

Negation, Coordination, and Quantifiers in Contextualized Language Models

Aikaterini-Lida Kalouli^{1*}, Rita Sevastjanova^{2*}, Christin Beck², and Maribel Romero²

¹ CIS - LMU Munich, kalouli@cis.lmu.de

² University of Konstanz, firstname.lastname@uni.kn

Abstract

With the success of contextualized language models, much research explores what these models really learn and in which cases they still fail. Most of this work focuses on specific NLP tasks and on the learning outcome. Little research has attempted to decouple the models' weaknesses from specific tasks and focus on the embeddings per se and their mode of learning. In this paper, we take up this research opportunity: based on theoretical linguistic insights, we explore whether the semantic constraints of function words are learned and how the surrounding context impacts their embeddings. We create suitable datasets, provide new insights into the inner workings of LMs vis-à-vis function words and implement an assisting visual web interface for qualitative analysis.

1 Introduction

Recently, tremendous progress has been observed in the development of contextualized language models (LM). After the introduction of contextualized embeddings like ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019), earlier static word embeddings (Mikolov et al., 2013; Pennington et al., 2014) have been sidelined, and new standards have been set for the state-of-the-art. Particularly, LMs have been shown to learn task-agnostic properties of language (e.g., Belinkov, 2018; Conneau et al., 2018; Peters et al., 2018) and linguistic properties that imitate the classical NLP pipeline (Tenney et al., 2019a). Despite this success, LMs cannot be taken to understand language the way humans do, as they fail to generalize on unseen data and tasks requiring compositionality (McCoy et al., 2019; Richardson et al., 2020).

Research efforts concentrate on shedding light on what these models learn and consequently how they can be improved. One strand of work relies on probing/diagnostic tasks, where classifiers are

trained on LMs to determine whether they can encode specific linguistic properties, e.g., Marvin and Linzen (2018); Tenney et al. (2019b); Hewitt and Manning (2019); Talmor et al. (2020). Another strand of research focuses on creating adversarial test sets with hard linguistic phenomena to observe where the models fail and thus reverse-engineer them, e.g., McCoy et al. (2019); Richardson et al. (2020); Nie et al. (2020). Both strands of research have approached their goals by employing specific tasks, e.g., Natural Language Inference, Question-Answering, etc., and this has revealed weaknesses of the current models, e.g., that negation is not treated according to its semantic nature. Nevertheless, there is still little work, e.g., Ethayarajh (2019); Sevastjanova et al. (2021), decoupling the weaknesses of LMs from specific NLP tasks and focusing on which of their inner workings are responsible for these weaknesses.

In this paper, we attempt to fill this gap by taking a closer look at function words, i.e., words with little lexical meaning. Function words have been traditionally ignored in NLP, being dismissed as stopwords. Since the rise of contextualized embeddings, function words are not dismissed anymore but do not receive any special treatment either; instead, they are contextually learned like any other word. However, the linguistically-motivated datasets created in this work and tested within the masked-language-modeling (MLM) task show that this context-based learning does not effectively capture the functionality of such words and, thus, that the weaknesses that LMs show in semantic tasks can be attributed to these ill-learned representations. Particularly, our work sheds light to the type and quality of masked-word predictions when **negation**, **coordination**, and **quantifiers** are involved. This work focuses on English masked LMs, but the basic findings should be extendable to other languages since the training process of most masked LMs is relatively similar.

* Equal contribution.

With this, the contributions of this work are four-fold. First, we show the semantic constraints of three of the most studied classes of function words – negation, coordination, and quantifiers. Second, we provide linguistically-motivated datasets that capture the semantic constraints of these function words. Based on these datasets, we offer new quantitative and qualitative insights into the treatment of function words in masked LMs. The qualitative insights are facilitated by an openly-available web interface that visualizes LM predictions and allows researchers to easily test their hypotheses.

2 Related Work

Most related work, i.e., probing studies and adversarial testing has focused on specific NLP tasks and on the learning outcome of the models. However, recently, research has also concentrated on exploring the learning procedure per se and understanding the LM inner workings as per the way of their training. With this, the study of contextualization has emerged. The embeddings are learned based on the co-occurrence of words in particular contexts. Thus, the LM embeddings are contextualized, i.e., a word has different representations based on the context it is found in. Particularly, [Ethayarajh \(2019\)](#) shows that the embeddings become more contextualized, i.e., more context-specific, in the upper layers of BERT, and that contextualization is not entirely driven by polysemy: (non-polysemous) stopwords such as ‘and’, ‘of’, ‘the’ and ‘to’ also become increasingly contextualized in the upper layers, and their representations are among the most context-specific ones. Rather, he finds that contextualization seems to be driven by the variety of contexts a word appears in. These findings are further explored by [Sevastjanova et al. \(2021\)](#), who show that contextualization is neither driven by polysemy nor by pure context variation, but by the combination of functionality, sense variation, syntactic variation, and semantic context variation: BERT can efficiently model polysemous, homonymous and monosemous words, and also words that appear in semantic contexts of high and low variation and independently of their polysemy. But it cannot model words that have a semi-functional/semi-content nature (e.g., modals, quantifiers, temporal adverbials). In this work, we take up this research direction and further explore what is exactly learned during the learning of function words and how their surrounding context influences their learned embeddings.

