

Controlled Evaluation of Grammatical Knowledge in Mandarin Chinese Language Models

Yiwen Wang¹ Jennifer Hu² Roger Levy² Peng Qian²

¹ John A. Paulson School of Engineering and Applied Sciences, Harvard University

² Department of Brain and Cognitive Sciences, MIT

yiwenwang@g.harvard.edu {jennhu, rplevy, pqian}@mit.edu

Abstract

Prior work has shown that structural supervision helps English language models learn generalizations about syntactic phenomena such as subject-verb agreement. However, it remains unclear if such an inductive bias would also improve language models' ability to learn grammatical dependencies in typologically different languages. Here we investigate this question in Mandarin Chinese, which has a logographic, largely syllable-based writing system; different word order; and sparser morphology than English. We train LSTMs, Recurrent Neural Network Grammars, Transformer language models, and Transformer-parameterized generative parsing models on two Mandarin Chinese datasets of different sizes. We evaluate the models' ability to learn different aspects of Mandarin grammar that assess syntactic and semantic relationships. We find suggestive evidence that structural supervision helps with representing syntactic state across intervening content and improves performance in low-data settings, suggesting that the benefits of hierarchical inductive biases in acquiring dependency relationships may extend beyond English.

1 Introduction

A rich collection of targeted linguistic evaluations has shown that neural language models can surprisingly learn many aspects of grammar from unlabeled linguistic input (e.g., Linzen et al., 2016; Guordava et al., 2018; Warstadt et al., 2020; Hu et al., 2020; Xiang et al., 2021). There is also growing evidence that explicit modeling of syntax helps neural network-based language models represent syntactic state and exhibit human-like processing behaviors of non-local grammatical dependencies, including number agreement (Kuncoro et al., 2018), negative polarity licensing, filler-gap dependencies (Wilcox et al., 2019; Hu et al., 2020), and garden-path effects (Futrell et al., 2019; Hu et al., 2020). However,

this line of research has focused primarily on the syntax of English. It is unclear to what extent structural supervision may help neural language models generalize for languages with differing typologies. Expanding these analyses beyond English has the potential to inform scientific questions about inductive biases for language acquisition, as well as practical questions about model architectures that approach language-independence (Bender, 2011).

Here, we perform a controlled case study of grammatical knowledge in Mandarin Chinese language models. The orthography and grammar of Chinese provide a useful testing ground given the differences from English and other Indo-European languages (see, e.g., Shopen 1985; Li et al. 2015). Whereas today's Indo-European languages like English generally use phone-based orthography, Chinese uses a logographic system where each character generally responds to a syllable. Most Mandarin Chinese words are one or two syllables, influencing the distribution of tokens. Grammatically, Chinese has almost no inflectional morphology, and corpus studies suggest that the average dependency length of Mandarin Chinese sentences is larger than that of English sentences (Jiang and Liu, 2015), with potential implications for language modeling. On the one hand, the need to track input across long dependencies may make structural supervision more beneficial for Mandarin Chinese language models; on the other hand, the prevalence of these dependencies may make it easier for them to learn to maintain non-local information without explicitly modeling syntax. Other fine-grained differences in typology also affect the types of syntactic tests that can be conducted. For example, since relative clauses precede the head noun in Chinese (unlike in English), we can manipulate the distance of a verb-object dependency by inserting relative clauses in between. These characteristics motivate our choice of Mandarin Chinese as a language for evaluating structurally supervised neural language models.

We design six classes of Mandarin test suites covering a range of syntactic and semantic relationships, some specific to Mandarin and some comparable to English. We train neural language models with differing inductive biases on two datasets of different sizes, and compare models’ performance on our targeted evaluation materials. While most prior work investigating syntactically guided language models has used Recurrent Neural Network Grammar models (Dyer et al., 2016) – potentially conflating structural supervision with a particular parameterization – this work further explores structured Transformer language models (Qian et al., 2021). Our results are summarized as follows. We find that structural supervision yields greatest performance advantages in low-data settings, in line with prior work on English language models. Our results also suggest a potential benefit of structural supervision in deriving garden-path effects induced by local classifier–noun mismatch, and in maintaining syntactic expectations across intervening content within a dependency relation. These findings suggest that the benefits of hierarchical inductive biases in acquiring dependency relationships may not be specific to English.

2 Targeted Linguistic Evaluation

Linguistic minimal pairs have been used to construct syntactic test suites in English (e.g., Linzen et al., 2016; Marvin and Linzen, 2018; Mueller et al., 2020; Davis and van Schijndel, 2020; Warstadt et al., 2020) and other languages such as Italian, Spanish, French, and Russian (Ravfogel et al., 2018; Gulordava et al., 2018; An et al., 2019; Mueller et al., 2020; Davis and van Schijndel, 2020). A minimal pair is formed by two sentences that differ in grammaticality or acceptability but are otherwise matched in structure and lexical content. (A) and (B) form an example English minimal pair differing in subject-verb agreement:

- (A) The man *drinks* coffee everyday.
- (B) The man *drink* coffee everyday.

The two sentences differ only at the main verb ‘drinks’/‘drink’, which must agree with its subject ‘The man’ in number. Since only the third-person singular form ‘drinks’ agrees with the subject, (A) is grammatical, whereas (B) is not.

The closest work to ours is the corpus of Chinese linguistic minimal pairs (CLiMP; Xiang et al., 2021), which provides a benchmark for testing

syntactic generalization of Mandarin Chinese language models. While CLiMP focuses on building a comprehensive challenge set, the current work performs controlled experiments to investigate the effect of structural supervision on language models’ ability to learn syntactic and semantic relationships. Moreover, the test items in CLiMP are (semi)-automatically generated, which may result in semantically anomalous sentences and introduce noise into the evaluation phase. In contrast, we manually construct the items in our test suites to sound as natural as possible.

3 Methods

3.1 Evaluation Paradigm

The general structure of our evaluation paradigm follows that of Wilcox et al. (2019) and Hu et al. (2020). We use surprisal as a linking function between the language model output and human expectations (Hale, 2001). Surprisal is defined as the inverse log probability of a word (w_i) conditioned on the preceding words in the same context ($w_1 \dots w_{i-1}$):

$$\text{surprisal}(w_i) = \log \frac{1}{p(w_i | w_1, \dots, w_{i-1})}$$

Our test suites take the form of a group of hand-written, controlled sentence sets. Each sentence set, or *test item*, contains at least two minimally differing sentences, and each sentence contains a *stimulus prefix* and a downstream *target region*. The content of the target region remains fixed across the sentence variants within the test item, while the content of the stimulus varies in a minimal manner that modulates the sentence’s grammaticality or acceptability. The target region, where we measure the surprisal output of a certain model, is underlined in all the example items described in Section 3.2.

For each test item, we measure success by computing the difference in surprisals assigned by the model to the target region, conditioned on the ungrammatical vs. grammatical stimulus prefixes. If the model successfully captures the dependency, it should be less surprised at the grammatical target region than the ungrammatical one, leading to a positive surprisal difference (ungrammatical – grammatical). If this criterion is satisfied, then the model achieves a success score of 1, and 0 otherwise. These binary scores are averaged over test suite items and/or classes to obtain accuracy scores.

