

Probing for targeted syntactic knowledge through grammatical error detection

Christopher Davis[†] Christopher Bryant[†] Andrew Caines[†]
Marek Rei^{‡†} Paula Buttery[†]

[†]ALTA Institute, Department of Computer Science & Technology,
University of Cambridge, U.K.

[‡]Imperial College London, U.K.

[†]{ccd38, cjb255, apc38, pjb48}@cam.ac.uk

[‡]marek.rei@imperial.ac.uk

Abstract

Targeted studies testing knowledge of subject-verb agreement (SVA) indicate that pre-trained language models encode syntactic information. We assert that if models robustly encode subject-verb agreement, they should be able to identify when agreement is correct and when it is incorrect. To that end, we propose grammatical error detection as a diagnostic probe to evaluate token-level contextual representations for their knowledge of SVA. We evaluate contextual representations at each layer from five pre-trained English language models: BERT, XLNET, GPT-2, ROBERTA, and ELECTRA. We leverage public annotated training data from both English second language learners and Wikipedia edits, and report results on manually crafted stimuli for subject-verb agreement. We find that masked language models linearly encode information relevant to the detection of SVA errors, while the autoregressive models perform on par with our baseline. However, we also observe a divergence in performance when probes are trained on different training sets, and when they are evaluated on different syntactic constructions, suggesting the information pertaining to SVA error detection is not robustly encoded.

1 Introduction

Recent work investigates whether linguistic information is encoded in pre-trained transformer-based language models (Peters et al., 2018; Devlin et al., 2019). Research using diagnostic methods (Shi et al., 2016; Alain and Bengio, 2017; Adi et al., 2017; Conneau et al., 2018; Hupkes et al., 2018) indicates models encode syntax via experiments targeting, for example, part-of-speech and dependency labelling (Tenney et al., 2019; Jawahar et al., 2019; Hewitt and Manning, 2019), while targeted syntactic evaluation studies show models encode a large amount of hierarchical syntactic information in tests for subject-verb agreement (Linzen et al., 2016; Marvin and Linzen, 2018; Goldberg,

2019). Although previous research has covered a large number of probing tasks (Tenney et al., 2019; Liu et al., 2019a), no one has yet fully explored grammatical error detection (GED) as a probe. We assert that the ability to detect ungrammatical tokens serves as a complementary evaluation to assess linguistic competence.

GED is a natural and complex NLP task that assesses a model’s ability to detect which tokens in a sentence are grammatically incorrect. Ungrammatical tokens may be categorised within a taxonomy¹ comprising three operational categories (replacement, unnecessary, and missing) and twenty-five categories based on parts-of-speech. For example:

- (1) [Replacement subject-verb agreement] *The train **are** a good option for long trips.*
- (2) [Replacement pronoun] *Everybody must have free time for **yourself**.*
- (3) [Missing determiner] *The birth of [**a**] new star.*
- (4) [Unnecessary preposition] *Public transport means travelling around [...] **by** using trains, buses, and planes.*

To do well in the task, a model must encode and make use of a wide array of linguistic information. For example, detecting subject-verb agreement errors in English tests a model’s capacity to identify i) verbs, ii) the subjects of the verbs, iii) the grammatical number (singular/plural) of both, and iv) whether their number agrees.

The above makes the task a very interesting test-bed for evaluating a model’s syntactic knowledge. We operationalise the GED task and train probes to detect replacement subject-verb agreement errors (as in Example 1) using contextual representations from different hidden layers from five pre-trained English language models – BERT (Devlin et al.,

¹Much recent research in GED uses error-type labels based on ERRANT (Bryant et al., 2017).

2019), XLNET (Yang et al., 2019), GPT-2 (Radford et al., 2019), ROBERTA (Liu et al., 2019b), and ELECTRA (Clark et al., 2020).

To ensure a robust and thorough evaluation, we leverage existing publicly annotated data from two domains for training: essays by learners of English as a second language from both the Cambridge English Write & Improve + LOCNESS (W&I) corpus (Bryant et al., 2019) and the First Certificate in English corpus (FCE) (Yannakoudakis et al., 2011), along with a corpus of automatically extracted edited sentences from native English Wikipedia edit histories (Grundkiewicz and Junczys-Dowmunt, 2014). For evaluation, we reframe the minimal-pair dataset from Marvin and Linzen (2018) to create targeted evaluation sets annotated for GED. In doing so, we demonstrate how existing minimal-pair datasets can be leveraged to create challenging and interpretable test sets for GED models² (Hu et al., 2020).

We find that ELECTRA, BERT, and ROBERTA linearly encode information for the detection of SVA errors in the contextual representations of verbs, however, we observe a gap in performance when probes are trained on data from different domains, implying the information is not encoded consistently or robustly. The results show consistent patterns across layers: both BERT and ELECTRA encode information related to SVA errors in the middle-to-late layers, while ROBERTA seems to encode information earlier in the model. Probes trained on representations from GPT-2 and XLNET (with uni- and bi-directional decoding) perform poorly on the evaluation set, indicating a fundamental difference from either the training objective or pre-training data. Finally, we show that GED probes can complement existing tools for syntactic evaluation: our results suggest that although neural language models perform well on targeted syntactic evaluation tasks, their encoding of SVA does not robustly extend to the detection of SVA errors.³

2 Token-level grammaticality

We motivate the use of GED-probes by first reviewing previous literature involving grammaticality judgements and tests for subject-verb agreement, then discuss the advantages in tests for GED.

Boolean acceptability judgements have long been used as a primary behavioural measure to observe humans’ grammatical knowledge (Chomsky, 1957; Pater, 2019), and have recently been employed in computational linguistics to evaluate grammatical knowledge in neural models. For example, Warstadt et al. (2019) train classifiers to predict sentence-level Boolean acceptability judgements on example sentences from the linguistics literature. As each sentence is designed to demonstrate a particular grammatical construction, performance on the task is interpreted as a reflection of the implicit knowledge of the classifier.

An alternative approach frames acceptability as a choice between minimal pairs of sentences – one grammatical and another ungrammatical, where the difference between the two is typically one or two tokens. Marvin and Linzen (2018) evaluate linguistic knowledge by testing whether a language model assigns higher probability to a grammatical sentence relative to its minimally different ungrammatical counterpart. Similar to Warstadt et al. (2019), fine-grained grammatical knowledge is evaluated by controlling the evaluation stimuli, with the hypothesis that models must have implicit knowledge of the underlying grammatical concept to succeed.