3 Negation, Coordination, and Quantifiers

3.1 Theoretical Linguistics

Traditionally in theoretical linguistics, there is a distinction between function and content words and several criteria to distinguish the two. The probably most popular criterion is that of semantic content: content words establish a specific semantic content and contribute to the principal meaning of a sentence; function words are rather “non-conceptual” and mainly fulfill some grammatical function, gluing content words together. Other criteria include membership openness (i.e., whether new members can be added to each of the two classes), the flexibility of syntactic attachment (i.e., whether they can combine with any syntactic phrase), and separability from complements (i.e., whether they can be detached from their lexical head) ([Corver and van Riemsdijk, 2001](#)). The distinction between the two classes is not always clear-cut as there are words that share both functional and lexical properties. Thus, it has often been argued that function and content words should rather form a quasi-continuum (‘squishiness’) ([Ross, 1972](#); [Emonds, 1985](#)). For example, prepositions are less functional than articles, e.g., some prepositions are associated with a locative or directional meaning, but they are also more functional than nouns or verbs, e.g., because they are inseparable from their content words.

In this work, we focus on three types of function words, which represent core notions of logic, mathematics, and human reasoning and ones that a state-of-the-art NLP system should efficiently be able to handle: the boolean notions of complement (negation), intersection (conjunction) and union (disjunction), and the notion of existential and universal quantification. Specifically, we study the negation markers *not* and *without*, the coordination markers *or* and *and*, and the existential quantifier *some* and the universal quantifiers *all* and *no*.

Negation Markers *Not* can be considered syntactic negation while *without*, being a preposition with some lexical content, can be considered lexical negation. Semantically, at the heart of negation lies the notion of inconsistency ([Ladusaw, 1996](#)): A given sentence and its negation, as in the episodic pair (1) and the generic pair (2), are inconsistent with each other, i.e., they have disjoint truth conditions (where e.g. MOTHER is short for $\{x : x \text{ is a mother}\}$ and \bar{A} is the complement of A):

- (1) Maria is a mother. $m \in \text{MOTHER}$
 Maria is not a mother. $m \in \overline{\text{MOTHER}}$
- (2) A mom is a mother. $\text{MOM} \subseteq \text{MOTHER}$
 A mom is not a mother. $\text{MOM} \subseteq \overline{\text{MOTHER}}$

Additionally, negation in natural language interacts with its clausemate elements for meaning composition in interesting ways. To name one case, negation is a well-known “hole” for presuppositions, i.e., it can negate the at-issue, propositional content of its clause but not its presupposed content (Karttunen and Peters, 1979). For example, the sentence *Joe isn’t sick with covid again* denies that Joe has covid now (at-issue content), but it does not deny that he had covid before (presupposed content). In this paper, we concentrate on the notion of inconsistency or disjointness and test whether current LMs have grasped this semantic functionality.

Coordination Markers Moving on to the coordination markers *and* and *or*, the former is used to form conjunctions and the latter disjunctions. In its run-of-the-mill boolean use, *and* semantically amounts to the intersection operation: Predicating a conjunctive property of an individual, e.g. Joe, amounts to asserting membership to the intersection of the two conjuncts, as in (3-a). In contrast, though the purely semantic content of *A or B* amounts to set union and thus to inclusive disjunction (meaning ‘A, B or both’), this meaning is typically strengthened to that of *exclusive disjunction* (i.e., ‘A or B and not both’) (Horn, 1972; Sauerland, 2004). This is linked to Hurford’s observation that a disjunction of shape *A or B* is odd if there is an entailment relation between A and B, as defined in (4) and illustrated in (3-b) (Hurford, 1974; Singh, 2008; Ippolito, 2020):

- (3) a. Joe is a dolphin and a mammal.
 $j \in \text{DOLPHIN} \cap \text{MAMMAL}$
 b. #Joe is a dolphin or a mammal.
 $j \in \text{DOLPHIN} \cup \text{MAMMAL} \wedge j \notin \text{DOLPHIN} \cap \text{MAMMAL}$
- (4) Hurford’s Constraint:
 A disjunction of the form *A or B* is infelicitous (i.e., #) if A entails B or vice-versa.

This means that *and* and *or* have different semantic signatures when it comes to the relation between their coordinated terms: *A and B* allows for A to entail B, whereas *A or B* prohibits it. We will test whether LMs are able to detect and learn these

different meaning signatures.

There are other interesting semantic properties of coordination markers in natural language that could be explored. To name just one, *and* has, in addition to its boolean meaning, a non-boolean meaning that roughly amounts to the creation of a new complex plural individual (Link, 2002). For example, (5) allows not only for the boolean, distributive reading (5-a) and also for the non-boolean, collective reading (5-b).

- (5) Jane and Paul built a castle.
 a. ‘Jane built a castle and Paul built a castle.’
 b. ‘Jane and Paul together built a castle.’

