Anything Goes? A Crosslinguistic Study of (Im)possible Language Learning in LMs

Anonymous ACL submission

Abstract

Do LLMs offer insights into human language learning? A common argument against this idea is that because their architecture and training paradigm are so vastly different from humans, LLMs can learn arbitrary inputs as easily as natural languages. In this paper, we test this claim by training LMs to model impossible or typologically unattested languages. Unlike previous work, which has focused exclusively on English, we conduct experiments on 12 natural languages from 4 language families. Our results show that while GPT-2 small can primarily distinguish attested languages from their impossible counterparts, it does not achieve perfect separation between all the attested languages and all the impossible ones. We further test whether GPT-2 small distinguishes typologically attested from unattested languages with different NP orders by manipulating word order based on Greenberg's Universal 20 and find that the model's perplexity scores do not distinguish attested vs. unattested word orders, as long as the unattested variants maintain constituency structure. These findings suggest that language models exhibit some human-like inductive biases, though these biases are weaker than those found in human learners.

1 Introduction

To what extent can LLMs serve as models of human language acquisition and processing? Some, such as Piantadosi (2023), argue that LLMs can function as comprehensive linguistic theories, challenging traditional symbolic generative approaches. On the other hand, critics maintain that the success of LLMs is largely irrelevant to human cognition due to fundamental differences in architecture and learning mechanisms (Chomsky et al., 2023; Fox and Katzir, 2024). Moreover, studies have shown that LLMs fail to acquire key aspects of linguistic knowledge, highlighting their limitations as models of human language (Fox and Katzir, 2024; Lan et al., 2024; Katzir, 2023; Dentella et al., 2024). A particularly compelling argument in this debate is that LLMs are highly flexible learners, capable of acquiring linguistic patterns beyond those learnable by humans, thus making the ability of LLMs to learn human languages uninformative for understanding human language acquisition (Chomsky and Moro, 2022; Moro, 2023; Moro et al., 2023).

We present data that favors a more moderate stance, in line with other researchers in this field (Futrell and Mahowald, 2025; Millière, 2024; Pater, 2019). We present new empirical evidence from the study of *impossible languages* (Kallini et al., 2024) in a multilingual setting. Our findings suggest that LLMs exhibit learning biases that align with certain aspects of human cognition while simultaneously displaying biases (or a lack thereof) that diverge from human language processing.

In this work, we examine both possible (attested or unattested) and impossible (unattested by definition) languages. Specifically, we define **attested languages** as the natural languages spoken by humans (e.g., English, German, and Chinese); **unattested languages** as languages constructed based on language universals identified in typological studies or generative grammar analysis; and **impossible languages** as those that humans cannot acquire and would never produce. Following Kallini et al. (2024), we have selected impossible variants because we take them to be uncontroversial examples of linguistic impossibility, such as languages with randomly shuffled word orders. To explore unattested languages, we draw from Greenberg's Universal 20 (Greenberg et al., 1963), which identifies unattested word order patterns in noun phrases (e.g., adjective-number-determiner-noun). While there is no direct evidence that such languages are unlearnable, previous studies suggest that typological feature frequencies correlate with learnability in human learners (Culbertson et al., 2020; Gentner and Bowerman, 2009; Saffran et al., 2008).

Regarding impossible language modeling, Kallini et al. (2024) provided initial evidence that GPT-2 small can distinguish between possible and impossible variants of English, suggesting that transformer models encode human-like linguistic biases (Futrell and Mahowald, 2025). However, their study was limited to English, leaving the question of whether this finding generalizes across languages. Furthermore, their focus on impossible languages leaves the study of unattested languages, which we take as an important testbed for LLM-human bias alignment, largely unexplored.

This paper is organized around two main research questions: (1) **Do LLMs encode human-like distinctions between attested and impossible languages?** Specifically, (a) Within each attested language, can LLMs correctly differentiate the attested language from its impossible variants? (b) Across different attested languages from multiple language families, can LLMs distinguish all attested languages from all impossible languages? (2) **Can LLMs recognize unattested languages as distinct from attested ones?** Specifically, does LLMs' ability to model unattested languages align with human typological biases?

Our findings reveal that GPT-2 small reliably distinguishes attested and impossible languages within each attested language (1a) but struggles to make this distinction across different languages (1b). It assigns lower perplexity to unattested languages when they preserve constituency and fixed word order (2), suggesting a preference for regular structures (Hudson Kam and Newport, 2005; Singleton and Newport, 2004).[1]

2 Related Work

2.1 Language Models & Cognitive Plausibility

The advancement of neural networks makes connectionism a widely adopted framework in cognitive language studies (e.g., Wilcox et al., 2023; Borenstein et al., 2024; Kirov and Cotterell, 2018). However, linguists remain divided on whether language models can meaningfully inform linguistic theories. On the one hand, language models have advanced psycholinguistics by serving as highly accurate probability estimators, and, in this capacity, have already been used for testing and refining Surprisal Theory (Goodkind and Bicknell, 2018; Oh and Schuler, 2023b,a; Kuribayashi et al., 2024), Uniform Information Density (Meister et al., 2021; Tsipidi et al., 2024), and other cognitive-linguistic theories and psychometrics (Pearl and Mis, 2011; Gibson et al., 2019; Kuribayashi et al., 2025). On the other hand, their limitations, including a lack of generalization (Yao and Koller, 2022; Kim and Linzen, 2020), the shortcomings of prompt-based approaches (Hu and Levy, 2023), and inconsistency with humans (de Dios-Flores et al., 2023; Davis and van Schijndel, 2020) suggest that, beyond their role as sophisticated estimators, they are limited as cognitive models.

The most relevant work to our study in this context is Kallini et al. (2024), which tests the hypothesis that LLMs cannot distinguish between possible and impossible languages (Chomsky et al., 2023; Moro et al., 2023). Their study relies on a 100M-word dataset from the BabyLM Challenge (Warstadt et al., 2023), focusing on systematically modified versions of English to investigate learnability and model performance. Using the language modeling task with English on GPT-2 small architecture and its impossible variants, Kallini et al. (2024) demonstrate that natural English is consistently easier to learn than its impossible counterparts, as reflected in lower perplexity scores. They conclude that the above critique of language models as cognitive models is largely invalid.

2.2 Multilingual Language Modeling

Whether languages vary in complexity remains a controversial topic, and linguists have taken different approaches to address this question (e.g., McWhorter, 2001, 2011; Newmeyer, 2021; Joseph and Newmeyer, 2012). While most generative linguists argue that Universal Grammar requires that all languages be equally complex, others have challenged this notion (Gil, 2008).[2]

Initial computational attempts to examine language complexity using language models were limited to RNN-based architectures (Cotterell et al., 2018; Mielke et al., 2019; Johnson et al., 2021) and n-grams (Koplenig and Wolfer, 2023). These studies suggest that language complexity correlates with morphological richness and the size of speaker populations. More recently, Arnett and Bergen (2025) investigated why morphologically rich languages are harder to model. By testing monolingual language models trained on carefully curated comparative datasets (Chang et al., 2024),

[1] Our code and data are available at `temporary-pseudo-url`.