3.2 Test Suites

We organize our materials into six classes of test suites, each of which assesses models' knowledge of a particular linguistic phenomenon.¹ Within each class, we develop individual test suites with different types of modifiers that intervene between the two ends of the dependency (including the no modifier case). For each test suite, we manually construct ~30 test items, taking care to maintain semantic plausibility wherever possible. For details about individual test suites, see Appendix A.1. While here we present example items of a specific type of modifier to introduce each test suite, full set of examples with other type of modifiers are included in Appendix A.2.

The test suite classes include both syntactic and semantic dependencies, some of which do not exist in languages that have been the focus of targeted evaluation, and thus have not yet been explored. MISSING OBJECT evaluates syntactic knowledge of argument structure, SUBORDINATION and GARDEN PATH SUBJECT/OBJECT assess representation of syntactic state, and CLASSIFIER-NOUN COMPATIBILITY and VERB-NOUN COMPATIBILITY evaluate a combination of syntactic and semantic factors (Balari, 1992). Looking cross-linguistically, one class assesses a phenomenon present in Mandarin but not English (CLASSIFIER-NOUN COMPATIBILITY); two classes assess an expectation-violation phenomenon that is present in both Mandarin and English but arises from different sources (GARDEN PATH SUBJECT/OBJECT); and three classes assess phenomena present in both languages (VERB-NOUN COMPATIBILITY, MISSING OBJECT, and SUBORDINATION).

For each phenomenon targeted by a given test suite class, the two components of the syntactic/semantic dependency often occur adjacently in a sentence. However, if a language model robustly represents the dependency, then it should maintain its expectations even when intervening content is present between the upstream and downstream ends of the dependency. We assess the robustness of the models' grammatical knowledge on each test suite class by inserting three commonly-used types of modifiers to create non-local dependencies: adjectives, subject-extracted relative clauses (SRCs) and their variants, and object-extracted relative clauses

(ORCs). The resulting set of test suites is described in greater detail in the following sections.

3.2.1 Classifier-Noun Compatibility

Classifiers are a special class of words in Chinese languages which are obligatorily used with numerals in a noun phrase. Each specific classifier in Mandarin is only compatible with a set of noun references that is largely semantically delimited. The general classifier “个”(CL_{GENERAL}), in contrast, is compatible with most nouns. In the CLASSIFIER-NOUN COMPATIBILITY test suites, we evaluate whether a model expects nouns from a semantically-compatible class over those from a semantically-incompatible class, given a specific classifier.

- (1.a) 孩子听到了 一 首 熟悉 的 歌曲 。
child hear PST one CL_{SONG} familiar DE song .
“The child heard a familiar song.”
- *(1.b) 孩子听到了 一 张 熟悉 的 歌曲 。
child hear PST one CL_{ALBUM} familiar DE song .
“The child heard a familiar song.”
- (1.c) 孩子听到了 一 张 熟悉 的 专辑 。
child hear PST one CL_{ALBUM} familiar DE album .
“The child heard a familiar album.”
- *(1.d) 孩子听到了 一 首 熟悉 的 专辑 。
child hear PST one CL_{SONG} familiar DE album .
“The child heard a familiar album.”

Example (1) shows a test item from the suite with adjectival modifiers. We also consider ORCs and SRCs as modifiers in this test suite class (see Appendix A.2.1). Here we consider the classifier “首”(CL_{SONG}), which is compatible with the noun “歌曲”(song) but not the noun “专辑”(album), and the classifier “张”(CL_{ALBUM}), which is compatible with the noun “专辑” but not the noun “歌曲”. The four variants (1.a-d) show four possible combinations of the two classifiers and the two nouns. Here the target region is the sentence-final noun together with the period. We also check that the two nouns compared within each test item have similar frequency in the training data. We measure surprisals at the sentence-final noun and the period. A human-like language model should assign lower surprisals to the target regions in (1.a) and (1.c), the items with an appropriate classifier-noun pair, and high surprisals to the target regions in (1.b) and (1.d), the items with mismatched classifier-noun pairs. In other words, we evaluate four pair-wise comparisons to see whether they meet the following criteria: (1.b) > (1.a), (1.d) > (1.c), (1.d) > (1.a),

¹All code and data, including test suites, can be found at <https://github.com/YiwenWang03/syntactic-generalization-mandarin>

and (1.b) > (1.c). We report mean accuracy averaged across all four pair-wise comparisons as a model’s accuracy on a given test suite.

3.2.2 Garden-Path Effects

Garden-path effects are a class of phenomena in human sentence processing, where the incremental parsing state of a sentence prefix needs to be reanalyzed as the comprehender processes a downstream disambiguator region (Bever, 1970). We construct a set of test suites evaluating whether models exhibit garden-path effects induced by locally mismatched classifier–noun pairs situated within a globally coherent sentence, inspired by previous human behavioral studies (Wu et al., 2018).

To illustrate the classifier-induced garden-path effect, consider examples (2.a) and (2.b):

- (2.a) 他离开了那间朋友开的工厂。
 he leave PST that CL_{BUILDING} friend start DE factory .
- (2.b) 他离开了那个朋友开的工厂。
 he leave PST that CL_{GENERAL} friend start DE factory .
 “He left the factory that the friend started.”

In (2.b), the general classifier “个”(CL_{GENERAL}) is compatible with the immediately following noun, “朋友”(friend), resulting in a garden-path interpretation that this noun is the object of the main-clause verb “离开”(leave). This interpretation is disconfirmed, however, by the next verb, “开”(start), which indicates that the noun “朋友” is actually the start of a relative clause preceding the main-clause object, “工厂”(factory) . This means that the relative clause verb “开” should be highly surprising in (2.b). In (2.a), in contrast, the specific classifier “间”(CL_{BUILDING}) is incompatible with “朋友” (though it is compatible with “工厂”), cueing the upcoming relative clause structure (Wu et al., 2018). A human-like language model should thus show a higher surprisal at the target verb “开” for (2.b) than (2.a).

We design test suites similar to the structure of example (2) and (3). Examples (2.a-b) show the structure of items in the GARDEN PATH OBJECT set, where an ORC modifies the *object* of the main clause verb and the target region is the verb immediately following the closest noun to the classifier. Examples (3.a-b) show the basic structure of items in the GARDEN PATH SUBJECT set, where an ORC modifies the sentence *subject*.

- (3.a) 那间朋友开的工厂倒闭了。
 that CL_{BUILDING} friend start DE factory close PST .
- (3.b) 那个朋友开的工厂倒闭了。
 that CL_{GENERAL} friend start DE factory closed PST .

“The factory that the friend started has closed.”

Here in (3.b), the garden-path interpretation is that the ORC’s subject “朋友” may be initially analyzed as the subject of the sentence during incremental processing; the target region of the garden-path effect is the disambiguating word “的”(DE) that ends the relative clause and precedes the head noun of the true main subject of the sentence, “工厂”.

The criterion for getting a test item correct is that the model shows lower surprisal at the target region “开” (start) in (2.a) than (2.b), and lower surprisal at the target region “的” in (3.a) than (3.b). Examples (2) and (3) have no modifiers in between the classifier stimulus and the target region. To manipulate the length of dependency, we also consider same types of modifiers as in CLASSIFIER–NOUN COMPATIBILITY in the full test suites.

