

# UnNatural Language Inference

Koustuv Sinha<sup>1,2,3</sup>, Prasanna Parthasarathi<sup>1,2</sup>, Joelle Pineau<sup>1,2,3</sup> and Adina Williams<sup>3</sup>

<sup>1</sup> School of Computer Science, McGill University, Canada

<sup>2</sup> Montreal Institute of Learning Algorithms (Mila), Canada

<sup>3</sup> Facebook AI Research (FAIR)

{koustuv.sinha, prasanna.parthasarathi, jpineau, adinawilliams}

@{mail.mcgill.ca, mail.mcgill.ca, cs.mcgill.ca, fb.com}

## Abstract

Recent investigations into the inner-workings of state-of-the-art large-scale pre-trained Transformer-based Natural Language Understanding (NLU) models indicate that they appear to know humanlike syntax, at least to some extent. We provide novel evidence that complicates this claim: we find that state-of-the-art Natural Language Inference (NLI) models assign the same labels to permuted examples as they do to the original, i.e. they are largely invariant to random word-order permutations. This behavior notably differs from that of humans; we struggle with ungrammatical sentences. To measure the severity of this issue, we propose a suite of metrics and investigate which properties of particular permutations lead models to be word-order invariant. In the MNLI dataset, for example, we find almost all (98.7%) examples contain at least one permutation which elicits the gold label. Models are sometimes even able to assign gold labels to permutations that they originally failed to predict correctly. We provide a comprehensive empirical evaluation of this phenomenon, and further show that this issue exists for both Transformers and pre-Transformer RNN / ConvNet based encoders, as well as across multiple languages (English and Mandarin Chinese). Our code and data are available at <https://github.com/facebookresearch/unlu>.

## 1 Introduction

Of late, large scale pre-trained Transformer-based (Vaswani et al., 2017) models—such as RoBERTa (Liu et al., 2019), BART (Lewis et al., 2020), and GPT-2 and -3 (Radford et al., 2019; Brown et al., 2020)—have exceeded recurrent neural networks’ performance on many NLU tasks (Wang et al., 2018, 2019). Several papers have even suggested that Transformers pretrained on a language modeling (LM) objective can capture syntactic informa-

Premise	Hypothesis	Predicted Label
Boats in daily use lie within feet of the fashionable bars and restaurants.	There are boats close to bars and restaurants.	E
restaurants and use feet of fashionable lie the in Boats within bars daily .	bars restaurants are There and to close boats .	E
He and his associates weren’t operating at the level of metaphor.	He and his associates were operating at the level of the metaphor.	C
his at and metaphor the of were He operating associates n’t level .	his the and metaphor level the were He at associates operating of .	C

Table 1: Examples from the MNLI Matched development set. Both the original example and the permuted one elicit the same classification label (entailment and contradiction respectively) from RoBERTa (large). A simple demo is provided in an associated [Google Colab notebook](#).

tion (Hewitt and Manning, 2019; Jawahar et al., 2019; Warstadt and Bowman, 2020; Wu et al., 2020), with their self-attention layers being capable of surprisingly effective learning (Rogers et al., 2020). In this work, we question such claims that current models “know syntax”.

Since there are many ways to investigate “syntax”, we must be clear on what we mean by the term. Knowing the syntax of a sentence means being sensitive to the *order of the words* in that sentence (among other things). Humans are sensitive to word order, so clearly, “language is not merely a bag of words” (Harris, 1954, p.156). Moreover, it is easier for us to identify or recall words presented in canonical orders than in disordered, ungrammatical sentences; this phenomenon is called the “*sentence superiority effect*” (Cattell 1886; Scheerer 1981; Toyota 2001; Baddeley et al. 2009; Snell and Grainger 2017, 2019; Wen et al. 2019, i.a.). In our estimation then, if one wants to claim that a model “knows syntax”, then they should minimally show that the model is sensitive to word order (at least

for e.g. English or Mandarin Chinese).

Generally, knowing the syntax of a sentence is taken to be a prerequisite for understanding what that sentence means (Heim and Kratzer, 1998). Models should have to know the syntax first then, if performing any particular NLU task that genuinely requires a humanlike understanding of meaning (cf. Bender and Koller 2020). Thus, if our models are as good at NLU as our current evaluation methods suggest, we should expect them to be sensitive to word order (see Table 1). We find, based on a suite of permutation metrics, that they are not.

We focus here on textual entailment, one of the hallmark tasks used to measure how well models understand language (Condoravdi et al., 2003; Dagan et al., 2005). This task, often also called Natural Language Inference (NLI; Bowman et al. 2015, i.a.), typically consists of two sentences: a premise and a hypothesis. The objective is to predict whether the premise entails the hypothesis, contradicts it, or is neutral with respect to it. We find rampant word order insensitivity in purportedly high performing NLI models. For nearly all premise-hypothesis pairs, **there are many permuted examples that fool the models** into providing the correct prediction. In case of MNLI, for example, the current state-of-the-art of 90.5% can be increased to **98.7%** merely by permuting the word order of test set examples. We even find drastically increased cross-dataset generalization when we reorder words. This is not just a matter of chance—we show that the model output probabilities are significantly different from uniform.

We verify our findings with three popular English NLI datasets—SNLI (Bowman et al., 2015), MultiNLI (Williams et al., 2018b) and ANLI (Nie et al., 2020)—and one Chinese one, OCNLI (Hu et al., 2020a). It is thus less likely that our findings result from some quirk of English or a particular tokenization strategy. We also observe the effect for various transformer architectures pre-trained on language modeling (BERT, RoBERTa, DistilBERT), and non-transformers, including a ConvNet, an In-Sent model, and a BiLSTM.

Our contributions are as follows: (i) we propose a suite of metrics (*Permutation Acceptance*) for measuring model insensitivity to word order (§3), (ii) we construct multiple permuted test datasets for measuring NLI model performance at a large scale (§5), (iii) we show that NLI models focus on words more than word order, but can partially

reconstruct syntactic information from words alone (§6), (iv) we show the problem persists on out-of-domain data, (v) we show that humans struggle with UnNatural Language Inference, underscoring the non-humanlikeness of SOTA models (§7), (vi) finally, we explore a simple maximum entropy-based method (§8) to encourage models not to accept permuted examples.

## 2 Related Work

Researchers in NLP have realized the importance of syntactic structure in neural networks going back to Tabor (1994). An early hand annotation effort on PASCAL RTE (Dagan et al., 2006) suggested that “syntactic information alone was sufficient to make a judgment” for roughly one third of examples (Vanderwende and Dolan, 2005). Anecdotally, large generative language models like GPT-2 or -3 exhibit a seemingly humanlike ability to generate fluent and grammatical text (Goldberg, 2019; Wolf, 2019). However, the jury is still out as to whether transformers genuinely acquire syntax.

**Models appear to have acquired syntax.** When researchers have peeked inside Transformer LM’s pretrained representations, familiar syntactic structure (Hewitt and Manning, 2019; Jawahar et al., 2019; Lin et al., 2019; Warstadt and Bowman, 2020; Wu et al., 2020), or a familiar order of linguistic operations (Jawahar et al., 2019; Tenney et al., 2019), has appeared. There is also evidence, notably from agreement attraction phenomena (Linzen et al., 2016) that transformer-based models pretrained on LM do acquire some knowledge of natural language syntax (Gulordava et al., 2018; Chrupala and Alishahi, 2019; Jawahar et al., 2019; Lin et al., 2019; Manning et al., 2020; Hawkins et al., 2020; Linzen and Baroni, 2021). Results from other phenomena (Warstadt and Bowman, 2020) such as NPI licensing (Warstadt et al., 2019a) lend additional support. The claim that LMs acquire some syntactic knowledge has been made not only for transformers, but also for convolutional neural nets (Bernardy and Lappin, 2017), and RNNs (Gulordava et al., 2018; van Schijndel and Linzen, 2018; Wilcox et al., 2018; Zhang and Bowman, 2018; Prasad et al., 2019; Ravfogel et al., 2019)—although there are many caveats (e.g., Ravfogel et al. 2018; White et al. 2018; Davis and van Schijndel 2020; Chaves 2020; Da Costa and Chaves 2020; Kodner and Gupta 2020).

