

Can current NLI systems handle German word order? Investigating language model performance on a new German challenge set of minimal pairs

Ines Reinig

Data and Web Science Group
Mannheim University
Germany

`ines.reinig@uni-mannheim.de`

Katja Markert

Institute of Computational Linguistics
Heidelberg University
Germany

`markert@cl.uni-heidelberg.de`

Abstract

Compared to English, German word order is freer and therefore poses additional challenges for natural language inference (NLI). We create WOGLI (Word Order in German Language Inference), the first adversarial NLI dataset for German word order that has the following properties: (i) each premise has an entailed and a non-entailed hypothesis; (ii) premise and hypotheses differ only in word order and necessary morphological changes to mark case and number. In particular, each premise and its two hypotheses contain exactly the same lemmata. Our adversarial examples require the model to use morphological markers in order to recognise or reject entailment. We show that current German autoencoding models fine-tuned on translated NLI data can struggle on this challenge set, reflecting the fact that translated NLI datasets will not mirror all necessary language phenomena in the target language. We also examine performance after data augmentation as well as on related word order phenomena derived from WOGLI. Our datasets are publically available at <https://github.com/ireinig/wogli>.

1 Introduction

German is endowed with a rather free word order (Bader and Portele, 2019), especially when it comes to ordering nominal arguments in a sentence. Currently, large German NLI datasets are only available as translations from other languages. For example, the training portion (392k pairs) of the German XNLI dataset (Conneau et al., 2018) is a machine translation of the English MultiNLI training set (Williams et al., 2018). The testing portion of German XNLI is a manual translation of 5k English premise-hypothesis pairs that were newly created by the authors of XNLI. Such translated sets do not necessarily mirror all German-specific linguistic phenomena, such as the freer German

word order.

We construct a new German challenge set named WOGLI (Word Order in German Language Inference). This dataset is handcrafted and does not stem from translation. It contains 16k premises where each premise is accompanied by one entailed (E) and one non-entailed (NE) hypothesis that both contain the same lemmata as the premise but change argument order. Morphological markers are indicative of subject and (direct) object, thus informing about the hypothesis’ entailment relationship to the premise. In other words, WOGLI serves as a test bed for current language models’ capabilities to distinguish subject from object in the context of German word order.

Our contributions are as follows:

1. We propose the first NLI dataset that specifically targets German word order phenomena.
2. We show that current German autoencoding models fine-tuned on the translated XNLI dataset can struggle on our proposed challenge set (Sections 4 and 5), tending to always predict entailment for both hypotheses.
3. We show that data augmentation can help performance on WOGLI but needs a considerable number of examples to work (Section 6).
4. We derive generalization sets including similar word order phenomena to WOGLI to investigate how the augmented models transfer to these datasets and show that German word order remains challenging in NLI (Section 7).

All our datasets are publically available¹.

2 German Word Order

The topological model. The topological model (Drach, 1937) describes regularities in German

¹<https://github.com/ireinig/wogli>

Clause	Order	Prefield	L brack.	Middlefield	R brack.	Count (% of accus.)
Main	SO	Peter	sieht	den Mann		231 (86%)
		Peter <i>Peter</i>	sees <i>sees</i>	the man _{ACC} <i>the man</i>		
	OS	Den Mann	sieht	Peter		38 (14%)
		The man _{ACC} <i>Peter</i>	sees <i>sees</i>	Peter <i>the man</i>		
Emb.	SO		dass	Peter den Mann	sieht	546 (99%)
			that <i>that</i>	Peter the man _{ACC} <i>Peter sees the man</i>	sees	
	OS		dass	den Mann Peter	sieht	6 (1%)
			that <i>that</i>	the man _{ACC} Peter <i>Peter sees the man</i>	sees	

Table 1: Examples for word order in declarative, active German main and embedded clauses with subject and (accusative) direct object arguments, with corpus statistics from Bader and Häussler (2010). As in the remainder of this paper, the subject is always bold. Transliterations and translations (in italics) are provided below each example.

word order, dependent on the concepts of *prefield* and *middlefield* for constituent positioning. In this model, so-called *left and right brackets* form “[t]he skeleton of the sentence” (Bader and Häussler, 2010, p. 719), while other fields are defined according to the position of the verb (Dürscheid, 2012).

Declarative main clauses, such as *Peter sieht den Mann* at the top of Table 1, have a verb-second order. The left bracket contains the finite verb and the prefield is filled with one constituent (Bader and Häussler, 2010; Dürscheid, 2012). In contrast, embedded clauses, such as *dass Peter den Mann sieht* in the bottom half of Table 1, have a verb-last order. In verb-last clauses, the left bracket is occupied by a subjunction, the right bracket by a finite verb or a verb complex, and other constituents are placed in the middlefield (Dürscheid, 2012).

While subject followed by object (SO) is viewed as the canonical word order, it is possible to place the object before the subject (OS) in both embedded and main clauses (Table 1). In the main clause either the subject or object is placed in the prefield, in embedded clauses both are placed in the middlefield but in varying order.

OS acceptability and minimal pairs. The marked OS order is more frequent in main clauses involving the prefield (Bader and Häussler, 2010) (around 14% of main clauses with accusative direct object) and in the active voice (Bader et al., 2017) (see data and examples in Table 1). Therefore, we construct our challenge set using only such clauses to raise acceptability of the marked OS word order examples. Even in the prefield, OS or-

der can vary in acceptability dependent on relative constituent weight (Siewierska, 1993) (shorter before longer), discourse properties such as givenness (Bader and Portele, 2019) (given before new) and semantic properties such as agency (Siewierska, 1993; Bader and Häussler, 2010) (animate before inanimate). As we focus on simple grammatical examples without further interference, however, all our constituents are short and all premises and hypotheses are single sentences. To ensure that entailed and non-entailed sentences are semantically plausible, all our constituents refer to persons.

German word order in XNLI. We extract hypotheses in the training portion of the translated German XNLI (henceforth, GXNLI-train) that are declarative main clauses with a length between 4 and 9 tokens. The 38,090 extracted clauses are in active voice and contain one subject NP and one direct object NP in accusative case. We exclude clauses that start with prepositions or adverbs to limit ourselves to prefield cases. Only 1.8% (698 clauses) of the extracted clauses are in OS order, compared to the 14% to be expected in a German corpus according to Bader and Häussler (2010). Additionally, a vast majority of the 698 OS clauses start with the same demonstrative pronoun object *das/this*, e.g. *Das werde ich tun/This I will do*, thus offering little variety. The extreme prevalence of the SO order in GXNLI-train hypotheses may be due to its translated nature.

3 WOGLI construction

Verb Collection. We collected 50 frequent German transitive verb types including agentive (such as *warnen/warn*), object-experiencer (such as *erschrecken/startle*) and subject-experiencer (such as *lieben/love*) verbs. All verbs can take animate (human) subjects as well as animate (human) direct objects, and all objects take the accusative case. All verbs are not symmetric, meaning that they do not lead to bidirectional entailments.² In addition, none of the verbs need to split prefixes when used in main clauses so that the resulting premises have a very simple SVO structure. All verbs occur at least 70 times in GXNLI-train. Consequently, any difficulties that a language model will experience are unlikely to be due to verb rarity.

