

Causal Analysis of Syntactic Agreement Neurons in Multilingual Language Models

Aaron Mueller[†], Yu Xia[‡], Tal Linzen[‡]

[†]Johns Hopkins University [‡]New York University
amueller@jhu.edu, yx1675@nyu.edu, linzen@nyu.edu

Abstract

Structural probing work has found evidence for latent syntactic information in pre-trained language models. However, much of this analysis has focused on monolingual models, and analyses of multilingual models have employed correlational methods that are confounded by the choice of probing tasks. In this study, we causally probe multilingual language models (XGLM and multilingual BERT) as well as monolingual BERT-based models across various languages; we do this by performing counterfactual perturbations on neuron activations and observing the effect on models’ subject-verb agreement probabilities. We observe where in the model and to what extent syntactic agreement is encoded in each language. We find significant neuron overlap across languages in autoregressive multilingual language models, but not masked language models. We also find two distinct layer-wise effect patterns and two distinct sets of neurons used for syntactic agreement, depending on whether the subject and verb are separated by other tokens. Finally, we find that behavioral analyses of language models are likely underestimating how sensitive masked language models are to syntactic information.

1 Introduction

Syntactic information is necessary for robust generalization in natural language processing tasks (for a case study using the natural language inference task, see McCoy et al. 2019). The success of pre-trained language models (LMs) such as RoBERTa (Liu et al., 2019) and GPT-3 (Brown et al., 2020) in many NLP tasks has prompted hypotheses that they accomplish their performance through structural representations induced during pre-training, rather than only lexical or positional representations (Manning et al., 2020); behavioral evidence for LMs’ syntactic abilities has been found in masked LMs (MLMs; Warstadt et al., 2020; Warstadt and Bowman, 2020; Goldberg, 2019) and

autoregressive LMs (ALMs; Hu et al., 2020). Evidence for structural representations has been reported for multilingual pre-trained LMs (Goldberg, 2019; Mueller et al., 2020) and in sequence-to-sequence models (Mueller et al., 2022).

Despite efforts to understand the structural information encoded by pre-trained LMs (Hewitt and Manning, 2019; Chi et al., 2020; Elazar et al., 2021; Ravfogel et al., 2021; Finlayson et al., 2021; *inter alia*), it remains unclear how and where multilingual models encode this information. Most multilingual probing studies are *correlational* and use dependency parsing or labeling as a proxy task indicative of syntactic information (Chi et al., 2020; Stańczak et al., 2022). This is problematic: Models do not need structural or word order information to achieve high performance on dependency labeling (Sinha et al., 2021), and training a parametric probing classifier introduces many confounds (Hewitt and Liang, 2019; Antverg and Belinkov, 2022).

Causal probing, however, enables non-parametric analyses of models through counterfactual interventions on inputs or model representations. Causal probing studies have argued for the existence of specific syntactic agreement neurons and units in neural language models (Finlayson et al., 2021; Lakretz et al., 2019; De Cao et al., 2021), but these studies have focused on monolingual models—usually (though not always) in English. Causal methods allow us to make stronger arguments about where and how syntactic agreement is performed in pre-trained LMs, and we can apply them to answer questions about the language specificity and construction specificity of syntactic agreement neurons.

In this study, we extend causal mediation analysis (Pearl, 2001; Robins, 2003; Vig et al., 2020) to multilingual language models, including an autoregressive LM and a masked LM. We also analyze a series of monolingual MLMs across languages. We employ the syntactic interventions approach

of Finlayson et al. (2021) on stimuli in languages typologically related to English, such that we can observe whether there exist syntax neurons that are shared across a set of languages that are all relatively high-resource and grammatically similar. Our contributions include the following:

1. We causally probe for syntactic agreement neurons in an autoregressive language model, XGLM (Lin et al., 2021); a masked language model, multilingual BERT (Devlin et al., 2019); and a series of monolingual BERT-based models. We find two distinct layer-wise effect patterns, depending on whether the subject and verb are separated by other tokens.
2. We quantify the degree of neuron overlap across languages and syntactic structures, finding that many neurons are shared across structures and fewer are shared across languages.
3. We analyze the sparsity of syntactic agreement representations for individual structures and languages, and find that syntax neurons are more sparse in MLMs than ALMs, but also that the degree of sparsity is similar across models and structures.

Our data and code are publicly available.¹

2 Related Work

Multilingual language modeling. Multilingual language models enable increased parameter efficiency per language, as well as cross-lingual transfer to lower-resource language varieties (Wu and Dredze, 2019). This makes both training and deployment more efficient when support for many languages is required. A common approach for training multilingual LMs is to concatenate training corpora for many languages into one corpus, often without language IDs (Conneau et al., 2020; Devlin et al., 2019).

These models present interesting opportunities for syntactic analysis: Do multilingual models maintain similar syntactic abilities despite a decreased number of parameters that can be dedicated to each language? Current evidence suggests slight interference effects, but also that identical models maintain much of their monolingual performance when trained on multilingual corpora (Mueller et al., 2020). Is syntactic agreement, in particular, encoded independently per language or

shared across languages? Some studies suggest that syntax is encoded in similar ways across languages (Chi et al., 2020; Stańczak et al., 2022), though these rely on correlational methods based on dependency parsing, which introduce confounds and may not rely on syntactic information *per se*.

Syntactic probing. Various behavioral probing studies have analyzed the syntactic behavior of monolingual and multilingual LMs (Linzen et al., 2016; Marvin and Linzen, 2018; Ravfogel et al., 2019; Mueller et al., 2020; Hu et al., 2020). Results from behavioral analyses are generally easier to interpret and present clearer evidence for *what* models’ preferences are given various contexts. However, these methods do not tell us *where* or *how* syntax is encoded.

A parallel line of work employs parametric probes. Here, a linear classifier or multi-layer perceptron probe is trained to map from a model’s hidden representations to dependency attachments and/or labels (Hewitt and Manning, 2019) to locate syntax-sensitive regions of a model. This approach has been applied in multilingual models (Chi et al., 2020), and produced evidence for parallel dependency encodings across languages. However, if such probes are powerful, they may learn the target task themselves rather than tap into an ability of the underlying model (Hewitt and Liang, 2019), leading to uninterpretable results. When controlling for this, even highly selective probes may not need access to syntactic information to achieve high structural probing performance (Sinha et al., 2021). There are further confounds when analyzing individual neurons using correlational methods; for example, probes may locate encoded information that is not actually used by the model (Antverg and Belinkov, 2022).