Quantifiers Semantically, quantifiers like *all/every*, *some* and *no* denote relations between the set P coming from the noun phrase headed by the quantifier and the set Q coming from the rest of the sentence (Barwise and Cooper, 1981):

- (6) a. $\llbracket \text{All Ps are Q} \rrbracket = 1$ iff $P \subseteq Q$
 b. $\llbracket \text{Some Ps are Q} \rrbracket = 1$ iff $P \cap Q \neq \emptyset$
 c. $\llbracket \text{No Ps are Q} \rrbracket = 1$ iff $P \cap Q = \emptyset$

Two points follow from these lexical entries. First, *all* and *no* are polar opposites (i.e., they are contrary –though not complementary– items (Ladusaw, 1996; Cruse, 2011)). This means that the statements *All Ps are Q* and *No Ps are Q* are inconsistent with each other. Second, the statement *All Ps are Q* technically entails the statement *Some Ps are Q* (under the assumption that set P is non-empty). However, the two quantifiers form part of the lexical scale $\langle \text{some, several, ... , many, most, all} \rangle$, where the leftmost item is the weakest, and the rightmost item is the strongest (Horn, 1972). As it typically happens with scales in natural language, pragmatic implicatures are routinely run on the weaker terms, leading to readings not entailed anymore by the stronger terms. In the case of *some*, its original weaker reading ‘some and possibly all’ in (6-b) is upgraded into the stronger pragmatic reading ‘some but not all’ in (7). Under this second reading, the statement *All Ps are Q* does no longer entail the statement *Some Ps are Q*; rather, the two are incompatible with each other.

- (7) $\llbracket \text{Some}_{str} \text{Ps are Q} \rrbracket = 1$ iff
 $P \cap Q \neq \emptyset$ and $P \not\subseteq Q$

For an LM understanding the meaning of the quantifiers *all*, *some*, and *no* and for a fixed value of P, the Q-embeddings (i.e., embeddings of the set Q) of the noun phrase *All Ps* should be radically different from the Q-embeddings of *No Ps*, since they are polar opposites (i.e., contraries). Similarly, the Q-embeddings of the noun phrase *All Ps* should differ substantially from the Q-embeddings of the pragmatically strengthened *Some Ps*, since they are incompatibles. In other words, in both cases, the Q-embedding of the quantificational noun phrases under comparison should be disjoint. We will test whether LMs show traces of semantic understanding in being able to create such disjoint predictions.

3.2 Computational Linguistics/NLP

Within computational linguistics, function words have mostly been treated as *stopwords*. The term was coined by Luhn (1960) to mean very common words that do not add much to the meaning of a text but ensure the structure of a sentence is sound. Historically, one of the main reasons for removing stopwords was to decrease the computational time for text mining (Huston and Croft, 2010) and help search engines to give better results. Nowadays, this reason for removing stopwords is not valid anymore since we have computationally more powerful hardware; nevertheless, various NLP tasks such as topic modeling and information retrieval, continue to use the practice as it is often argued to improve performance (e.g., Fan et al. (2019); Sarica and Luo (2020)).

With the rise of distributional semantics and word embeddings, more research focused on the nature of these words and the special treatment required (e.g., (Bernardi et al., 2013; Hermann et al., 2013; Linzen et al., 2016; Tang et al., 2016)). More recently, the special nature of function words has also interested the LM community. In one of the probing tasks proposed by Kassner and Schütze (2020) it is shown that LMs cannot distinguish between negated (“Birds cannot [MASK]”) and non-negated (“Birds can [MASK]”) questions. The researchers insert the negation *not* into the LAMA dataset (Petroni et al., 2019) and create positive and negative cloze questions. By comparing the predictions of LMs on these pairs, they find that the positive and negative predicted fillers have high overlap and correlation, i.e., models are equally likely to generate true (e.g. *birds can fly*) and incorrect statements (e.g., *birds cannot fly*). High

correlation means that the models do not understand negation; correct answers for positive and negative questions are expected to be disjoint. Interestingly, the researchers show that BERT can learn both positive and negative facts correctly if they occur in training but still fails to generalize to unseen (positive and negative) sentences.

4 Experiments

In this work we focus on the three functional categories discussed and set out to complement the existing research on the treatment of function words within contextualized LMs. First, we create datasets that can be used to evaluate how LMs have learned to capture the semantic constraints of these functional categories. Second, we generate quantitative insights into the contextualization of function words, which we qualitatively investigate through a user-friendly web interface, which also allows researchers to evaluate their own datasets.

4.1 Data

Motivated by the experiment performed by Kassner and Schütze (2020) on the negated LAMA dataset and the theoretical linguistic research about the type of data needed for this kind of exploration, we create two datasets of different types.¹ For the creation of the datasets, we combine existing resources: the family and location relations of the analogy dataset used by Mikolov et al. (2013), the English occupations dataset of the European Skills and Competences, Qualifications and Occupations (ESCO) initiative² and 5 common-sense relations found in ConceptNet (Speer et al., 2016) (*has-a*, *used-for*, *is-a*, *has-property*, *capable-of*). Based on these resources, we create templates which are then used to produce the sentences of our datasets. Examples for both datasets can be found in the Appendix. More details are given in the following.

Inconsistent Dataset The Inconsistent Dataset aims at revealing whether the predictions in sentences with inconsistent meaning are indeed disjoint. Disjointness is factored in through inconsistent function words, as described in Section 3.1. Parts of the inconsistent sentences are masked and the goal is to detect overlapping predictions within each pair. An overlap means that the LM has not

¹Can be found under: <https://function-words.lingvis.io/>

²ec.europa.eu/esco/portal/download

learned the functionality of the inconsistent functional words. The Inconsistent Dataset consists of 1272 pairs and can be split into three subsets: the negation subset (534), the coordination subset (500), and the quantifiers subset (238).