[2] See Newmeyer (2021) for a more thorough discussion.

they found that morphological features alone could not predict language learnability when training data size was controlled.

While valuable, previous studies often rely on comparative corpora, introducing inconsistencies across languages. Even with parallel corpora (Mielke et al., 2019), studies are limited by small datasets and outdated models. Our study addresses these gaps using a larger parallel corpus and modern transformer architectures.

3 Data and Implementation Details

3.1 Parallel Data Construction: OPUS12 and OPUS30

A key challenge in addressing our questions is that different languages texts drawn from different sources will have different amounts of information. To control for this, we construct two sentence-aligned multilingual parallel corpora to ensure that all languages in our dataset match the content. This allows us to isolate the effect of how information is conveyed, specifically, the role of grammar, in learnability differences between languages.

We name the two parallel corpora **OPUS12** and **OPUS30**, gathering aligned sentences from five corpora available on OPUS (Tiedemann, 2012): NLLB (Schwenk et al., 2021), TED2020 (Reimers and Gurevych, 2020), the Bible (Christodouloupoulos and Steedman, 2015), OpenSubtitles (Lison and Tiedemann, 2016), and CCAligned (El-Kishky et al., 2020). Since overlap among languages decreases as more languages are included, we decide to select a minimum of 10M words in English as a standard for our parallel corpora. 10M words also correspond to the amount of input of children's first 2 to 5 years of development (Warstadt et al., 2023).

OPUS12 is a 12-language multilingual sentence-aligned corpus[3]. There are around 10M words in the case of English. OPUS30 contains 30 languages with a much smaller size: 48K sentences with 0.7M words. While the two datasets share overlapping languages, their sentences do not overlap, making OPUS30 a suitable test set for additional language modeling experiments.

After deduplicating and removing English sentences from non-English data split using FastText (Joulin et al., 2017), we report the statistics of our corpora in Table 1.

[3]The languages and their typological information are listed in Appendix C.

Data Source	OPUS12		OPUS30	
	# Sent	# Word	# Sent	# Word
NLLB	5K	0.1M	16	368
TED2020	164K	2.9M	11K	182K
Bible	40K	1M	14K	324K
OpenSubtitles	680K	4.5M	15K	60K
CCAligned	117K	1.6M	8K	111K
Overall	1M	10.1M	48K	0.7M

Table 1: Data sources of OPUS12 and OPUS30. The word counts are based on the English data. See Appendix C for licensing information.

3.2 Validation Experiment

To ensure the reliability of our findings presented in the remainder of this paper, we replicate experiments in Kallini et al. (2024) using a scaled-down version of their original corpus (10M words). We find a perfect rank correlation between our results and Kallini et al. (2024) (Spearman's $\rho = \mathbf{1}, p < 0.001$). The results and detailed statistical analyses can be found in the Appendix A.

3.3 Model Architecture & Training

In our experiments, following Kallini et al. (2024), we trained standard GPT-2 small models for each language and evaluated its performance based on the perplexity over a parallel test split of 10K randomly sampled sentences. Due to limited computational resources, we trained each model using 3 random seeds instead of the 5 used in the original study and reduced the maximum training steps from 2000 to 1200 to avoid overfitting and adjusted the warmup steps proportionally to 120.[4]

3.4 Multilingual Tokenization

Given our multilingual experiments, tokenization is crucial for fair comparison. To avoid bias toward Latin-script languages, which are overrepresented in our study, we opted against using a multilingual tokenizer with a shared vocabulary.

Previous monolingual experiments either set the vocabulary size of tokenizers to be the same across languages (Arnett and Bergen, 2025) or applied the formula $0.4 \times |V|$ (Koplenig et al., 2023; Mielke et al., 2019), where $|V|$ represents the number of unique word types. We conducted a series of pilot experiments on tokenization and found neither

[4]We did not experiment with alternative warmup steps, as Kallini et al. (2024) demonstrated that changing the warmup schedule does not affect the ranking of perplexities for impossible language models.

Group	Language	Definition
Ours	SHUFFLE_LOCAL (W=2)	The sentence is reordered with every two tokens reversed in order.
	REVERSE_FULL	Every word is reversed in order in a sentence.
K+	SHUFFLE_DETERMINISTIC (S=84)	The sentence is deterministically shuffled by length with seed 84.
	SHUFFLE_DETERMINISTIC (S=57)	The sentence is deterministically shuffled by length with seed 57.
	SHUFFLE_DETERMINISTIC (S=21)	The sentence is deterministically shuffled by length with seed 21.
	SHUFFLE_LOCAL (W=10)	The sentence is deterministically shuffled in local window size being 10.
	SHUFFLE_LOCAL (W=5)	The sentence is deterministically shuffled in local window size being 5.
	SHUFFLE_LOCAL (W=3)	The sentence is deterministically shuffled in local window size being 3.
	SHUFFLE_EVEN_ODD	The sentence is reordered with even-indexed tokens first, then odd-indexed.

Table 2: Overview of impossible languages in our Experiment1 and Experiment2. **K+** languages are borrowed from Kallini et al. (2024) and the rest are new variants introduced in our experiments.

approach suitable for our experimental design (Details can be found in Appendix B).

Given these challenges, we decided to use pretrained tokenizers. The rationale behind this choice is that when the tokenizer training data is sufficiently large and diverse, the resulting tokenization scheme should be equally good across languages, as long as the tokenizer algorithm and hyperparameters (e.g., vocabulary size, subword strategy) remain the same.[5] Some may argue that the BPE algorithm might not be optimized for agglutinative languages such as Turkish, which makes the cross-linguistic comparison unfair. However, much literature on cross-linguistic LM comparison adopts BPT tokenizers (e.g., Mielke et al., 2019; Arnett and Bergen, 2025). As an additional check, we use token counts per word (TCW; reported in Appendix E Table 7) to measure the morphological complexity of a language and report the correlation between the perplexity and TCW. The results show the correlation is not significant (see Section5), suggesting that the morphological complexity of a language does not substantially impact its learnability.

While defining *sufficiently large and diverse* is difficult, we consider the size of the training data for GPT2 (Radford et al., 2019) as a reference point, as English was a high-resource language even in 2019 when the paper was published. We believe that this data size is sufficient to minimize differences that tokenization will make across languages.

When selecting pretrained tokenizers, we primarily use **monolingual BPE** tokenizers.[6] Our goal is to maintain a relatively consistent vocabulary size across languages, though we make exceptions for Romanian, Arabic, and Chinese due to limited model availability. The training data for all other languages is at least as large as the English corpus.

4 Experiment 1: Attested vs. Impossible Languages in an Intra-Language Modeling Setting

4.1 Impossible Languages

In this experiment, we use the deterministic shuffled languages from Kallini et al. (2024) along with two new variants (see Table 2). We include shuffled languages because (1) Kallini et al. (2024) identify them as the *most* impossible languages in their language possibility ranking, and (2) their difficulty is also indirectly supported by empirical studies showing that both adults and children exhibit a regularization bias, meaning they tend to acquire grammars with minimal variation (Newmeyer, 2005; Singleton and Newport, 2004).

