

Undersensitivity in Neural Reading Comprehension

Johannes Welbl^{†*} Pasquale Minervini[†] Max Bartolo[†]
Pontus Stenetorp[†] Sebastian Riedel[†]

[†]University College London

{j.welbl,p.minervini,m.bartolo,p.stenetorp,s.riedel}@cs.ucl.ac.uk

Abstract

Current reading comprehension methods generalise well to in-distribution test sets, yet perform poorly on adversarially selected data. Prior work on adversarial inputs typically studies model oversensitivity: semantically invariant text perturbations that cause a model’s prediction to change. Here we focus on the complementary problem: excessive prediction undersensitivity, where input text is meaningfully changed but the model’s prediction does not, even though it should. We formulate an adversarial attack which searches among semantic variations of the question for which a model erroneously predicts the same answer, and with even higher probability. We demonstrate that models trained on both *SQuAD2.0* and *NewsQA* are vulnerable to this attack, and then investigate data augmentation and adversarial training as defences. Both substantially decrease adversarial vulnerability, which generalises to held-out data and held-out attack spaces. Addressing undersensitivity furthermore improves model robustness on the previously introduced ADDSENT and ADDONESSENT datasets, and models generalise better when facing train/evaluation distribution mismatch: they are less prone to overly rely on shallow predictive cues present only in the training set, and outperform a conventional model by as much as 10.9% F_1 .

1 Introduction

Neural networks can be vulnerable to adversarial input perturbations (Szegedy et al., 2013; Kurakin et al., 2016). In Natural Language Processing (NLP), which operates on discrete symbol sequences, adversarial attacks can take a variety of forms (Ettinger et al., 2017; Alzantot et al., 2018) including character perturbations (Ebrahimi

et al., 2018), semantically invariant reformulations (Ribeiro et al., 2018b; Iyyer et al., 2018a) or – specifically in Reading Comprehension (RC) – adversarial text insertions (Jia and Liang, 2017; Wang and Bansal, 2018). A model’s inability to handle adversarially chosen input text puts into perspective otherwise impressive generalisation results for in-distribution test sets (Seo et al. (2017); Yu et al. (2018); Devlin et al. (2019); *inter alia*) and constitutes an important caveat to conclusions drawn regarding a model’s comprehension abilities.

While semantically invariant text transformations can remarkably alter a model’s predictions, the complementary problem of model *undersensitivity* is equally troublesome: a model’s text input can often be drastically changed in meaning while retaining the original prediction. In particular, previous works (Feng et al., 2018; Ribeiro et al., 2018a; Welbl et al., 2020) show that even after deletion of all but a small fraction of input words, models often produce the same output. However, such reduced inputs are usually unnatural to a human reader, and it is both unclear what behaviour we should expect from natural language models evaluated on unnatural text, and how to use such unnatural inputs to improve models.

In this work we explore RC model undersensitivity for natural language questions, and we show that addressing undersensitivity not only makes RC models more sensitive where they should be, but also less reliant on shallow predictive cues. Fig. 1 shows an example for a BERT LARGE model (Devlin et al., 2019) trained on *SQuAD2.0* (Rajpurkar et al., 2018) that is given a text and a comprehension question, i.e. *What was Fort Caroline renamed to after the Spanish attack?* which it correctly answers as *San Mateo* with 98% probability. Altering this question, however, can increase model probability for this same prediction to 99%, although the new question is unanswerable given the same con-

*Now at DeepMind.