The negation subset contains pairs with positive and negative sentences, whereby the negation can be syntactic or lexical. Concerning syntactic negation, the subset contains examples with the copula verb *be* and its negation, and transitive verbs and their negation with *not* and *no*. As far as lexical negation is concerned, inconsistent sentences are created with the copula verb *be* and transitive verbs and based on the preposition *with-without*. The coordination subset contains a single type of inconsistency by using the markers *or* and *and* and masking the second part of the coordination. The quantifiers subset has two types of opposing contexts: the first one is formulated through *all-no* and the second one through *all-some*. Both transitive and intransitive sentences are considered.

Semantic Dataset Although the Disjoint Dataset goes beyond the negated LAMA in that it contains an additional type of negation (lexical negation) and two further types of inconsistency (coordination and quantifiers), it can indicate whether the LMs learn the semantic constraints of these functional categories only *indirectly*, i.e., by capturing the predictions’ overlap. Thus, a different “semantic” dataset is required, on which we can directly judge the correctness of the predictions. This can be made clearer with an example such as *A mother is not a [MASK]*: based on the semantic constraints of negation discussed in Section 3.1 and the generic reading ‘For any x , if x is a mom, then x is not a MASK’, predicted masks should not be one of the words *mom, grandmother, grandma, granddaughter, bride, wife, woman, niece, stepmother, daughter, aunt, etc.* Another example is the sentence *John was born in Berlin or in [MASK]*, where the predicted masks should not contain the word *Germany*; again, see constraints in Section 3.1. We call these words *forbidden* although they could be valid sentences in some (e.g., figurative) contexts and could theoretically have been seen during LM training. However, since these sentences are logically incorrect and do not represent the prototypical concepts of things, they should not be among the most likely predictions of the LM – since LMs learn based on the occurrences of things, literal, prototypical meanings should have been encountered more of-

ten. Thus, if the most likely predictions contain the forbidden words, it means that the model has neither learned the functionality of the corresponding operator nor any proper world knowledge, e.g., that a *mother* is the same as a *mom* or that *Berlin* is part of *Germany*.

The creation of such a dataset is particularly challenging because the examples need to be chosen in a way that the necessary semantics is captured within the words themselves and not based on the co-occurrence of the words. Only in this way can we reliably evaluate whether LMs have learned something about the actual semantic constraints of the functional markers or they simply reproduce common concordances of the training data. This means that an example such as *John lives in Berlin or in [MASK]* is not suitable because any predicted word could be right. Thus, to produce such examples, we systematically create examples related to concepts for which we can selectively define invalid predictions according to the particular concept’s characteristics. We choose family relations and occupations for syntactic negation, animals and their main body parts and activities for lexical negation, and capital countries and animals and their hypernyms for coordination. For the selected quantifiers, no “semantic” examples with forbidden answers could be created because of the very nature of quantifiers, i.e., even the universal *all* is not strong enough to create logical invalid examples in the real world. The Semantic Dataset contains a total of 2780 sentences with 187 examples of syntactic and 123 examples of lexical negation and 2470 examples of coordination. The size of the dataset is comparable to existing datasets used in other related experiments (e.g., the negated LAMA dataset by Kassner and Schütze (2020)).

4.2 Models

For our study, we use the *huggingface* (Wolf et al., 2020) implementation of three pretrained LMs, BERT (Devlin et al., 2019)³, RoBERTa (Liu et al., 2019)⁴, and ALBERT (Lan et al., 2019)⁵, which have been shown to achieve state-of-the-art results on the GLUE, RACE, and SQUAD benchmarks. Our two datasets are input to each of the three models to extract the masked predictions, layerwise attentions, and the word embeddings of all words

³<https://huggingface.co/bert-base-uncased>

⁴<https://huggingface.co/roberta-base>

⁵<https://huggingface.co/albert-base-v2>

of each sentence. The embeddings are taken from layer 11, as the higher layers of models like BERT have been shown to mostly capture semantic properties, while the last layer has been found to be very close to the actual classification task and thus to be less suitable (Jawahar et al., 2019; Lin et al., 2019). Word piece embeddings are merged to their corresponding word embeddings through averaging. Based on the predictions, the word embeddings, and attentions, we calculate the following:

- the cosine similarity between the embedding vector of the predicted word to each other word of the sentence
- the layerwise average attention of the predicted word to each other word of the sentence
- for the Inconsistent Dataset: the overlap of the predictions between the two inconsistent versions of the sentences; overlap in the first returned prediction (`overlap@1`), in the first 5 (`overlap@5`), and first 10 (`overlap@10`) returned predictions
- for the Semantic Dataset: the percentage of examples containing at least one forbidden word within the 1st prediction (`forbidden@1`), the first 5 (`forbidden@5`), and the first 10 predictions (`forbidden@10`).

4.3 Web Interface

In addition to the quantitative results, i.e., prediction overlaps and forbidden predictions, we implement a web interface⁶ that visually shows the predictions, the layerwise average attentions, and the word similarities. The interface provides qualitative insights into similarity patterns that are common for different prediction outcomes.

