

How Does Data Corruption Affect Natural Language Understanding Models? A Study on GLUE Datasets

Aarne Talman^{*§} Marianna Apidianaki[◇]
Stergios Chatzikyriakidis^{†‡§} Jörg Tiedemann^{*}

^{*}Department of Digital Humanities, University of Helsinki

{name.surname}@helsinki.fi

[◇]Department of Computer and Information Science, University of Pennsylvania

marapi@seas.upenn.edu

[†]Department of Philology, University of Crete

stergios.chatzikyriakidis@uoc.gr

[‡]Centre of Linguistic Theory and Studies in Probability, FLoV, University of Gothenburg

[§]Basement AI

Abstract

A central question in natural language understanding (NLU) research is whether high performance demonstrates the models' strong reasoning capabilities. We present an extensive series of controlled experiments where pre-trained language models are exposed to data that have undergone specific corruption transformations. These involve removing instances of specific word classes and often lead to non-sensical sentences. Our results show that performance remains high on most GLUE tasks when the models are fine-tuned or tested on corrupted data, suggesting that they leverage other cues for prediction even in non-sensical contexts. Our proposed data transformations can be used to assess the extent to which a specific dataset constitutes a proper testbed for evaluating models' language understanding capabilities.

1 Introduction

The super-human performance of recent Transformer-based pre-trained language models (Devlin et al., 2019; Liu et al., 2019) on natural language understanding (NLU) tasks has raised scepticism regarding the quality of the benchmarks used for evaluation (Wang et al., 2018, 2019). There is increasing evidence that these datasets contain annotation artefacts and other statistical irregularities that can be leveraged by machine learning models to perform the tasks (Gururangan et al., 2018; Poliak et al., 2018b; Tsuchiya, 2018; Glockner et al., 2018; Talman and Chatzikyriakidis, 2019; Pham et al., 2020; Talman et al., 2021). These studies have so far largely focused on the natural language inference (NLI) and textual entailment tasks. The scope of our work is wider, in the sense that we address all but one NLU tasks

	Sentence 1	Sentence 2
paraphrase	<i>Easynews Inc. was subpoenaed late last week by the FBI, which was seeking account information related to the uploading of the virus to the ISP's Usenet news group server.</i>	<i>Easynews Inc. said Monday that it was cooperating with the FBI in trying to locate the person who uploaded the virus to a Usenet news group hosted by the ISP.</i>
non-paraphrase	<i>Arison said Mann may have been one of the pioneers of the world music movement and he had a deep love of Brazilian music.</i>	<i>Arison said Mann was a pioneer of the world music movement – well before the term was coined – and he had a deep love of Brazilian music.</i>

Table 1: Example sentence pairs from the corrupted MRPC training dataset where all instances of nouns have been removed.

comprised in the GLUE benchmark, specifically: linguistic acceptability (COLA), paraphrasing (MRPC and QQP), sentiment prediction (SST-2), and semantic textual similarity (STS-B).

We present a series of experiments where the datasets used for model training and evaluation undergo a number of corruption transformations, which involve removing specific word classes from the data. We remove words pertaining to a specific class (e.g., nouns, verbs), instead of random words, to see the relative importance of word classes for the NLU tasks. For instance, verbs arguably play a significant role in sentence level semantics and removing them is expected to have a bigger impact on the GLUE scores, compared to say determiners.

The transformations seriously affect the quality of the sentences found in the datasets, making them in many cases unintelligible (cf. examples in Table 1); a decrease in performance for mod-

	Task	Baseline	Metric
COLA	The Corpus of Linguistic Acceptability (Warstadt et al., 2018)	64.05	Matthew’s correlation
MNLI-M	Multi-Genre Natural Language Inference (Williams et al., 2018)	87.89	accuracy
MRPC	Microsoft Research Paraphrase Corpus (Dolan and Brockett, 2005)	88.73	accuracy
QNLI	Question Natural Language Inference (Rajpurkar et al., 2016)	92.64	accuracy
QQP	Quora Question Pairs	91.32	accuracy
RTE	Recognizing Textual Entailment (Dagan et al., 2006)	70.04	accuracy
SST-2	The Stanford Sentiment Treebank (Socher et al., 2013)	94.61	accuracy
STS-B	Semantic Textual Similarity Benchmark (Cer et al., 2017)	90.08	Pearson correlation