3.2.3 Verb–Noun Compatibility

Similar to CLASSIFIER–NOUN COMPATIBILITY, VERB–NOUN COMPATIBILITY is also a group of semantic test suites, assessing the consistency between a transitive verb² and its direct object noun. (4) shows an example with an adjective modifier in between the verb and its object, where (4.b) is semantically inconsistent since the word “阅读”(read) does not match the object noun “电脑”(computer).

- (4.a) 我修理了这个新的电脑。
 I fix PST this CL new DE computer .
 “I have fixed this new computer.”
- *(4.b) 我阅读了这个新的电脑。
 I read PST this CL new DE computer .
 “I have read this new computer.”

The stimulus is the transitive verb, and the target region is the object noun and the period (which encapsulates the possibility of an incomplete but potentially grammatical sentence). We insert adjectives, ORCs and SRCs modifiers (same as in CLASSIFIER–NOUN COMPATIBILITY) between the verb and object. The expected behavior is that the surprisal at the target region being lower in the semantically consistent variant (here, (4.a)).

3.2.4 Missing Object

Next, we turn to phenomena primarily characterized by syntactic expectations. The first of these test suite classes, MISSING OBJECT, assesses models’ ability to track a direct object required by a

²We test the transitivity of the verbs with a Tregex search (Levy and Andrew, 2006) on the CTB dataset, along with human judgments from native Mandarin speakers.

transitive main verb. Consider (5.a) and (5.b) (no modifier case):

- (5.a) 记者 采访 了 科学家 。
 journalist interview PST scientist
 “The journalist interviewed the scientist.”
- *(5.b) 记者 采访 了 。
 journalist interview PST
 “The journalist interviewed.”

(5.a) is grammatical and (5.b) is not, since the main verb “采访”(interview) requires a downstream direct object. To test if models learn this dependency, we record the model’s surprisal at the sentence-final period “。”。 The model should be more surprised to see “。” in (5.b) since it is less likely to end a sentence without an object needed by the verb. Note that this is a case where we assess human-likeness of an autoregressive language model by whether it is *temporarily confused* as to the structural interpretation midway through the sentence, as evidenced by its next-word predictions.

To continue our investigation of long-distance dependencies, we add three types of modifiers: single SRC, coordinated SRCs, and embedded SRCs, exemplified in Appendix A.2.5, respectively.³ We expect the insertion of modifiers after the main verb to increase difficulty, as the model must track the verb-object dependency over a greater amount of content. In addition, the parallel and hierarchical SRCs are longer and more syntactically complex than the simple SRCs.

3.2.5 Subordination

Finally, our SUBORDINATION test suites assess the ability of a model to maintain global expectations for a main clause while inside a local subordinate clause. For example, consider (6.a) and (6.b) (no modifier case):

- (6.a) 如果他不 尝试, 他将 失去 机会 。
 if he NEG try, he will lose opportunity
 “If he doesn’t try, he will lose the opportunity.”
- *(6.b) 如果他不 尝试 。
 if he NEG try
 “If he doesn’t try.”

In this case, we test the surprisal at the sentence-final period. If the model correctly represents the gross syntactic state within the subordinate clause,

³We do not consider the conventional modifier types used in the other test suite classes because adding an ORC can be confounding — the language model might be surprised at the ungrammatical target region due to the RC verb instead of the main transitive verb. Therefore, we instead focus on SRCs for this particular class.

then it should assign higher surprisal to the period in sentences like (6.b) than in sentences like (6.a). We include modifiers (same types as in CLASSIFIER–NOUN COMPATIBILITY) before the subject noun inside the matrix clause.

3.3 Models

To investigate the effect of syntax modeling in learning the dependencies described in Section 3.2, we train four classes of neural language models by crossing two types of parameterization with two types of supervision. Two of our model classes are trained for vanilla next-word-prediction: Long Short-Term Memory networks (LSTM; Hochreiter and Schmidhuber, 1997) and Transformers (Vaswani et al., 2017). The remaining two model classes are based on the LSTM and Transformer architectures, but explicitly incorporate syntactic structure during training: Recurrent Neural Network Grammars (RNNG; Dyer et al., 2016) and Transformer-parameterized parsing-as-language-modelling models (PLM; Qian et al., 2021). While prior work on structural supervision in English language models has focused primarily on RNNGs, both RNNGs and PLMs are joint probabilistic models of terminal word sequences along with the corresponding constituency parses. Thus, they both explicitly model syntax (in contrast to their vanilla language modeling counterparts), while featuring different parameterizations.

We use the PyTorch implementation of the LSTM (Paszke et al., 2019). The Transformer and PLM models are based on the HuggingFace GPT-2 architecture (Conneau and Lample, 2019). While we use the model architecture equivalent to the size of pre-trained GPT-2, we do not use the pre-trained tokenizer. All of our models are trained on a pre-tokenized Mandarin corpus and share the same vocabulary for each training dataset. Model sizes are reported in Table 3a in Appendix B.

For the LSTM and Transformer models, we calculate the surprisal at the target region by taking the negative log of the model’s predicted conditional probability. We estimate the RNNGs and PLMs’ word surprisals with word-synchronous beam search (Stern et al., 2017), following Hale et al. (2018) and Wilcox et al. (2019). The action beam size is 100 and the word beam size is 10. For regions with multi-token content, we sum over the probabilities of each token.

As a baseline, we additionally implement an n -

gram model with Kneser-Ney Smoothing (Kneser and Ney, 1995) using the SRILM toolkit (Stolcke, 2002). For cases where the smoothed n -gram model assigns identical probabilities to the target region across different conditions in a test item, we do tie-breaking by randomly flipping a fair coin to determine the outcome for that particular item.

3.4 Corpus Data

We consider two datasets to explore how training data size affects models ability to acquire grammatical knowledge (similar to Hu et al., 2020). The LSTM and Transformer models are trained on the raw text only, whereas the RNNG and PLM models are trained with additional syntactic annotations.

Chinese Treebank (CTB) The Chinese Treebank (CTB 9.0; Xue, Nianwen et al., 2016) is a Chinese language corpus annotated with Penn Treebank-style (Marcus et al., 1993) constituency parses. We use the Newswire, Magazine articles, Broadcast news, Broadcast conversations, and Weblogs sections, as we expect these sources to contain well-formed sentences with a variety of syntactic constructions. We follow the split defined by Shao et al. (2017) to construct training, development, and test sets.

Xinhua News Data To investigate the effects of increased training data size on models’ syntactic generalization, we create a larger corpus combining CTB with a subset from the Xinhua News corpus (Wu, Zhibiao, 1995).⁴ The Xinhua corpus contains metadata and content for 406K Mandarin news articles, collected from three mainstream media sources. Only the article contents are used for our training purposes. The texts from Xinhua corpus are first split into sentences and then tokenized into words with SpaCy (Honnibal et al., 2020). We then obtain sentence parses with the Berkeley Neural parser (Kitaev and Klein, 2018; Kitaev et al., 2019). We filter out extremely long sentences (>100 tokens) and map tokens occurring less than twice in the training data to fine-grained UNK tokens. Appendix C reports full statistics of the training corpora.