Since all languages are deterministically shuffled, the original ones (=attested ones) can be recovered from their variants through another deterministic function. If LLMs function as non-human-like pattern recognizers as Chomsky et al. (2023); Moro et al. (2023) argue, they should be able to learn these deterministic languages.

4.2 Results & Discussion

The results are presented in Figure 1. First, in all languages except Italian, the perplexity of the attested language is lower than all its impossible variants. For Italian, SHUFFLE_LOCAL (W=2) yields a slightly lower perplexity than natural

[5] Although tokenization quality, measured by metrics like compression (Schmidt et al., 2024) and Rényi entropy (Zouhar et al., 2023), has been linked to language modeling performance (e.g., Liang et al., 2023; Goldman et al., 2024), recent studies challenge this connection (Arnett and Bergen, 2025).

[6] However, for Chinese, we follow previous studies (Mielke et al., 2019) and use the Chinese-BERT tokenizer.

Figure 1: Attested individual Language vs. their corresponding counterparts with a 95% confidence interval over 3 random seeds tested on 10k sentences from OPUS30.

Italian, though the difference is not statistically significant (Mann-Whitney U test: W = 63, p = 0.353). Additionally, attested languages exhibit smaller error bars compared to their impossible counterparts, indicating more stable learning. Third, we observe a consistent learnability pattern across languages. For each language, smaller shuffling windows result in lower perplexity. Moreover, most languages perturbed with SHUFFLE_DETERMINISTIC are harder to model than SHUFFLE_LOCAL ones. We assume this is because SHUFFLE_DETERMINISTIC languages are shuffled based on sequence length and autoregressive models do not have direct access to this information when predicting the next word, making language modeling harder.

We also conducted a Spearman's Ranking Correlation test between the results on OPUS30 English and those from Kallini et al. (2024)'s experiments. We observe that the ranking of our English impossible variants aligns perfectly with that reported by Kallini et al. (2024) when the rank variation within random seeds is ignored (ρ = 1, p = 0.0027). Additionally, the observation that the smaller the window size for locally shuffled English, the lower the perplexity aligns with Kallini et al. (2024)'s experiment results.

Based on these findings, we answer the first sub-question: **language models can (largely) distinguish between each attested language and their corresponding impossible counterparts.**

5 Experiment 2: Attested vs. Impossible Languages in an Inter-Language Modeling Setting

In this experiment, we gather the learning results of all possible and impossible languages to see if there is a separation boundary between them. If GPT-2 small can distinguish between possible and impossible languages, we expect the former's perplexity to be lower than the latter's.

The results of different language models are shown in Figure 2.[7] Not every language is equally

[7]To highlight the overlap of perplexity between attested languages and impossible ones, we zoom in on the lower perplexity range while displaying higher perplexity values in a separate, compressed section with a break in the y-axis.

Langs	Attested		Example
	Typo	Theo	
PERTURB_NNDA	NO	NO	`She enjoyed books three the fantastically interesting a lot .`
PERTURB_ANND	NO	NO	`She enjoyed fantastically interesting three books the a lot .`
PERTURB_DANN	FEW	YES	`She enjoyed the fantastically interesting books three a lot .`
DPERTURB_DNAN	MANY	YES	`She enjoyed the three fantastically interesting books a lot .`
PERTURB_DNNA	MANY	YES	`She enjoyed the books three fantastically interesting a lot .`
NP_RANDOM	NO	NO	`She enjoyed books fantastically three interesting the a lot .`

Table 3: List of NP-perturbations with corresponding categories and examples.

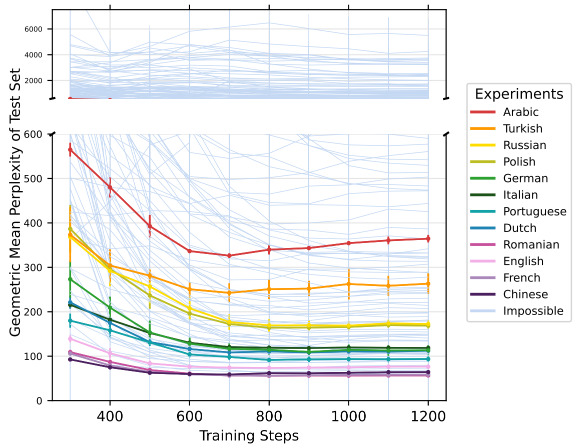

Figure 2: Attested natural languages vs. impossible languages with a 95% confidence interval over 3 random seeds. The x-axis represents the training steps, and the y-axis shows the perplexity on the test split. All the impossible languages are marked in light blue.

easy to learn: Chinese is the easiest language, while Arabic is the hardest, followed by Turkish and Russian. We also observe a moderate positive correlation between the average number of tokens per word (TCW) and perplexity of each of the last checkpoints in 11 languages (Chinese is excluded because the BERT tokenizer is a character-level tokenizer), as indicated by a Spearman's rank test ($\rho = 0.569$), but it is not significant ($p = 0.067$). This finding aligns with the observation by Arnett and Bergen (2025) that there is no significant difference in language modeling difficulty of agglutinative vs. fusional languages when the amount of information is controlled. Although all the attested languages are distributed at the bottom of the graph, we see some impossible languages fall between these attested languages. For example, Russian, Turkish and Arabic all show higher perplexity than English perturbed with shuffle_local (w=3). This means that for GPT-2 small, these impossible languages are are just as hard to learn as the attested languages.

To quantify the extent GPT-2 small can distinguish attested from impossible languages, we train a linear SVM classifier with the average perplexity value across the three random seeds of each checkpoint as features. The classifier reaches 0.70 macro F1 score averaged over 10-folds cross-validation.

Based on this experiment, we answer the second sub-question posed in our paper: **Although language models tend to learn attested languages better than impossible ones, their perplexity does not distinguish all attested languages from all impossible languages overall.**

6 Experiment 3: Attested vs. Unattested Languages

In this experiment, we investigate how well language models can learn **unattested languages**, languages whose structure is conceivable according to rules of grammar or morphology, but which have not been attested. While unattested languages are not necessarily unlearnable (e.g., Tsimpli and Smith, 1995), prior research suggests a link between typological feature frequency, cognitive biases, and language learnability (e.g., Gentner and Bowerman, 2009; Culbertson et al., 2012; Culbertson and Newport, 2015; Culbertson et al., 2020).