Table 2: Baseline results obtained for different GLUE tasks with RoBERTa-base and the relevant metric.

els fine-tuned on these corrupted datasets would, thus, be expected. High performance would, instead, indicate that the models rely on lexical cues that remain after corruption, and possibly on other dataset artefacts, to perform a task without necessarily understanding the meaning of the processed utterances.

Our results show that performance after the corruptions remains high for most GLUE tasks, suggesting that the models leverage other cues for prediction even in non-sensical contexts.

2 Related Work

Annotation artefacts and statistical biases in NLI datasets are easily leveraged by the models and can guide prediction (Lai and Hockenmaier, 2014; Marelli et al., 2014; Poliak et al., 2018a; Gururangan et al., 2018). Examples include explicit negation being indicative of contradiction, and generic nouns suggesting entailment. Artefacts are also present in other types of datasets, for example in the ROC Story dataset where models can provide story endings without looking at the actual stories (Schwartz et al., 2017; Cai et al., 2017). Several works have proposed more challenging and cleaner NLI datasets where artefacts have been removed (McCoy et al., 2019). An efficient way to do this is using adversarial filtering (Nie et al., 2020; Zellers et al., 2018). The superior quality of the resulting NLI datasets is confirmed by Talman et al. (2021) in a series of experiments where it is shown that data corruption affects these higher quality datasets to a greater extent than previous datasets.

This work follows the same experimental direction where text perturbations serve to explore the sensitivity of language models to specific phenomena (Futrell et al., 2019; Ettinger, 2020; Taktaševa et al., 2021; Dankers et al., 2021). It has been shown, for example, that shuffling word order causes significant performance drops on a wide range of QA tasks (Si et al., 2019; Sugawara et al.,

2019), but that state-of-the-art NLU models are not sensitive to word order (Pham et al., 2020; Sinha et al., 2021). Syntax-based perturbations have also been studied in relation to robustness and faithfulness of machine translation models (Parthasarathi et al., 2021).

We add to this line of research by applying data corruption transformations that involve removing entire word classes (Talman et al., 2021) to all but one GLUE tasks.¹ We interpret high performance of models fine-tuned and/or tested on corrupted datasets as an indication of the presence of lexical cues, and possibly artefacts, guiding prediction, since the meaning of the corrupted utterances is often hard to recover.

3 Datasets and Corruptions

In our experiments, we address eight tasks included in the General Language Understanding Evaluation (GLUE) benchmark for the English language (Wang et al., 2018): CoLa, MNLI, MRPC, QNLI, QQP, RTE, SST-2, STS-B. Following Talman et al. (2021), we corrupt the training and development sets available for these tasks by removing words of specific word classes.² We use the development sets for evaluation, since annotated test data have not been made publicly available.³ We create three configurations for each task: (a) CORRUPT-TRAIN: fine-tuning on the corrupted training set, evaluation on the original development set; (b) CORRUPT-TEST: fine-tuning on the original training set, evaluation on the corrupted test set; (c) CORRUPT-TRAIN AND TEST: training and evaluation on corrupted data. The corruption procedure

¹We exclude WNLI as its development dataset was designed to be adversarial (Wang et al., 2018) and hence the corruptions do not have any impact on this dataset when evaluating with the development set.

²We annotate the original texts with universal part of speech (POS) tags using the NLTK library (<https://www.nltk.org/>) and the averaged perceptron tagger.

³For MNLI, we use the matched development set (Williams et al., 2018).

