

Cross-Linguistic Syntactic Evaluation of Word Prediction Models

Aaron Mueller¹ Garrett Nicolai^{1†} Panayiota Petrou-Zeniou²

Natalia Talmina² Tal Linzen^{1,2}

¹Department of Computer Science

²Department of Cognitive Science

Johns Hopkins University

{amueller, gnicola2, ppetrou1, talmina, tal.linzen}@jhu.edu

Abstract

A range of studies have concluded that neural word prediction models can distinguish grammatical from ungrammatical sentences with high accuracy. However, these studies are based primarily on monolingual evidence from English. To investigate how these models' ability to learn syntax varies by language, we introduce CLAMS (Cross-Linguistic Assessment of Models on Syntax), a syntactic evaluation suite for monolingual and multilingual models. CLAMS includes subject-verb agreement challenge sets for English, French, German, Hebrew and Russian, generated from grammars we develop. We use CLAMS to evaluate LSTM language models as well as monolingual and multilingual BERT. Across languages, monolingual LSTMs achieved high accuracy on dependencies without attractors, and generally poor accuracy on agreement across object relative clauses. On other constructions, agreement accuracy was generally higher in languages with richer morphology. Multilingual models generally underperformed monolingual models. Multilingual BERT showed high syntactic accuracy on English, but noticeable deficiencies in other languages.

1 Introduction

Neural networks can be trained to predict words from their context with much greater accuracy than the architectures used for this purpose in the past. This has been shown to be the case for both recurrent neural networks (Mikolov et al., 2010; Sundermeyer et al., 2012; Jozefowicz et al., 2016) and non-recurrent attention-based models (Devlin et al., 2019; Radford et al., 2019).

To gain a better understanding of these models' successes and failures, in particular in the domain of syntax, proposals have been made for testing the

models on subsets of the test corpus where successful word prediction crucially depends on a correct analysis of the structure of the sentence (Linzen et al., 2016). A paradigmatic example is subject-verb agreement. In many languages, including English, the verb often needs to agree in number (here, singular or plural) with the subject (asterisks represent ungrammatical word predictions):

(1) The key to the cabinets is/*are next to the coins.

To correctly predict the form of the verb (underlined), the model needs to determine that the head of the subject of the sentence—an abstract, structurally defined notion—is the word *key* rather than *cabinets* or *coins*.

The approach of sampling challenging sentences from a test corpus has its limitations. Examples of relevant constructions may be difficult to find in the corpus, and naturally occurring sentences often contain statistical cues (confounds) that make it possible for the model to predict the correct form of the verb without an adequate syntactic analysis (Gulordava et al., 2018). To address these limitations, a growing number of studies have used constructed materials, which improve experimental control and coverage of syntactic constructions (Marvin and Linzen, 2018; Wilcox et al., 2018; Futrell et al., 2019; Warstadt et al., 2019a).

Existing experimentally controlled data sets—in particular, those targeting subject-verb agreement—have largely been restricted to English. As such, we have a limited understanding of the effect of the cross-linguistic variability in neural networks' syntactic prediction abilities. In this paper, we introduce the Cross-Linguistic Assessment of Models on Syntax (CLAMS) data set, which extends the subject-verb agreement component of the Marvin and Linzen (2018) challenge set to French, German, Hebrew and Russian. By focusing on a single lin-

[†] Work done while at Johns Hopkins University. Now in the University of British Columbia's Linguistics Department.

guistic phenomenon in related languages,¹ we can directly compare the models' performance across languages. We see the present effort as providing a core data set that can be expanded in future work to improve coverage to other languages and syntactic constructions. To this end, we release the code for a simple grammar engineering framework that facilitates the creation and generation of syntactic evaluation sets.²

We use CLAMS to test two hypotheses. First, we hypothesize that a multilingual model would show transfer across languages with similar syntactic constructions, which would lead to improved syntactic performance compared to monolingual models. In experiments on LSTM language models (LMs), we do not find support for this hypothesis; contrarily, accuracy was lower for the multilingual model than the monolingual ones. Second, we hypothesize that language models would be better able to learn hierarchical syntactic generalizations in morphologically complex languages (which provide frequent overt cues to syntactic structure) than in morphologically simpler languages (Gulordava et al., 2018; Lorimor et al., 2008; McCoy et al., 2018). We test this using LSTM LMs we train, and find moderate support for this hypothesis.

In addition to our analysis of LSTM LMs, we demonstrate the utility of CLAMS for testing pre-trained word prediction models. We evaluate multilingual BERT (Devlin et al., 2019), a bidirectional Transformer model trained on a multilingual corpus, and find that this model performs well on English, has mixed syntactic abilities in French and German, and performs poorly on Hebrew and Russian. Its syntactic performance in English was somewhat worse than that of monolingual English BERT, again suggesting that interference between languages offsets any potential syntactic transfer.

2 Background and Previous Work

2.1 Word Prediction Models

Language models (LMs) are statistical models that estimate the probability of sequences of words—or, equivalently, the probability of the next word of the sentence given the preceding ones. Currently, the most effective LMs are based on neural networks that are trained to predict the next word in a

¹English, French, German and Russian are all Indo-European languages, and (Modern) Hebrew syntax exhibits European areal influence (for different perspectives, see Wexler 1990; Zuckermann 2006; Zeldes 2013).

²<https://github.com/aaronmueller/clams>

large corpus. Neural LMs are commonly based on LSTMs (Hochreiter and Schmidhuber, 1997; Sundermeyer et al., 2012) or non-recurrent attention-based architectures (Transformers, Vaswani et al. 2017). The results of existing studies comparing the performance of the two architectures on grammatical evaluations are mixed (Tran et al., 2018; van Schijndel et al., 2019), and the best reported syntactic performance on English grammatical evaluations comes from LMs trained with explicit syntactic supervision (Kuncoro et al., 2018, 2019). We focus our experiments in the present study on LSTM-based models, but view CLAMS as a general tool for comparing LM architectures.

A generalized version of the word prediction paradigm, in which a bidirectional Transformer-based encoder is trained to predict one or more words in arbitrary locations in the sentence, has been shown to be an effective pre-training method in systems such as BERT (Devlin et al., 2019). While there are a number of variations on this architecture (Raffel et al., 2019; Radford et al., 2019), we focus our evaluation on the pre-trained English BERT and multilingual BERT.

2.2 Acceptability Judgments

Human acceptability judgments have long been employed in linguistics to test the predictions of grammatical theories (Chomsky, 1957; Schütze, 1996). There are a number of formulations of this task; we focus on the one in which a speaker is expected to judge a contrast between two minimally different sentences (a minimal pair). For instance, the following examples illustrate the contrast between grammatical and ungrammatical subject-verb agreement on the second verb in a coordination of short (2a) and long (2b) verb phrases; native speakers of English will generally agree that the first underlined verb is more acceptable than the second one in this context.

(2) *Verb-phrase coordination:*

- a. The woman laughs and talks/*talk.
- b. My friends play tennis every week and then get/*gets ice cream.