**Models appear to struggle with syntax.** Several works have cast doubt on the extent to which NLI models in particular know syntax (although each work adopts a slightly different idea of what “knowing syntax” entails). For example, McCoy et al. (2019) argued that the knowledge acquired by models trained on NLI (for at least some popular datasets) is actually not as syntactically sophisticated as it might have initially seemed; some transformer models rely mainly on simpler, non-humanlike heuristics. In general, transformer LM performance has been found to be patchy and variable across linguistic phenomena (Dasgupta et al., 2018; Naik et al., 2018; An et al., 2019; Ravichander et al., 2019; Jeretic et al., 2020). This is especially true for syntactic phenomena (Marvin and Linzen, 2018; Hu et al., 2020b; Gauthier et al., 2020; McCoy et al., 2020; Warstadt et al., 2020), where transformers are, for some phenomena and settings, worse than RNNs (van Schijndel et al., 2019). From another angle, many have explored architectural approaches for increasing a network’s sensitivity to syntactic structure (Chen et al., 2017; Li et al., 2020). Williams et al. (2018a) showed that learning jointly to perform NLI and to parse resulted in parse trees that match no popular syntactic formalisms. Furthermore, models trained explicitly to differentiate acceptable sentences from unacceptable ones (i.e., one of the most common syntactic tests used by linguists) have, to date, come nowhere near human performance (Warstadt et al., 2019b).

**Insensitivity to Perturbation.** Most relatedly, several concurrent works (Pham et al., 2020; Alleman et al., 2021; Gupta et al., 2021; Sinha et al., 2021; Parthasarathi et al., 2021) investigated the effect of word order permutations on transformer NNs. Pham et al. (2020) is very nearly a proper subset of our work except for investigating additional tasks (i.e. from the GLUE benchmark of Wang et al. 2018) and performing a by-layer-analysis. Gupta et al. (2021) also relies on the GLUE benchmark, but additionally investigates other types of “destructive” perturbations. Our contribution differs from these works in that we additionally include the following: we (i) outline theoretically-informed predictions for how models *should be expected* to react to permuted input (we outline a few options), (ii) show that permuting can “flip” an incorrect prediction to a correct one, (iii) show that the problem isn’t specific to Transformers, (iv) show that the problem persists on out of domain data, (v) offer

a suite of flexible metrics, and (vi) analyze *why* models might be accepting permutations (BLEU and POS-tag neighborhood analysis). Finally, we replicate our findings in another language. While our work (and Pham et al.; Gupta et al.) only permutes data during fine-tuning and/or evaluation, recently Sinha et al. explored the sensitivity during pre-training, and found that models trained on n-gram permuted sentences perform remarkably close to regular MLM pre-training. In the context of generation, Parthasarathi et al. (2021) crafted linguistically relevant perturbations (on the basis of part-of-speech tagging and dependency parsing) to evaluate whether permutation hinders automatic machine translation models. Relatedly, but not for translation, Alleman et al. (2021) investigated a smaller inventory of perturbations with emphasis on phrasal boundaries and the effects of n-gram perturbations on different layers in the network.

**NLI Models are very sensitive to words.** NLI models often over-attend to particular words to predict the correct answer (Gururangan et al., 2018; Clark et al., 2019). Wallace et al. (2019) show that some short sequences of non-human-readable text can fool many NLU models, including NLI models trained on SNLI, into predicting a specific label. In fact, Ettinger (2020) observed that for one of three test sets, BERT loses some accuracy in word-perturbed sentences, but that there exists a subset of examples for which BERT’s accuracy remains intact. If performance isn’t affected (or if permutation helps, as we find it does in some cases), it suggests that these state-of-the-art models actually perform somewhat similarly to bag-of-words models (Blei et al., 2003; Mikolov et al., 2013).

### 3 Our Approach

As we mentioned, linguists generally take syntactic structure to be necessary for humans to know what sentences mean. Many also find the NLI task to a very promising approximation of human natural language understanding, in part because it is rooted in the tradition of logical entailment. In the spirit of propositional logic, sentence meaning is taken to be truth-conditional (Frege, 1948; Montague, 1970; Chierchia and McConnell-Ginet, 1990; Heim and Kratzer, 1998). That is to say that understanding a sentence is equivalent to knowing the actual conditions of the world under which the sentences would be (judged) true (Wittgenstein, 1922). If grammatical sentences are required for

sentential inference, as per a truth conditional approach (Montague, 1970), then permuted sentences should be meaningless. Put another way, the meanings of highly permuted sentences (if they exist) are not propositions, and thus those sentences don't have truth conditions. Only from their truth conditions of sentences can we tell if a sentence entails another. In short, the textual entailment task is technically undefined in our "unnatural" setting.

Since existing definitions don't immediately extend to UnNatural Language Inference (UNLI), we outline several hypothetical *systematic* ways that a model might perform, had it been sensitive to word order. We hypothesize two models that operate on the first principles of NLI, and one that doesn't. In the first case, Model A deems permuted sentences meaningless (devoid of truth values), as formal semantic theories of human language would predict. Thus, it assigns "neutral" to every permuted example. Next, Model B does not deem permuted sentences meaningless, and attempts to understand them. Humans find understanding permuted sentences difficult (see our human evaluations in §7). Model B could also similarly struggle to decipher the meaning, and just equally sample labels for each example (i.e., assigns equal probability mass to the outcome of each label). Finally, we hypothesize a non-systematic model, Model C, which attempts to treat permuted sentences as though they weren't permuted at all. This model could operate similarly as bag-of-words (BOW), and thus always assign the same label to the permuted examples as it would to the un-permuted examples. If the model failed to assign the original gold label to the original unpermuted examples, it will also fail to assign the original gold label to its permutations; it will never get higher accuracy on permuted examples than on unpermuted ones.

We find in our experiments that the state-of-the-art Transformer-based NLI models (as well as pre-Transformer class of models) do not perform like any of the above hypothetical models. They perform closest to Model C, but are, in some cases, actually able to achieve *higher* accuracy on permuted examples. To better quantitatively describe this behaviour, we introduce our suite of **Permutation Acceptance** metrics that enable us to quantify how accepting models are of permuted sentences.

## 4 Methods

**Constructing the permuted dataset.** For a given dataset  $D$  having splits  $D_{\text{train}}$  and  $D_{\text{test}}$ , we first train an NLI model  $M$  on  $D_{\text{train}}$  to achieve comparable accuracy to what was reported in the original papers. We then construct a randomized version of  $D_{\text{test}}$ , which we term as  $\hat{D}_{\text{test}}$  such that: for each example  $(p_i, h_i, y_i) \in D_{\text{test}}$  (where  $p_i$  and  $h_i$  are the premise and hypothesis sentences of the example respectively and  $y_i$  is the gold label), we use a permutation operator  $\mathcal{F}$  that returns a list  $(\hat{P}_i, \hat{H}_i)$  of  $q$  permuted sentences ( $\hat{p}_i$  and  $\hat{h}_i$ ), where  $q$  is a hyperparameter.  $\mathcal{F}$  essentially permutes all positions of the words in a given sentence (i.e., either in premise or hypothesis) with the restriction that *no words maintain their original position*. In our initial setting, we do not explicitly control the placement of the words relative to their original neighbors, but we analyze clumping effects in §5.  $\hat{D}_{\text{test}}$  now consists of  $|D_{\text{test}}| \times q$  examples, with  $q$  different permutations of hypothesis and premise for each original test example pair. If a sentence  $S$  (e.g.,  $h_i$ ) contains  $w$  words, then the total number of available permutations of  $S$  are  $(w-1)!$ , thus making the output of  $\mathcal{F}$  a list of  $\binom{(w-1)!}{q}$  permutations in this case. For us, the space of possible outputs is larger, since we permute  $p_i$  and  $h_i$  separately (and ignore examples for which any  $|S| \leq 5$ ).

**Defining Permutation Acceptance.** The choice of  $q$  naturally allows us to analyze a statistical view of the predictability of a model on the permuted sentences. To that end, we define the following notational conventions. Let  $\mathcal{A}$  be the original accuracy of a given model  $M$  on a dataset  $D$ , and  $c$  be the number of examples in a dataset which are marked as correct according to the standard formulation of accuracy for the original dataset (i.e., they are assigned the ground truth label). Typically  $\mathcal{A}$  is given by  $\frac{c}{|D_{\text{test}}|}$  or  $\frac{c}{|D_{\text{dev}}|}$ .