Rather than evaluating sentence-level scores, Linzen et al. (2016) compare predicted probabilities assigned to target verbs in minimal pair sentences, where each sentence in a pair uses a different form of the verb. Goldberg (2019) extends this to masked language models where he replaces a target verb with the [MASK] token and feeds the entire sentence to a BERT model. A model is considered successful, and thereby has knowledge related to SVA, if it assigns higher probability to the correct form of the verb.

Our work differs from the above in three important ways. First, we don’t assume to know where the incorrect token is – the probe is trained to detect errors for all tokens in a sentence, given each token’s contextual representation. This is a more fine-grained evaluation compared to sentence-level judgements and tests whether probes know where the error is located. Second, instead of targeting information in the masked token, we investigate whether the model implicitly encodes SVA information in a token’s contextual representation. Third, we test for knowledge of SVA without comparing to the counterpart token or sentence. We argue that if a model has knowledge of SVA, it should

²In principle these minimal-pair datasets can also be used to evaluate grammatical error correction systems.

³We release our code at <https://github.com/chrisdavis90/ged-syntax-probing>

Syntactic Construction	Example
Simple agreement	The author <u>laughs/laugh</u> *
Agreement in a sentential complement	The bankers knew the officer <u>smiles/smile</u> *
Agreement across a prepositional phrase	The farmer near the parents <u>smiles/smile</u> *
Agreement across a subject relative clause	The officers that love the skater <u>smile/smiles</u> *
Short verb-phrase coordination	The senator smiles and <u>laughs/laugh</u> *
Long verb-phrase coordination	The manager writes in a journal every day and <u>likes/like</u> * to watch television shows
Agreement across an objective relative clause	The farmer that the parents love <u>swims/swim</u> *
Agreement within an objective relative clause	The farmer that the parents <u>love/loves</u> * swims

Table 1: Examples for the main syntactic constructions from the subject-verb-agreement stimuli from [Marvin and Linzen \(2018\)](#). **Bold** indicates the subject-noun, and underlined tokens indicate the grammatical/ungrammatical* verb.

be able to detect SVA-errors without requiring a comparison.

3 Data

3.1 Second language learner corpora

Following previous work in grammatical error correction and detection, we use the Cambridge English Write & Improve + LOCNESS corpus ([Bryant et al., 2019](#)) and the First Certificate in English ([Yannakoudakis et al., 2011](#)), hereinafter W&I-FCE.⁴

The edit annotations in these corpora were pre-processed and standardised using the ERRANT annotation framework ([Bryant et al., 2017](#)). One advantage of this framework is that error types are modular, and consist of “operation” + “main” type tags. This provides us flexibility to target grammatical errors at different levels of granularity. E.g. all NOUN errors or only R:NOUN for replacement nouns. In addition, we can take advantage of the corrections provided with each edit annotation to control the number and variation of grammatical errors. Since we focus only on replacement subject-verb-agreement errors, R:VERB:SVA, we correct all other error types and keep only those sentences containing at least one grammatical error. This leaves 1936 sentences for training and 142 sentences for validation.

3.2 Dataset of Wikipedia edits

As an alternative to the learner corpora, we additionally experiment with a corpus of automatically extracted edited sentences from native English Wikipedia edit histories (WIKED) ([Grundkiewicz](#)

[and Junczys-Dowmunt, 2014](#)). We use the clean and preprocessed version of English Wikipedia edits, consisting of ~29 million sentences.⁵ We follow the same procedure as above and retain only sentences with R:VERB:SVA errors by correcting all other error types, and keep only the sentences containing at least one error. This leaves ~233K sentences, from which we sample five training sets each with 1936 sentences each to match the amount of sentences in the learner corpora after processing, and 5839 sentences for the validation set. Statistics for both corpora are given in [Appendix A](#). We refer to the sampled training sets as WIKED-S.

3.3 Minimal-pair datasets

We use the manually constructed subject-verb agreement stimuli from [Marvin and Linzen \(2018\)](#) (M&L) to evaluate the GED-probes – this enables a more controlled evaluation compared to the naturally occurring sentences in W&I-FCE and WIKED. The dataset consists of seven main syntactic constructions, shown in [Table 1](#). In addition to those shown in the table, sentences with multiple nouns (except for those testing VP coordination) include instances with two nouns and one acts as a distractor, potentially agreeing with the verb even though it is not the subject:

- (5)
- a. The farmer near the parent smiles/smile*.
 - b. The farmer near the parents smiles/smile*.
 - c. The farmers near the parent smiles*/smile.
 - d. The farmers near the parents smiles*/smile.

In the above sentences, the verbs marked with an asterisk are ungrammatical. The dataset also expands on sentences testing agreement with object relative clauses: agreement is tested across

⁴Public data for W&I and the FCE are available at: <https://www.cl.cam.ac.uk/research/nl/bea2019st#data>

⁵<https://github.com/snukky/wikiedits>

and within the clause, using animate and inanimate main subjects, and with and without the *that*-complementizer.

We process the M&L minimal pairs to create token-level GED annotations, where the ungrammatical verbs are tagged as R:VERB:SVA.⁶ We include all of the ungrammatical and grammatical sentences for evaluation – to be successful, the GED-probe should recognise when the agreement is correct and label all tokens as grammatical.⁷ Finally, we capitalise the first word in each sentence and add a full stop if one doesn't already exist. [Appendix B](#) contains details about the processed dataset.

4 Experiment 1: Per Layer Probes

We first investigate whether models encode SVA-errors by examining probing performance at each layer; we want to test whether models encode this information in the final layer, the token representation, but also how this information develops across the layers. We then break-down performance by syntactic construction to better understand how the SVA-error encoding generalizes. Finally, we carry out a follow-up experiment to investigate the impact of training data size and verb frequency.

4.1 Experimental setup

For each model, we extract contextual representations for every token in a sentence for every one of the twelve layers in the model. We then train a linear probe (Tenney et al., 2019; Hupkes et al., 2018; Liu et al., 2019a) per layer to predict whether each token-level contextual representation is ungrammatical. We train two versions of each probe: one trained using W&I-FCE, and the other using WIKED-S. Every probe is evaluated on the M&L stimuli

Since the probe is trained to detect R:VERB:SVA errors, high probing performance would indicate the probe has learned to extract features to identify subject-verb agreement errors from the contextual representations of the verbs. This implicitly includes sub-tasks to identify the verb, the subject noun, the number of both the verb and noun, and

⁶In principle, any minimal-pair dataset can be converted to token-level annotations using ERRANT, but not all grammatical errors map cleanly to ERRANT categories. For example, replacement pronouns (R:PRON) includes reflexive anaphor gender- and number- agreement errors.