In computational linguistics, acceptability judgments have been used extensively to assess the grammatical abilities of LMs (Linzen et al., 2016; Lau et al., 2017). For the minimal pair paradigm, this is done by determining whether the LM assigns a higher probability to the grammatical member of

the minimal pair than to the ungrammatical member. This paradigm has been applied to a range of constructions, including subject-verb agreement (Marvin and Linzen, 2018; An et al., 2019), negative polarity item licensing (Marvin and Linzen, 2018; Jumelet and Hupkes, 2018), filler-gap dependencies (Chowdhury and Zamparelli, 2018; Wilcox et al., 2018), argument structure (Kann et al., 2019), and several others (Warstadt et al., 2019a).

To the extent that the acceptability contrast relies on a single word in a particular location, as in (2), this approach can be extended to bidirectional word prediction systems such as BERT, even though they do not assign a probability to the sentence (Goldberg, 2019). As we describe below, the current version of CLAMS only includes contrasts of this category.

An alternative use of acceptability judgments in NLP involves training an encoder to classify sentences into acceptable and unacceptable, as in the Corpus of Linguistic Acceptability (CoLA, Warstadt et al. 2019b). This approach requires supervised training on acceptable and unacceptable sentences; by contrast, the prediction approach we adopt can be used to evaluate any word prediction model without additional training.

2.3 Grammatical Evaluation Beyond English

Most of the work on grammatical evaluation of word prediction models has focused on English. However, there are a few exceptions, which we discuss in this section. To our knowledge, all of these studies have used sentences extracted from a corpus rather than a controlled challenge set, as we propose. Gulordava et al. (2018) extracted English, Italian, Hebrew, and Russian evaluation sentences from a treebank. Dhar and Bisazza (2018) trained a multilingual LM on a concatenated French and Italian corpus, and tested whether grammatical abilities transfer across languages. Ravfogel et al. (2018) reported an in-depth analysis of LSTM LM performance on agreement prediction in Basque, and Ravfogel et al. (2019) investigated the effect of different syntactic properties of a language on RNNs’ agreement prediction accuracy by creating synthetic variants of English. Finally, grammatical evaluation has been proposed for machine translation systems for languages such as German and French (Sennrich, 2017; Isabelle et al., 2017).

3 Grammar Framework

To construct our challenge sets, we use a lightweight grammar engineering framework that we term **attribute-varying grammars** (AVGs). This framework provides more flexibility than the hard-coded templates of Marvin and Linzen (2018) while avoiding the unbounded embedding depth of sentences generated from a recursive context-free grammar (CFG, Chomsky 1956). This is done using *templates*, which consist of *preterminals* (which have *attributes*) and *terminals*. A *vary statement* specifies which preterminal attributes are varied to generate ungrammatical sentences.

Templates define the structure of the sentences in the evaluation set. This is similar to the expansions of the *S* nonterminal in CFGs. Preterminals are similar to nonterminals in CFGs: they have a left-hand side which specifies the name of the preterminal and the preterminal’s list of attributes, and a right-hand side which specifies all terminals to be generated by the preterminal. However, they are non-recursive and their right-hand sides may not contain other preterminals; rather, they must define a list of terminals to be generated. This is because we wish to generate all possible sentences given the template and preterminal definitions; if there existed any recursive preterminals, there would be an infinite number of possible sentences. All preterminals have an attribute list which is defined at the same time as the preterminal itself; this list is allowed to be empty. A terminal is a token or list of space-separated tokens.

The `vary` statement specifies a list of preterminals and associated attributes for each. Typically, we only wish to vary one preterminal per grammar such that each grammatical case is internally consistent with respect to which syntactic feature is varied. The following is a simple example of an attribute-varying grammar:

```
vary: V[]
S[]    → je V[1,s]
V[1,s] → pense
V[2,s] → penses
V[1,p] → pensons
V[2,p] → pensez
```

Preterminals are **blue** and attributes are **orange**. Here, the first statement is the `vary` statement. This is followed by a template, with the special *S* keyword in **red**. All remaining statements are preterminal definitions. All attributes are spec-

ified within brackets as comma-separated lists; these may be multiple characters and even multiple words long, so long as they do not contain commas. The output of this AVG is as follows (True indicates that the sentence is grammatical):

True	je pense
False	je penses
False	je pensons
False	je pensez

This particular grammar generates all possible verb forms because the attribute list for V in the vary statement is empty, which means that we may generate any V regardless of attributes. One may change which incorrect examples are generated by changing the vary statement; for example, if we change V[] to V[1], we would only vary over verbs with the 1 (first-person) attribute, thus generating *je pense* and **je pensons*. One may also add multiple attributes within a single vary preterminal (implementing a logical AND) or multiple semicolon-separated vary preterminals (a logical OR). Changing V[] to V[1, s] in the example above would generate all first-person singular V terminals (here, *je pense*). If instead we used V[1] ; V[s], this would generate all V terminals with either first-person and/or singular attributes (here, *je pense*, **je penses*, and **je pensons*).

4 Syntactic Constructions

We construct grammars in French, German, Hebrew and Russian for a subset of the English constructions from Marvin and Linzen (2018), shown in Figure 1. These are implemented as AVGs by native or fluent speakers of the relevant languages who have academic training in linguistics.³

A number of the constructions used by Marvin and Linzen are English-specific. None of our languages besides English allow relative pronoun dropping, so we are unable to compare performance across languages on reduced relative clauses (*the author the farmers like smile/*smiles*). Likewise, we exclude Marvin and Linzen’s sentential complement condition, which relies on the English-specific ability to omit complementizers (*the bankers knew the officer smiles/*smile*).

The Marvin and Linzen (2018) data set includes two additional structure-sensitive phenomena other than subject-verb agreement: reflexive anaphora

³The German grammar was created by a non-native speaker but was then validated by native speakers.

Simple Agreement:

The author laughs/*laugh.

Across a Prepositional Phrase:

The farmer near the parents smiles/*smile.

Across a Subject Relative Clause:

The officers that love the skater *smiles/smile.

Short Verb Phrase Coordination:

The senator smiles and laughs/*laugh.

Long Verb Phrase Coordination:

The manager writes in a journal every day and likes/*like to watch television shows.

Across Object Relative Clause:

The farmer that the parents love swims/*swim.

Within Object Relative Clause:

The farmer that the parents *loves/love swims.

Figure 1: Syntactic constructions used in CLAMS. Only English examples are shown; for examples in other languages, see Appendix A. Ungrammatical forms are marked with asterisks.

and negative polarity item licensing. We do not include reflexive anaphora, as our languages vary significantly in how those are implemented. French and German, for example, do not distinguish singular from plural third-person reflexive pronouns. Similarly, negative polarity items (NPIs) have significantly different distributions across languages, and some of our evaluation languages do not even have items comparable to English NPIs (Giannakidou, 2011).

We attempt to use translations of all terminals in Marvin and Linzen (2018). In cases where this is not possible (due to differences in LM vocabulary across languages), we replace the word with another in-vocabulary item. See Appendix D for more detail on vocabulary replacement procedures.

For replicability, we observe only third-person singular vs. plural distinctions (as opposed to all possible present-tense inflections) when replicating the evaluation sets of Marvin and Linzen (2018) in any language.