Data	CORRUPT-TRAIN	Δ	CORRUPT-TEST	Δ	CORRUPT-TRAIN AND TEST	Δ
COLA-NOUN	39.72	-24.34	17.75	-46.30	34.33	-29.73
MNLI-M-NOUN	85.64	-2.24	72.85	-15.04	77.46	-10.42
MRPC-NOUN	86.27	-2.45	82.35	-6.37	80.15	-8.58
QNLI-NOUN	89.13	-3.51	71.02	-21.62	82.02	-10.62
QQP-NOUN	86.69	-4.63	72.57	-18.75	84.17	-7.16
RTE-NOUN	47.29	-22.74	53.79	-16.25	47.29	-22.74
SST-2-NOUN	94.04	-0.57	87.27	-7.34	88.76	-5.85
STS-B-NOUN	81.67	-8.41	56.12	-33.96	63.52	-26.56
COLA-VERB	23.26	-40.79	4.30	-59.75	20.22	-43.83
MNLI-M-VERB	86.95	-0.94	77.61	-10.28	80.32	-7.57
MRPC-VERB	85.54	-3.19	85.54	-3.19	85.05	-3.68
QNLI-VERB	92.00	-0.64	87.41	-5.24	90.15	-2.49
QQP-VERB	89.49	-1.84	86.01	-5.31	89.05	-2.27
RTE-VERB	65.34	-4.69	65.70	-4.33	65.34	-4.69
SST-2-VERB	93.69	-0.92	89.33	-5.28	89.56	-5.05
STS-B-VERB	87.63	-2.46	85.54	-4.54	86.22	-3.86

Table 3: Example results for the RoBERTa-base model fine-tuned on CORRUPT-TRAIN and tested on the original evaluation set (columns 2 and 3); fine-tuned on the original data and tested on CORRUPT-TEST; fine-tuned on CORRUPT-TRAIN and tested on CORRUPT-TEST (columns 6 and 7). Δ is the difference to the baseline scores obtained by RoBERTa-base on the original dataset, given in Table 2.

involves removing all instances of a specific word class from the corresponding dataset (ADJ, ADV, CONJ, DET, NOUN, NUM, PRON, VERB). We label the corrupted datasets by indicating the class of the words that have been removed (e.g., COLA-NOUN, QNLI-VERB). Given the possible combinations of tasks, datasets and corruptions, we end up with 192 setups for our experiments.

Note that the resulting sentence fragments do not constitute propositions. Although not ideal, this is not necessarily problematic for tasks such as sentiment analysis. For inference, the assumption that the task can only be performed at the propositional level is a strong claim, especially given that examples which are not propositions are abundant in existing benchmarks such as MNLI (e.g., examples extracted from dialogue).

4 Models

We fine-tune the pre-trained RoBERTa-base model (Liu et al., 2019) from the Huggingface Transformers library (Wolf et al., 2020a) in each of our 192 configurations. We use the same fine-tuning and evaluation set up for all the experiments. We retrieve the GLUE datasets using the Huggingface Datasets library (Wolf et al., 2020b). We fine-tune the models for 3 epochs, using a batch size of 32 and a learning rate of 0.00002.

5 Results

The baseline results using the original (non-corrupted) datasets are shown in Table 2. Given the large number of configurations, we only report the exact evaluation results for the -NOUN and -VERB settings in Table 3, as these content word




Figure 1: Impact of specific data corruptions in the CORRUPT-TRAIN setting. The columns correspond to the removed word class and the rows to the GLUE tasks.

classes arguably contribute a lot to the meaning of utterances. For the remaining configurations, we visualise the effect of the corruptions using heatmaps that show the difference in performance compared to the baseline results (Figures 1 to 3).