Causal probing has recently become more common for interpreting various phenomena in neural models of language. Lakretz et al. (2019) and Lakretz et al. (2021) search for syntax-sensitive units in English and Italian monolingual LSTMs by intervening directly on activations and evaluating syntactic agreement performance. Vig et al. (2020) propose causal mediation analysis for locating neurons and attention heads implicated in gender bias in pre-trained language models; this method involves intervening directly on the inputs or on individual neurons. Finlayson et al. (2021) extend this approach to implicate neurons in syntactic agreement. This study extends their data and

¹<https://github.com/aaronmueller/multilingual-lm-intervention>

method to multilingual stimuli and models.

Other causal probing work uses interventions on model representations, rather than inputs. This includes amnesic probing (Elazar et al., 2021), where part-of-speech and dependency information is deleted from a model using iterative nullspace projection (INLP; Ravfogel et al., 2020). Ravfogel et al. (2021) employ INLP to understand how relative clause boundaries are encoded in BERT.

3 Methods

3.1 Causal Metrics

We first define terms to represent the quantities we measure before and after the intervention. We are interested in the impact of an intervention \mathbf{x} on a model’s preference $y_{\mathbf{x}}$ for grammatical inflections over ungrammatical ones. We start with the original input, on which we apply the null intervention: This represents performing no change to the original input. Given prompt u and verb v , we first calculate the following ratio:

$$y_{\text{null}}(u, v) = \frac{p(v_{\text{pl}} | u_{\text{sg}})}{p(v_{\text{sg}} | u_{\text{sg}})} \quad (1)$$

Here, u_{sg} represents a prompt that would require a singular verb inflection v_{sg} at the [MASK] for the sentence to be grammatical; for example, “The **doctor** near the **cars** [MASK] it”. v_{sg} is the third-person singular present inflection of verb v , and v_{pl} is the plural present inflection; for example, $v_{\text{sg}} = \text{“observes”}$ and $v_{\text{pl}} = \text{“observe”}$. Note that this ratio has the incorrect inflection as the numerator; this entails that if the model computes agreement correctly, we will have $y < 1$.

We now define the swap-number intervention, where the grammatical number of u is flipped (resulting in “The **doctors** near the **cars** [MASK] it” for the previous example). This results in the following expression for y :

$$y_{\text{swap-number}}(u, v) = \frac{p(v_{\text{pl}} | u_{\text{pl}})}{p(v_{\text{sg}} | u_{\text{pl}})} \quad (2)$$

Now, the numerator is the correct inflection, so we expect $y > 1$.

As we are interested in the contribution of individual model components to the model’s overall preference for correct inflections, we focus on *indirect effects*, where we perform interventions on individual model components and observe the

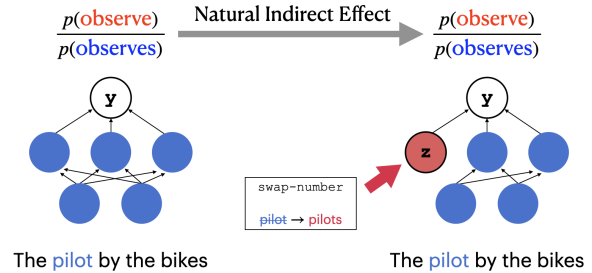


Figure 1: Example of computing the natural indirect effect (NIE). We change a neuron’s activation to what it would have been if we had intervened on the prompt, then measure the relative change in y .

change in y . In particular, we measure the **natural indirect effect** (NIE), as follows.

We intervene on an individual neuron \mathbf{z} . We change \mathbf{z} ’s original activation given u and v (denoted $\mathbf{z}_{\text{null}}(u, v)$) to the activation it *would have* taken if we had performed the intervention on u (denoted $\mathbf{z}_{\text{swap-number}}(u, v)$). The rest of the neurons retain their original activations. “Natural” here refers to the fact that our intervention changes the activation \mathbf{z} to the value it would have in another natural setting u' , rather than setting it to some predefined constant (such as 0) that it may or may not obtain given natural inputs. We measure the relative change in y after applying the intervention (see Figure 1 for a visual example):

$$\begin{aligned} \overline{\text{NIE}}(\text{swap-number}, \text{null}; y, \mathbf{z}) &= \\ \mathbb{E}_{u,v} \left[\frac{y_{\text{null}, \mathbf{z}_{\text{swap-number}}(u,v)}(u, v) - y_{\text{null}}(u, v)}{y_{\text{null}}(u, v)} \right] &= \\ \mathbb{E}_{u,v} \left[\frac{y_{\text{null}, \mathbf{z}_{\text{swap-number}}(u,v)}(u, v)}{y_{\text{null}}(u, v)} - 1 \right] & \quad (3) \end{aligned}$$

If a neuron encodes useful information for syntactic agreement, we expect y to *increase* after the intervention, making the numerator positive. Positive NIEs indicate that a neuron encodes preferences for correct verb inflections, and negative NIEs indicate that the neuron prefers incorrect inflections. The closer the NIE is to 0, the less of a contribution a neuron makes to syntactic agreement in either direction.

3.2 Models

Finlayson et al. (2021) analyzed a series of monolingual autoregressive language models (ALMs): GPT-2 (Radford et al., 2019), TransformerXL (Dai

Model	Layers	Neurons	Parameters
BERT	12	9126	110M
mBERT	12	9126	110M
GPT-2	24	25600	345M
XGLM	24	25600	564M

Table 1: The size of each model used in this study. Each monolingual BERT variant (including the RoBERTa-based CamemBERT) has the same number of layers, neurons, and parameters as BERT.

et al., 2019), and XLNet (Yang et al., 2019). Here, we apply their analysis approach to multilingual models. Multilingual ALMs are rare in the literature; to our knowledge, the only ALM designed to be multilingual is XGLM (Lin et al., 2021),² which we employ in this study.

