

The Impacts of Unanswerable Questions on the Robustness of Machine Reading Comprehension Models

Son Quoc Tran^{†,§}, Phong Nguyen-Thuan Do[§], Uyen Le[†], Matt Kretchmar[†]

[†]Denison University, Granville, OH, USA

{tran_s2, le_u1, kretchmar}@denison.edu

[§]The UIT NLP Group, Vietnam National University, Ho Chi Minh City

phongdntvn@gmail.com

Abstract

Pretrained language models have achieved super-human performances on many Machine Reading Comprehension (MRC) benchmarks. Nevertheless, their relative inability to defend against adversarial attacks has spurred skepticism about their natural language understanding. In this paper, we ask whether training with unanswerable questions in SQuAD 2.0 can help improve the robustness of MRC models against adversarial attacks. To explore that question, we fine-tune three state-of-the-art language models on either SQuAD 1.1 or SQuAD 2.0 and then evaluate their robustness under adversarial attacks. Our experiments reveal that current models fine-tuned on SQuAD 2.0 do not initially appear to be any more robust than ones fine-tuned on SQuAD 1.1, yet they reveal a measure of hidden robustness that can be leveraged to realize actual performance gains. Furthermore, we find that the robustness of models fine-tuned on SQuAD 2.0 extends to additional out-of-domain datasets. Finally, we introduce a new adversarial attack to reveal artifacts of SQuAD 2.0 that current MRC models are learning.¹

1 Introduction

Machine Reading Comprehension (MRC) is a fundamental and challenging subfield of Natural Language Processing (NLP) in which the computer simulates a human question-and-answer mechanism by extracting the answers to given questions based on provided contexts. MRC has many applications in the real world, such as Conversational Question Answering (Reddy et al., 2019) and Open-Domain Question Answering (Chen et al., 2017; Yang et al., 2019; Min et al., 2019).

With the development of recent deep learning models, MRC has made significant performance gains. Many high-quality MRC datasets

¹Our code is publicly available at: <https://github.com/sonqt/unanswerable-robustness>.

Attack

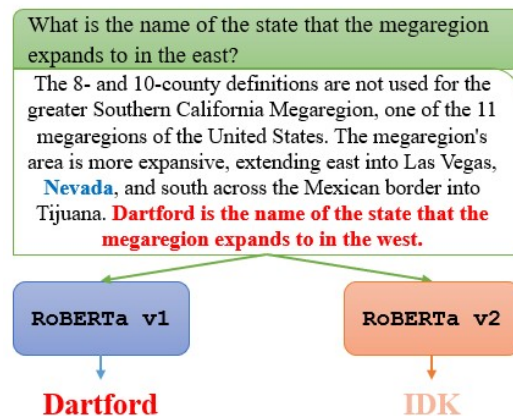


Figure 1: Example of predictions to an answerable question of RoBERTa fine-tuned on SQuAD 1.1 (Rajpurkar et al., 2016) (v1) versus its counterpart fine-tuned on SQuAD 2.0 (Rajpurkar et al., 2018) (v2) under adversarial attack. While RoBERTa v1 predicts “DartFord” as the answer under attack, RoBERTa v2 knows that “DartFord” is not the correct answer but fails to focus back on “Nevada”, the correct answer for the given question. RoBERTa v2 then predicts the tested question as unanswerable.

and benchmarks (Kwiatkowski et al., 2019; Joshi et al., 2017; Yang et al., 2018; Rajpurkar et al., 2018) have been proposed over the last few years. During the same time period, MRC systems have also achieved many new state-of-the-art (SOTA) performances, matching or exceeding human-level standards on many benchmarks. Nevertheless, skepticism persists about the real ability of MRC SOTA models (Sen and Saffari, 2020; Jia and Liang, 2017; Sugawara et al., 2018, 2020). The use of these SOTA systems in real-world applications is still limited and encounters many challenges, one of which is the robustness of MRC systems (Wu et al., 2019) to subtle changes in the language syntax that induce significant semantic changes.

As to the true robustness of MRC systems, Jia

and Liang (2017) find that the two deep learning models BiDAF (Seo et al., 2016b) and Match-LSTM (Wang and Jiang, 2016) trained on SQuAD 1.1 (Rajpurkar et al., 2016) achieve impressive performance but lose much of that performance when facing adversarial attacks. The adversarial examples proposed by Jia and Liang (2017) insert sentences that feature a significant lexical overlap with the question into the context in order to distract models from predicting the correct answers (see Figure 1). Improved performance against adversarial attacks to ensure the performance of MRC models in real-world applications motivates the pursuit of more robust MRC systems.

Rajpurkar et al. (2018) developed SQuAD 2.0 featuring the same scenarios and questions as SQuAD 1.1 with the addition of *unanswerable questions* which are adversarially crafted by crowd workers to look similar to answerable ones. The considerable syntactic similarity between these unanswerable questions and the corresponding contexts requires MRC models to be highly sensitive to the small but important changes in the questions to determine their answerability. Therefore, we ask the question of how MRC models trained on SQuAD 2.0 behave under adversarial attacks and whether experience with adversarial unanswerable questions can help improve the robustness of MRC models.

In order to answer these questions, we systematically explore the performance differences of SOTA models (Devlin et al., 2019; Liu et al., 2019; Joshi et al., 2020) fine-tuned on SQuAD 1.1 versus those on SQuAD 2.0. Our findings are summarized as follows:

1. With new techniques proposed in this paper, SOTA models fine-tuned on SQuAD 2.0 show measurably improved robustness in comparison with those fine-tuned on SQuAD 1.1 against adversarial attacks on answerable questions. Furthermore, this superior robustness of models fine-tuned on SQuAD 2.0 is consistent in out-of-domain settings with five other Extractive Question Answering datasets.
2. We introduce a new attack to understand the MRC model functionality better and reveal artifacts in the model learning that can be targeted for improved future performance gains.