We focus on Greenberg's Universal 20 (Greenberg et al., 1963), which suggests that certain determiner-adjective-number-noun orders in an NP are universally unattested. As a well-studied typological phenomenon, Universal 20 serves as a good testbed for comparing human learners and language models, complementing experimental findings that show harmonic NP orders (i.e., the dependents always precedes/follow the head; e.g., NUM-ADJ-NOUN and NOUN-ADJ-NUM) are easier to learn than non-harmonic ones (e.g., NUM-NOUN-ADJ or ADJ-NOUN-NUM) (Culbertson and Newport, 2015, 2017; Culbertson et al., 2020). One influential hypothesis, the Typological Prevalence Hypothesis, proposes that more common typological patterns

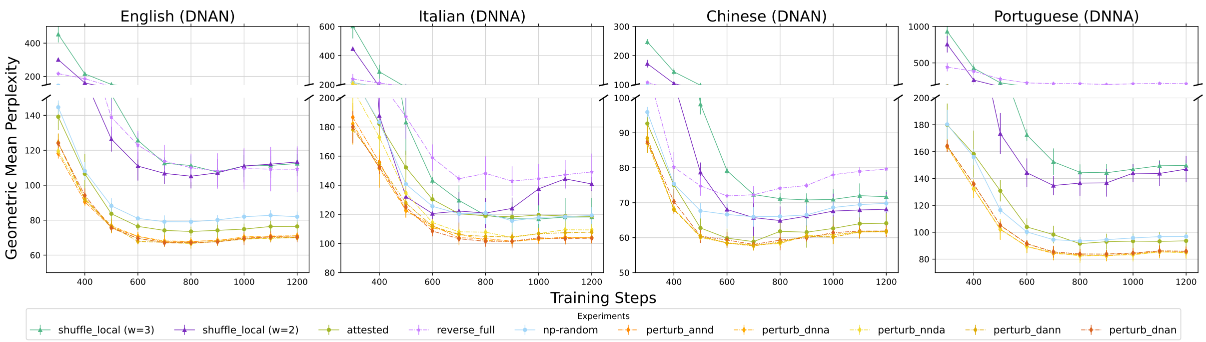

Figure 3: Attested individual natural languages vs. their corresponding impossible languages with a 95% confidence interval over 3 random seeds. The x-axis represents the training steps, and the y-axis shows the perplexity on the test split. Different language types are distinguished using distinct color palettes.

are easier to learn (Gentner and Bowerman, 2009). Therefore, we predict that if language models exhibit similar biases as humans, we expect a gradient of difficulty in learning different NP orders, with some unattested configurations posing greater challenges than others.

Among the 24 theoretically possible orders of adjectives, nouns, determiners, and numbers, we select five combinations, covering cases classified as FEW, MANY, and ZERO in Cinque (2005)'s typological analysis.[8] In this experiment, we only permute words within NPs. If the perplexity of these permuted languages is similar to that of attested languages, it suggests two possible reasons: (1) Language models can learn these unattested languages; (2) the number of words in NP may be a small number with respect to the entire data size, and hence NP-internal perturbation introduces a much smaller noise compared to the entire data perturbation we used in previous experiments, which may not significantly affect the learnability of a language. To rule out the latter possibility, we also construct a control condition in which words corresponding to these POS categories are randomly shuffled within NPs. This language serves as a baseline, indicating the extent to which NP-internal permutations influence the learnability of a language.

Examples of perturbed NP word orders and their typological information are listed in Table 3 and their word orders are reported below:

- PERTURB_NNDA: NOUN>NUM>DET>ADJ.
- PERTURB_ANND: ADJ>NUM>NOUN>DET.
- PERTURB_DANN: DET>ADJ>NOUN>NUM.
- PERTURB_DNAN: DET>NUM>ADJ>NOUN, typical of English and Chinese.
- PERTURB_DNNA: DET>NUM>NOUN>ADJ, typical of Italian and Portuguese.
- NP_RANDOM: Random permutation of ADJ, NOUN, NUM, and DET within NPs.

Since identifying NP structures requires a constituency parser, we use Stanza (Qi et al., 2020) to parse raw text. Stanza provides constituency parsing for only Chinese, Portuguese, English, and Italian, with acceptable accuracy (>0.85)[9], so we limit our analysis to these four languages. As different parsers are trained on distinct treebanks with varying annotation guidelines, we select POS tags based on each treebank's guidelines. Details are provided in Appendix F.

Results Our results are visualized in Figure 3 (bottom subgraph). Surprisingly, all five NP-perturbed languages exhibit lower perplexity compared to their attested counterparts across all four languages. Two of these (NNDA and ANND) are unattested in typological studies and are ruled out by generative approaches (Cinque, 2005), but we do not observe a significant difference in perplexity between the three languages with attested NP orders and the two languages with unattested orders.

When these POS tags are shuffled within NPs, perplexity increases, reaching or exceeding the perplexity of the attested languages. This rules out the possibility that limited perturbations simply do not affect model training.

To summarize, this experiment shows that, unlike humans, **language models fail to show a gra-**

[8] Although Cinque (2005) seeks to explain why ZERO languages really are "underivable" under the minimalist program we refer to them as *unattested* to contrast them with the impossible languages of the previous section, i.e., ones that involve random shuffling or reversed word order.

[9] https://stanfordnlp.github.io/stanza/constituency.html

dient of difficulty in learning different NP orders according to their typological prevalence.

Discussion: Why Can't LMs Distinguish Between Attested and Unattested Languages? To make sense of these results, we propose two key factors that influence LM learning outcomes: *randomness* and *constituency structure*. By *randomness*, we refer to whether the perturbation function produces a perturbed text that can be deterministically recovered to its original form. By *constituency structure*, we mean whether the phrase structures of the original language are preserved in the perturbed version.

Regarding randomness, string distributions with higher entropy are harder to learn. This explains why NP-perturbed unattested languages show lower perplexity than attested languages and NP_RANDOM variants. The reasoning is that our perturbation procedure enforces a strict ordering procedure, which may be (sometimes) violated in the original attested language. For example, although English is a DNAN language, certain constructions such as the DANN (DET-ADJ-NUM-NOUN; e.g., *a beautiful five days in Austin*) does not follow the dominant pattern. Once POS tags orders are normalized within NPs, the resulting constructions become more predictable. Therefore, all normalized NPs, including our unattested NPs, may have lower overall entropy, which could explain why they are easier to learn. In fact, the normalized DNAN, which has the same word order as English, shows lower perplexity than the original, unnormalized English; and the same applies to our other languages in this experiment.

Regarding constituency structure, we hypothesize that disrupting constituency weakens local dependency relations within phrase structures. This explains why all impossible languages in the previous experiments, despite maintaining a deterministic order, still results in higher perplexity than NP-perturbed languages (Figure 3). Similarly, this may also explain the higher perplexity of count-based grammars in Kallini et al. (2024): the insertion of morphological markers disrupts phrase structure integrity. One exception is REVERSE_FULL, which preserves constituency structure while maintaining a deterministic order. As shown in Figure 3, REVERSE_FULL exhibits perplexity closer to unattested languages. We hypothesize that this exception may be due to other factors, such as *information density* (Clark et al., 2023). Since reversing word order alters information flow, it may obscure more accessible information. For example, predicting a pronoun given a preceding noun is easier than predicting a noun given a preceding pronoun, potentially increasing difficulty for the LM. One shortcoming of this experiment is that it was based entirely on NP perturbations(Kallini et al., 2024). Future work could investigate whether similar effects occur when POS order or constituency structure is disrupted on a larger scale.

