%	40.52	42.48%
+ 112-f2291	<b>77.98</b>	10.42%	47.62	<b>17.89%</b>	55.59	<b>16.37%</b>	42.02	<b>29.47%</b>
+ 113-f3517	77.83	<b>10.22%</b>	<b>47.95</b>	20.36%	<b>57.20</b>	19.26%	<b>42.38</b>	32.76%

Table 1: Cross-lingual generalization of translation-initiation features discovered on en-zh data. Activating feature 112-f2291 or 113-f3517 consistently improves COMET scores and drastically reduces hallucination rates across all tested language pairs, demonstrating their language-agnostic, task-initiating function.

erality. These features were discovered solely on English-to-Chinese (en-zh) data. To validate that they represent a language-agnostic function, we tested their causal effect on linguistically distant language pairs: En-Ja, En-Ru, and En-Ar. We performed a multiplicative intervention by amplifying each feature’s activation during inference and evaluate the impact on COMET and hallucination rates.

The results, presented in Table 1, demonstrate remarkable cross-lingual generalization. Amplifying either feature robustly improves translation performance across all four language pairs. Most notably, the intervention dramatically and universally slashes the hallucination rate, often by nearly half (e.g., from 19.15% to 10.42% in en-zh). The fact that features identified on a single language pair can universally enhance translation quality and faithfulness strongly suggests they do not encode language-specific knowledge. Instead, they encode a more abstract, language-agnostic instruction: to *initiate translation*.

**Interpreting Features through Output Statistics** To elucidate the interpretable function of our identified features, we conducted a statistical analysis of the model’s generated text. We first examined the baseline outputs and, through frequency analysis, observed that the model often uses a recurring set of tokens to preface its translations. We refer to these as “translation-framing tokens,” a sample of which is provided for each language in Figure 5.

We then measured the generation frequency of these specific tokens, comparing the baseline model against outputs generated with our feature amplification intervention. The results, presented in Figure 6, are unambiguous. Amplifying the feature’s activation leads to a significant and consistent increase in the emission rate of these pre-identified framing tokens across all four languages. For instance, the rate for Arabic tokens surged from 46.4% to 77.1%.

This establishes a direct, observable link: the feature’s activation is semantically correlated with an increased probability of generating these specific translation-framing tokens. This finding provides a clear, mechanistic explanation for the feature’s function. By promoting the emission of tokens that explicitly initiate the task, the feature forces the model to commit to translation at the outset. This initial commitment preempts off-topic drift, providing a direct causal pathway to the observed reduction in hallucinations and the corresponding improvements in translation quality.

zh: ["zh:", "Translation:", "Chinese (Simplified):"]  
ja: ["Japanese Translation:", "Translation:", "ja", "日本語訳:"]  
ru: ["Перевод:", "Translation:", "ru:", "Russian:"]  
ar: ["Translation:", "Arabic:", "Arabic Translation:", "ترجمة:"]

Figure 5: The curated lists of “translation-framing tokens” used for our analysis across the four target languages.

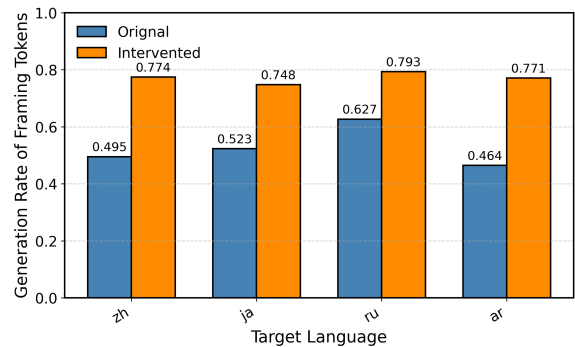


Figure 6: Rate of generating explicit “translation-framing tokens” with and without feature intervention.

### Application: Mechanistic Data Selection for Efficient Fine-Tuning

While direct causal intervention validates a feature’s function, its computational overhead is impractical for deployment (Li et al. 2024; Zhang et al. 2025). We therefore explore a more pragmatic application: using our identified feature as a signal for data selection (Xia et al. 2024) to improve fine-tuning efficiency. Our rationale is that the initiation feature’s activation reflects the model’s innate ability to recognize the translation task. Samples that fail to activate this feature are “mechanistically hard” or confusing to the model. We hypothesize that selectively fine-tuning on these confusing samples make task initiation features robust, which means the inner metric can be an effective teaching signal rather than solely external metrics.