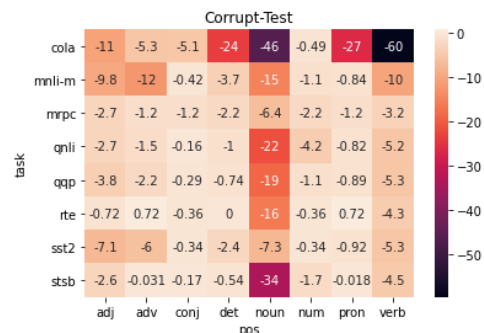


Figure 2: Impact of specific data corruptions in the CORRUPT-TEST setting for each task.

Our results for the -NOUN and -VERB corruptions in CORRUPT-TRAIN (Table 3), and for all

Original Sentences	CORRUPT-TEST-NOUN	CORRUPT-TEST-ADJ	Labels	Gold label
<i>An unclassifiably awful study in self - and audience-abuse.</i>	<i>an unclassifiably awful in - and.</i>	<i>an unclassifiably study in self - and audience-abuse.</i>	positive	negative
<i>It proves quite compelling as an intense, brooding character study.</i>	<i>it proves quite compelling as an intense, brooding.</i>	<i>it proves quite as an, brooding character study.</i>	positive	positive

Table 4: Labels assigned by NRCLex to sentences from the SST-2 CORRUPT-TEST-NOUN/-ADJ datasets.

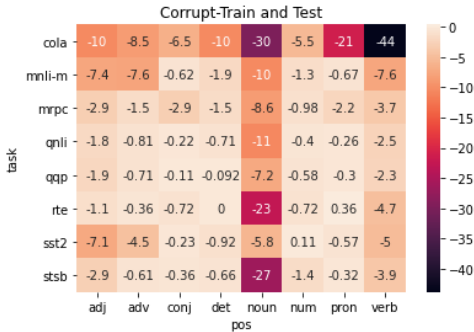


Figure 3: Impact of specific data corruptions in the CORRUPT-TRAIN AND TEST setting for each task.

configurations in Figure 1, show a notable decrease in performance on COLA and RTE, especially when nouns are removed. The impact on MNLi-M and QNLI datasets is small, confirming previous findings regarding the presence of annotation artefacts and lexical cues that can guide model prediction. Our results suggest that this is the case also in other GLUE datasets, such as MRPC and SST-2, where the models still manage to perform fairly well compared to the baseline when fine-tuned on corrupted data.

Our CORRUPT-TEST results in Table 3 and in Figure 2 show that removing nouns from the data used for evaluation has a much larger impact across tasks, compared to CORRUPT-TRAIN. The biggest drop in performance is observed on COLA, MNLi-M and STS-B. However, accuracy on MRPC and SST-2 is still very high, suggesting that good performance does not require sentence-level understanding but can be achieved by relying on lexical cues present in the data. In the CORRUPT-TRAIN AND TEST setting (Table 3 and Figure 3), we observe the biggest drop in performance on COLA, MNLi-M and STS-B, and a lower impact on QNLI, QQP and SST-2.

6 Discussion and Analysis

6.1 Lexical Cues

Our results show that model performance in many tasks is marginally affected by the imposed corruptions which, however, in many cases alter the mean-

Word class	Dataset	Accuracy
NOUN	CORRUPT-TEST	14.7%
NOUN	original	34.1%
VERB	CORRUPT-TEST	31.1%
VERB	original	66.4%

Table 5: Accuracy of RoBERTa-BASE in predicting a masked word in the MRPC development set.

ing of utterances. We conduct additional analyses aimed at identifying the lexical cues that remain after corruption and can guide model prediction. We focus on MRPC (Microsoft Research Paraphrase Corpus) and SST-2 (Stanford Sentiment Treebank), where the impact of CORRUPT-TEST transformations was the smallest.

MRPC addresses the paraphrase relationship between sentence pairs. We explore the semantic similarity of the information that remains after corruption. Our assumption is that if a sentence pair (from which nouns or verbs have been removed) still contains synonyms or longer paraphrases, this can guide the model towards detecting a similarity or entailment relationship. For this analysis, we use the unigram paraphrases in the L (large) package of PPDB 2.0 (Pavlick et al., 2015). We find that in the CORRUPT-TEST-NOUN MRPC dataset, 76% of the sentence pairs for which the model made correct predictions still include a lexical paraphrase.

