

# Syntactic Data Augmentation Increases Robustness to Inference Heuristics

Junghyun Min<sup>1</sup> R. Thomas McCoy<sup>1</sup> Dipanjan Das<sup>2</sup> Emily Pitler<sup>2</sup> Tal Linzen<sup>1</sup>

<sup>1</sup>Department of Cognitive Science, Johns Hopkins University, Baltimore, MD

<sup>2</sup>Google Research, New York, NY

{jmin10, tom.mccoy, tal.linzen}@jhu.edu

{dipanjang, epitler}@google.com

## Abstract

Pretrained neural models such as BERT, when fine-tuned to perform natural language inference (NLI), often show high accuracy on standard datasets, but display a surprising lack of sensitivity to word order on controlled challenge sets. We hypothesize that this issue is not primarily caused by the pretrained model’s limitations, but rather by the paucity of crowd-sourced NLI examples that might convey the importance of syntactic structure at the fine-tuning stage. We explore several methods to augment standard training sets with syntactically informative examples, generated by applying syntactic transformations to sentences from the MNLI corpus. The best-performing augmentation method, subject/object inversion, improved BERT’s accuracy on controlled examples that diagnose sensitivity to word order from 0.28 to 0.73, without affecting performance on the MNLI test set. This improvement generalized beyond the particular construction used for data augmentation, suggesting that augmentation causes BERT to recruit abstract syntactic representations.

## 1 Introduction

In the supervised learning paradigm common in NLP, a large collection of labeled examples of a particular classification task is randomly split into a training set and a test set. The system is trained on this training set, and is then evaluated on the test set. Neural networks—in particular systems pretrained on a word prediction objective, such as ELMo (Peters et al., 2018) or BERT (Devlin et al., 2019)—excel in this paradigm: with large enough pretraining corpora, these models match or even exceed the accuracy of untrained human annotators on many test sets (Raffel et al., 2019).

At the same time, there is mounting evidence that high accuracy on a *test set* drawn from the

same distribution as the training set does not indicate that the model has mastered the *task*. This discrepancy can manifest as a sharp drop in accuracy when the model is applied to a different dataset that illustrates the same task (Talmor and Berant, 2019; Yogatama et al., 2019), or as excessive sensitivity to linguistically irrelevant perturbations of the input (Jia and Liang, 2017; Wallace et al., 2019).

One such discrepancy, where strong performance on a standard test set did not correspond to mastery of the task as a human would define it, was documented by McCoy et al. (2019b) for the Natural Language Inference (NLI) task. In this task, the system is given two sentences, and is expected to determine whether one (the premise) entails the other (the hypothesis). Most if not all humans would agree that NLI requires sensitivity to syntactic structure; for example, the following sentences do not entail each other, even though they contain the same words:

- (1) The lawyer saw the actor.
- (2) The actor saw the lawyer.

McCoy et al. constructed the HANS challenge set, which includes examples of a range of such constructions, and used it to show that, when BERT is fine-tuned on the MNLI corpus (Williams et al., 2018), the fine-tuned model achieves high accuracy on the test set drawn from that corpus, yet displays little sensitivity to syntax; the model wrongly concluded, for example, that (1) entails (2).

We consider two explanations as to why BERT fine-tuned on MNLI fails on HANS. Under the **Representational Inadequacy Hypothesis**, BERT fails on HANS because its pretrained representations are missing some necessary syntactic information. Under the **Missed Connection Hypothesis**, BERT extracts the relevant syntactic information from the input (cf. Goldberg 2019;

Tenney et al. 2019), but it fails to *use* this information with HANS because there are few MNLI training examples that indicate how syntax should support NLI (McCoy et al., 2019b). It is possible for both hypotheses to be correct: there may be some aspects of syntax that BERT has not learned at all, and other aspects that have been learned, but are not applied to perform inference.

The Missed Connection Hypothesis predicts that augmenting the training set with a small number of examples from one syntactic construction would teach BERT that the task requires it to use its syntactic representations. This would not only cause improvements on the construction used for augmentation, but would also lead to generalization to other constructions. In contrast, the Representational Inadequacy Hypothesis predicts that to perform better on HANS, BERT must be taught how each syntactic construction affects NLI from scratch. This predicts that larger augmentation sets will be required for adequate performance and that there will be little generalization across constructions.

This paper aims to test these hypotheses. We constructed augmentation sets by applying syntactic transformations to a small number of examples from MNLI. Accuracy on syntactically challenging cases improved dramatically as a result of augmenting MNLI with only about 400 examples in which the subject and the object were swapped (about 0.1% of the size of the MNLI training set). Crucially, even though only a single transformation was used in augmentation, accuracy increased on a range of constructions. For example, BERT’s accuracy on examples involving relative clauses (e.g., *The actors called the banker who the tourists saw* → *The banker called the tourists*) was 0.33 without augmentation, and 0.83 with it. This suggests that our method does not overfit to one construction, but taps into BERT’s existing syntactic representations, providing support for the Missed Connection Hypothesis. At the same time, we also observe limits to generalization, supporting the Representational Inadequacy Hypothesis in those cases.