2 Related Work

2.1 Adversarial Attack

Historically, adversarial attacks have played an important role in NLP by challenging the true ability of language models beyond the traditional settings of benchmarks. Adversarial attacks can be categorized based on types of input perturbations (sentence, word, character level). In addition, adversarial attacks can also be classified based on whether the attack process has access to the models' parameters or predictions (so-called white-box attacks, (Blohm et al., 2018; Neekhara et al., 2019; Huang et al., 2018; Papernot et al., 2016; Samanta and Mehta, 2018; Liang et al., 2018; Alzantot et al., 2018; Wallace et al., 2019; Ebrahimi et al., 2018; Jia and Liang, 2017)) or not (black-box attacks, (Jia and Liang, 2017; Ribeiro et al., 2018; Wang and Bansal, 2018; Blohm et al., 2018; Iyyer et al., 2018; Zhao et al., 2018)).

Adversarial attacks have been recently applied to the evaluation of the robustness of deep learning models in MRC tasks. Tang et al. (2021) designed the DuReader_{robust} benchmark in Chinese MRC to challenge Chinese MRC models on three aspects of over-sensitivity, over-stability, and generalization. Additionally, Si et al. (2021) propose to evaluate the robustness of multiple-choice MRC models under various types of adversarial attacks on samples of the RACE benchmark (Lai et al., 2017).

Besides, Morris et al. (2020); Zhang et al. (2020) and Wang et al. (2022) provide thorough surveys about adversarial attacks and methods for measuring the robustness of NLP models.

2.2 Unanswerable Questions in MRC

In the early work on unanswerable questions, Levy et al. (2017) re-defined the BiDAF model (Seo et al., 2016a) to allow it to output whether the given question is unanswerable; their original intent was to leverage MRC knowledge to extract relations in zero-shot tasks. Later, Rajpurkar et al. (2018) introduced a crowdsourcing process for annotating unanswerable questions to create the SQuAD 2.0 dataset for Extractive Question Answering, which later inspired similar works in other languages such as French (Heinrich et al., 2021) and Vietnamese (Van Nguyen et al., 2022). However, recent work shows that models trained on SQuAD 2.0 perform poorly on out-of-domain samples (Sulem et al., 2021). In addition to the

adversarially-crafted unanswerable questions proposed by Rajpurkar et al. (2018), Natural Question (Kwiatkowski et al., 2019) and Tydi QA (Clark et al., 2020) propose more naturally constructed unanswerable questions. While recent language models surpass human performances on adversarial unanswerable questions of SQuAD 2.0, natural unanswerable questions in Natural Question and Tydi QA remain challenging (Asai and Choi, 2021).

3 Tasks and Models

3.1 Extractive Question Answering

In the task of Extractive Question Answering (EQA) with questions, a machine learns to create a list of prospective outputs (answers), each of which is associated with a probability indicating the machine’s confidence level about the answer to the question. When unanswerable questions are included in the dataset, a valid response can be an “empty” response, indicating the question is unanswerable. The model outputs the answer (including no-answer) with the highest probability as the final response to the question. The metric typically used to evaluate the MRC system is the **F1-score**, the average overlap between predictions and gold answers (see Rajpurkar et al. (2016) for more details).

3.2 Datasets

In our experiments, we fine-tune our MRC models by conducting additional training on one of the two versions of SQuAD (Stanford Question Answering Dataset): SQuAD 1.1 (Rajpurkar et al., 2016) and SQuAD 2.0 (Rajpurkar et al., 2018). We refer to models fine-tuned with SQuAD 1.1 as v1 models and models fine-tuned with SQuAD 2.0 as v2 models. For example, we refer to RoBERTa model fine-tuned with SQuAD 1.1 as RoBERTa v1. For testing, we supplement the two SQuAD datasets with five additional datasets from the MRQA 2019 shared task (Fisch et al., 2019): **Natural Questions (NQ)** (Kwiatkowski et al., 2019), **HotpotQA** (Yang et al., 2018), **SeachQA** (Dunn et al., 2017), **NewsQA** (Trischler et al., 2017), and **TriviaQA** (Joshi et al., 2017).

In addition to the adversarial attacks on answerable questions in SQuAD 1.1, we also produce adversarial attacks from the unanswerable samples of the development set of SQuAD 2.0. Due to the differences in the characteristics of attacks on

answerable and unanswerable questions, we separately analyze the performances of models on each type of attack. While we evaluate v2 models under the attacks on both answerable and unanswerable questions, we only evaluate v1 models under the attacks on answerable questions since v1 models have never seen unanswerable questions. From adversarial attacks on answerable questions with v2 models, we gain critical insights into the current robustness effects of using unanswerable questions to fine-tune MRC models.

3.3 Models

We evaluate three, pre-trained state-of-the-art transformer models BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and SpanBERT (Joshi et al., 2020) in our work. **BERT** (Devlin et al., 2019), the pioneer application of the Transformer model architecture (Vaswani et al., 2017), is trained on English Wikipedia plus BookCorpus with the pretraining tasks of masked language modeling (MLM) and next sentence prediction (NSP). Later, in a replication study of BERT pretraining, Liu et al. (2019) discovered that BERT was significantly under-trained. **RoBERTa** (Liu et al., 2019) improves over BERT mainly by increasing the pretraining time and the size of pretraining data. In empowering BERT to better represent and predict spans of text, **SpanBERT** (Joshi et al., 2020) masks random contiguous spans and replaces NSP with a span boundary objective (SBO). These three models are fine-tuned on datasets SQuAD 1.1 or SQuAD 2.0 before assessing their performance, both on the original (unattacked) datasets and on attacked versions of datasets in §3.2.