⁷While the evaluation stimuli consists of minimal pairs, the training data does not.

that their number disagrees. Furthermore, as this is a token labelling probe, high probing performance would indicate the pre-trained model has encoded the relevant features in the contextual representations of the verbs.

We evaluate probes using F_1 on the evaluation stimuli from M&L containing an equal number of grammatical and ungrammatical sentences.⁸ We compare probes to a VERB-ONLY baseline which incorrectly tags all verbs as ungrammatical. The number of verbs per sentence varies across syntactic constructions; constructions with one verb have an equal number of grammatical and ungrammatical verbs, and therefore have a baseline score of 0.67. Constructions with two and three verbs have scores of 0.40 and 0.30, respectively. Evaluating the baseline over all constructions yields a score of 0.43.

We evaluate five pre-trained models: BERT-BASE-CASED, GPT-2 (small), ROBERTA-BASE, XLNET with both uni-directional (XLNET-UNI) and bi-directional XLNET-BI decoding, and ELECTRA-BASE (discriminator). As all five models use sub-word tokenisation, we follow Liu et al. (2019a) and use the last sub-word unit for token classification. We train the probes for 50 epochs with a patience of 10 epochs for early stopping based on in-domain validation sets.

Four of the five models were pre-trained with a language modelling objective: either masked language modelling (MLM) or autoregressive language modelling (ALM). Whereas ELECTRA is the exception – the *replaced token detection* training objective is somewhat aligned with GED and therefore we may expect representations to encode grammatically discriminative information.⁹ Indeed, Yuan et al. (2021) find ELECTRA outperforms BERT when fine-tuned for binary GED targeting a wide range of error-types. BERT and ROBERTA also detect replaced tokens during training, but only on 1.5% of tokens.

⁸This departs from $F_{0.5}$ used in the GED literature, which was motivated from educational applications where high precision is preferred over recall because false-positives can be more harmful for language learners compared to false-negatives.

⁹The ELECTRA discriminator model is trained to detect substituted tokens in a grammatical sentence, where an original token is substituted with a plausible alternative from a masked language model.

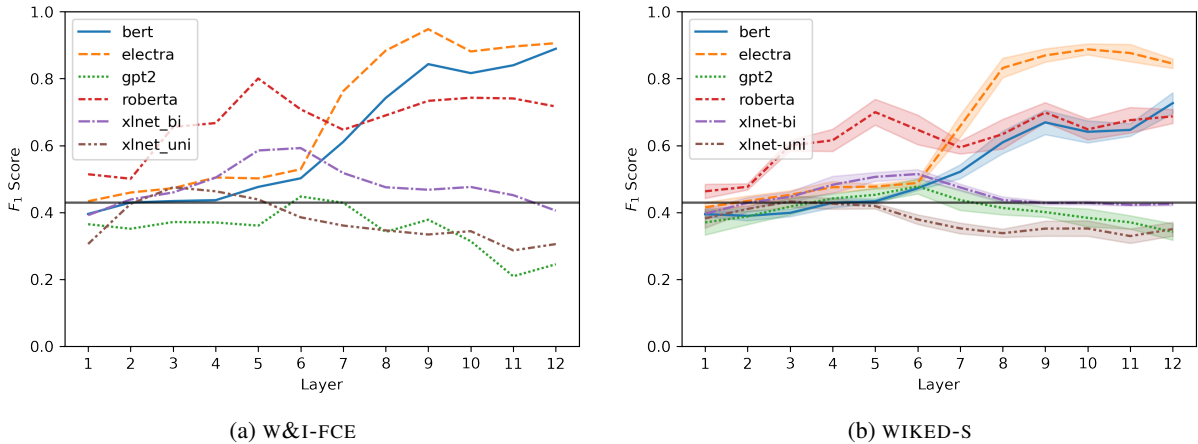


Figure 1: F_1 scores for probes trained on contextual representations at different layers from BERT, ELECTRA, ROBERTA, XLNET with bidirectional decoding (XLNET-BI), XLNET with unidirectional decoding (XLNET-UNI), and GPT-2. **1a** and **1b** show results for probes trained on W&I-FCE and the WIKED-S training sets respectively. The VERB-ONLY baseline scores are illustrated using grey horizontal lines. **1b** shows the mean and standard deviation across the five training sets (§3.2). All probes are evaluated on the M&L stimuli (§3.3).

Model	W&I-FCE		WIKED-S	
	Layer	F_1	Layer	F_1
ELECTRA	9	0.95	10	$0.89_{\pm 0.02}$
BERT	12	0.89	12	$0.73_{\pm 0.03}$
ROBERTA	5	0.80	5	$0.70_{\pm 0.04}$
XLNET-BI	6	0.59	6	$0.51_{\pm 0.01}$
XLNET-UNI	3	0.48	3	$0.43_{\pm 0.02}$
GPT-2	6	0.45	6	$0.48_{\pm 0.02}$
VERB-ONLY	-	0.43	-	0.43

Table 2: Top F_1 scores on the M&L evaluation set, for probes trained on either W&I-FCE or WIKED-S. Scores for probes trained on WIKED-S are reported as the *mean* ± 1 standard deviation over the five sampled training sets.

4.2 Results

Figure 1 shows F_1 scores for probes trained on W&I-FCE and WIKED-S, and evaluated on the M&L stimuli.¹⁰ For the latter, we plot the mean and standard deviation evaluated over the five sampled training sets. We illustrate the VERB-ONLY baseline score with a grey horizontal line. Table 2 shows layers which obtained the top F_1 score per model.

The figure and table shows ELECTRA encodes the most salient information for SVA error detection, with probes obtaining maximum scores of 0.95 and 0.89 ($\sigma=0.02$) when trained on W&I-

FCE and WIKED-S, respectively. Though this may not be surprising given the *replaced token detection* pre-training objective, it illustrates that probes trained on representations from a model capable of SVA error detection can obtain high performance using both training sets.