## 2 Background

HANS is a template-generated challenge set designed to test whether NLI models have adopted three syntactic heuristics. First, the **lexical overlap heuristic** is the assumption that any time all of the words in the hypothesis are also in the premise, the label should be *entailment*. In the MNLI training

set, this heuristic often makes correct predictions, and almost never makes incorrect predictions. This may be due to the process by which MNLI was generated: crowdworkers were given a premise and were asked to generate a sentence that contradicts or entails the premise. To minimize effort, workers may have overused lexical overlap as a shortcut to generating entailed hypotheses. Of course, the lexical overlap heuristic is not a generally valid inference strategy, and it fails on many HANS examples; e.g., as discussed above, *the lawyer saw the actor* does not entail *the actor saw the lawyer*.

HANS also includes cases that are diagnostic of the **subsequence heuristic** (assume that a premise entails any hypothesis which is a contiguous subsequence of it) and the **constituent heuristic** (assume that a premise entails all of its constituents). While we focus on counteracting the lexical overlap heuristic, we will also test for generalization to the other heuristics, which can be seen as particularly challenging cases of lexical overlap. Examples of all constructions used to diagnose the three heuristics are given in Tables A.5, A.6 and A.7.

Data augmentation is often employed to increase robustness in vision (Perez and Wang, 2017) and language (Belinkov and Bisk, 2018; Wei and Zou, 2019), including in NLI (Minervini and Riedel, 2018; Yanaka et al., 2019). In many cases, augmentation with one kind of example improves accuracy on that particular case, but does not generalize to other cases, suggesting that models overfit to the augmentation set (Jia and Liang, 2017; Ribeiro et al., 2018; Iyyer et al., 2018; Liu et al., 2019). In particular, McCoy et al. (2019b) found that augmentation with HANS examples generalized to a different word overlap challenge set (Dasgupta et al., 2018), but only for examples similar in length to HANS examples. We mitigate such overfitting to superficial properties by generating a diverse set of corpus-based examples, which differ from the challenge set both lexically and syntactically. Finally, Kim et al. (2018) used a similar augmentation approach to ours but did not study generalization to types of examples not in the augmentation set.

## 3 Generating Augmentation Data

We generate augmentation examples from MNLI using two syntactic transformations: INVERSION, which swaps the subject and object of the source sentence, and PASSIVIZATION. For each of these transformations, we had two families of augmenta-

---

Original MNLi example:
There are 16 El Grecos in this small collection. →
This small collection contains 16 El Grecos.
Inversion (original premise):
There are 16 El Grecos in this small collection. ↗
16 El Grecos contain this small collection.
Inversion (transformed hypothesis):
This small collection contains 16 El Grecos. ↗
16 El Grecos contain this small collection.
Passivization (transformed hypothesis; <i>non-entailment</i> ):
This small collection contains 16 El Grecos. ↗
This small collection is contained by 16 El Grecos.
Random shuffling with a random label:
16 collection small El contains Grecos This. ↗/→
collection This Grecos El small 16 contains.

---

Table 1: A sample of syntactic augmentation strategies, with gold labels (→: *entailment*; ↗: *non-entailment*). For the full list, see Table A.1 in the Appendix.

tion sets. The ORIGINAL PREMISE strategy keeps the original MNLi premise and transforms the hypothesis; and TRANSFORMED HYPOTHESIS uses the original MNLi hypothesis as the new premise, and the transformed hypothesis as the new hypothesis (see Table 1 for examples, and §A.2 for details). We experimented with three augmentation set sizes: small (101 examples), medium (405) and large (1215). All augmentation sets were much smaller than the MNLi training set (297k).<sup>1</sup>

We did not attempt to ensure the naturalness of the generated examples; e.g., in the INVERSION transformation, *The carriage made a lot of noise* was transformed into *A lot of noise made the carriage*. In addition, the labels of the augmentation dataset were somewhat noisy; e.g., we assumed that INVERSION changed the correct label from *entailment* to *neutral*, but this is not necessarily the case (if *The buyer met the seller*, it is likely that *The seller met the buyer*). As we show below, this noise did not hurt accuracy on MNLi.

Finally, we included a random shuffling condition, in which an MNLi premise and its hypothesis were both randomly shuffled, with a random label. We used this condition to test whether a syntactically uninformed method could teach the model that, when word order is ignored, no reliable inferences can be made.

<sup>1</sup>The augmentation sets and the code used to generate them are available at <https://github.com/aatlantise/syntactic-augmentation-nli>.