Let  $\text{Pr}_M(\hat{P}_i, \hat{H}_i)_{\text{cor}}$  then be the percentage of  $q$  permutations of an example  $(p_i, h_i)$  assigned the ground truth label  $y_i$  by  $M$ :

$$\text{Pr}_M(\hat{P}_i, \hat{H}_i)_{\text{cor}} = \frac{1}{q} \sum_{(\hat{p}_j \in \hat{P}_i, \hat{h}_j \in \hat{H}_i)} ((M(\hat{p}_j, \hat{h}_j) = y_i) \rightarrow 1) \quad (1)$$

To get an overall summary score, we let  $\Omega_x$  be the percentage of examples  $(p_i, h_i) \in D_{\text{test}}$  for which  $\text{Pr}_M(\hat{P}_i, \hat{H}_i)_{\text{cor}}$  exceeds a predetermined threshold  $0 < x < 1$ . Concretely, a given example will count



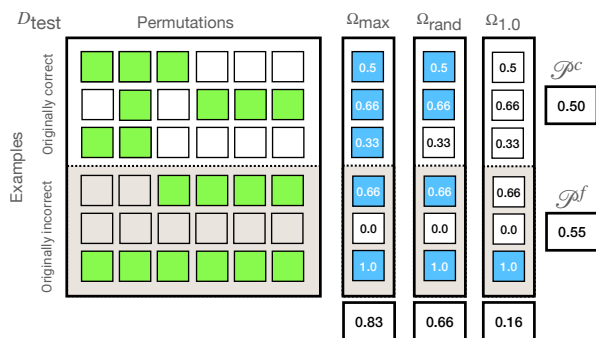


Figure 1: Graphical representation of the Permutation Acceptance class of metrics. Given a sample test set  $D_{\text{test}}$  with six examples, three of which originally predicted correctly (model predicts gold label), three incorrectly (model fails to predict gold label), with  $n = 6$  permutations,  $\Omega_{\text{max}}, \Omega_{\text{rand}}, \Omega_{1.0}, \mathcal{P}^c$  and  $\mathcal{P}^f$  are provided. Green boxes indicate permutations accepted by the model. Blue boxes mark examples that crossed each threshold and were used to compute the corresponding metric.

as correct according to  $\Omega_x$  if more than  $x$  percent of its permutations ( $\hat{P}_i$  and  $\hat{H}_i$ ) are assigned  $y_i$  by the model  $M$ . Mathematically,

$$\Omega_x = \frac{1}{|D_{\text{test}}|} \sum_{(p_i, h_i) \in D_{\text{test}}} (\Pr_M((\hat{P}_i, \hat{H}_i)_{\text{cor}} > x) \rightarrow 1). \quad (2)$$

There are two specific cases of  $\Omega_x$  that we are most interested in. First, we define  $\Omega_{\text{max}}$  or the **Maximum Accuracy**, where  $x = 1/|D_{\text{test}}|$ . In short,  $\Omega_{\text{max}}$  gives the percentage of examples  $(p_i, h_i) \in D_{\text{test}}$  for which there is *at least one* permutation  $(\hat{p}_j, \hat{h}_j)$  that model  $M$  assigns the gold label  $y_i$ <sup>1</sup>. Second, we define  $\Omega_{\text{rand}}$ , or **Random Baseline Accuracy**, where  $x = 1/m$  or chance probability (for balanced  $m$ -way classification, where  $m = 3$  in NLI). This metric is less stringent than  $\Omega_{\text{max}}$ , as it counts an example if at least *one third* of its permutations are assigned the gold label (hence provides a lower-bound relaxation). See Figure 1 for a graphical representation of  $\Omega_x$ .

We also define  $D^f$  to be the list of examples originally marked incorrect according to  $\mathcal{A}$ , but are now deemed correct according  $\Omega_{\text{max}}$ .  $D^c$  is the list of examples originally marked correct according to  $\mathcal{A}$ . Thus, we should expect  $D^f < D^c$  for models that have high accuracy. Additionally, we define  $\mathcal{P}^c$  and  $\mathcal{P}^f$ , as the dataset average percentage of permutations which predicted the gold label, when the

<sup>1</sup>Theoretically,  $\Omega_{\text{max}} \rightarrow 1$  if the number of permutations  $q$  is large. Thus, in our experiments we set  $q = 100$ .

examples were originally correct ( $D^c$ ) and when the examples were originally incorrect ( $D^f$ ) as per  $\mathcal{A}$  (hence, flipped) respectively.

$$\mathcal{P}^c = \frac{1}{|D^c|} \sum_{i=0}^{|D^c|} M(\hat{P}_i, \hat{H}_i)_{\text{cor}} \quad (3)$$

$\mathcal{P}^f$  is defined similarly by replacing  $D^c$  by  $D^f$ . Note that for a classic BOW model,  $\mathcal{P}^c = 100$  and  $\mathcal{P}^f = 0$ , because it would rely on the words alone (not their order) to make its classification decision. Since permuting removes no words, BOW models should come to the same decisions for permuted examples as for the originals.

## 5 Results

We present results for two types of models: (a) Transformer-based models and (b) Non-Transformer Models. In (a), we investigate the state-of-the-art pre-trained models such as RoBERTa-Large (Liu et al., 2019), BART-Large (Lewis et al., 2020) and DistilBERT (Sanh et al., 2020). For (b) we consider several recurrent and convolution based neural networks, such as InferSent (Conneau et al., 2017), Bidirectional LSTM (Collobert and Weston, 2008) and ConvNet (Zhao et al., 2015). We train all models on MNLI, and evaluate on in-distribution (SNLI and MNLI) and out-of-distribution datasets (ANLI). We independently verify results of (a) using both our fine-tuned model using HuggingFace Transformers (Wolf et al., 2020) and pre-trained checkpoints from FairSeq (Ott et al., 2019) (using PyTorch Model Hub). For (b), we use the InferSent codebase. We sample  $q = 100$  permutations for each example in  $D_{\text{test}}$ , and use 100 seeds for each of those permutations to ensure full reproducibility. We drop examples from test sets where we are unable to compute *all unique* randomizations, typically these are examples with sentences of length of less than 6 tokens.<sup>2</sup>

**Models accept many permuted examples.** We find  $\Omega_{\text{max}}$  is very high for models trained and evaluated on MNLI (in-domain generalization), reaching **98.7%** on MNLI dev. and test sets (in RoBERTa, compared to  $\mathcal{A}$  of 90.6% (Table 2). Recall, human accuracy is approximately 92% on MNLI dev., Nangia and Bowman 2019). This shows that there exists at least one permutation (usually many more)

<sup>2</sup>Code, data, and model checkpoints will be available at <https://github.com/facebookresearch/unli>.

