

# Large Language Models Share Representations of Latent Grammatical Concepts Across Typologically Diverse Languages

Jannik Brinkmann<sup>1,2</sup> Chris Wendler<sup>3</sup> Christian Bartelt<sup>2</sup> Aaron Mueller<sup>1,4</sup>  
<sup>1</sup>Northeastern University <sup>2</sup>TU Clausthal <sup>3</sup>EPFL <sup>4</sup>Technion – IIT

## Abstract

Human bilinguals often use similar brain regions to process multiple languages, depending on when they learned their second language and their proficiency. In large language models (LLMs), how are multiple languages learned and encoded? In this work, we explore the extent to which LLMs share representations of morphosyntactic concepts such as grammatical number, gender, and tense across languages. We train sparse autoencoders on Llama-3-8B and Aya-23-8B, and demonstrate that abstract grammatical concepts are often encoded in feature directions shared across many languages. We use causal interventions to verify the multilingual nature of these representations; specifically, we show that ablating only multilingual features decreases classifier performance to near-chance across languages. We then use these features to precisely modify model behavior in a machine translation task; this demonstrates both the generality and selectivity of these feature’s roles in the network. Our findings suggest that even models trained predominantly on English data can develop robust, cross-lingual abstractions of morphosyntactic concepts.

## 1 Introduction

In the brains of human bilinguals, syntax processing may occur in similar regions for their first and second language, depending on factors like when the second language was learned (Cargnelli et al., 2019), language proficiency (Połczyńska and Bookheimer, 2021), among many other factors (Sulpizio et al., 2020; Costa and Sebastián-Gallés, 2014). In multilingual language models (LMs; Shannon, 1948), how apt is the analogy of shared processing to human bilinguals? If we desire parameter-efficiency, we might want multilingual representations of concepts such as grammatical number, rather than many language-specific repre-

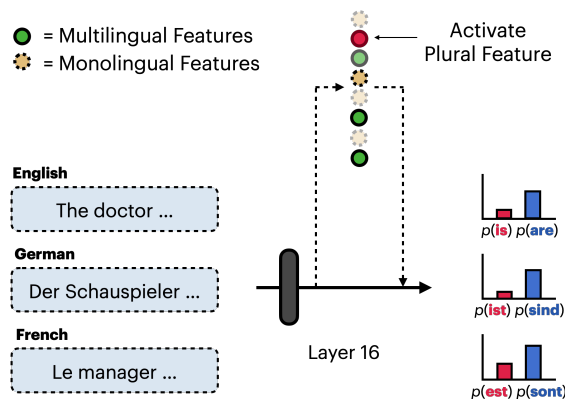


Figure 1: Using sparse autoencoders, we find that language models share representations of grammatical concepts across languages. By intervening on these multilingual representations, we can change the model behavior given inputs in different languages. For example, we can make the model predict plural verbs in *different* languages by activating the *same* plural feature.

sentations of the same concept.

Past work has emphasized language-balanced pretraining corpora (e.g., Conneau et al., 2020; Xue et al., 2021; Mohammadshahi et al., 2022), such that an LM could be said to have many primary languages. However, many of the best-performing multilingual LMs are now primarily English models, trained on over 90% English text (e.g., Dubey et al., 2024). Why do these models perform so well in non-English languages? We hypothesize that these models learn **generalizable abstractions** that enable more efficient learning of new languages. In other words, being able to deeply characterize a smaller distribution could allow models to acquire more robust abstractions that may generalize more effectively to a wider distribution post-hoc; in contrast, more balanced corpora (wider distributions) could encourage the model to start by learning language-specific abstractions, which are po-

We release code, data, and autoencoders at [github.com/jannik-brinkmann/multilingual-features](https://github.com/jannik-brinkmann/multilingual-features).

tentially never merged into higher-level language-invariant concepts. For small models, on a synthetic language pair (English and a token-level duplication of it), Schäfer et al. show that imbalanced datasets lead to higher sample efficiency and higher overall performance in both languages in the low-data regime.<sup>1</sup>

In large language models, multilinguality has both practical and theoretical scientific value. High performance on natural language processing tasks across many languages increases the impact and inclusivity of language technologies, and in a more compute-efficient manner than would be possible by training a large series of monolingual models. Moreover, multilingual models are hypothesized to be able to outperform monolingual models in low-resource languages: past work has emphasized the importance of **cross-lingual transfer**, where knowledge in one language is shared with another language whose corpora did not contain the information of interest (Libovický et al., 2020; Chang et al., 2022; Hua et al., 2024). More abstractly, the extent to which abstract concepts are shared across languages addresses a key question: what kinds of multi-lingual processing mechanisms are acquirable from exposure to large-scale text distributions? It would be more parameter-efficient to learn generalizable abstractions that apply across languages, as opposed to redundantly learning the same grammatical feature for each language separately. Moreover, the existence of generalizable abstractions informs debates on memorization versus generalization: cross-linguistic generalization suggests a more sophisticated and broad application of particular concepts.

Past work has largely investigated multilingual language models using behavioral/benchmarking analyses (Gupta and Srikumar, 2021; Raganato et al., 2020, *inter alia*) or neuron-level mechanistic analyses (Stanczak et al., 2022; De Varda and Marelli, 2023). Stanczak et al. observe that masked language models often share neurons across languages for a particular concept, but causally verifying this is difficult, as counterfactual interventions to neurons often affect the processing of irrelevant concepts (see §2). To overcome this challenge, we make our units of causal analysis *sparse autoencoder features*, which have been shown to be more *monosemantic*, and therefore more human-interpretable (Cunningham et al., 2024; Bricken

<sup>1</sup>For real languages, we expect that features will be shared across languages to a much greater extent.

et al., 2023; Rajamanoharan et al., 2024). This enables more precise causal interventions (Marks et al., 2025; Mueller et al., 2024), and allows us to qualitatively verify that features of interest are truly sensitive to the hypothesized feature, rather than a spurious but close-enough feature.

In this work, we train a set of sparse autoencoders on the intermediate activations of Llama-3-8B (Dubey et al., 2024) and Aya-3-8B (Aryabumi et al., 2024) and locate (massively) multilingual features for various morphosyntactic concepts. We design experiments to quantify the degree to which these concepts are shared across languages, and validate their role in the models generations. Our results reveal that language models share morphosyntactic concept representations across typologically diverse languages, and that the internal *lingua franca* of large language models may not be English words *per se*, but rather concepts.<sup>2</sup>

## 2 Background

**Feature disentanglement using sparse autoencoders.** The features underlying model computation are not guaranteed to be aligned to neuron bases; they may instead be represented in a distributed manner, such that there is a many-to-many relationship between neurons and concepts (Hinton et al., 1986; Smolensky, 1986). In practice, single neurons<sup>3</sup> are often *polysemantic*; in other words, they activate on a range of seemingly unrelated concepts (Bolukbasi et al., 2021; Elhage et al., 2022b). For example, Bricken et al. (2023) observe that a single neuron in a small LM responds to a mixture of academic citations, English dialogue, HTTP requests, and Korean text. To address this, Cunningham et al. (2024); Bricken et al. (2023) propose sparse autoencoders (SAEs) as a scalable technique for unsupervised discovery of interpretable feature directions in neural networks. Given an activation vector  $\mathbf{x} \in \mathbb{R}^{d_{\text{model}}}$ , the autoencoder computes a decomposition

$$\mathbf{x} = \hat{\mathbf{x}} + \boldsymbol{\epsilon}(x) = \mathbf{b} + \sum_i f_i(\mathbf{x})\mathbf{v}_i + \boldsymbol{\epsilon}(x) \quad (1)$$

into an approximate reconstruction  $\hat{\mathbf{x}}$  as a linear combination of features  $\mathbf{v}_i$ . The features  $\mathbf{v}_i \in \mathbb{R}^{d_{\text{model}}}$  are unit vectors, the feature activations

<sup>2</sup>Though we note that these concepts are likely to be biased toward English-like representations, as discussed in Wendler et al. (2024).