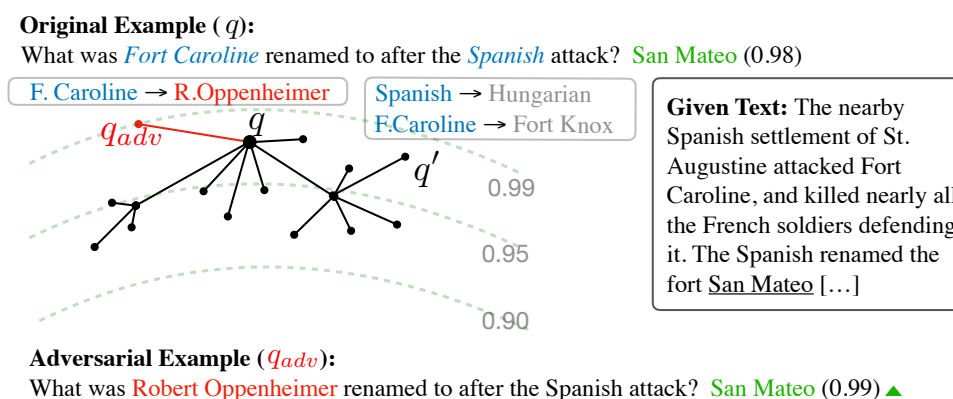


Figure 1: Method Overview: Adversarial search over semantic variations of RC questions, producing unanswerable questions for which the model retains its predictions with even higher probability.

text. That is, we observe an increase in probability despite removing relevant question information and replacing it with new, irrelevant content.

We formalise the process of finding such questions as an adversarial search in a discrete space arising from perturbations of the original question. There are two types of discrete perturbations we consider, based on part-of-speech tags and named entities, with the aim of obtaining grammatical and semantically consistent alternative questions that do not accidentally have the same correct answer. We find that *SQuAD2.0* and *NewsQA* models can be attacked on a substantial proportion of samples.

The observed undersensitivity correlates negatively with in-distribution test set performance metrics (EM/ F_1), suggesting that this phenomenon – where present – is indeed a reflection of a model’s lack of question comprehension. When training models to defend against undersensitivity attacks with data augmentation and adversarial training, we observe that they can generalise their robustness to held out evaluation data without sacrificing in-distribution test set performance. Furthermore, the models improve on the adversarial datasets proposed by Jia and Liang (2017), and behave more robustly in a learning scenario that has dataset bias with a train / evaluation distribution mismatch, increasing performance by up to 10.9% F_1 .

List of Contributions: i) We propose a new type of adversarial attack exploiting the undersensitivity of neural RC models to input changes, and show that contemporary models are vulnerable to it; ii) We compare data augmentation and adversarial training as defences, and show their effectiveness at reducing undersensitivity errors on both held-out data and held-out perturbations without sacrificing nominal test performance; iii) We demonstrate that

the resulting models generalise better on the adversarial datasets of Jia and Liang (2017), and in the biased data setting of Lewis and Fan (2019).

2 Related Work

Adversarial Attacks in NLP Adversarial examples have been studied extensively in NLP – see Zhang et al. (2019) for a recent survey. Yet automatically generating adversarial inputs is non-trivial, as altering a single word can change the semantics of an instance or render it incoherent. Prior work typically considers semantic-invariant input transformations to which neural models are *oversensitive*. For instance, Ribeiro et al. (2018b) use a set of simple perturbations such as replacing *Who is* with *Who’s*. Other semantics-preserving perturbations include typos (Hosseini et al., 2017), the addition of distracting sentences (Jia and Liang, 2017; Wang and Bansal, 2018), character-level adversarial perturbations (Ebrahimi et al., 2018; Belinkov and Bisk, 2018), and paraphrasing (Iyyer et al., 2018b). In this work, we focus on the complementary problem of *undersensitivity* of neural RC models to semantic perturbations of the input. Our method is based on the idea that modifying, for instance, the named entities in a question can completely change its meaning and, as a consequence, the question should become unanswerable given the context. Our approach does not assume white-box access to the model, as do e.g. Ebrahimi et al. (2018) and Wallace et al. (2019).