We train ten types of language models, crossing model architecture (n -gram, LSTM, RNNG, Transformer, and PLM) with dataset (Chinese Tree-

bank and hybrid Xinhua dataset).⁵ Each model type is trained with multiple random seeds,⁶ and results are reported as averages across these instances. Model perplexity scores are reported in Table 3b.

4 Results

We begin by reporting the overall performance of the models on the test suite classes introduced in Section 3.2. Figure 1 shows accuracy scores averaged across test suites within each class.⁷ First, we note that the n -gram baseline overall performs the worst among all language models, which matches our expectations since syntactic dependencies beyond the 5-token window are difficult for the model to capture.

Turning to the neural models, we assess the effects of training data size and architecture by examining the mean accuracy scores across test suites for each model. We fit separate linear mixed-effects models comparing the effects of data size, using the `lme4` package in R (Bates et al., 2015). The dependent variable is the mean accuracy score (Figure 2). For each language model type, the main effect is a binary indicator of whether the model is trained on the CTB dataset or the larger Xinhua dataset. We include test suite class, modifier type, and model seed as random factors with random intercepts and slopes for the accuracy score. Across test suite classes, the Transformer-Xinhua models outperform their smaller CTB counterparts ($p < .05$), but the effect of data size is less clear for the LSTM, RNNG and PLM models.

Comparing different model architectures, we find that in the SUBORDINATION test suite class, RNNGs trained on the smaller CTB dataset achieve comparable performance to LSTMs trained on the larger Xinhua dataset ($p = .938$; linear mixed-effects model with model type as the main factor). Furthermore, for both VERB-NOUN COMPATIBILITY and SUBORDINATION, the RNNGs perform better than the LSTMs and Transformers when trained on the smaller CTB dataset (see Figure 1). These results suggest that an inductive bias for learning hierarchical structure may help in low-data Chinese settings, consistent with prior

⁵We denote the language models trained on the hybrid Xinhua dataset with the suffix “-Xinhua”.

⁶We train RNNG models for 2 random seeds, and all other neural models for 3 random seeds.

⁷For numerical accuracy scores, see Table 6 in Appendix E. For results on individual test suites, see Figure 7 in Appendix F.

⁴We include CTB to guarantee that this larger dataset covers the CTB vocabulary.




Figure 1: Accuracy by test suite class and model. For each test suite class, we average the accuracy across different modifier types (including the no modifier case). Error bars denote 95% CIs of the mean accuracy score.

	CTB		Xinhua	
	RNNG vs. LSTM	PLM vs. Transformer	RNNG vs. LSTM	PLM vs. Transformer
CLASSIFIER-NOUN	-	-	-	✓**
GARDEN PATH OBJECT	-	✓**	-	✗*
GARDEN PATH SUBJECT	-	-	-	-
VERB-NOUN	✓**	-	✗*	-
MISSING OBJECT	-	-	-	✗***
SUBORDINATION	✓**	-	✓**	-

Table 1: Comparison between language models that perform explicit syntax modeling (RNNGs and PLMs) and their vanilla counterparts (LSTMs and Transformers). ✓ represents statistically significant improvement in structurally supervised language models, and ✗ represents the opposite direction. *: $p \leq .05$, **: $p \leq .01$, ***: $p \leq .001$.




Figure 2: Mean accuracy scores by model type across all test suites.

experiments conducted in English language models (Hu et al., 2020).

Overall, we find only suggestive benefits of structural supervision. While prior work in English (Kuncoro et al., 2018; Wilcox et al., 2019; Hu et al., 2020; Qian et al., 2021) has shown that RNNGs and PLMs can improve upon their vanilla LSTM and Transformer counterparts, the improvement is relatively smaller for Mandarin. To compare the performance of models with and without structural supervision, we fit a generalized linear mixed-effect model on the binary outcome (whether or not the model predicts a positive surprisal difference)

for each test item, within each combination of test suite class, training data, and model parameterization. We consider a binary indicator of whether or not the model performs syntax modeling explicitly as the main factor, and include test item, modifier type and model seed as random factors. Table 1 summarizes the results. For the CTB-only models, the structurally supervised models (RNNG and PLM) achieve accuracy either significantly greater or comparable to the corresponding vanilla models (LSTM and Transformer) across all test suite classes. However, the pattern is less clear for the Xinhua-trained models: the structurally supervised models lead to both gains and losses in accuracy compared to the vanilla models. We conjecture that word segmentation and parsing errors in the automatically-annotated Xinhua dataset might have affected the learning process of the model. In addition, the Xinhua training data explored in this work is still not very large in size (<10 million tokens), so it could be that further benefits of syntactic supervision may be more pronounced with much larger training datasets. Nevertheless, the suggestive benefits of explicitly modelling syntax with very small amounts of data could have im-




Figure 3: Accuracy on MISSING OBJECT as function of modifier complexity, for each model class.

plications for language modeling in low-resource settings.

We also assess whether our models are better at capturing syntactic or semantic relationships. To do this, we group the six classes of test suites into three categories: syntactic dependency (MISSING OBJECT and SUBORDINATION), semantic relationship (CLASSIFIER–NOUN COMPATIBILITY and VERB–NOUN COMPATIBILITY), and a hybrid capturing semantically-driven syntactic state representations (GARDEN PATH SUBJECT and GARDEN PATH OBJECT). We find that the average accuracy score is higher in the syntactic test suites than in the semantic test suites ($p < .001$),⁸ suggesting that the language models in our study – including those with no explicit syntax modeling – find it easier to learn syntactic dependencies than semantic relationships.


4.1 Robustness to Intervening Content

Next, we investigate the effect of structural supervision on tracking dependencies across intervening content. We focus our analysis on MISSING OBJECT,⁹ as the modifiers considered in these test suites can be ordered according to their syntactic complexity (no SRC < single SRC < coordinated SRCs < embedded SRCs).


Figure 3 shows the models’ performance on this test suite class as a function of modifier complexity, ranging from least to most difficult along the horizontal axis of each subplot. The vanilla LSTMs and Transformers clearly degrade in performance as the intervening materials between the stimulus and the target region grow in length and complex-

⁸See Appendix D.1 for details.

⁹The modifiers used in the other test suite classes (adjective, ORC, SRC) are not as directly comparable, since they vary in multiple dimensions, not just complexity.



(a) GARDEN PATH OBJECT



(b) GARDEN PATH SUBJECT

Figure 4: Mean surprisal difference at target region (matched classifier – mismatched classifier).

ity. In contrast, it is also visually apparent that the RNNG models do not degrade sharply as modifiers get longer and more complex. We fit linear mixed-effects models to investigate the relationship between modifier type and accuracy for each language model.¹⁰ Our results appear to confirm that both the RNNGs and PLMs do not significantly degrade on the simple SRC and coordinated SRC modifiers (compared to the no-modifier baseline). For the most complex modifier (embedded SRC), all models suffer in accuracy, but the magnitude of this effect is smaller for the RNNGs and PLMs compared to the LSTMs and Transformers. Taken together, our results suggest that while structural supervision does not give the language models

¹⁰See Appendix D.2 for details.

a significant advantage compared to their vanilla counterparts in the accuracy scores, it seems to help the model maintain syntactic expectations despite the intervention of syntactically complex content.