Multilingual MLMs are much more common. We focus on multilingual BERT (Devlin et al., 2019). We were unable to analyze XLM-R (Conneau et al., 2020), a more recent multilingual MLM that performs better than mBERT on certain benchmarks, since its tokenizer splits a large proportion of our nouns and verbs into multiple tokens, which greatly constrained the stimuli we could use. In future work, we intend to address this issue by developing methods that enable multi-token interventions, as well as calibrated comparisons across variable-length sequences.

We also analyze a series of monolingual MLMs—one for each language included in our sample. Four of these models were based on BERT: BERT (English), GermanBERT,³ BERTje (Dutch; de Vries et al., 2019), and FinnishBERT (Virtanen et al., 2019). Our French MLM, CamemBERT (Martin et al., 2020), is based on RoBERTa (Liu et al., 2019), which is very similar to BERT.

3.3 Materials

We translate the stimuli from Finlayson et al. (2021) (Figure 2) to French, German, Dutch, and Finnish. Since the subjects and verbs on which we intervene must be one token each,⁴ we are restricted

²GPT-3 (Brown et al., 2020) is technically multilingual, as its training corpus contains data from other languages. However, it was not designed with multilinguality in mind, and the vast majority of its training data is English.

³<https://www.deepset.ai/german-bert>

⁴It is not clear how to compare the probability of variable-length sequences in masked language models, and autoregressive language models tend to prefer sequences containing fewer tokens. There have been attempts to compare variable-length sequence probabilities using iterative approaches (e.g.,

Simple Agreement:
The athlete investigates/*investigate...

Across Prepositional Phrase:
The manager behind the bikes
observes/*observe...

Across Object Relative Clause:
The farmers that the parent loves
*confuses/confuse...

Figure 2: Constructions used in this study, grouped by whether the subject and verb are adjacent. We use a subset of constructions from Finlayson et al. (2021), directly translating the stimuli to French, German, Dutch, and Finnish. See Appendix A for examples of each structure in each language.

to very frequent words in the pre-training corpus which do not get split into subwords by a model’s tokenizer. This limits us to high-resource language varieties—and as most of the top languages in mBERT and XGLM’s pre-training corpora are Indo-European, this also limits the typological range of this method. A virtue of our sample of languages, however, is that it allows us to study whether neurons are shared across typologically similar languages, where shared neurons and similar layer-wise effect patterns are most likely to occur: If syntactic agreement neurons are not shared across similar languages, they are unlikely to be shared across *any* languages.

For each structure, we sample up to 200 sentences. If there are fewer than 200 sentences where the subjects and verbs are single tokens, we take the entire set of valid stimuli. When we use the original stimuli from Finlayson et al. (2021), we often have very few sentences where the subjects and verbs are single tokens. Thus, we also create short-word versions of the stimuli, where we use shorter and more common words (e.g., instead of "managers" or "observe", we can use "cats" or "see"). Our results are consistent when using the original and short nouns and verbs; see Appendix B.

The original stimuli were generated from a grammar given a list of manually selected terminals. By generating artificial stimuli and *not* sampling sentences from a corpus, we partially control for memorized sequences or token collocations in the pre-training corpus.

⁵Schick and Schütze, 2021), though this generally requires fine-tuning to work properly.




Figure 3: Natural indirect effects for the top 5% of neurons in each layer for monolingual masked language models. There are two distinct layer-wise NIE contours in each language, depending on whether the subject and verb are separated by other tokens (as in ‘across a relative clause’ and ‘across a prepositional phrase’ structures) or not (as in ‘simple agreement’).

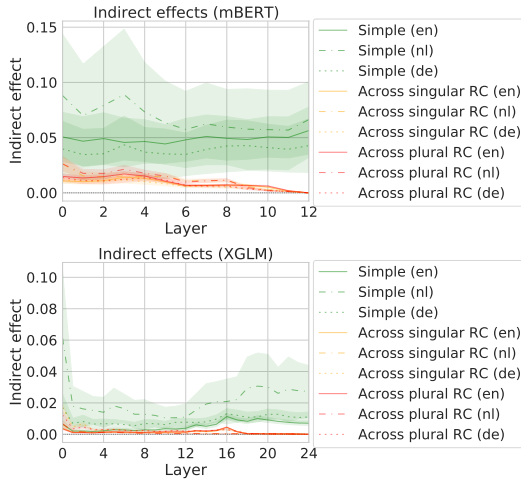


Figure 4: Natural indirect effects for mBERT (top) and XGLM (bottom) for Germanic languages. There are two distinct layer-wise NIE patterns in each language. NIE patterns for the same structure look very similar across languages.

Finlayson et al. (2021) found two distinct layer-wise NIE patterns for syntactic agreement: one when the subject and verb are adjacent (the *short-range* effect), and another when they are separated by any number of tokens (the *long-range* effect). To understand whether the short-range effect is due to preferences for frequent bigrams (rather than specifically grammatical subject-verb bigrams), we also design a bigram swap intervention. We use high-mutual-information adjective-noun English bigrams as the original inputs and intervene by randomly swapping the first or second words in the bigram with words from a different bigram. For example, given the bigrams **coaxial cable** and **police officer**, we can define $y_{\text{null}} = \frac{p(\text{officer}|\text{coaxial})}{p(\text{cable}|\text{coaxial})}$ and $y_{\text{swap-bigram}} = \frac{p(\text{officer}|\text{police})}{p(\text{cable}|\text{police})}$. Then we can compute the NIE as in Equation 3.

Finally, to test whether separate neurons are used

for short- and long-range token collocations in general, we also define short- and long-range *semantic plausibility* baselines, where nouns are associated with stereotypical adjectives (e.g., **square T.V.** and **red apple**). The short-range semantic plausibility intervention is the same as for the bigram intervention: We compute the probability ratio of the first and second noun in a pair of bigrams before and after swapping the adjective. For long-range semantic plausibility, the prompt u is “The **T.V./apple** is”, and v is the probability ratio of the *adjectives* before and after swapping the nouns.