5 Experimental Setup

5.1 Corpora

Following Gulordava et al. (2018), we download recent Wikipedia dumps for each of the languages,

strip the Wikipedia markup using WikiExtractor,⁴ and use TreeTagger⁵ to tokenize the text and segment it into sentences. We eliminate sentences with more than 5% unknown words.

Our evaluation is within-sentence rather than across sentences. Thus, to minimize the availability of cross-sentential dependencies in the training corpus, we shuffle the preprocessed Wikipedia sentences before extracting them into train/dev/test corpora. The corpus for each language consists of approximately 80 million tokens for training, as well as 10 million tokens each for development and testing. We generate language-specific vocabularies containing the 50,000 most common tokens in the training and development set; as is standard, out-of-vocabulary tokens in the training, development, and test sets are replaced with `<unk>`.

5.2 Training and Evaluation

We experiment with recurrent LMs and Transformer-based bidirectional encoders. LSTM LMs are trained for each language using the best hyperparameters in van Schijndel et al. (2019).⁶ We will refer to these models as *monolingual* LMs. We also train a *multilingual* LSTM LM over all of our languages. The training set for this model is a concatenation of all of the individual languages’ training corpora. The validation and test sets are concatenated in the same way, as are the vocabularies. We use the same hyperparameters as the monolingual models (Footnote 6). At each epoch, the corpora are randomly shuffled before batching; as such, each training batch consists with very high probability of sentences from multiple languages.

To obtain LSTM accuracies, we compute the total probability of each of the sentences in our challenge set, and then check within each minimal set whether the grammatical sentence has higher probability than the ungrammatical one. Because the syntactic performance of LSTM LMs has been found to vary across weight initializations (McCoy et al., 2018; Kuncoro et al., 2019), we report mean accuracy over five random initializations for each

⁴<https://github.com/attardi/wikiextractor>

⁵<https://www.cis.uni-muenchen.de/~schmid/tools/TreeTagger/>

⁶ Specifically, we use 2-layer word-level LSTMs with 800 hidden units in each layer, 800-dimensional word embeddings, initial learning rate 20.0 (annealed after any epoch in which validation perplexity did not improve relative to the previous epoch), batch size 20, and dropout probability 0.2.

LM. See Appendix C for standard deviations across runs on each test construction in each language.

We evaluate the syntactic abilities of multilingual BERT (mBERT, Devlin et al. 2019) using the approach of Goldberg (2019). Specifically, we mask out the focus verb, obtain predictions for the masked position, and then compare the scores assigned to the grammatical and ungrammatical forms in the minimal set. We use the scripts provided by Goldberg⁷ without modification, with the exception of using `bert-base-multilingual-cased` to obtain word probabilities. This approach is not equivalent to the method we use to evaluate LSTM LMs, as LSTM LMs score words based only on the left context, whereas BERT has access to left and right contexts. In some cases, mBERT’s vocabulary does not include the focus verbs that we vary in a particular minimal set. In such cases, if either or both verbs were missing, we skip that minimal set and calculate accuracies without the sentences contained therein.

6 Results

6.1 LSTMs

The overall syntactic performance of the monolingual LSTMs was fairly consistent across languages (Table 1 and Figure 2). Accuracy on short dependencies without attractors—Simple Agreement and Short VP Coordination—was close to perfect in all languages. This suggests that all monolingual models learned the basic facts of agreement, and were able to apply them to the vocabulary items in our materials. At the other end of the spectrum, performance was only slightly higher than chance in the Across an Object Relative Clause condition for all languages except German, suggesting that LSTMs tend to struggle with center embedding—that is, when a subject-verb dependency is nested within another dependency of the same kind (Marvin and Linzen, 2018; Noji and Takamura, 2020).

There was higher variability across languages in the remaining three constructions. The German models had almost perfect accuracy in Long VP Coordination and Across Prepositional Phrase, compared to accuracies ranging between 0.76 and 0.87 for other languages in those constructions. The Hebrew, Russian, and German models showed very high performance on the Across Subject Relative Clause condition: ≥ 0.88 compared to 0.6–0.71

⁷<https://github.com/yoavg/bert-syntax>

	English		French		German		Hebrew		Russian	
	Mono	Multi	Mono	Multi	Mono	Multi	Mono	Multi	Mono	Multi
Test Perplexity	57.90	66.13	35.48	57.40	46.31	61.06	48.78	61.85	35.09	54.61
Simple agreement	1.00	1.00	1.00	1.00	1.00	0.96	0.95	0.96	0.91	0.75
VP coordination (short)	0.94	0.96	0.97	0.85	0.99	1.00	1.00	0.95	0.98	0.92
VP coordination (long)	0.76	0.69	0.85	0.72	0.96	0.73	0.84	0.70	0.86	0.72
Across subject rel. clause	0.60	0.63	0.71	0.70	0.94	0.74	0.91	0.84	0.88	0.86
Within object rel. clause	0.89	0.79	0.99	0.99	0.74	0.69	1.00	0.88	0.95	0.88
Across object rel. clause	0.55	0.52	0.52	0.52	0.81	0.74	0.56	0.54	0.60	0.57
Across prepositional phrase	0.63	0.61	0.74	0.63	0.89	0.82	0.88	0.82	0.76	0.61
Average accuracy	0.77	0.74	0.83	0.78	0.90	0.81	0.88	0.81	0.85	0.76

Table 1: LSTM LM test perplexities and accuracies on CLAMS across languages for the language-specific monolingual models and for our multilingual model. Results are averaged across five random initializations. Chance accuracy is 0.5. Boldfaced numbers indicate the model that achieved the highest performance on a given construction *across languages*.

in other languages (recall that all our results are averaged over five runs, so this pattern is unlikely to be due to a single outlier).

With each of these trends, German seems to be a persistent outlier. This could be due to its marking of cases in separate article tokens—a unique feature among the languages evaluated here—or some facet of its word ordering or unique capitalization rules. In particular, subject relative clauses and object relative clauses have the same word order in German, but are differentiated by the case markings of the articles and relative pronouns. More investigation will be necessary to determine the sources of this deviation.

For most languages and constructions, the multilingual LM performed worse than the monolingual LMs, even though it was trained on five times as much data as each of the monolingual ones. Its average accuracy in each language was at least 3 percentage points lower than that of the corresponding monolingual LMs. Although all languages in our sample shared constructions such as prepositional phrases and relative clauses, there is no evidence that the multilingual LM acquired abstract representations that enable transfer across those languages; if anything, the languages interfered with each other. The absence of evidence for syntactic transfer across languages is consistent with the results of [Dhar and Bisazza \(2020\)](#), who likewise found no evidence of transfer in an LSTM LM trained on two closely related languages (French and Italian). One caveat is that the hyperparameters we chose for all of our LSTM LMs were based on a monolingual LM ([van Schijndel et al., 2019](#)); it is

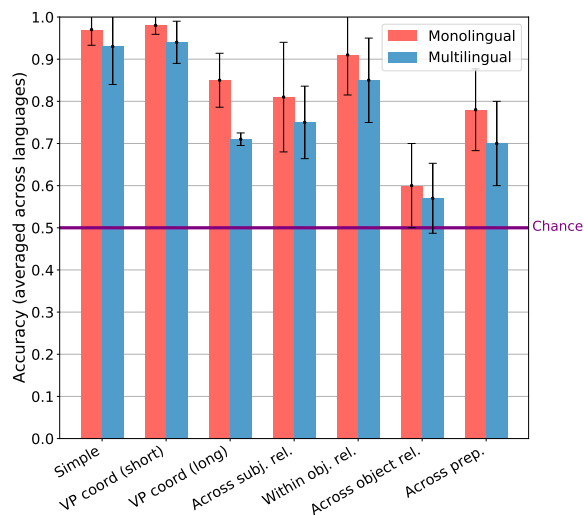


Figure 2: Mean accuracy (bars) and standard deviation (whiskers) for LSTM LMs over all languages for each stimulus type. **Note:** these are means over languages per-case, whereas the numbers in Table 1 are means over cases per-language.

possible that the multilingual LM would have been more successful if we had optimized its hyperparameters separately (e.g., it might benefit from a larger hidden layer).