4.2 Garden-Path Effects

Building upon the CLASSIFIER–NOUN COMPATIBILITY results, we investigate whether a mismatched local classifier–noun pair may serve as an early cue for the upcoming RC structure, inducing a garden path effect. Figure 1 shows that the neural models systematically perform better on GARDEN PATH OBJECT than GARDEN PATH SUBJECT. We conjecture that neural language models may implicitly predict an ORC modifying the subject noun regardless of the type of classifiers. Therefore, the language models may be more prepared to see an ORC modifying the subject “工厂” in (3.b) than in (2.b) with the object “工厂”.

To gain a better understanding of model performance on these two test suite classes, we examine the average difference in target-region surprisal values between sentences with and without local classifier–noun mismatch (Figure 4). The average difference gives detailed information on how each language model processes the garden path region, which is complementary to the binary success/failure score achieved by a model on a given test item. Figure 4a shows that the neural models have a positive average surprisal difference across test items for GARDEN PATH OBJECT. Furthermore, the magnitude of this difference increases with the inclusion of the larger Xinhua dataset, suggesting that with more data, models become more confident in taking the incongruence between the classifier and the noun as a pre-RC cue.¹¹ On the other hand, recall that all models perform rather poorly on GARDEN PATH SUBJECT (Figure 1). Figure 4b shows that only the RNNs trained on the Xinhua corpus output a statistically significant positive average surprisal difference ($p < .001$; one-sample t -test). PLM-Xinhua, although not statistically significant, has a positive mean surprisal difference as well. This is due to the fact that the magnitude of the surprisal differences predicted by the RNN-Xinhua and PLM-Xinhua models is greater when they exhibit the predicted garden-path effects, and smaller when they do not follow the predicted direction. Therefore, structural supervi-

¹¹This result also accords with prior findings that classifiers facilitate object-modifying RC processing (Wu et al., 2014, 2018).

sion may help models represent syntactic state in a more human-like way.

5 Conclusion

This work evaluates Mandarin Chinese language models on six grammatical relationships, including syntactic dependencies and semantic compatibilities. We use Mandarin as a case study for analyzing how the potential advantages of explicit syntax modeling (as performed by RNNs and PLMs) generalize from English to a typologically different language. Although structural supervision does not boost the sequential model’s learning in all relationships tested in this study, we find that it does allow the RNN and PLM models to learn dependencies robust to increasingly complex modifiers, as seen in the MISSING OBJECT test suites. Compared to the vanilla sequence-based LSTM and Transformer models, explicit syntactic modeling also seems to help with grammatical generalization in settings with small training data. We also find that Mandarin syntactic dependencies (such as tracking gross syntactic state within a subordinate clause) tend to be easier to learn than semantic dependencies (such as the compatibility between classifiers and nouns). This study is one of the first steps towards understanding the role structural inductive biases may play in learning semantic and syntactic relationships in typologically diverse languages.

Acknowledgements

We thank three anonymous reviewers for their feedback. J.H. acknowledges support from an NSF Graduate Research Fellowship under Grant No. 1745302. This work was supported by the MIT–IBM AI Research Lab and MIT’s Quest for Intelligence.

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 2888–2899, Hong Kong, China. Association for Computational Linguistics.
- Sergio Balari. 1992. Two types of agreement. In *Catalan Working Papers in Linguistics*, volume 2, pages 1–43.

- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48.
- Emily M. Bender. 2011. [On Achieving and Evaluating Language-Independence in NLP](#). *Linguistic Issues in Language Technology*, 6(3).
- Thomas G. Bever. 1970. The cognitive basis for linguistic structures. *Cognition and the Development of Language*, 279(362):1–61.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual Language Model Pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. [Recurrent neural network grammars](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 199–209, San Diego, California. Association for Computational Linguistics.
- 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.
- 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.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- John Hale, Chris Dyer, Adhiguna Kuncoro, and Jonathan Brennan. 2018. [Finding syntax in human encephalography with beam search](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2727–2736, Melbourne, Australia. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. [A systematic assessment of syntactic generalization in neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Jingyang Jiang and Haitao Liu. 2015. The effects of sentence length on dependency distance, dependency direction and the implications—based on a parallel English–Chinese dependency treebank. *Language Sciences*, 50:93–104.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. [Multi-lingual constituency parsing with self-attention and pre-training](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- Reinhard Kneser and Hermann Ney. 1995. Improved backing-off for m-gram language modeling. In *International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 181–184. IEEE.
- 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.
- Roger Levy and Galen Andrew. 2006. [Tregex and tsurgeon: tools for querying and manipulating tree data structures](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Yen-hui Audrey Li, Andrew Simpson, and Wei-Tien Dylan Tsai. 2015. *Chinese Syntax in a Cross-linguistic Perspective*. Oxford Studies in Comparative Syntax. Oxford University Press.
- 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.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330.
- 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.
- Aaron Mueller, Garrett Nicolai, Panayiota Petrou-Zeniou, Natalia Talmina, and Tal Linzen. 2020. [Cross-linguistic syntactic evaluation of word prediction models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5523–5539, Online. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An Imperative Style, High-Performance Deep Learning Library](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Peng Qian, Tahira Naseem, Roger Levy, and Ramón Fernandez Astudillo. 2021. [Structural guidance for transformer language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3735–3745, Online. 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.
- Yan Shao, Christian Hardmeier, Jörg Tiedemann, and Joakim Nivre. 2017. Character-based joint segmentation and POS tagging for Chinese using bidirectional RNN-CRF. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 173–183.
- Timothy Shopen. 1985. *Language Typology and Syntactic Description: Volume 1, Clause Structure*. Language Typology and Syntactic Description. Cambridge University Press.
- Mitchell Stern, Daniel Fried, and Dan Klein. 2017. [Effective inference for generative neural parsing](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1695–1700, Copenhagen, Denmark. Association for Computational Linguistics.
- Andreas Stolcke. 2002. SRILM: An extensible language modeling toolkit. In *Seventh International Conference on Spoken Language Processing*.
- 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*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8.
- Ethan Wilcox, Peng Qian, Richard Futrell, Miguel Ballesteros, and Roger Levy. 2019. [Structural supervision improves learning of non-local grammatical dependencies](#). 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 3302–3312, Minneapolis, Minnesota. Association for Computational Linguistics.
- Fuyun Wu, Elsi Kaiser, and Shrvan Vasishth. 2018. Effects of early cues on the processing of chinese relative clauses: Evidence for experience-based theories. *Cognitive science*, 42:1101–1133.
- Fuyun Wu, Yingyi Luo, and Xiaolin Zhou. 2014. [Building Chinese relative clause structures with lexical and syntactic cues: Evidence from visual world eye-tracking and reading times](#). *Language, Cognition and Neuroscience*, 29(10):1205–1226.
- Wu, Zhibiao. 1995. [Mandarin Chinese news text LDC95T13](#). *Linguistic Data Consortium*.
- Beilei Xiang, Changbing Yang, Yu Li, Alex Warstadt, and Katharina Kann. 2021. [CLiMP: A benchmark for Chinese language model evaluation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2784–2790, Online. Association for Computational Linguistics.
- Xue, Nianwen, Zhang, Xiuhong, Jiang, Zixin, Palmer, Martha, Xia, Fei, Chiou, Fu-Dong, and Chang, Meiyu. 2016. [Chinese Treebank 9.0 LDC2016T13](#). *Linguistic Data Consortium*.
- Meilin Zhan and Roger Levy. 2018. [Comparing theories of speaker choice using a model of classifier production in Mandarin Chinese](#). 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 1997–2005, New Orleans, Louisiana.
Association for Computational Linguistics.