Undersensitivity Jacobsen et al. (2019) demonstrated classifier undersensitivity in computer vision. Niu and Bansal (2018) investigated undersensitivity in dialogue models and addressed the problem with a max-margin training approach. Ribeiro

et al. (2018a) describe a general model diagnosis tool to identify minimal feature sets that are sufficient for a model to form high-confidence predictions. Feng et al. (2018) showed that it is possible to reduce inputs to minimal input word sequences without changing a model’s predictions. Welbl et al. (2020) investigated formal verification against undersensitivity to text deletions. We see our work as a continuation of these lines of inquiry, with a particular focus on undersensitivity in RC. In contrast to Feng et al. (2018) and Welbl et al. (2020) we consider concrete alternative questions, rather than arbitrarily reduced input word sequences, and furthermore address the observed phenomenon using dedicated training objectives, in contrast to Feng et al. (2018) and Ribeiro et al. (2018a) who simply highlight it. Gardner et al. (2020) and Kaushik et al. (2020) also recognise the problem of models learning shallow but successful heuristics, and propose counterfactual data annotation paradigms as prevention. The perturbations used in this work define such counterfactual samples. Their composition does not require additional annotation efforts, and we furthermore adapt an adversarial perspective on the choice of such samples. Finally, one of the methods we evaluate for defending against undersensitivity attacks is a form of data augmentation that has similarly been used for de-biasing NLP models (Zhao et al., 2018; Lu et al., 2018).

Concurrent work on model CHECKLIST evaluation (Ribeiro et al., 2020) includes an *invariance test* which also examines model undersensitivity. In contrast to CHECKLIST, our work focuses with more detail on the analysis of the invariance phenomenon, the automatic generation of probing samples, an investigation of concrete methods to overcome undesirably invariant model behaviour, and shows that adherence to invariance tests leads to more robust model generalisation.

Unanswerable Questions in RC Rajpurkar et al. (2018) proposed the *SQuAD2.0* dataset, which includes over 43,000 human-curated unanswerable questions. *NewsQA* is a second dataset with unanswerable questions, in the news domain (Trischler et al., 2017). Training on these datasets should conceivably result in models with an ability to tell whether questions are answerable or not; we will however see that this does not extend to adversarially chosen unanswerable questions. Hu et al. (2019) address unanswerability of questions from a given text using additional verification steps. Other

approaches have shown the benefit of synthetic data to improve performance in *SQuAD2.0* (Zhu et al., 2019; Alberti et al., 2019). In contrast to prior work, we demonstrate that despite improving performance on test sets that include unanswerable questions, the problem persists when adversarially choosing from a larger space of questions.

3 Methodology

Problem Overview Consider a discriminative model f_θ , parameterised by a collection of vectors θ , which transforms an input \mathbf{x} into a prediction $\hat{y} = f_\theta(\mathbf{x})$. In our task, $\mathbf{x} = (t, q)$ is a given text t paired with a question q about this text. The label y is the answer to q where it exists, or a *NoAnswer* label where it cannot be answered.¹

In an RC setting, the set of possible answers is large, and predictions \hat{y} should be dependent on \mathbf{x} . And indeed, randomly choosing a different input (t', q') usually changes the model prediction \hat{y} . However, there exist many examples where the prediction erroneously remains stable; the goal of the attack formulated here is to find such cases. Formally, the goal is to discover inputs \mathbf{x}' for which f_θ still erroneously predicts $f_\theta(\mathbf{x}') = f_\theta(\mathbf{x})$, even though \mathbf{x}' is not answerable from the text.

Identifying suitable candidates for \mathbf{x}' can be achieved in manifold ways. One approach is to search among a large question collection, but we find this to only rarely be successful; an example is shown in Table 8, Appendix E. Generating \mathbf{x}' , on the other hand, is prone to result in ungrammatical or otherwise ill-formed text. Instead, we consider a perturbation space $\mathcal{X}_\mathcal{T}(\mathbf{x})$ spanned by perturbing original inputs \mathbf{x} using a perturbation function family \mathcal{T} : $\mathcal{X}_\mathcal{T}(\mathbf{x}) = \{T_i(\mathbf{x}) \mid T_i \in \mathcal{T}\}$. This space $\mathcal{X}_\mathcal{T}(\mathbf{x})$ contains alternative model inputs derived from \mathbf{x} . Ideally the transformation function family \mathcal{T} is chosen such that the correct label of these inputs is changed, and the question becomes unanswerable. We will later search within $\mathcal{X}_\mathcal{T}(\mathbf{x})$ to find inputs \mathbf{x}' which erroneously retain the same prediction as \mathbf{x} : $\hat{y}(\mathbf{x}) = \hat{y}(\mathbf{x}')$.