We observe a divergence in performance between MLM-probes and ALM-probes; the MLM-probes tend to perform better, obtaining maximum scores between 0.70 and 0.89 F_1 , while the ALM-probes don’t score above 0.59 on either training set. In fact, probes trained on representations from GPT-2 and XLNET-UNI often perform worse than the VERB-ONLY baseline at 0.43 and don’t score above 0.50 F_1 . These results imply that GPT-2 and XLNET-UNI representations do not linearly encode enough information to differentiate between grammatical and ungrammatical verbs in SVA.

The MLM-ALM performance gap could be due to the language model directionality: the two unidirectional models (GPT-2 and XLNET-UNI) do perform the worst, but this fails to account for the performance of XLNET-BI – a bi-directional language model which does not perform much better. It may be that the MLM training objective helps to imbue contextual representations with information useful for detecting SVA errors, but we cannot discount the inclusion of the *replaced token detection* objective, even though it is rarely included. Finally, we note key differences in the pre-training data used by the models: BERT and ELECTRA use the BooksCorpus and English Wikipedia, GPT-2 uses

¹⁰We additionally evaluate probes against the other error types in the W&I dataset and verify that probes only detect SVA errors. Probes trained on either BERT or ELECTRA obtain mean scores of 0.04 ($\sigma=0.04$), verifying that information extracted by the probe is isolated to subject-verb agreement errors.

web-scraped data, and XLNET and ROBERTA use a combination of BooksCorpus, Wikipedia, and web-scraped data.

When we examine performance across layers we see that ELECTRA- and BERT-probes follow a similar trajectory, with performance on par with the baseline from layers 1-5 and higher performance only in layers 8-12. SVA-error information is highest in the final layer for BERT (the token representation) while layers 9 and 10 seem to encode the most useful information for ELECTRA. In contrast, probes trained with representations from ROBERTA peak at layer 5 but have mostly consistent performance until 12, apart from a drop in layer 7.

The results for BERT support those from [Jawahar et al. \(2019\)](#), where they find probes encode elements of syntax in the middle to late layers. [Liu et al. \(2019a\)](#) also find that later layers obtain the best performance for a general GED probe, though they experiment on a single dataset (FCE ([Yanakoudakis et al., 2011](#))) and include all grammatical error types. Recent results from [Lasri et al. \(2022b\)](#) show that removing *number* information from nouns at different layers has a detrimental effect on the number-agreement task up until layer 8. They hypothesise that some transfer of noun-*number* takes place in the previous layers. Our results seem to support this: low probing performance in layers 1-5 (when the noun-*number* has yet to be transferred to the verb), and high probing performance in layers 8-12 after transfer has taken place. Interestingly, performance for ELECTRA takes the same shape, suggesting that the pre-training objective (*replace token detection* versus MLM) may not have an important role in how models encode SVA information, especially given that ROBERTA, trained using masked-language modelling, displays a different pattern across layers. Finally, if GED-probe-performance can be taken as a proxy for noun-*number* transfer, then results for ROBERTA-probes suggest that noun-*number* information is transferred to the target verb earlier, possibly due to the more robust optimization.

Turning to probe performance across training sets, we find that probes trained on W&I-FCE consistently perform better than those trained on WIKED-S, for all models tested apart from GPT-2. This could be due to a domain mismatch, where learner writing may be more similar to the M&L stimuli than data from WIKED-S. Though, at the very least this indicates that information pertaining

to SVA errors is not consistently encoded.

Previous work finds some evidence that BERT’s representations encode knowledge of SVA ([Goldberg, 2019](#); [Jawahar et al., 2019](#); [Newman et al., 2021](#); [Lasri et al., 2022a](#)) but that this knowledge is based on heuristics rather than robust SVA rule learning ([Chaves and Richter, 2021](#); [McCoy et al., 2019](#)). Our results indicate that information encoded in contextual representations extends to the detection of SVA errors in ELECTRA, BERT, and to a lesser extent, ROBERTA. However, we find the encoding is not robust across domains, supporting the heuristic-learning claim. These results illustrate the importance of utilising training and evaluation datasets from disparate domains to evaluate probes.

4.3 Results per syntactic construction

We break down the performance of probes for each syntactic construction in the M&L dataset. Figure 2 illustrates F_1 scores for probes trained on W&I-FCE and WIKED-S using representations from each layer of BERT and ROBERTA.¹¹ The VERB-ONLY baseline is shown as grey horizontal lines. For brevity, we present results on sentences with simple agreement, sentential complements, prepositional phrases, subject relative clauses, and object relative clauses (agreement within and across the clause). Results for the other models and syntactic constructions are included in Appendix E.

The trends across layers observed in Figure 1 are generally consistent within syntactic constructions: probes trained on BERT representations improve in the later layers, while layers 5 and 8-12 seem to be the most salient for probes trained on ROBERTA.

Probes for both models detect SVA errors in the simple agreement constructions – with only the BERT-probe trained on WIKED-S scoring less than 0.90 F_1 . We find BERT-probes perform well for most constructions, especially in layer 12, but performance for ROBERTA-probes drops for sentences with subject-relative clauses, prepositional phrases, and object-relative clauses (agreement within the clause), particularly when trained on WIKED-S data. This suggests that the token representation from BERT models, potentially used in downstream tasks, already encodes a lot of information related to SVA before any fine-tuning.

When comparing performance between probes across training sets, we observe a noticeable differ-

¹¹We select BERT and ROBERTA because they are both MLMs, whereas the ELECTRA-discriminator is trained using a *replaced token detection* objective.