These findings also suggest that test perplexity and subject-verb agreement accuracy in syntactically complex contexts are not strongly correlated cross-linguistically. This extends one of the results of [Kuncoro et al. \(2019\)](#), who found that test perplexity and syntactic accuracy were not necessarily strongly correlated within English. Finally, the multilingual LM’s perplexity was always higher than that of the monolingual LMs. At

	English	French	German	Hebrew	Russian
Simple agreement	1.00	1.00	0.95	0.70	0.65
VP coordination (short)	1.00	1.00	0.97	0.91	0.80
VP coordination (long)	0.92	0.98	1.00	0.73	—
Across subject relative clause	0.88	0.57	0.73	0.61	0.70
Within object relative clause	0.83	—	—	—	—
Across object relative clause	0.87	0.86	0.93	0.55	0.67
Across prepositional phrase	0.92	0.57	0.95	0.62	0.56

Table 2: Multilingual BERT accuracies on CLAMS. If a hyphen is present, this means that all focus verbs for that particular language and construction were out-of-vocabulary. Chance accuracy is 0.5.

first glance, this contradicts the results of Östling and Tiedemann (2017), who observed lower perplexity in LMs trained on a small number of very similar languages (e.g., Danish, Swedish, and Norwegian) than in LMs trained on just one of those languages. However, their perplexity rose precipitously when trained on more languages and/or less-related languages—as we have here.

6.2 BERT and mBERT

Table 2 shows mBERT’s accuracies on all stimuli. Performance on CLAMS was fairly high in the languages that are written in Latin script (English, French and German). On English in particular, accuracy was high across conditions, ranging between 0.83 and 0.88 for sentences with relative clauses, and between 0.92 and 1.00 for the remaining conditions. Accuracy in German was also high: above 0.90 on all constructions except Across Subject Relative Clause, where it was 0.73. French accuracy was more variable: high for most conditions, but low for Across Subject Relative Clause and Across Prepositional Phrase.

In all Latin-script languages, accuracy on Across an Object Relative Clause was much higher than in our LSTMs. However, the results are not directly comparable, for two reasons. First, as we have mentioned, we followed Goldberg (2019) in excluding the examples whose focus verbs were not present in mBERT’s vocabulary; this happened frequently (see Appendix D for statistics). Perhaps more importantly, unlike the LSTM LMs, mBERT has access to the right context of the focus word; in Across Object Relative Clause sentences (*the farmers that the lawyer likes smile/*smiles.*), the period at the end of the sentence may indicate to a bidirectional model that the preceding word (*smile/smiles*) is part of the main clause rather than the relative clause, and should therefore agree with *farmers* rather than *lawyer*.

In contrast to the languages written in Latin script, mBERT’s accuracy was noticeably lower on Hebrew and Russian—even on the Simple Agreement cases, which do not pose any syntactic challenge. Multilingual BERT’s surprisingly poor syntactic performance on these languages may arise from the fact that mBERT’s vocabulary (of size 110,000) is shared across all languages, and that a large proportion of the training data is likely in Latin script. While Devlin et al. (2019) reweighted the training sets for each language to obtain a more even distribution across various languages during training, it remains the case that most of the largest Wikipedias are written in languages which use Latin script, whereas Hebrew script is used only by Hebrew, and the Cyrillic script, while used by several languages, is not as well-represented in the largest Wikipedias.

We next compare the performance of monolingual and multilingual BERT. Since this experiment is not limited to using constructions that appear in all of our languages, we use additional constructions from Marvin and Linzen (2018), including reflexive anaphora and reduced relative clauses (i.e., relative clauses without *that*). We exclude their negative polarity item examples, as the two members of a minimal pair in this construction differ in more than one word position.

The results of this experiment are shown in Table 3. Multilingual BERT performed better than English BERT on Sentential Complements, Short VP Coordination, and Across a Prepositional Phrase, but worse on Within an Object Relative Clause, Across an Object Relative Clause (no relative pronoun), and in Reflexive Anaphora Across a Relative Clause. The omission of the relative pronoun *that* caused a sharp drop in performance in mBERT, and a milder drop in English BERT. Otherwise, both models had similar accuracies on other stimuli.

	Mono	Multi
SUBJECT-VERB AGREEMENT		
Simple	1.00	1.00
In a sentential complement	0.83	1.00
VP coordination (short)	0.89	1.00
VP coordination (long)	0.98	0.92
Across subject rel. clause	0.84	0.88
Within object rel. clause	0.95	0.83
Within object rel. clause (no <i>that</i>)	0.79	0.61
Across object rel. clause	0.89	0.87
Across object rel. clause (no <i>that</i>)	0.86	0.64
Across prepositional phrase	0.85	0.92
Average accuracy	0.89	0.87
REFLEXIVE ANAPHORA		
Simple	0.94	0.87
In a sentential complement	0.89	0.89
Across a relative clause	0.80	0.74
Average accuracy	0.88	0.83

Table 3: English BERT (base) and multilingual BERT accuracies on the English stimuli from Marvin and Linzen (2018). Monolingual results are taken from Goldberg (2019).

These results reinforce the finding in LSTMs that multilingual models generally underperform monolingual models of the same architecture, though there are specific contexts in which they can perform slightly better.

6.3 Morphological Complexity vs. Accuracy

Languages vary in the extent to which they indicate the syntactic role of a word using overt morphemes. In Russian, for example, the subject is generally marked with a suffix indicating nominative case, and the direct object with a different suffix indicating accusative case. Such case distinctions are rarely indicated in English, with the exception of pronouns (*he* vs. *him*). English also displays significant syncretism: morphological distinctions that are made in some contexts (e.g., *eat* for plural subjects vs. *eats* for singular subjects) are neutralized in others (*ate* for both singular and plural subjects). We predict that greater morphological complexity, which is likely to correlate with less syncretism, will provide more explicit cues to hierarchical syntactic structure,⁸ and thus result in increased overall accuracy on a given language.

To measure the morphological complexity of a

⁸For more evidence that explicit cues to structural information can aid syntactic performance, see Appendix B.

language, we use the C_{WALS} metric of Bentz et al. (2016): $\frac{\sum_{i=1}^n f_i}{n}$. This is a typological measure of complexity based on the World Atlas of Language Structures (WALS, Dryer and Haspelmath 2013), where f_i refers to a morphological feature value normalized to the range $[0, 1]$.⁹ This essentially amounts to a mean over normalized values of quantified morphological features. Here, n is 27 or 28 depending on the number of morphological categorizations present for a given language in WALS.