4 Results

4.1 Layer-wise NIE contours are similar across languages

We present indirect effects for monolingual masked language models (Figure 3), as well as mBERT and XGLM (Figure 4). Here, we select the top 5% of neurons per layer by NIE. In each language, whether in a monolingual or multilingual MLM or ALM, **there are two distinct layer-wise NIE effect patterns for number agreement**: one for short-range dependencies and one for long-range dependencies. This agrees with the findings of Finlayson et al. (2021) on autoregressive English LMs. However, these effects look more distinct across monolingual models, whereas **multilingual models exhibit more similar layer-wise NIE patterns across languages**. In other words, monolingual models accomplish syntactic agreement in different layers and neurons depending on the language (even though these languages are typologically similar), but in multilingual models agreement computations implicate the same layers across languages. This does not necessarily mean that the same individual neuron are being used cross-linguistically in multilingual models (we explore this question

in more detail in §4.2.1); rather, the model may simply be learning similar layer-wise strategies for each language.

While prior work finds that syntactic agreement is easier to learn in languages that have more explicit morphological cues to hierarchical structure (Ravfogel et al., 2019; Mueller et al., 2020),⁵ this does not necessarily imply that different agreement mechanisms are learned in such languages. We find similar layer-wise NIEs in mBERT across each language we consider, including Finnish, a non-Indo-European (specifically, Uralic) language.

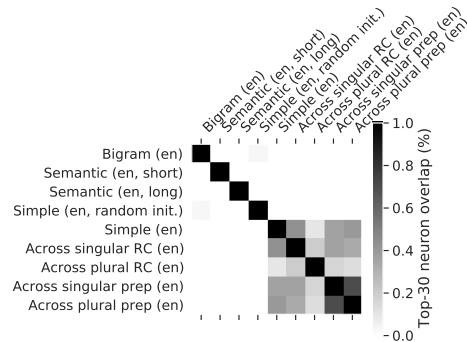
4.2 Syntax neurons are shared across structures, but not with semantic baselines

Here, we analyze to what extent the same neurons are implicated across syntactic structures and languages in mBERT. For each structure, we take the top 30 neurons by indirect effect (from any layer); we then compute the proportion of such high-NIE neurons that are shared across structures.


First, we investigate to what extent the neurons that have high NIE for the syntactic structures are selective to syntax. We do so by computing the overlaps in English between neurons with high NIE for syntactic structures and the neurons with high NIE for our bigram and semantic plausibility baselines. We find that the top syntactic agreement neurons for any structure are *not* shared with the neurons implicated in semantic plausibility or bigram collocation (Figure 5). In other words, **the neurons used for syntactic agreement are specific to agreement** and do not track common bigrams more generally.

Figure 5 also shows that **neurons are shared across syntactic structures**, providing evidence for an abstract notion of syntactic agreement encoded in mBERT that is separate from the individual structures that the model is presented with. However, the varying extents of overlap indicate that there are also neurons specialized to particular structures. To further contextualize these overlap proportions, we also compute overlaps for simple agreement in a randomly initialized mBERT, as a baseline. This experiment yields near-zero overlaps, indicating that the overlaps across structures we obtain for mBERT and XGLM are unlikely to

⁵Explicit case marking correlates well with performance on syntactic evaluations (Ravfogel et al., 2019), so we would expect German and Finnish to exhibit different results if these cues give rise to different agreement computations.



(a) mBERT



(b) XGLM

Figure 5: Neuron overlap across structures (including baselines) in English for (a) mBERT and (b) XGLM. There is zero or near-zero overlap between the baselines and all syntactic agreement structures, whereas overlap is relatively high (and statistically significant) for all other structures.

be due to random chance.⁶

4.2.1 Neurons are shared across languages in autoregressive language models

The overlap in neurons across languages (Figure 6) is significant for all structures in XGLM. For mBERT, overlap is significant between “across a PP” structures and other long-distance agreement structures, but not for any other structure pairs. Note that in XGLM, the diagonal is no darker than most other squares; in other words, there is *not* more cross-lingual neuron overlap for the same syntactic structure relative to other structures. **These may be generic cross-lingual syntax neurons which are not specialized to any particular structure or language.** We found in §4.2 that there is almost no overlap between syntactic agreement neurons and bigram collocation/semantic plausibility neurons in English, which is further evidence

⁶For reference, the probability of at least one neuron being shared between two random samples of 30 neurons in (m)BERT-base is $1 - \frac{\binom{9984-30}{30}}{\binom{9984}{30}} \approx .086$.

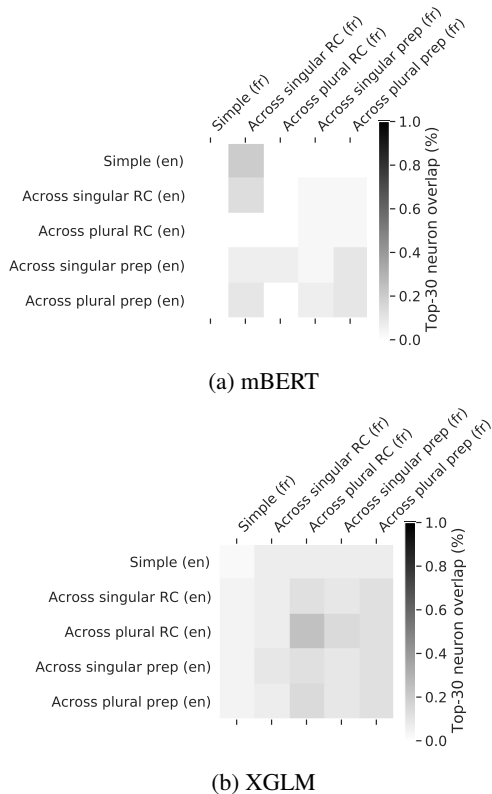


Figure 6: Cross-lingual neuron overlaps for the top 30 neurons by NIE in (a) mBERT and (b) XGLM. We present English-French overlaps; overlaps between other language pairs look similar (see Appendix C). The overlap percentages in (b) are significantly higher than random chance. Overlaps for most structure pairs in (a) are not significant, except for overlaps between ‘Across a preposition’ structures and other long-range agreement structures.