Noun Collection. We collected 144 noun types describing humans that function as direct object or subject in our premises/hypotheses. These include 38 masculine common nouns such as *Gast/guest*, each of which was seen at least 10 times in GXNLI-train and 24 feminine common nouns such as *Lehrerin/(female) teacher*. We collected feminine common nouns by searching for the suffix *in* in GXNLI-train, which often indicates female persons in German. The unbalanced masculine-feminine split is due to the automatic translation of GXNLI-train as gender-neutral English job descriptions, for example *doctor*, are most frequently translated via the German male form, e.g. *Arzt* instead of the female form *Ärztin*³. We also collected 41 female and 41 male first names that occur at least 10 times in GXNLI-train. The 144 noun types yield 181 different noun surface forms (nominative/accusative, plural/singular).

Premise and Hypothesis Generation. We automatically generated German premises as declarative, present tense, main clauses in the active voice with SVO structure (see lines 1 and 5) in Table 2). Each SVO premise is accompanied by two hypotheses. H1-SO (NE) exchanges object and subject including changing S/O case markers and potentially verb number markers. Therefore, similarly to English, this change leads to non-entailment, as

²For example, for the symmetric verb *heiraten/marry*, X marries Y would entail Y marries X, which would not allow us to automatically derive non-entailed hypotheses.

³We could have made up the shortfall by including more feminine forms, even if they do not occur in GXNLI-train, but we consider it more important for this study to keep lexical differences to the fine-tuning set minimal.

the premise *The doctor warns the client* and the corresponding H1 *The client warns the doctor* illustrate. We call this subset WOGLI-SO, as the new subject precedes the object. H2-OS (E) simply swaps argument order but keeps case and number markers intact, leading to a sentence synonymous to the premise but with marked OS word order. The resulting set of entailed hypotheses is called WOGLI-OS. Table 2 shows two full examples with case and number marking.

We have 17 patterns due to combinations of different argument NPs, including masculine and feminine proper names and common nouns as well as singular and plural arguments. Subjects/objects are either a simple proper name (such as *Maria*) or consist of an article⁴ and a common noun, e.g. *der Arzt/the doctor*. Consequently, each sentence always has a length of four or five words. A list of all 17 patterns is provided in Table 6 in the Appendix; we exclude the patterns in Table 7 in the Appendix as they generate ambiguous hypotheses, due to the absence of disambiguating morphological markers. The 17 patterns in WOGLI can be divided into two groups: 5 **all-singular** patterns that combine two singular nominal arguments (see first example in Table 2) and 12 **singular-plural** patterns in which one argument is singular and the other one is plural (see second example in Table 2). In all 9 patterns involving a masculine singular NP, (i) masculine determiners and (ii) masculine common nouns belonging to the weak declension type⁵ carry morphological markers of case. Proper nouns never change surface forms. Additionally, in all **singular-plural** patterns, verb number agreement with the subject always leads to a change in the verb's surface form between E and NE hypotheses.

WOGLI statistics. We generate 1,000 premises per pattern by randomly selecting an appropriate subject/object and verb from our lists, leading to 17,000 possible premises. As in random generation, some premises are generated twice, we deduplicate and are left with 16,971 premises. H1-SO (NE) and H2-OS (E) are deterministically generated from the premises, leading to 33,942 sentence pairs.

⁴We used the articles *ein* (indef.), *der* (def.) and *dieser* (demonstrative), as well as their feminine and plural forms.

⁵The six masculine common nouns in WOGLI that belong to the weak declension type are *Kunde/Kunden/client*, *Student/Studenten/student*, *Journalist/Journalisten/journalist*, *Patient/Patienten/patient*, *Soldat/Soldaten/soldier* and *Zeuge/Zeugen/witness*. The remaining masculine nouns, e.g. *Anwalt/lawyer*, maintain the same surface forms in nominative and accusative.

Premise	<u>Der/Dieser/Ein</u> _{NOM-SG-M} <u>The/This/A</u> _{NOM-SG-M} <i>The/This/A</i>	Arzt doctor <i>doctor</i>	warnt _{SG} warns _{SG} <i>warns</i>	<u>den/diesen/einen</u> _{ACC-SG-M} <u>the/this/a</u> _{ACC-SG-M} <i>the/this/a</i>	Kunden [†] client <i>client</i>
H1-SO (NE)	<u>Der/Dieser/Ein</u> _{NOM-SG-M} <u>The/This/A</u> _{NOM-SG-M} <i>The/This/A</i>	Kunde [†] client <i>client</i>	warnt _{SG} warns _{SG} <i>warns</i>	<u>den/diesen/einen</u> _{ACC-SG-M} <u>the/this/a</u> _{ACC-SG-M} <i>the/this/a</i>	Arzt doctor <i>doctor</i>
H2-OS (E)*	<u>Den/Diesen/Einen</u> _{ACC-SG-M} <u>The/This/A</u> _{ACC-SG-M} <i>The/This/A</i>	Kunden [†] client <i>doctor</i>	warnt _{SG} warns _{SG} <i>warns</i>	<u>der/dieser/ein</u> _{NOM-SG-M} <u>the/this/a</u> _{NOM-SG-M} <i>the/this/a</i>	Arzt doctor <i>client</i>
H3-OS (NE)*	<u>Den/Diesen/Einen</u> _{ACC-SG-M} <u>The/This/A</u> _{ACC-SG-M} <i>The/This/A</i>	Arzt doctor <i>client</i>	warnt _{SG} warns _{SG} <i>warns</i>	<u>der/dieser/ein</u> _{NOM-SG-M} <u>the/this/a</u> _{NOM-SG-M} <i>the/this/a</i>	Kunde [†] client <i>doctor</i>
Premise	<u>Der/Dieser/Ein</u> _{NOM-SG-M} <u>The/This/A</u> _{NOM-SG-M} <i>The/This/A</i>	Minister minister <i>minister</i>	empfeht _{SG} recommends _{SG} <i>recommends</i>	<u>die/diese</u> _{ACC-PL-F} <u>the/these</u> _{ACC-PL-F} <i>the/these</i>	Autorinnen authors <i>authors</i>
H1-SO (NE)	<u>Die/Diese</u> _{NOM-PL-F} <u>The/These</u> _{NOM-PL-F} <i>The/These</i>	Autorinnen authors <i>authors</i>	empfehlen _{PL} recommend _{PL} <i>recommend</i>	<u>den/diesen/einen</u> _{ACC-SG-M} <u>the/this/a</u> _{ACC-SG-M} <i>the/this/a</i>	Minister minister <i>minister</i>
H2-OS (E)*	<u>Die/Diese</u> _{ACC-PL-F} <u>The/These</u> _{ACC-PL-F} <i>The/This/A</i>	Autorinnen authors <i>minister</i>	empfeht _{SG} recommends _{SG} <i>recommends</i>	<u>der/dieser/ein</u> _{NOM-SG-M} <u>the/this/a</u> _{NOM-SG-M} <i>the/these</i>	Minister minister <i>authors</i>
H3-OS (NE)*	<u>Den/Diesen/Einen</u> _{ACC-SG-M} <u>The/This/A</u> _{ACC-SG-M} <i>The/These</i>	Minister minister <i>authors</i>	empfehlen _{PL} recommend _{PL} <i>recommend</i>	<u>die/diese</u> _{NOM-PL-F} <u>the/these</u> _{NOM-PL-F} <i>the/this/a</i>	Autorinnen authors <i>minister</i>

Table 2: Two examples of WOGLI premise-hypothesis pairs, one for the pattern `sing_masc.v.sing_masc` and one for the pattern `sing_masc.v.pl_fem`. Underlined words have different surface forms in NE and E hypotheses and carry distinguishing morphological markers of case and/or number. Nouns belonging to the weak declension type are identified by [†]. Hypotheses H3 are not part of WOGLI proper but will be used in a generalization set called WOGLI-OS-hard as they demand to both process marked OS word order as well as recognising non-entailment in the face of high word overlap. As in the remainder of this paper, hypotheses with a marked word order are identified by an asterisk.