A Individual Test Suites

A.1 List of Test Suites

See Table 2 for a summary of test suite classes constructed in this study.

A.2 Complementary Test Suite Examples

In this section, we provide complementary test items for those in Section 3.2, constructing a full set of examples for each test suite class. Note that except for MISSING OBJECT where we consider variants of SRCs as the modifier (SRC, coordinated SRC, and embedded SRC), the general modifiers types are adjective, ORC and SRC.

A.2.1 Classifier–Noun Compatibility

No modifier:

- (7.a) 孩子听到了一首 歌曲。
child hear PST one CL_{SONG} song.
“The child heard a song.”
- * (7.b) 孩子听到了一张 歌曲。
child hear PST one CL_{ALBUM} song.
“The child heard a song.”
- (7.c) 孩子听到了一张 专辑。
child hear PST one CL_{ALBUM} album.
“The child heard an album.”
- * (7.d) 孩子听到了一首 专辑。
child hear PST one CL_{SONG} album.
“The child heard an album.”

ORC as modifier:

- (8.a) 孩子听到了一首 他熟悉的歌曲。
child hear PST one CL_{SONG} he familiar DE song.
“The child heard a song that he is familiar with.”
- * (8.b) 孩子听到了一张 他熟悉的歌曲。
child hear PST one CL_{ALBUM} he familiar DE song.
“The child heard a song that he is familiar with.”
- (8.c) 孩子听到了一张 他熟悉的专辑。
child hear PST one CL_{ALBUM} he familiar DE album.
“The child heard an album that he is familiar with.”
- * (8.d) 孩子听到了一首 他熟悉的专辑。
child hear PST one CL_{SONG} he familiar DE album.
“The child heard an album that he is familiar with.”

SRC as modifier:

- (9.a) 孩子听到了一首 令他熟悉的歌曲。
child hear PST one CL_{SONG} make he familiar DE song.
“The child heard a song that makes him feel familiar.”

- * (9.b) 孩子听到了一张 令他熟悉的歌曲。
child hear PST one CL_{ALBUM} make he familiar DE song.
“The child heard a song that makes him feel familiar.”
- (9.c) 孩子听到了一张 令他熟悉的专辑。
child hear PST one CL_{ALBUM} make he familiar DE album.
“The child heard an album that makes him feel familiar.”
- * (9.d) 孩子听到了一首 令他熟悉的专辑。
child hear PST one CL_{SONG} make he familiar DE album.
“The child heard an album that makes him feel familiar.”

A.2.2 Garden-Path Object

Adjectival modifier:

- (10.a) 他离开了那间 负责 的 朋友 开 的工厂。
he leave PST that CL_{BUILDING} conscientious DE friend start DE factory .
- (10.b) 他离开了那个 负责 的 朋友 开 的工厂。
he leave PST that CL_{GENERAL} conscientious DE friend start DE factory .
“He left the factory that the conscientious friend started.”

ORC as modifier:

- (11.a) 他离开了那间 我尊敬 的 朋友 开 的工厂。
he leave PST that CL_{BUILDING} I respect DE friend start DE factory .
- (11.b) 他离开了那个 我尊敬 的 朋友 开 的工厂。
he leave PST that CL_{GENERAL} I respect DE friend start DE factory .
“He left the factory that the friend whom I respect started.”

SRC as modifier:

- (12.a) 他离开了那间 帮助过 我的 朋友 开 的工厂。
he leave PST that CL_{BUILDING} help PST I DE friend start DE factory .
- (12.b) 他离开了那个 帮助过 我的 朋友 开 的工厂。
he leave PST that CL_{GENERAL} help PST I DE friend start DE factory .
“He left the factory that the friend who helped me before started.”

A.2.3 Garden-Path Subject

Adjectival modifier:

- (13.a) 那间 负责 的 朋友 开 的工厂 倒闭 了。
that CL_{BUILDING} conscientious DE friend start DE factory close PST .

Test Suite Class	Modifier Type	# Test Items	Prior Work in English	Prior Work in Chinese
CLASSIFIER–NOUN COMPATIBILITY	None	30	-	Zhan and Levy (2018); Xiang et al. (2021)
CLASSIFIER–NOUN COMPATIBILITY	Adjective	30	-	Xiang et al. (2021)
CLASSIFIER–NOUN COMPATIBILITY	Object-Extracted RC	30	-	Xiang et al. (2021)
CLASSIFIER–NOUN COMPATIBILITY	Subject-Extracted RC	30	-	Xiang et al. (2021)
GARDEN PATH OBJECT	None	31	-	Wu et al. (2018)
GARDEN PATH OBJECT	Adjective	31	-	-
GARDEN PATH OBJECT	Object-Extracted RC	31	-	-
GARDEN PATH OBJECT	Subject-Extracted RC	31	-	-
GARDEN PATH SUBJECT	None	31	Futrell et al. (2019); Hu et al. (2020)	-
GARDEN PATH SUBJECT	Adjective	31	-	-
GARDEN PATH SUBJECT	Object-Extracted RC	31	Hu et al. (2020)	-
GARDEN PATH SUBJECT	Subject-Extracted RC	31	Hu et al. (2020)	-
VERB–NOUN COMPATIBILITY	None	31	-	-
VERB–NOUN COMPATIBILITY	Adjective	31	-	-
VERB–NOUN COMPATIBILITY	Object-Extracted RC	31	-	-
VERB–NOUN COMPATIBILITY	Subject-Extracted RC	31	-	-
MISSING OBJECT	None	30	Warstadt et al. (2020)	-
MISSING OBJECT	Subject-Extracted RC	30	-	-
MISSING OBJECT	Coordinated Subject-Extracted RC	30	-	-
MISSING OBJECT	Embedded Subject-Extracted RC	30	-	-
SUBORDINATION	None	30	Futrell et al. (2019); Hu et al. (2020)	-
SUBORDINATION	Adjective	30	-	-
SUBORDINATION	Object-Extracted RC	30	Futrell et al. (2019); Hu et al. (2020)	-
SUBORDINATION	Subject-Extracted RC	30	Futrell et al. (2019); Hu et al. (2020)	-

Table 2: Summary of individual test suites used in our experiments.

(13.b) 那个负责的朋友开的
that CL_{GENERAL} conscientious DE friend start DE
工厂倒闭了。
factory close PST .
“The factory that the conscientious friend started has closed.”

ORC as modifier:

(14.a) 那间我尊敬的朋友开的
that CL_{BUILDING} I respect DE friend start DE
工厂倒闭了。
factory close PST .
(14.b) 那个我尊敬的朋友开的
that CL_{GENERAL} I respect DE friend start DE
工厂倒闭了。
factory close PST .
“The factory that the friend whom I respect started has closed.”

SRC as modifier:

(15.a) 那间帮助过我的朋友开的
that CL_{BUILDING} help PST I DE friend start DE
工厂倒闭了。
factory close PST .
(15.b) 那个帮助过我的朋友开的
that CL_{GENERAL} help PST I DE friend start DE
工厂倒闭了。
factory closed PST .
“The factory that the friend who helped me before started has closed.”