Part-of-Speech (PoS) Perturbations We first consider the perturbation space $\mathcal{X}_{\mathcal{T}_P}(\mathbf{x})$ with PoS perturbations \mathcal{T}_P of the original question: we swap individual tokens with other, PoS-consistent alternative tokens, drawing from large collections of tokens of the same PoS types. For example, we might

¹Unanswerable questions are part of, e.g. the *SQuAD2.0* and *NewsQA* datasets, but not *SQuAD1.1*.

alter the question “*Who patronized the monks in Italy?*” to “*Who betrayed the monks in Italy?*” by replacing the past tense verb “*patronized*” with “*betrayed*”. There is however no guarantee that the altered question will require a different answer (e.g. due to synonyms). Even more so – there might be type clashes or other semantic inconsistencies. We perform a qualitative analysis to investigate the extent of this problem and find that, while a valid concern, for the majority of attackable samples (at least 51%) there exist attacks based on correct well-formed questions (see Section 5).

Named Entity (NE) Perturbations The space $\mathcal{X}_{\mathcal{T}_E}(\mathbf{x})$ of the transformation family \mathcal{T}_E is created by substituting NE mentions in the question with different type-consistent NE, derived from a large set E . For example, a question “*Who patronized the monks in Italy?*” could be altered to “*Who patronized the monks in Las Vegas?*”, replacing the geopolitical entity “*Italy*” with “*Las Vegas*”, chosen from E . Altering NE often changes the question specifics and alters the answer requirements, which are unlikely to be satisfied from what is stated in the given text, given the broad nature of entities in E . While perturbed questions are not guaranteed to be unanswerable or require a different answer, we will in a later analysis see that for the large majority of cases (at least 84%) they do.

Undersensitivity Attacks Thus far we have described methods for perturbing questions. We will search in the resulting perturbation spaces $\mathcal{X}_{\mathcal{T}_P}(\mathbf{x})$ and $\mathcal{X}_{\mathcal{T}_E}(\mathbf{x})$ for inputs \mathbf{x}' for which model prediction remains constant. However, we pose a slightly stronger and more conservative requirement to rule out cases where the prediction is retained, but with lower probability: f_θ should assign a higher probability to the same prediction $\hat{y}(\mathbf{x}) = \hat{y}(\mathbf{x}')$ than for the original input:

$$P(\hat{y} | \mathbf{x}') > P(\hat{y} | \mathbf{x}) \quad (1)$$

To summarise, we are searching in a perturbation space for altered questions which result in a higher model probability to the same answer as the original input question. If we have found an altered question that satisfies inequality (1), then we have identified a successful attack, which we will refer to as an *undersensitivity attack*.

Adversarial Search in Perturbation Space In its simplest form, a search for an adversarial attack in the previously defined spaces amounts to a search over a list of single lexical alterations for the

maximum (or any) higher prediction probability. We can however repeat the replacement procedure multiple times, arriving at texts with larger lexical distance to the original question. For example, in two iterations of PoS-consistent lexical replacement, we can alter “*Who was the duke in the battle of Hastings?*” to inputs like “*Who was the duke in the expedition of Roger?*” The space of possibilities grows combinatorially, and with increasing perturbation radius it becomes computationally infeasible to comprehensively cover the full perturbation space arising from iterated substitutions. To address this, we follow Feng et al. (2018) and apply beam search to narrow the search space, and seek to maximise the difference $\Delta = P(\hat{y} | \mathbf{x}') - P(\hat{y} | \mathbf{x})$. Beam search is conducted up to a pre-specified maximum perturbation radius ρ , but once \mathbf{x}' with $\Delta > 0$ has been found, we stop the search.