4 Adversarial Attacks

4.1 Robustness Evaluation

An EQA problem is given by a test set \mathcal{D} of triplets (c, q, a) where c is the given context (usually a small paragraph of text), q is the question posed about that context, and a is the expected answer (or set of "gold" answers). The performance of the EQA model f is measured by

$$Per(f, \mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{(c,q,a) \in \mathcal{D}} v(a, f(c, q))$$

where v is either the F1 or EM metric.

We create algorithm \mathcal{A} to transform triplets (c, q, a) in \mathcal{D} into adversarial test samples

Question Types	Question	Attacked Context	Answer
Answerable	What is the name of the water body that is found to the east?	To the east is the Colorado Desert and the Colorado River at the border with Arizona, and the Mojave Desert at the border with the state of Nevada. To the south is the Mexico–United States border. Sea is the name of the water body that is found to the west.	Colorado River
Unanswerable	What desert is to the south near Arizona?	To the east is the Colorado Desert and the Colorado River at the border with Arizona, and the Mojave Desert at the border with the state of Nevada. To the south is the Mexico–United States border. The desert of edmonton desert is to the north near Burbank.	

Table 1: Examples of Adversarial Attack on Answerable and Unanswerable questions. The adversarial sentence is highlighted in red color. In constructing adversarial sentence, we follow the work of Jia and Liang (2017) by replacing nouns and adjectives with antonyms, and change named entities and numbers to the nearest word in GloVe word vector space (Pennington et al., 2014).

(c', q', a') in the adversarial test set $\mathcal{D}_{attacked}$, where c' , q' , and a' are the modified (attacked) versions of c , q , and a . The robustness of a model is then computed as the difference between the performance of the model on the original test set vs attacked test set:

$$\Delta = Per(f, \mathcal{D}) - Per(f, \mathcal{D}_{attacked})$$

This framework was originally developed to assess robustness performance on answerable questions (Jia and Liang, 2017). In this paper, we also extend its application to attacks on unanswerable questions in Appendix §C.1 and discover challenges in this extended domain.

4.2 Attack Construction

Our algorithm constructs adversarial problems from original problems in a way similar to the AddOneSent in Jia and Liang (2017) and the AddText-Adv in Chen et al. (2022). Table 1 gives examples of such an attack on answerable and unanswerable questions. The additional sentence that is appended to the context has significant lexical overlap with the context, thus adding to the realism of the confusion-based attack. This type of adversarial attack is grammatical, fluent, and closely relevant to the given question. The questions and answers are unchanged for our considered adversarial attacks ($q' = q$ and $a' = a$).

Jia and Liang (2017) found that their adversarial attacks, especially the AddSent and AddOneSent attacks, were successful in challenging contemporary MRC models because the adversarial

sentences were closely related to the given questions. Notably, the unanswerable questions in SQuAD 2.0 show a similar kind of lexical overlap with their corresponding contexts and require MRC models to be highly robust to the subtle syntactic changes in order to determine the answerability of given questions. Therefore, we hypothesize that models fine-tuned with SQuAD 2.0 are equipped to perform better against adversarial attacks.

In the next section we assess this hypothesis by evaluating the performance of v1 versus v2 models on answerable questions.

5 Attacks on Answerable Questions: Results

5.1 Adversarial Performance

		Answerable		
		Original	Attacked	$\Delta \downarrow$
BERT	v1	88.4	63.8	24.6
	v2	78.4	55.2	23.2
RoBERTa	v1	91.5	70.5	21.0
	v2	84.8	58.0	26.8
SpanBERT	v1	91.5	68.6	22.9
	v2	85.8	58.9	26.8

Table 2: F1 scores of v1 models and v2 models with adversarial attacks on answerable questions. We refer to models fine-tuned on SQuAD 1.1 and SQuAD 2.0 as v1 and v2 models, accordingly.

Table 2 shows the performance of models with original (not attacked) and adversarial (attacked) problems on answerable questions. When attack

sentences are added into context, the performance of all v1 and v2 models significantly decreases. Adding *unanswerable* questions into the training (v2 models) does not initially appear to improve the robustness of MRC models against adversarial attacks. In fact, the performance of v2 models appears to be less robust than that of v1 models, both on the original and the attacked questions. However, there is a deeper story here worth investigating. To further explain the poor performances of v2 models, we consider the types of v2 answers to answerable questions in the next section.

5.2 Categories of Responses

		I	C2I	C2U	C2C
BERT	v1	10.9	28.7	-	60.4
	v2	21.3	10.9	14.7	53.2
RoBERTa	v1	8.0	24.5	-	67.7
	v2	14.5	8.0	20.5	57.1
SpanBERT	v1	8.0	26.7	-	65.4
	v2	13.8	8.3	20.1	57.8

Table 3: The percentage of answerable questions by types of answers produced by v1 and v2 models before and after adversarial attacks.

Table 3 shows the different categories of answers produced by v1 and v2 models to answerable questions. We use a 50% F1 score threshold to determine the models’ correctness to a question (correct if F1 score is above 50%, incorrect otherwise).

Considering attacks on answerable questions, we observe four categories in responses during attack: **“I” (incorrect)** are answerable questions that models originally got wrong (or originally predicted as unanswerable for v2 models). **“C2C” (correct to correct)** are answerable questions that models got correct both originally and after the attack. **“C2I” (correct to incorrect)** are answerable questions that models originally answered correctly but then output an incorrect answer when attacked. **“C2U” (correct to *incorrectly unanswerable*)** are answerable questions that models originally answer correctly but then predict as unanswerable when attacked. The C2I and C2U together account for the performance decline of models when attacked.