## 4 Experimental setup

We added each augmentation set separately to the MNLi training set, and fine-tuned BERT on each resulting training set. Further fine-tuning details are in Appendix A.1. We repeated this process for five random seeds for each combination of augmentation strategy and augmentation set size, except for the most successful strategy (INVERSION + TRANSFORMED HYPOTHESIS), for which we had 15 runs for each augmentation size. Following McCoy et al. (2019b), when evaluating on HANS, we merged the neutral and contradiction labels produced by the model into a single *non-entailment* label.

For both ORIGINAL PREMISE and TRANSFORMED HYPOTHESIS, we experimented with using each of the transformations separately, and with a combined dataset including both inversion and passivization. We also ran separate experiments with only the passivization examples with an *entailment* label, and with only the passivization examples with a *non-entailment* label. As a baseline, we used 100 runs of BERT fine-tuned on the unaugmented MNLi (McCoy et al., 2019a).

We report the models’ accuracy on HANS, as well as on the MNLi development set (MNLi test set labels are not publicly available). We did not tune any parameters on this development set. All of the comparisons we discuss below are significant at the  $p < 0.01$  level (based on two-sided t-tests).

## 5 Results

Accuracy on MNLi was very similar across augmentation strategies and matched that of the unaugmented baseline (0.84), suggesting that syntactic augmentation with up to 1215 examples does not harm overall performance on the dataset. By contrast, accuracy on HANS varied significantly, with most models performing worse than chance (which is 0.50 on HANS) on *non-entailment* examples, suggesting that they adopted the heuristics (Figure 1). The most effective augmentation strategy, by a large margin, was inversion with a transformed hypothesis. Accuracy on the HANS word overlap cases for which the correct label is *non-entailment*—e.g., *the doctor saw the lawyer* ↗ *the lawyer saw the doctor*—was 0.28 without augmentation, and 0.73 with the large version of this augmentation set. Simultaneously, this strategy **decreased** BERT’s accuracy on the cases where the heuristic makes the correct prediction (*The tourists by the actor called the authors* → *The tourists called the authors*); in

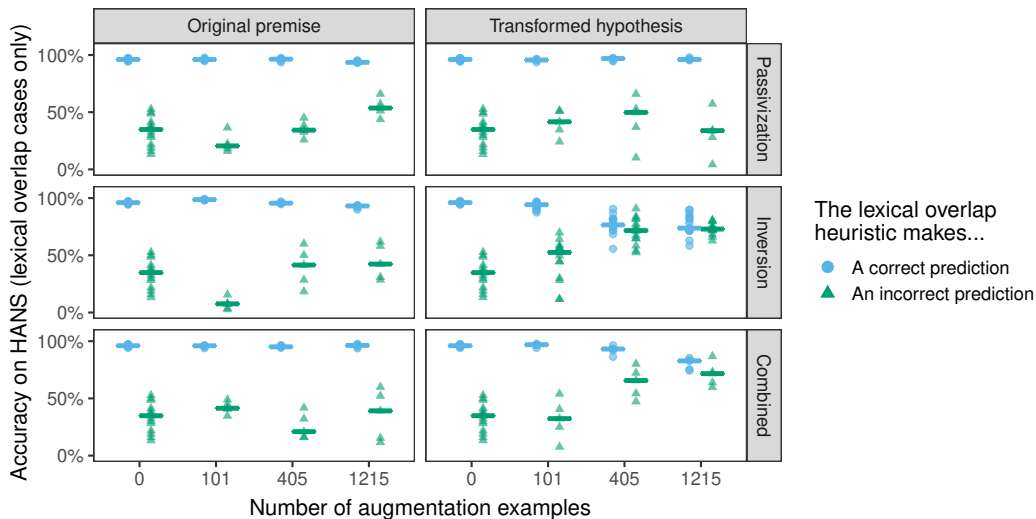


Figure 1: Comparison of syntactic augmentation strategies. Dots represent accuracy on the HANS examples that diagnose the lexical overlap heuristic, as produced by each of the runs of BERT fine-tuned on MNLI combined with each augmentation data set. Horizontal bars indicate median accuracy across runs. Chance accuracy is 0.5.

fact, the best model’s accuracy was similar across cases where lexical overlap made correct and incorrect predictions, suggesting that this intervention prevented the model from adopting the heuristic.

The random shuffling method did not improve over the unaugmented baseline, suggesting that syntactically-informed transformations are essential (Table A.2). Passivization yielded a much smaller benefit than inversion, perhaps due to the presence of overt markers such as the word *by*, which may lead the model to attend to word order only when those are present. Intriguingly, even on the passive examples in HANS, inversion was more effective than passivization (large inversion augmentation: 0.13; large passivization augmentation: 0.01). Finally, inversion on its own was more effective than the combination of inversion and passivization.

We now analyze in more detail the most effective strategy, inversion with a transformed hypothesis. First, this strategy is similar on an abstract level to the HANS subject/object swap category, but the two differ in vocabulary and some syntactic properties; despite these differences, performance on this HANS category was perfect (1.00) with medium and large augmentation, indicating that BERT benefited from the high-level syntactic structure of the transformation. For the small augmentation set, accuracy on this category was 0.53, suggesting that 101 examples are insufficient to teach BERT that subjects and objects cannot be freely swapped. Conversely, tripling the augmentation size from

medium to large had a moderate and inconsistent effect across HANS subcases (see Appendix A.3 for case-by-case results); for clearer insight about the role of augmentation size, it may be necessary to sample this parameter more densely.