All word lists with GXNLI-train frequencies and translations can be found in our Github repository. Each of the 50 verb types appears between 308 and 383 times (mean: 339.4 times) in the 16,971 premises. They also appear 20 times on average per pattern in the premises. Table 8 in the Appendix gives noun statistics for WOGLI.

4 Experiments on WOGLI

Models. We use two German models and one multilingual BERT model:

- BERT-base⁶ is a cased base BERT model pre-trained by the MDZ Digital Library team on 16GB of German-language text.
- GBERT-large⁷ is a BERT model pre-trained on 163.4GB of data (Chan et al., 2020), using the same cased vocabulary as BERT-base.⁸

⁶<https://huggingface.co/dbmdz/bert-base-german-cased>

⁷<https://huggingface.co/deepset/gbert-large>

⁸Other large-scale German language models such as Gott-

- mBERT-base⁹ is a cased BERT model pre-trained on 104 languages (Devlin et al., 2019).

Since models were fine-tuned on GXNLI-train in a three-class setting, we merge contradiction and neutral into non-entailed predictions for evaluations on WOGLI. Fine-tuning details are provided in Section C of the Appendix.

Results (see Table 3). As a sanity check, we first test our models on GXNLI-test. Our models’ performances on GXNLI-test are broadly in line with published work. Conneau et al. (2020) achieve an accuracy of 81.2% on GXNLI-test with a monolingual BERT-base model, higher than our 76.67%. However, their model uses a larger vocabulary (40k, ours: 31k) and was pre-trained on a larger corpus (up to 60GB, ours: 16GB). This particular model is unfortunately not available. Other prior work concentrates on multilingual models: GBERT-large’s BERT (Scheible et al., 2020) and GELECTRA (Chan et al., 2020) are of similar size and downstream performance. Thus, we use GBERT-large as a representative.

⁹<https://huggingface.co/bert-base-multilingual-cased>

size is smaller than mT5-base (580m parameters) by Xue et al. (2021), but its performance of 84.65% on GXNLI-test exceeds the one reported for mT5-base (81.6%). Devlin et al. (2019) achieve an accuracy of 75.9% with mBERT-base (*Translate Train Cased*)¹⁰, in line with ours.

On WOGLI, both base models completely fail, labeling almost all instances as entailments. GBERT-large performs a bit better, suggesting that the language model’s scale plays a role in its ability on WOGLI. However, it still shows a strong tendency for the entailment class and the results are not robust across runs. Our vocabulary is frequent and present in GXNLI-train and our sentences have a very simple grammar. Therefore, the models’ poor performances on WOGLI suggest that not all German-specific linguistic phenomena are represented in the translated GXNLI-train, similar to our GXNLI word order analysis in Section 2.¹¹

5 Error analysis

All analyses in this section are carried out on ensemble predictions (majority vote of the 5 runs) of the strongest model in Table 3, GBERT-large. The ensemble model reaches an accuracy of 57.82% on WOGLI and 27.41% on WOGLI-SO.

5.1 Fluency

We measure the correlation of model performance and linguistic acceptability, approximating the latter via pseudo-loglikelihood (Salazar et al., 2020). WOGLI premises have an average PLL of -30.54 (SD: 8.318). H1-SO (NE) hypotheses have an average PLL of -30.56 (SD: 8.287), while H2-OS (E) hypotheses are less fluent due to marked word order, with an average PLL of -36.53 (SD: 8.535). GBERT-large performs worse on SO (NE) pairs than on the less fluent OS (E) pairs; fluency thus does not play an important role in the model’s performance on WOGLI. Instead, the lexical overlap heuristic (Naik et al., 2018; McCoy et al., 2019; Gururangan et al., 2018) is a possible reason for the degradation on non-entailed pairs.

¹⁰Results on XNLI are provided in the corresponding GitHub repository: <https://github.com/google-research/bert/blob/master/multilingual.md#results>.

¹¹One question that arises is whether even larger models or models pretrained on substantially more data will solve the problem. Other monolingual models for German are sparse. We therefore ran two large, publically available, multilingual model checkpoints fine-tuned on XNLI on WOGLI. They also do not perform well (see Section D in the Appendix).

5.2 Performance by subject and object properties

We now focus on WOGLI-SO (NE) only as this is the part of the dataset where the models fail.

Gender. Regarding the gender of arguments in WOGLI, we formulate the following hypothesis:

- A1** SO hypotheses with masculine subjects (objects) are easier to classify than the ones with feminine subjects (objects).

A1 can be explained by (a) the presence of gender bias due to translation in GXNLI-train (see Section 3) or (b) morphological differences between masculine and feminine NPs.

Performance on instances in WOGLI-SO (NE) with masculine common noun subjects is indeed significantly higher than for feminine common noun subjects. The same holds for common noun objects (see also Table 10 in the Appendix). However, this does not transfer to proper names. Gender bias in GXNLI-train (a) as an explanation for **A1** is therefore unlikely.

Morphological differences between feminine and masculine NPs (b), however, are a possible explanation for **A1**. Feminine articles and common nouns have the same surface forms in accusative/nominative. Masculine articles and common nouns, however, can bear morphological case markers. The masculine singular articles *der*, *ein* and *dieser* are the only articles in WOGLI to change surface forms in the accusative to *den*, *einen* and *diesen*. Additionally, singular masculine common nouns belonging to the weak declension type also carry case markers. Morphological markers in some masculine NPs could thus be helpful for the model to distinguish subject from object.

Referential properties of subjects/objects. In prefield SO sentences, definite NPs tend to precede indefinite NPs (Weber and Müller, 2004), probably because indefinite constituents are often new and definite constituents are often given (Chafe, 1976). Although XNLI and WOGLI do not contain discourse context, preference for SO sentences with definite before indefinite NPs might be encapsulated in pretraining data. We thus hypothesize that:

- A2** SO hypotheses in which a definite NP precedes an indefinite NP are easier to classify.