Model	Eval. Dataset	$\mathcal{A}$	$\Omega_{\max}$	$\mathcal{P}^c$	$\mathcal{P}^f$	$\Omega_{\text{rand}}$
RoBERTa-Large	MNLI_m.dev	0.906	0.987	0.707	0.383	0.794
	MNLI_mm.dev	0.901	0.987	0.707	0.387	0.790
	SNLI.dev	0.879	0.988	0.768	0.393	0.826
	SNLI_test	0.883	0.988	0.760	0.407	0.828
	A1*	0.456	0.897	0.392	0.286	0.364
	A2*	0.271	0.889	0.465	0.292	0.359
	A3*	0.268	0.902	0.480	0.308	0.397
Mean	0.652	0.948	0.611	<b>0.351</b>	0.623	
BART-Large	MNLI_m.dev	0.902	0.989	0.689	0.393	0.784
	MNLI_mm.dev	0.900	0.986	0.695	0.399	0.788
	SNLI.dev	0.886	0.991	0.762	0.363	0.834
	SNLI_test	0.888	0.990	0.762	0.370	0.836
	A1*	0.455	0.894	0.379	0.295	0.374
	A2*	0.316	0.887	0.428	0.303	0.397
	A3*	0.327	0.931	0.428	0.333	0.424
Mean	<b>0.668</b>	<b>0.953</b>	0.592	<b>0.351</b>	<b>0.634</b>	
DistilBERT	MNLI_m.dev	0.800	0.968	0.775	0.343	0.779
	MNLI_mm.dev	0.811	0.968	0.775	0.346	0.786
	SNLI.dev	0.732	0.956	0.767	0.307	0.731
	SNLI_test	0.738	0.950	0.770	0.312	0.725
	A1*	0.251	0.750	0.511	0.267	0.300
	A2*	0.300	0.760	0.619	0.265	0.343
	A3*	0.312	0.830	0.559	0.259	0.363
Mean	0.564	0.883	<b>0.682</b>	0.300	0.575	
InferSent	MNLI_m.dev	0.658	0.904	0.842	0.359	0.712
	MNLI_mm.dev	0.669	0.905	0.844	0.368	0.723
	SNLI.dev	0.556	0.820	0.821	0.323	0.587
	SNLI_test	0.560	0.826	0.824	0.321	0.600
	A1*	0.316	0.669	0.425	0.395	0.313
	A2*	0.310	0.662	0.689	0.249	0.330
	A3*	0.300	0.677	0.675	0.236	0.332
Mean	<b>0.481</b>	0.780	0.731	<b>0.322</b>	0.514	
ConvNet	MNLI_m.dev	0.631	0.926	0.773	0.340	0.684
	MNLI_mm.dev	0.640	0.926	0.782	0.343	0.694
	SNLI.dev	0.506	0.819	0.813	0.339	0.597
	SNLI_test	0.501	0.821	0.809	0.341	0.596
	A1*	0.271	0.708	0.648	0.218	0.316
	A2*	0.307	0.725	0.703	0.224	0.356
	A3*	0.306	0.798	0.688	0.234	0.388
Mean	0.452	<b>0.817</b>	<b>0.745</b>	0.291	0.519	
BiLSTM	MNLI_m.dev	0.662	0.925	0.800	0.351	0.711
	MNLI_mm.dev	0.681	0.924	0.809	0.344	0.724
	SNLI.dev	0.547	0.860	0.762	0.351	0.598
	SNLI_test	0.552	0.862	0.771	0.363	0.607
	A1*	0.262	0.671	0.648	0.271	0.340
	A2*	0.297	0.728	0.672	0.209	0.328
	A3*	0.304	0.731	0.656	0.219	0.331
Mean	0.472	0.814	0.731	0.301	<b>0.520</b>	

Table 2: Statistics for Transformer-based models trained on MNLI corpus (Williams et al., 2018b). The highest values are bolded (red indicates the model most insensitive to permutation) per metric and per model class (Transformers and non-Transformers). A1\*, A2\* and A3\* refer to the ANLI dev. sets (Nie et al., 2020).

Model	$\mathcal{A}$	$\Omega_{\max}$	$\mathcal{P}^c$	$\mathcal{P}^f$	$\Omega_{\text{rand}}$
RoBERTa-Large	<b>0.784</b>	<b>0.988</b>	0.726	<b>0.339</b>	<b>0.773</b>
InferSent	0.573	0.931	0.771	0.265	0.615
ConvNet	0.407	0.752	<b>0.808</b>	0.199	0.426
BiLSTM	0.566	0.963	0.701	0.271	0.611

Table 3: Results on evaluation on OCNLI Dev set. All models are trained on OCNLI corpus (Hu et al., 2020a). Bold marks the highest value per metric (red shows the model is insensitive to permutation).

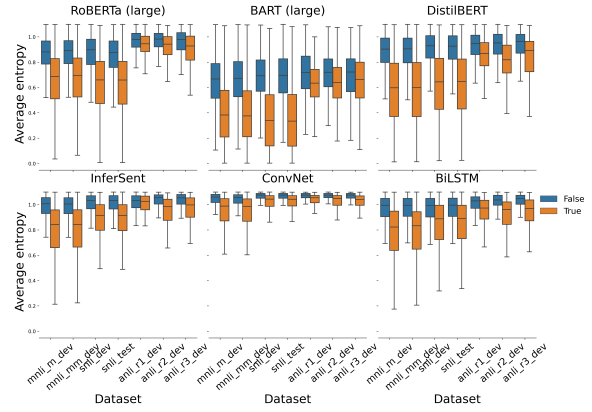


Figure 2: Average entropy of model confidences on permutations that yielded the correct results for Transformer-based models (top) and Non-Transformer-based models (bottom). Results are shown for  $D^c$  (orange) and  $D^f$  (blue). The boxes show the quartiles of the entropy distributions.

for almost all examples in  $D_{\text{test}}$  such that model  $M$  predicts the gold label. We also observe high  $\Omega_{\text{rand}}$  at 79.4%, showing that there are many examples for which the models outperform even a random baseline in accepting permuted sentences (see Appendix E for more  $\Omega$  values.)

Evaluating out-of-domain generalization with ANLI dataset splits resulted in an  $\Omega_{\max}$  value that is notably higher than  $\mathcal{A}$  (89.7%  $\Omega_{\max}$  for RoBERTa compared to 45.6%  $\mathcal{A}$ ). As a consequence, we encounter many *flips*, i.e., examples where the model is unable to predict the gold label, but at least one permutation of that example is able to. However, recall this analysis expects us to know the gold label upfront, so this test can be thought of as running a word-order probe test on the model until the model predicts the gold label (or give up by exhausting our set of  $q$  permutations). For out-of-domain generalization,  $\Omega_{\text{rand}}$  decreases considerably (36.4%  $\Omega_{\text{rand}}$  on A1), which means fewer permutations are accepted by the model. Next, recall that a classic bag-of-words model would have  $\mathcal{P}^c = 100$  and  $\mathcal{P}^f = 0$ . No model performs strictly like a classic bag of words although they do perform somewhat BOW-like ( $\mathcal{P}^c \gg \mathcal{P}^f$  for all test splits, Figure 5). We find this BOW-likeness to be higher for certain non-Transformer models, (InferSent) as they exhibit higher  $\mathcal{P}^c$  (84.2% for InferSent compared to 70.7% for RoBERTa on MNLI).

**Models are very confident.** The phenomenon we observe would be of less concern if the correct label prediction was just an outcome of chance,

which could occur when the entropy of the log probabilities of the model output is high (suggesting uniform probabilities on entailment, neutral and contradiction labels, recall Model B from §3). We first investigate the model probabilities for the Transformer-based models on the permutations that lead to the correct answer in Figure 2. We find overwhelming evidence that model confidences on in-distribution datasets (MNLI, SNLI) are highly skewed, resulting in low entropy, and it varies among different model types. BART proves to be the most skewed Transformer-based model. This skewness is not a property of model capacity, as we observe DistilBERT log probabilities to have similar skewness as RoBERTa (large) model, while exhibiting lower  $\mathcal{A}$ ,  $\Omega_{\max}$ , and  $\Omega_{\text{rand}}$ .

For non-Transformers whose accuracy  $\mathcal{A}$  is lower, the  $\Omega_{\max}$  achieved by these models are also predictably lower. We observe roughly the same relative performance in the terms of  $\Omega_{\max}$  (Figure 5 and Appendix Table 2) and Average entropy (Figure 2). However, while comparing the averaged entropy of the model predictions, it is clear that there is some benefit to being a worse model—non-Transformer models are not as overconfident on randomized sentences as Transformers are. High confidence of Transformer models can be attributed to the *overthinking* phenomenon commonly observed in deep neural networks (Kaya et al., 2019) and BERT-based models (Zhou et al., 2020).

**Similar artifacts in Chinese NLU.** We extended the experiments to the Original Chinese NLI dataset (Hu et al., 2020a, OCNLI), and re-used the pre-trained RoBERTa-Large and InferSent (non-Transformer) models on OCNLI. Our findings are similar to the English results (Table 3), thereby suggesting that the phenomenon is not just an artifact of English text or tokenization.

**Other Results.** We investigated the effect of sentence length (which correlates with number of possible permutations; Appendix A), and hypothesis-only randomization (models exhibit similar phenomenon even when only hypothesis is permuted; Appendix C).