To test this, we conducted two sets of fine-tuning experiments. First, we validated the core principle of mechanistic selection by using a model’s own feature signals for self-tuning. Second, we tested the transferability of these mechanistic insights by using features from a small model to select data for larger models. Additionally, we investigated the impact of different data selection ratios on the results.

Model	Method	Data Count	COMET $\uparrow$	Halluc. Rate $\downarrow$
<i>Part 1: Self-Tuning Efficiency (Using Model’s Own Features)</i>				
<b>Gemma-2-2B-IT</b>	Original	-	73.62	19.15%
	S0: Random	20k	82.49	3.62%
	S1: High-Quality	20k	83.32	2.12%
	S2: High-Loss	20k	82.14	4.32%
	S3: Mechanistic (Ours)	20k	<b>83.37</b>	<b>0.90%</b>
<b>LLaMA-3.1-1B-IT</b>	Original	-	57.61	32.24%
	S0: Random	20k	75.88	4.62%
	S1: High-Quality	20k	76.35	4.21%
	S2: High-Loss	20k	73.41	6.57%
	S3: Mechanistic (Ours)	20k	<b>77.92</b>	<b>2.39%</b>
<i>Part 2: Transfer-Tuning Scalability (Using Features from Gemma-2-2B)</i>				
<b>Gemma-2-9B-IT</b>	Original	-	79.50	12.50%
	S0: Random	50k	85.36	4.21%
	S1: High-Quality	50k	84.17	4.97%
	S2: High-Loss	50k	83.38	8.51%
	S3: Mechanistic (Ours)	50k	<b>86.48</b>	<b>0.60%</b>
<b>LLaMA-3.2-8B-IT</b>	Original	-	80.60	10.00%
	S0: Random	50k	86.20	0.20%
	S1: High-Quality	50k	<b>86.69</b>	<b>0.10%</b>
	S2: High-Loss	50k	84.51	0.20%
	S3: Mechanistic (Ours)	50k	86.34	0.30%

Table 2: Fine-tuning performance using different data selection strategies.

### Part 1: Validating the Efficacy of Mechanistic Selection.

In this experiment, we tested whether using a model’s own internal signals can improve its fine-tuning efficiency. We fine-tuned Gemma-2-2B-IT and LLaMA-3.1-1B-IT on targeted subsets of a 100k en-zh data pool. For each model, we used its corresponding open-source SAE to identify its translation initiation feature. We then compared three data selection strategies:

- (S0) **Random:** Selecting samples randomly.
- (S1) **High-Quality:** Selecting samples with the highest reference COMET score, which are generally high-quality for fine-tuning (Wang et al. 2024).
- (S2) **High-Loss:** A hard-mining baseline selecting high-COMET samples that also incur the highest training loss (Cao, Kang, and Sun 2023).
- (S3) **Mechanistic (Ours):** Selecting high-COMET samples that exhibit the lowest activation of the model’s own initiation features. We used the final set of high-consistency features produced by Step 3 of our framework.

The results, detailed in the top half of Table 2, confirm the effectiveness of our approach. For both Gemma-2-2B-IT and LLaMA-3.1-1B-IT, our mechanistic selection (S3) achieves a superior balance of performance, nearly matching the COMET scores of the best strategies while significantly reducing hallucination rates compared to other selection heuristics. This is particularly evident for Gemma-2-2B-IT, where our method achieves the performance of the full dataset with a fraction of the data, demonstrating high data efficiency. This confirms that a feature’s activation level

can be a reliable proxy for the “instructive difficulty” of a sample.

### Part 2: Testing the Transferability of Mechanistic Insights.

A key question is whether mechanistic insights are portable across model scales and families. Here, we used the initiation feature identified in the small Gemma-2-2B-IT to select data for fine-tuning two larger models: Gemma-2-9B-IT (intra-family transfer) and LLaMA-3.2-8B-IT (cross-family transfer).

The results reveal a clear pattern of family-specific transferability.

