

Understanding by Understanding Not: Modeling Negation in Language Models

Arian Hosseini

Mila/Université de Montréal
Montréal, Canada

arian.hosseini9@gmail.com

Siva Reddy

Mila/McGill University
Montréal, Canada

Dzmitry Bahdanau

Element AI
a ServiceNow Company
Montréal, Canada

R Devon Hjelm

Mila/Université de Montréal
and Microsoft Research
Montréal, Canada

Alessandro Sordoni

Microsoft Research
Montréal, Canada

Aaron Courville

Mila/Université de Montréal
Montréal, Canada

Abstract

Negation is a core construction in natural language. Despite being very successful on many tasks, state-of-the-art pre-trained language models often handle negation incorrectly. To improve language models in this regard, we propose to augment the language modeling objective with an unlikelihood objective that is based on negated generic sentences from a raw text corpus. By training BERT with the resulting combined objective we reduce the mean top 1 error rate to 4% on the negated LAMA dataset. We also see some improvements on the negated NLI benchmarks.

1 Introduction

Negation is an important property in many language understanding tasks, such as sentiment analysis, question answering, knowledge base completion and natural language inference (Kassner and Schütze, 2019; Naik et al., 2018). While Pre-trained Language Models (PLMs) such as BERT pushed the state-of-the-art on these tasks (Devlin et al., 2019; Petroni et al., 2019), they fail dramatically on instances that require understanding negation.

Kassner and Schütze (2019) show that current PLMs cannot correctly distinguish between the negated and non-negated forms of fill-in-the-blank tests. For instance, when asked to predict the [MASK] token in sentences such as “The capital of Cuba is [MASK]” and “The capital of Cuba is not [MASK]”, BERT often generate the same answer “Havana”, indicating that it may not be appropriately modeling the distribution of negative sentences. Additional evidence is given by the fact that, when fine-tuned on natural language inference tasks, PLMs tend to mis-classify examples which

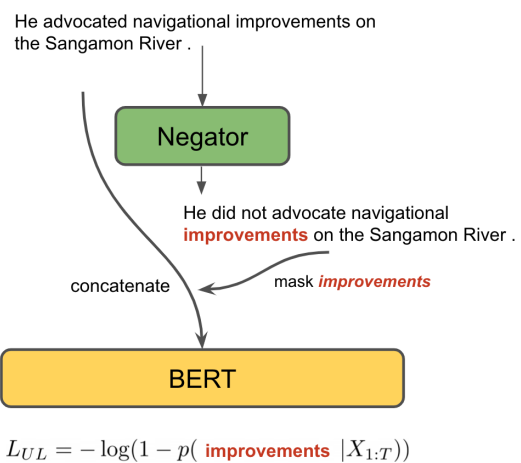


Figure 1: An overview of the unlikelihood objective. A generic sentence is negated using our data augmentation method and an unlikelihood token is chosen and replaced with [MASK]. This new sentence is concatenated with the original sentence and fed into the model. The unlikelihood loss is computed using $p(\text{improvements})$ from the language modeling head of BERT.

contain *not* or *no* as contradiction when the true label is neutral or entailment (Naik et al., 2018). Recently, Hossain et al. (2020b) proposed new natural language inference test sets to specifically target the model’s understanding of negation and show that current state-of-the-art models perform poorly on these test sets.

In this work, we investigate whether we can alleviate the modeling bias of PLMs on negated sentences. Our approach is composed of two core contributions: i) a syntactic data augmentation scheme to automatically generate negated sentences; ii) a new training paradigm, dubbed *unlikelihood training with reference* (Fig. 1), based on the recently proposed unlikelihood training (Welleck

et al., 2020).

At first, we generate a large number of negated sentences by negating sentences mined from an openly available text corpus (Wikipedia). Our sentence negator uses the dependency parse of the sentence, part of speech tags, and morphological features of each word in the sentence and deterministically negates the sentence. Given a negated version of a sentence, we replace its object with the [MASK] token and use unlikelihood training to make the object unlikely under the PLM distribution (e.g. we minimize the probability of “*improvements*” as depicted in Fig. 1). Importantly, in order to ensure that the negated sentence is *factually* false, we use the positive sentence as context (i.e., as a reference) for the unlikelihood prediction task. Concretely, we provide the concatenation of the positive and the masked negated sentence as input to the PLM. Our method can be thought of a type data augmentation, which has been shown to be effective at improving robustness across many tasks in language, such as text classification (Wei and Zou, 2019), natural language inference (Min et al., 2020; McCoy et al., 2019) and semantic parsing (Andreas, 2019).