## 6 Analyzing Syntactic Structure Associated with Tokens

A natural question to ask following our findings: what is it about particular permutations that leads models to accept them? Since the permutation oper-

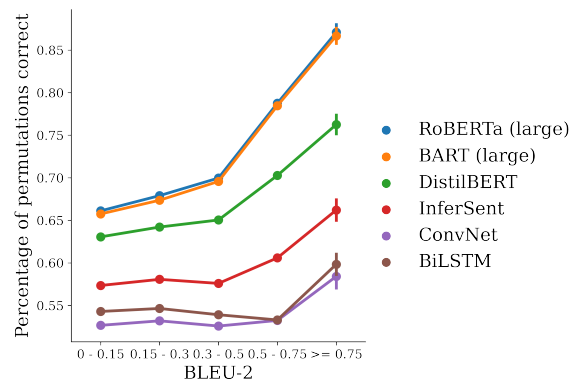


Figure 3: BLEU-2 score versus acceptability of permuted sentences across all test datasets. RoBERTa and BART performance is similar but differs considerably from the performance of non-Transformer-based models, such as InferSent and ConvNet.

ation is drastic and only rarely preserves local word relations, we first investigate whether there exists a relationship between Permutation Acceptance scores and local word order preservation. Concretely, we compare bi-gram word overlap (BLEU-2) with the percentage of permutations that are deemed correct (Figure 3).<sup>3</sup> Although the probability of a permuted sentence to be predicted correctly does appear to track BLEU-2 score (Figure 3), the percentage of examples which were assigned the gold label by the Transformer-based models is still higher than we would expect from permutations with lower BLEU-2 (66% for the lowest BLEU-2 range of 0 – 0.15), suggesting preserved relative word order alone cannot explain the high permutation acceptance rates.

Thus, we find that local order preservation does correlate with Permutation Acceptance, but it doesn’t fully explain the high Permutation Acceptance scores. We now further ask whether  $\Omega$  is related to a more abstract measure of local word relations, i.e., part-of-speech (POS) neighborhood.

Many syntactic formalisms, like Lexical Functional Grammar (Kaplan and Bresnan, 1995; Bresnan et al., 2015, LFG), Head-drive Phrase Structure Grammar (Pollard and Sag, 1994, HPSG) or Lexicalized Tree Adjoining Grammar (Schabes et al., 1988; Abeille, 1990, LTAG), are “lexicalized”, i.e.,

<sup>3</sup>We observe, due to our permutation process, the maximum BLEU-3 and BLEU-4 scores are negligibly low ( $< 0.2$  BLEU-3 and  $< 0.1$  BLEU-4), already calling into question the hypothesis that n-grams are the sole explanation for our finding. Because of this, we only compare BLEU-2 scores. Detailed experiments on specially constructed permutations that cover the entire range of BLEU-3 and BLEU-4 is provided in Appendix D.

individual words or morphemes bear syntactic features telling us which other words they can combine with. For example, “buy” could be associated with (at least) two lexicalized syntactic structures, one containing two noun phrases (as in *Kim bought cheese*), and another with three (as in *Lee bought Logan cheese*). We speculate that our NLI models might accept permuted examples at high rates, because they are (perhaps noisily) reconstructing the original sentence from abstract, word-anchored information about common neighbors.

To test this, we POS-tagged  $D_{\text{train}}$  using 17 Universal Part-of-Speech tags (using spaCy, Honnibal et al. 2020). For each  $w_i \in S_i$ , we compute the occurrence probability of POS tags on tokens in the *neighborhood* of  $w_i$ . The neighborhood is specified by the radius  $r$  (a symmetrical window  $r$  tokens from  $w_i \in S_i$  to the left and right). We denote this sentence level probability of neighbor POS tags for a word  $w_i$  as  $\psi_{\{w_i, S_i\}}^r \in \mathcal{R}^{17}$  (see an example in Figure 7 in the Appendix). Sentence-level word POS neighbor scores can be averaged across  $D_{\text{train}}$  to get a type level score  $\psi_{\{w_i, D_{\text{train}}\}}^r \in \mathcal{R}^{17}, \forall w_i \in D_{\text{train}}$ . Then, for a sentence  $S_i \in D_{\text{test}}$ , for each word  $w_i \in S_i$ , we compute a **POS mini-tree overlap score**:

$$\beta_{\{w_i, S_i\}}^k = \frac{1}{k} \left| \underset{\text{argmax}_k \psi_{\{w_i, D_{\text{train}}\}}^r}{\cap} \underset{\text{argmax}_k \psi_{\{w_i, S_i\}}^r}{\cap} \right| \quad (4)$$

Concretely,  $\beta_{\{w_i, S_i\}}^k$  computes the overlap of top- $k$  POS tags in the neighborhood of a word  $w_i$  in  $S$  with that of the train statistic. If a word has the same mini-tree in a given sentence as it has in the training set, then the overlap would be 1. For a given sentence  $S_i$ , the aggregate  $\beta_{\{S_i\}}^k$  is defined by the average of the overlap scores of all its words:  $\beta_{\{S_i\}}^k = \frac{1}{|S_i|} \sum_{w_i \in S_i} \beta_{\{w_i, S_i\}}^k$ , and we call it a POS minitree *signature*. We can also compute the POS minitree signature of a permuted sentence  $\hat{S}_i$  to have  $\beta_{\{\hat{S}_i\}}^k$ . If the permuted sentence POS signature comes close to that of the true sentence, then their ratio (i.e.,  $\beta_{\{\hat{S}_i\}}^k / \beta_{\{S_i\}}^k$ ) will be close to 1. Also, since POS signature is computed with respect to the train distribution, a ratio of  $> 1$  indicates that the permuted sentence is closer to the overall train statistic than to the original unpermuted sentence in terms of POS signature. If high overlap with the training distribution correlates with percentage of permutations deemed correct, then our models treat words as if they project syntactic minitrees.




Figure 4: POS Tag Mini Tree overlap score and percentage of permutations which the models assigned the gold-label.

We investigate the relationship with percentage of permuted sentences accepted with  $\beta_{\{\hat{S}_i\}}^k / \beta_{\{S_i\}}^k$  in Figure 4. We observe that the POS Tag Minitree hypothesis holds for Transformer-based models, RoBERTa, BART and DistilBERT, where the percentage of accepted pairs increase as the sentences have higher overlap with the un-permuted sentence in terms of POS signature. For non-Transformer models such as InferSent, ConvNet, and BiLSTM models, the POS signature ratio to percentage of correct permutation remains the same or decreases, suggesting that the reasoning process employed by these models does not preserve local abstract syntax structure (i.e., POS neighbor relations).

## 7 Human Evaluation

We expect humans to struggle with UNLI, given our intuitions and the sentence superiority findings (but see Mollica et al. 2020). To test this, we presented two experts in NLI (one a linguist) with permuted sentence pairs to label.<sup>4</sup> Concretely, we draw equal number of examples from MNLI Matched dev set (100 examples where RoBERTa predicts the gold label,  $D^c$  and 100 examples where it fails to do so,  $D^f$ ), and then permute these examples using  $\mathcal{F}$ . The experts were given no additional information (recall that it is common knowledge that NLI is a roughly balanced 3-way classification task). Unbeknownst to the experts, all permuted sentences in the sample were actually accepted by the RoBERTa (large) model (trained on MNLI dataset). We observe that the experts performed

<sup>4</sup>Concurrent work by Gupta et al. (2021) found that untrained crowdworkers accept NLI examples that have been subjected to different kinds of perturbations at roughly most frequent class levels—i.e., only 35% of the time.



Evaluator	Accuracy	Macro F1	Acc on $D^c$	Acc on $D^f$
X	0.581 $\pm$ 0.068	0.454	0.649 $\pm$ 0.102	0.515 $\pm$ 0.089
Y	0.378 $\pm$ 0.064	0.378	0.411 $\pm$ 0.098	0.349 $\pm$ 0.087

Table 4: Human (expert) evaluation on 200 permuted examples from the MNLI matched development set. Half of the permuted pairs contained shorter sentences and the other, longer ones. All permuted examples were assigned the gold label by RoBERTa-Large.

much worse than RoBERTa (Table 4), although their accuracy was a bit higher than random. We also find that for both experts, accuracy on permutations from  $D^c$  was higher than on  $D^f$ , which verifies findings that showed high word overlap can give hints about the ground truth label (Dasgupta et al., 2018; Poliak et al., 2018; Gururangan et al., 2018; Naik et al., 2019).