- **Intra-Family Transfer (Gemma  $\rightarrow$  Gemma):** As shown for Gemma-2-9B-IT, the benefits of mechanistic selection are highly pronounced. Using data selected via the smaller model’s feature allows the larger model to nearly match the performance of the full dataset. In contrast, the standard high-loss heuristic proves unreliable, leading to a substantial increase in hallucinations.
- **Cross-Family Transfer (Gemma  $\rightarrow$  LLaMA):** The mechanistic insight does not transfer effectively. For LLaMA-3.2-8B-IT, the high-quality data baseline (S1) performs best, and our feature-based selection strategy offers no clear advantage.

These findings yield a crucial insight: while the underlying neural circuits for tasks like translation appear highly portable across different scales *within the same model family*, they are not universally interchangeable across different architectural foundations. This demonstrates that our feature-based signal is a robust tool for efficient, family-specific model specialization.

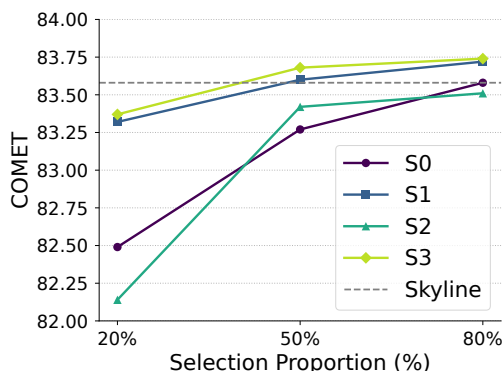


Figure 7: COMET score versus data selection proportion on the Gemma-2-2B-IT model.

**Part 3: Impact of Data Selection Ratio.** To further probe the data efficiency of our mechanistic selection, we analyzed how performance scales with the amount of selected data. We focused on the Gemma-2-2B-IT model and vary the fine-tuning data budget, selecting 20%, 50%, and 80% of the 100k data pool using the four strategies (S0-S3). We benchmarked against a “Skyline” performance, which is the COMET score achieved by fine-tuning on the full 100k dataset (83.58).

The results, depicted in Figure 7, highlight the superior data efficiency of our mechanistic approach (S3). Across all data proportions, S3 consistently outperforms or matches the best-performing baselines, particularly excelling in the low-data 20% regime. Most notably, with just 50% of the data, our method’s COMET score (83.68) **surpasses the full-dataset skyline** (83.58), demonstrating that it can curate a subset more effective than the entire dataset. This contrasts sharply with the high-loss (S2) and random (S0) baselines, which lag significantly. This finding strongly validates that our feature-based signal identifies a highly potent and compact training subset, maximizing performance while minimizing data requirements.

## Related Work

### Large Language Models for Translation

Large Language Models (LLMs) have surpassed traditional Neural Machine Translation (NMT) systems (Hendy et al. 2023; Li et al. 2023), leveraging document-level context (Sizov et al. 2024; Wang et al. 2023) and world knowledge while offering superior controllability over style and linguistic variations via prompting (Zhang, Haddow, and Birch 2023; Zhu, Cui, and Xiong 2024). Current specialization methods primarily involve Supervised Fine-Tuning (SFT) on parallel data (Zheng et al. 2024; Wang et al. 2025b) or Continual Pre-training to enhance multilingual foundations (Zhu, Jian, and Xiong 2025; Wang et al. 2025a).

Despite these advances, the innate translation ability of LLMs, which emerges even without explicit parallel training (Stap et al. 2024; Ji et al. 2024), remains a key area of investigation. A prevailing hypothesis has attributed this to

“incidental bilingualism” within the pre-training corpus (Li and Flanigan 2024), where non-parallel multilingual data inadvertently contains sentence-level translation pairs or code-switched text, providing implicit translation signals (Briakou, Cherry, and Foster 2023).

However, such data-centric analyses are prohibitively expensive at scale. In contrast, our work leverages interpretability to probe the model’s internal mechanisms directly. We aim to locate and understand the specific neural circuits, or features, that drive translation.