For our negation experiments, we fine-tune pre-trained BERT with our new objective and a knowledge distillation objective. We test our model on the *negated LAMA dataset* (Kassner and Schütze, 2019), which is the negated version of knowledge probing dataset LAMA, introduced in Petroni et al. (2019). Our model achieves a mean error rate of 4% (a improvement of 5 points) on the negated LAMA dataset while maintaining the performance on the original LAMA dataset without any direct training on the negated LAMA sentences. We also fine-tune BERT on RTE (Dagan et al., 2005; Bar-Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009), SNLI (Bowman et al., 2015) and MNLI (Williams et al., 2018) tasks and achieve better results on the language inference benchmark including negation from (Hossain et al., 2020b).

2 Related Work

Pre-trained language models have shown impressive results across many tasks, such as question answering (Alberti et al., 2019) and natural language inference (Liu et al., 2019). These models are also known to encode factual and common-sense knowledge (Radford et al., 2019; Petroni et al., 2019; Bosselut et al., 2019). Despite these abilities, Kass-

ner and Schütze (2019) found that these models fail at understanding negation through analysing negated factual statements.

Extensive literature looks at the linguistic knowledge learned by language models (McCoy et al., 2019; Jumelet and Hupkes, 2018; Gulordava et al., 2018; Marvin and Linzen, 2018; Tenney et al., 2019; Warstadt and Bowman, 2019; Talmor et al., 2019). Recent work has also studied the shortcomings in negation scope detection (Jumelet and Hupkes, 2018; Fancellu et al., 2016, 2017; Morante and Daelemans, 2009; Li and Lu, 2018; Zhao and Bethard, 2020; Chen, 2019) and focus detection (Shen et al., 2019; Zou et al., 2014, 2015; Hossain et al., 2020a). Naik et al. (2018) and McCoy et al. (2019) systematically study the linguistic abilities of these models using NLI, and show that these models rely on erroneous syntactic heuristics. Our work is in this spirit for negations.

Noji and Takamura (2020) propose taking advantage of negative examples and unlikelihood in the training of language models to increase their syntactic abilities. Similarly, Min et al. (2020) show the effectiveness of syntactic data augmentation in the case of robustness in NLI. Neither of these works focus on negations.

3 Syntactic Negation Augmentation

We generate the negated versions of sentences using a syntactic augmentation method. The method gets as input the dependency parse of the sentence, POS tags and morphological information of each word and negates the sentence using a set of rules. Each rule has a dependency tree regular expression pattern (Semgrep; Chambers et al. 2007). We use Semgrep patterns to identify different syntactic templates, and then transform the sentence based on a list of actions defined in the rule. These actions can be *move*, *replace*, *insert* and *lemmatize*. The unlikelihood token which will be discussed later is also chosen using Semgrep patterns (see Appendix C for some examples).

We use Stanza (Qi et al., 2020) to get the dependency parse of the sentences, parts of speech tags, lemma, and morphological features of the words. We also filter out sentences with more than 20 words.

To test the coverage of our Semgrep patterns, we randomly sampled 930 sentences from Wikipedia. Only 31 of them did not match any of our Semgrep patterns (See table 8 in Appendix B for the number

Model	SQuAD	ConceptNet	T-REx	Google-RE
BERT	13.53	15.65	29.10	10.24
BERT + KL	13.64	15.64	29.28	10.27
BERTNOT	13.97	15.49	29.25	10.31

Table 1: Mean precision at $k = 1$ ($p @ 1$) for original LAMA queries (higher is better) of pre-trained BERT, BERT trained with distillation objective, and BERT with unlikelihood and distillation objectives (BERTNOT, sec 4.2). The scores are averaged across 3 runs.