## 8 Training by Maximizing Entropy

We propose an initial attempt to mitigate the effect of correct prediction on permuted examples. As we observe in §5, model entropy on permuted examples is significantly lower than expected. Neural networks tend to output higher confidence than random for even unknown inputs (Gandhi and Lake, 2020), which might be an underlying cause of the high Permutation Acceptance.

An ideal model would be ambivalent about randomized ungrammatical sentences. Thus, we train NLI models baking in the principle of mutual exclusivity (Gandhi and Lake, 2020) by maximizing model entropy. Concretely, we fine-tune RoBERTa on MNLI while maximizing the entropy ( $\mathbf{H}$ ) on a subset of  $n$  randomized examples  $((\hat{p}_i, \hat{r}_i)$ , for each example  $(p, h)$  in MNLI. We modify the loss function as follows:

$$\mathcal{L} = \underset{\theta}{\operatorname{argmin}} \sum_{((p,h),y)} y \log(p(y|(p, h); \theta)) + \sum_{i=1}^n \mathbf{H} \left( y | (\hat{p}_i, \hat{h}_i); \theta \right) \quad (5)$$

Using this maximum entropy method ( $n = 1$ ), we find that the model improves considerably with respect to its robustness to randomized sentences, all while taking no hit to accuracy (Table 5). We observe that no model reaches a  $\Omega_{\max}$  score close to 0, suggesting further room to explore other methods for decreasing models’ Permutation Acceptance. Similar approaches have also proven useful (Gupta et al., 2021) for other tasks as well.

Eval Dataset	$\mathcal{A}$ (V)	$\mathcal{A}$ (ME)	$\Omega_{\max}$ (V)	$\Omega_{\max}$ (ME)
MNLI_m_dev	0.905	0.908	0.984	0.328
MNLI_mm_dev	0.901	0.903	0.985	0.329
SNLI_test	0.882	0.888	0.983	0.329
SNLI_dev	0.879	0.887	0.984	0.333
ANLI_r1_dev	0.456	0.470	0.890	0.333
ANLI_r2_dev	0.271	0.258	0.880	0.333
ANLI_r3_dev	0.268	0.243	0.892	0.334

Table 5: NLI Accuracy ( $\mathcal{A}$ ) and Permutation Acceptance metrics ( $\Omega_{\max}$ ) of RoBERTa when trained on MNLI dataset using vanilla (V) and Maximum Random Entropy (ME) method.

## 9 Future Work & Conclusion

We show that state-of-the-art models do not rely on sentence structure the way we think they should: NLI models (Transformer-based models, RNNs, and ConvNets) are largely insensitive to permutations of word order that corrupt the original syntax. We also show that reordering words can cause models to flip classification labels. We do find that models seem to have learned some syntactic information as is evidenced by a correlation between preservation of abstract POS neighborhood information and rate of acceptance by models, but these results do not discount the high rates of Permutation Acceptance, and require further verification. Coupled with the finding that humans cannot perform UNLI at all well, the high rate of permutation acceptance that we observe leads us to conclude that current models do not yet “know syntax” in the fully systematic and humanlike way we would like them to.

A few years ago, Manning (2015) encouraged NLP to consider “the details of human language, how it is learned, processed, and how it changes, rather than just chasing state-of-the-art numbers on a benchmark task.” We expand upon this view, and suggest one particular future direction: we should train models not only to do well on clean test data, but also to not to overgeneralize to corrupted input.

## Acknowledgments

Thanks to Omar Agha, Dzmitry Bahdanau, Sam Bowman, Hagen Blix, Ryan Cotterell, Emily Dinan, Michal Drozdal, Charlie Lovering, Nikita Nangia, Alicia Parrish, Grusha Prasad, Roy Schwartz, Shagun Sodhani, Anna Szabolsci, Alex Warstadt, Jackie Chi-kit Cheung, Timothy O’Donnell and members of McGill MCQLL lab for many invaluable comments and feedback on early drafts.

## References

- Anne Abeille. 1990. [Lexical and syntactic rules in a Tree Adjoining Grammar](#). In *28th Annual Meeting of the Association for Computational Linguistics*, pages 292–298, Pittsburgh, Pennsylvania, USA. Association for Computational Linguistics.
- Matteo Alleman, Jonathan Mamou, Miguel A Del Rio, Hanlin Tang, Yoon Kim, and SueYeon Chung. 2021. [Syntactic perturbations reveal representational correlates of hierarchical phrase structure in pretrained language models](#). *arXiv preprint arXiv:2104.07578*.
- Aixiu An, Peng Qian, Ethan Wilcox, and Roger Levy. 2019. [Representation of constituents in neural language models: Coordination phrase as a case study](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2888–2899, Hong Kong, China. Association for Computational Linguistics.
- Alan D Baddeley, Graham J Hitch, and Richard J Allen. 2009. [Working memory and binding in sentence recall](#). *Journal of Memory and Language*.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Jean-Phillipe Bernardy and Shalom Lappin. 2017. [Using deep neural networks to learn syntactic agreement](#). In *Linguistic Issues in Language Technology, Volume 15, 2017*. CSLI Publications.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent dirichlet allocation](#). *JMLR*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Joan Bresnan, Ash Asudeh, Ida Toivonen, and Stephen Wechsler. 2015. *Lexical-functional syntax*. John Wiley & Sons.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- James McKeen Cattell. 1886. [The time it takes to see and name objects](#). *Mind*, os-XI(41):63–65.
- Rui Chaves. 2020. [What don’t RNN language models learn about filler-gap dependencies?](#) In *Proceedings of the Society for Computation in Linguistics 2020*, pages 1–11, New York, New York. Association for Computational Linguistics.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Gennaro Chierchia and Sally McConnell-Ginet. 1990. *Meaning and grammar: An Introduction to Semantics*. Cambridge, Ma: MIT Press.
- Grzegorz Chrupała and Afra Alishahi. 2019. [Correlating neural and symbolic representations of language](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2952–2962, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Ronan Collobert and Jason Weston. 2008. [A unified architecture for natural language processing: Deep neural networks with multitask learning](#). In *ICML*.
- Cleo Condoravdi, Dick Crouch, Valeria de Paiva, Reinhard Stolle, and Daniel G. Bobrow. 2003. [Entailment, intensionality and text understanding](#). In *Proceedings of the HLT-NAACL 2003 Workshop on Text Meaning*, pages 38–45.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Jillian Da Costa and Rui Chaves. 2020. [Assessing the ability of transformer-based neural models to represent structurally unbounded dependencies](#). In *Proceedings of the Society for Computation in Linguistics 2020*, pages 12–21, New York, New York. Association for Computational Linguistics.

- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The pascal recognising textual entailment challenge](#). In *Machine Learning Challenges Workshop*. Springer.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The PASCAL recognising textual entailment challenge. In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising textual entailment*. Springer.
- Ishita Dasgupta, Demi Guo, Andreas Stuhlmüller, Samuel J Gershman, and Noah D Goodman. 2018. [Evaluating compositionality in sentence embeddings](#). In *Proceedings of Annual Meeting of the Cognitive Science Society*.
- Forrest Davis and Marten van Schijndel. 2020. [Recurrent neural network language models always learn English-like relative clause attachment](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1979–1990, Online. Association for Computational Linguistics.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Gottlob Frege. 1948. Sense and reference. *The philosophical review*.
- Kanishk Gandhi and Brenden M Lake. 2020. [Mutual exclusivity as a challenge for deep neural networks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 14182–14192. Curran Associates, Inc.
- Jon Gauthier, Jennifer Hu, Ethan Wilcox, Peng Qian, and Roger Levy. 2020. [SyntaxGym: An online platform for targeted evaluation of language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 70–76, Online. Association for Computational Linguistics.
- Yoav Goldberg. 2019. [Assessing BERT’s syntactic abilities](#). *arXiv preprint arXiv:1901.05287*.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.
- Ashim Gupta, Giorgi Kvernadze, and Vivek Srikumar. 2021. [BERT & family eat word salad: Experiments with text understanding](#). *AAAI*.
- 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.
- Zellig S Harris. 1954. Distributional structure. *Word*.
- Robert Hawkins, Takateru Yamakoshi, Thomas Griffiths, and Adele Goldberg. 2020. [Investigating representations of verb bias in neural language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4653–4663, Online. Association for Computational Linguistics.
- Irene Heim and Angelika Kratzer. 1998. *Semantics in generative grammar*. Blackwell Oxford.
- 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.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence Moss. 2020a. [OCNLI: Original Chinese Natural Language Inference](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3512–3526, Online. Association for Computational Linguistics.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020b. [A systematic assessment of syntactic generalization in neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamé 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.
- Paloma Jeretic, Alex Warstadt, Suvrat Bhooshan, and Adina Williams. 2020. [Are natural language inference models IMPPRESSive? Learning IMPLICature and PRESupposition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8690–8705, Online. Association for Computational Linguistics.