In the interface, the user can select one of the three models (i.e., BERT, ALBERT, or RoBERTa) to explore its predictions. Sentences belonging to one dataset are grouped together, whereby disjoint sentence pairs from the Inconsistent Dataset are placed underneath each other for better comparability. The visualization of predictions consists of multiple elements, e.g., see Figure 1. Each prediction of a masked word is displayed as a row in the visualization (i.e., 10 rows for 10 predictions). On the left, we display the prediction’s probability through a horizontal bar representation. Next to the probability, we display the sentence tokens

⁶<https://function-words.lingvis.io/>

visualized as colored rectangles. The color of the rectangle represents its cosine similarity or its layerwise average attention to the predicted word of the particular sentence. The darker the color, the higher the similarity/attention. Note that ALBERT and RoBERTa generally show higher cosine similarities between the predicted and the surrounding context words than BERT, i.e., they are displayed with darker colors on the visualized figures. On the right, we display the predicted word. To support the analysis of prediction overlaps as well as prediction of forbidden words, we color the words that overlap or are forbidden, respectively, in red color.

5 Results and Discussion

The results of our experiments can be found in Tables 1 and 2. Table 1 shows the percentages of overlap of the inconsistent predictions within the first @1, @5, and @10 predictions. Table 2 shows the percentages of at least one forbidden word being included within the first @1, @5, and @10 predictions of the models.

model	Incons. Dataset	@1	@5	@10
BERT	coord	50	53	57
	neg	46	46	47
	quant	27	38	40
	all	41	45	48
ALBERT	coord	61	80	81
	neg	42	42	42
	quant	27	28	31
	all	43	50	51
RoBERTa	coord	40	58	63
	neg	38	40	41
	quant	25	29	31
	all	34	42	45

Table 1: Inconsistent Dataset: percentage of inconsistent pairs with overlapping predictions within the first x predictions.

model	Semantic Dataset	@1	@5	@10
BERT	synNeg	41	56	60
	lexNeg	51	72	73
	coord	32	73	87
	all	41	67	73
ALBERT	synNeg	14	51	60
	lexNeg	38	63	66
	coord	6	31	47
	all	19	53	58
RoBERTa	synNeg	48	64	71
	lexNeg	21	31	42
	coord	27	68	82
	all	32	54	65

Table 2: Semantic Dataset: percentage of sentences in which there is at least one forbidden prediction within the first x predictions.

5.1 Findings

In the following, we describe findings made using quantitative and qualitative evaluation methods.

Same predictions for inconsistent pairs: We can reproduce and extend the findings by Kassner and Schütze (2020). The first prediction (overlap@1) of disjoint pairs overlaps in 41% of the cases for BERT, 43% for ALBERT, and 34% for RoBERTa. The first 5 predictions (overlap@5) overlap in 45% for BERT, 50% for ALBERT, and 42% for RoBERTa, while the first 10 predictions (overlap@10) overlap in 48% for BERT, 51% for ALBERT, and 45% for RoBERTa. Specifically, overlap@10 is worse in the coordination subset in all 3 models and best in the quantifiers subset. This overlap shows that the models neither learn nor consider the functional nature of these markers.

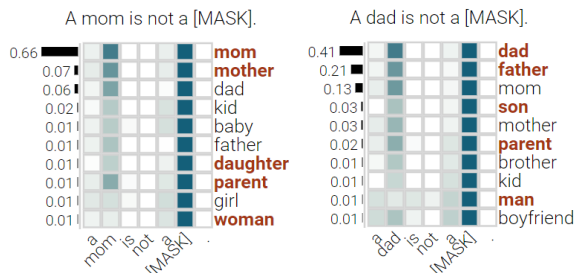


Figure 1: BERT: Similarity visualization of the first 10 predictions of the negation sentences *A mom is not a [MASK]* - *A dad is not a [MASK]*.

Forbidden predictions for semantic pairs: forbidden@1 lies at 41% for BERT, 19% for ALBERT, and 32% for RoBERTa. forbidden@5 is at 67% for BERT, 53% for ALBERT, and 54% for RoBERTa, while forbidden@10 lies at 73% for BERT, 58% for ALBERT, and 65% for RoBERTa. For BERT and RoBERTa the coordination subset seems to be the hardest. The easiest for BERT is the syntactic negation, and for RoBERTa – the lexical negation. In contrast, for ALBERT the easiest is the coordination subset, and the hardest is the lexical negation subset. These results might create the impression that BERT has the most difficulty in this task and that newer models such as ALBERT perform better. However, the further findings and the interpretation following shall shed light on this preliminary impression. Also, considering the embarrassingly easy examples included in our datasets, the amount of forbidden predictions is alarming for all datasets (see Figure 1, Figure 2⁷).

⁷Note that the sentence *The painter does not paint* is not an impossible sentence, and thus an LM might have seen it.




Figure 2: BERT: Similarity visualization of the first 10 predictions of the negation sentences *A singer does not [MASK]* - *A painter does not [MASK]*.

[ALBERT] Similar predictions independently from context: For negation and quantifiers, RoBERTa and ALBERT perform better than BERT. However, if we examine the examples more closely, we see that ALBERT predicts similar words no matter the exact sentence (see Figure 3). For the Inconsistent Dataset, all positive sentences have similar predictions to each other and all negative ones as well. It is doubtful that the training data contained such concordances and this raises the question if this model is contextualized in the same way BERT is.