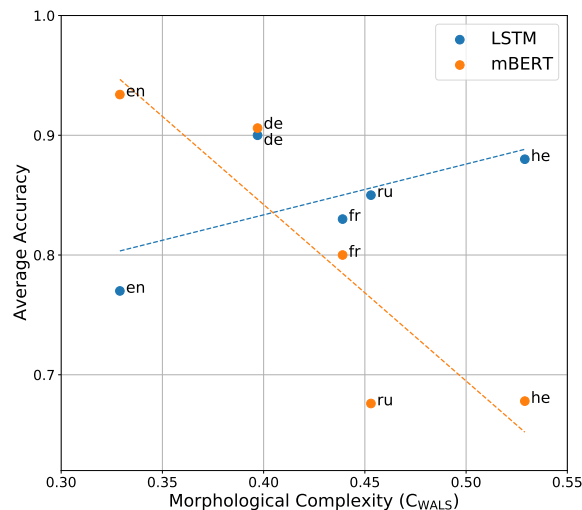


Figure 3: Morphological complexities against average accuracies per-language for LSTMs and mBERT.

Does the morphological complexity of a language correlate with the syntactic prediction accuracy of LMs trained on that language? In the LSTM LMs (Table 1), the answer is generally yes, but not consistently. We see higher average accuracies for French than English (French has more distinct person/number verb inflections), higher for Russian than French, and higher for Hebrew than Russian (Hebrew verbs are inflected for person, number, and gender). However, German is again an outlier: despite its notably lower complexity than Hebrew and Russian, it achieved a higher average accuracy. The same reasoning applied in Section 6.1 for German’s deviation from otherwise consistent trends applies to this analysis as well.

Nonetheless, the Spearman correlation between morphological complexity and average accuracy including German is 0.4; excluding German, it is 1.0. Because we have the same amount of training data per-language in the same domain, this could point to the importance of having explicit cues to lin-

⁹For example, if WALS states that a language has negative morphemes, f_{28} is 1; otherwise, f_{28} is 0.

guistic structure such that models can learn that structure. While more language varieties need to be evaluated to determine whether this trend is robust, we note that this finding is consistent with that of Ravfogel et al. (2019), who compared English to a synthetic variety of English augmented with case markers and found that the addition of case markers increased LSTM agreement prediction accuracy.

We see the opposite trend for mBERT (Table 2): if we take the average accuracy over all stimulus types for which we have scores for all languages—i.e., all stimulus types except Long VP Coordination and Within an Object Relative Clause—then we see a correlation of $\rho = -0.9$. In other words, accuracy is likely to *decrease* with increasing morphological complexity. This unexpected inverse correlation may be an artifact of mBERT’s limited vocabulary, especially in non-Latin scripts. Morphologically complex languages have more unique word types. In some languages, this issue can be mitigated to some extent by splitting the word into subword units, as BERT does; however, the effectiveness of such a strategy would be limited at best in a language with non-concatenative morphology such as Hebrew. Finally, we stress that the exclusion of certain stimulus types and the differing amount of training data per-language act as confounding variables, rendering a comparison between mBERT and LSTMs difficult.

7 Conclusions

In this work, we have introduced the CLAMS data set for cross-linguistic syntactic evaluation of word prediction models, and used it to evaluate monolingual and multilingual versions of LSTMs and BERT. The design conditions of Marvin and Linzen (2018) and our cross-linguistic replications rule out the possibility of memorizing the training data or relying on statistical correlations/token collocations. Thus, our findings indicate that LSTM language models can distinguish grammatical from ungrammatical subject-verb agreement dependencies with considerable overall accuracy across languages, but their accuracy declines on some constructions (in particular, center-embedded clauses). We also find that multilingual neural LMs in their current form do not show signs of transfer across languages, but rather harmful interference. This issue could be mitigated in the future with architectural changes to neural LMs (such as better handling of morphology), more principled combinations of languages

(as in Dhar and Bisazza 2020), or through explicit separation between languages during training (e.g., using explicit language IDs).

Our experiments on BERT and mBERT suggest (1) that mBERT shows signs of learning syntactic generalizations in multiple languages, (2) that it learns these generalizations better in some languages than others, and (3) that its sensitivity to syntax is lower than that of monolingual BERT. It is possible that its performance drop in Hebrew and Russian could be mitigated with fine-tuning on more data in these languages.

When evaluating the effect of the morphological complexity of a language on the LMs’ syntactic prediction accuracy, we found that recurrent neural LMs demonstrate better hierarchical syntactic knowledge in morphologically richer languages. Conversely, mBERT demonstrated moderately better syntactic knowledge in morphologically *simpler* languages. Since CLAMS currently includes only five languages, this correlation should be taken as very preliminary. In future work, we intend to expand the coverage of CLAMS by incorporating language-specific and non-binary phenomena (e.g., French subjunctive vs. indicative and different person/number combinations, respectively), and by expanding the typological diversity of our languages.

Acknowledgments

This material is based on work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1746891. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the other supporting agencies. Additionally, this work was supported by a Google Faculty Research Award to Tal Linzen, and by the United States–Israel Binational Science Foundation (award 2018284).

References

- Aixiu An, Peng Qian, Ethan Wilcox, and Roger Levy. 2019. Representation of constituents in neural language models: Coordination phrase as a case study. 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 2881–2892, Hong Kong, China. Association for Computational Linguistics.

- Christian Bentz, Tatyana Ruzsics, Alexander Koplenig, and Tanja Samardžić. 2016. [A comparison between morphological complexity measures: Typological data vs. language corpora](#). In *Proceedings of the Workshop on Computational Linguistics for Linguistic Complexity (CLALC)*, pages 142–153, Osaka, Japan. The COLING 2016 Organizing Committee.
- Noam Chomsky. 1956. Three models for the description of language. *IRE Transactions on Information Theory*, 2(3):113–124.
- Noam Chomsky. 1957. *Syntactic Structures*. Mouton, The Hague.
- Shammur Absar Chowdhury and Roberto Zamparelli. 2018. RNN simulations of grammaticality judgments on long-distance dependencies. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 133–144. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Prajit Dhar and Arianna Bisazza. 2018. [Does syntactic knowledge in multilingual language models transfer across languages?](#) In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 374–377, Brussels, Belgium. Association for Computational Linguistics.
- Prajit Dhar and Arianna Bisazza. 2020. [Understanding cross-lingual syntactic transfer in multilingual recurrent neural networks](#). arXiv preprint 2003.14056.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. 2019. [Neural language models as psycholinguistic subjects: Representations of syntactic state](#). 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 32–42, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anastasia Giannakidou. 2011. Negative and positive polarity items: Variation, licensing, and compositionality. In *Semantics: An international handbook of natural language meaning*, pages 1660–1712. Berlin: Mouton de Gruyter.
- Yoav Goldberg. 2019. [Assessing BERT’s syntactic abilities](#). arXiv preprint 1901.05287.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.
- Yiding Hao. 2020. [Attribution analysis of grammatical dependencies in lstms](#). arXiv preprint 2005.00062.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Pierre Isabelle, Colin Cherry, and George Foster. 2017. [A challenge set approach to evaluating machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2486–2496, Copenhagen, Denmark. Association for Computational Linguistics.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. [Exploring the limits of language modeling](#). arXiv preprint 1602.02410.
- Jaap Jumelet and Dieuwke Hupkes. 2018. [Do language models understand anything? On the ability of LSTMs to understand negative polarity items](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 222–231, Brussels, Belgium. Association for Computational Linguistics.
- Katharina Kann, Alex Warstadt, Adina Williams, and Samuel R. Bowman. 2019. [Verb argument structure alternations in word and sentence embeddings](#). In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 287–297.
- Adhiguna Kuncoro, Chris Dyer, John Hale, Dani Yogatama, Stephen Clark, and Phil Blunsom. 2018. [LSTMs can learn syntax-sensitive dependencies well, but modeling structure makes them better](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1436, Melbourne, Australia. Association for Computational Linguistics.
- Adhiguna Kuncoro, Chris Dyer, Laura Rimell, Stephen Clark, and Phil Blunsom. 2019. [Scalable syntax-aware language models using knowledge distillation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3472–3484, Florence, Italy. Association for Computational Linguistics.
- Jey Han Lau, Alexander Clark, and Shalom Lappin. 2017. [Grammaticality, acceptability, and probability: A probabilistic view of linguistic knowledge](#). *Cognitive Science*, (5):1202–1247.