### Sparse Autoencoders for Interpretability

Sparse Autoencoders (SAEs) are a key technique in mechanistic interpretability (Sharkey et al. 2025) designed to address the superposition hypothesis (Elhage et al. 2022; Gurnee et al. 2023). This hypothesis posits that dense neural networks often encode multiple, distinct semantic concepts within the activations of a single neuron, making interpretation at both the neuron and embedding levels challenging. SAEs mitigate this issue by decomposing these dense representations into sparse, more semantically meaningful features, thereby enhancing model transparency and interpretability (Cunningham et al. 2023; Gao et al. 2024).

Current research on SAEs generally falls into two categories. The first focuses on theoretical exploration and the automated identification of general computational mechanisms within models (Dunefsky, Chlenski, and Nanda 2024). For instance, SAEs have been used to analyze and reproduce the Indirect Object Identification (IOI) circuit (Makelov 2024) and to determine if a model is “aware” of specific knowledge entities (Ferrando et al. 2024). The second category focuses on identifying task-specific features, using their derived activation signals to precisely intervene on and steer model behavior (Templeton et al. 2024; Farrell, Lau, and Conmy 2024; Han et al. 2025).

While prior SAE research has focused on simpler tasks, the complex mechanisms of translation remain underexplored. Our work addresses this gap by systematically applying SAEs to identify, validate, and leverage the key “trigger features” of the endogenous translation circuitry within LLMs, extending mechanistic interpretability into a more practical and complex domain.

## Conclusion

In this work, we have decoded a core component of the emergent translation capabilities in Large Language Models. We introduce a framework leveraging Sparse Autoencoders to isolate a “translation initiation” features and demonstrate through causal interventions that it governs the model’s ability to perform translation faithfully. Moving beyond analysis, we translate this mechanistic insight into a practical data selection strategy that significantly improves fine-tuning efficiency and reduces hallucinations. Our findings also reveal that these circuits are portable across scales within a model family, but not across different architectures. This research provides a blueprint for how understanding the internal workings of LLMs can lead directly to more robust and efficient model specialization.

## Acknowledgments

The present research was supported by the National Key Research and Development Program of China (Grant No. 2023YFE0116400). We would like to thank the anonymous reviewers for their insightful comments.