- Ronald M Kaplan and Joan Bresnan. 1995. Formal system for grammatical representation. *Formal Issues in Lexical-Functional Grammar*.
- Yiğitcan Kaya, Sanghyun Hong, and Tudor Dumitras. 2019. [Shallow-Deep Networks: Understanding and mitigating network overthinking](#). In *Proceedings of the 2019 International Conference on Machine Learning (ICML)*, Long Beach, CA.
- Jordan Kodner and Nitish Gupta. 2020. [Overestimation of syntactic representation in neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1757–1762, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Peiguang Li, Hongfeng Yu, Wenkai Zhang, Guangluan Xu, and Xian Sun. 2020. [SA-NLI: A supervised attention based framework for natural language inference](#). *Neurocomputing*.
- 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.
- Tal Linzen and Marco Baroni. 2021. [Syntactic structure from deep learning](#). *Annual Review of Linguistics*.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- 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](#). *arXiv*.
- Christopher D Manning. 2015. Computational linguistics and deep learning. *Computational Linguistics*.
- Christopher D. Manning, Kevin Clark, John Hewitt, Urvashi Khandelwal, and Omer Levy. 2020. [Emergent linguistic structure in artificial neural networks trained by self-supervision](#). *Proceedings of the National Academy of Sciences*, 117(48):30046–30054.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- R. Thomas McCoy, Junghyun Min, and Tal Linzen. 2020. [BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 217–227, Online. 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.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Francis Mollica, Matthew Siegelman, Evgeniia Datchek, Steven T Piantadosi, Zachary Mineroff, Richard Futrell, Hope Kean, Peng Qian, and Evelina Fedorenko. 2020. [Composition is the core driver of the language-selective network](#). *Neurobiology of Language*, 1(1):104–134.
- Richard Montague. 1970. Universal grammar. *Theoria*.
- Aakanksha Naik, Abhilasha Ravichander, Carolyn Rose, and Eduard Hovy. 2019. [Exploring numeracy in word embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3374–3380, Florence, Italy. 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.
- Nikita Nangia and Samuel R. Bowman. 2019. [Human vs. muppet: A conservative estimate of human performance on the GLUE benchmark](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4566–4575, Florence, Italy. Association for Computational Linguistics.



- 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.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Prasanna Parthasarathi, Koustuv Sinha, Joelle Pineau, and Adina Williams. 2021. [Sometimes we want translationese](#). *arXiv preprint arXiv:2104.07623*.
- Thang M Pham, Trung Bui, Long Mai, and Anh Nguyen. 2020. [Out of order: How important is the sequential order of words in a sentence in natural language understanding tasks?](#) *arXiv preprint arXiv:2012.15180*.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis only baselines in natural language inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Carl Pollard and Ivan A Sag. 1994. *Head-driven phrase structure grammar*. University of Chicago Press.
- Grusha Prasad, Marten van Schijndel, and Tal Linzen. 2019. [Using priming to uncover the organization of syntactic representations in neural language models](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 66–76, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*.
- Shauli Ravfogel, Yoav Goldberg, and Tal Linzen. 2019. [Studying the inductive biases of RNNs with synthetic variations of natural languages](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3532–3542, Minneapolis, Minnesota. Association for Computational Linguistics.
- Shauli Ravfogel, Yoav Goldberg, and Francis Tyers. 2018. [Can LSTM learn to capture agreement? the case of Basque](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 98–107, Brussels, Belgium. Association for Computational Linguistics.
- Abhilasha Ravichander, Aakanksha Naik, Carolyn Rose, and Eduard Hovy. 2019. [EQUATE: A benchmark evaluation framework for quantitative reasoning in natural language inference](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 349–361, Hong Kong, China. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#).
- Yves Schabes, Anne Abeille, and Aravind K. Joshi. 1988. [Parsing strategies with ‘lexicalized’ grammars: Application to Tree Adjoining Grammars](#). In *Coling Budapest 1988 Volume 2: International Conference on Computational Linguistics*.
- Eckart Scheerer. 1981. Early german approaches to experimental reading research: The contributions of wilhelm wundt and ernst meumann. *Psychological Research*.
- Marten van Schijndel and Tal Linzen. 2018. [A neural model of adaptation in reading](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4704–4710, Brussels, Belgium. Association for Computational Linguistics.
- Marten van Schijndel, Aaron Mueller, and Tal Linzen. 2019. [Quantity doesn’t buy quality syntax with neural language models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5831–5837, Hong Kong, China. Association for Computational Linguistics.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021. [Masked language modeling and the distributional hypothesis: Order word matters pre-training for little](#). *arXiv preprint arXiv:2104.06644*.
- Joshua Snell and Jonathan Grainger. 2017. The sentence superiority effect revisited. *Cognition*.
- Joshua Snell and Jonathan Grainger. 2019. Word position coding in reading is noisy. *Psychonomic bulletin & review*, 26(2):609–615.
- Whitney Tabor. 1994. *Syntactic innovation: A connectionist model*. Ph.D. thesis.

- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers 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.
- Hiroshi Toyota. 2001. [Changes in the constraints of semantic and syntactic congruity on memory across three age groups](#). *Perceptual and Motor Skills*.
- Masatoshi Tsuchiya. 2018. [Performance impact caused by hidden bias of training data for recognizing textual entailment](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Lucy Vanderwende and William B Dolan. 2005. [What syntax can contribute in the entailment task](#). In *Machine Learning Challenges Workshop*. Springer.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). 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 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *NeurIPS*.
- 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.
- Alex Warstadt and Samuel R Bowman. 2020. [Can neural networks acquire a structural bias from raw linguistic data?](#) In *Proceedings of the 42nd Annual Virtual Meeting of the Cognitive Science Society*.
- Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, Sheng-Fu Wang, Jason Phang, Anhad Mohanane, Phu Mon Htut, Paloma Jeretic, and Samuel R. Bowman. 2019a. [Investigating BERT’s knowledge of language: Five analysis methods with NPIs](#). 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 2877–2887, Hong Kong, China. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohanane, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019b. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Yun Wen, Joshua Snell, and Jonathan Grainger. 2019. [Parallel, cascaded, interactive processing of words during sentence reading](#). *Cognition*.
- Aaron Steven White, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2018. [Lexicosyntactic inference in neural models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4717–4724, Brussels, Belgium. Association for Computational Linguistics.
- Ethan Wilcox, Roger Levy, Takashi Morita, and Richard Futrell. 2018. [What do RNN language models learn about filler-gap dependencies?](#) In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 211–221, Brussels, Belgium. Association for Computational Linguistics.
- Adina Williams, Andrew Drozdov, and Samuel R. Bowman. 2018a. [Do latent tree learning models identify meaningful structure in sentences?](#) *Transactions of the Association for Computational Linguistics*, 6:253–267.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018b. [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.
- Ludwig Wittgenstein. 1922. *Tractatus Logico-Philosophicus*. Harcourt, Brace & Company, Inc.
- Thomas Wolf. 2019. [Some additional experiments extending the tech report” assessing bert’s syntactic abilities” by yoav goldberg](#). Technical report, HuggingFace.
- Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. 2020. [Transformers: State-of-the-art natural language processing](#). In *EMNLP: System Demonstrations*.

- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020. [Perturbed masking: Parameter-free probing for analyzing and interpreting BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4166–4176, Online. Association for Computational Linguistics.
- Kelly Zhang and Samuel Bowman. 2018. [Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361, Brussels, Belgium. Association for Computational Linguistics.
- Han Zhao, Zhengdong Lu, and Pascal Poupart. 2015. [Self-adaptive hierarchical sentence model](#). In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 4069–4076.
- Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. 2020. [Bert loses patience: Fast and robust inference with early exit](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18330–18341. Curran Associates, Inc.





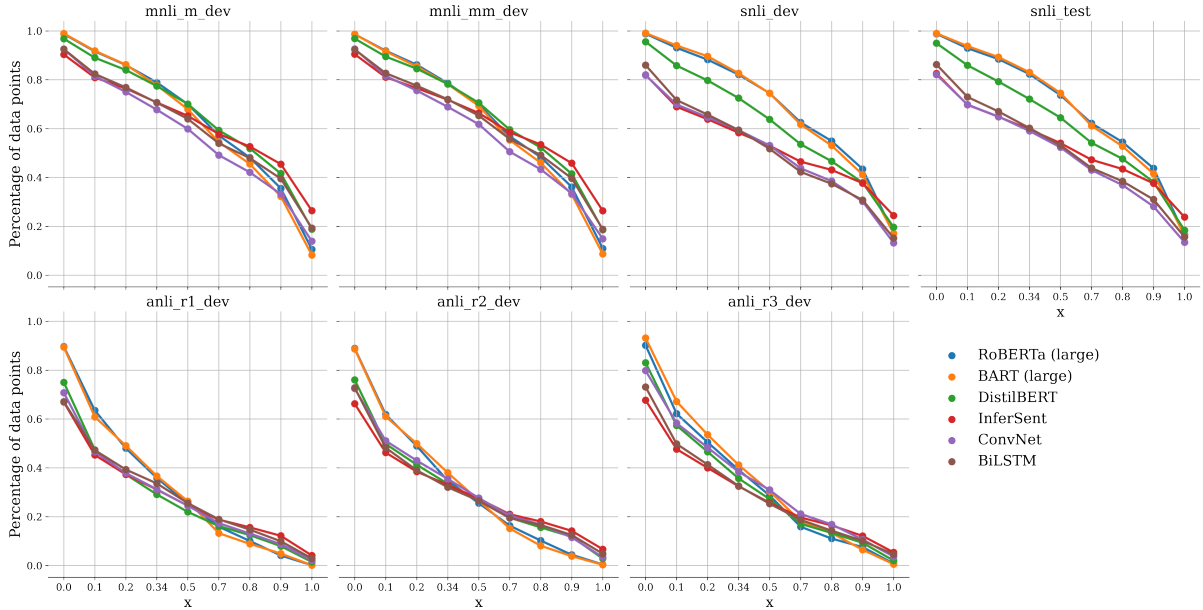


Figure 8:  $\Omega_x$  threshold for all datasets with varying  $x$  and computing the percentage of examples that fall within the threshold. The top row consists of in-distribution datasets (MNLI, SNLI) and the bottom row contains out-of-distribution datasets (ANLI)

## E Effect of the threshold of $\Omega_x$ in various test splits

We defined two variations of  $\Omega_x$ ,  $\Omega_{\max}$  and  $\Omega_{\text{rand}}$ , but theoretically it is possible to define any arbitrary threshold percentage  $x$  to evaluate the unnatural language inference mechanisms of different models. In Figure 8 we show the effect of different thresholds, including  $\Omega_{\max}$  where  $x = 1/|D_{\text{test}}|$  and  $\Omega_{\text{rand}}$  where  $x = 0.34$ . We observe for in-distribution datasets (top row, MNLI and SNLI splits), in the extreme setting when  $x = 1.0$ , there are more than 10% of examples available, and more than 25% in case of InferSent and DistilBERT. For out-of-distribution datasets (bottom row, ANLI splits) we observe a much lower trend, suggesting generalization itself is the bottleneck in permuted sentence understanding.

## F Training with permuted examples

In this section, we hypothesize that if the NLU models are mostly insensitive to word order, then training using permuted examples should also yield the same or comparable accuracy as training using grammatically correct data (i.e., the *standard setup*). To test this, we train Transformer-based models on top of  $\hat{D}_{\text{train}}$ , which is computed by applying  $\mathcal{F}$  on each example of  $D_{\text{train}}$  for  $q = 1$  times. This ensures a control case where we keep the same amount of training data as the standard setup (such

Eval Data	RoBERTa		BART		DistilBERT	
	$\mathcal{A}$	$\hat{\mathcal{A}}$	$\mathcal{A}$	$\hat{\mathcal{A}}$	$\mathcal{A}$	$\hat{\mathcal{A}}$
MNLI Matched	0.906	0.877	0.902	0.862	0.800	0.760
MNLI Mismatched	0.901	0.878	0.900	0.869	0.811	0.769
SNLI Dev	0.879	0.870	0.886	0.854	0.732	0.719
SNLI Test	0.883	0.873	0.888	0.859	0.738	0.719
ANLI R1 (Dev)	0.456	0.367	0.455	0.336	0.251	0.250
ANLI R2 (Dev)	0.271	0.279	0.316	0.293	0.300	0.290
ANLI R3 (Dev)	0.268	0.271	0.327	0.309	0.312	0.312

Table 6: Statistics for Transformer-based models when trained on permuted MNLI corpus. We compare the accuracy for both models trained on unpermuted data ( $\mathcal{A}$ ) and the permuted data ( $\hat{\mathcal{A}}$ ). We use original test sets during inference.

that models does not benefit from data augmentation). We also ensure that we use the same hyperparameters while training as with the standard setup. Concretely,  $\hat{D}_{\text{train}}$  consists of  $n$  hypothesis-premise pairs from MNLI training data, where each example is a permuted output of the original pair.

We present the results of such training in Table 6, and compare the accuracy ( $\hat{\mathcal{A}}$ ) with that of the standard setup ( $\mathcal{A}$ ). Note, during inference for all the models we use the un-permuted examples. As we can see, models perform surprisingly close to the original accuracy  $\mathcal{A}$  even when trained with *ungrammatical sentences*. This adds further proof to the BOW nature of NLU models.




Figure 9: Comparing the effect between randomizing both premise and hypothesis and only hypothesis on two Transformer-based models, RoBERTa and BART (For more comparisons please refer to Appendix). In 9(a), we observe the difference of  $\Omega_{\max}$  is marginal in in-distribution datasets (SNLI, MNLI), while hypothesis-only randomization is worse for out-of-distribution datasets (ANLI). In 9(b), we compare the mean number of permutations which elicited correct response, and naturally the hypothesis-only randomization causes more percentage of randomizations to be correct.




Figure 10: Relation of BLEU-2/3/4 scores against the acceptability of clumped-permuted sentences across all test datasets on all models.

Dataset	Test Examples	Used Examples
MNLI Matched	9815	6655
MNLI Mismatched	9832	7449
SNLI Dev	9842	3697
SNLI Test	9824	3671
ANLI R1 (Dev)	1000	756
ANLI R2 (Dev)	1000	709
ANLI R3 (Dev)	1200	754

Table 7: Dataset statistics used in this paper for inference. ‘Used Examples’ provides the number of premise-hypothesis pairs for the dataset which we selected for inference (i.e., examples where at least 100 unique permutations were possible).

## G Reproducibility Checklist

As per the prescribed Reproducibility Checklist, we provide the information of the following:

- *A clear description of the mathematical setting, algorithm and/or model:* We provide details of models used in §5
- *Description of the computing infrastructure used:* We used 8 NVIDIA V100 32GB GPUs to train the models and perform all necessary inferences. We didn’t run hyperparameter tuning for Transformer-based models as we used the published hyperparameters from the original models.
- *Average runtime for each approach:* On an average, each model inference experiment consistine of 100 permutations for each example takes roughly 1 hour to complete.
- *Relevant statistics of the datasets used:* We provide the statistics of the datasets used in Table 7.
- *Explanation of any data that were excluded, and all pre-processing steps:* We exclude examples where either the hypothesis and premise consists of less than 6 tokens. This way, we ensure that we have 100 unique permutations for each example.
- *Link to downloadable version of data and code:* We provide downloadable version of our data and code at <https://github.com/facebookresearch/unlu>.