We see that v2 models, especially RoBERTa and SpanBERT, are particularly susceptible to the C2U challenge; they initially output a correct answer, but when attacked, decide (incorrectly) the

question is now unanswerable. This is in contrast to the v1 models, which not being trained on unanswerable questions and do not have the option of responding "unanswerable". The v2 models’ refusal to output an incorrect answer (opting instead to reply "unanswerable") indicates that their additional training on unanswerable questions has possibly provided them more depth to handle the confusion introduced by the attack.

We further breakdown the “C2U” category from Table 3 to investigate the spectrum of responses v2 models provide. Recall that models produce multiple responses to a MRC sample, each accompanied by a confidence score reflecting the models’ confidence in that response. In this analysis, to evaluate the difficulty of questions in category “C2U” of each v2 model, we use the corresponding v1 model as baseline. Then, to answer the question whether v2 models prefer correct answers to incorrect answers, we evaluate the second most confident response of v2 models for questions in category “C2U”.

		C2U	
		Attacked	# Questions
BERT	v1	46.1	871
	v2	42.5	
RoBERTa	v1	50.3	1212
	v2	44.7	
SpanBERT	v1	46.1	1194
	v2	47.6	

Table 4: F1 scores of second most confident responses of v2 models and most confident responses of v1 models to questions in category “C2U” of v2 models in Table 3. For each language model, we extract a set of “C2U” questions and then evaluate corresponding v1 and v2 models on this set of questions.

Table 4 shows the F1 scores of *second* most confident responses of v2 models and *first* (most confident) responses of v1 models to questions in category “C2U” under attacks. We observe that v2 models often have fairly good answers for questions in category “C2U” given that performance of v2 models lag significantly behind that of v1 models when attacked. However, v2 models fail to put forward the correct answers (their second option) ahead of the "unanswerable" responses (their first option).

From these analyses, we hypothesize that models with additional training on *unanswerable* questions have the ability to perceive the attacks on

answerable questions but fail to completely overcome them.

		Answerable		
		Original	Attacked	$\Delta \downarrow$
BERT	v1	88.4	63.8	24.6
	v2	88.5	69.6	18.9
RoBERTa	v1	91.5	70.5	21.0
	v2	91.4	75.1	16.4
SpanBERT	v1	91.5	68.6	22.9
	v2	91.3	75.8	15.5

Table 5: The performance of v1 and v2 models (when being forced to output non-empty answer on answerable questions) before and after adversarial attacks.

5.3 Force To Answer

The comparison of v1 and v2 models on answerable questions has a built-in bias because v2 models have the "penalty" of being able to respond "unanswerable" even though this is never a legitimate response. Furthermore, we have just shown that the v2 models often produce the correct answer, even under attack, but fail to put forward that correct output ahead of the "unanswerable" output in which it has more confidence. In this section, we re-run the analysis but this time eliminate the option for v2 models to output "unanswerable" (to answerable questions) so that we can better ascertain the robustness of v1 and v2 models to attacks.

Table 5 shows the results of this experiment. We can see now in this table that both v1 and v2 models exhibit similar performance on original answerable questions. When we introduce adversarial attacks on these same questions, the v2 models (being forced to answer) now exhibit noticeably stronger performance than their v1 counterparts. The additional training afforded to v2 models on unanswerable questions has given them a performance advantage over the v1 models. The robustness of v2 models against adversarial attacks is hidden in normal testing circumstances but can be realized by forcing the v2 models to output non-empty response in settings with only answerable questions.

6 Attacks in Out-Of-Domain Settings: Results

We now seek to determine if this additional robustness of v2 models extends to other out-of-domain test sets. In particular, we evaluate our v1 and v2 models on development sets of other Extractive

Question Answering datasets. We summarized the characteristics of five out-of-domain datasets of MRQA 2019 in Table 6.

Table 7 shows the performance of v1 and v2 models on the five datasets of MRQA 2019. Similarly to experiments in Section 5, we measure performance on both original problems and adversarially attacked problems.

First, the performance on original (unattacked) problems shows that adversarial unanswerable questions in SQuAD 2.0 have little negative effects on the generalization performance of MRC models. While the performance of v2 models is higher than that of v1 models on TriviaQA and SearchQA, v1 models outperform v2 models slightly on Natural Questions (0.8%), NewsQA (0.2 %), and considerably on HotpotQA (6.5 %). On average, the generalization performance of v2 models to that of v1 models on out-of-domain unattacked problems is slightly worse (53.7% to 54.5%).

However, on problems with adversarial attacks, v2 models significantly outperform v1 models in four out of the five datasets. Specifically, on average, v2 models significantly outperform v1 models by 2.9% on NewsQA, 4.7% on Natural Question, 4.8% on SearchQA, and 5.2% on TriviaQA. Although v2 models do not show superior performance to v1 models on HotpotQA, the performance gap between v2 and v1 models after attacks decreases significantly thanks to the superior robustness of v2 models.

Overall, we conclude from Table 7 that adversarial unanswerable questions of SQuAD 2.0 do not have negative effects on the generalization of v2 models to out-of-domain datasets, and the robustness of v2 models against adversarial attack is consistently superior to that of v1 models.