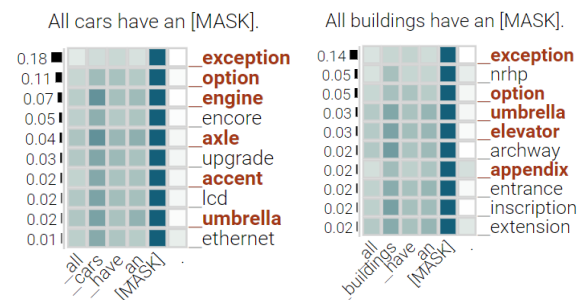


Figure 3: ALBERT: Similarity visualization of the first 10 predictions of the sentences *All cars have an [MASK]* - *All buildings have an [MASK]*.

[ALBERT, RoBERTa] Predictions correlate only with the predicate: For the lexical negation subset of the Semantic Dataset, ALBERT’s predictions are not based on the entire context, i.e., the predicate, the subject, and any function words, but rather correlate only with the predicate of the sentences. For example, all sentences containing the verb *fly* lead to the same predictions, no matter whether it is a *fly*, an *owl* or a *bird* flying (see Figure 4).

Similar observations can be made for RoBERTa.

However, as noted in Section 4.1, LMs should have mainly learned prototypical notions, and thus such predictions should not occur within the 10 most probable ones.




Figure 4: ALBERT: Similarity visualization of the first 10 predictions of the lexical negation sentences *A fly flies without [MASK]* - *A bird flies without [MASK]*.

It has been shown that LMs can learn stereotypical associations reasonably well, e.g., that *walk* is related to *shoes* as in the given example in Figure 5. These associations suggest that LMs are capable of learning commonsense reasoning – knowledge accepted by the majority of people, e.g., how the world works (Bhargava and Ng, 2022). Although the given examples suggest that the models are capable to learn such associations, the visualizations reveal that the predictions are stronger related to the sentence’s predicate than the subject. The same goes for the predicates such as *swim* or *see*.

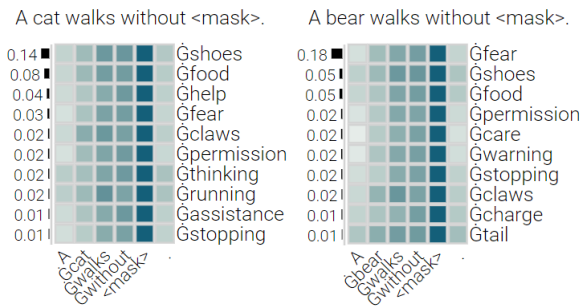


Figure 5: RoBERTa: Similarity visualization of the first 10 predictions of the lexical negation sentences *A cat walks without <mask>* - *A bear walks without <mask>*.

[BERT] Prediction quality dependent on specific named entities: Taking a closer look at the coordination partition of the Semantic Dataset and the location examples specifically, we observe a high similarity between the named entity of the location already contained in the sentence and the predicted word, which is also a named entity in 100% of the cases (see Figure 6). This not only indicates that the presence of named entities has a strong influence on the learning outcome because the remaining context is ignored, but also that the specific named entities have an impact on the predictions’ quality. Particularly, some country/state-capital combinations lead to more forbidden predic-

tions than others. For example, the US states *Texas*, *California*, *Arizona*, *Florida* are more often within the *forbidden@1* than states such as *Indiana*, *Tennessee*, *Minnesota*. This might suggest that the training data of BERT contained more instances of the former combinations, and the model learned a strong relation between these named entity pairs, ignoring any other (functional or lexical) words.




Figure 6: BERT: Similarity visualization of the first 10 predictions of the sentences *Tina comes from Bern or [MASK]* - *Peter was born in Memphis or in [MASK]*.

5.2 Insights and Interpretation

In the following, we describe some potential reasons why current LMs fail in making linguistically correct word predictions.

Function words ignored in semantically related contexts: We find that if the masked word has some semantic relation with the other main concept of the sentence – most often the subject, then the functional word embeddings have low similarity to the masked word. In contrast, if no semantic relation can be established, there is a higher similarity to the function words. For example, in Figure 7, we can see that all predictions that are related to (animals’) body parts and have some semantic relation with the subject *insects* do not exhibit any cosine similarity to the quantifiers, i.e., the first cell of the matrix is white, while for irrelevant predictions such as *nothing*, *died*, *eaten* there is a similarity between the predicted word and the quantifier. This indicates that in many cases the predictions are dominated by semantically rich words and the model ignores other functional operators.