Relation to Attacks in Prior Work Note that this type of attack stands in contrast to other attacks based on small, *semantically invariant* input perturbations which investigate oversensitivity problems. Such semantic *invariance* comes with stronger requirements and relies on synonym dictionaries (Ebrahimi et al., 2018) or paraphrases harvested from back-translation (Iyyer et al., 2018a), which are both incomplete and noisy. Our attack is instead focused on *undersensitivity*, i.e. where the model is stable in its prediction even though it should not be. Consequently the requirements are not as difficult to fulfil when defining perturbation spaces that *alter* the question meaning, and one can rely on sets of entities and PoS examples automatically extracted from a large text collection.

4 Experiments: Model Vulnerability

Training and Dataset Details We next conduct experiments using the attacks laid out above to investigate model undersensitivity. We attack the BERT model fine-tuned on *SQuAD2.0*, and measure to what extent the model exhibits undersensitivity to adversarially chosen inputs. Our choice of BERT is motivated by the currently widespread adoption of its variants across the NLP field, and empirical success across a wide range of datasets. *SQuAD2.0* per design contains unanswerable questions; models are thus trained to predict a *NoAnswer* option where a comprehension question cannot be answered. As the test set is unavailable, we split off 5% from the original training set for development purposes and retain the remaining




Figure 2: BERT LARGE on *SQuAD2.0*: vulnerability to noisy attacks on held out data for differently sized attack spaces (parameter η) and different beam search depth (perturbation radius ρ).

95% for training, stratified by articles. The original *SQuAD2.0* development set is then used as evaluation data, where the model reaches 73.0%EM and 76.5%F₁; we will compute the undersensitivity attacks on this entirely held out part of the dataset. Appendix B provides further training details.

Attack Details To compute the perturbation spaces, we gather large collections of NE and PoS expressions across types that define the perturbation spaces \mathcal{T}_E and \mathcal{T}_P , which we gather from the Wikipedia paragraphs used in the *SQuAD2.0* training set, with the pretrained taggers in *spaCy*, and the Penn Treebank tag set for PoS. This results on average in 5,126 different entities per entity type, and 2,337 different tokens per PoS tag. When computing PoS perturbations, we found it useful to disregard perturbations of particular PoS types that often led to only minor changes or incorrectly formed expressions, such as punctuation or determiners; details on these can be found in Appendix A. As the number of possible perturbations to consider is potentially very large, we limit the beam search at each step to a maximum of η randomly chosen type-consistent entities from E , or tokens from P , and re-sample these throughout the search. We use a beam width of $b = 5$, resulting in a bound to the total computation spent on adversarial search of $b \cdot \rho \cdot \eta$ model evaluations per sample, where ρ is the perturbation “radius” (maximum search depth).

We quantify vulnerability to the described attacks by measuring the fraction of evaluation samples for which at least one undersensitivity attack is found given a computational search budget, disregarding cases where a model predicts *NoAnswer*.²

²Altering such samples likely retains their unanswerability.

Results Fig. 2 shows plots for adversarial error rates on *SQuAD2.0* for both perturbation types. We observe that attacks based on PoS perturbations can already for very small search budgets ($\eta = 32$, $\rho = 1$) reach more than 60% attack success rates, and this number can be raised to 95% with a larger computational budget. For perturbations based on NE substitution, we find overall lower attack success rates, but still find that more than half of the samples can successfully be attacked with the budgets tested. Note that where attacks were found, we observed that there often exist multiple alternatives satisfying inequality 1.

These findings demonstrate that BERT is not necessarily considering the entire content of a comprehension question given to it, and that even though trained to tell when questions are unanswerable, the model often fails when facing adversarially selected unanswerable questions.