Evaluation set	BERT-base (110m)	GBERT-large (335m)	mBERT-base (172m)
GXNLI-test	76.65 (0.41)	84.65 (0.163)	75.16 (0.552)
WOGLI	50.16 (0.133)	57.68 (1.86)	50.01 (0.015)
WOGLI-SO (NE)	0.33 (0.269)	27.42 (7.828)	0.02 (0.029)
WOGLI-OS (E)*	100 (0.005)	87.94 (4.171)	100 (0.0)

Table 3: Accuracies for two German and one multilingual model on GXNLI-test and WOGLI, averaged over 5 runs. All are trained on GXNLI-train. Accuracies are computed for 3 classes in GXNLI-test and 2 classes in WOGLI.

We separate WOGLI constituents into definite and indefinite following Prince (1992): definite and demonstrative articles, as well as proper names are markers of definiteness, while indefinite articles point to indefiniteness. We then separate WOGLI-SO (16,971 pairs) into two groups: **preferred** (14,671 pairs) and **dispreferred** (2,300 pairs). Pairs in the **dispreferred** group are opposed to the aforementioned discourse hierarchy in that indefinite constituents precede definite constituents in the SO hypothesis. Pairs in the **preferred** group form the three other possible cases: definite precedes indefinite, definite precedes definite and indefinite precedes indefinite; these cases are not in opposition to the hierarchy.

GBERT-large achieves an accuracy of 29.85% on the SO pairs in the **preferred** group but only 11.78% on the **dispreferred** group (difference significant at 1% significance level, z-test for proportions). Therefore we can confirm **A2**.

Number. Lastly, we analyse the role of verb number agreement in classifying WOGLI-SO (NE) instances. As explained in Section 3, WOGLI patterns either combine only singular arguments (**all-singular**) or a singular and a plural argument (**singular-plural**). Only in the latter group of patterns, subject-verb agreement leads to a change in the verb’s surface form from the premise to the H1-SO (NE) hypothesis (see *empfehlen/recommend_{PL}* vs. *empfehlt/recommends_{SG}* in the second example in Table 2). We investigate the importance of verb number agreement for classifier performance by separating WOGLI-SO (16,971 pairs) into two groups, **all-singular** (4,997 pairs) and **singular-plural** (11,974 pairs).

GBERT-large achieves an accuracy of 36.66% on the SO pairs in the **all-singular** group and 23.54% on the **singular-plural** group (difference significant at 1% significance level, z-test for proportions). Thus the number switch in the verb occurring in **singular-plural** SO hypotheses is not a particularly helpful cue for the classifier.

6 Data augmentation

Following McCoy et al. (2019) and Min et al. (2020) on data augmentation with challenge sets, we hypothesize that augmenting GXNLI-train with a WOGLI subset can be helpful.

We sample 1,037 premises and their corresponding E/NE hypotheses from WOGLI, resulting in 2,074 training instances. Each of the 17 patterns occurs 61 times. All 50 verb lemmas are represented, each appearing between 18 and 25 times. All 181 noun forms appear at least once.¹²

We concatenate these WOGLI instances with GXNLI-train, name the resulting augmented training set GXNLI+1037 and shuffle it before fine-tuning GBERT-large 10 times on this augmented training set. We evaluate on the remaining 31,868 WOGLI instances, named WOGLI-test-1037. This augmented training set allows GBERT-large to classify WOGLI almost perfectly, while maintaining its performance on GXNLI-test (Table 4).

Smaller augmentation size. We fine-tune GBERT-large on a shuffled concatenation of GXNLI-train and only 102 WOGLI premises sampled in a stratified manner from the aforementioned 1,037 premises along with both their corresponding NE and E hypotheses. Each one of the 17 patterns appears 6 times and each one of the 50 verb lemmas appears at least once and at most 4 times. Due to the small augmentation size, it is not possible to ensure representation of all 181 nouns, with 73 not appearing. We evaluate on the remaining 33,738 WOGLI pairs, named WOGLI-test-102. The smaller augmentation size yields a model that performs worse and less robustly on WOGLI test instances (Table 4).

7 Generalization experiments

McCoy et al. (2019) investigate whether augmented models improved by simply memorizing the seen

¹²The nouns occur in varying frequencies due to the small size of the augmentation set.

Evaluation set	GXNLI+1037	Evaluation set	GXNLI+102
GXNLI-test	84.7 (0.301)	GXNLI-test	84.78 (0.244)
WOGLI-test-1037	99.98 (0.016)	WOGLI-test-102	86.04 (4.091)
WOGLI-SO-test-1037 (NE)	99.99 (0.008)	WOGLI-SO-test-102 (NE)	87.57 (6.428)
WOGLI-OS-test-1037 (E)*	99.97 (0.03)	WOGLI-OS-test-102 (E)*	84.52 (4.45)

(a) Larger augmentation

(b) Smaller augmentation

Table 4: Accuracy for GBERT-large fine-tuned on GXNLI-train augmented with WOGLI instances. Results are averaged over 10 runs and computed in a 3-class (GXNLI-test) or a 2-class (WOGLI-test) setting.

templates. To do so, they evaluate them on pairs from unseen patterns. Inspired by this setup, we study the models’ generalization capabilities by evaluating them on four new evaluation sets that share structural and lexical similarities with the WOGLI pairs that were seen during fine-tuning.

7.1 Construction of generalization sets

Pronoun subjects: WOGLI-p-subject. We replace the premise subject in WOGLI by a personal pronoun (*He warns the client*). Correspondingly, the H1-SO (NE) hypothesis then has the pronoun as the object (*The client warns him*) whereas the entailed H2-OS (E) hypothesis just swaps word order with regards to the premise (*The client warns he*) (see also Table 11 in the Appendix). To focus on the pronominalization change, the same 17 patterns, verb lemmas, proper nouns and common nouns are also used in WOGLI-p-subject. In addition to the previously mentioned morphological markers of case and/or verb number occurring in WOGLI sentences (Section 3), the masculine singular pronoun *er/he* (nominative) in WOGLI-p-subject changes surface form in the accusative case (*ihn/him*). Feminine and plural pronouns (*sie/she/her/they/them*) in WOGLI-p-subject, however, do not change surface form. Some WOGLI premises can become duplicates after replacing the subject by a personal pronoun. Consider the two premises *Die Ärzte warnen den Kunden/The doctors_{masc} warn the client* and *Die Ärztinnen warnen den Kunden/The doctors_{fem} warn the client*. After replacing the subject, both premises lead to the new premise *Sie warnen den Gast/They warn the guest*, since plural masculine and plural feminine nominative personal pronouns have the same surface form in German. We keep only one version for such duplicates. The new generalization set contains 13,802 unique premises, or a total of 27,604 pairs.

Dative: WOGLI-dative. We collect a new list of 22 transitive verbs that require dative instead of accusative objects. All verbs are not symmetric, which ensures that NE hypotheses always have the correct gold label. Each verb lemma appears at least 17 times in GXNLI-train. We use the same 144 noun types as in WOGLI to generate new instances. The premises again have SVO structure, and H1 has SO (NE) and H2 has OS (E) structure. Therefore the instances are completely parallel to WOGLI apart from the case of the object.