<sup>3</sup>We use “neuron” to refer to a single dimension of any hidden representation vector in a model.

$f_i(\mathbf{x}) \in \mathbb{R}$  are a sparse set of coefficients,  $\mathbf{b} \in \mathbb{R}^{d_{\text{model}}}$  is a bias term, and  $\epsilon(x) \in \mathbb{R}^{d_{\text{model}}}$  is the approximation error.

Bricken et al. (2023) train SAEs by minimizing an L2 reconstruction error and an L1 regularization term to promote sparsity. However, the L1 regularization term introduces biases that can harm the accuracy of the reconstruction. Therefore, we use the Gated SAE architecture (Rajamanoharan et al., 2024), which separates the functionality of selecting the features to use and estimating the activation magnitude. We provide additional details about Gated SAEs and their training in App. D.

A key idea underlying SAE-based interpretability is that one can reinterpret the language model’s (LM) internal computations in terms of the SAE features. By applying the decomposition in Eq. 1 to hidden states  $\mathbf{x}$  in the LM and folding the SAEs into the forward pass,<sup>4</sup> we can express model representations as a combination of SAE feature activations  $f_i$  and reconstruction errors  $\epsilon$ . This reframing enables gradient-based attribution patching techniques (Syed et al., 2024) to identify which SAE features have causal influence on particular aspects of model behavior.

**Finding causally relevant features using attribution patching.** Causal interpretability methods aim to locate components in neural networks (such as attention heads or neurons) responsible for particular behaviors (Mueller et al., 2024). This requires constructing a distribution  $\mathcal{D}$  over pairs of inputs  $(x_{\text{clean}}, x_{\text{patch}})$ , where  $x_{\text{clean}}$  is a prompt on which the behaviour occurs, and  $x_{\text{patch}}$  is a reference prompt on which the behaviour does not occur. Then, we can quantify the importance of components using a target metric  $m : \mathbb{R}^V \rightarrow \mathbb{R}$ , e.g., the difference in log-probabilities between a correct token continuation and minimally differing incorrect token continuation. In other words, we quantify the importance of a component on the model behaviour as the expected change in  $\mathcal{L}$  when replacing the components’s activation on the clean prompt with its value on the patch prompt. This is known as the **indirect effect** (IE; Pearl, 2001).

An exact but compute-intensive method for causal analysis is activation patching (Vig et al., 2020), where we observe the change in model behaviour when performing a counterfactual intervention to a component. However, this requires a separate forward pass for each component. Therefore,

<sup>4</sup>As described in detail in Marks et al. (2025).

we use attribution patching (AtP; Kramar et al., 2024) to estimate  $\hat{\text{IE}}$ , a linear approximation of the IE computed as a first-order Taylor expansion:

$$\hat{\text{IE}}_{\text{atp}}(m; \mathbf{a}; x_{\text{clean}}, x_{\text{patch}}) = \nabla_{\mathbf{a}} m|_{\mathbf{a}=\mathbf{a}_{\text{clean}}} (\mathbf{a}_{\text{patch}} - \mathbf{a}_{\text{clean}}) \quad (2)$$

Here,  $m$  is the target metric, usually defined as the logit difference between a correct and minimally differing incorrect token completion given context  $x_{\text{clean}}$ ,  $\mathbf{a}_{\text{clean}}$  is the activation of component  $\mathbf{a}$  given  $x_{\text{clean}}$ , and  $x_{\text{patch}}$  is a minimally differing context that changes the correct answer. Thus, the intuition is that we can approximate the causal contribution of  $\mathbf{a}$  to  $m$  by taking the slope of  $m$  with respect to  $\mathbf{a}$  (where this is approximated as the gradient of  $m$  at  $\mathbf{a}^5$ ) and multiplying this by the change in activations in  $\mathbf{a}$ .

AtP requires two forward passes and one backward pass to compute an estimate score for *all* components in parallel. In practice, we employ a more expensive but more accurate approximation based on integrated gradients (Sundararajan et al., 2017). For details, please see App. C.

**Probing classifiers.** Linear probes have frequently been used to locate representations of morphosyntactic features (e.g., Hewitt and Manning, 2019; Giulianelli et al., 2018; Chi et al., 2020). Given labeled classification data, we may train logistic regressions to map from the activations of frozen intermediate layers of pre-trained models (Belinkov, 2022) to the task labels. We similarly train classifiers to predict morphosyntactic concepts. We then quantify the extent of cross-lingual feature sharing for morphosyntactic concepts, and later evaluate the selectivity of targeted morphosyntactic concept interventions.

### 3 Multilingual Features

In this section, we investigate the extent to which morphosyntactic concepts are shared (or redundantly encoded) across languages. We first measure the overlap between the most causally influential sparse features for a given concept across languages. Then, we investigate whether ablating the multilingual features leads to a consistent decrease in classifier performance across languages.

<sup>5</sup>The gradient can be viewed as a local estimate of how much the metric we backpropagate from would be affected by changing the activation of the component.




Figure 2: Proportion of features shared across languages (intersection over union) among the top 32 features for each morphosyntactic concept. A significant fraction of the morphosyntactic concept representations are shared across languages in both Llama-3-8B and Aya-23-8B.

### 3.1 Experimental Setup

**Models.** We consider Llama-3-8B (Dubey et al., 2024) and Aya-23-8B (Aryabumi et al., 2024). While Llama 3 was trained primarily on English data, Aya was explicitly developed to support 23 typologically diverse and resource-diverse languages.

**Data.** We select 23 languages from Universal Dependencies 2.1 (UD; Nivre et al., 2017), a multilingual treebank containing dependency-parsed sentences. This corresponds to the 23 languages that Aya-23 was trained on (see App. B). Each word in each sentence in UD is annotated with its part of speech and morphosyntactic features, as defined in the UniMorph schema (Kirov et al., 2018).

**Training probing classifiers.** We train a probing classifier for each combination of morphosyntactic concept (e.g., gender) and language (e.g., English). We filter for morphosyntactic concepts that are annotated in at least six languages. During training, the inputs to the classifiers are residual stream activations from the end of the model’s middle layer ( $L = 16$ ), as this layer is expected to capture abstract, high-level and human-interpretable features (Elhage et al., 2022a; Lad et al., 2024). Following Marks et al. (2025), we pool non-padding token activations by summing them; then, we fit a logistic regression to the pooled representations.

**Training sparse autoencoders.** We train Gated SAEs (Rajamanoharan et al., 2024) for both Llama-3-8B and Aya-23-8B. We collect activations over 250 million tokens of The Pile (Gao et al., 2020) extracted from the residual stream at the same layer at which we train the probing classifiers. To verify the quality of the SAEs, we measure the proportion of loss recovered when replacing the model acti-

vations with SAE reconstructions thereof. While the training dataset is primarily composed of English texts, we verify that it also reconstructs other languages effectively using the mC4 dataset (Xue et al., 2021; see App. D).

### 3.2 Results

**Representations of morphosyntactic concepts are shared across languages.** First, we measure the extent of overlap across languages for the top features encoding a morphosyntactic category. For this, we perform attribution patching (§2), where the target metric is the logit of the probing classifiers we trained in §3.1. This identifies the set of causally relevant features that are most influential on the probe’s predictions. We observe that a small fraction of features dominate the causal effect, with a long tail of features that have a small effect. Therefore, for each language and concept, we select the top-32 most informative features per probe and compute the overlap across languages within a concept. We present the results in Figure 2.

We observe significant overlap of up to 50% in the features identified for the same grammatical concepts across typologically diverse languages. For example, feature 22860 is among the most influential features across *all* 15 languages that inflect for grammatical gender. The overlap is greater for morphosyntactic concepts present in a larger number of languages, such as grammatical number. We also find that the extent of overlap is largely consistent across Llama-3-8B and Aya-23-8B, with only a small number of exceptions (e.g. negative polarity). Thus, despite the lack of a one-to-one mapping of grammatical concepts across languages, there exist highly multilingual representations of abstract morphosyntactic concepts.