In a side experiment we also investigated undersensitivity attacks using NE perturbations on *SQuAD1.1*, which proves even more vulnerable with an adversarial error rate of 70% already using $\eta = 32$; $\rho = 1$ (compared to 34% on *SQuAD2.0*). While this demonstrates that undersensitivity is also an issue for *SQuAD1.1*, the unanswerable question behaviour is not really well-defined, rendering results difficult to interpret. On the other hand, the notable drop between the datasets demonstrates the effectiveness of the unanswerable questions added during training in *SQuAD2.0*.

5 Analysis of Vulnerable Samples

Qualitative Analysis of Attacks The attacks are potentially noisy, and the introduced substitutions are by no means guaranteed to result in semanti-

Original / Modified Question	Prediction	Annotation	Scores
What city in Victoria is called the cricket ground of Australia the Delhi Metro Rail Corporation Limited ?	Melbourne	valid	0.63 / 0.75
What were the annual every year carriage fees for the channels?	£30m	same answer	0.95 / 0.97
What percentage of Victorians are Christian Girlish ?	61.1%	valid	0.92 / 0.93
Which plateau is the left part achievement of Warsaw on?	moraine	sem. inconsist.	0.52 / 0.58

Table 1: Example adversarial questions (original), (attack), together with their annotation as either a valid counterexample or other type. Top 2: Named Entity perturbations. Bottom 2: PoS perturbations.

	PoS	NE
Valid attack	51%	84%
Syntax error	10%	6%
Semantically incoherent	24%	5%
Same answer	15%	5%

Table 2: Analysis of undersensitivity attack samples for both PoS and named entity (NE) perturbations.

cally meaningful and consistent expressions, or require a different answer than the original. To gauge the extent of this we inspect 100 successful attacks conducted at $\rho = 6$ and $\eta = 256$ on *SQuAD2.0*, both for PoS and NE perturbations. We label them as either: i) Having a syntax error (e.g. *What would platform lower if there were fewer people?*). These are mostly due to cascading errors stemming from incorrect NE / PoS tag predictions. ii) Semantically incoherent (*Who built the monks?*). iii) Questions that require the same correct answer as the original, e.g. due to a paraphrase. iv) Valid attacks: Perturbed questions that would either demand a different answer or are unanswerable given the text (e.g. *When did the United States / Tuvalu withdraw from the Bretton Woods Accord?*)

Table 1 shows several example attacks along with their annotations, and in Table 2 the respective proportions are summarised. We observe that a non-negligible portion of questions has some form of syntax error or incoherent semantics, especially for PoS perturbations. Questions with the identical correct answer are comparatively rare. Finally, about half (51%) of the attacks in PoS, as well as 84% for NE are valid questions that should either have a different answer, or are unanswerable.

Overall the NE perturbations result in cleaner questions than PoS perturbations, which suffer from semantic inconsistencies in about a quarter

(24%) of the cases. While these questions have some sort of inconsistency (e.g. *What year did the case go before the supreme court?* vs. a perturbed version *What scorer did the case go before the supreme court?*), it is remarkable that the model assigns higher probabilities to the original answer even when faced with incoherent questions, casting doubt on the extent to which the replaced question information is used to determine the answer. Since NE-based attacks have a substantially larger fraction of valid, well-posed questions, we will focus our study on these for the remainder of this paper.

Characterising Successfully Attacked Samples We observe that models are vulnerable to undersensitivity adversaries, yet not all samples are successfully attacked. This raises the question of what distinguishes samples that can and cannot be attacked. We thus examine various characteristics, aiming to understand model vulnerability causes.

First, successfully attacked questions produce lower original prediction probabilities, on average 72.9% vs. 83.8% for unattackable questions. That is, there exists a direct inverse link between a model’s original prediction probability and sample vulnerability to an undersensitivity attack. The adversarially chosen questions had an average probability of 78.2% – a notable gap to the original questions. It is worth noting that search halted once a single question with higher probability was found; continuing the search increases the respective probabilities.

Second, vulnerable samples are furthermore less likely to be given the correct prediction overall. Concretely, evaluation metrics for vulnerable examples are 56.4%/69.6% EM/F₁, compared to 73.0%/76.5% on the whole dataset (-16.6% and -6.9% EM/F₁).