SST-2 involves detecting the sentiment expressed in individual sentences. We use the NRCLex tool⁴ to measure the sentiment expressed by lexical cues in the CORRUPT-TEST sentences for which model predictions are correct. Given that sentiment can be expressed in a text by words pertaining to different grammatical categories, we explore whether lexical cues indicating the polarity of the text still remain after removing instances of a specific word class. In Table 4, we show the labels predicted by NRCLex for corrupted test sentences, where the nouns and adjectives have been dropped. We observe that even if sentences become non-sensical

⁴NRCLex is based on the expanded version of the NRC Word-Emotion Association Lexicon (Mohammad and Turney, 2010, 2013). We only use the ‘positive’ and ‘negative’ keys.

after corruption, it is still possible to detect the (positive or negative) polarity of the sentences from the remaining words. Relying on these lexical cues, RoBERTa often manages to predict the correct sentiment. Specifically, according to the NRCLex predictions, the correct sentiment is still present in 383 out of 761 corrupted sentences where RoBERTa made correct predictions in the CORRUPT-TEST-NOUN setting. If both nouns and adjectives are removed (CORRUPT-TEST-NOUN-ADJ), NRCLex detects that the correct sentiment is still present in 125 out of the 672 examples that were correctly predicted by RoBERTa.

6.2 Can RoBERTa Guess the Missing Tokens?

As RoBERTa has been pre-trained using a Masked Word Prediction task, it is reasonable to ask if high model performance with our corrupted datasets could be due to the model’s ability to “fill in the gaps” and predict the missing words. To test this, in each sentence of the MRCP development set, we replace the first token that is aimed by a specific corruption procedure (-NOUN/VERB) with the [MASK] token. We do this in the original sentence (by removing only the first noun/verb instance) and in the corrupted sentence (where all other nouns/verbs are missing). For example, from the first sentence in Table 4, we generate two cloze-task queries in the -NOUN setting:

- (a) An unclassifiably awful [MASK] in self - and audience-abuse.
- (b) An unclassifiably awful [MASK] in - and.

We use these queries to test RoBERTa’s token prediction capability. As shown in Table 5, it is easier to predict the masked token in the original sentences, but the model is still able to make correct predictions in the corrupted sentences. This could partly explain the high performance observed for MRPC in the corrupted setting (cf. Section 5).

7 Conclusion

We apply a set of controllable corruption transformations to the datasets of NLU tasks in the GLUE benchmark, and study their impact on model performance. The proposed transformations are generic enough to be applicable to other NLU tasks, and can enrich the available artillery for dataset quality assessment in terms of how efficiently they trigger and test the language understanding capabilities of the models. Our results indicate that understanding

the meaning of utterances is not required for high performance in most GLUE tasks. This finding suggests caution in interpreting leaderboard results and in the conclusions that can be drawn regarding the language understanding capabilities of the models. We make our code available⁵ in order to promote the application of these tests to other NLU datasets, and to favour the development of benchmarks addressing the actual capability of the models to reason about language.

Acknowledgments



During the course of this project, Marianna Apidianaki and Jörg Tiedemann were supported by the FoTran project, funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 771113). Stergios Chatzikyriakidis is supported by grant 2014-39 from the Swedish Research Council, which funds the Centre for Linguistic Theory and Studies in Probability (CLASP) in the Department of Philosophy, Linguistics, and Theory of Science at the University of Gothenburg. We thank the reviewers for their thoughtful comments and valuable suggestions.

References

- Zheng Cai, Lifu Tu, and Kevin Gimpel. 2017. *Pay Attention to the Ending: Strong Neural Baselines for the ROC Story Cloze Task*. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 616–622, Vancouver, Canada. Association for Computational Linguistics.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. *SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation*. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- 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*, pages 177–190. Springer.