that these may be more general syntactic agreement neurons. Nonetheless, overlap is very low across languages compared to across structures within a given language. Thus, **in autoregressive language models, syntactic agreement neurons can be language-specific or cross-lingual, but most are language-specific. For masked language models, syntactic agreement neurons are rarely shared across languages.**

4.3 Neuron sparsity differs across structures, but not across languages

What proportion of LMs’ neurons encode subject-verb agreement? The sparsity of syntax neurons in pre-trained models may vary depending on which language and structure we observe. Lakretz et al. (2019) and Lakretz et al. (2021) found that agreement neurons are sparse in LSTMs, but it is not clear whether this would hold for MLMs or large

Language	Model	% Neurons for TE	% Neurons for Max. NIE
en	BERT	1.0%	5.8%
	mBERT	1.0%	8.7%
	GPT-2	17.5%	25.0%
	XGLM	4.5%	16.5%
fr	CamemBERT	6.7%	10.6%
	mBERT	3.8%	29.8%
	XGLM	3.5%	24.0%
de	GermanBERT	1.0%	8.7%
	mBERT	1.0%	6.7%
	XGLM	1.5%	18.0%
nl	BERTje	1.0%	5.8%
	mBERT	1.0%	2.9%
	XGLM	0.5%	37.5%
fi	FinnishBERT	1.0%	4.8%

Table 2: Neuron sparsities for the ‘simple agreement’ structure across languages and models. Multilingual models do not necessarily encode syntax more sparsely than monolingual models. Sparsities are generally consistent across languages for the same model.


Transformer-based ALMs. Given our consistent results across languages, we hypothesize that the neuron sparsity of subject-verb agreement will be similar across monolingual models. Given the consistent distinction thus far in how neurons encode short- and long-range agreement, we also hypothesize that neuron sparsity will differ between agreement distances. Due to lower parameterization per language in multilingual models, however, we hypothesize that multilingual models encode agreement more sparsely than monolingual models.

We measure sparsity by iteratively selecting the top k neurons by NIE, intervening on them simultaneously, and computing the natural indirect effect after performing the swap-number intervention. We continue sampling k more neurons and computing NIEs until we have selected all neurons; the NIE after intervening on all neurons is equivalent to the *total effect* (TE).⁷ Computing effects for each neuron and each structure is computationally expensive, so we use $k = 128$ for XGLM and GPT-2 (0.5% of neurons selected at a time) and $k = 96$ for (m)BERT ($\approx 1.0\%$ of neurons selected at a time).


We report two metrics: (1) the percentage of neurons at which we see the maximum NIE, and (2) the minimum percentage of neurons required for the NIE to reach the TE of the model. These correspond to the peak NIE and the point at which the NIEs cross the dashed line in Figure 7.

For ‘simple agreement’ (Table 2), the proportion of neurons to reach the TE is generally small, espe-

⁷Intuitively, the TE can be thought of as the preference of the model as a whole for correct verbs over incorrect verbs.



(a) Simple agreement



(b) Across a singular RC

Figure 7: Indirect effects when intervening on increasing numbers of neurons in XGLM. The dashed line represents the total effect. For ‘simple agreement’, there exists a set of neurons that strongly prefers grammatical completions; however, there are many more neurons that have weak preferences against them, and this results in the model as a whole having weak preferences for correct verb inflections. For ‘across a singular RC’, however, almost every set of neurons seems to have preferences for grammatical inflections.

cially for MLMs. However, the TE itself is often a couple orders of magnitude smaller for MLMs than for ALMs; thus, these percentages are not comparable across model architectures.

The proportion of neurons required to achieve the maximum NIE is typically lower for MLMs than ALMs. In other words, **syntax neurons are more sparse in masked language models than autoregressive language models.**⁸

The percentage of neurons to reach the maximum NIE does not significantly differ across monolingual and multilingual models, however. This means that **multilingual models do not consistently encode syntactic agreement in a more sparse way than monolingual models.** This and our neuron overlap results suggest that multilingual models encode syntactic information in a similar way to monolingual models (including the proportion of neurons sensitive to syntax), though

⁸French is an exception: there are more syntax-sensitive neurons in both monolingual and multilingual models.

most syntax-sensitive neurons tend to be language-specific rather than shared across languages.

Sparsity also differs across syntactic structures. For ‘simple agreement’, NIEs peak at around 5–20% of neurons. For ‘across a singular RC’, the addition of every k neurons almost always increases the NIEs. **Long-range syntactic information seems to be distributed throughout the majority of neurons in XGLM, but short-range syntactic information is more sparsely encoded.**

These numbers hide more interesting trends, however. The TEs for mBERT are often close to 0 across structures, while the maximum NIEs are in the hundreds for those same structures.⁹ This has interesting implications for interpreting behavioral analyses: studies such as Hu et al. (2020) and Mueller et al. (2020) suggest that mBERT does not have strong syntax-sensitive preferences compared to autoregressive language models, and the low TEs we observe support this. However, this obscures that **there are actually many neurons in mBERT which are highly sensitive to syntactic agreement**, as indicated by the high maximum NIEs: we observe weak agreement preferences in the model as a whole because there are many more neurons which have weak preferences *against* syntactic agreement (i.e., small negative NIEs), perhaps because those neurons are specializing in other phenomena (e.g., token collocations or semantic agreement). Thus, behavioral analyses of model behavior may be underestimating the sensitivity of models to syntactic phenomena, for there is negative interference from neurons that prefer non-syntax-sensitive completions.

5 Discussion

We observe two distinct layer-wise NIE patterns for syntactic agreement, depending on whether the subject and verb are adjacent or separated by other tokens. This extends the findings of Finlayson et al. (2021) to multilingual MLMs and ALMs, as well as monolingual MLMs in various languages. Going beyond their findings, we ruled out the possibility that these neurons do not simply track semantic plausibility or bigram collocations more generally. While this is not conclusive evidence that these neurons are specialized to syntax, evidence from other behavioral and probing studies also supports

⁹The effect contours for mBERT have a similar contour to those in Figure 7, though the TE (≈ 0) and maximum NIE (≈ 340) for ‘simple agreement’ are far smaller.

the existence of neurons focused on syntax (Hewitt and Manning, 2019; Elazar et al., 2021; Goldberg, 2019). De Cao et al. (2021) found neurons focused *purely* on syntax, while Tucker et al. (2022) found redundantly encoded syntactic information across neurons. It is not clear how much of the neuron overlap we observe is due to redundantly encoded information, but future work could investigate this.