7 Conclusion

In this paper, we extend the initial work of Kallini et al. (2024) to a broader multilingual context. Our experiments provide mixed results that complement those of Kallini et al. (2024). We find that while GPT-2 small can distinguish between attested languages and their impossible variants, its learning outcomes do not separate all attested and all unattested or impossible languages. That being said, LMs do tend to learn attested languages better, on average, than impossible languages, and we achieve a separability of 0.7 between the two classes based on the models' perplexity. Finally, we observe that some unattested languages show lower perplexity than their attested counterparts even though they exhibit NP orderings that flout Greenberg's Universal 20.

What to make of these results in the context of our original question–whether LLMs can serve as cognitive models? While our results show that GPT-2 does not behave as we might expect from a fully human-like learner, they also demonstrate that it has a soft preference for attested over impossible languages. Skeptics have previously likened LLMs to a bad theory of physics in which "anything goes."[10] In line with Kallini et al. (2024), our results demonstrate that these models do not instantiate an "anything goes" hypothesis. Rather, their incremental data-processing architectures represent a useful starting point for studying human language processing. Refining models to achieve stronger alignment with people is possible, and will likely lead to lasting insights about human cognitive architecture.

[10] Chomsky, quoted from an email to Gary Marcus: *You can't go to a physics conference and say: I've got a great theory. It accounts for everything and is so simple it can be captured in two words: "Anything goes." All known and unknown laws of nature are accommodated, no failures. Of course, everything impossible is accommodated also.*

8 Limitations

We acknowledge that our experiments rely on GPT-2 Small, which may not generalize to larger models. This choice was made for two reasons: (1) running experiments across multiple languages is computationally expensive; (2) we aimed for comparability with Kallini et al. (2024). Future work could explore whether our findings hold for larger models. Additionally, the dataset used for training the language model is relatively small. This is a deliberate trade-off between data size and linguistic diversity. While a larger dataset might yield more robust results, our approach ensures broader typological coverage. Lastly, in our experiments on unattested languages, we generated synthetic data by perturbing languages based on Universal 20. However, linguistic correlations extend beyond word order universals. For instance, Greenbergian correlations (Dryer, 1992) suggest that verb-object order often correlates with other features such as adposition-noun phrase order and determiner-noun phrase order. Future work will explore more nuanced perturbations to better capture such cross-linguistic dependencies.

9 Ethical Statement

Our research adheres to ethical guidelines in data collection, model development, and evaluation. We use publicly available datasets, ensuring that no private or personally identifiable information is included. Our dataset selection prioritizes linguistic diversity while maintaining data transparency. Regarding computational resources, we use GPT-2 small trained on A-100 and V-100 GPUs. Each experiment on each language took around 10-12 hours. Finally, our research is intended for advancing linguistic understanding in computational models and does not facilitate any malicious applications. We encourage responsible usage and open discussions on the ethical implications of NLP research.

References

Wissam Antoun, Fady Baly, and Hazem Hajj. 2021. AraGPT2: Pre-trained transformer for Arabic language generation. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 196–207, Kyiv, Ukraine (Virtual). Association for Computational Linguistics.

Catherine Arnett and Benjamin Bergen. 2025. Why do language models perform worse for morphologically complex languages? In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6607–6623, Abu Dhabi, UAE. Association for Computational Linguistics.

Florbela Barreto, António Branco, Eduardo Ferreira, Amália Mendes, Maria Fernanda Bacelar do Nascimento, Filipe Nunes, and João Ricardo Silva. 2006. Open resources and tools for the shallow processing of Portuguese: The TagShare project. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).

Nadav Borenstein, Anej Svete, Robin Chan, Josef Valvoda, Franz Nowak, Isabelle Augenstein, Eleanor Chodroff, and Ryan Cotterell. 2024. What languages are easy to language-model? a perspective from learning probabilistic regular languages. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15115–15134, Bangkok, Thailand. Association for Computational Linguistics.

Tyler A. Chang, Catherine Arnett, Zhuowen Tu, and Benjamin K. Bergen. 2024. Goldfish: Monolingual language models for 350 languages. *Preprint*, arXiv:2408.10441.

Noam Chomsky and Andrea Moro. 2022. *The secrets of words*. MIT Press.

Noam Chomsky, Lan Roberts, and Jeffrey Watumull. 2023. Noam Chomsky: The false promise of ChatGPT. *The New York Times*. Accessed: 2024-12-16.

Christos Christodouloupoulos and Mark Steedman. 2015. A massively parallel corpus: the Bible in 100 languages. *Language resources and evaluation*, 49:375–395.

Guglielmo Cinque. 2005. Deriving Greenberg's Universal 20 and its exceptions. *Linguistic inquiry*, 36(3):315–332.

Thomas Hikaru Clark, Clara Meister, Tiago Pimentel, Michael Hahn, Ryan Cotterell, Richard Futrell, and Roger Levy. 2023. A cross-linguistic pressure for Uniform Information Density in word order. *Transactions of the Association for Computational Linguistics*, 11:1048–1065.

Ryan Cotterell, Sabrina J. Mielke, Jason Eisner, and Brian Roark. 2018. Are all languages equally hard to language-model? In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 536–541, New Orleans, Louisiana. Association for Computational Linguistics.

Jennifer Culbertson, Julie Franck, Guillaume Braquet, Magda Barrera Navarro, and Inbal Arnon. 2020. A learning bias for word order harmony: Evidence from speakers of non-harmonic languages. *Cognition*, 204:104392.

Jennifer Culbertson and Elissa L Newport. 2015. Harmonic biases in child learners: In support of language universals. *Cognition*, 139:71–82.

Jennifer Culbertson and Elissa L Newport. 2017. Innovation of word order harmony across development. *Open Mind*, 1(2):91–100.

Jennifer Culbertson, Paul Smolensky, and Géraldine Legendre. 2012. Learning biases predict a word order universal. *Cognition*, 122(3):306–329.

Forrest Davis and Marten van Schijndel. 2020. Recurrent neural network language models always learn English-like relative clause attachment. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1979–1990, Online. Association for Computational Linguistics.

Iria de Dios-Flores, Juan Garcia Amboage, and Marcos Garcia. 2023. Dependency resolution at the syntax-semantics interface: psycholinguistic and computational insights on control dependencies. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 203–222, Toronto, Canada. Association for Computational Linguistics.

Rodolfo Delmonte, Antonella Bristot, and Sara Tonelli. 2007. Vit-venice italian treebank: Syntactic and quantitative features. In *Sixth International Workshop on Treebanks and Linguistic Theories*, volume 1, pages 43–54. Northern European Association for Language Technol.

Vittoria Dentella, Fritz Guenther, and Evelina Leivada. 2024. Language in vivo vs. in silico: Size matters but larger language models still do not comprehend language on a par with humans. *arXiv preprint arXiv:2404.14883*.

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.

Matthew S Dryer. 1992. The Greenbergian word order correlations. *Language*, 68(1):81–138.

Stefan Dumitrescu. 2024. Gpt-neo romanian 780m. Hugging Face. Available at: https://huggingface.co/dumitrescustefan/gpt-neo-romanian-780m.

Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. 2020. CCAligned: A massive collection of cross-lingual web-document pairs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 5960–5969, Online. Association for Computational Linguistics.

Danny Fox and Roni Katzir. 2024. Large language models and theoretical linguistics. *Theoretical Linguistics*, 50(1-2):71–76.