Although inversion was the only transformation in this augmentation set, performance also improved dramatically on constructions other than subject/object swap (Figure 2); for example, the models handled examples involving a prepositional phrase better, concluding, for instance, that *The judge behind the manager saw the doctors* does not entail *The doctors saw the manager* (unaugmented: 0.41; large augmentation: 0.89). There was a much more moderate, but still significant, improvement on the cases targeting the subsequence heuristic; this smaller degree of improvement suggests that contiguous subsequences are treated separately from lexical overlap more generally. One exception was accuracy on “NP/S” inferences, such as *the managers heard the secretary resigned*  $\rightarrow$  *The managers heard the secretary*, which improved dramatically from 0.02 (unaugmented) to 0.50 (large augmentation). Further improvements for subsequence cases may therefore require augmentation with examples involving subsequences.

A range of techniques have been proposed over the past year for improving performance on HANS. These include syntax-aware models (Moradshahi et al., 2019; Pang et al., 2019), auxiliary models designed to capture pre-defined shallow heuristics so that the main model can focus on robust strategies



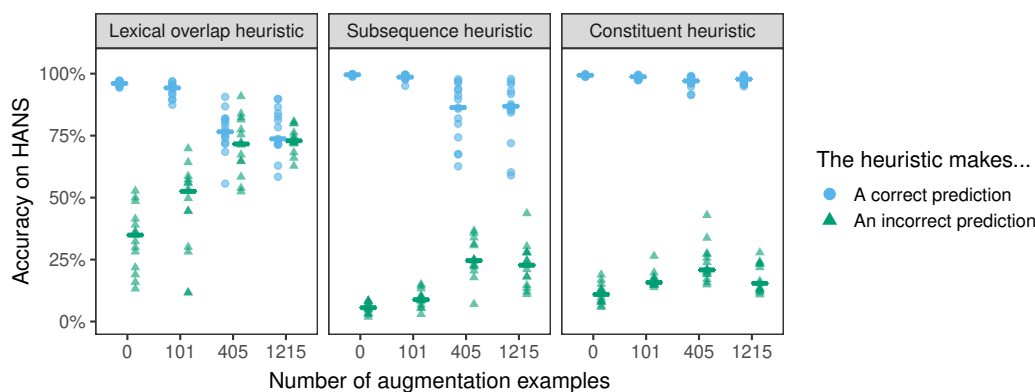


Figure 2: Augmentation using subject/object inversion with a transformed hypothesis. Dots represent the accuracy on HANS examples diagnostic of each of the heuristics, as produced by each of the 15 runs of BERT fine-tuned on MNLI combined with each augmentation data set. Horizontal bars indicate median accuracy across runs.

(Clark et al., 2019; He et al., 2019; Mahabadi and Henderson, 2019), and methods to up-weight difficult training examples (Yaghoobzadeh et al., 2019). While some of these approaches yield higher accuracy on HANS than ours, including better generalization to the constituent and subsequence cases (see Table A.4), they are not directly comparable: our goal is to assess how the prevalence of syntactically challenging examples in the training set affects BERT’s NLI performance, without modifying either the model or the training procedure.

## 6 Discussion

Our best-performing strategy involved augmenting the MNLI training set with a small number of instances generated by applying the subject/object inversion transformation to MNLI examples. This yielded considerable generalization: both to another domain (the HANS challenge set), and, more importantly, to additional constructions, such as relative clauses and prepositional phrases. This supports the Missed Connection Hypothesis: a small amount of augmentation with one construction induced abstract syntactic sensitivity, instead of just “inoculating” the model against failing on the challenge set by providing it with a sample of cases from the same distribution (Liu et al., 2019).

At the same time, the inversion transformation did not completely counteract the heuristic; in particular, the models showed poor performance on passive sentences. For these constructions, then, BERT’s pretraining may not yield strong syntactic representations that can be tapped into with a small nudge from augmentation; in other words, this may be a case where our Representational Inadequacy Hypothesis holds. This hypothesis predicts

that pretrained BERT, as a word prediction model, struggles with passives, and may need to learn the properties of this construction specifically for the NLI task; this would likely require a much larger number of augmentation examples.

The best-performing augmentation strategy involved generating premise/hypothesis pairs from a single source sentence—meaning that this strategy does not rely on an NLI corpus. The fact that we can generate augmentation examples from any corpus makes it possible to test if very large augmentation sets are effective (with the caveat, of course, that augmentation sentences from a different domain may hurt performance on MNLI itself).