A.2.4 Verb–Noun Compatibility

No modifier:

(16.a) 我修理了这个电脑。
I fix PST this CL computer .
“I have fixed this computer.”
*(16.b) 我阅读了这个电脑。
I read PST this CL computer .
“I have read this computer.”

ORC as modifier:

(17.a) 我修理了这个父亲使用的电脑。
I fix PST this CL father use DE computer .
“I have fixed this computer that the father uses.”
*(17.b) 我阅读了这个父亲使用的电脑。
I read PST this CL father use DE computer .
“I have read this computer that the father uses.”

SRC as modifier:

(18.a) 我修理了这个计算公式的
I fix PST this CL calculate formula DE
电脑。
computer .
“I have fixed this computer that calculates formulas.”
*(18.b) 我阅读了这个计算公式的
I read PST this CL calculate formula DE
电脑。
computer .
“I have read this computer that calculates formulas.”

A.2.5 Missing Object

SRC as modifier:

(19.a) 记者采访了研发产品的
journalist interview PST develop product DE
科学家。
scientist .
“The journalist interviewed the scientist who developed the product.”
*(19.b) 记者采访了研发产品。
journalist interview PST develop product .
“The journalist interviewed who developed the product.”

Coordinated SRCs as modifier:

(20.a) 记者采访了研发产品并且获奖的科学家。
journalist interview PST develop product and win
了奖的科学家。
PST prize DE scientist .

- “The journalist interviewed the scientist who developed the product and won a prize.”
- *(20.b) 记者 采访 了 研发 产品 并且
journalist interview PST develop product and
获 了 奖 。
win PST prize .
“The journalist interviewed who developed the product and won a prize.”

Embedded RCs as modifier:

- (21.a) 记者 采访 了 研发 帮助 老人 的
journalist interview PST develop help elderly DE
产 品 的 科 学 家 。
product DE scientist .
“The journalist interviewed the scientist who developed the product that helps the elderly.”
- *(21.b) 记者 采访 了 研发 帮助 老人 的
journalist interview PST develop help elderly DE
产 品 。
product .
“The journalist interviewed who developed the product that helps the elderly.”

A.2.6 Subordination

Adjectival modifier:

- (22.a) 如果 内 向 的 他 不 尝 试 ， 他 将 失 去
if introverted DE he NEG try , he will lose
机 会 。
opportunity .
“If he who is introverted doesn’t try, he will lose the opportunity.”
- *(22.b) 如果 内 向 的 他 不 尝 试 。
if introverted DE he NEG try .
“If he who is introverted doesn’t try.”

ORC as modifier:

- (23.a) 如果 父 亲 期 待 的 他 不 尝 试 ， 他 将
if father expect DE he NEG try , he will
失 去 机 会 。
lose opportunity .
“If he who the father has expectations of doesn’t try, he will lose the opportunity.”
- *(23.b) 如果 父 亲 期 待 的 他 不 尝 试 。
if father expect DE he NEG try .
“If he who the father has expectations of doesn’t try.”

SRC as modifier:

- (24.a) 如 果 没 有 工 作 的 他 不 尝 试 ， 他 将 失 去
if NEG job DE he NEG try , he will lose
机 会 。
opportunity .
“If he who has no jobs doesn’t try, he will lose the opportunity.”
- *(24.b) 如 果 没 有 工 作 的 他 不 尝 试 。
if NEG job DE he NEG try .
“If he who has no jobs doesn’t try.”

Model	# layers	# hidden units	Emb size
LSTM	2	256	256
RNNG	2	256	256
Transformer	12	768	768
PLM	12	768	768

(a)

Model	CTB	Xinhua
<i>n</i> -gram	330	332
LSTM	161	190
RNNG	227	194
Transformer	234	170
PLM	297	244

(b)

Table 3: (a) Model architecture size. (b) Perplexity results on CTB test data.

B Model Information

We find that the perplexity score reported in Table 3b is comparatively high for RNNG compared to that reported in Dyer et al. (2016). This may be because the CTB data we use includes some informal and spoken language, such as weblogs and broadcast conversations.

C Corpus Statistics

Corpus	# Tokens	Vocab Size
CTB	974K	27K
Xinhua	7M	88K

Table 4: Statistics of training corpora.

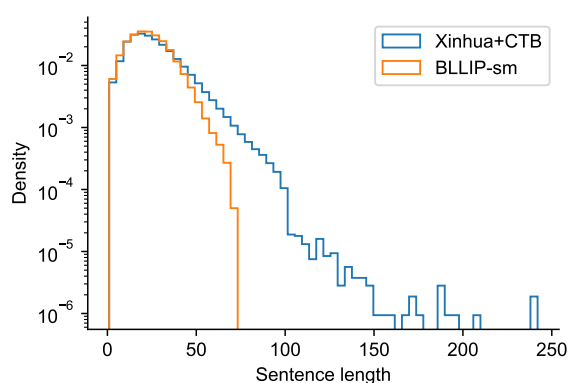


Figure 5: Comparing the distribution of sentence length in Mandarin and English corpora.

Figure 5 shows differences in the distribution of sentence lengths between the Mandarin corpus and a comparable English corpus of similar register. The maximum sentence length is 242 words for

Xinhua+CTB, and only 70 words for the English BLLIP-sm corpus used to train Hu et al.’s (2020) models. The distribution of sentence lengths also has a heavier tail for the Mandarin corpus than the English corpus.

D Additional Analysis

D.1 Comparing Syntactic and Semantic Test Suites

In this section, we focus on comparing how well models learn about syntactic dependency and semantic compatibility.

Recall that we group CLASSIFIER–NOUN COMPATIBILITY and VERB–NOUN COMPATIBILITY as semantic relationships, and MISSING OBJECT and SUBORDINATION as syntactic dependency. We compute the accuracy scores of these two categories, as shown in Figure 6. Here we exclude the n -gram models since their performances deviate from the other language models to a great extent. Adding modifiers between the stimulus and the target region seems to shrink the gap a bit. And this again might be due to the fact that for MISSING OBJECT, we intentionally make the intervening contents increasingly hard to learn, dragging down the average accuracy score for syntactic dependencies.




Figure 6: Accuracy on semantic and syntactic test suite classes, with and without intervening content. n -gram results are excluded.

Similar to testing the effect of data size, we determine the statistical significance here with a linear mixed-effects model on the accuracy score. We predict accuracy with a binary indicator of whether or not the test is in the syntactic group that we define, including model type and test item as random factors. We find that syntactic relationships seem to be easier to learn than semantic ones regardless of the intervening content ($p < .001$).

D.2 Intervening Content in MISSING OBJECT

In this section, we discuss our statistical analysis on language models’ robustness against intervening contents.

We fit separate linear mixed-effects models for each language model with the accuracy score as the dependent variable, and modifier type as the predictor. We include model seed and data size as random factors, both with random intercepts and random slopes. Table 5 summarizes our results. Recall that in MISSING OBJECT, we consider four types of modifiers: None, single SRC, coordinated SRCs, and embedded SRCs. Each cell in Table 5 represents coefficients of a particular modifier with respect to the None modifier baseline for a particular language model class. For single SRC and coordinated SRC modifiers, neither of the RNNG or PLM models show significant degradation in the accuracy score. All models suffer from the embedded SRC modifier, with negative coefficients that are all statistically significant. However, RNNGs and PLMs seem to be affected the least, with larger coefficients than vanilla LSTMs and Transformers. This suggests that structural supervision helps language model to learn syntactic dependencies that are more robust against the intervening content.