Figure 3: Examples of the activation patterns of selected features in Llama-3-8B that correspond to cross-lingual representations of grammatical concepts. For example, we locate features that indicate the presence of plural nouns or features that indicate past tense across languages.




Figure 4: Performance of the probing classifiers before and after ablating features. Specifically, we test ablating (a) all monolingual features, (b) all multilingual features, and (c) the upper quartile of the *most multilingual* features. We find that the classifiers crucially rely on massively multilingual features to predict the presence of morphosyntactic features.

**Effect of causal interventions on multilingual features are consistent across languages.** How important are multilingual features compared to monolingual features? We evaluate the importance of multilingual features by measuring the impact of ablating the multilingual features on the performance of the probing classifiers. We define multilingual features as those in the top feature set across at least two languages. The results (Figure 4) suggest that the classifiers crucially rely on multilingual features to predict the presence of morphosyntactic concepts. This suggests that large language models—including those trained primarily on English—learn to rely on shared representations to detect particular concepts, rather than relying on language-specific representations.

We also observe that while many features are shared across at least two languages, there is a small set of features shared across a large number of languages (App. E). Therefore, we investigate ablating only *massively multilingual* features. Specifically, we order features by the number of

languages for which they are causally relevant, and the average strength of their causal effect. Then, we select the upper quartile of features in this ranking—i.e., the most multilingual features—and only ablate those. For Llama 3, the performance after ablating only massively multilingual features drops to 64% on average. Notably, when ablating *all* multilingual features, the performance drops to 61%—a difference of only 3%, despite ablating four times as many features. This suggests that the most massively multilingual features explain most of the probes’ behavior.

**Feature overlap across concepts.** We now investigate the extent of feature overlap across grammatical concepts. This provides insights into the specificity of the functional roles of features. Specifically, we analyze the overlap in the sets of multilingual features associated with different concept-value pairs (e.g., masculine gender and singular number).<sup>6</sup> For each pair, we measure the fraction of multilingual features shared, and compute the mean overlap across all pairs.

We find that the mean overlap across concepts is 13.9% ( $\pm 10.7\%$ ). This relatively low (but significant) overlap is expected, as some linguistic concepts are jointly morphologically realized. For instance, singular number and masculine gender share 40.6% of their features; this is intuitive, given the frequent joint inflection of adjectives and pronouns for both number and gender. For example, in French and Hebrew, “*il*” and “*הוא*” are masculine singular pronouns, where changing either number *or* gender alters the form. Similarly, adjectives in these languages have distinct forms for combinations of grammatical gender and number.

<sup>6</sup>We use *concept* to refer to the variable—e.g., gender or number—and *concept-value* to refer to specific values that the variable can take—e.g., masculine gender or singular number.

### Multilingual features are human-interpretable.

One benefit of sparse features over neurons is their interpretability. This provides a qualitative way to verify that the features are truly relevant to the target concept, rather than related solely in some spurious correlational manner. Thus, we manually inspect a subset of the identified features, finding that many of them are intuitively meaningful and interpretable across different languages. We present the most multilingual features for two common concepts and a selection of their activation patterns across typologically diverse languages in Figure 3.

## 4 Demonstrating Functional Selectivity: A Case Study in Machine Translation

Another benefit of sparse features is the level of fine-grained control they provide over language model behaviors (Templeton et al., 2024). Indeed, feature-based interventions tend to be more effective than neuron-based methods at precisely modifying how language models generalize without destroying task performance (Marks et al., 2025). We leverage this property in a machine translation setting to (i) demonstrate the **functional selectivity** of the discovered morphosyntactic features (i.e., to show that intervening with them has few side effects on other concepts), and (ii) provide additional evidence that their causal role is highly **general** (i.e., that it is effective in this MT setting and not just on the probes). To evaluate (i) and (ii), we use feature steering (Templeton et al., 2024), where we clamp features of interest to counterfactual activations. These activations can be either significantly higher or lower than the values observed in practice. This approach is similar to prior work on understanding the functional role of individual neurons in image models (Bau et al., 2019).

As translation is a difficult task that does not always have a single clear answer, we restrict to a smaller subsample of concepts and languages such that we can manually verify the baseline quality of translations, and then the efficacy of the interventions. We restrict our analysis to tense (Past vs. Present), gender (Masc. vs. Fem.), and number (Sing. vs. Plural), as these concepts are shared across most languages we consider. We analyze English, French, and German; we additionally include Turkish, but restrict our focus to number and tense here, as Turkish does not inflect for grammatical gender. We provide details on the prompt format and translation performance in Appendix F.

## 4.1 Experimental Setup

**Data.** We design datasets consisting of minimal pairs of inputs that differ only with respect to the presence of a grammatical concept (see Table 3 in App. G). For example, we generate counterfactual pairs that elicit singular or plural verbs based on the grammatical number of the subject:

- a. The **parents** near the cars → **were**
- b. The **parent** near the cars → **was**

This is an adaptation and translation of data from Arora et al. (2024). Then, we check whether the probability of the grammatical sequence is higher for each pair. In our experiments, we only consider sentences where Llama-3-8B is capable of making grammatically correct predictions. We find that it is almost always correct.

**Feature intervention.** We follow Templeton et al. (2024) and intervene in the model by setting a single SAE feature to a multiple of its maximum activation value observed in the dataset.<sup>7</sup> Specifically, we decompose the residual stream activations  $\mathbf{x}$  into the sum of two components, the SAE reconstruction and the reconstruction error (see Eq. 1). Then, we replace the SAE term with a modified SAE reconstruction in which we set the activation value of a specific feature to a counterfactual value, but leave the error term unchanged. We then run the forward pass on the network using this modified residual stream. We only apply the intervention to the last token position of each generation step to prevent degenerate outputs.

To select the feature to intervene with, we manually inspect feature activation patterns. We select one feature per concept from the set of multilingual features identified in §3 where the activation values are most interpretable. Similar to Templeton et al. (2024), we explore different multiples and choose the one that works best empirically.

## 4.2 Quantitative Results

**Efficacy.** We define the efficacy of the intervention as the proportion of examples where the probe’s label flips after the intervention. To measure this, we generate a translation with and without the intervention. Then, in new forward passes with no interventions, we give the model these translations, pool their representations, and measure whether the probe’s (from §3.1) labels differ

<sup>7</sup>For this purpose, we record the maximum activations of each feature across 20M tokens from mC4 (Xue et al., 2021).




Figure 5: Efficacy in flipping the model behavior on our dataset when translating while intervening on a *single* multilingual feature per concept. For each concept, we translate a sentence containing some concept (e.g., present tense) from some source language to another language while intervening on a feature, and measure the number of times the model generates a translation containing the counterfactual concept value (e.g. past tense). In each setting, we intervene on a single feature and measure the success rate over 64 examples. Results aggregated across translation directions; see App. F.3 for separate results per direction.




Figure 6: Fraction of times an intervention on a grammatical concept exclusively impacts that concept, indicating the degree of cross-concept interference. We find that interference is generally very rare. Note that Turkish does not inflect for grammatical gender.

between the two translations. We find that intervening on a single feature is often sufficient to change the models prediction towards the intended concept (e.g., a masculine instead of feminine pronoun; see Figure 5). However, for some concepts, intervening on a single feature is often insufficient.

**Selectivity.** We also evaluate the selectivity of the interventions. Specifically, we want interventions to *only* flip the labels on the concept that we intervene on, and not others. We operationalise this using the probing classifiers trained in §3.1. As before, we generate two translations: one with and without feature interventions to a given concept. Then, we compute selectivity as the proportion of translation pairs for which none of the other concept probes changed their predictions. Our results (Figure 6) suggest that the interventions are generally selective to the concept.