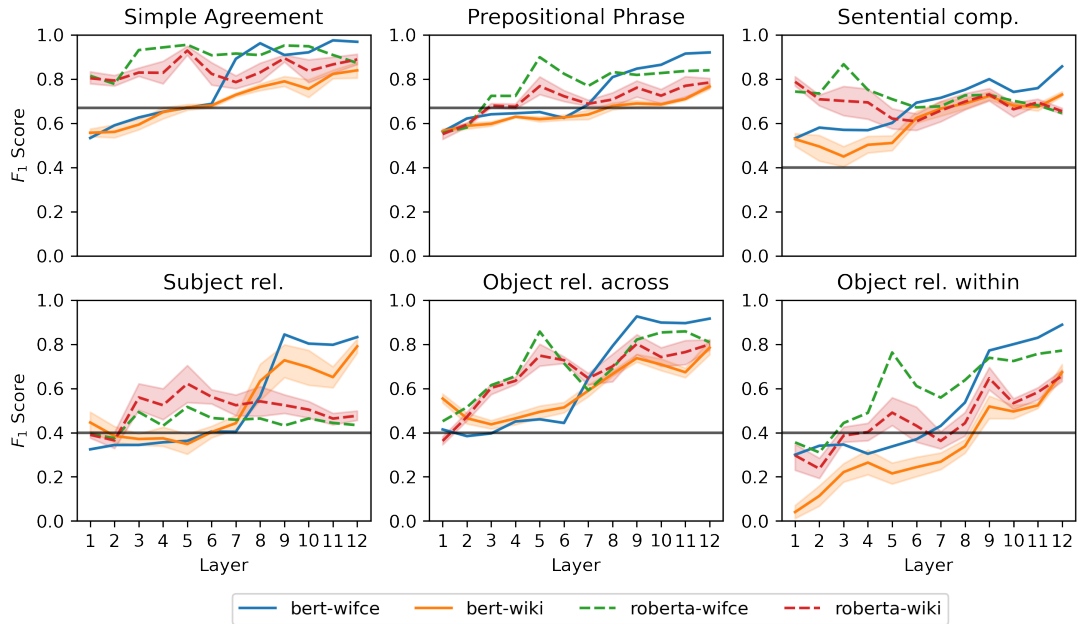


Figure 2: F_1 scores for probes trained on contextual representations from BERT and ROBERTA, using both the W&I-FCE and WIKED-S training sets. The probes are evaluated on M&L stimuli. The VERB-ONLY baseline is illustrated using grey horizontal lines.

ence in performance between BERT-probes trained on W&I-FCE versus those trained on WIKED-S, most evidently in five out of the six syntactic constructions shown. On the other hand, probes trained on ROBERTA don’t always display a performance gap – performance is more comparable between probes on sentences testing simple agreement, sentential complements, and agreement across object relative clauses. For sentences with subject relative clauses we find the probes trained on WIKED-S outperform those trained on W&I-FCE. These results may indicate that information pertaining to the detection of SVA errors is more robustly encoded in ROBERTA than BERT – that is, the information is more invariant to the choice of probe training set.

5 Experiment 2: Generalisation to unseen verbs

In our first experiment, we observe that although the MLMs encode more information for SVA error detection compared to the ALMs, the information does not always generalize across domains or syntactic constructions. We carry out a follow-up experiment to investigate whether the information generalizes across verbs using probes trained on BERT representations and WIKED data. We focus on layers 6 to 12 as these were the layers where performance was above the baseline. There are 13 target verbs in the M&L stimuli, of which “to be” is

the most frequent with 946 occurrences in WIKED-S. The remaining verbs appear very infrequently – for example, eight verbs have frequencies less than 30. To test generalization across verbs we remove all sentences from the training and development sets which contain a verb from the M&L stimuli except for “to be”. We then re-sample sentences from the full WIKED data to maintain 1936 sentences as in the first experiment. Due to the infrequency of many verbs and to ensure a more thorough evaluation, we also increase the training set size by 4- and 8-times to yield training sets with 7744 and 15488 sentences, respectively. We refer to the three sizes as small, medium, and large. This results in paired training sets: for each training set size, there is one set “with M&L verbs” and a set “without M&L verbs”. We sample each training set five times and report the mean and standard deviation over the samples. For example, we sample five training sets with 7744 sentences “with M&L verbs”, and another five “without M&L verbs”. Since we are only interested in the performance of 12 verbs, we modify the evaluation stimuli to remove a) sentences containing only “to be” verbs, and for sentences with multiple verbs we remove the “to be” token from evaluation. For example, in the sentence “The movie the security guards like is good”, we remove “is” from evaluation.

Figure 3 illustrates the F_1 scores: the left and

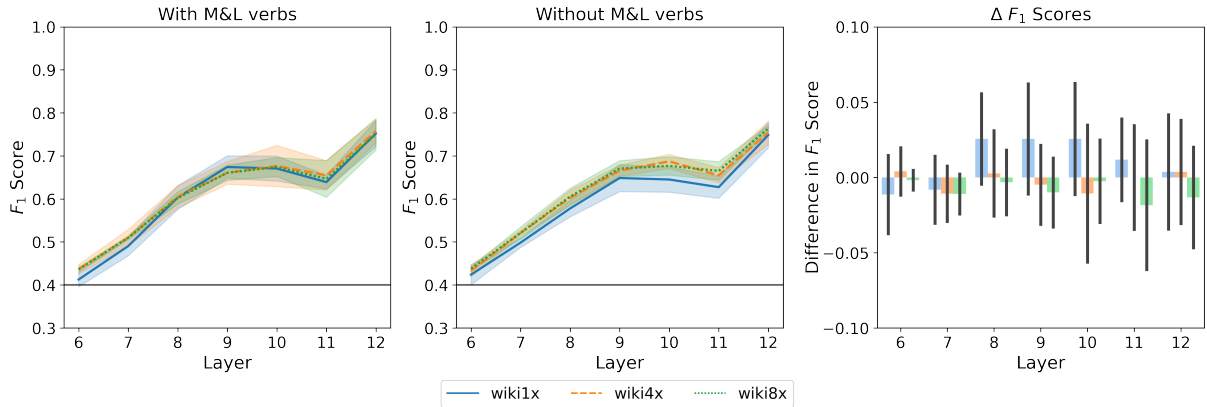


Figure 3: F_1 scores for probes trained on BERT representations, with varying amounts of data (left and centre); $1x=1936$, $4x=7744$, $8x=15488$ sentences. The plots show the mean and standard deviation across five training sets (§5). On the left, probes are trained on data including M&L verbs, while the centre shows scores for probes trained on data without M&L verbs (apart from “is” and “are”, as described in §5). The right-hand plot shows the mean and standard deviation of differences in F_1 scores between probes trained on the two types of dataset: with and without M&L verbs. The VERB-ONLY baseline is shown as a grey horizontal line.

centre plots show results for probes trained on data with and without M&L verbs, respectively. The plot on the right presents the mean and standard deviation for the pairwise differences between probes trained on datasets of the same size. The right-hand plot shows that for the smallest training set size we observe a slight benefit when including the M&L verbs, but this is limited to ~ 0.05 F_1 and restricted to layers 8-10. We generally observe no difference in performance for the medium and large training sets, indicating that SVA-error information does generalize across verbs. Furthermore, we observe no difference when comparing results across training set sizes in the left and central plots, except for probes trained on the small training set without M&L verbs. These results indicate that SVA-error information is linearly accessible and generalizable across across the verbs we test, even when probes are trained with limited data. Future work may expand the investigation to cover more verbs, though we expect performance to deteriorate as verbs become infrequent in the pre-training data (Wei et al., 2021).