Contextualization of function words: The visualizations show that if all words of the sentence have similar (high or low) similarities to the predicted word, then the predictions are mostly neither inconsistent nor forbidden – even if they do not belong to the average common sense. This suggests that if the masked word is predicted considering the whole context, it should have a similarity to

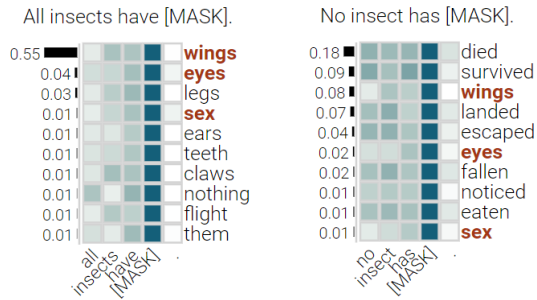


Figure 7: BERT: Similarity visualization of the first 10 predictions of the disjoint pair *All insects have [MASK]* - *No insect has [MASK]*.

all words in the context. Related work supports this assumption (Ethayarajh, 2019; Sevastjanova et al., 2021): function words have high similarity to the rest of the words in higher layers because they are highly contextualized, i.e., they become very context specific. However, our findings show that this mode of learning leads to the negligence of function words: in some cases, the contextualization of function words is indeed high, i.e., their similarity to the predicted word is high, and then no inconsistent predictions arise. But when the sentence contains words that are semantically related, the model gets distracted, the function words are glossed over and their similarity to the predicted word drops. In consequence, their semantic nature is not captured.

Attention confirms insights: Visualizing attention can show whether the observed behavior is specific to cosine similarity or can also be retraced in attention patterns. Indeed, we observe that throughout the layers the first word receives the most attention no matter its part-of-speech. Function words, such as *not* and *or* in Figure 8, receive higher attention only in the middle layers, which have been shown to mainly capture syntactic properties (Jawahar et al., 2019). Thus, it seems that the semantic constraints of these words, which should be captured in the higher layers, are indeed left unaccounted for.

Models are sensitive to minor input changes: While extracting predictions for RoBERTa model, we noticed its high sensitivity to the provided input. Even additional space(s) between tokens in the input context can produce different prediction probabilities, ranks, or even different predictions. Thus, we want to sensitize researchers to be careful when working on similar experiments in order to avoid such potential, undesired errors.




Figure 8: BERT: Attention visualization of 2 predictions of the sentences *A mother is not a [MASK]* and *Joe comes from Athens or [MASK]*. For each prediction, rows represent the layerwise average attention from the predicted word to the words (columns) in the sentence.

6 Conclusion

This paper presented the extent to which LMs learn semantic constraints of function words. Based on theoretical linguistic literature, we created new datasets for testing three functional classes, and showed that popular masked LMs make problematic predictions. The visualizations in the developed web interface highlighted potential reasons for this poor performance. In this work, we set out to shed light on a subset of function word categories. Future work shall look into other function word categories that ought to be challenging for LMs and uncover additional reasons for the poor model performances.

Acknowledgments

We thank the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for funding within the project BU 1806/10-2 “Questions Visualized” of the FOR2111 and project D02 “Evaluation Metrics for Visual Analytics in Linguistics” (Project ID: 251654672 – TRR 161).

References

- Jon Barwise and Robin Cooper. 1981. [Generalized Quantifiers and Natural Language](#). In J. Kulas, J.H. Fetzer, and T.L. Rankin, editors, *Philosophy, Language, and Artificial Intelligence. Studies in Cognitive Systems*, pages 241–301. Springer, Dordrecht.
- Yonatan Belinkov. 2018. *On internal language representations in deep learning: An analysis of machine translation and speech recognition*. Ph.D. thesis, Massachusetts Institute of Technology.
- Raffaella Bernardi, Georgiana Dinu, Marco Marelli, and Marco Baroni. 2013. [A relatedness benchmark](#)

- to test the role of determiners in compositional distributional semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 53–57, Sofia, Bulgaria. Association for Computational Linguistics.
- Prajwal Bhargava and Vincent Ng. 2022. **Commonsense Knowledge Reasoning and Generation with Pre-trained Language Models: A Survey**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(11):12317–12325.
- 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.
- Norbert Corver and Henk van Riemsdijk. 2001. *Semilexical Categories: The Function of Content Words and the Content of Function Words*. De Gruyter Mouton, Berlin, New York.
- Alan Cruse. 2011. *Meaning in language: An introduction to semantics and pragmatics*. Oxford University Press, Oxford.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Joseph E. Emonds. 1985. *A Unified Theory of Syntactic Categories*. De Gruyter Mouton, Berlin, Boston.
- Kawin Ethayarajh. 2019. **How Contextual are Contextualized Word Representations? Comparing the Geometry of BERT, ELMo, and GPT-2 Embeddings**. 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 55–65, Hong Kong, China. Association for Computational Linguistics.
- Angela Fan, Finale Doshi-Velez, and Luke Miratrix. 2019. **Assessing topic model relevance: Evaluation and informative priors**. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 12(3):210–222.
- Karl Moritz Hermann, Edward Grefenstette, and Phil Blunsom. 2013. **“not not bad” is not “bad”: A distributional account of negation**. In *Proceedings of the Workshop on Continuous Vector Space Models and their Compositionality*, pages 74–82, Sofia, Bulgaria. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. **A Structural Probe for Finding Syntax in Word Representations**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Larry Horn. 1972. *On the Semantic Properties of Logical Operators in English*. Ph.D. thesis, UCLA.
- James R. Hurford. 1974. Exclusive or inclusive disjunction. *Foundations of Language*, pages 409–411.
- Samuel Huston and W. Bruce Croft. 2010. **Evaluating verbose query processing techniques**. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10*, page 291–298, New York, NY, USA. Association for Computing Machinery.
- Michela Ippolito. 2020. Varieties of Sobel sequences. *Linguistics and Philosophy*, pages 633–671.
- Ganesh Jawahar, Benoît Sagot, and Djamel Seddah. 2019. **What Does BERT Learn about the Structure of Language?** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Lauri Karttunen and Stanley Peters. 1979. Conventional implicature. In C.K. Oh and D. Dineen, editors, *Pre-supposition*, pages 1–56. Academic Press, New York.
- Nora Kassner and Hinrich Schütze. 2020. **Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.
- William A. Ladusaw. 1996. Negation and polarity items. In S. Lappin, editor, *The Handbook of Contemporary Linguistics Theory*, pages 321–341. Blackwell, Oxford.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. **ALBERT: A lite BERT for self-supervised learning of language representations**. *CoRR*, abs/1909.11942.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. **Open Sesame: Getting inside BERT’s Linguistic Knowledge**. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- Godehard Link. 2002. **The logical analysis of plurals and mass terms: A lattice-theoretical approach**. In P. Portner and B. Partee, editors, *Formal semantics: The essential readings*, pages 127–146. Blackwell, Oxford / Malden, MA.