Model	SQuAD	ConceptNet	T-REx	Google-RE
BERT	8.61	2.24	21.42	3.76
BERT + KL	4.97	1.19	21.77	3.99
BERTNOT	2.10	0.73	11.86	1.10

Table 2: Mean top 1 error rate for negated LAMA queries (lower is better) of pre-trained BERT, BERT trained with distillation objective, and BERT with unlikelihood and distillation objectives (BERTNOT, sec 4.2). The scores are averaged across 3 runs.

of matches for each rule in our rule set for these 930 sentences). In addition, to get a better sense of the correctness of our method, 100 random sentences (from Wikipedia) were negated and reviewed by a native English speaker. The precision for these negations is 94.00%. Table 7 in Appendix B shows examples of original and negated sentences.

4 Unlikelihood Training With Reference

4.1 Reference setup

Applying unlikelihood to a word in any random sentence is problematic, unless the sentence is a factual statement (e.g. unlikelihood on *improvements* in “He did not advocate navigational **improvements** on the Sangamon River.” in Fig 1 is problematic as this sentence is not grounded in reality). Moreover, using solely factual sentences limits the application of this method.¹ To be able to use any generic (not necessarily factual) sentence and pick an unlikelihood token in it, there needs to be some sort of grounding or context. In this setup, each training example is of the form $\langle \textit{sentence A}, \textit{sentence B} \rangle$ where *sentence A* is the reference for *sentence B*, and provides the grounding or context for it.

4.2 Unlikelihood and knowledge distillation

The unlikelihood loss has recently been proposed by Welleck et al. (2020) to mitigate the problem of repetition in neural text generation. Noji and Takamura (2020) also adopted this loss to penalize the desirability of an incorrect token in a sentence.

¹We did try to apply unlikelihood without any context or reference, but as expected it performed poorly for both LAMA and negated LAMA. See appendix E.

We adopt this method to penalize the likelihood of a token in *sentence B* that makes this sentence contradictory with the reference *sentence A*.

- (1) A Humans have a rational soul.
B Humans do not have a rational **soul**.

In the example 1, assuming that *sentence A* is true, we want the model to avoid assigning “soul” in *sentence B* a high probability. To this end, the probability of the unlikelihood token $x_u = \text{“soul”}$ is penalized with the unlikelihood loss L_{UL} as:

$$L_{UL}(x_u) = -\log(1 - p(x_u|x_{1:T})), \quad (1)$$

where $x_{1:T}$ is the whole input sequence (*sentence A* concatenated with *sentence B* which is the negated version of *sentence A* as illustrated in Fig 1). To have a balanced augmentation data set, we also include examples where *sentence B* is the copy of *sentence A* and therefore not contradictory with it. In this context, we want the model to perform as it was untouched (before any fine-tuning). The KL divergence knowledge distillation loss is used for these examples on the same token:

- (2) A Humans have a rational soul.
B Humans have a rational [MASK].

The loss L_{KL} for token $x_l = \text{“[MASK]”}$ is written as:

$$L_{KL}(x_l) = D_{KL}(p_{LM} || p) \quad (2)$$

where p_{LM} is the probability distribution over the vocabulary for the masked token x_l under the LM before any fine-tuning.

In our experiments, we use the BERT-base model and further train it with two objectives, the unlikelihood objective (Eq. 1) and the knowledge

Query	Top 3 words with log probs from BERT	Top 3 words with log probs from BERTNOT
iOS is developed by [MASK].	Apple (-1.8), Google (-2.6), Microsoft (-2.8)	Apple (-1.8), Google (-2.5), Microsoft (-2.7)
iOS is not developed by [MASK].	Apple (-1.8), Google (-2.6), Microsoft (-2.8)	Microsoft (-1.8), Google (-2.4), Apple (-3.1)
The majority of the amazon forest is in [MASK].	Brazil (-2.6), Bolivia (-2.7), Madagascar (-3.1)	Brazil (-2.9), Bolivia (-3.1), Mexico (-3.2)
The majority of the amazon forest is not in [MASK].	cultivation (-1.0), Brazil (-3.5), Mexico (-3.5)	cultivation (-2.0), Mexico (-4.1), France (-4.3)
Charles Nodier died in [MASK].	Paris (-1.35), Rome (-3.2), office (-3.4)	Paris (-1.5), Rome (-3.3), France (-3.6)
Charles Nodier did not die in [MASK].	Paris (-2.4), office (-2.7), France (-2.8)	vain (-3.5), error (-4.0), doubt (-4.5)
Mac OS is developed by [MASK].	Apple (-1.9), Microsoft (-2.0), Intel (-2.0)	Apple (-2.0), Microsoft (-2.0), Intel (-2.1)
Mac OS is not developed by [MASK].	Apple (-1.3), Microsoft (-1.5), IBM (-2.3)	Microsoft (-2.1), IBM (-2.7), itself (-3.4)