In these dative constructions, 24 patterns are possible (Table 6 in the Appendix). Each pattern appears 150 times in WOGLI-dative and each verb lemma appears between 132 and 182 times in the premises. All possible noun surface forms appear between 6 and 81 times in the premises. We generate 3,600 premises, or 7,200 pairs in total. Table 12 in the Appendix shows an example. In WOGLI-dative, all determiners (singular and plural, feminine and masculine) change surface forms by case. Additionally, plural masculine common nouns change surface forms in the dative if they do not end with -n in the nominative¹³. As in WOGLI, singular masculine nouns of the weak declension type and verbs in singular-plural patterns also change surface forms.

Ditransitive verbs. We collect 21 ditransitive verbs (such as *schicken/send* and *verheimlichen/conceal*), each of which appears at least 6 times in GXNLI-train. Verbs are grouped into 5 semantic categories (giving, taking, sending, communication, secret). Subjects and indirect objects of the verbs are compatible with the semantic class human, so that we can reuse the 144 noun types from WOGLI.

¹³Masculine nouns ending with -n in plural nominative are: *Kunden/clients*, *Professoren/professors*, *Studenten/students*, *Mentoren/mentors*, *Patienten/patients*, *Soldaten/soldiers*, *Journalisten/journalists*, *Zeugen/witnesses*. These nouns maintain the same surface forms in plural dative.

For direct objects, we use a new list of 54 common nouns, appearing at least 15 times in GXNLI-train. They are grouped with the verb semantic categories so that resulting premises/hypotheses are meaningful (thus, you can combine the direct object *Identität/identity* with *secret* verbs but not with *sending* verbs). The direct object is always preceded by a definite article.

Ditransitive premises follow the preferred word order SiO: subject-verb-IndirectObject-DirectObject (*The waitresses give the merchant the cake*). Very similar to WOGLI, the not-entailed H1 hypothesis swaps the underlying arguments of subject and indirect object, adapting case and number (*The merchant gives the waitresses the cake*) whereas the entailed H2 hypothesis reorders subject and indirect object into the marked iOS word order without changing the meaning by keeping case and number markers intact (*The merchant give the waitresses the cake*). The direct object is not affected. An example is shown in Table 13 in the Appendix. With respect to morphological markers, WOGLI-ditransitive follows the same surface form changes between E and NE hypotheses as WOGLI-dative, since indirect objects in WOGLI-ditransitive are in dative case.

Ditransitives allow for 24 unique patterns. We allow each pattern to appear 1,000 times leading to 12,000 premises or 24,000 pairs.

WOGLI-OS-hard (NE). Neither WOGLI nor the previous generalization datasets contain instances where the marked OS word order leads to non-entailment, i.e. where you have to recognise non-entailment in the face of high word overlap while at the same time processing a rare word order. Therefore we create a third hypothesis H3-OS (NE) for each WOGLI premise where we similar to H1 invert the underlying arguments but present this changed meaning in OS word order. Two examples are given as H3-OS (NE) in Table 2. This is possible for all 17 WOGLI patterns. Pairing H3 with each WOGLI premise leads to 16,971 new non-entailed pairs. All premises as well as all lexical items have been seen in normal WOGLI.

7.2 Generalization results

We evaluate GBERT-large fine-tuned on GXNLI without augmentation (GXNLI+0) as well as fine-tuned on GXNLI+1037 and on GXNLI+102 on our four generalization sets. Results are in Table 5.

GXNLI+1037 transfers very well to WOGLI-p-

subject, while GXNLI+102 reaches an accuracy of only 59.03% on SO instances and is less robust. Thus, even for simple pronoun replacement a relatively large augmentation size is needed. A similar picture emerges for WOGLI-dative. Since WOGLI-dative contains more patterns than WOGLI, we investigate whether GXNLI+102’s poor performance is only observable in patterns that were not seen during fine-tuning but find no preference for seen or unseen patterns.

With respect to ditransitive pairs, GXNLI+1037 has almost perfect accuracy and GXNLI+102 reaches its best generalization set performance, reaching similar results as on standard WOGLI.

We hypothesized that generalization to H3-OS (NE) in WOGLI-OS-hard is the most difficult as it contains both marked word order and non-entailment, whereas (i) in GXNLI, the marked word order is very rare (see Section 2) and (ii) in WOGLI, the marked word order has always been seen with the entailment class, potentially tripping up an augmented model that could have learnt this hypothesis-only fact. This turns out to be true: GXNLI+0 classifies basically all WOGLI-OS-hard (NE) examples wrongly as entailment and performs even worse than the same model on the original WOGLI-SO (NE) non-entailed examples (see the 27.42% in Table 3). With substantial augmentation (GXNLI+1037), performance is slightly better but the results are still both very low and unstable.

Our generalization experiments show that (i) the augmentation set needs to be sufficiently large for successful generalization to new NLI pairs that are structurally similar to WOGLI and (ii) models exposed to WOGLI do not necessarily generalize well to some related datasets at all. As German word order is quite intricate and will have additional variations for embedded or non-declarative clauses this means training datasets need to be very large and varied to learn German word order.

8 Related work

Many English adversarial NLI datasets have been proposed. Some of these (Dasgupta et al., 2018; Kim et al., 2018; Nie et al., 2019; McCoy et al., 2019), like us, include minimal pairs with a high word overlap between premise and hypotheses. Kim et al. (2018), for example, change argument order to generate non-entailments so that “understanding” word order is necessary to solve these. However, in WOGLI, changes in argument order

Evaluation set	GXNLI+0	GXNLI+102	GXNLI+1037
WOGLI-p-subject-test	53.23 (1.715)	77.97 (6.653)	98.89 (0.957)
WOGLI-p-subject-SO-test (NE)	7.34 (4.16)	59.03 (13.353)	97.78 (1.916)
WOGLI-p-subject-OS-test (E)*	99.12 (0.74)	96.91 (1.822)	99.99 (0.008)
WOGLI-dative-test	58.11 (2.789)	79.4 (5.446)	94.72 (0.565)
WOGLI-dative-SO-test (NE)	17.76 (6.223)	60.87 (11.227)	91.49 (1.421)
WOGLI-dative-OS-test (E)*	98.47 (0.776)	97.93 (0.534)	97.96 (0.504)
WOGLI-ditransitive-test	73.93 (6.327)	92.59 (4.634)	99.58 (0.143)
WOGLI-ditransitive-SiO-test (NE)	50.23 (13.11)	86.55 (9.41)	99.62 (0.261)
WOGLI-ditransitive-iOS-test (E)*	97.63 (0.635)	98.63 (0.591)	99.55 (0.276)
WOGLI-OS-hard (NE)*	0.15 (0.082)	0.77 (0.75)	23.45 (15.985)

Table 5: Accuracy on generalization sets, averaged over 5 runs for GXNLI+0 and over 10 runs for remaining models

generate entailed *and* non-entailed hypotheses, depending on keeping or changing corresponding morphology. The more fixed English word order does not allow for flexibility to that degree.