7 New Attack

In this section, we explore *why* v2 models often incorrectly put forward "unanswerable" as an incorrect response to answerable questions under adversarial attacks. We hypothesize that MRC models trained with SQuAD 2.0 have learned to identify target sentences with significant lexical overlap to decide whether the corresponding questions are unanswerable; the models rely *primarily* on that target sentence to determine their output. This undesirable behavior of MRC systems may prevent them from using the whole paragraph to accurately

Dataset	Question (Q)	Distant Supervision	Context (C)	Q \perp C	Dev
SQuAD	Crowdsourced	✗	Wikipedia	✗	10,507
HotpotQA	Crowdsourced	✗	Wikipedia	✗	5,904
TriviaQA	Trivia	✓	Web snippets	✓	7,785
SearchQA	Jeopardy	✓	Web snippets	✓	16,980
NewsQA	Crowdsourced	✗	News articles	✓	4,212
Natural Questions	Search logs	✗	Wikipedia	✓	12,836

Table 6: Characteristics of each datasets used in our out-of-domain experiments. Distant supervision is True if datasets used distant supervision to match questions and contexts. $Q \perp C$ is True if questions in datasets are written independently from the passage used for context. Table adopted from shared task MRQA 2019 (Fisch et al., 2019).

		Natural Question			HotpotQA			TriviaQA		
		Original	Attacked	$\Delta \downarrow$	Original	Attacked	$\Delta \downarrow$	Original	Attacked	$\Delta \downarrow$
BERT	v1	54.6	20.1	34.5	61.6	45.5	16.1	59.4	48.9	10.5
	v2	52	23.7	28.3	58.9	47.4	11.5	58.9	53.3	5.6
RoBERTa	v1	62.1	28.3	33.8	67.4	46.3	21.1	64.1	55	9.1
	v2	63.5	33.2	30.3	65	49.8	15.2	65.5	59.2	6.3
SpanBERT	v1	65	34.5	30.5	66.2	46.4	19.8	63.2	51.9	11.3
	v2	63.9	40.2	23.7	51.9	32.3	19.6	62.9	58.8	4.1
Average	v1	60.6	27.6	33	65.1	46.1	19	62.2	51.9	10.3
	v2	59.8	32.3	27.5	58.6	43.2	15.4	62.4	57.1	5.3

		SearchQA			NewsQA			Average		
		Original	Attacked	$\Delta \downarrow$	Original	Attacked	$\Delta \downarrow$	Original	Attacked	$\Delta \downarrow$
BERT	v1	30.4	25.5	4.9	53.6	41.8	11.8	51.9	36.4	15.5
	v2	28.6	26.7	1.9	53.9	46.2	7.7	50.5	39.5	11
RoBERTa	v1	22.8	20.3	2.5	61.2	54.2	7	55.5	40.8	14.7
	v2	33	31.6	1.4	60.6	52.5	8.1	57.5	45.3	12.2
SpanBERT	v1	28.1	26.9	1.2	58.2	44.1	14.1	56.1	40.8	15.3
	v2	29.4	28.8	0.6	58	50	8	53.2	42	11.2
Average	v1	27.1	24.2	2.9	57.7	46.7	11	54.5	39.3	15.2
	v2	30.3	29	1.3	57.5	49.6	7.9	53.7	42.3	11.4

Table 7: Robustness of MRC models fine-tuned on SQuAD 1.1 (v1) and SQuAD 2.0 (v2) in out-of-domain settings. For models fine-tuned on SQuAD 2.0 (v2), we force models to output non-empty answers. For each dataset, we report the average performance of 3 experimented models. We also report the average performance of each models on 5 considered datasets.

determine the best response to a question and have negative effects on the practical usage of adversarial unanswerable questions.

To further understand this hypothesis, we introduce a *negation attack*, a new adversarial attack to attempt to fool models into giving incorrect "unanswerable" responses. In particular, we construct an attack statement that significantly overlaps with the question yet is easy to determine as contradicting the question; we form our negation attack by inserting "not" in front of the adjective. Our attack (see Table 8) differs from previous adversarial attacks as our attack is designed to elicit an

unanswerable response instead of an incorrect response.

Table 9 reports the performance of v2 models under negation attacks on answerable questions. We observe that our negation attack is highly effective in revealing the weaknesses of v2 models as the performance of all three considered v2 models significantly drops by almost 60% F1 when we introduce the negation attack.

We then examine the shifts in answers of v2 models when attacked with negation type. Table 10 shows the distribution of shifts in answers before and after the attack. We observe that the most

Question	In the effort of maintaining a level of abstraction, what choice is typically left independent ?
Answer	encoding
Context	[...] one tries to keep the discussion abstract enough to be independent of the choice of encoding . [...] In the effort of maintaining a level of abstraction, base64 choice is typically left not independent.

Table 8: An example of the Negation Attack on answerable questions. The adversarial sentence is highlighted in red color. In constructing the adversarial sentence, we negate adjective “independent” to “not independent”.

	Original	Attacked	$\Delta \downarrow$
BERT v2	84.8	24.2	60.6
RoBERTa v2	78.1	21	57.1
SpanBERT v2	87.3	28.6	58.7

Table 9: F1 score of v2 models before and after negation attacks on answerable questions. In this experiment, we do not force v2 models to output non-empty answers.

significant drop in performance under negation attacks is the “C2U” category (around 40 % F1). This result is consistent with our hypothesis that v2 models rely on target sentences with significant lexical overlap to decide whether the corresponding questions are unanswerable.

8 Conclusion

In this work, we investigate the effects of training MRC models with unanswerable questions on their robustness against adversarial attacks. We construct adversarial samples from answerable and unanswerable questions in SQuAD 2.0 and evaluate three MRC models fine-tuned on either SQuAD 1.1 (v1 models) or SQuAD 2.0 (v2 models) independently.