- 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.
- Heidi Lorimor, Kathryn Bock, Ekaterina Zalkind, Alina Sheyman, and Robert Beard. 2008. [Agreement and attraction in Russian](#). *Language and Cognitive Processes*, 23(6):769–799.
- 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.
- R. Thomas McCoy, Robert Frank, and Tal Linzen. 2018. [Revisiting the poverty of the stimulus: Hierarchical generalization without a hierarchical bias in recurrent neural networks](#). In *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, pages 2093–2098, Austin, TX.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Proceedings of the 11th Annual Conference of the International Speech Communication Association (INTERSPEECH 2010)*, pages 1045–1048, Makuhari, Chiba, Japan.
- Hiroshi Noji and Hiroya Takamura. 2020. An analysis of the utility of explicit negative examples to improve the syntactic abilities of neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Seattle, Washington, USA. Association for Computational Linguistics.
- Robert Östling and Jörg Tiedemann. 2017. [Continuous multilinguality with language vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 644–649, Valencia, Spain. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv preprint 1910.10683*.
- 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, Yoav Goldberg, and Francis Tyers. 2018. [Can LSTM learn to capture agreement? the case of Basque](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 98–107, Brussels, Belgium. Association for Computational Linguistics.
- Marten van Schijndel, Aaron Mueller, and Tal Linzen. 2019. [Quantity doesn’t buy quality syntax with neural language models](#). 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 5835–5841, Hong Kong, China. Association for Computational Linguistics.
- Carson T Schütze. 1996. *The empirical base of linguistics: Grammaticality judgments and linguistic methodology*. University of Chicago Press.
- Rico Sennrich. 2017. [How grammatical is character-level neural machine translation? Assessing MT quality with contrastive translation pairs](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 376–382, Valencia, Spain. Association for Computational Linguistics.
- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. 2012. LSTM neural networks for language modeling. In *Proceedings of the 13th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 194–197.
- Ke Tran, Arianna Bisazza, and Christof Monz. 2018. [The importance of being recurrent for modeling hierarchical structure](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4731–4736, Brussels, Belgium. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2019a. [BLiMP: A benchmark of linguistic minimal pairs for English](#). *arXiv preprint 1912.00582*.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019b. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Paul Wexler. 1990. *The schizoid nature of modern Hebrew: A Slavic language in search of a Semitic past*. Wiesbaden: Otto Harrassowitz Verlag.

Ethan Wilcox, Roger Levy, Takashi Morita, and Richard Futrell. 2018. [What do RNN language models learn about filler-gap dependencies?](#) In *Proceedings of the 2018 EMNLP Workshop Black-boxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 211–221, Brussels, Belgium. Association for Computational Linguistics.

Amir Zeldes. 2013. Is Modern Hebrew Standard Average European? The view from European. *Linguistic Typology*, 17(3):439–470.

Ghil'ad Zuckermann. 2006. A new vision for Israel Hebrew: Theoretical and practical implications of analyzing Israel's main language as a semi-engineered Semito-European hybrid language. *Journal of Modern Jewish Studies*, 5(1):57–71.

A Linguistic Examples

This section provides examples of the syntactic structures included in the CLAMS dataset across languages. For Hebrew, we transliterate its original right-to-left script into the left-to-right Latin script; this makes labeling and glossing more consistent across languages. Hebrew was **not** transliterated in the training/development/test corpora or in the evaluation sets. In all examples, (a) is English, (b) is French, (c) is German, (d) is Hebrew, and (e) is Russian.

The first case is simple agreement. This simply involves agreeing a verb with its adjacent subject, which should pose little challenge for any good language model regardless of syntactic knowledge.

(3) Simple Agreement:

- a. The surgeons laugh/*laughs.
- b. Le pilote parle / *parlent.
The pilot laughs / *laugh.
- c. Der Schriftsteller spricht / *sprechen.
The writer speaks / *speak.
- d. Ha meltsar yashen / yeshenim.
The server sleeps / *sleep.
- e. Врачи говорят / *говорит.
Doctors speak / *speaks.

Short verb-phrase coordination introduces some slight distance between the subject and verb, though the presence of the previous verb should give a model a clue as to which inflection should be more probable.

(4) VP coordination (short):

- a. The author swims and smiles/*smile.
- b. Les directeurs parlent et déménagent /
The directors talk and move /
*déménage.
*moves.

- c. Der Polizist schwimmt und lacht /
The police.officer swims and laughs /
*lachen.
*laugh.
- d. Ha tabaxim rokdim ve soxim / *soxe.
The cooks dance and swim / *swims.
- e. Профессор старый и читает / *читают.
Professor is.old and reads / *read.

Long verb-phrase coordination is similar, but makes each verb phrase much longer to introduce more distance and attractors between the subject and target verb.

(5) VP coordination (long):

- a. The teacher knows many different foreign languages and likes/*like to watch television shows.
- b. L' agriculteur écrit dans un journal tous les jours et préfère / *préfèrent jouer au tennis avec des collègues.
The farmer writes in a journal all the days and prefers / *prefer to.play at.the tennis with some colleagues.
- c. Die Bauern sprechen viele verschiedene Sprachen und sehen / *sieht gern Fernsehprogramme.
The farmers speak many various languages and watch / *watches gladly TV.shows.
- d. Ha tabax ohev litspot be toxniot televizya ve gar / *garim be merkaz ha ir.
The cook likes to.watch in shows TV and lives / *live in center the city.
- e. Автор знает много иностранных языков и любит / *любят смотреть телепередачи.
Author knows many foreign languages and likes / *like to.watch TV.shows.