Ultimately, it would be desirable to have a model with a strong inductive bias for using syntax across language understanding tasks, even when overlap heuristics lead to high accuracy on the training set; indeed, it is hard to imagine that a human would ignore syntax entirely when understanding a sentence. An alternative would be to create training sets that adequately represent a diverse range of linguistic phenomena; crowdworkers’ (rational) preferences for using the simplest generation strategies possible could be counteracted by approaches such as adversarial filtering (Nie et al., 2019). In the interim, however, we conclude that data augmentation is a simple and effective strategy to mitigate known inference heuristics in models such as BERT.

## Acknowledgments

This research was supported by a gift from Google, NSF Graduate Research Fellowship No. 1746891, and NSF Grant No. BCS-1920924. Our experiments were conducted using the Maryland Advanced Research Computing Center (MARCC).

## References

- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *International Conference on Learning Representations*.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. [Don't take the easy way out: Ensemble based methods for avoiding known dataset biases](#). 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 4067–4080, Hong Kong, China. Association for Computational Linguistics.
- Ishita Dasgupta, Demi Guo, Andreas Stuhlmüller, Samuel J. Gershman, and Noah D. Goodman. 2018. [Evaluating compositionality in sentence embeddings](#). In *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, pages 1596–1601, Madison, WI.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yoav Goldberg. 2019. [Assessing BERT's syntactic abilities](#). *arXiv preprint arXiv:1901.05287*.
- He He, Sheng Zha, and Haohan Wang. 2019. [Unlearn dataset bias in natural language inference by fitting the residual](#). In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 132–142, Hong Kong, China. Association for Computational Linguistics.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). 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 1875–1885. Association for Computational Linguistics.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031. Association for Computational Linguistics.
- Juho Kim, Christopher Malon, and Asim Kadav. 2018. [Teaching syntax by adversarial distraction](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 79–84, Brussels, Belgium. Association for Computational Linguistics.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019. [Inoculation by fine-tuning: A method for analyzing challenge datasets](#). 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 2171–2179, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rabeeh Karimi Mahabadi and James Henderson. 2019. [Simple but effective techniques to reduce biases](#). *arXiv preprint arXiv:1909.06321*.
- R. Thomas McCoy, Junghyun Min, and Tal Linzen. 2019a. [BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance](#). *arXiv preprint arXiv:1911.02969*.
- R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019b. [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.
- Pasquale Minervini and Sebastian Riedel. 2018. [Adversarially regularising neural NLI models to integrate logical background knowledge](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 65–74, Brussels, Belgium. Association for Computational Linguistics.
- Mehrad Moradshahi, Hamid Palangi, Monica S. Lam, Paul Smolensky, and Jianfeng Gao. 2019. [HUBERT Untangles BERT to Improve Transfer across NLP Tasks](#). *arXiv preprint arXiv:1910.12647*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2019. [Adversarial NLI: A new benchmark for natural language understanding](#). *arXiv preprint arXiv:1910.14599*.
- Deric Pang, Lucy H. Lin, and Noah A. Smith. 2019. [Improving natural language inference with a pre-trained parser](#). *arXiv preprint arXiv:1909.08217*.
- Luis Perez and Jason Wang. 2017. [The effectiveness of data augmentation in image classification using deep learning](#). *arXiv preprint arXiv:1712.04621*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv preprint arXiv:1910.10683*.

- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. [Semantically equivalent adversarial rules for debugging NLP models](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. Association for Computational Linguistics.
- Alon Talmor and Jonathan Berant. 2019. [MultiQA: An empirical investigation of generalization and transfer in reading comprehension](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4911–4921, Florence, Italy. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? Probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- 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.
- Jason Wei and Kai Zou. 2019. [EDA: Easy data augmentation techniques for boosting performance on text classification tasks](#). 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 6381–6387, Hong Kong, China. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Yadollah Yaghoobzadeh, Remi Tachet, T. J. Hazen, and Alessandro Sordani. 2019. [Robust natural language inference models with example forgetting](#). *arXiv preprint arXiv:1911.03861*.
- Hitomi Yanaka, Koji Mineshima, Daisuke Bekki, Kentaro Inui, Satoshi Sekine, Lasha Abzianidze, and Johan Bos. 2019. [HELP: A dataset for identifying shortcomings of neural models in monotonicity reasoning](#). In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (\*SEM 2019)*, pages 250–255, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, and Phil Blunsom. 2019. [Learning and evaluating general linguistic intelligence](#). *arXiv preprint arXiv:1901.11373*.

## A Appendix

### A.1 Fine-tuning details

We used `bert-base-uncased` for all experiments. As is standard, we fine-tuned this pretrained model on MNLI by training a linear classifier to predict the label from the CLS token’s final layer embedding, while continuing to update BERT’s parameters (Devlin et al., 2019). The order of training examples was reshuffled for each model. All models were trained for three epochs.