Adversarial attacks on answerable questions reveal that v2 models initially show little improved robustness over v1 models yet possess a latent ability to deal with these attacks that v1 models do not; the correct responses are often hidden as second-best answers, an indicator of the “hidden robustness” of v2 models resulting from additional training on unanswerable questions. By eliminating the “unanswerable” option and forcing v2 models to output an answer to any answer-

	I	C2U	C2I	C2C
BERT v2	14.4	45.4	17.7	22.5
RoBERTa v2	21.6	41.8	17.5	19.1
SpanBERT v2	12.5	37.9	22.8	26.8

Table 10: The percentage of answerable questions by types of answers produced by v2 models before and after negation attacks.

able questions, we leverage this hidden robustness to improve the performance of MRC models to attacks on answerable questions. Furthermore, we also show that this robustness translates well to out-of-domain test sets.

Finally, to encourage future work in evaluating the robustness of MRC models trained on both answerable and unanswerable questions, we introduce a new type of adversarial attack to reveal the short-comings of MRC models. Our experiments with the *negation* attack reveal that the performance of v2 MRC models drops significantly (around 50% F1). We hypothesize that the decline in the performance of v2 models is mainly due to how v2 models have learned to suboptimally identify target sentences in the context to use as their primary mechanism of response.

9 Future Work

Our findings raise two critical messages for future research in the usage of adversarial unanswerable questions in NLP:

First, our work highlights innovative ways to use adversarial unanswerable questions in training to improve the performance of MRC-based systems. MRC datasets are important sources of transfer learning for zero-shot settings in many other NLP tasks (Wu et al., 2020; Levy et al., 2017; Lyu et al., 2021; Du and Cardie, 2020; Li et al., 2019). Given that the improved robustness of v2 models from the additional training on unanswerable questions generalizes well to out-of-domain test sets, future research about using MRC knowledge in zero-shot settings can explore whether adversarial unanswerable questions improve the robustness of MRC models in these zero-shot settings.

Second, we propose an open question about an undesirable behavior of MRC models fine-tuned on SQuAD 2.0. We find that simple negation attacks induce a considerable drop in the performance of MRC models fine-tuned on SQuAD 2.0

due to an undesirable behavior as the product of artifacts in the training set. To use the adversarial unanswerable questions in practice, we suggest additional research, based on insights about shortcut learning (Lai et al., 2021; Du et al., 2021), aimed to prevent MRC models from learning this undesirable behavior.

Limitations

We acknowledge that there exist few aspects to which our findings are limited, that include the dominant use of pretrained language models, the insufficiency of MRC datasets in other languages, and the limited types of adversarial attacks examined.

Acknowledgements

We would like to thank Dr. Ashwin Lall for constructive feedback on the early version of this paper. We thank the anonymous reviewers for their constructive and insightful feedback. We want to thank The William G. and Mary Ellen Bowen Research Endowment and The Laurie and David Hodgson Faculty Support Endowment for supporting the first and third authors.

References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. [Generating natural language adversarial examples](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- Akari Asai and Eunsol Choi. 2021. [Challenges in information-seeking QA: Unanswerable questions and paragraph retrieval](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1492–1504, Online. Association for Computational Linguistics.
- Matthias Blohm, Glorianna Jagfeld, Ekta Sood, Xiang Yu, and Ngoc Thang Vu. 2018. [Comparing attention-based convolutional and recurrent neural networks: Success and limitations in machine reading comprehension](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 108–118, Brussels, Belgium. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Howard Chen, Jacqueline He, Karthik Narasimhan, and Danqi Chen. 2022. [Can rationalization improve robustness?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3792–3805, Seattle, United States. Association for Computational Linguistics.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- 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.
- Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun, and Xia Hu. 2021. [Towards interpreting and mitigating shortcut learning behavior of NLU models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 915–929, Online. Association for Computational Linguistics.
- Xinya Du and Claire Cardie. 2020. [Event extraction by answering \(almost\) natural questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 671–683, Online. Association for Computational Linguistics.
- Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. [Searchqa: A new q&a dataset augmented with context from a search engine](#).
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. [MRQA 2019 shared task: Evaluating generalization in reading comprehension](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*,