Richard Futrell and Kyle Mahowald. 2025. How linguistics learned to stop worrying and love the language models. *arXiv preprint arXiv:2501.17047*.

Dedre Gentner and Melissa Bowerman. 2009. Why some spatial semantic categories are harder to learn than others: The typological prevalence hypothesis. *Crosslinguistic approaches to the psychology of language: Research in the tradition of Dan Isaac Slobin*, 465:480.

Edward Gibson, Richard Futrell, Steven P Piantadosi, Isabelle Dautriche, Kyle Mahowald, Leon Bergen, and Roger Levy. 2019. How efficiency shapes human language. *Trends in cognitive sciences*, 23(5):389–407.

David Gil. 2008. How complex are isolating languages. *Language*.

Omer Goldman, Avi Caciularu, Matan Eyal, Kris Cao, Idan Szpektor, and Reut Tsarfaty. 2024. Unpacking tokenization: Evaluating text compression and its correlation with model performance. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2274–2286, Bangkok, Thailand. Association for Computational Linguistics.

Adam Goodkind and Klinton Bicknell. 2018. Predictive power of word surprisal for reading times is a linear function of language model quality. In *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)*, pages 10–18, Salt Lake City, Utah. Association for Computational Linguistics.

Joseph H Greenberg et al. 1963. Some universals of grammar with particular reference to the order of meaningful elements. *Universals of language*, 2:73–113.

Yeb Havinga. 2023. GPT Neo 1.3B pre-trained on cleaned Dutch mC4. Accessed: 2024-12-10.

Jennifer Hu and Roger Levy. 2023. Prompting is not a substitute for probability measurements in large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5040–5060, Singapore. Association for Computational Linguistics.

Carla L Hudson Kam and Elissa L Newport. 2005. Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language learning and development*, 1(2):151–195.

iGeniusAI. 2024. Italia-9b-instruct-v0.1. https://huggingface.co/iGeniusAI/Italia-9B-Instruct-v0.1. Accessed: 2024-06-17.

Tamar Johnson, Kexin Gao, Kenny Smith, Hugh Rabagliati, and Jennifer Culbertson. 2021. Investigating the effects of i-complexity and e-complexity on the learnability of morphological systems. *J. Lang. Model.*, 9.

John E Joseph and Frederick J Newmeyer. 2012. 'All languages are equally complex'. *Historiographia linguistica*, 39.

Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.

Julie Kallini, Isabel Papadimitriou, Richard Futrell, Kyle Mahowald, and Christopher Potts. 2024. Mission: Impossible language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14691–14714, Bangkok, Thailand. Association for Computational Linguistics.

Roni Katzir. 2023. Why large language models are poor theories of human linguistic cognition: A reply to piantadosi. *Biolinguistics*, 17:1–12.

H. Toprak Kesgin, M. Kaan Yuce, Eren Dogan, M. Egemen Uzun, Atahan Uz, H. Emre Seyrek, Ahmed Zeer, and M. Fatih Amasyali. 2024. Introducing cosmosGPT: Monolingual training for Turkish language models. In *2024 International Conference on INnovations in Intelligent SysTems and Applications (INISTA)*, page 1–6. IEEE.

Najoung Kim and Tal Linzen. 2020. COGS: A compositional generalization challenge based on semantic interpretation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9087–9105, Online. Association for Computational Linguistics.

Christo Kirov and Ryan Cotterell. 2018. Recurrent neural networks in linguistic theory: Revisiting pinker and prince (1988) and the past tense debate. *Transactions of the Association for Computational Linguistics*, 6:651–665.

Alexander Koplenig and Sascha Wolfer. 2023. Languages with more speakers tend to be harder to (machine-) learn. *Scientific Reports*, 13(1):18521.

Alexander Koplenig, Sascha Wolfer, and Peter Meyer. 2023. A large quantitative analysis of written language challenges the idea that all languages are equally complex. *Scientific Reports*, 13(1):15351.

Tatsuki Kuribayashi, Yohei Oseki, and Timothy Baldwin. 2024. Psychometric predictive power of large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1983–2005, Mexico City, Mexico. Association for Computational Linguistics.

Tatsuki Kuribayashi, Yohei Oseki, Souhaib Ben Taieb, Kentaro Inui, and Timothy Baldwin. 2025. Large language models are human-like internally. *arXiv preprint arXiv:2502.01615*.

Nur Lan, Emmanuel Chemla, and Roni Katzir. 2024. Large language models and the argument from the poverty of the stimulus. *Linguistic Inquiry*, pages 1–56.

Julien Launay, E.l. Tommasone, Baptiste Pannier, François Boniface, Amélie Chatelain, Alessandro Cappelli, Iacopo Poli, and Djamé Seddah. 2022. PAGnol: An extra-large French generative model. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4275–4284, Marseille, France. European Language Resources Association.

Davis Liang, Hila Gonen, Yuning Mao, Rui Hou, Naman Goyal, Marjan Ghazvininejad, Luke Zettlemoyer, and Madian Khabsa. 2023. XLM-V: Overcoming the vocabulary bottleneck in multilingual masked language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13142–13152, Singapore. Association for Computational Linguistics.

Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).

Ricardo Lopes, Joao Magalhaes, and David Semedo. 2024. GlórIA: A generative and open large language model for Portuguese. In *Proceedings of the 16th International Conference on Computational Processing of Portuguese - Vol. 1*, pages 441–453, Santiago de Compostela, Galicia/Spain. Association for Computational Lingustics.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330.

John H McWhorter. 2001. The worlds simplest grammars are creole grammars. *Linguistic Typology*, 5(2-3):125–166.

John H McWhorter. 2011. *Linguistic simplicity and complexity: Why do languages undress?*, volume 1. Walter de Gruyter.

Clara Meister, Tiago Pimentel, Patrick Haller, Lena Jäger, Ryan Cotterell, and Roger Levy. 2021. Revisiting the Uniform Information Density hypothesis. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 963–980, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Sabrina J. Mielke, Ryan Cotterell, Kyle Gorman, Brian Roark, and Jason Eisner. 2019. What kind of language is hard to language-model? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4975–4989, Florence, Italy. Association for Computational Linguistics.

Raphaël Millière. 2024. Language models as models of language. *arXiv preprint arXiv:2408.07144*.

Andrea Moro. 2023. Embodied syntax: impossible languages and the irreducible difference between humans and machines. *Sistemi intelligenti*, 35(2):321–328.

Andrea Moro, Matteo Greco, and Stefano F Cappa. 2023. Large languages, impossible languages and human brains. *Cortex*, 167:82–85.

Frederick J Newmeyer. 2005. *Possible and probable languages*. Oxford: Oxford.

Frederick J Newmeyer. 2021. Complexity and relative complexity in generative grammar. *Frontiers in Communication*, 6:614352.

Byung-Doh Oh and William Schuler. 2023a. Transformer-based language model surprisal predicts human reading times best with about two billion training tokens. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1915–1921, Singapore. Association for Computational Linguistics.