Regarding adversarial NLI datasets for German, Hartmann et al. (2021) investigate negation but do not work on word order. Tikhonova et al. (2022) propose NLI diagnostic datasets for French, German and Swedish. Sentence pairs are manually translated from the Russian TERRa dataset (Shavrina et al., 2020) as well as from the diagnostic dataset of GLUE (Wang et al., 2018). We inspected a random 100 hypotheses of the German TERRa dataset, none of which were in marked word order. The translated GLUE benchmark is annotated with linguistic features relevant for entailment such as lexical semantics, logic, and predicate-argument structure. Only the predicate-argument structure examples include a handful where word order of arguments has been inverted between premise and hypothesis. However, resulting hypotheses were often ambiguous and — in our opinion — wrongly annotated as not-entailed. Consider the premise *John zerbrach das Fenster/John broke the window* and the hypothesis *Das Fenster hat John eingeschlagen*, which is ambiguous between *The window_{NOM} broke John_{ACC}* (SO order, NE) OR *The window_{ACC} broke John_{NOM}* (OS order, E). This is annotated as non-entailment in the dataset, assuming SO order with an implausible semantic reading, whereas the marked word order with a plausible semantic reading leads to entailment.

Unlike us, both datasets do not emphasise word order. They are also based on translations and therefore rarely contain OS hypotheses.

9 Conclusion

We created WOGLI, a new NLI challenge set, in order to examine the challenges brought by the freer German word order. Premises, entailed and not-entailed hypotheses contain exactly the same lemmata; the two hypotheses differ only in word order and morphological changes but change label. Three current BERT-based models fine-tuned on GXNLI-train struggle on WOGLI pairs. This poor performance mirrors the fact that translated NLI training sets such as GXNLI do not incorporate all required linguistic phenomena that are specific to the target language, German. We find that the number of WOGLI pairs for augmentation during fine-tuning must be sufficiently high in order to (i) learn WOGLI and (ii) generalize to other WOGLI-like pairs. Even with a larger augmentation set and a large pretrained model, a generalization set that differs more from WOGLI, such as WOGLI-OS-hard (NE), remains difficult.

In future experiments, we will expand WOGLI datasets to contain additional variation, such as tense variation, more complex sentence structure (additional arguments and adjuncts, active/passive), more complex constituent structure and other sentence types (non-declarative, embedded). This will also allow us to conduct more fine-grained error analyses regarding the hierarchies that influence the linearization of arguments and thus word order.

Acknowledgments

Ines Reinig is funded by the German Research Foundation (DFG) under the UNCOVER project (PO 1900/7-1 and RE 3536/3-1).

References

- Markus Bader, Emilia Ellsiepen, Vasiliki Koukouloti, and Yvonne Portele. 2017. Filling the prefield: Findings and challenges. In *DGfS 2016 workshop “V2 in grammar and processing: Its causes and its consequences”*, pages 27–49.
- Markus Bader and Jana Häussler. 2010. Word order in German: A corpus study. *Lingua*, 120(3):717–762.
- Markus Bader and Yvonne Portele. 2019. Givenness and the licensing of object-first order in German: The effect of referential form. In *Proceedings of Linguistic Evidence 2018: Experimental Data Drives Linguistic Theory*. Universität Tübingen.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Wallace Chafe. 1976. *Givenness, contrastiveness, definiteness, subjects, topics, and point of view*, pages 25–55. Academic Press, New York.
- Branden Chan, Stefan Schweter, and Timo Möller. 2020. [German’s next language model](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6788–6796, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Ishita Dasgupta, Demi Guo, Andreas Stuhlmüller, Samuel J. Gershman, and Noah D. Goodman. 2018. Evaluating compositionality in sentence embeddings. In *Proceedings of the Fortieth Annual Conference of the Cognitive Science Society*.
- 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.
- Erich Drach. 1937. *Grundgedanken der Deutschen Satzlehre*. M. Diesterweg.
- Christa Dürscheid. 2012. *Syntax: Grundlagen und Theorien*, volume 6. Vandenhoeck & Ruprecht.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Mareike Hartmann, Miryam de Lhoneux, Daniel Herscovich, Yova Kementchedjhieva, Lukas Nielsen, Chen Qiu, and Anders Søgaard. 2021. [A multilingual benchmark for probing negation-awareness with minimal pairs](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 244–257, Online. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- Juho Kim, Christopher Malon, and Asim Kadav. 2018. [Teaching syntax by adversarial distraction](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 79–84, 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.
- Junghyun Min, R. Thomas McCoy, Dipanjan Das, Emily Pitler, and Tal Linzen. 2020. [Syntactic data augmentation increases robustness to inference heuristics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2339–2352, Online. Association for Computational Linguistics.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress test evaluation for natural language inference](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

- Yixin Nie, Yicheng Wang, and Mohit Bansal. 2019. Analyzing compositionality-sensitivity of NLI models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6867–6874.
- Ellen F Prince. 1992. The ZPG letter: Subjects, definiteness, and information-status. *Discourse description: diverse analyses of a fund raising text*, pages 295–325.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrina Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Raphael Scheible, Fabian Thomczyk, Patric Tippmann, Victor Jaravine, and Martin Boeker. 2020. GottBERT: a pure German language model. *arXiv preprint arXiv:2012.02110*.
- Tatiana Shavrina, Alena Fenogenova, Emelyanov Anton, Denis Shevelev, Ekaterina Artemova, Valentin Malykh, Vladislav Mikhailov, Maria Tikhonova, Andrey Chertok, and Andrey Evlampiev. 2020. [RussianSuperGLUE: A Russian language understanding evaluation benchmark](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4717–4726, Online. Association for Computational Linguistics.
- Anna Siewierska. 1993. [On the interplay of factors in the determination of word order](#). In Joachim Jacobs, Arnim von Stechow, Wolfgang Sternefeld, and Theo Vennemann, editors, *An International Handbook of Contemporary Research*, pages 826–846. De Gruyter Mouton, Berlin • New York.
- Mirac Suzgun, Luke Melas-Kyriazi, and Dan Jurafsky. 2022. Prompt-and-rerank: A method for zero-shot and few-shot arbitrary textual style transfer with small language models. *arXiv preprint arXiv:2205.11503*.
- Maria Tikhonova, Vladislav Mikhailov, Dina Pisarevskaya, Valentin Malykh, and Tatiana Shavrina. 2022. Ad astra or astray: Exploring linguistic knowledge of multilingual BERT through NLI task. *Natural Language Engineering*, pages 1–30.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Andrea Weber and Karin Müller. 2004. Word order variation in German main clauses: A corpus analysis. In *20th International Conference on Computational Linguistics*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

A SVO patterns

Table 6 lists the patterns that we use to build WOGLI and generalization sets. Table 7 lists the 8 patterns that we exclude from WOGLI. The noun phrases in the SO and OS form have the same morphological surface form and the verb also has the same form in both word orders. Therefore, SO and OS meaning are not distinguishable.

B WOGLI statistics

Table 8 shows counts and ratios for the subject and object roles in WOGLI. The different noun categories shown in the first column generally take the subject role as often as they take the object role.