### A.2 Generating augmentation examples

The following list describes the augmentation strategies we used. Table A.1 illustrates all of these strategies as applied to a particular source sentence. Note that inversion generally changes the meaning of the sentence (*the detective followed the suspect* refers to a different event from *the suspect followed the detective*), but passivization on its own does not (*the detective followed the suspect* refers to the same event as *the suspect was followed by the detective*).

- Inversion (original premise): For a source example  $(p, h, \rightarrow)$ , generate  $(p, \text{INV}(h), \rightarrow)$ , where `INV` returns the source sentence with the subject and object switched. Ignore source examples whose label is  $\rightarrow$ .
- Inversion (transformed hypothesis): For a source  $(p, h)$  (with any label), discard the premise  $p$  and generate  $(h, \text{INV}(h), \rightarrow)$ .
- Passivization (original premise): For a source  $(p, h)$  (with any label), generate  $(p, \text{PASS}(h))$ , with the same label, where `PASS` returns the passive version of the source sentence (without changing its meaning).
- Passivization (transformed hypothesis): For a source  $(p, h)$ , discard the premise  $p$ , and generate two examples, one with an *entailment* label— $(h, \text{PASS}(h), \rightarrow)$ —and one with a *non-entailment* label— $(h, \text{PASS}(\text{INV}(h)), \rightarrow)$ .

We identified transitive sentences in MNLI that could serve as source sentences using the constituency parses provided with MNLI, excluding the noisier TELEPHONE genre. We did so by searching for matrix S nodes with exactly one NP daughter of the VP, where the subject and the object were both full noun phrases (i.e., neither were a personal pronoun such as *me*), and where the verb lemma

was not *be* or *have*. We kept the original tense of the verb, and modified its agreement features if necessary (e.g., *the movie stars Matt Dillon and Gary Sinise* was transformed into *Matt Dillon and Gary Sinise star the movie*).

The size of the largest augmentation set was 1215 for all strategies. This size was determined based on the largest augmentation dataset we could generate from MNLI for the inversion with original premise strategy using the procedure mentioned above. For fair comparison, we kept the same size even for strategies where we could have generated a larger dataset. We also created a Medium dataset by randomly sampling 405 of the cases identifying using the procedure above, as well as a small dataset with 101 examples. We performed this process only once for each strategy: as such, runs varied only in the classifier’s weight initialization and the order of examples but not in the augmentation examples included in training.

To create the Combined augmentation dataset, we concatenated the inversion and passivization datasets, then randomly discarded half of the examples (to match the size of the combined dataset with the others). As with the other datasets, we only did this once: the Combined augmentation set was the same across runs. One consequence of this procedure is that the number of passivization and inversion examples was not exactly identical.

### A.3 Detailed Results

The following tables provide the detailed results of our experiments. Table A.2 shows each strategy’s mean accuracy on MNLI, as well on the HANS cases that diagnose each of the three heuristics (the Lexical Overlap Heuristic, the Subsequence Heuristic, and the Constituent Heuristic), for which the correct label is *non-entailment* ( $\rightarrow$ ). Table A.3 zooms in on the best-performing augmentation strategy—subject/object inversion with a transformed hypothesis—on BERT’s accuracy on HANS, both when the correct label is *entailment* ( $\rightarrow$ ) and when the label is *non-entailment* ( $\rightarrow$ ). Finally, the last three tables detail the effect of augmentation by inversion with a transformed hypothesis on each of the 30 HANS subcases, broken down by the heuristic that they were designed to diagnose: the Lexical Overlap Heuristic (Table A.5), the Subsequence Heuristic (Table A.6), and the Constituent Heuristic (Table A.7).



---

**Original**

There are 16 El Grecos in this small collection.  $\rightarrow$   
This small collection contains 16 El Grecos.

---

**Inversion**

Original premise:  
There are 16 El Grecos in this small collection.  $\leftrightarrow$   
16 El Grecos contain this small collection.

Transformed hypothesis:  
This small collection contains 16 El Grecos.  $\leftrightarrow$   
16 El Grecos contain this small collection.

---

**Passivization**

Original premise:  
There are 16 El Grecos in this small collection.  $\rightarrow$   
16 El Grecos are contained by this small collection.

Transformed hypothesis (*entailment* label):  
This small collection contains 16 El Grecos.  $\rightarrow$   
16 El Grecos are contained by the small collection.

Transformed hypothesis (*non-entailment* label):  
This small collection contains 16 El Grecos.  $\leftrightarrow$   
This small collection is contained by 16 El Grecos.

---

**Random shuffling** (with random label)

are collection. small El this in 16 There Grecos  $\leftrightarrow/\rightarrow$   
collection This Grecos El small 16 contains.

---

Table A.1: Syntactic augmentation strategies (full table).