A consistent trend across our experiments is that ALMs encode syntactic agreement in a distinct way from MLMs. In ALMs, there is more cross-lingual and cross-structure neuron overlap than in MLMs; more similar layer-wise effect patterns across structures and languages (though they are still distinct); and a greater proportion of neurons which are sensitive to agreement. This could be partially explained by ALMs’ left-to-right processing of natural language input, which more closely resembles incremental inputs to human learners. MLMs are able to perform syntactic agreement (Hu et al., 2020; Goldberg, 2019), but their fill-in-the-blank pre-training objectives may induce distinct representations of sentence structure as compared to models that process or predict inputs incrementally.

Why do we observe different indirect effect contours for short- and long-range agreement? Perhaps syntactic agreement is encoded using a single mechanism, but the way that syntactic information is used for predicting output tokens depends on the structure of the input or prior output tokens. Alternatively, there could be two completely distinct agreement mechanisms that function in different ways entirely. While our findings do not disambiguate between these possibilities (or some other separate type or amount of mechanisms), future work could employ methods like those in Meng et al. (2022) to observe this distinction more explicitly. The findings of Meng et al. (2022) suggest that the model regions that are implicated in *saying* something are distinct from those implicated in *knowing* something—that is, knowledge retrieval and predicting particular tokens are separate mechanisms in pre-trained language models. Perhaps their method could be extended to study syntactic agreement, such that we can better understand what, exactly, these distinct indirect effect trends represent.

6 Conclusions

We have used causal mediation analysis to observe which neurons track syntactic agreement in multi-

lingual pre-trained language models, and in which layers they are concentrated. We found two distinct layer-wise contours for syntactic agreement regardless of the language, multilinguality, or architecture of the model (§4.1); that syntax-sensitive neurons are shared across languages in autoregressive language models (§4.2.1); and that the neuron sparsity of syntactic agreement is similar in monolingual and multilingual models (§4.3). We also found that behavioral analyses of masked language models obscure the extent to which their neurons are sensitive to syntactic agreement (§4.3).

Acknowledgments

We thank the members of NYU’s Computation and Psycholinguistics Lab for valuable feedback on earlier versions of this work.

This material is based upon work supported by the National Science Foundation (NSF) under Grant No. BCS-2114505. Aaron Mueller was supported by a National Science Foundation Graduate Research Fellowship (Grant #1746891). This work was also supported by supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

References

- Omer Antverg and Yonatan Belinkov. 2022. [On the pitfalls of analyzing individual neurons in language models.](#)
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners.](#) In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- 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.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. **Transformer-XL: Attentive language models beyond a fixed-length context**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Nicola De Cao, Leon Schmid, Dieuwke Hupkes, and Ivan Titov. 2021. **Sparse interventions in language models with differentiable masking**. *CoRR*, arXiv:2112.06837.
- Wietse de Vries, Andreas van Cranenburgh, Arianna Bisazza, Tommaso Caselli, Gertjan van Noord, and Malvina Nissim. 2019. **BERTje: A Dutch BERT model**. *CoRR*, arXiv:1912.09582.
- 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.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2021. **Amnesic Probing: Behavioral Explanation with Amnesic Counterfactuals**. *Transactions of the Association for Computational Linguistics*, 9:160–175.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. 2021. **Causal analysis of syntactic agreement mechanisms in neural language models**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1828–1843, Online. Association for Computational Linguistics.
- Yoav Goldberg. 2019. **Assessing BERT’s syntactic abilities**. arXiv preprint 1901.05287.
- John Hewitt and Percy Liang. 2019. **Designing and interpreting probes with control tasks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. 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.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. **A systematic assessment of syntactic generalization in neural language models**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Yair Lakretz, Dieuwke Hupkes, Alessandra Vergallito, Marco Marelli, Marco Baroni, and Stanislas Dehaene. 2021. **Mechanisms for handling nested dependencies in neural-network language models and humans**. *Cognition*, page 104699.
- Yair Lakretz, German Kruszewski, Theo Desbordes, Dieuwke Hupkes, Stanislas Dehaene, and Marco Baroni. 2019. **The emergence of number and syntax units in LSTM language models**. 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 11–20, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shrutu Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona T. Diab, Veselin Stoyanov, and Xian Li. 2021. **Few-shot learning with multilingual language models**. *CoRR*, arXiv:2112.10668.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. **Assessing the ability of LSTMs to learn syntax-sensitive dependencies**. *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **RoBERTa: A robustly optimized BERT pretraining approach**. *CoRR*, arXiv:1907.11692.
- Christopher D Manning, Kevin Clark, John Hewitt, Urvasi Khandelwal, and Omer Levy. 2020. **Emergent linguistic structure in artificial neural networks trained by self-supervision**. *Proceedings of the National Academy of Sciences*, 117(48):30046–30054.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. **CamemBERT: a tasty French language model**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.

- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual knowledge in gpt](#).
- Aaron Mueller, Robert Frank, Tal Linzen, Luheng Wang, and Sebastian Schuster. 2022. [Coloring the blank slate: Pre-training imparts a hierarchical inductive bias to sequence-to-sequence models](#). *CoRR*, arXiv:2203.09397.
- Aaron Mueller, Garrett Nicolai, Panayiota Petrou-Zeniou, Natalia Talmina, and Tal Linzen. 2020. [Cross-linguistic syntactic evaluation of word prediction models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5523–5539, Online. Association for Computational Linguistics.
- Judea Pearl. 2001. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 411–420.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI Blog*.
- Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. [Null it out: Guarding protected attributes by iterative nullspace projection](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online. Association for Computational Linguistics.
- Shauli Ravfogel, Yoav Goldberg, and Tal Linzen. 2019. [Studying the inductive biases of RNNs with synthetic variations of natural languages](#). 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 3532–3542, Minneapolis, Minnesota. Association for Computational Linguistics.
- Shauli Ravfogel, Grusha Prasad, Tal Linzen, and Yoav Goldberg. 2021. [Counterfactual interventions reveal the causal effect of relative clause representations on agreement prediction](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 194–209, Online. Association for Computational Linguistics.
- James M. Robins. 2003. Semantics of causal DAG models and the identification of direct and indirect effects. *Oxford Statistical Science Series*, pages 70–82.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021. [Masked language modeling and the distributional hypothesis: Order word matters pre-training for little](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2888–2913, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Karolina Stańczak, Edoardo Ponti, Lucas Torroba Henrigen, Ryan Cotterell, and Isabelle Augenstein. 2022. [Same neurons, different languages: Probing morphosyntax in multilingual pre-trained models](#).
- Mycal Tucker, Tiwalayo Eisape, Peng Qian, Roger Levy, and Julie Shah. 2022. [When does syntax mediate neural language model performance? evidence from dropout probes](#).
- 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.
- Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski, Filip Ginter, and Sampo Pyysalo. 2019. [Multilingual is not enough: BERT for Finnish](#).
- Alex Warstadt and Samuel R. Bowman. 2020. [Can neural networks acquire a structural bias from raw linguistic data? In Proceedings of the 42nd Annual Meeting of the Cognitive Science Society](#), Online. Cognitive Science Society.
- Alex Warstadt, Yian Zhang, Xiaocheng Li, Haokun Liu, and Samuel R. Bowman. 2020. [Learning which features matter: RoBERTa acquires a preference for linguistic generalizations \(eventually\)](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 217–235, Online. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. [Beto, Bentz, Becas: The surprising cross-lingual effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language*

Processing (EMNLP-IJCNLP), pages 833–844, Hong Kong, China. Association for Computational Linguistics.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [XLNet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

A Example Sentences

Here, we present examples of each syntactic structure we observe in each language.

(1) *Simple agreement (English):*

̣ The woman observes/*observe.

(2) *Simple agreement (French):*

̣
L' homme approuve/*approuvent.
The man approves/*approve.

(3) *Simple agreement (German):*

̣
Der Arzt weiß/*wissen.
The physician knows/*know.

(4) *Simple agreement (Dutch):*

̣
De schrijver begrijpt/*begrijpen.
The writer understands/*understand.

(5) *Simple agreement (Finnish):*

̣
Täti ymmärtää/*ymmärtävät.
Aunt understands/*understand.
“The aunt understands/*understand.”

For each of the following syntactic structures containing a grammatical number attractor, we separate structures by whether the attractor is singular or plural. For concision, we simply present examples of each structure without separating out examples by the number of the attractor. Note that Finnish mainly uses *postpositions* rather than prepositions; the attractor still intervenes between the main subject and its verb, but the order of the preposition and noun phrase is different compared to the Indo-European languages we consider.

(6) *Across a relative clause (English):*

̣ The woman that the guards like observes/*observe.

(7) *Across a relative clause (French):*

̣
L' homme que le chef suit
The man that the boss follows

approuve/*approuvent.
approves/*approve.

(8) *Across a relative clause (German):*

̣
Der Arzt den die Tiere vergeben
The physician that the animals forgive
weiß/*wissen.
knows/*know.

(9) *Across a relative clause (Dutch):*

̣
De schrijver die de ouder roept
The writer that the parent calls
begrijpt/*begrijpen.
understands/*understand.

(10) *Across a relative clause (Finnish):*

̣
Täti jota luistelijat kehuvat
Aunt that skaters praise
ymmärtää/*ymmärtävät.
understands/*understand.

“The aunt that the skaters praise understands/*understand.”

(11) *Across a prepositional phrase (English):*

̣ The woman behind the cars observes/*observe.

(12) *Across a prepositional phrase (French):*

̣
L' homme devant le chat
The man in-front-of the cat
approuve/*approuvent.
follows approves/*approve.

(13) *Across a prepositional phrase (German):*

̣
Der Arzt nahe den Äpfeln weiß/*wissen.
The physician near the apples knows/*know.

(14) *Across a prepositional phrase (Dutch):*

̣
De schrijver achter de fiets
The writer behind the bike
begrijpt/*begrijpen.
understands/*understand.

(15) *Across a postpositional phrase (Finnish):*

̣
Täti puiden lähellä ymmärtää/*ymmärtävät.
Aunt trees near understands/*understand.

“The aunt near the trees understands/*understand.”




Figure 8: Natural indirect effects for the top 5% of neurons in each layer for monolingual masked language models. The indirect effect contours we observe do not vary significantly when replacing the nouns and verbs with shorter, more frequent words—except in layer 11 of CamemBERT.

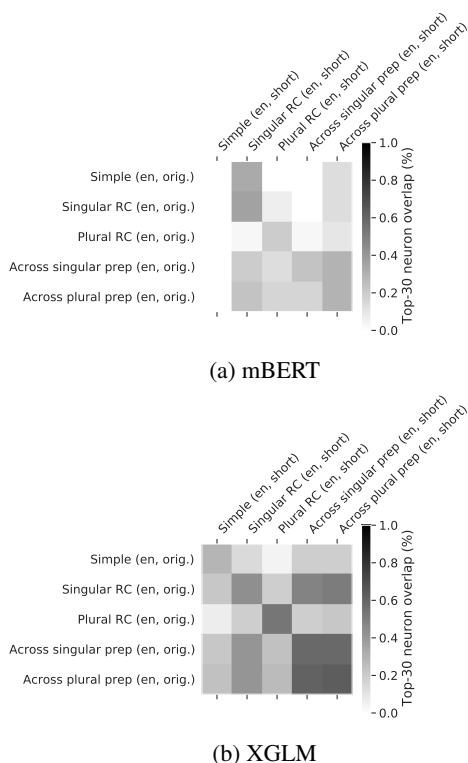


Figure 9: Neuron overlap across structures (including baselines) in English for (a) mBERT and (b) XGLM. There is significant overlap between the original stimuli and short-word stimuli, though this is more the case for XGLM than mBERT.