C Fine-tuning details

The input sequence, consisting of the premise-hypothesis pair, is encoded using the given BERT model. The final hidden state of the special [CLS] token constitutes the aggregate representation of the input sequence, following Devlin et al. (2019). This representation is then passed through a dropout layer and a linear classification layer, which maps it to the three-label classification space. All models were fine-tuned for three epochs, with linear warmup over 6% of the first steps and a maximum sequence length of 128. BERT-base and mBERT-base were fine-tuned with a batch size of 16 and a learning rate of $5e^{-5}$. GBERT-large was fine-tuned with a batch size of 32 and a learning rate of $5e^{-6}$. Regarding mBERT-base, fine-tuning on the translated training set is comparable to the *Translate Train* setup in Conneau et al. (2018). We also experimented with fine-tuning mBERT-base on the English MNLI training set, similarly to the *Zero Shot* setup in Conneau et al. (2018), but found better validation set accuracy using the translated training set. GBERT-large was fine-tuned on one NVIDIA A6000 GPU. Base models were fine-tuned on one NVIDIA T4 GPU. Fine-tuning took approximately 2 hours per run.

D WOGLI results for other models

Table 9 shows results on WOGLI for publicly available checkpoints (single runs) of two larger multilingual models: XLM-RoBERTa-large¹⁴ (Conneau et al., 2020) (XLM-R) and the

¹⁴<https://huggingface.co/joeddav/xlm-roberta-large-xnli>

generative encoder-decoder model mT5-large¹⁵ (Xue et al., 2021). These two models have considerably more parameters than GBERT-large. According to the respective model cards:

- XLM-R was fine-tuned on the concatenation of the MNLI training set and the XNLI validation and test sets.
- mT5-large was fine-tuned on the MNLI and the XTREME XNLI¹⁶ (Hu et al., 2020) training sets.

Both large models perform much better than the two base models (see Table 3 in the paper), which suggests again that model scale is relevant on WOGLI. However, they do not achieve higher overall accuracies than GBERT-large (average: 57.68%). Interestingly, mT5-large performs best on WOGLI-SO, but struggles substantially on WOGLI-OS, often labeling these pairs as non-entailments.

ChatGPT: discussion. In a small-scale experiment, we evaluate the ability of the recently made available research preview for the chatbot ChatGPT (February 13 version) by OpenAI¹⁷ on WOGLI. This chatbot is based on the autoregressive GPT-3 model (Brown et al., 2020), as opposed to autoencoding models such as BERT, and has recently drawn a lot of attention in the AI community. We attempted to obtain classifications from ChatGPT on a WOGLI subset consisting of 51 WOGLI-SO and 51 WOGLI-OS pairs. However, we observed (i) a strong prompt-dependence (Suzgun et al., 2022), as even minor changes in the prompt’s phrasing lead to different answers by the chatbot and (ii) overall inconsistent results across multiple instances of showing the model the same sets of pairs. Due to the inconsistency of these preliminary results, we leave it to future work to assess ChatGPT’s capabilities on WOGLI in a more systematic manner and for a range of different prompt styles.

E Error analysis: performance by gender

Table 10 provides more detailed results for the analysis discussed in Section 5.2.

¹⁵<https://huggingface.co/alan-turing-institute/mt5-large-finetuned-mnli-xtreme-xnli>

¹⁶This version of the XNLI dataset contains different machine translations than the original XNLI dataset: https://www.tensorflow.org/datasets/catalog/xtreme_xnli

¹⁷<https://openai.com/blog/chatgpt>

Dative, Ditransitive	WOGLI, WOGLI-OS-hard (NE)
pnoun.v.sing.masc	pnoun.v.sing.masc
pnoun.v.plural.masc	pnoun.v.plural.masc
pnoun.v.plural.fem	pnoun.v.plural.fem
pnoun.v.sing.fem	
plural.masc.v.pnoun	plural.masc.v.pnoun
plural.masc.v.sing.masc	plural.masc.v.sing.masc
plural.masc.v.sing.fem	plural.masc.v.sing.fem
plural.masc.v.plural.fem	
plural.masc.v.plural.masc	
plural.fem.v.sing.masc	plural.fem.v.sing.masc
plural.fem.v.sing.fem	plural.fem.v.sing.fem
plural.fem.v.pnoun	plural.fem.v.pnoun
plural.fem.v.plural.fem	
plural.fem.v.plural.masc	
sing.masc.v.sing.masc	sing.masc.v.sing.masc
sing.masc.v.plural.masc	sing.masc.v.plural.masc
sing.masc.v.plural.fem	sing.masc.v.plural.fem
sing.masc.v.sing.fem	sing.masc.v.sing.fem
sing.masc.v.pnoun	sing.masc.v.pnoun
sing.fem.v.sing.masc	sing.fem.v.sing.masc
sing.fem.v.plural.fem	sing.fem.v.plural.fem
sing.fem.v.plural.masc	sing.fem.v.plural.masc
sing.fem.v.pnoun	
sing.fem.v.sing.fem	

Table 6: Exhaustive list of patterns used to build WOGLI-dative, WOGLI-ditransitive (24 patterns) and WOGLI (17 patterns). WOGLI-OS-hard (NE) uses the same patterns as WOGLI.

F Generalization sets

Tables 11, 12 and 13 provide examples for pairs created for generalization sets.

Pattern	Premise	Hypothesis	Label
sing_fem.v.pnoun	Die Freundin begrüßt David . The friend _{CASE?-SING-FEM} greets David _{CASE?-SING-MASC}	David begrüßt die Freundin . David _{CASE?-SING-MASC} greets the friend _{CASE?-SING-FEM}	?
pnoun.v.sing_fem	David begrüßt die Freundin . David _{CASE?-SING-MASC} greets the friend _{CASE?-SING-FEM}	Die Freundin begrüßt David . The friend _{CASE?-SING-FEM} greets David _{CASE?-SING-MASC}	?
pnoun.v.pnoun	Walter begrüßt David . Walter _{CASE?-SING-MASC} greets David _{CASE?-SING-MASC}	David begrüßt Walter . David _{CASE?-SING-MASC} greets Walter _{CASE?-SING-MASC}	?
sing_fem.v.sing_fem	Die Mitbewohnerin begrüßt die Freundin . The flatmate _{CASE?-SING-FEM} greets the friend _{CASE?-SING-FEM}	Die Freundin begrüßt die Mitbewohnerin . The friend _{CASE?-SING-FEM} greets the flatmate _{CASE?-SING-FEM}	?
plural_fem.v.plural_fem	Die Freundinnen begrüßen die Mitbewohnerinnen . The friends _{CASE?-PL-FEM} greet the flatmates _{CASE?-PL-FEM}	Die Mitbewohnerinnen begrüßen die Freundinnen . The flatmates _{CASE?-PL-FEM} greet the friends _{CASE?-PL-FEM}	?
plural_masc.v.plural_masc	Die Freunde begrüßen die Mitbewohner . The friends _{CASE?-PL-MASC} greet the flatmates _{CASE?-PL-MASC}	Die Mitbewohner begrüßen die Freunde . The flatmates _{CASE?-PL-MASC} greet the friends _{CASE?-PL-MASC}	?
plural_masc.v.plural_fem	Die Freunde begrüßen die Mitbewohnerinnen . The friends _{CASE?-PL-MASC} greet the flatmates _{CASE?-PL-FEM}	Die Mitbewohnerinnen begrüßen die Freunde . The flatmates _{CASE?-PL-FEM} greet the friends _{CASE?-PL-MASC}	?
plural_fem.v.plural_masc	Die Freundinnen begrüßen die Mitbewohner . The friends _{CASE?-PL-FEM} greet the flatmates _{CASE?-PL-MASC}	Die Mitbewohner begrüßen die Freundinnen . The flatmates _{CASE?-PL-MASC} greet the friends _{CASE?-PL-FEM}	?