## References

- Briakou, E.; Cherry, C.; and Foster, G. 2023. Searching for Needles in a Haystack: On the Role of Incidental Bilingualism in PaLM’s Translation Capability. In *ACL*, 9432–9452.
- Cao, Y.; Kang, Y.; and Sun, L. 2023. Instruction mining: High-quality instruction data selection for large language models. *arXiv preprint arXiv:2307.06290*, 1(3): 6.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; et al. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3): 1–45.
- Cunningham, H.; Ewart, A.; Riggs, L.; Huben, R.; and Sharkey, L. 2023. Sparse autoencoders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600*.
- Deng, R.; Hu, X.; Gilberti, M.; Storks, S.; Taxali, A.; Angstadt, M.; Sripada, C.; and Chai, J. 2025. Sparse Feature Coactivation Reveals Composable Semantic Modules in Large Language Models. *arXiv preprint arXiv:2506.18141*.
- Deutsch, D.; Briakou, E.; Caswell, I.; Finkelstein, M.; Gallor, R.; Juraska, J.; Kovacs, G.; Lui, A.; Rei, R.; Riesa, J.; Rijnhwani, S.; Riley, P.; Salesky, E.; Trabelsi, F.; Winkler, S.; Zhang, B.; and Freitag, M. 2025. WMT24++: Expanding the Language Coverage of WMT24 to 55 Languages & Dialects. *arXiv:2502.12404*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv e-prints*, arXiv–2407.
- Dunefsky, J.; Chlenski, P.; and Nanda, N. 2024. Transcoders find interpretable llm feature circuits. *Advances in Neural Information Processing Systems*, 37: 24375–24410.
- Elhage, N.; Hume, T.; Olsson, C.; Schiefer, N.; Henighan, T.; Kravec, S.; Hatfield-Dodds, Z.; Lasenby, R.; Drain, D.; Chen, C.; et al. 2022. Toy models of superposition. *arXiv preprint arXiv:2209.10652*.
- Farrell, E.; Lau, Y.-T.; and Conmy, A. 2024. Applying Sparse Autoencoders to Unlearn Knowledge in Language Models. In *Neurips Safe Generative AI Workshop 2024*.
- Ferrando, J.; Obeso, O.; Rajamanoharan, S.; and Nanda, N. 2024. Do I Know This Entity? Knowledge Awareness and Hallucinations in Language Models. *arXiv preprint arXiv:2411.14257*.
- Fu, T.; Liu, L.; Cai, D.; Huang, G.; Shi, S.; and Yan, R. 2023. The reasonableness behind unreasonable translation capability of large language model. In *The Twelfth International Conference on Learning Representations*.
- Gao, L.; la Tour, T. D.; Tillman, H.; Goh, G.; Troll, R.; Radford, A.; Sutskever, I.; Leike, J.; and Wu, J. 2024. Scaling and evaluating sparse autoencoders. *arXiv preprint arXiv:2406.04093*.
- Geiger, A.; Wu, Z.; Potts, C.; Icard, T.; and Goodman, N. 2024. Finding alignments between interpretable causal variables and distributed neural representations. In *Causal Learning and Reasoning*, 160–187. PMLR.
- Gogoulou, E.; Zahra, S.; Guillou, L.; Dürlich, L.; and Nivre, J. 2025. Can LLMs Detect Intrinsic Hallucinations in Paraphrasing and Machine Translation? *arXiv preprint arXiv:2504.20699*.
- Gurnee, W.; Nanda, N.; Pauly, M.; Harvey, K.; Troitskii, D.; and Bertsimas, D. 2023. Finding neurons in a haystack: Case studies with sparse probing. *arXiv preprint arXiv:2305.01610*.
- Han, Z.; Wu, X.; Shi, D.; Jin, R.; and Xiong, D. 2025. Towards a Unified Paradigm of Concept Editing in Large Language Models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 18456–18472.
- Hendy, A.; Abdelrehim, M.; Sharaf, A.; Raunak, V.; Gabr, M.; Matsushita, H.; Kim, Y. J.; Afify, M.; and Awadalla, H. H. 2023. How good are gpt models at machine translation? a comprehensive evaluation. *arXiv preprint arXiv:2302.09210*.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; et al. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2): 1–55.
- Ji, B.; Duan, X.; Zhang, Y.; Wu, K.; and Zhang, M. 2024. Zero-shot prompting for llm-based machine translation using in-domain target sentences. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Kocmi, T.; Avramidis, E.; Bawden, R.; Bojar, O.; Dvorkovich, A.; Federmann, C.; Fishel, M.; Freitag, M.; Gowda, T.; Grundkiewicz, R.; et al. 2024. Findings of the WMT24 general machine translation shared task: The LLM era is here but MT is not solved yet. In *Proceedings of the Ninth Conference on Machine Translation*, 1–46.
- Li, C.; and Flanigan, J. 2024. Task contamination: Language models may not be few-shot anymore. In *AAAI*, volume 38, 18471–18480.
- Li, Q.; Liu, X.; Tang, Z.; Dong, P.; Li, Z.; Pan, X.; and Chu, X. 2024. Should We Really Edit Language Models? On the Evaluation of Edited Language Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Li, S.; Wei, X.; Zhu, S.; Xie, J.; Yang, B.; and Xiong, D. 2023. MMNMT: Modularizing multilingual neural machine translation with flexibly assembled MoE and dense blocks. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, 4978–4990.
- Lieberum, T.; Rajamanoharan, S.; Conmy, A.; Smith, L.; Sonnerat, N.; Varma, V.; Kramár, J.; Dragan, A.; Shah, R.; and Nanda, N. 2024. Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2. *arXiv preprint arXiv:2408.05147*.