Model	Single SRC	Coordinated SRC	Embedded SRC
LSTM	-0.06*	-0.14***	-0.3***
RNNG	-0.016	-0.025	-0.175***
Transformer	-0.1**	-0.13*	-0.35***
PLM	0.05*	-0.01	-0.2***

Table 5: Coefficients of the modifiers in MISSING OBJECT with the condition of no modifier as the baseline. *: $p \leq .05$, **: $p \leq .01$, ***: $p \leq .001$.

E Accuracy by Model and Test Suite Class

See Table 6 for a summary of accuracy scores achieved by each model across test suite classes.

F Results on Individual Test Suites

Figure 7 shows language models’ accuracy scores on the six test suite classes: CLASSIFIER–NOUN COMPATIBILITY, GARDEN PATH OBJECT, GARDEN PATH SUBJECT, VERB–NOUN COMPATIBILITY, MISSING OBJECT, and SUBORDINATION. On the x-axis, we have four types of modifiers tested on that specific test suite class.

Model	CLASSIFIER–NOUN	GARDEN PATH OBJECT	GARDEN PATH SUBJECT	VERB–NOUN	MISSING OBJECT	SUBORDINATION
<i>n</i> -gram	0.552	0.508	0.484	0.484	0.575	0.675
LSTM	0.598	0.659	0.320	0.624	0.847	0.789
RNNG	0.609	0.690	0.359	0.714	0.838	0.854
Transformer	0.603	0.672	0.376	0.581	0.761	0.817
PLM	0.615	0.750	0.419	0.589	0.775	0.836
<i>n</i> -gram-Xinhua	0.594	0.516	0.516	0.500	0.700	0.500
LSTM-Xinhua	0.650	0.691	0.355	0.782	0.819	0.850
RNNG-Xinhua	0.636	0.750	0.367	0.714	0.854	0.913
Transformer-Xinhua	0.708	0.720	0.363	0.755	0.792	0.931
PLM-Xinhua	0.746	0.656	0.395	0.745	0.692	0.908

Table 6: Accuracy score by model and test suite class. Blue color denotes the best score within the CTB dataset.

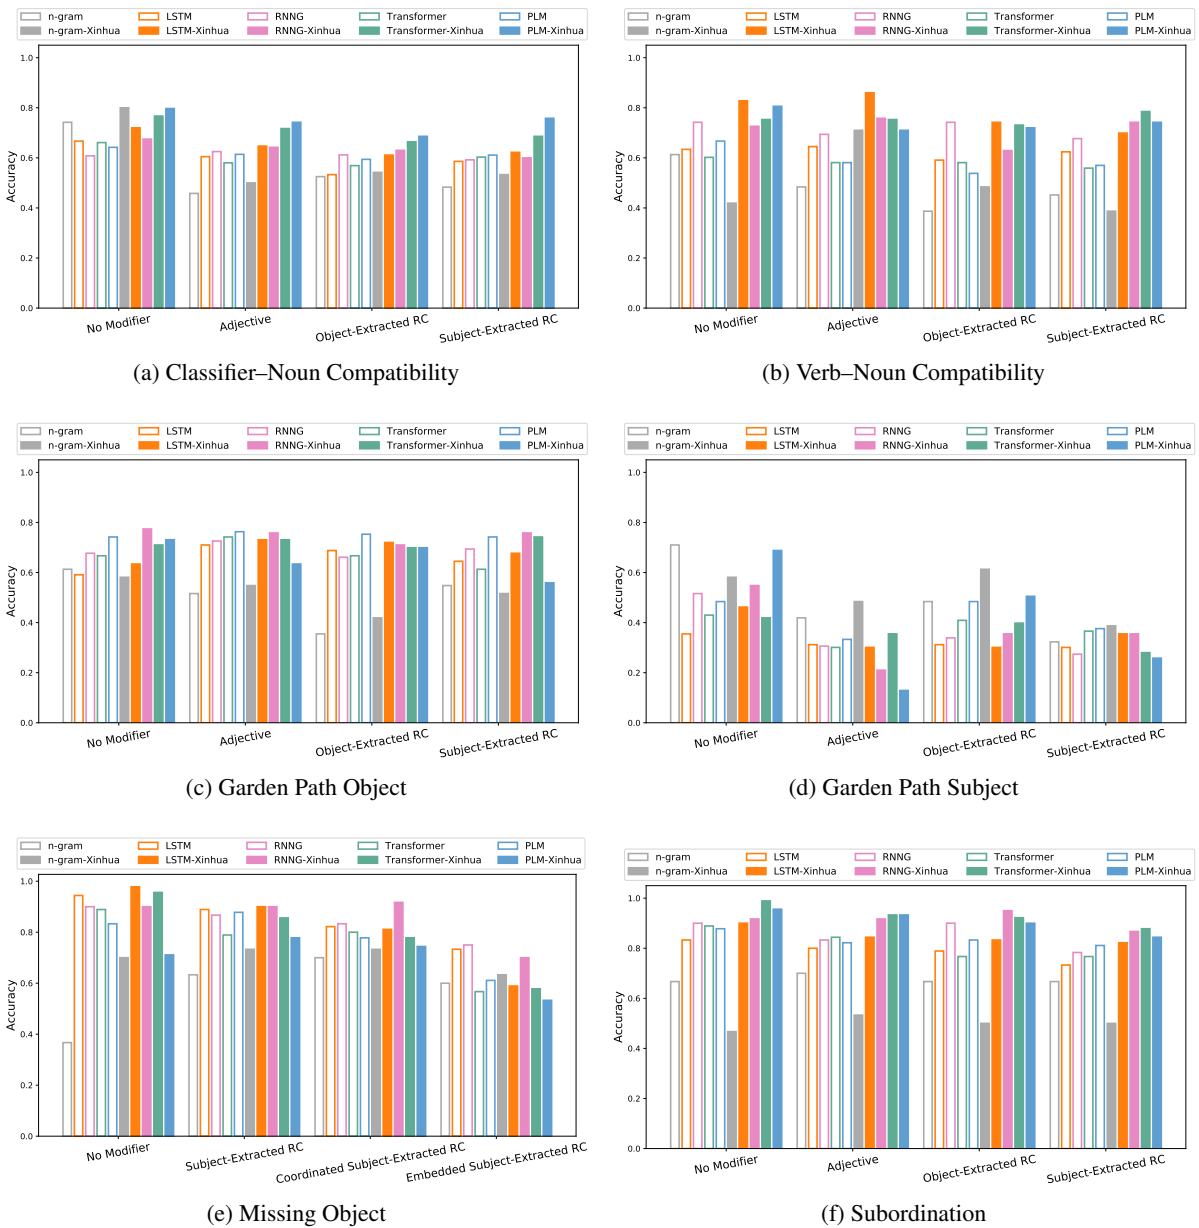


Figure 7: Accuracy on individual test suites used in our experiments.