B Invariance to Short- and Long-Word Stimuli

When using the stimuli from [Finlayson et al. \(2021\)](#), most of the subjects and verbs are split into multiple tokens. These are generally long and relatively infrequent nouns and verbs like “managers” and “observe”. We could use more stimuli if we replace each word with words that are shorter and more frequent in pre-training corpora, such as “cats” and “see”.

Will these lexical replacements change the trends we observe? We observe the layer-wise natural indirect effect of the top neurons in each layer for the original stimuli and the short-word stimuli to see if lexical replacements have an effect on the way neurons encode syntactic agreement in monolingual BERT models. Our results (Figure 8) are nearly identical for the original stimuli and the short stimuli. A notable exception is layer 11 of CamemBERT, where indirect effects are so large that the rest of the effects are dwarfed by comparison. However, when excluding this result, indirect effect contours look similar between original and short stimuli.

We also compare the extent of neuron overlap between original and short stimuli for multilingual BERT and XGLM. Our results (Figure 9) show a relatively high degree of overlap, especially for XGLM. However, overlap is somewhat lower than when we use only one stimulus type (Figure 5). Ideally, overlap should be nearly 100% along the diagonal of both matrices if these neurons account only for syntactic agreement rather than specific lexical items, so these results suggest that lexical (and not syntactic) features may account for a notable proportion of the neuron overlap we observe in our previous experiments. Alternatively, it could mean that these neurons attend both syntactic *and* lexical information. Nonetheless, overlaps are still significant and indirect effects still look similar when swapping our nouns and verbs, so it is likely that models are picking up on some abstraction for syntactic agreement that generalizes across specific token sequences.

These results suggest that the neuron-level effects we observe are not simply spurious lexical correlations. More significantly, this is further evidence that **the neuron-level effects we observe are**

not word-level effects, but some more abstract structural feature(s) that the model has learned.

C Neuron Overlap Across Languages: Full Results

Here, we present neuron overlaps across languages for mBERT and XGLM (Figure 10). As in §4.2.1, we present overlaps for the top 30 neurons (in *any* layer of the model) per structure per language. As before, we find that neuron overlap is generally greater in autoregressive LMs than masked LMs.

Neuron overlaps are most prominent between English and French; while not typologically the most closely related language pair, English and French share a great deal of vocabulary and have similar SVO word orders when pronominal objects are not present. German, meanwhile, uses SOV with V2 in main clauses.

D Limitations

Perhaps the greatest limitation of our method—and many other causal probing methods (Vig et al., 2020; Finlayson et al., 2021; Ravfogel et al., 2021)—is that we are limited to stimuli where the subjects on which we intervene and the competing verb forms are one token each. This greatly limits the range of subjects and verbs (and languages) we can consider in this study, especially for more multilingual models where a greater proportion of words are split into subwords by the tokenizer. Models may use a different mechanism altogether to calculate the probability of two competing verbs given the presence or lack of a morpheme like {-s} which expresses number information, and our method would not allow us to understand where and how models are performing this kind of agreement. While one can, in theory, compare the probability of variable-length token sequences in autoregressive language models, there is no principled way to do this in masked language models. And in practice, autoregressive language models tend to prefer shorter sequences. Future work could consider probing methods which allow for variable-length span predictions.

There are also more general issues with probing individual neurons. Complex phenomena like syntactic agreement are likely to be encoded in *sets of* neurons, rather than individual neurons; indeed, we find evidence for this in §4.3. This means that analyzing individual neurons can result in oversimplified understandings of where and how certain

phenomena are encoded and used. Future causal probing work could focus on non-parametric methods which allow one to probe multiple neurons simultaneously, such that we may causally implicate model *regions* rather than just individual components like neurons or attention heads.

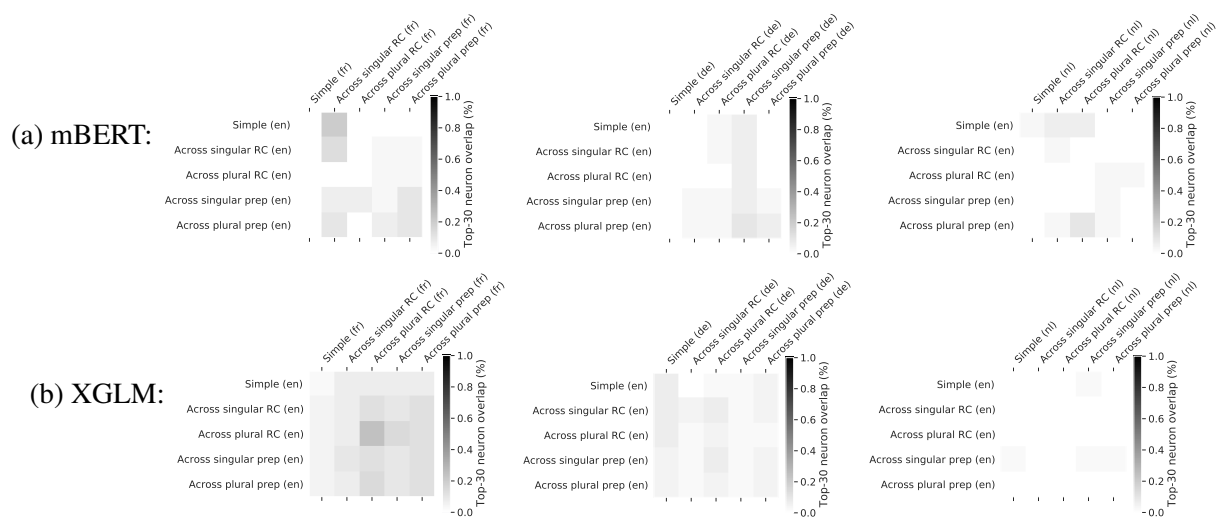


Figure 10: Neuron overlap for the top 30 neurons in mBERT (top row) and XGLM (bottom row). We show overlaps between English and French (left), German (center), and Dutch (right).