Table 3: Examples from BERT base before and after training it with the unlikelihood (UL) and KL divergence knowledge distillation (KL) objectives (BERTNOT). Queries are from LAMA and negated LAMA.

distillation objective (Eq. 2). We also use original Wikipedia sentences for the latter to prevent catastrophic forgetting of language modeling. The probability of the unlikelihood token $p(x_u|x_{1:T})$ and the distribution for masked token x_l are computed using the language modeling head of the BERT model by replacing x_u and x_l in the input sequences with the [MASK] token. Examples for each objective are sampled uniformly. We will refer to our model as BERTNOT.

5 Experiments

We report our main results on LAMA and Negated LAMA for knowledge base completion. The cloze statements from LAMA are facts or commonsense knowledge generated from either subject-relation-object triples (X, rel, Y) or question-answers pairs. The cloze statements for the triples are generated using a template for each relation which includes the placeholders X and Y (e.g. “X is located in Y”). X is replaced for the subject and Y is replaced with the [MASK] token to be predicted by the model. In the question-answer pairs, the answer is replaced with [MASK] token. The facts in the LAMA dataset are from multiple sources: 1) Google-RE relations, namely “place of birth”, “date of birth” and “place of death”; 2) T-REx, a subset of Wikidata triples with 41 relations (ElSahar et al., 2018); 3) ConceptNet with 16 relations (Li et al., 2016); 4) SQuAD, a subset of 305 context-insensitive questions manually rephrased as cloze-style questions (Rajpurkar et al., 2016). Negated LAMA was created by manually negating the templates or questions (Kassner and Schütze, 2019). Following Petroni et al. (2019) we use mean precision at k ($P @ k$) for LAMA. For negated LAMA we report mean top 1 error rate.

5.1 Knowledge Base Completion

As discussed in section 4.2, we train a pre-trained BERT base cased model for 5 epochs, with 20k examples for each objective, a maximum sequence length of 128 and a learning rate of 1e-5. To see the effects of the unlikelihood objective more clearly, we also train a pre-trained BERT base cased model with only the KL knowledge distillation objective with the same data and hyper-parameters.

Tables 1 and 2 respectively show the mean precision at rank 1 (averaged over all the relations) for LAMA, and mean top 1 error rate for negated LAMA queries.² The mean error rate on the negated LAMA queries decreases to below 4% while the results on original LAMA stay the same. These results are achieved without any direct training on LAMA queries (negated or non-negated). Table 3 shows the top 3 predicted words for a pre-trained BERT model and the model trained with our method. Pre-trained BERT seems to ignore negation and mostly predict based on the subject of the query, but the prediction probability in the negated queries seems to be generally lower. Our method is as good as the vanilla model (BERT) on original queries. For the negated queries, our model predictions are far-superior than the vanilla model. We also tried out method on BERT-large. See appendix E for results and discussion.

5.2 Natural Language Inference

We fine-tune our model with a language inference objective on RTE, SNLI and MNLI tasks. Table 4 shows the accuracies on the original development splits and the new splits from Hossain et al. (2020b) containing negation for each task. We used the hyper-parameters from Hossain et al. (2020b) to fine-tune all of our models.

²Baseline scores differ slightly from Petroni et al. (2019). We were unable to get the same results with their code.

Model	RTE		SNLI		MNLI	
	dev	w/neg	dev	w/neg	dev	w/neg
BERT	70.04 ± 1.57	65.47 ± 3.63	89.47 ± 0.18	44.18 ± 0.67	82.95 ± 0.18	60.62 ± 1.32
BERTNOT	69.68 ± 1.88	74.47 ± 0.29	89.00 ± 0.10	45.96 ± 0.41	84.31 ± 2.29	60.89 ± 0.31

Table 4: Accuracies on original development splits (dev) and new splits containing negation from Hossain et al. (2020b) (w/neg) for RTE, SNLI and MNLI (matched genres) tasks. Results are averaged across 3 runs.