### 4.3 Qualitative Results

The quantitative results provide insights into the generalisation of the identified features to a differ-

ent distribution, and allow us to quantify the success rate of the interventions. However, the dataset is neither naturalistic nor diverse. Therefore, we perform the same interventions on the Flores-101 dataset (Goyal et al., 2022), which provides aligned translations across 101 languages. We present a selection of qualitative examples in Figure 7.

However, we acknowledge that the intervention was not successful in some cases, and sometimes led to degenerate generations. Concept- and language-specific tuning over the activation values was also required. These limitations highlight common problems with feature steering methods, including that activations must often be scaled to significantly larger values than would naturally be observed (Templeton et al., 2024). This, combined with intervening on multiple features simultaneously, can result in unexpected or degenerate outputs (Shabalin et al., 2023).

## 5 Related Work

**Multilingual language modeling.** Multilingual language models are trained on multiple languages simultaneously, typically with an explicitly designed balance between languages in the pretraining corpus. For example, XLM-R (Conneau et al., 2020), M2M (Mohammadshahi et al., 2022), mT5 (Xue et al., 2021), and Aya (Aryabumi et al., 2024) are trained with corpora balanced across many languages. In contrast, Llama 3 (Dubey et al., 2024) is trained on a much more English-centric distribution: over 90% of the corpus is English text. Generally, the motivation behind balanced corpora during multilingual pretraining is to encourage cross-lingual transfer. We find that this balance may not be necessary for multilingual rep-

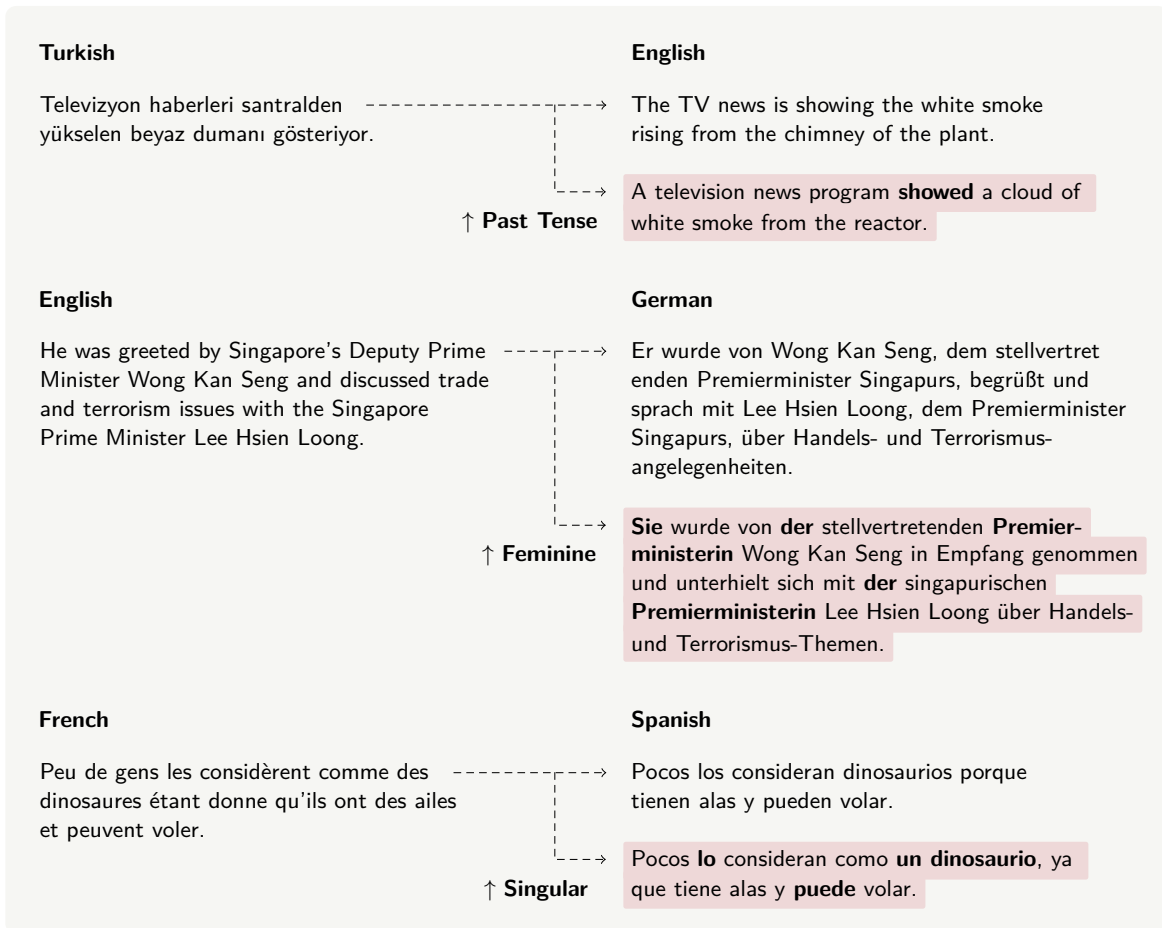


Figure 7: Comparison of translations with (highlighted in red) and without intervening on multilingual features encoding a morphosyntactic concept. The words indicating that the model flipped its behavior are highlighted in bold. The input texts are sampled from the Flores-101 dataset.

representations when the corpora are sufficiently large.

**Interpreting multilingual language models.** A central question in multilingual language modeling is whether LMs develop universal concept representations, disentangled from specific languages. Dhar and Bisazza (2021) study different forms of cross-lingual transfer in recurrent neural networks and found that exposing our LMs to a related language does not always increase grammatical knowledge in the target language. Ruzzetti et al. (2023) study the similarity of monolingual encoder models, finding higher similarity between models trained for typologically similar languages. Wendler et al. (2024) find that languages models trained on unbalanced, English-dominated corpora operate in a concept space that lies closer to English than to other languages. Ferrando (2024) find that mechanisms of subject-verb agreement are consistent across languages. Dumas et al. (2024) observe that concepts (e.g., “car”) are encoded disentangled from specific languages and can be transferred between

them. Similarly, Stanczak et al. (2022) find that probes trained on different languages read the same concepts from the same neurons in smaller-scale masked language models such as mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020), while Feng et al. (2024) observe that their propositional probes generalise to Spanish. De Varda and Marelli (2023) extend this study by investigating the cross-lingual consistency of individual neurons responding to syntactic phenomena. Our work extends these findings to large-scale models (including one not explicitly intended to be multilingual), and shows stronger causal evidence for the functional selectivity of these conceptual representations.

## 6 Discussion and Conclusion

Language models trained primarily on English data perform surprisingly well in other languages. Prior work has emphasized the importance of balanced pretraining corpora, so as to not overfit to a sin-



gle language (Conneau et al., 2020; Mohammadshahi et al., 2022); however, more recent studies suggest that the size of the pretraining corpus—not balance—may be more important for cross-lingual generalization (Jiang et al., 2023; Dubey et al., 2024; Schäfer et al.). Why is this effective? Our results provide evidence that large-scale pretraining, even with imbalanced corpora, induces equally cross-lingualistically generalizable grammatical abstractions as pretraining with balanced corpora. This is evidence in favor of a hypothesis discussed in Wendler et al. (2024): the internal *lingua franca* of large language models may not be English words *per se*, but rather concepts. That said, it is still reasonable to assume these concepts are likely biased toward how English handles them in models pre-trained on imbalanced corpora.

How generalizable should grammatical concept representations be? While it is more parameter-efficient to share representations across languages, this may lead to biases if two languages assign different social or semantic connotations to—or simply distribute differently—the same grammatical concepts. Indeed, there is not a one-to-one relationship between concepts across languages: different languages may have differing numbers of values for the same concept (e.g., Finnish has many more grammatical cases than German), and may use the same categories in different ways (e.g., the same nouns often have different genders in different languages, even within language families). More language-specific features could lead to less bias, but would require more parameters to represent. It is not clear what the optimal point between these two ends of generalization are.