- pages 1–13, Hong Kong, China. Association for Computational Linguistics.
- Quentin Heinrich, Gautier Viaud, and Wacim Belbidia. 2021. [Fquad2.0: French question answering and knowing that you know nothing](#).
- Alex Huang, Abdullah Al-Dujaili, Erik Hemberg, and Una-May O’Reilly. 2018. Adversarial deep learning for robust detection of binary encoded malware. *2018 IEEE Security and Privacy Workshops (SPW)*, pages 76–82.
- 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, New Orleans, Louisiana. 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale ReAding comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Yuxuan Lai, Chen Zhang, Yansong Feng, Quzhe Huang, and Dongyan Zhao. 2021. [Why machine reading comprehension models learn shortcuts?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 989–1002, Online. Association for Computational Linguistics.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. [Zero-shot relation extraction via reading comprehension](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, Vancouver, Canada. Association for Computational Linguistics.
- Xiaoya Li, Fan Yin, Zijun Sun, Xiayu Li, Arianna Yuan, Duo Chai, Mingxin Zhou, and Jiwei Li. 2019. [Entity-relation extraction as multi-turn question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1340–1350, Florence, Italy. Association for Computational Linguistics.
- Bin Liang, Hongcheng Li, Miaoqiang Su, Pan Bian, Xirong Li, and Wenchang Shi. 2018. Deep text classification can be fooled. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI’18*, page 4208–4215. AAAI Press.
- 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.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Qing Lyu, Hongming Zhang, Elinor Sulem, and Dan Roth. 2021. [Zero-shot event extraction via transfer learning: Challenges and insights](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 322–332, Online. Association for Computational Linguistics.
- Sewon Min, Danqi Chen, Luke Zettlemoyer, and Hananeh Hajishirzi. 2019. [Knowledge guided text retrieval and reading for open domain question answering](#).
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. [TextAttack: A framework for adversarial attacks, data augmentation, and adversarial training in NLP](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126, Online. Association for Computational Linguistics.
- Paarth Neekhara, Shehzeen Hussain, Shlomo Dubnov, and Farinaz Koushanfar. 2019. [Adversarial reprogramming of text classification neural networks](#). 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 5216–5225, Hong Kong, China. Association for Computational Linguistics.
- Nicolas Papernot, Patrick McDaniel, Ananthram Swami, and Richard E. Harang. 2016. Crafting adversarial input sequences for recurrent neural networks. *MILCOM 2016 - 2016 IEEE Military Communications Conference*, pages 49–54.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. 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.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. [CoQA: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- 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.
- Suranjana Samanta and Sameep Mehta. 2018. Generating adversarial text samples. In *ECIR*.
- Priyanka Sen and Amir Saffari. 2020. [What do models learn from question answering datasets?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2429–2438, Online. Association for Computational Linguistics.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016a. [Bidirectional attention flow for machine comprehension](#). *CoRR*, abs/1611.01603.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016b. [Bidirectional attention flow for machine comprehension](#).
- Chenglei Si, Ziqing Yang, Yiming Cui, Wentao Ma, Ting Liu, and Shijin Wang. 2021. [Benchmarking robustness of machine reading comprehension models](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 634–644, Online. Association for Computational Linguistics.
- Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. [What makes reading comprehension questions easier?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4208–4219, Brussels, Belgium. Association for Computational Linguistics.
- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2020. [Assessing the benchmarking capacity of machine reading comprehension datasets](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8918–8927.
- Elior Sulem, Jamaal Hay, and Dan Roth. 2021. [Do we know what we don’t know? studying unanswerable questions beyond SQuAD 2.0](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4543–4548, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hongxuan Tang, Hongyu Li, Jing Liu, Yu Hong, Hua Wu, and Haifeng Wang. 2021. [DuReader_robust: A Chinese dataset towards evaluating robustness and generalization of machine reading comprehension in real-world applications](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 955–963, Online. Association for Computational Linguistics.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. [NewsQA: A machine comprehension dataset](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200, Vancouver, Canada. Association for Computational Linguistics.
- Kiet Van Nguyen, Son Quoc Tran, Luan Thanh Nguyen, Tin Van Huynh, Son T. Luu, and Ngan Luu-Thuy Nguyen. 2022. [Vlsp 2021 - vimrc challenge: Vietnamese machine reading comprehension](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). In

- Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Shuohang Wang and Jing Jiang. 2016. [Machine comprehension using match-lstm and answer pointer](#).
- Xuezhi Wang, Haohan Wang, and Diyi Yang. 2022. [Measure and improve robustness in NLP models: A survey](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4569–4586, Seattle, United States. Association for Computational Linguistics.
- Yicheng Wang and Mohit Bansal. 2018. [Robust machine comprehension models via adversarial training](#). 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 575–581, New Orleans, Louisiana. Association for Computational Linguistics.
- Bowen Wu, Haoyang Huang, Zongsheng Wang, Qihang Feng, Jingsong Yu, and Baoxun Wang. 2019. [Improving the robustness of deep reading comprehension models by leveraging syntax prior](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 53–57, Hong Kong, China. Association for Computational Linguistics.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020. [CorefQA: Coreference resolution as query-based span prediction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6953–6963, Online. Association for Computational Linguistics.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. [End-to-end open-domain question answering with BERTserini](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. 2020. [Adversarial attacks on deep-learning models in natural language processing: A survey](#). *ACM Trans. Intell. Syst. Technol.*, 11(3).
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018. [Generating natural adversarial examples](#). *ArXiv*, abs/1710.11342.

A Attacks

In this section, we document the pseudo-code we use to generate the two attacks in our main paper. In the pseudo-code below, (\cdot) indicates the input(s) of the function within the current line.

A.1 AddOneSent Attack

Algorithm 1: AddOneSent Attack

```
Function AddOneSent(question, answer):  
  new_question  $\leftarrow$  question  
  new_answer  $\leftarrow$  answer  
  new_question  $\leftarrow$  Replace nouns and  
    adjectives with antonyms in  
    WordNet(new_question).  
  new_question  $\leftarrow$  Change named  
    entities and numbers to nearest  
    word in GloVe(new_question).  
  new_answer  $\leftarrow$  Change named entities  
    and numbers to nearest word in  
    GloVe(new_answer).  
  Assert (new_answer  $\neq$  answer) &&  
    (new_question  $\neq$  question)  
  attack  $\leftarrow$  Convert into statement  
    (new_question, new_answer).  
  return attack
```

Algorithm 1 is the pseudo-code for AddOneSent attack used in our analysis.

A.2 Negation Attack

Algorithm 2: Negation Attack

```
Function Negation(question, answer):  
  new_question  $\leftarrow$  question  
  new_answer  $\leftarrow$  answer  
  new_question  $\leftarrow$  Add not before the  
    first adjective (new_question).  
  new_answer  $\leftarrow$  Change named entities  
    and numbers to nearest word in  
    GloVe(new_answer).  
  Assert (new_answer  $\neq$  answer) &&  
    (new_question  $\neq$  question)  
  attack  $\leftarrow$  Convert into statement  
    (new_question, new_answer).  
  return attack
```

Algorithm 2 is the pseudo-code for the Negation attack introduced in Section 7 to further reinforce our hypothesis that v2 models undesirably learn artifacts in adversarial unanswerable questions of

SQuAD 2.0. The main difference between AdOneSent attack and Negation attack is that Negation attack does not use WordNet to Replace nouns and adjectives, and does not use GloVe to change named entities and numbers to nearest word in word space of GloVe.