	Premise	Hypothesis	T	B	BN
1	It does not use the first day of the first month of the Lunar Year as the start of the Chinese New Year.	The Chinese New Year’s Day falls on the first day of the first month of the Lunar Year.	N	E	N
2	The prosecutor told the court that the incident had caused "distress" to one of the children.	The prosecutor did not tell the court that "distress" in one of the children is associated with the incident.	N	E	N
3	Green cards are not becoming more difficult to obtain.	Green card is now difficult to receive.	N	E	N
4	Moog’s synthesiser, which bears his name, revolutionised music from the 1960s onwards, and was used by bands like The Beatles and The Doors.	Moog’s instruments were not used by The Beatles and The Doors among others.	N	N	E
5	The board of Marks & Spencer will not take another look at Philip Green’s increased takeover offer.	Philip Green does not try to take over Marks & Spencer.	E	E	N
6	Albert Sabin developed an oral, attenuated (live) vaccine, which, with Salk’s discovery, did not bring polio under control.	Polio is not under control in the world.	E	E	N

Table 5: Examples from the new split from Hossain et al. (2020b) containing negation for RTE. T, B and BN denote true label, BERT’s prediction and BERTNOT’s prediction respectively. E and N are used for entailment and not entailment labels.

Our model achieves superior results on RTE (low-resource setting) and slightly better accuracies on SNLI and MNLI (high-resource setting) on all the new splits containing negation, while keeping roughly the same scores on the original dev splits. We conjecture that fine-tuning on large-amounts of data (SNLI and MNLI) may have resulted in catastrophic forgetting of the negation knowledge, decreasing the gap between BERT and BERTNOT. We tried to alleviate the catastrophic forgetting by mixing in some unlikelihood training and knowledge distillation along the NLI training, but that did not help. You can see these results for MNLI in appendix D. We leave further exploration of better fine-tuning objectives while preserving the pretrained knowledge for future work.

Table 5 shows some of the examples of the new RTE split containing negation from Hossain et al. (2020b), along with the predictions from BERT and BERTNOT. Examples 4 and 6 show the failure cases of BERTNOT. As it can be seen, for the fifth example, the true label is incorrect, but BERTNOT predicts the correct label for this pair of premise and hypothesis.

6 Conclusion

In this work, we propose a combination of the unlikelihood objective with a reference based setup for input sentences to model negation. This allows us to utilize generic sentences, and negate them with our data augmentation method to be used as examples for the unlikelihood objective. Our method notably improves the error rate on the negated LAMA dataset while keeping the same performance on the original LAMA queries.

We also test our method on the original development sets and new splits containing negation from Hossain et al. (2020b) of RTE, SNLI and MNLI tasks. We see large improvements on the negated splits in low-resource setting (RTE) and slight improvements in high-resource setting (SNLI and MNLI), while also maintaining similar results as BERT on original splits.

References

- Chris Alberti, Kenton Lee, and Michael Collins. 2019. [A BERT baseline for the natural questions](#). *CoRR*, abs/1901.08634.
- Jacob Andreas. 2019. [Good-enough compositional data augmentation](#). *CoRR*, abs/1904.09545.
- Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro, and Danilo Giampiccolo. 2006. The second pascal recognising textual entailment challenge. *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2009. B.: The fifth pascal recognizing textual entailment challenge. In *Proceedings of TAC 9*.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. [COMET: Commonsense transformers for automatic knowledge graph construction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- 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.
- Nathanael Chambers, Daniel M. Cer, Trond Grenager, David Hall, Chloé Kiddon, Bill MacCartney, Marie-Catherine de Marneffe, Daniel Ramage, Eric Yeh, and Christopher D. Manning. 2007. Learning alignments and leveraging natural logic. In *ACL-PASCAL@ACL*, pages 165–170. Association for Computational Linguistics.
- Long Chen. 2019. [Attention-based deep learning system for negation and assertion detection in clinical notes](#). *International Journal of Artificial Intelligence & Applications*, 10:1–9.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The PASCAL recognising textual entailment challenge](#). In *Machine Learning Challenges, Evaluating Predictive Uncertainty, Visual Object Classification and Recognizing Textual Entailment, First PASCAL Machine Learning Challenges Workshop, MLCW 2005, Southampton, UK, April 11-13, 2005, Revised Selected Papers*, volume 3944 of *Lecture Notes in Computer Science*, pages 177–190. Springer.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics.
- Hady ElSahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon S. Hare, Frédérique Laforest, and Elena Simperl. 2018. [T-rex: A large scale alignment of natural language with knowledge base triples](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Federico Fancellu, Adam Lopez, and Bonnie Webber. 2016. [Neural networks for negation scope detection](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 495–504, Berlin, Germany. Association for Computational Linguistics.
- Federico Fancellu, Adam Lopez, Bonnie Webber, and Hangfeng He. 2017. [Detecting negation scope is easy, except when it isn't](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 58–63, Valencia, Spain. Association for Computational Linguistics.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- 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.
- Md Mosharaf Hossain, Kathleen Hamilton, Alexis Palmer, and Eduardo Blanco. 2020a. [Predicting the focus of negation: Model and error analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8389–8401, Online. Association for Computational Linguistics.
- Md Mosharaf Hossain, Venelin Kovatchev, Pranoy Dutta, Tiffany Kao, Elizabeth Wei, and Eduardo Blanco. 2020b. [An analysis of natural language inference benchmarks through the lens of negation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9106–9118. Association for Computational Linguistics.
- Jaap Jumelet and Dieuwke Hupkes. 2018. [Do language models understand anything? on the ability of LSTMs to understand negative polarity items](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 222–231, Brussels, Belgium. Association for Computational Linguistics.