⁵<https://github.com/Helsinki-NLP/nlu-dataset-diagnostics>

- Verna Dankers, Elia Bruni, and Dieuwke Hupkes. 2021. [The paradox of the compositionality of natural language: a neural machine translation case study](#). *arXiv preprint arXiv:2108.05885*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of ACL*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically Constructing a Corpus of Sentential Phrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- 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.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. 2019. [Neural language models as psycholinguistic subjects: Representations of syntactic state](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 32–42, Minneapolis, Minnesota. Association for Computational Linguistics.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI Systems with Sentences that Require Simple Lexical Inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics.
- 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.
- Alice Lai and Julia Hockenmaier. 2014. [Illinois-LH: A Denotational and Distributional Approach to Semantics](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 329–334, Dublin, Ireland.
- 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 preprint arXiv:1907.11692*.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. [A SICK cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- 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.
- Saif Mohammad and Peter Turney. 2010. [Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon](#). In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34, Los Angeles, CA. Association for Computational Linguistics.
- Saif M. Mohammad and Peter D. Turney. 2013. [Crowdsourcing a Word–Emotion Association Lexicon](#). *Computational Intelligence*, 29.
- 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.
- Prasanna Parthasarathi, Koustuv Sinha, Joelle Pineau, and Adina Williams. 2021. [Sometimes We Want Ungrammatical Translations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3205–3227, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. [PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics.
- 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, Aparajita Haldar, Rachel Rudinger, J. Edward Hu, Ellie Pavlick, Aaron Steven White, and Benjamin Van Durme. 2018a. [Collecting Diverse Natural Language Inference Problems for Sentence Representation Evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 67–81, Brussels, Belgium. Association for Computational Linguistics.

- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018b. [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.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Roy Schwartz, Maarten Sap, Ioannis Konstas, Leila Zilles, Yejin Choi, and Noah A. Smith. 2017. [The Effect of Different Writing Tasks on Linguistic Style: A Case Study of the ROC Story Cloze Task](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 15–25, Vancouver, Canada. Association for Computational Linguistics.
- Chenglei Si, Shuohang Wang, Min-Yen Kan, and Jing Jiang. 2019. [What does BERT learn from multiple-choice reading comprehension datasets?](#) *arXiv preprint arXiv:1910.12391*.
- 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*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2019. [Assessing the Benchmarking Capacity of Machine Reading Comprehension Datasets](#). *arXiv preprint arXiv:1911.09241*.
- Ekaterina Taktasheva, Vladislav Mikhailov, and Ekaterina Artemova. 2021. [Shaking Syntactic Trees on the Sesame Street: Multilingual Probing with Controllable Perturbations](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 191–210, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Aarne Talman, Marianna Apidianaki, Stergios Chatzikyriakidis, and Jörg Tiedemann. 2021. [NLI Data Sanity Check: Assessing the Effect of Data Corruption on Model Performance](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 276–287, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Aarne Talman and Stergios Chatzikyriakidis. 2019. [Testing the Generalization Power of Neural Network Models across NLI Benchmarks](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 85–94, Florence, Italy. Association for Computational Linguistics.
- 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).
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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 BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2018. [Neural Network Acceptability Judgments](#). *arXiv preprint arXiv:1805.12471*.
- 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020a. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Thomas Wolf, Quentin Lhoest, Patrick von Platen, Yacine Jernite, Mariama Drame, Julien Plu, Julien Chaumond, Clement Delangue, Clara Ma, Abhishek Thakur, Suraj Patil, Joe Davison, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angie McMillan-Major, Simon Brandeis, Sylvain Gugger,

François Lagunas, Lysandre Debut, Morgan Funtowicz, Anthony Moi, Sasha Rush, Philipp Schmid, Pierric Cistac, Victor Muštar, Jeff Boudier, and Anna Tordjmann. 2020b. [Datasets](#). *GitHub*. Note: <https://github.com/huggingface/datasets>, 1.

Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.