More practically, the extent to which concepts are shared across languages has implications for multilingual downstream NLP tasks. For example, it helps explain the phenomenon that preference tuning for toxicity mitigation in a single language generalizes to other languages (Li et al., 2024). Thus, we might reasonably expect various types of interventions and model editing approaches to generalize from a single language to other languages.

## Limitations

**Steering with SAEs.** We use SAE features to intervene on the models internal state. Similar to Templeton et al. (2024), we find that steering with SAEs typically requires clamping feature activations to values outside their observed range over

the training dataset. However, clamping feature activations to extreme values often causes the model to produce nonsensical generations, e.g., repeating the same token indefinitely. SAEs have also been shown to perform worse for fine-grained steering model behavior than alternatives (Wu et al., 2025), such as difference-in-means (Marks and Tegmark, 2024). Thus, while our results demonstrate that the identified features can be causally responsible for model behavior, optimally steering model behavior is an open problem that future work should address.

**Distinction between model understanding and generation.** In our experiments, the linear probes measure whether it is possible to linearly separate grammatical concepts in the model’s representations, whereas the machine translation experiments measure the effect of feature directions on model generation. Past work (Meng et al., 2022; Orgad et al., 2025; Gottesman and Geva, 2024) has observed a distinction between a model encoding a concept versus being able to use that concept during generation (i.e., a “knowing vs. saying” distinction). Anecdotally, we observed that some features with a high causal effect on the linear probe do not necessarily have a similar effect on the model’s generation. We hypothesize that this is due to an intrinsic distinction between concept understanding/representation and concept generation.

**Human biases in feature interpretation.** One must be cautious in applying human concepts to LM representations. It is probable that LMs deploy distinct concept spaces from humans; thus, human explanations of neurons or features are likely to be biased in ways that can result in suboptimal predictions of when the neuron or feature will activate (cf. Huang et al., 2023).

**Non-linear features.** Sparse autoencoders are a method for unsupervised search for features. However, they will only recover a feature if it is *linearly encoded*. Thus, if a concept is encoded in a non-linear manner, then sparse features may struggle to recover the full range of that variable. Thus, sparse features are generally best-suited to encoding binary relations, whereas some features may be encoded in more complex arrangements such as circular shapes (cf. Engels et al. 2025; Csordás et al. 2024).

## Acknowledgements

Jannik Brinkmann is supported by the German Federal Ministry for Economic Affairs and a grant from the Long-Term Future Fund. Chris Wendler is supported by a grant from Open Philanthropy. Aaron Mueller is supported by a postdoctoral fellowship under the Zuckerman STEM Leadership Program.

## References

- Aryaman Arora, Dan Jurafsky, and Christopher Potts. 2024. [CausalGym: Benchmarking causal interpretability methods on linguistic tasks](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14638–14663, Bangkok, Thailand. Association for Computational Linguistics.
- Viraat Aryabumi, John Dang, Dwarak Talupuru, Saurabh Dash, David Cairuz, Hangyu Lin, Bharat Venkitesh, Madeline Smith, Kelly Marchisio, Sebastian Ruder, Acyr Locatelli, Julia Kreutzer, Nick Frosst, Phil Blunsom, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. [Aya 23: Open weight releases to further multilingual progress](#). *Preprint*, arXiv:2405.15032.
- David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. 2019. [GAN dissection: Visualizing and understanding generative adversarial networks](#). In *International Conference on Learning Representations*.
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Tolga Bolukbasi, Adam Pearce, Ann Yuan, Andy Coenen, Emily Reif, Fernanda Viégas, and Martin Wattenberg. 2021. [An interpretability illusion for BERT](#). *Preprint*, arXiv:2104.07143.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. 2023. [Towards monosemanticity: Decomposing language models with dictionary learning](#). *Transformer Circuits Thread*.
- Elisa Cargnelutti, Barbara Tomasino, and Franco Fabbro. 2019. [Language brain representation in bilinguals with different age of appropriation and proficiency of the second language: A meta-analysis of functional imaging studies](#). *Frontiers in human neuroscience*, 13:445658.
- Tyler Chang, Zhuowen Tu, and Benjamin Bergen. 2022. [The geometry of multilingual language model representations](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 119–136, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ethan A. Chi, John Hewitt, and Christopher D. Manning. 2020. [Finding universal grammatical relations in multilingual BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5577, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Albert Costa and Núria Sebastián-Gallés. 2014. [How does the bilingual experience sculpt the brain?](#) *Nature Reviews Neuroscience*, 15(5):336–345. Research Support, Non-U.S. Gov't; Review.
- Róbert Csordás, Christopher Potts, Christopher D Manning, and Atticus Geiger. 2024. [Recurrent neural networks learn to store and generate sequences using non-linear representations](#). In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 248–262, Miami, Florida, US. Association for Computational Linguistics.
- Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, Robert Huben, and Lee Sharkey. 2024. [Sparse autoencoders find highly interpretable features in language models](#). In *The Twelfth International Conference on Learning Representations*.
- Andrea Gregor De Varda and Marco Marelli. 2023. [Data-driven cross-lingual syntax: An agreement study with massively multilingual models](#). *Computational Linguistics*, 49(2):261–299.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Prajit Dhar and Arianna Bisazza. 2021. [Understanding cross-lingual syntactic transfer in multilingual recurrent neural networks](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 74–85, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Alonso, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Milon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Celebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raporthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whit-

ney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papanikos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhota, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks,

- Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiao Cheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The Llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Clément Dumas, Veniamin Veselovsky, Giovanni Monea, Robert West, and Chris Wendler. 2024. [How do Llamas process multilingual text? A latent exploration through the patchscope](#). In *ICML 2024 Workshop on Mechanistic Interpretability*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Neel Nanda, Tom Henighan, Scott Johnston, Sheer ElShowk, Nicholas Joseph, Nova DasSarma, Ben Mann, Danny Hernandez, Amanda Askell, Kamal Ndousse, Andy Jones, Dawn Drain, Anna Chen, Yuntao Bai, Deep Ganguli, Liane Lovitt, Zac Hatfield-Dodds, Jackson Kernion, Tom Conerly, Shauna Kravec, Stanislav Fort, Saurav Kadavath, Josh Jacobson, Eli Tran-Johnson, Jared Kaplan, Jack Clark, Tom Brown, Sam McCandlish, Dario Amodei, and Christopher Olah. 2022a. [Softmax linear units](#). *Transformer Circuits Thread*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. 2022b. Toy models of superposition. *Transformer Circuits Thread*.
- Joshua Engels, Eric J Michaud, Isaac Liao, Wes Gurnee, and Max Tegmark. 2025. [Not all language model features are linear](#). In *The Thirteenth International Conference on Learning Representations*.
- Jiahai Feng, Stuart Russell, and Jacob Steinhardt. 2024. Monitoring latent world states in language models with propositional probes. *arXiv preprint arXiv:2406.19501*.
- Javier Ferrando. 2024. [On the similarity of circuits across languages: a case study on the subject-verb agreement task](#). In *ICML 2024 Workshop on Mechanistic Interpretability*.
- Jaden Fried Fiotto-Kaufman, Alexander Russell Lof-tus, Eric Todd, Jannik Brinkmann, Koyena Pal, Dmitrii Troitskii, Michael Ripa, Adam Belfki, Can Rager, Caden Juang, Aaron Mueller, Samuel Marks, Arnab Sen Sharma, Francesca Lucchetti, Nikhil Prakash, Carla E. Brodley, Arjun Guha, Jonathan Bell, Byron C Wallace, and David Bau. 2025. [NNSight and NDIF: Democratizing access to foundation model internals](#). In *The Thirteenth International Conference on Learning Representations*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The Pile: An 800GB dataset of diverse text for language modeling](#). *Preprint*, arXiv:2101.00027.
- Mario Giulianelli, Jack Harding, Florian Mohnert, Dieuwke Hupkes, and Willem Zuidema. 2018. [Under the hood: Using diagnostic classifiers to investigate and improve how language models track agreement information](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 240–248, Brussels, Belgium. Association for Computational Linguistics.
- Daniela Gottesman and Mor Geva. 2024. [Estimating knowledge in large language models without generating a single token](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3994–4019, Miami, Florida, USA. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). volume 10, pages 522–538, Cambridge, MA. MIT Press.
- Ashim Gupta and Vivek Srikumar. 2021. [X-fact: A new benchmark dataset for multilingual fact checking](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 675–682, Online. Association for Computational Linguistics.



- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- G. E. Hinton, J. L. McClelland, and D. E. Rumelhart. 1986. *Distributed representations*, page 77–109. MIT Press, Cambridge, MA, USA.
- Tianze Hua, Tian Yun, and Ellie Pavlick. 2024. [mOthello: When do cross-lingual representation alignment and cross-lingual transfer emerge in multilingual models?](#) In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1585–1598, Mexico City, Mexico. Association for Computational Linguistics.
- Jing Huang, Atticus Geiger, Karel D’Oosterlinck, Zhengxuan Wu, and Christopher Potts. 2023. [Rigorously assessing natural language explanations of neurons](#). In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 317–331, Singapore. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Adam Karvonen, Benjamin Wright, Can Rager, Rico Angell, Jannik Brinkmann, Logan Riggs Smith, Claudio Mayrink Verdun, David Bau, and Samuel Marks. 2024. [Measuring progress in dictionary learning for language model interpretability with board game models](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, G eraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sabrina J. Mielke, Arya McCarthy, Sandra K ubler, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. [UniMorph 2.0: Universal Morphology](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Janos Kramar, Tom Lieberum, Rohin Shah, and Neel Nanda. 2024. [AtP\\*: An efficient and scalable method for localizing LLM behaviour to components](#). *Preprint*, arXiv:2403.00745.
- Vedang Lad, Wes Gurnee, and Max Tegmark. 2024. [The remarkable robustness of LLMs: Stages of inference?](#) In *ICML 2024 Workshop on Mechanistic Interpretability*.
- Xiaochen Li, Zheng Xin Yong, and Stephen Bach. 2024. [Preference tuning for toxicity mitigation generalizes across languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13422–13440, Miami, Florida, USA. Association for Computational Linguistics.
- Jindřich Libovick y, Rudolf Rosa, and Alexander Fraser. 2020. [On the language neutrality of pre-trained multilingual representations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1663–1674, Online. Association for Computational Linguistics.
- Aleksandar Makelov, Georg Lange, and Neel Nanda. 2024. [Towards principled evaluations of sparse autoencoders for interpretability and control](#). In *ICLR 2024 Workshop on Secure and Trustworthy Large Language Models*.
- Sam Marks and Aaron Mueller. 2024. [Dictionary learning](#). [https://github.com/saprmarks/dictionary\\_learning/tree/main](https://github.com/saprmarks/dictionary_learning/tree/main).
- Samuel Marks, Can Rager, Eric J Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. 2025. [Sparse feature circuits: Discovering and editing interpretable causal graphs in language models](#). In *The Thirteenth International Conference on Learning Representations*.
- Samuel Marks and Max Tegmark. 2024. [The geometry of truth: Emergent linear structure in large language model representations of true/false datasets](#). In *First Conference on Language Modeling*.
- Kevin Meng, David Bau, Alex J Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems*.
- Alireza Mohammadshahi, Vassilina Nikoulina, Alexandre Berard, Caroline Brun, James Henderson, and Laurent Besacier. 2022. [SMaLL-100: Introducing shallow multilingual machine translation model for low-resource languages](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8348–8359, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Aaron Mueller, Jannik Brinkmann, Millicent Li, Samuel Marks, Koyena Pal, Nikhil Prakash, Can Rager, Aruna Sankaranarayanan, Arnab Sen Sharma, Jiuding Sun, Eric Todd, David Bau, and Yonatan Belinkov. 2024. [The quest for the right mediator: A history, survey, and theoretical grounding of causal interpretability](#). *Preprint*, arXiv:2408.01416.
- Joakim Nivre,  eljko Agi c, Lars Ahrenberg, Lene Antonsen, Maria Jesus Aranzabe, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, John Bauer, Kepa Ben-goetxea, Riyaz Ahmad Bhat, Eckhard Bick, Victoria

- Bobicev, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Aljoscha Burchardt, Marie Candito, Gauthier Caron, Gülşen Cebiroğlu Eryiğit, Giuseppe G. A. Celano, Savas Cetin, Fabricio Chalub, Jinho Choi, Silvie Cinková, Çağrı Çöltekin, Miriam Connor, Elizabeth Davidson, Marie-Catherine de Marnette, Valeria de Paiva, Arantza Diaz de Ilarraza, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Marhaba Eli, Ali Elkahky, Tomaž Erjavec, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökirmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta Gonzáles Saavedra, Matias Grioni, Normunds Grūzītis, Bruno Guillaume, Nizar Habash, Jan Hajič, Jan Hajič jr., Linh Hà Mỹ, Kim Harris, Dag Haug, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Radu Ion, Elena Irimia, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Hüner Kaşıkara, Hiroshi Kanayama, Jenna Kanerva, Tolga Kayadelen, Václava Kettnerová, Jesse Kirchner, Natalia Kotsyba, Simon Krek, Veronika Laipala, Lorenzo Lambertino, Tatiana Lando, John Lee, Phng Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, Nikola Ljubešić, Olga Loginova, Olga Lyahevskaya, Teresa Lynn, Vivien Macketanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Cătălina Măranduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Gustavo Mendonça, Niko Miekka, Anna Missilä, Cătălin Mititelu, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Shinsuke Mori, Bohdan Moskalevskyi, Kadri Muischnek, Kaili Müürisep, Pinkey Nainwani, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Lng Nguyễn Thị, Huyèn Nguyễn Thị Minh, Vitaly Nikolaev, Hanna Nurmi, Stina Ojala, Petya Osenova, Robert Östling, Lilja Övrelid, Elena Pascual, Marco Passarotti, Cenel-Augusto Perez, Guy Perrier, Slav Petrov, Jussi Pitulainen, Emily Pitler, Barbara Plank, Martin Popel, Lauma Pretkalniņa, Prokopis Prokopidis, Tiina Puolakainen, Sampo Pyysalo, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Larissa Rinaldi, Laura Rituma, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Benoît Sagot, Shadi Saleh, Tanja Samardžić, Manuela Sanguinetti, Baiba Saulīte, Sebastian Schuster, Djamel Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Zsolt Szántó, Dima Taji, Takaaki Tanaka, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Uřešová, Larraitz Uribe, Hans Uszkoreit, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villamonte de la Clergerie, Veronika Vincze, Lars Wallin, Jonathan North Washington, Mats Wirén, Tak-sum Wong, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, and Hanzhi Zhu. 2017. [Universal dependencies 2.1](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. 2025. [LLMs know more than they show: On the intrinsic representation of LLM hallucinations](#). In *The Thirteenth International Conference on Learning Representations*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Judea Pearl. 2001. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, UAI'01*, page 411–420, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Monika M. Połczyńska and Susan Y. Bookheimer. 2021. [General principles governing the amount of neuroanatomical overlap between languages in bilinguals](#). *Neuroscience Biobehavioral Reviews*, 130:1–14.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Alessandro Raganato, Tommaso Pasini, Jose Camacho-Collados, and Mohammad Taher Pilehvar. 2020. [XLWiC: A multilingual benchmark for evaluating semantic contextualization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7193–7206, Online. Association for Computational Linguistics.
- Senthooran Rajamanoharan, Arthur Conmy, Lewis Smith, Tom Lieberum, Vikrant Varma, Janos Kra-mar, Rohin Shah, and Neel Nanda. 2024. [Improving](#)