Table 7: The 8 patterns that we exclude from WOGLI and WOGLI-OS-hard (NE)

Noun	# types	Subject count			Object count			Subject/Object ratio			
		Mean	Min	Max	Mean	Min	Max	Mean	Median	Min	Max
Masc pnoun	41	36.8 (4.22)	27	49	36.5 (6.89)	25	53	1.05	1.0	0.58	1.72
Fem pnoun	41	36.3 (5.71)	27	52	36.5 (6.18)	23	49	1.02	1.0	0.63	1.57
Masc sing. cnoun	38	131.4 (12.23)	108	163	131.5 (10.38)	114	158	1.01	1.0	0.78	1.38
Masc pl. cnoun	38	78.7 (9.23)	60	97	78.8 (8.86)	53	99	1.01	1.01	0.75	1.36
Fem sing. cnoun	24	124.8 (11.05)	103	143	124.7 (7.90)	109	137	1.01	1.0	0.81	1.25
Fem pl. cnoun	24	124.7 (12.08)	100	144	124.8 (12.40)	101	145	1.01	0.98	0.77	1.29

Table 8: Average counts and subject to object ratios for different groups of nouns in WOGLI. For example, masculine proper nouns are subjects 36.8 times and objects 36.5 times on average. Values in parentheses are standard deviations.

Evaluation set	XLM-R (550m)	mT5-large (1.2b)
WOGLI	55.42	52.26
WOGLI-SO (NE)	46.2	68.7
WOGLI-OS (E)*	64.64	35.82

Table 9: Accuracies on WOGLI for two larger multilingual models. Results are for single runs.

Constituent	Argument		Conditional probability		Signif.
	Premise	Hypo (SO)	Definition	Value (%)	
Common noun	Subject	Object	$p(\text{correct} \mid \text{SO, m. cnoun psubj})$	28.17	95%
			$p(\text{correct} \mid \text{SO, f. cnoun psubj})$	26.22	
Common noun	Object	Subject	$p(\text{correct} \mid \text{SO, m. cnoun pobj})$	33.11	99%
			$p(\text{correct} \mid \text{SO, f. cnoun pobj})$	22.42	
Proper noun	Subject	Object	$p(\text{correct} \mid \text{SO, m. pnoun psubj})$	27.50	n.s.
			$p(\text{correct} \mid \text{SO, f. pnoun psubj})$	27.98	
Proper noun	Object	Subject	$p(\text{correct} \mid \text{SO, m. pnoun pobj})$	23.25	n.s.
			$p(\text{correct} \mid \text{SO, f. pnoun pobj})$	21.06	

Table 10: Conditional probabilities for the correctness of predictions given the subject’s or the object’s gender. The rightmost column indicates a significant difference between compared proportions at the 99% or 95% confidence level, or no significance (n.s.), using a z-test for the equality of two proportions.

Premise	Er _{NOM-SG-M}	warnt _{SG}	<u>den</u> _{ACC-SG-M}	Gast
	He _{NOM-SG-M}	warns _{SG}	<u>the</u> _{ACC-SG-M}	guest
	<i>He</i>	warns	<i>the</i>	<i>guest</i>
H1-SO (NE)	Der _{NOM-SG-M}	Gast	warnt _{SG}	<u>ihn</u> _{ACC-SG-M}
	The _{NOM-SG-M}	guest	warns _{SG}	<u>him</u> _{ACC-SG-M}
	<i>The</i>	<i>guest</i>	warns	<i>him</i>
H2-OS (E)*	<u>Den</u> _{ACC-SG-M}	Gast	warnt _{SG}	er _{NOM-SG-M}
	<u>The</u> _{ACC-SG-M}	guest	warns _{SG}	he _{NOM-SG-M}
	<i>He</i>	warns	<i>the</i>	<i>guest</i>

Table 11: Examples of WOGLI-p-subject pairs. Just as in WOGLI, the entailed hypothesis has a marked word order.

Premise	Ein _{NOM-SG-M}	Richter	gratuiert _{SG}	<u>diesen</u> _{DAT-PL-M}	<u>Beratern</u>
	A _{NOM-SG-M}	judge	congratulates _{SG}	<u>these</u> _{DAT-PL-M}	consultants
	<i>A</i>	<i>judge</i>	congratulates	<i>these</i>	<i>consultants</i>
H1-SO (NE)	Diese _{NOM-PL-M}	Berater	gratulieren _{PL}	<u>einem</u> _{DAT-SG-M}	Richter
	These _{NOM-PL-M}	consultants	congratulate _{PL}	<u>a</u> _{DAT-SG-M}	judge
	<i>The</i>	<i>consultants</i>	congratulate	<i>a</i>	<i>judge</i>
H2-OS (E)*	<u>Diesen</u> _{DAT-PL-M}	<u>Beratern</u>	gratuiert _{SG}	ein _{NOM-SG-M}	Richter
	<u>These</u> _{DAT-PL-M}	consultants	congratulates _{SG}	a _{NOM-SG-M}	judge
	<i>A</i>	<i>judge</i>	congratulates	<i>these</i>	<i>consultants</i>

Table 12: Examples of WOGLI-dative pairs. Just as in WOGLI, the entailed hypothesis has a marked word order.

Premise	Die _{NOM-PL-F}	Kellnerinnen	geben _{PL}	<u>einem</u> _{DAT-SG-M}	Händler	den	Kuchen
	The _{NOM-PL-F}	waitresses	give _{PL}	<u>a</u> _{DAT-SG-M}	merchant	the	cake
	<i>The</i>	<i>waitresses</i>	give	<i>the</i>	<i>cake</i>	<i>to a</i>	<i>merchant</i>
H1-SiO (NE)	Ein _{NOM-SG-M}	Händler	gibt _{SG}	<u>den</u> _{DAT-PL-SG}	Kellnerinnen	den	Kuchen
	A _{NOM-SG-M}	merchant	gives _{SG}	<u>the</u> _{DAT-PL-SG}	waitresses	the	cake
	<i>A</i>	<i>merchant</i>	gives	<i>the</i>	<i>cake</i>	<i>to the</i>	<i>waitresses</i>
H2-iOS (E)*	<u>Einem</u> _{DAT-SG-M}	Händler	geben _{PL}	die _{NOM-PL-F}	Kellnerinnen	den	Kuchen
	<u>A</u> _{DAT-SG-M}	merchant	give _{PL}	the _{NOM-PL-F}	waitresses	the	cake
	<i>The</i>	<i>waitresses</i>	give	<i>the</i>	<i>cake</i>	<i>to a</i>	<i>merchant</i>

Table 13: Examples of WOGLI-ditransitive pairs. Just as in WOGLI, the entailed hypothesis has a marked word order.