- Nora Kassner and Hinrich Schütze. 2019. [Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly.](#)
- Hao Li and Wei Lu. 2018. [Learning with structured representations for negation scope extraction.](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 533–539, Melbourne, Australia. Association for Computational Linguistics.
- Xiang Li, Aynaz Taheri, Lifu Tu, and Kevin Gimpel. 2016. [Commonsense knowledge base completion.](#) In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach.](#) *CoRR*, abs/1907.11692.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference.](#) In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3428–3448. Association for Computational Linguistics.
- Junghyun Min, R. Thomas McCoy, Dipanjan Das, Emily Pitler, and Tal Linzen. 2020. [Syntactic data augmentation increases robustness to inference heuristics.](#) *CoRR*, abs/2004.11999.
- Roser Morante and Walter Daelemans. 2009. [A meta-learning approach to processing the scope of negation.](#) In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning, CoNLL 2009, Boulder, Colorado, USA, June 4-5, 2009*, pages 21–29. ACL.
- Aakanksha Naik, Abhilasha Ravichander, Norman M. Sadeh, Carolyn Penstein Rosé, and Graham Neubig. 2018. [Stress test evaluation for natural language inference.](#) In *COLING*, pages 2340–2353. Association for Computational Linguistics.
- Hiroshi Noji and Hiroya Takamura. 2020. [An analysis of the utility of explicit negative examples to improve the syntactic abilities of neural language models.](#) *CoRR*, abs/2004.02451.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. [Language models as knowledge bases?](#) In *EMNLP/IJCNLP (1)*, pages 2463–2473. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages.](#) *CoRR*, abs/2003.07082.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners.](#)
- 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, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Longxiang Shen, Bowei Zou, Yu Hong, Guodong Zhou, Qiaoming Zhu, and AiTi Aw. 2019. [Negative focus detection via contextual attention mechanism.](#) 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 2019, Hong Kong, China, November 3-7, 2019*, pages 2251–2261. Association for Computational Linguistics.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2019. [olmpics - on what language model pre-training captures.](#) *CoRR*, abs/1912.13283.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [Bert rediscovered the classical nlp pipeline.](#) In *Association for Computational Linguistics*.
- Alex Warstadt and Samuel R. Bowman. 2019. [Grammatical analysis of pretrained sentence encoders with acceptability judgments.](#) *CoRR*, abs/1901.03438.
- Jason W. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 6381–6387. Association for Computational Linguistics.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training.](#) In *ICLR*. OpenReview.net.
- 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.

Yiyun Zhao and Steven Bethard. 2020. [How does BERT's attention change when you fine-tune? an analysis methodology and a case study in negation scope.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4729–4747, Online. Association for Computational Linguistics.

Bowei Zou, Guodong Zhou, and Qiaoming Zhu. 2014. [Negation focus identification with contextual discourse information.](#) In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 522–530. The Association for Computer Linguistics.