- Makelov, A. 2024. Sparse autoencoders match supervised features for model steering on the ioi task. In *ICML 2024 Workshop on Mechanistic Interpretability*.
- Rei, R.; Stewart, C.; Farinha, A. C.; and Lavie, A. 2020. COMET: A neural framework for MT evaluation. *arXiv preprint arXiv:2009.09025*.
- Sharkey, L.; Chughtai, B.; Batson, J.; Lindsey, J.; Wu, J.; Bushnaq, L.; Goldowsky-Dill, N.; Heimersheim, S.; Ortega, A.; Bloom, J.; et al. 2025. Open Problems in Mechanistic Interpretability. *arXiv preprint arXiv:2501.16496*.
- Shen, T.; Jin, R.; Huang, Y.; Liu, C.; Dong, W.; Guo, Z.; Wu, X.; Liu, Y.; and Xiong, D. 2023. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*.
- Sizov, F.; España-Bonet, C.; van Genabith, J.; Xie, R.; and Chowdhury, K. D. 2024. Analysing translation artifacts: A comparative study of LLMs, NMTs, and human translations. In *Proceedings of the Ninth Conference on Machine Translation*, 1183–1199.
- Stap, D.; Hasler, E.; Byrne, B.; Monz, C.; and Tran, K. M. 2024. The Fine-Tuning Paradox: Boosting Translation Quality Without Sacrificing LLM Abilities. In *ACL*, 6189–6206.
- Team, G.; Riviere, M.; Pathak, S.; Sessa, P. G.; Hardin, C.; Bhupatiraju, S.; Hussenot, L.; Mesnard, T.; Shahriari, B.; Ramé, A.; et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Templeton, A.; Conerly, T.; Marcus, J.; Lindsey, J.; Bricken, T.; Chen, B.; Pearce, A.; Citro, C.; Ameisen, E.; Jones, A.; Cunningham, H.; Turner, N. L.; McDougall, C.; MacDiarmid, M.; Freeman, C. D.; Sumers, T. R.; Rees, E.; Batson, J.; Jermyn, A.; Carter, S.; Olah, C.; and Henighan, T. 2024. Scaling monosemanticity: Extracting interpretable features from Claude 3 Sonnet. Transformer Circuits Thread.
- Wang, H.; Xu, L.; Liu, H.; Liu, Y.; Zhao, X.; Zeng, B.; Shao, L.; Wang, L.; Luo, W.; and Zhang, K. 2025a. Beyond Single-Reward: Multi-Pair, Multi-Perspective Preference Optimization for Machine Translation. *arXiv preprint arXiv:2510.13434*.
- Wang, H.; Xu, L.; Liu, H.; Liu, Y.; Zhao, X.; Zeng, B.; Wang, L.; Luo, W.; and Zhang, K. 2025b. Marco Large Translation Model at WMT2025: Transforming Translation Capability in LLMs via Quality-Aware Training and Decoding. In *Proceedings of the Tenth Conference on Machine Translation*, 587–593.
- Wang, J.; Zhang, B.; Du, Q.; Zhang, J.; and Chu, D. 2024. A survey on data selection for llm instruction tuning. *arXiv preprint arXiv:2402.05123*.
- Wang, L.; Lyu, C.; Ji, T.; Zhang, Z.; Yu, D.; Shi, S.; and Tu, Z. 2023. Document-Level Machine Translation with Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 16646–16661.
- Wu, X.; Liu, H.; Zhou, J.; Zhao, X.; Xu, L.; Wang, L.; Luo, W.; and Zhang, K. 2025. Challenging Multilingual LLMs: A New Taxonomy and Benchmark for Unraveling Hallucination in Translation. *arXiv preprint arXiv:2510.24073*.
- Xia, M.; Malladi, S.; Gururangan, S.; Arora, S.; and Chen, D. 2024. Less: Selecting influential data for targeted instruction tuning. *arXiv preprint arXiv:2402.04333*.
- Zhang, B.; Haddow, B.; and Birch, A. 2023. Prompting large language model for machine translation: A case study. In *International Conference on Machine Learning*, 41092–41110. PMLR.
- Zhang, T.; Fang, J.; Jiang, H.; Bi, B.; Wang, X.; and He, X. 2025. Explainable and efficient editing for large language models. In *Proceedings of the ACM on Web Conference 2025*, 1963–1976.
- Zheng, J.; Hong, H.; Liu, F.; Wang, X.; Su, J.; Liang, Y.; and Wu, S. 2024. Fine-tuning large language models for domain-specific machine translation. *arXiv preprint arXiv:2402.15061*.
- Zhu, S.; Cui, M.; and Xiong, D. 2024. Towards robust in-context learning for machine translation with large language models. In *LREC-COLING 2024*, 16619–16629.
- Zhu, S.; Jian, D.; and Xiong, D. 2025. A Survey of Multilingual Neural Machine Translation Based on Sparse Models. *Tsinghua Science and Technology*, 30(6): 2399–2418.