	MNLi			Overlap			Subsequence			Constituent		
	<i>S</i>	<i>M</i>	<i>L</i>	<i>S</i>	<i>M</i>	<i>L</i>	<i>S</i>	<i>M</i>	<i>L</i>	<i>S</i>	<i>M</i>	<i>L</i>
<b>Original premise</b>												
Inversion	.84	.84	.84	.07	.40	.44	.01	.06	.12	.06	.09	.12
Passivization	.84	.84	.84	.23	.35	.54	.04	.05	.09	.13	.11	.15
Combined	.84	.84	.84	.42	.25	.36	.07	.05	.04	.14	.15	.12
<b>Transformed hypothesis</b>												
Inversion	.84	.84	.84	<b>.46</b>	<b>.71</b>	<b>.73</b>	<b>.09</b>	<b>.25</b>	.23	<b>.17</b>	.23	.18
Passivization	.84	.84	.84	.41	.43	.31	.06	.06	.07	.13	.15	.17
Combined	.84	.84	.84	.32	.64	.71	.06	.13	<b>.28</b>	.15	<b>.26</b>	<b>.22</b>
Pass. (only pos)	.84	.84	.84	.30	.20	.29	.04	.04	.05	.10	.13	.11
Pass. (only neg)	.84	.84	.85	.36	.45	.39	.06	.06	.06	.15	.13	.13
Random shuffling	.84	.84	.84	.26	.19	.35	.05	.05	.06	.15	.14	.14
Unaugmented	.84			.28			.05			.13		

Table A.2: Accuracy of models trained using each augmentation strategy when evaluated on HANS examples diagnostic of each of the three heuristics—lexical overlap, subsequence and constituent—for which the correct label is *non-entailment* ( $\nrightarrow$ ). Augmentation set sizes are *S* (101 examples), *M* (405) and *L* (1215). Chance performance is 0.5.

Subset of HANS	Label	Unaugmented	Small	Medium	Large
MNLi	All	0.84	0.84	0.84	0.84
Subject/object swap	$\nrightarrow$	0.19	0.53	1.00	1.00
All other	$\rightarrow$	0.96	0.93	0.77	0.77
lexical overlap	$\nrightarrow$	0.30	0.44	0.64	0.66
Subsequence	$\rightarrow$	0.99	0.99	0.84	0.85
	$\nrightarrow$	0.05	0.09	0.25	0.23
Constituent	$\rightarrow$	0.99	0.98	0.97	0.97
	$\nrightarrow$	0.13	0.17	0.23	0.18

Table A.3: Effect on HANS accuracy of augmentation using subject/object inversion with a transformed hypothesis. Results are shown for BERT fine-tuned on the MNLi training set augmented with the three size of augmentation sets (101, 405 and 1215 examples), as well as for BERT fine-tuned on the unaugmented MNLi training set.

Architecture or training method	Overall	<i>Entailment</i>			<i>Non-entailment</i>		
		<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>
Baseline (McCoy et al., 2019a)	0.57	0.96	0.99	0.99	0.28	0.05	0.13
Learned-Mixin + H (Clark et al., 2019)	0.69	0.68	0.84	0.81	0.77	0.45	0.60
DRiFt-HAND (He et al., 2019)	0.66	0.77	0.71	0.76	0.71	0.41	0.61
Product of experts (Mahabadi and Henderson, 2019)	0.67	0.94	0.96	0.98	0.62	0.19	0.30
HUBERT + (Moradshahi et al., 2019)	0.63	0.96	1.00	0.99	0.70	0.04	0.11
MT-DNN + LF (Pang et al., 2019)	0.61	0.99	0.99	0.94	0.07	0.07	0.13
BiLSTM forgettables (Yaghoobzadeh et al., 2019)	0.74	0.77	0.91	0.93	0.82	0.41	0.61
<b>Ours:</b>							
Inversion (transformed hypothesis), small	0.60	0.93	0.99	0.98	0.46	0.09	0.17
Inversion (transformed hypothesis), medium	0.63	0.77	0.84	0.97	0.71	0.25	0.23
Inversion (transformed hypothesis), large	0.62	0.77	0.85	0.97	0.73	0.23	0.18
Combined (transformed hypothesis), medium	0.65	0.92	0.96	0.98	0.64	0.13	0.26

Table A.4: HANS accuracy from various architectures and training methods, broken down by the heuristic that the example is diagnostic of and by its gold label, as well as overall accuracy on HANS. All but MT-DNN + LF use BERT as base model. *L*, *S*, and *C* stand for lexical overlap, subsequence, and constituent heuristics, respectively. Augmentation set sizes are  $n = 101$  for small,  $n = 405$  for medium, and  $n = 1215$  for large.