- Tal Linzen, Emmanuel Dupoux, and Benjamin Spector. 2016. [Quantificational features in distributional word representations](#). In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 1–11, Berlin, Germany. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- H. P. Luhn. 1960. [Key word-in-context index for technical literature \(kwic index\)](#). *American Documentation*, 11(4):288–295.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted Syntactic Evaluation of Language Models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *Proceedings of Workshop at ICLR*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global Vectors for Word Representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Matthew 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.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Kyle Richardson, Hai Hu, Lawrence S. Moss, and Ashish Sabharwal. 2020. [Probing Natural Language Inference Models through Semantic Fragments](#). In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 8713–8721. AAAI Press.
- John R Ross. 1972. The category squish: Endstation hauptwort. In *Chicago Linguistic Society*, volume 8, pages 316–328.
- Serhad Sarica and Jianxi Luo. 2020. [Stopwords in technical language processing](#). *CoRR*, abs/2006.02633.
- Uli Sauerland. 2004. [Scalar Implicatures in Complex Sentences](#). *Linguistics and Philosophy*, pages 367–391.
- Rita Sevastjanova, Aikaterini-Lida Kalouli, Christin Beck, Hanna Schäfer, and Mennatallah El-Assady. 2021. [Explaining Contextualization in Language Models using Visual Analytics](#). 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 464–476, Online. Association for Computational Linguistics.
- Raj Singh. 2008. [On the interpretation of disjunction: Asymmetric, incremental, and eager for inconsistency](#). *Linguistics and Philosophy*, pages 245–260.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2016. [Conceptnet 5.5: An open multilingual graph of general knowledge](#). *CoRR*, abs/1612.03975.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2020. [oLMpics-on what language model pre-training captures](#). *Transactions of the Association for Computational Linguistics*, 8:743–758.
- Gongbo Tang, Gaoqi Rao, Dong Yu, and Endong Xun. 2016. [Can we neglect function words in word embedding?](#) In *Natural Language Understanding and Intelligent Applications - 5th CCF Conference on Natural Language Processing and Chinese Computing, NLPCC 2016, and 24th International Conference on Computer Processing of Oriental Languages, IC-CPOL 2016, Kunming, China, December 2-6, 2016, Proceedings*, volume 10102 of *Lecture Notes in Computer Science*, pages 541–548. Springer.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. [BERT Rediscovered the Classical NLP Pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.

Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. *What do you learn from context? Probing for sentence structure in contextualized word representations*. In *International Conference on Learning Representations*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

A Appendix A: Sample of Datasets

		Examples
Inconsistent Dataset	negation	Cairo is not located in [MASK]. Cairo is located in [MASK].
		A guitar does not have [MASK]. A guitar has [MASK].
		A chair has no [MASK]. A chair has [MASK].
		Maria is a mother without a [MASK]. Maria is a mother with a [MASK].
		A cat sees without [MASK]. A cat sees with [MASK].
	coord.	Joe is a dolphin or an [MASK]. Joe is a dolphin and an [MASK].
		Tina is a bird or an [MASK]. Tina is a bird and an [MASK].
	quantifiers	All cars have an [MASK]. No car has an [MASK].
		All cooks [MASK]. No cook [MASK].
		Some shoes have [MASK]. All shoes have [MASK].
Semantic Dataset	syn. negation	A mom is not a [MASK]. forbidden: mom, mother, grandmother, grandma, granddaughter, bride, wife, woman, niece, stepmother, daughter, aunt
		A designer does not [MASK]. forbidden: design
		A guitar does not have [MASK]. forbidden: strings
	lex. negation	A bird flies without [MASK]. forbidden: wings
		John is a father without a [MASK]. forbidden: child
		Peter is a brother without a [MASK]. forbidden: sibling
	coordination	Mark was born in Athens or in [MASK]. forbidden: Greece
		Tina died in Beijing or in [MASK]. forbidden: China
		George comes from Berlin or [MASK]. forbidden: Germany
		Kate is a cat or an [MASK]. forbidden: animal

Table 3: Sample sentences from each dataset and each subset. The complete datasets can be found under: <https://function-words.lingvis.io/>.