- sparse decomposition of language model activations with gated sparse autoencoders. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Elena Sofia Ruzzetti, Federico Ranaldi, Felicia Logozzo, Michele Mastromattei, Leonardo Ranaldi, and Fabio Massimo Zanzotto. 2023. Exploring linguistic properties of monolingual BERTs with typological classification among languages. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14447–14461, Singapore. Association for Computational Linguistics.
- Anton Schäfer, Shauli Ravfogel, Thomas Hofmann, Tiago Pimentel, and Imanol Schlag. Language imbalance can boost cross-lingual generalisation. *Preprint*, arXiv:2404.07982.
- Stepan Shabalín, Dmitrii Kharlapenko, Arthur Conmy, and Neel Nanda. 2023. SAE features for refusal and sycophancy steering vectors. LessWrong.
- Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423.
- P. Smolensky. 1986. *Neural and conceptual interpretation of PDP models*, page 390–431. MIT Press, Cambridge, MA, USA.
- Karolina Stanczak, Edoardo Ponti, Lucas Torroba Henning, Ryan Cotterell, and Isabelle Augenstein. 2022. Same neurons, different languages: Probing morphosyntax in multilingual pre-trained models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1589–1598, Seattle, United States. Association for Computational Linguistics.
- Simone Sulpizio, Nicola Del Maschio, Davide Fedeli, and Jubin Abutalebi. 2020. Bilingual language processing: A meta-analysis of functional neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, 108:834–853.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 3319–3328. JMLR.org.
- Aaquib Syed, Can Rager, and Arthur Conmy. 2024. Attribution patching outperforms automated circuit discovery. In *The 7th BlackboxNLP Workshop*.
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. 2024. Scaling monosemanticity: Extracting interpretable features from Claude 3 Sonnet. *Transformer Circuits Thread*.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. In *Advances in Neural Information Processing Systems*, volume 33, pages 12388–12401. Curran Associates, Inc.
- 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.
- Benjamin Wright and Lee Sharkey. 2024. Addressing feature suppression in sparse autoencoders. LessWrong post.
- Zhengxuan Wu, Aryaman Arora, Atticus Geiger, Zheng Wang, Jing Huang, Dan Jurafsky, Christopher D. Manning, and Christopher Potts. 2025. AxBench: Steering LLMs? Even simple baselines outperform sparse autoencoders. *Preprint*, arXiv:2501.17148.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

## A Implementation Details

### A.1 Infrastructure

The experiments were run on a single server with 8 NVIDIA RTX A6000 48 GB GPUs with CUDA Version 12.4 and an AMD EPYC 7413 24-Core Processor. The total runtime for training the probes and autoencoders was less than a week.

### A.2 Libraries

To extract activations and intervene on the model, we use `nnsight` (Fiotto-Kaufman et al., 2025). For training the linear probes, we use `scikit-learn` (Pedregosa et al., 2011). For training the SAEs, we use `dictionary-learning` (Marks and Mueller, 2024), itself based on `nnsight`.

## B Languages

In our experiments, we consider the set of languages which Aya-23 (Aryabumi et al., 2024) was trained on. This includes: Arabic, Chinese, Czech, Dutch, English, French, German, Greek, Hebrew,



Hindi, Indonesian, Italian, Japanese, Korean, Persian, Polish, Portuguese, Romanian, Russian, Spanish, Turkish, Ukrainian, and Vietnamese.

## C Attribution Patching

Here, we elaborate on activation patching, attribution patching, and the modified attribution patching approach we use in our experiments.

*Activation patching* is a method for quantifying the causal importance of a model component on a model behavior. It involves replacing the activation of some component in the model during one model forward pass with the activations of the same component from a different forward pass. This requires constructing a distribution  $\mathcal{D}$  over pairs of inputs  $(x_{\text{clean}}, x_{\text{patch}})$ , where  $x_{\text{clean}}$  is a prompt on which the behaviour occurs, and  $x_{\text{patch}}$  is a reference prompt on which the behaviour does not occur. For example, to study subject-verb agreement one could construct a datasets of pairs such as:

- a. The **dog**  $\rightarrow$  **barks**
- b. The **dogs**  $\rightarrow$  **bark**

Then, we can define the contribution of a node  $n \in N$ , where  $N$  is the number of components in the model, to the model’s behaviour as the counterfactual expected impact of replacing that node on the clean prompt with its value on the patch prompt (Meng et al., 2022). In practice, performing activation patching for all components requires  $\mathcal{O}(N)$  forward passes, which does not scale efficiently (Mueller et al., 2024). Therefore, we use a faster linear approximation of the contribution of a node; this is *attribution patching*, which computes a first-order Taylor expansion of this metric as described in §2.

In our experiments, we use the attribution patching approach proposed by Marks et al. (2025) which employs a more expensive but more accurate approximation based on integrated gradients:

$$\hat{\text{IE}}_{\text{ig}}(m; \mathbf{a}_{\text{clean}}, \mathbf{a}_{\text{patch}}) = \left( \sum_{\alpha} \nabla_{\mathbf{a}} m|_{\alpha \mathbf{a}_{\text{clean}} + (1-\alpha) \mathbf{a}_{\text{patch}}} \right) (\mathbf{a}_{\text{patch}} - \mathbf{a}_{\text{clean}}) \quad (3)$$

where the sum in (3) ranges over  $K = 10$  equally-spaced  $\alpha \in \{0, \frac{1}{K}, \dots, \frac{K-1}{K}\}$ . While this cannot be done in parallel for two nodes when one is downstream of another, it can be done in parallel for arbitrarily many nodes which do not depend on each other. Thus the additional cost of computing  $\hat{\text{IE}}_{\text{ig}}$  over  $\hat{\text{IE}}_{\text{atp}}$  scales linearly in  $K$ .

## D Sparse Autoencoders

### D.1 Gated Sparse Autoencoders

The Gated SAE architecture, as introduced in Bricken et al. (2023) is defined by encoder weights  $W_e \in \mathbb{R}^{m \times n}$ , decoder weights  $W_d \in \mathbb{R}^{n \times m}$  with columns constrained to have a  $L_2$ -norm of 1, and biases  $\mathbf{b}_e \in \mathbb{R}^m$ ,  $\mathbf{b}_d \in \mathbb{R}^n$ . Given an input  $\mathbf{x} \in \mathbb{R}^n$ , the SAE computes

$$\mathbf{f}(\mathbf{x}) = \text{ReLU}(W_e(\mathbf{x} - \mathbf{b}_d) + \mathbf{b}_e) \quad (4)$$

$$\hat{\mathbf{x}} = W_d \mathbf{f}(\mathbf{x}) + \mathbf{b}_d \quad (5)$$

where  $\mathbf{f}(\mathbf{x})$  is the vector of feature activations, and  $\hat{\mathbf{x}}$  is the reconstruction. This model is trained using the following loss term:

$$\mathcal{L}_{\text{standard}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\text{train}}} \left[ \|\mathbf{x} - \hat{\mathbf{x}}\|_2 + \lambda \|\mathbf{f}(\mathbf{x})\|_1 \right]. \quad (6)$$

for some hyperparameter  $\lambda > 0$  controlling sparsity.