Subcase	Unaugmented	Small	Medium	Large
Subject-object swap <i>The senators mentioned the artist. → The artist mentioned the senators.</i>	0.19	0.53	1.00	1.00
Sentences with PPs <i>The judge behind the manager saw the doctors. → The doctors saw the manager.</i>	0.41	0.61	0.81	0.89
Sentences with relative clauses <i>The actors called the banker who the tourists saw. → The banker called the tourists.</i>	0.33	0.53	0.77	0.83
Passives <i>The senators were helped by the managers. → The senators helped the managers.</i>	0.01	0.04	0.29	0.13
Conjunctions <i>The doctors saw the presidents and the tourists. → The presidents saw the tourists.</i>	0.45	0.59	0.69	0.81
Untangling relative clauses <i>The athlete who the judges saw called the manager. → The judges saw the athlete.</i>	0.98	0.94	0.74	0.76
Sentences with PPs <i>The tourists by the actor called the authors. → The tourists called the authors.</i>	1.00	0.98	0.85	0.86
Sentences with relative clauses <i>The actors that danced encouraged the author. → The actors encouraged the author.</i>	0.99	0.98	0.89	0.89
Conjunctions <i>The secretaries saw the scientists and the actors. → The secretaries saw the actors.</i>	0.83	0.78	0.68	0.66
Passives <i>The authors were supported by the tourists. → The tourists supported the authors.</i>	1.00	0.99	0.67	0.67

Table A.5: Subject/object inversion with a transformed hypothesis: results for the HANS subcases that are diagnostic of the lexical overlap heuristic, for four training regimens—unaugmented (trained only on MNLI), and with small ( $n = 101$ ), medium ( $n = 405$ ) and large ( $n = 1215$ ) augmentation sets. Chance performance is 0.5. Top: cases in which the gold label is *non-entailment*. Bottom: cases in which the gold label is *entailment*.



Subcase	Unaugmented	Small	Medium	Large
NP/S <i>The managers heard the secretary resigned. → The managers heard the secretary.</i>	0.02	0.03	0.47	0.50
PP on subject <i>The managers near the scientist shouted. → The scientist shouted.</i>	0.12	0.21	0.21	0.23
Relative clause on subject <i>The secretary that admired the senator saw the actor. → The senator saw the actor.</i>	0.07	0.13	0.14	0.13
MV/RR <i>The senators paid in the office danced. → The senators paid in the office.</i>	0.00	0.01	0.05	0.02
NP/Z <i>Before the actors presented the doctors arrived. → The actors presented the doctors.</i>	0.06	0.09	0.41	0.25
Conjunctions <i>The actor and the professor shouted. → The professor shouted.</i>	0.98	0.96	0.87	0.86
Adjectives <i>Happy professors mentioned the lawyer. → Professors mentioned the lawyer.</i>	1.00	1.00	0.92	0.91
Understood argument <i>The author read the book. → The author read.</i>	1.00	0.99	0.97	0.97
Relative clause on object <i>The artists avoided the actors that performed. → The artists avoided the actors.</i>	0.99	0.98	0.70	0.71
PP on object <i>The authors called the judges near the doctor. → The authors called the judges.</i>	1.00	1.00	0.75	0.79

Table A.6: Subject/object inversion with a transformed hypothesis: results for the HANS subcases diagnostic of the subsequence heuristic, for four training regimens—unaugmented (trained only on MNLI), and with small ( $n = 101$ ), medium ( $n = 405$ ) and large ( $n = 1215$ ) augmentation sets. Top: cases in which the gold label is *non-entailment*. Bottom: cases in which the gold label is *entailment*.

Subcase	Unaugmented	Small	Medium	Large
Embedded under preposition <i>Unless the senators ran, the professors recommended the doctor. → The senators ran.</i>	0.41	0.43	0.57	0.49
Outside embedded clause <i>Unless the authors saw the students, the doctors resigned. → The doctors resigned.</i>	0.00	0.01	0.02	0.01
Embedded under verb <i>The tourists said that the lawyer saw the banker. → The lawyer saw the banker.</i>	0.17	0.25	0.28	0.22
Disjunction <i>The judges resigned, or the athletes saw the author. → The athletes saw the author.</i>	0.01	0.01	0.04	0.03
Adverbs <i>Probably the artists saw the authors. → The artists saw the authors.</i>	0.06	0.13	0.25	0.13
Embedded under preposition <i>Because the banker ran, the doctors saw the professors. → The banker ran.</i>	0.96	0.94	0.94	0.95
Outside embedded clause <i>Although the secretaries slept, the judges danced. → The judges danced.</i>	1.00	1.00	0.99	0.99
Embedded under verb <i>The president remembered that the actors performed. → The actors performed.</i>	0.99	0.99	0.98	0.97
Conjunction <i>The lawyer danced, and the judge supported the doctors. → The lawyer danced.</i>	1.00	1.00	0.98	0.99
Adverbs <i>Certainly the lawyers advised the manager. → The lawyers advised the manager.</i>	1.00	1.00	0.93	0.96

Table A.7: Subject/object inversion with a transformed hypothesis: results for the HANS subcases diagnostic of the constituent heuristic, for four training regimens—unaugmented (trained only on MNLI), and with small ( $n = 101$ ), medium ( $n = 405$ ) and large ( $n = 1215$ ) augmentation sets. Chance performance is 0.5. Top: cases in which the gold label is *non-entailment*. Bottom: cases in which the gold label is *entailment*.