Bowei Zou, Guodong Zhou, and Qiaoming Zhu. 2015. [Unsupervised negation focus identification with word-topic graph model.](#) In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1632–1636. The Association for Computational Linguistics.

A Training details

Here are the hyper-parameters used in our fine-tunings.

Task	Epochs	Batch Size	Learning Rate	Weight Decay
Unlikelihood training	5	32	1e-5	N/A
RTE	50	32	2e-5	N/A
SNLI	3	32	1e-5	0.1
MNLI	3	32	2e-5	N/A

Table 6: Hyper-parameters

Algorithm 1 shows the details of further training the BERT base cased model with the unlikelihood and knowledge distillation objectives.

```
number of training steps      :T
for  $i \leftarrow 1$  to  $T$  do
   $L_{UL} \leftarrow$  compute unlikelihood loss with contradictory  $\langle \textit{sentence A}, \textit{sentence B} \rangle$  pairs;
   $L_{KL} \leftarrow$  compute knowledge distillation loss with non-contradictory  $\langle \textit{sentence A}, \textit{sentence B} \rangle$  pairs;
   $g_1 \leftarrow$  compute gradient of  $\gamma L_{UL} + (1 - \gamma)L_{KL}$ ;
  update the parameters with  $g_1$ ;
   $L_{KL} \leftarrow$  compute knowledge distillation loss with sentences from Wikipedia;
   $g_2 \leftarrow$  compute gradient of  $L_{KL}$ ;
  update the parameters with  $g_2$ ;
end
```

Algorithm 1: Details of the training procedure of BERTNOT. The unlikelihood loss and knowledge distillation loss are first computed with the $\langle \textit{sentence A}, \textit{sentence B} \rangle$ inputs. These inputs are contradictory for the UL loss, and non-contradictory for knowledge distillation (sec 4.2). We use $\gamma = 0.4$ in our experiments to sum these losses and compute the gradient g_1 . Then, we compute the knowledge distillation loss for inputs sampled from Wikipedia. These inputs do not have our reference based format. The parameters are updated again using the gradient from this knowledge distillation loss (g_2).

B Examples of negated sentences

Here are some examples and details of our syntactic negation method.

	Original	Negated	Unlikelihood Token
1	That tournament helped demonstrate the high caliber of play in women’s soccer.	That tournament did not help demonstrate the high caliber of play in women’s soccer.	tournament
2	The attributes of this vector (length and direction) characterize the rotation at that point.	The attributes of this vector (length and direction) do not characterize the rotation at that point.	rotation
3	This was broadcast live on Norway’s main national TV carrier NRK.	This was not broadcast live on Norway’s main national TV carrier NRK.	Norway
4	The latter may occur implicitly through the use of a construct like DEFVAR or DEFPARAMETER.	The latter may not occur implicitly through the use of a construct like DEFVAR or DEFPARAMETER.	latter
5	When Arjuna was fighting Karna, the latter’s chariot’s wheels sank into the ground.	When Arjuna was fighting Karna, the latter’s chariot’s wheels did not sink into the ground.	wheels
6	It also prohibits or restricts the use of certain accounts held at financial institutions.	It also does not prohibit or restricts the use of certain accounts held at financial institutions.	use

Table 7: Examples of original and negated sentences with the chosen unlikelihood token. Examples 5 and 6 are incorrect negations since *sank* in example 5 and *restricts* in example 6 are incorrect word forms in the negated context.

Rule Name	# of Sentences Matched
simple past	315
simple present	295
Imperative	93
present with auxiliary verb	37
past perfect	35
copula statements	34
present with modal	24
already negated with not	14
NPI words (anywhere, anyone, etc)	5
negative words (no, nobody, etc)	4
other	13

Table 8: Number of matches for each rule in our rule set over 930 sentences used to analyze the syntactic negation.