Byung-Doh Oh and William Schuler. 2023b. Why does surprisal from larger transformer-based language models provide a poorer fit to human reading times? *Transactions of the Association for Computational Linguistics*, 11:336–350.

Joe Pater. 2019. Generative linguistics and neural networks at 60: Foundation, friction, and fusion. *Language*, 95(1):e41–e74.

Lisa Pearl and Benjamin Mis. 2011. How far can indirect evidence take us? Anaphoric one revisited. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33.

Steven T Piantadosi. 2023. Modern language models refute chomsky's approach to language. *From fieldwork to linguistic theory: A tribute to Dan Everett*, pages 353–414.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Nils Reimers and Iryna Gurevych. 2020. Making monolingual sentence embeddings multilingual using knowledge distillation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.

Jenny Saffran, Marc Hauser, Rebecca Seibel, Joshua Kapfhamer, Fritz Tsao, and Fiery Cushman. 2008. Grammatical pattern learning by human infants and cotton-top tamarin monkeys. *Cognition*, 107(2):479–500.

Craig W Schmidt, Varshini Reddy, Haoran Zhang, Alec Alameddine, Omri Uzan, Yuval Pinter, and Chris Tanner. 2024. Tokenization is more than compression. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 678–702, Miami, Florida, USA. Association for Computational Linguistics.

Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, and Angela Fan. 2021. CCMatrix: Mining billions of high-quality parallel sentences on the web. 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 6490–6500, Online. Association for Computational Linguistics.

Stefan Schweter. 2020. German GPT-2 model (dbmdz/german-gpt2). `https://huggingface.co/dbmdz/german-gpt2`. Accessed: 2025-01-05.

Jenny L Singleton and Elissa L Newport. 2004. When learners surpass their models: The acquisition of american sign language from inconsistent input. *Cognitive psychology*, 49(4):370–407.

Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).

Ianthi-Maria Tsimpli and Neil Smith. 1995. *Mind of a Savant: Language, Learning and Modularity*. Blackwell.

Eleftheria Tsipidi, Franz Nowak, Ryan Cotterell, Ethan Wilcox, Mario Giulianelli, and Alex Warstadt. 2024. Surprise! Uniform Information Density isn't the whole story: Predicting surprisal contours in long-form discourse. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18820–18836, Miami, Florida, USA. Association for Computational Linguistics.

Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, and Ryan Cotterell. 2023. Findings of the BabyLM challenge: Sample-efficient pretraining on developmentally plausible corpora. In *Proceedings* of the BabyLM Challenge at the 27th Conference on *Computational Natural Language Learning*, pages 1–34, Singapore. Association for Computational Linguistics.

Ethan G. Wilcox, Tiago Pimentel, Clara Meister, Ryan Cotterell, and Roger P. Levy. 2023. Testing the predictions of surprisal theory in 11 languages. *Transactions of the Association for Computational Linguistics*, 11:1451–1470.

Michał Wojczulis and Dariusz Kłeczek. 2021. papuGaPT2 - Polish GPT2 language model.

Naiwen Xue, Fei Xia, Fu-Dong Chiou, and Marta Palmer. 2005. The penn chinese treebank: Phrase structure annotation of a large corpus. *Natural language engineering*, 11(2):207–238.

Yuekun Yao and Alexander Koller. 2022. Structural generalization is hard for sequence-to-sequence models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5048–5062, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Dmitry Zmitrovich, Aleksandr Abramov, Andrey Kalmykov, Vitaly Kadulin, Maria Tikhonova, Ekaterina Taktasheva, Danil Astafurov, Mark Baushenko, Artem Snegirev, Tatiana Shavrina, Sergei S. Markov, Vladislav Mikhailov, and Alena Fenogenova. 2024. A family of pretrained transformer language models for Russian. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 507–524, Torino, Italia. ELRA and ICCL.

Vilém Zouhar, Clara Meister, Juan Gastaldi, Li Du, Mrinmaya Sachan, and Ryan Cotterell. 2023. Tokenization and the noiseless channel. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5184–5207, Toronto, Canada. Association for Computational Linguistics.

A Experiment Results of Replicating Kallini et al. (2024)

We implement the training and evaluation following the same experiment setting from Kallini et al. (2024). The result is shown in Figure 4.

We calculate Spearman's rank correlation between our results for the **shuffled* languages and those of Kallini et al. (2024) at every 200-step interval from 400 to 1,200. The Spearman's ρ is consistently **1** ($p < 0.001$), indicating perfect agreement between the rankings, showing that 10M words are sufficient enough to replicate the language modeling experiments using 100M words conducted by Kallini et al. (2024).

B Tokenization Pilot Experiments and Results

In our experiments, where we trained tokenizers for each language using 10M words (around 60MB data), testing vocabulary sizes ranging from 30K to 80K in increments of 10K, we observed two key findings: (1) tokenizers trained with around 60MB data resulted in unstable language modeling outcomes, and (2) different languages require distinct optimal vocabulary sizes. These results are shown in Figure 5. Additionally, agglutinative languages like Turkish, with their large number of unique tokens, made large vocabulary sizes impractical. For instance, Turkish has three times the number of unique words as English (467K vs. 140K), and applying $0.4 \times |V|$ would result in a vocabulary size of 186K, which is too large for efficient language model training with the limited data available and a small model.

C Details of OPUS12 and OPUS30

The typological features of languages used in the two corpora are reported in Table 5. The licensing terms vary depending on their original sources, listed below.

- NLLB: ODC-By
- TED2020: CC BY–NC–ND 4.0 International; for details, see the official website.
- Bible: CC0 1.0
- OpenSubtitles: GNU General Public License v3.0
- MultiCCAligned: unknown; see the official website.

D Tokenizers

Table 6 shows the details of the tokenizers we use in the experiments.

E TCW

The TCW is reported in Table 7.

F POS tags of each treebank

Different constituency parsers are trained with different treebanks. We select POS-tags that are relevant to the four word classes. The detailed POS-tags for each language can be found in Table 4.

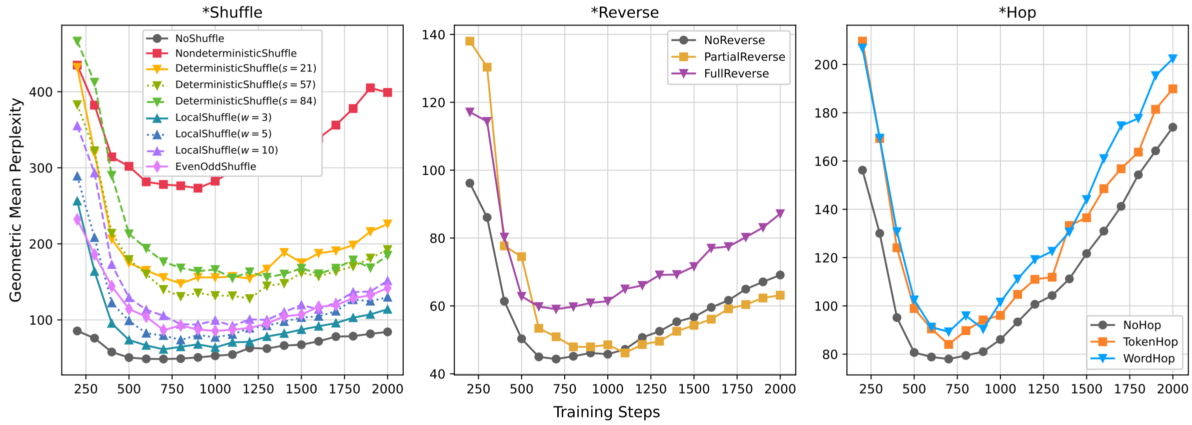

Figure 4: Replication of (Kallini et al., 2024) with 10M words from BabyLM Challenge dataset (strict-small track)