6 Discussion

Our experiments test whether information for SVA errors is implicitly encoded in the contextual representations of verbs, but they don’t provide any indication as to *how* the information is encoded: is grammaticality encoded atomically or compositionally? Furthermore, we note that selecting the “erroneous token” can be an ambiguous choice

between the noun and the verb, for example in “The authors laughs”. Yet, the probes we evaluate never tag the nouns. This could indicate that a) the probes learn to only tag verbs, and/or b) that SVA-grammaticality is disparately encoded between nouns and verbs. A compositional account of grammatical encoding is a plausible explanation given the results provided in Lasri et al. (2022b) – that nouns and verbs have different encodings for *number*. We plan to investigate how grammaticality is encoded in future work, both in pre-trained language models as well as models trained specifically for GED.

7 Conclusion

We analyse whether pre-trained transformer-based language models implicitly encode knowledge of SVA errors using GED probes. We carry out a thorough evaluation on five models, using two public training sets from different domains, and evaluate on a manually constructed evaluation set. This enables us to get a more complete and reliable picture of a models’ performance.

Grammatical error detection is a challenging and linguistically aligned task to assess the knowledge of neural language models; we show that GED-probes can be used as a complementary analysis tool to evaluate a models’ linguistic capabilities.

Our results show that ELECTRA, BERT, and ROBERTA encode information for SVA-error detection, but GPT-2 and XLNET do not. For BERT and ROBERTA, we find that the SVA-error encoding is not robust across all syntactic constructions

or training set domains, though we do find some evidence that the encoding generalizes across verbs for BERT. Furthermore, a layer-wise analysis reveals the final layers in ELECTRA and BERT are the most salient for SVA-error detection.

Acknowledgements

We thank Guy Aglionby, Rami Aly, Paula Czarnowska, and Tiago Pimentel for helpful discussions and feedback on early drafts of this work. We also thank the anonymous reviewers for their helpful feedback. This paper reports on research supported by Cambridge University Press & Assessment. We thank the NVIDIA Corporation for the donation of the Titan X Pascal GPU used in this research.

References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. [Fine-grained analysis of sentence embeddings using auxiliary prediction tasks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Guillaume Alain and Yoshua Bengio. 2017. [Understanding intermediate layers using linear classifier probes](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Christopher Bryant, Mariano Felice, Øistein E Andersen, and Ted Briscoe. 2019. The BEA-2019 shared task on grammatical error correction. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75.
- Christopher Bryant, Mariano Felice, and Edward Briscoe. 2017. Automatic annotation and evaluation of error types for grammatical error correction. Association for Computational Linguistics.
- Rui P Chaves and Stephanie N Richter. 2021. Look at that! BERT can be easily distracted from paying attention to morphosyntax. *Proceedings of the Society for Computation in Linguistics*, 4(1):28–38.
- Noam Chomsky. 1957. *Syntactic structures*. Mouton: The Hague, The Netherlands.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators](#). In *International Conference on Learning Representations*.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single \$\&\#\&^*\$ vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Yoav Goldberg. 2019. Assessing BERT’s syntactic abilities. *arXiv preprint arXiv:1901.05287*.
- Roman Grundkiewicz and Marcin Junczys-Dowmunt. 2014. [The WikEd Error Corpus: A corpus of corrective Wikipedia edits and its application to grammatical error correction](#). In *Advances in Natural Language Processing – Lecture Notes in Computer Science*, volume 8686, pages 478–490. Springer.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Dieuwke Hupkes, Sara Veldhoen, and Willem Zuidema. 2018. Visualisation and ‘diagnostic classifiers’ reveal how recurrent and recursive neural networks process hierarchical structure. *Journal of Artificial Intelligence Research*, 61:907–926.
- Ganesh Jawahar, Benoît Sagot, and Djamel Seddah. 2019. [What does BERT learn about the structure of language?](#) In *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy.
- Karim Lasri, Alessandro Lenci, and Thierry Poibeau. 2022a. Does BERT really agree? fine-grained analysis of lexical dependence on a syntactic task. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2309–2315.

- Karim Lasri, Tiago Pimentel, Alessandro Lenci, Thierry Poibeau, and Ryan Cotterell. 2022b. Probing for the usage of grammatical number. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8818–8831.
- 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.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew E Peters, and Noah A Smith. 2019a. Linguistic knowledge and transferability of contextual representations.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448.
- Benjamin Newman, Kai-Siang Ang, Julia Gong, and John Hewitt. 2021. Refining targeted syntactic evaluation of language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3710–3723.
- Joe Pater. 2019. Generative linguistics and neural networks at 60: Foundation, friction, and fusion. *Language*, 95(1):e41–e74.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. Does string-based neural MT learn source syntax? In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 1526–1534.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei, Dan Garrette, Tal Linzen, and Ellie Pavlick. 2021. Frequency effects on syntactic rule learning in transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 932–948.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading ESOL texts. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 180–189, Portland, Oregon, USA. Association for Computational Linguistics.
- Zheng Yuan, Shiva Taslimipour, Christopher Davis, and Christopher Byrant. 2021. Multi-Class Grammatical Error Detection for Correction: A Tale of Two Systems. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP-2021)*.

A Corpus statistics

	Corpus	# sentences	μ sent. length (σ)	μ errors per sentence (σ)
Original	FCE-train	28K	16 (11)	1.9 (2.5)
	W&I-train	34K	18 (12)	2.0 (2.8)
	BEA-dev	4K	20 (12)	1.9 (2.8)
	WikEd (total)	28M	22 (12)	1.6 (1.6)
Processed	FCE-train	626	25 (14)	1.1 (0.3)
	W&I-train	1310	27 (21)	1.1 (0.3)
	BEA-dev	142	26 (19)	1.1 (0.3)
	WikEd-train	1936	24 (12)	1.0 (0.2)
	WikEd-dev	5839	23 (11)	1.0 (0.2)

Table 3: Corpus statistics.

B Marvin & Linzen statistics

Statistics per construction.