The  $L_1$  penalty introduces biases that can harm the accuracy of the reconstruction, as the loss can be decreased by trading-off reconstruction for lower L1 (Wright and Sharkey, 2024). To address this, Rajamanoharan et al. (2024) introduced a architectural modification that separates (i) the selection of dictionary elements to use in a reconstruction, and (ii) estimating the coefficients of these elements. This results in the following gated architecture:

$$\pi_{\text{gate}}(\mathbf{x}) := W_{\text{gate}}(\mathbf{x} - \mathbf{b}_d) + \mathbf{b}_{\text{gate}}$$

$$\tilde{\mathbf{f}}(\mathbf{x}) := \mathbb{I}[\pi_{\text{gate}}(\mathbf{x}) > 0] \odot$$

$$\text{ReLU}(W_{\text{mag}}(\mathbf{x} - \mathbf{b}_d) + \mathbf{b}_{\text{mag}})$$

$$\hat{\mathbf{x}}(\tilde{\mathbf{f}}(\mathbf{x})) = W_d \tilde{\mathbf{f}}(\mathbf{x}) + \mathbf{b}_d$$

where  $\mathbb{I}[\cdot > 0]$  is the Heaviside step function and  $\odot$  denotes elementwise multiplication. Then, the loss function uses  $\hat{\mathbf{x}}_{\text{frozen}}$ , a frozen copy of the decoder:

$$\mathcal{L}_{\text{gated}} := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\text{train}}} \left[ \|\mathbf{x} - \hat{\mathbf{x}}(\tilde{\mathbf{f}}(\mathbf{x}))\|_2^2 + \lambda \|\text{ReLU}(\pi_{\text{gate}}(\mathbf{x}))\|_1 + \|\mathbf{x} - \hat{\mathbf{x}}_{\text{frozen}}(\text{ReLU}(\pi_{\text{gate}}(\mathbf{x})))\|_2^2 \right] \quad (7)$$

### D.2 Training Parameters

### D.3 Performance Across Languages

We evaluate the performance of the SAE trained on Llama-3-8B using the multilingual C4 (mC4) dataset (Xue et al., 2021), itself a multilingual extension of the C4 dataset (Raffel et al., 2020). Table



Parameter	Value
Optimiser	Adam
Learning Rate	0.001
L1 Coefficient	0.505
Expansion Factor	8
Number of Token	250 Million
Batch Size	512
Warmup Steps	1,000

Table 1: Training parameters of our sparse autoencoders for both Llama-3-8B and Aya-23-8B.

2 presents the *loss recovered* when replacing the residual stream activations with the SAE reconstructions across a range of languages. For this purpose, we compute the loss recovered as

$$\mathcal{L}_{\text{recovered}} = \frac{\mathcal{L}_{\text{reconstructed}} - \mathcal{L}_{\text{zero}}}{\mathcal{L}_{\text{original}} - \mathcal{L}_{\text{zero}}} \quad (8)$$

where  $\mathcal{L}_{\text{original}}$  is the cross entropy (CE) loss of the language model without intervention,  $\mathcal{L}_{\text{zero}}$  is the CE loss of the model when zero ablating the activations, and  $\mathcal{L}_{\text{reconstructed}}$  is the loss when replacing the activations with the SAE reconstructions. This is a standard extrinsic metric to measure autoencoder quality (Mueller et al., 2024). However, it does not measure the intrinsic interpretability of the SAE features, which is not possible to measure in the absence of ground truth features (Karvonen et al., 2024; Makelov et al., 2024).

Language	Loss Recovered
Chinese	0.8122
Dutch	0.9004
English	0.9257
Finnish	0.8582
French	0.9285
German	0.9026
Indonesian	0.9205
Japanese	0.8759
Turkish	0.8522

Table 2: Evaluation of the sparse autoencoder trained on Llama-3-8B across different subsets on the AllenAI C4 dataset. The table shows the fraction of loss recovered when replacing residual stream activations with sparse autoencoders reconstructions across nine languages.

## E Features Across Languages

We measured the number of languages for which the top-32 features for each concept were strongly influential. Across grammatical concepts, the number of languages that a feature is influential for roughly follows a power law distribution. Below, we present the distributions for a selection of common concepts (see Figures 8, 9, 10). We find that for most concepts, there are multiple features that are shared across nearly all languages. Note that the maximum number of languages varies across concepts, as not all languages inflect for a given concept nor are annotated for it in Universal Dependencies.




Figure 8: Distribution of the number of languages across which a given feature associated with masculine gender is shared (Llama 3).




Figure 9: Distribution of the number of languages across which a given feature associated with past tense is shared (Llama 3).




Figure 10: Distribution of the number of languages across which a given feature associated with accusative case is shared (Llama 3).




Figure 11: Distribution of the number of languages across which a given feature associated with masculine gender is shared (Aya-23).




Figure 12: Distribution of the number of languages across which a given feature associated with past tense is shared (Aya-23).




Figure 13: Distribution of the number of languages across which a given feature associated with accusative case is shared (Aya-23).

## F Machine Translation Details

### F.1 Prompt Format

We present an example of our 2-shot translation prompt in Figure 14.

### F.2 Translation Performance

We evaluate the ability of Llama-3-8B and Aya-23-8B to translate languages in-context using the Flores-101 dataset (Goyal et al., 2022). Specifically, we measure the BLEU score (Papineni et al., 2002) using the SacreBLEU implementation (Post, 2018). We prompt models with 2 ground-truth exemplar translations, and then give it the test sentence in the source language (see Figure 14). Re-

sults are presented in Figure 15 and 16.

The governor’s office said nineteen of the injured were police officers. \\ Das Gouverneursamt erklärte, dass von den Verletzten neun Polizisten waren.

He produced over 1,000 stamps for Sweden and 28 other countries. \\ Er produzierte über 1.000 Briefmarken für Schweden und 28 weitere Länder.

The truck driver, who is aged 64, was not injured in the crash. \\

Figure 14: Example 2-shot translation prompt between English and German. Note that the double backslashes are *literals*, not newlines. We found this prompt format to work best empirically in 2-shot settings in initial experiments.



Figure 15: SacreBLEU scores of Llama-3-8B.




Figure 16: SacreBLEU scores of Aya-23-8B.




Figure 17: Efficacy in flipping the model behavior on our counterfactual dataset when translating between languages and intervening on a *single* multilingual feature per concept. For each concept, we translate a sentence from some source language (e.g. German) where some concept (e.g. present tense) to another language (e.g. English) and measure the number of times the model starts to predict some alternative concept value (e.g. past tense). In each of these settings, we intervene on a single feature and measure the success rate over 64 examples.

### F.3 Directional Translation Results

In Fig. 5, we present averages for language pairs translating in both directions (e.g., English to French and French to English). Here, we provide additional details about the results for each translation direction separately (see Fig. 17).

### G Counterfactual Dataset

Our counterfactual dataset considers three concepts (gender, number, tense) and three languages (English, French, German). The dataset format is inspired by CausalGym (Arora et al., 2024) and the templates are in part taken from that dataset. We present example counterfactuals in Table 3.

Morphosyntactic Concept	Language	Example Input	Example Output
Gender (Masc. vs. Fem.)	English	When <b>Sam</b> was at the park	<b>he</b>
		When <b>Sarah</b> was at the park	<b>she</b>
	French	<b>Louis</b> a accepté, parce	<b>qu’il</b>
		<b>Charlotte</b> a accepté, parce	<b>qu’elle</b>
Number (Sing. vs. Plur.)	German	<b>Lukas</b> lachte, weil	<b>er</b>
		<b>Hannah</b> lachte, weil	<b>sie</b>
	English	The <b>scientist</b>	<b>discovers</b>
		The <b>scientists</b>	<b>discover</b>
Tense (Pres. vs. Past)	French	<b>Le docteur</b>	<b>est</b>
		<b>Les docteurs</b>	<b>sont</b>
	German	<b>Der Schauspieler</b>	<b>ist</b>
		<b>Die Schauspieler</b>	<b>sind</b>
Tense (Pres. vs. Past)	Turkish	<b>Hemşire</b> sonraki adımlardan emin	<b>değil</b>
		<b>Hemşireler</b> sonraki adımlardan emin	<b>değiller</b>
	English	When Mary <b>is</b> in the bed, she	<b>sleeps</b>
		When Mary <b>was</b> in the bed, she	<b>slept</b>
Tense (Pres. vs. Past)	French	Louis <b>marche</b> parce qu’il	<b>est</b>
		Louis <b>marchait</b> parce qu’il	<b>était</b>
	German	Lukas <b>geht</b> , weil er	<b>ist</b>
		Lukas <b>ging</b> , weil er	<b>war</b>
Tense (Pres. vs. Past)	Turkish	Hasan <b>arıyor</b> çünkü o	<b>hastadır</b>
		Hasan <b>aradı</b> çünkü o	<b>hastaydı</b>

Table 3: Examples from our counterfactual dataset.