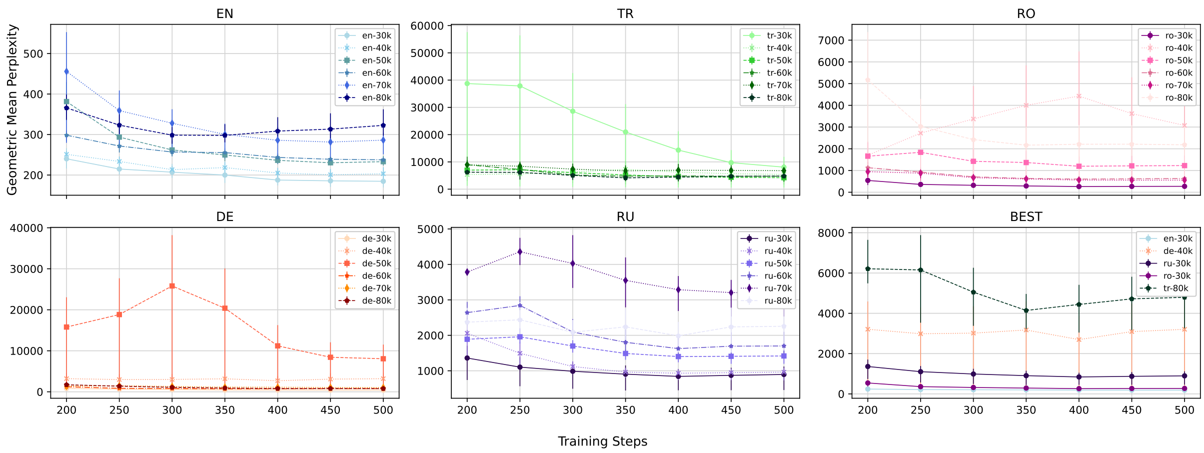

Figure 5: Perplexity results on the development set (10K sentences) for five languages (EN, TR, RO, DE, RU), trained on a 10M-sentence training set across different vocabulary sizes. Error bars represent the first and last quartiles (25% and 75%) of the results. A plot for the optimized vocabulary size (labeled 'BEST') is also included, showing high variance for TR and RU even with optimized vocabulary size.

Language	Treebank	POS-tags			
		DET	NUM	ADJ	NOUN
English	Penn Treebank (Marcus et al., 1993)	DT, PRP$, PDT, POS	QP, $, CD	RB, ADJP, JJR, JJS, JJ	NN, NNS, NNP, NNPS
Italian	VIT(Delmonte et al., 2007)	DET	NUM, SQ	ADJ, SA	NOUN, PRON, PROPN, SYM, X
Chinese	CTB 3.0(Xue et al., 2005)	DT, M, CLP, DP	CD, OD, QP	JJ, ADJP, DNP, DEC, DEG	NN, NP, NR, NT, PRP, PN, FW
Portuguese	Cintil (Barreto et al., 2006)	DET, D, DEM, POSS, POSS'	QNT, QNT', NUM, PERCENTP, PERCENTP', CARD, CARD'	ADJ, AP	N', NOUN, PRON

Table 4: POS-tag categories across languages

Language	Family	Word Order	Morphology
OPUS12			
English	Indo-European (Germanic)	SVO	Analytic
German	Indo-European (Germanic)	No dominant	Fusional
Russian	Indo-European (Slavonic)	SVO	Fusional
Romanian	Indo-European (Romance)	SVO	Fusional
Turkish	Turkic (Altaic)	SOV	Agglutinative
Dutch	Indo-European (Germanic)	No dominant	Fusional
Polish	Indo-European (Slavonic)	SVO	Fusional
Portuguese	Indo-European (Romance)	SVO	Fusional
Italian	Indo-European (Romance)	SVO	Fusional
French	Indo-European (Romance)	SVO	Fusional
Chinese	Sino-Tibetan	SVO	Analytic
Arabic	Afro-Asiatic (Semitic)	VSO	Root-based (nonconcatenative)
OPUS30			
Spanish	Indo-European (Romance)	SVO	Fusional
Czech	Indo-European (Slavonic)	SVO	Fusional
Bulgarian	Indo-European (Slavonic)	SVO	Fusional
Slovak	Indo-European (Slavonic)	SVO	Fusional
Serbian	Indo-European (Slavonic)	SVO	Fusional
Croatian	Indo-European (Slavonic)	SVO	Fusional
Ukrainian	Indo-European (Slavonic)	SVO	Fusional
Danish	Indo-European (Germanic)	SVO	Fusional
Swedish	Indo-European (Germanic)	SVO	Fusional
Greek	Indo-European (Hellenic)	No dominant	Fusional
Persian	Indo-European (Indo-Iranian)	SVO	Fusional
Lithuanian	Indo-European (Baltic)	SVO	Fusional
Vietnamese	Austroasiatic	SVO	Analytic
Hebrew	Afro-Asiatic (Semitic)	VSO	Root-based (nonconcatenative)
Hungarian	Uralic	SVO	Agglutinative
Indonesian	Austronesian	SVO	Analytic
Japanese	Japonic	SOV	Agglutinative
Korean	Koreanic	SOV	Agglutinative

Table 5: Typological features of the OPUS12 and OPUS30 corpora, with OPUS30 including 18 additional languages beyond those in OPUS12.

Language	|Vocab|	|Training|	Reference
Arabic	64,000	77GB	Antoun et al. (2021)
Turkish	50,257	100GB	Kesgin et al. (2024)
Russian	50,257	450GB	Zmitrovich et al. (2024)
Polish	50,257	47GB	Wojczulis and Kłeczek (2021)
German	50,000	51GB	Schweter (2020)
Italian	50,176	Trillions toks	iGeniusAI (2024)
Portegese	50,258	35B tokens	Lopes et al. (2024)
Dutch	50,257	151GB	Havinga (2023)
Romanian	64,000	40GB	Dumitrescu (2024)
English	50,257	40GB	Radford et al. (2019)
French	50,262	130GB	Launay et al. (2022)
Chinese	21,128	300GB	Devlin et al. (2019)

Table 6: Tokenizers, vocabulary sizes, and training data sizes, and references for each language tested in our experiments.

LANGS	AR	TR	RU	PL	DE	IT
TCW	2.19	2.05	2.05	1.9	1.61	1.40
LANGS	PT	NL	RO	EN	FR	
TCW	1.68	1.51	1.81	1.45	1.67	

Table 7: TCW per language by each of their pretrained tokenizer