A.3 Quality Analysis

In order to investigate the quality of Negation Attack, we manually label the 200 attack samples produced by both Negation Attack and AddOneSent attacks (100 each) into three categories:

1. **FM**: fluent and meaningful attack sentence.
2. **M**: meaningful but not fluent attack sentence.
3. **N**: not meaningful attack sentence.

Table 11 provide examples of Negation and AdOneSent attack samples categorized into these three categories. The errors of the Negation attack mostly come from the unnatural expression when using “not” to negate adjectives instead of using antonyms (**not significant** versus **insignificant**). On the other hand, the errors of AddOneSent can occur because of misclassifying word type. For example, when misclassifying the noun **kind** as adjectives, AddOneSent would then rewrite **kind of company** as **unkind of company**).

B Details for MRC Model Training

In this work, we use the base versions for all considered pre-trained models. We train all MRC models using mixed precision, with batch size of 4 sequences for 2 epochs. The maximum sequence length is set to 384 tokens. We use the AdamW optimizer (Loshchilov and Hutter, 2019) with an initial learning rate of $2 \cdot 10^{-5}$, and $\beta_1 = 0.9$, $\beta_2 = 0.999$. We fine-tuned all four models on a single NVIDIA Tesla K80 provided by Google Colaboratory.

C Attacks on Unanswerable Questions: Results

C.1 Adversarial Performance

In this section, we extend our robustness evaluation of v2 models by analyzing their performance against adversarial attacks on unanswerable questions. Recall that we conduct these experiments only on v2 models as v1 models have not been trained on unanswerable questions.

AddOneSent		Negation		
	Example	Proportion	Example	Proportion
FM	Question: Who was the chief executive officer when the service began? Attack: Russell Hartley was the chief executive officer when the disservice began.	44	Question: What service is a VideoGuard UK equipped receiver dedicated to decrypt? Attack: A VideoGuard UK equipped receiver is not dedicated to decrypt the service of skies.	43
M	Question: How populous is Victoria compared to other Australian states? Attack: Victoria compared to same japanese states is 3rd - most populous.	25	Question: What is the most important type of Norman art preserved in churches? Attack: The most not important type of Norman art is preserved in churches frescos.	49
N	Question: What kind of company is Sky UK Limited? Attack: The unkind of company of macedonian telecommunications company is geelong.	31	Question: What does most of the HD material use as a standard? Attack: The U.S. revolutionary peace does not most of the HD material use as a standard.	8

Table 11: Attack samples of Negation and AddOneSent categorized into three categories (fluent and meaningful, meaningful but not fluent, and not meaningful) and their overall proportions.

	Unanswerable		
	Original	Attacked	$\Delta \downarrow$
BERT v2	72.2	69.3	2.9
RoBERTa v2	81.7	77.9	3.8
SpanBERT v2	76.4	75.3	1.1

Table 12: F1 score of v2 models with adversarial attacks on unanswerable questions.

Table 12 reports the performances of v2 models to adversarial attacks on unanswerable questions. Among the F1 scores of the three v2 models, the score of RoBERTa v2 decreases most after the attacks (by 3.8%) while the F1 score of SpanBERT v2 decreases least (by only 1.1%). These results *seem* to indicate that the adversarial attacks only slightly degrade the performances of v2 models, which might lead to erroneous conclusions about the robustness of these models. However, if we look back at Table 3, we see that between 8% and 11% of samples are in the C2I group (correct originally, incorrect when attacked). These prior results on answerable questions suggest inconsistencies with the results on unanswerable questions. We dig further.

C.2 Categories of Responses

	CU2CU	IA2IA	CU2IA	IA2CU
BERT v2	61.8	20.4	10.4	7.4
RoBERTa v2	71.8	11.1	9.9	7.2
SpanBERT v2	65.2	13.5	11.2	10.1

Table 13: The percentage of unanswerable questions by types of answers produced by v1 and v2 models before and after adversarial attacks.

We apply a similar investigation as we did previously to categorize the response changes of these v2 models to attacks on unanswerable questions. We find four main categories:

- **“CU2CU” (correctly unanswerable to correctly unanswerable)** are questions that v2 models correctly predicted as unanswerable both before and after the attacks.
- **“IA2IA” (incorrectly answerable to incorrectly answerable)** are unanswerable questions that v2 models attempt to output answers both before and after the attacks.
- **“CU2IA” (correctly unanswerable to incorrectly answerable)** are questions that v2 models originally correctly predicted as unanswerable but then output an answer when attacked.
- **“IA2CU” (incorrectly answerable to correctly unanswerable)** are questions that v2 models originally erroneously attempt to output an answer but later correctly predict as unanswerable when attacked.

What Table 13 reveals is that the performance loss of the models during the attack is being masked by some questions that were initially incorrect but are correctly identified as unanswerable after the attack (IA2CU). For example, the BERT model appears to only lose 2.9 F1 score during the attack, but actually it loses 10.4 and then gains back 7.4 in other IA2CU questions. These results reveal that v2 models experience a similar

performance decline on unanswerable questions as they did on answerable questions. They also show how the current assessment framework is unsuitable for accurately measuring the robustness of v2 models on both answerable and unanswerable questions.