C Example rules for transforming a sentence into its negation

Original Sentence	Rule	Negated Sentence
Nowhere in his confession did he mention the Monteagle letter.	<pre>{ "name": "aux before subj", "pattern": "{\$;tag:/VB.*/}=A >/ advmod cc/ {word:/never nobody no nothing nowhere neither Never Nobody No Nothing Nowhere Neither/}=npiword >/aux.*/ ({= B \$++ {}=subject) >/nsbj.*/ { }=subject ?>obj {tag:/NN.*/}= object", "actions": [{ "type": "move", "to_move": "B", "anchor": "A", "position": "before" }, { "type": "replace", "token": "", "to_replace": "npiword" }] }</pre>	in his confession he did mention the Monteagle letter.
Many fonts then made the right leg vertical.	<pre>{ "name": "simple past", "pattern": "{\$;cpos:/.*Tense=Past.*/ }=A >/nsbj csubj/=E {}=subject ?>obj {tag:/NN.*/}=object", "actions": [{ "type": "insert", "token": "did", "rel": "AUX", "anchor": "A", "position": "before" }, { "type": "insert", "token": "not", "rel": "ADV", "anchor": "A", "position": "before" }, { "type": "lemmatize" }] }</pre>	Many fonts then did not make the right leg vertical.

Table 9: Examples of how the syntactic negation augmentation method works. For the first sentence, the matched rule has two actions, *move* and *replace*. The *move* action has moved the token **B** = *did* before token **A** = *mention*. The *replace* action has replaced **npiword** = *Nowhere* with an empty token, which means removing this token. The token **object** = *letter* is chosen as the unlikely token in this sentence.

In the second sentence, the matched rule has three actions, two *inserts* and one *lemmatize* action. The *insert* actions, add the tokens “did not” before **A** = *made*, and the token **A** = *made* is replaced with its lemma by the *lemmatize* action. The token **object** = *leg* is chosen as the unlikely token in the negated sentence.

D Mixing negation unlikelihood training and knowledge distillation with NLI training

In order to reduce the catastrophic forgetting behavior of the model during NLI training, we added the unlikelihood, knowledge distillation and MLM objectives to the original NLI classification objective and trained the model with the same hyper-parameters for the MNLI task. We also trained one version with only the original NLI classification objective and the MLM objective. As the results in table 10 show, this method did not improve the scores for development split and the new split containing negation from Hossain et al. (2020b) for MNLI.

Model	MNLI	
	dev	w/neg
BERTNOT + UL + KL + MLM + NLI obj	81.17	60.20
BERTNOT + MLM + NLI obj	81.42	62.00

Table 10: Accuracies on original development split (dev) and new split containing negation from Hossain et al. (2020b) (w/neg) for MNLI (matched genres) task.

E Supplementary Results

Model	lr	SQuAD	ConceptNet	T-REx	Google-RE
BERTNOT without reference setup	1e-5	13.86	15.65	29.54	10.29
BERT-large	1e-5	16.83	19.26	30.76	10.93
BERTNOT-large	1e-5	14.19	19.14	32.09	11.02
BERTNOT-large	5e-5	15.18	16.97	30.71	10.62
BERTNOT-large	1e-4	11.55	13.58	28.41	9.25

Table 11: Mean precision at $k = 1$ ($p @ 1$) for original LAMA queries (higher is better) of BERT with unlikelihood and distillation objectives without references for sentences, BERT-large, and BERT-large with unlikelihood and distillation objectives with different learning rates.

Model	lr	SQuAD	ConceptNet	T-REx	Google-RE
BERTNOT without reference setup	1e-5	5.96	1.34	21.54	3.73
BERT-large	1e-5	7.95	1.67	22.97	4.13
BERTNOT-large	1e-5	8.28	1.87	23.49	4.22
BERTNOT-large	5e-5	8.28	2.20	24.05	4.09
BERTNOT-large	1e-4	4.97	1.47	20.86	3.60

Table 12: Mean top 1 error rate for negated LAMA queries (lower is better) of BERT with unlikelihood and distillation objectives without references for sentences, BERT-large, and BERT-large with unlikelihood and distillation objectives with different learning rates.

As the results in table 12 show, pre-trained BERT-large performs worse than pre-trained BERT-base on negated LAMA queries. We decreased the batch-size to be able to fine-tune BERT-large. As the scores for negated LAMA queries from table 12 show, fine-tuning BERT-large with our method using the same or slightly larger learning rate does not improve the results. We observe a decrease in the mean top 1 error rates for negated LAMA queries when we use a larger learning rate ($1e - 5$), but this also hinders the performance of the model on the original LAMA queries (table 11). This requires some hyper-parameter tuning and further investigation.