Now we have more complex structures that require some form of structural knowledge if a model is to obtain the correct predictions with more than random-chance accuracy. Agreement across a subject relative clause involves a subject with an attached relative clause containing a verb and object, followed by the main verb. Here, the attractor is the object in the relative clause. (An attractor is an intervening noun between a noun and its associated finite verb which might influence a human's or model's decision as to which inflection to choose. This might be of the same person and number, or, in more difficult cases, a different person and/or number. It does not necessarily need to occur between the noun and its associated verb, though this

does tend to render this task more difficult.)

(6) *Across a subject relative clause:*

- a. The officers that love the chef are/*is old.
- b. Les chirurgiens qui détestent le garde
The surgeons that hate the guard
retournent / *retourne.
return / *returns
- c. Der Kunde, der die Architekten hasst, ist /
The customer that the architects hates is /
*sind klein.
*are short.
- d. Ha menahel she ma'arits et ha shomer
The manager who admires ACC the guard
rats / *ratsim.
runs / *run.
- e. Пилоты, которые понимают агентов,
Pilots that understand agents
говорят / *говорит.
speak / *speaks.

Agreement within an object relative clause requires the model to inflect the proper verb inside of an object relative clause; the object relative clause contains a noun and an associated transitive verb whose object requirement is filled by the relative pronoun. The model must choose the proper verb inflection given the noun within the relative clause as opposed to the noun outside of it. This may seem similar to simple agreement, but we now have an attractor which appears before the noun of the target verb.

(7) *Within an object relative clause:*

- a. The senator that the executives love/*loves laughs.
- b. Les professeurs que le chef admire /
The professors that the boss admires /
*admirent parlent.
*admire talk.
- c. Die Polizisten, die der Bruder hasst, /
The police.officers that the brother hates /
*hassen, sind alt
*hate are old.
- d. Ha menahel she ha nahag ma'aritz /
The manager that the driver admires /
*ma'aritsim soxe.
*admire swims.
- e. Сенаторы, которых рабочие ищут, /
Senators that workers seek /
*ищет, ждали.
*seeks wait.

Agreement across an object relative clause is similar, but now the model must choose the correct inflection for the noun outside of the relative clause.

This requires the model to capture long-range dependencies, and requires it to have the proper structural understanding to ignore the relative clause when choosing the proper inflection for the focus verb.

(8) *Across an object relative clause:*

- a. The senator that the executives love laughs/*laugh.
- b. Les professeurs que le chef admire parlent /
The professors that the boss admires talk /
*parle.
*talks.
- c. Der Senator, den die Tänzer mögen, spricht /
The senator that the dancers like speaks /
*sprechen.
*speak.
- d. Ha katsin she ha zamar ohev soxe /
The officer that the singer likes swims /
*soxim.
*swim.
- e. Фермеры, которых танцоры хотят,
Farmers that dancers want
большие / *большой.
are.big / *is.big.

Finally, agreement across a prepositional phrase entails placing a prepositional phrase after the subject; the prepositional phrase contains an attractor, which makes choosing the correct inflection more difficult.

(9) *Across a prepositional phrase:*

- a. The consultants behind the executive smile/*smiles.
- b. Les clients devant l' adjoint sont / *est
The clients in.front.of the deputy are / *is
vieux.
old.
- c. Der Lehrer neben den Ministern lacht /
The teacher next.to the ministers laughs /
*lachen.
*laugh.
- d. Ha meltsarim leyad ha zamarim nos'im /
The servers near the singers drive /
*nose'a.
*drives.
- e. Режиссёры перед агентами
Directors in.front.of agents
маленькие / *маленький.
are.small / *is.small.

Some of the constructions used by [Marvin and Linzen \(2018\)](#) could not be replicated across languages. This includes reflexive anaphora, where none of our non-English languages use quite the

	English		French		German		Russian	
	Mono	Multi	Mono	Multi	Mono	Multi	Mono	Multi
Simple agreement	—	-0.02	—	-0.01	—	+0.02	+0.02	—
VP coordination (short)	-0.01	—	+0.01	+0.14	-0.02	-0.01	-0.03	-0.01
VP coordination (long)	-0.03	+0.01	+0.04	-0.02	-0.06	+0.07	+0.04	+0.02
Across subject rel. clause	+0.24	+0.07	+0.23	+0.15	-0.03	+0.13	+0.02	+0.01
Within object rel. clause	—	-0.04	—	-0.07	—	-0.02	-	-0.03
Across object rel. clause	+0.09	+0.02	+0.05	+0.03	+0.01	+0.09	+0.01	-
Across prepositional phrase	+0.18	+0.11	+0.20	+0.20	+0.03	+0.03	+0.03	+0.02
Average accuracy	+0.06	+0.03	+0.07	+0.05	-0.01	+0.05	+0.01	+0.03

Table 4: Gains (positive, blue) and losses (negative, red) in LSTM LM accuracies on CLAMS after capitalizing the first character of each evaluation example. Differences are relative to the results in Table 1. Results are averaged across five random initializations.

same syntactic structures as English (or even to each other) when employing reflexive verbs and pronouns. Some do not even have separate reflexive pronouns for third-person singular and plural distinctions (like French and German). Moreover, the English reflexive examples rely on the syncretism between past-tense verbs for any English person and number,¹⁰ whereas other languages often have different surface forms for different person and number combinations in the past tense. This would give the model a large clue as to which reflexive is correct. Thus, any results on reflexive anaphora would not be comparable cross-linguistically. See example (10) below for English, French, and German examples of the differences in reflexive syntax.

(10) *Reflexive anaphora across relative clause:*

- a. The author that the guards like injured himself/*themselves.
- b. L' auteur que les gardes aiment s'
The author that the guards like REFL.3
est blessé / *se sont
has.3S injured.S.MASC / REFL.3 have.3P
blessés.
injured.P.MASC
- c. Der Autor, den die Wächter mögen,
The author that the guards like
verletzte sich / *verletzten sich.
injured.3S REFL.3 / injured.3P REFL.3

B The Importance of Capitalization

As discovered in Hao (2020), capitalizing the first character of each test example improves the per-

¹⁰For example, regardless of whether the subject is singular, plural, first- or third-person, etc., the past-tense of *see* is always *saw*.

formance of language models in distinguishing grammatical from ungrammatical sentences in English. To test whether this finding holds cross-linguistically, we capitalize the first character of each of our test examples in all applicable languages. Hebrew has no capital-/lower-case distinction, so it is excluded from this analysis.

Table 4 contains the results and relative gains or losses of our LSTM language models on the capitalized stimuli compared to the lowercase ones. For all languages except German, we see a notable increase in the syntactic ability of our models. For German, we see a small drop in overall performance, but its performance was already exceptionally high in the lowercase examples (perhaps due to its mandatory capitalization of all nouns).

An interesting change is that morphological complexity no longer correlates with the overall syntactic performance across languages ($\rho = 0.2$). Perhaps the capitalization acts as an explicit cue to syntactic structure by delineating the beginning of a sentence, thus supplanting the role of morphological cues in aiding the model to distinguish grammatical sentences.