Construction	# sentences	μ sent. length (σ)
Simple agr.	280	4.57 (0.49)
In sent. comp.	3360	7.57 (0.49)
Across prep.	44800	8.85 (1.17)
Across subj. rel.	22400	8.77 (0.64)
Short VP coord	1680	7.14 (0.64)
Long VP coord	800	14.40 (0.49)
Across obj. rel.	44800	9.18 (0.86)
Across obj. rel. (no comp)	44800	8.18 (0.86)
Within obj. rel.	44800	9.18 (0.86)
Within obj. rel. (no comp)	44800	8.18 (0.86)

Table 4: Details for the evaluation stimuli from (Marvin and Linzen, 2018).

C Results for probes trained on W&I-FCE

Model	1	2	3	4	5	6	7	8	9	10	11	12
BERT	0.40	0.43	0.43	0.44	0.48	0.50	0.61	0.74	0.84	0.82	0.84	0.89
ELECTRA	0.43	0.46	0.47	0.50	0.50	0.53	0.76	0.88	0.95	0.88	0.90	0.91
ROBERTA	0.51	0.50	0.66	0.67	0.80	0.71	0.65	0.69	0.73	0.74	0.74	0.72
GPT-2	0.37	0.35	0.37	0.37	0.36	0.45	0.43	0.34	0.38	0.31	0.21	0.25
XLNET-BI	0.39	0.44	0.46	0.50	0.59	0.59	0.52	0.48	0.47	0.48	0.45	0.41
XLNET-UNI	0.31	0.43	0.48	0.46	0.44	0.39	0.36	0.35	0.33	0.34	0.29	0.31

Table 5: F_1 scores for probes trained on contextual representations at different layers from BERT, ELECTRA, ROBERTA, XLNET with bidirectional decoding, XLNET with unidirectional decoding, and GPT-2. Probes were trained on learner data described in §3.1, and evaluated on the [Marvin and Linzen \(2018\)](#) stimuli (§3.3).

D Results for probes trained on WIKED-S

Model	1	2	3	4	5	6	7	8	9	10	11	12
BERT	0.39	0.39	0.40	0.43	0.43	0.47	0.53	0.62	0.68	0.65	0.65	0.73
ELECTRA	0.42	0.43	0.45	0.48	0.48	0.49	0.66	0.83	0.87	0.89	0.88	0.84
ROBERTA	0.46	0.48	0.60	0.62	0.70	0.65	0.59	0.63	0.70	0.65	0.68	0.69
GPT-2	0.37	0.39	0.42	0.44	0.45	0.48	0.44	0.41	0.40	0.38	0.37	0.34
XLNET-UNI	0.38	0.41	0.43	0.43	0.42	0.38	0.35	0.34	0.35	0.35	0.33	0.35
XLNET-BI	0.40	0.43	0.45	0.48	0.51	0.51	0.48	0.44	0.43	0.43	0.42	0.43

Table 6: F_1 scores for probes trained on contextual representations at different layers from BERT, ELECTRA, ROBERTA, XLNET with bidirectional decoding, XLNET with unidirectional decoding, and GPT-2. Probes were trained on wikipedia data described in §3.2, and evaluated on the [Marvin and Linzen \(2018\)](#) stimuli (§3.3).