Overall, it seems quite beneficial to capitalize one's test sentences before feeding them to a language model if one wishes to improve syntactic accuracy. The explanation given by Hao (2020) is that *The* essentially only appears sentence-initially, thus giving the model clues as to which noun (typically the token following *The*) is the subject. Conversely, *the* has a more varied distribution, as it may appear before essentially any noun in subject or object position; thus, it gives the model fewer

	English		French		German		Hebrew		Russian	
	Mono	Multi	Mono	Multi	Mono	Multi	Mono	Multi	Mono	Multi
Simple agreement	.00	.00	.00	.00	.00	.02	.01	.01	.01	.07
VP coordination (short)	.01	.00	.01	.05	.02	.00	.01	.01	.02	.02
VP coordination (long)	.06	.08	.05	.09	.04	.07	.06	.06	.04	.06
Across subject rel. clause	.06	.02	.05	.05	.04	.07	.03	.03	.03	.04
Within object rel. clause	.01	.02	.01	.01	.03	.04	.01	.03	.04	.02
Across object rel. clause	.05	.02	.01	.01	.09	.06	.01	.01	.03	.02
Across prepositional phrase	.02	.02	.02	.02	.06	.03	.03	.04	.02	.01

Table 5: Standard deviation of LSTM LM performance across five random weight initializations for all languages and stimulus types.

cues as to which noun agrees with a given verb. This would explain the larger score increase for English and French (which employ articles in a similar fashion in CLAMS), as well as the milder increase for Russian (which does not have articles). However, it does not explain the decrease in performance on German. A deeper investigation of this trend per-language could reveal interesting trends about the heuristics employed by language models when scoring syntactically complex sentences.

C Performance Variance

Previous work has found the variance of LSTM performance in syntactic agreement to be quite high (McCoy et al., 2018; Kuncoro et al., 2019). In Table 5, we provide the standard deviation of accuracy over five random initializations on all CLAMS languages and stimulus types. This value never exceeds 0.1, and tends to only exceed 0.05 in more difficult syntactic contexts.

For syntactic contexts without attractors, the standard deviation is generally low. In more difficult cases like Across a Subject Relative Clause and Long VP Coordination, we see far higher variance. In Across an Object Relative Clause, however, the standard deviation is quite low despite this being the case on which language models struggled most; this is likely due to the consistently at-chance performance on this case, further showcasing the difficulty of learning syntactic agreements in such contexts.

On cases where German tended to deviate from the general trends seen in other languages, we see our highest standard deviations. Notably, the performance of German LMs in Across an Object Relative Clause and Across a Prepositional Phrase varies far more than other languages for the same stimulus type.

D Evaluation Set Sizes

Here, we describe the size of the various evaluation set replications. These will differ for the LSTMs, BERT, and mBERT, as the two latter models sometimes do not contain the varied focus verb for a particular minimal set.

Table 6 displays the number of minimal sets per language and stimulus type (with animate nouns only) in our evaluation sets; the total number of sentences (grammatical *and* ungrammatical) is the number of minimal sets times two. These are also the number of examples that the LSTM is evaluated on. We do not include inanimate-noun cases in our evaluations for now, since these are much more difficult to replicate cross-linguistically. Indeed, grammatical gender is a confounding variable which—according to preliminary experiments—does have an effect on model performance. Additionally, Hebrew has differing inflections depending on the combination of the subject and object noun genders, which means that we rarely have all needed inflections in the vocabulary.

We have differing numbers of examples per-language for similar cases. The reasoning for this is two-fold: (1) direct translations do not exist for all English items in the evaluation set of Marvin and Linzen (2018), so we often must decide between multiple possibilities. In cases where there are two translations of a noun that could reasonably fit, we use both; if we have multiple possibilities for a given verb, we use only one—the most frequent of the possible translations. If no such translation exists for a given noun or verb, we pick a different word that is as close to the English token is possible in the same domain.

Reason (2) is that many of the nouns and verbs in the direct translation of the evaluation sets do not appear in the language models’ vocabularies. Thus,

	English	French	German	Hebrew	Russian
Simple agreement	140	280	140	140	280
VP coordination (short)	840	980	980	980	980
VP coordination (long)	400	500	500	500	500
Across subject rel. clause	11200	11200	11200	11200	10080
Within object rel. clause	11200	11200	11200	11200	11200
Across object rel. clause	11200	11200	11200	11200	11200
Across prepositional phrase	16800	14000	12600	5600	5880

Table 6: Number of minimal sets for all languages and stimulus types using animate nouns.

	English		French	German	Hebrew	Russian
	Mono	Multi				
SUBJECT-VERB AGREEMENT						
Simple agreement	120	80	40	100	20	80
In a sentential complement	1440	960	-	-	-	-
VP coordination (short)	720	480	140	700	140	280
VP coordination (long)	400	240	100	300	100	0
Across subject rel. clause	9600	6400	1600	5406	1600	2880
Within object rel. clause	15960	5320	0	0	0	0
Within object rel. clause (no <i>that</i>)	15960	5320	-	-	-	-
Across object rel. clause	19680	16480	1600	5620	1600	3200
Across object rel. clause (no <i>that</i>)	19680	16480	-	-	-	-
Across prepositional phrase	19440	14640	2000	9000	800	1680
REFLEXIVE ANAPHORA						
Simple	280	280	-	-	-	-
In a sentential complement	3360	3360	-	-	-	-
Across a rel. clause	22400	22400	-	-	-	-

Table 7: Number of minimal sets used by BERT (English monolingual only) and mBERT for evaluation. The number of monolingual English examples is the same as in Goldberg (2019). Hyphens indicate non-replicable stimulus types, and 0 indicates that all focus verbs for a given stimulus type were out-of-vocabulary.

some nouns or focus verbs would effectively be `<unk>`s if left in, rendering that particular example unusable. In such cases, if a given noun/verb is not the vocabulary, we pick a similar noun from the same domain if one exists; if a similar item does not exist in the vocabulary, we choose some common noun in that language’s vocabulary that has not already been used in the evaluation set.

We use a similar process to add new verbs, but sometimes, third-person singular and plural inflections of similar verbs did not exist in the vocabulary. In such cases, we used a similar verb if possible (e.g., ‘dislike’ would be reasonably similar in distribution and meaning to ‘hate’), but if no such similar verb exists in the vocabulary, we do not replace it. A similar process is used for closed classes like prepositions: if no sufficient replacement exists in

the vocabulary, it is not replaced.

Table 7 contains the number of examples used by BERT and mBERT to calculate examples. Important to note is that for these evaluations, we use stimulus types containing both animate *and* inanimate nouns to better match Goldberg (2019)’s experimental setup; this is why we have more examples for English in this table than for the LSTM evaluations. Including or excluding inanimate nouns was found to make no significant difference in the final scores (for BERT or mBERT) regardless, since the performance of the model never diverges by more than 0.02 for animate vs. inanimate stimulus types.

The variation in the number of examples across languages is due to many of the focus verbs not being in mBERT’s vocabulary. We see the lowest

coverage in general for Hebrew and (surprisingly) French; this is likely due to Hebrew script being a rarer script in mBERT and due to many of French's most common tokens being split into subwords, respectively. Russian also has relatively low coverage, having 0 in-vocabulary target verbs for long VP coordination. None of our languages except English had any target verbs for Within an Object Relative Clause.