E Results across syntactic constructions

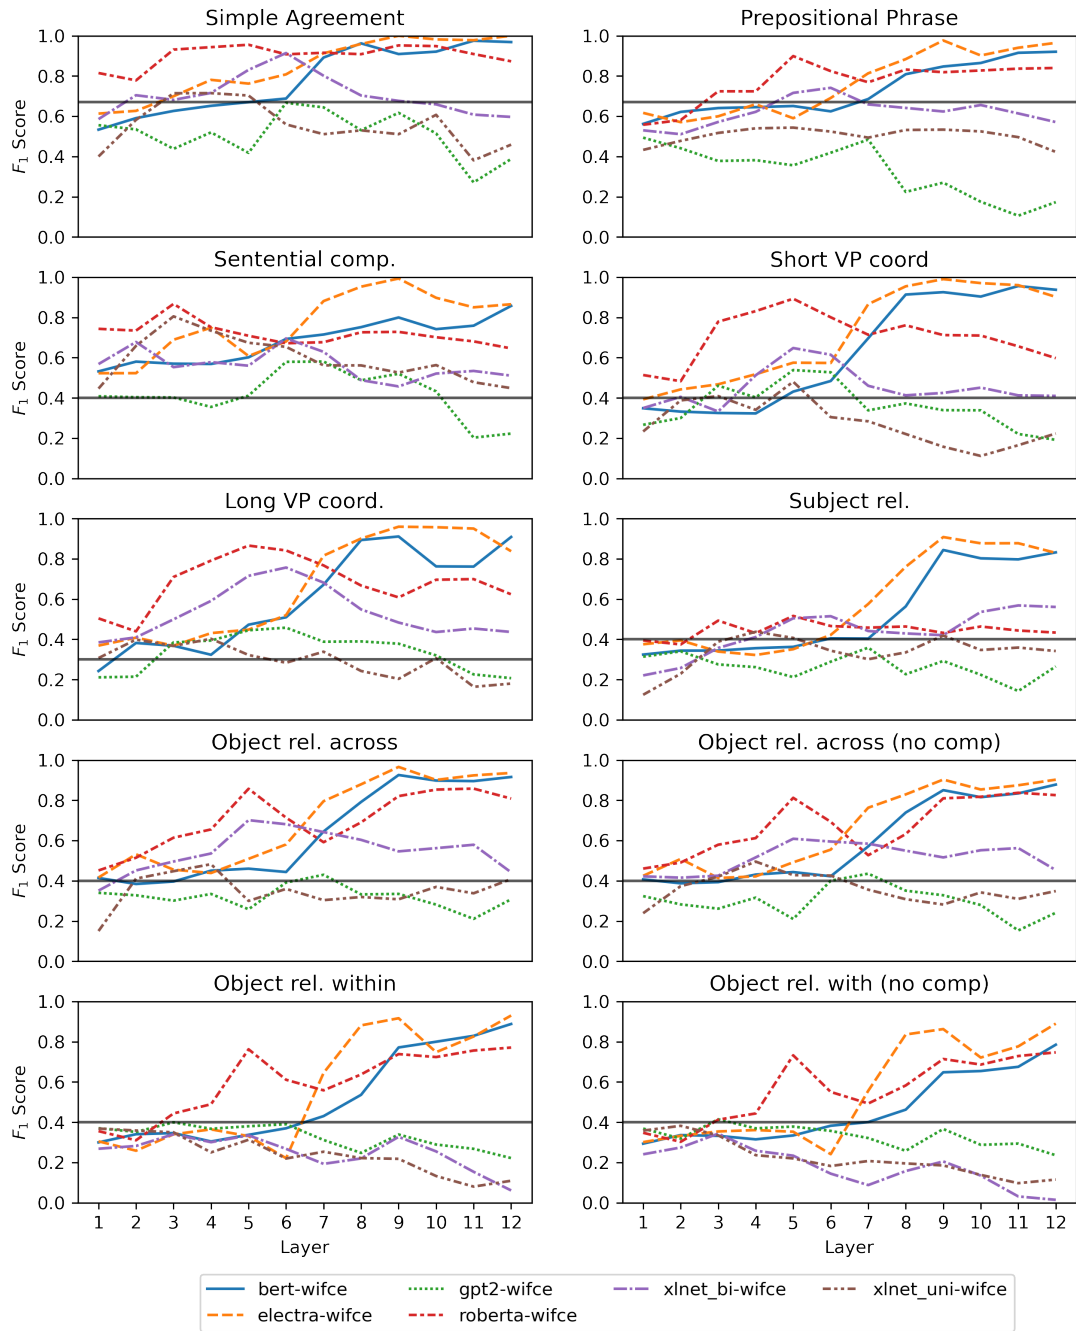


Figure 4: F₁ scores for probes trained the w&i-FCE and training set. The probes are evaluated on M&L stimuli. The VERB-ONLY baseline is illustrated using grey horizontal lines.

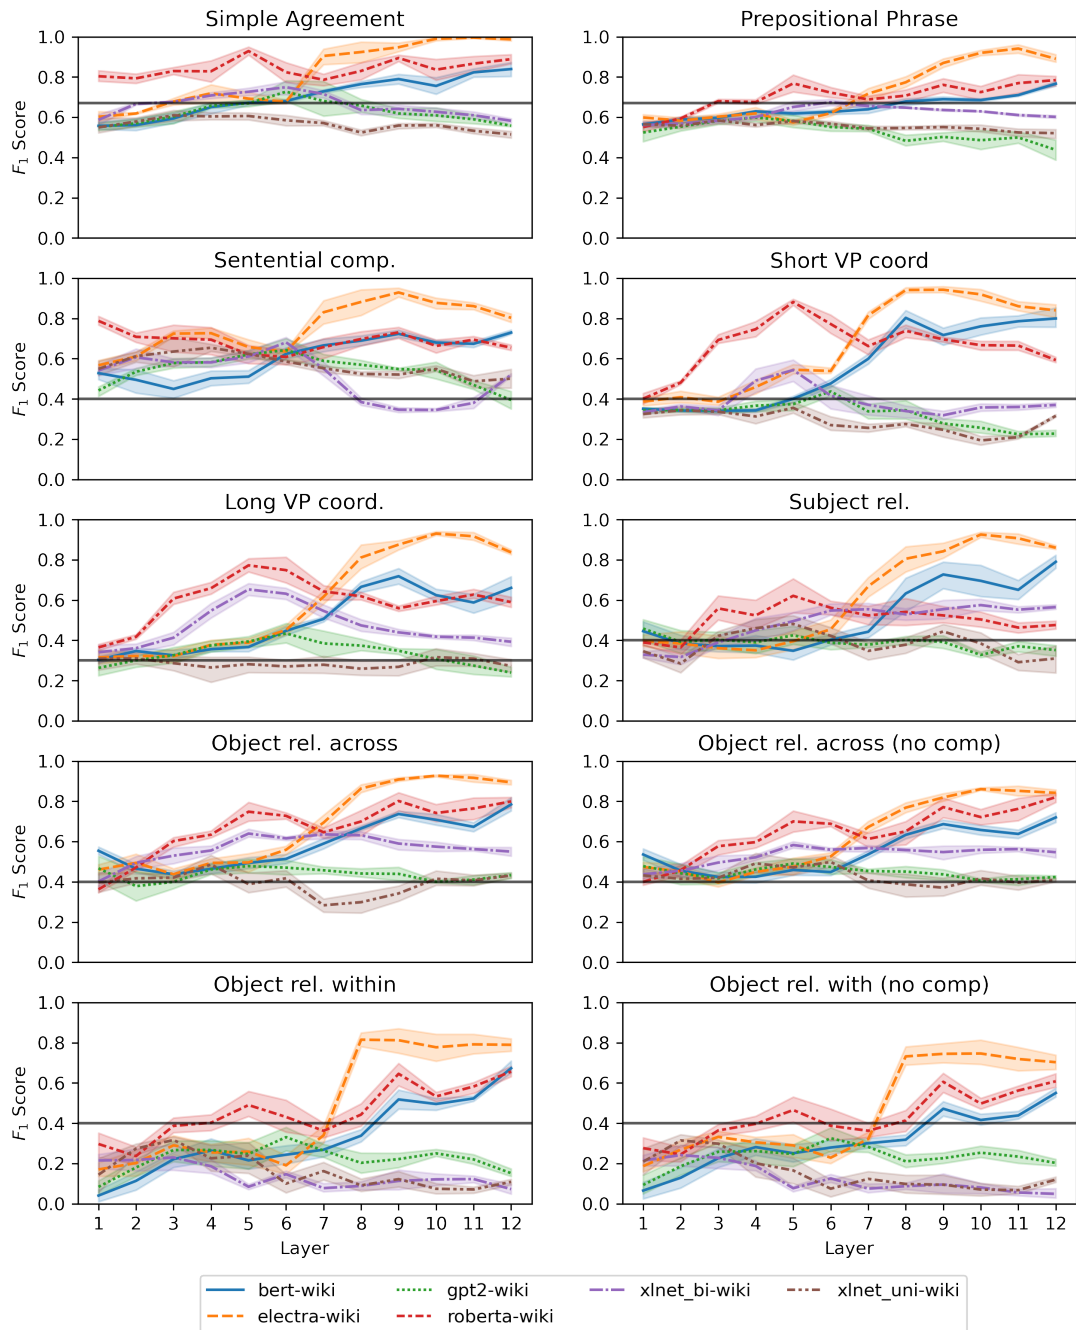


Figure 5: F₁ scores for probes trained the WIKED and training set. The probes are evaluated on M&L stimuli. The VERB-ONLY baseline is illustrated using grey horizontal lines.