Attackable questions have on average 12.3 tokens, whereas unattackable ones are slightly shorter

Adv. budget η		Undersensitivity Error Rate				HasAns		NoAns	Overall	
		@32	@64	@128	@256	EM	F ₁	EM/F1	EM	F ₁
SQuAD	BERT LARGE	44.0	50.3	52.7	54.7	70.1	77.1	76.0	73.0	76.5
	+ Data Augment.	4.5	9.1	11.9	18.9	66.1	72.2	80.7	73.4	76.5
	+ Adv. Training	11.0	15.9	22.8	28.3	69.0	76.4	77.1	73.0	76.7
NewsQA	BERT BASE	34.2	34.7	36.4	37.3	41.6	53.1	61.6	45.7	54.8
	+ Data Augment.	7.1	11.6	17.5	20.8	41.5	53.6	62.1	45.8	55.3
	+ Adv. Training	20.1	24.1	26.9	29.1	39.0	50.4	67.1	44.8	53.9

Table 3: Breakdown of undersensitivity error rate overall (lower is better), and standard performance metrics (EM, F₁; higher is better) on different subsets of *SQuAD2.0* and *NewsQA* evaluation data, all in [%].




Figure 3: Named entity type characteristics of attackable vs. unattackable samples.

with on average 11.1 tokens.

Next we considered the distribution of different question types (*What*, *Who*, *When*, ...) for both attackable and unattackable samples and did not observe notable differences apart from the single most frequent question type *What*; it is a lot more prevalent among the unattacked questions (56.4%) than among successfully attacked questions (42.1%). This is by far the most common question type, and furthermore one that is comparatively open-ended and does not prescribe particular type expectations to its answer, as e.g., a *Where* question would require a location. A possible explanation for the prevalence of the *What* questions among unsuccessfully attacked samples is that the model cannot rely on type constraints alone to arrive at its predictions (Sugawara et al., 2018), and is thus less prone to such exploitation - see Section 6 for a more in-depth analysis.

Finally, Figure 3 shows a histogram of the 10 most common NE tags appearing in unsuccessfully attacked samples, and the corresponding fraction of replaced entities in successfully attacked samples. Besides one exception, the distributions are remarkably similar: Undersensitivity can be induced for a variety of entity types used in the perturbation. Notably, questions with geopolitical entities (*GPE*)

are particularly error-prone. A possible explanation can be provided by observations regarding (non-contextualised) word vectors, which cluster geopolitical entities (e.g. countries) together, thus making them harder to distinguish for a model operating on these embeddings (Mikolov et al., 2013).

6 Robustness to Undersensitivity Attacks

We will now investigate methods for mitigating excessive model undersensitivity. Prior work has considered both data augmentation and adversarial training for more robust models; we will conduct experiments with both. Adding a robustness objective can negatively impact standard test metrics (Tsipras et al., 2019), and it should be noted that there exists a natural trade-off between performance on one particular test set and performance on a dataset of adversarial inputs. We perform data augmentation and adversarial training by adding a corresponding loss term to the log-likelihood training objective: $\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{llh}}(\Omega) + \lambda \cdot \mathcal{L}_{\text{llh}}(\Omega')$ where Ω is the standard training data, fit with a discriminative log-likelihood objective, Ω' either a set of augmentation data points, or of successful adversarial attacks where they exist, and $\lambda \in \mathbb{R}_+$ a hyperparameter. In data augmentation, we randomly sample perturbed input questions, whereas in adversarial training we perform an adversarial search to identify them. In both cases, alternative data points in Ω' will be fit to a *NULL* label to represent the *NoAnswer* prediction - again using a log-likelihood objective. We continuously update Ω' throughout training to reflect adversarial samples with respect to the current model. We conduct experiments on both *SQuAD2.0* and *NewsQA*; details on training and hyperparameters can be found in the appendix.

Experimental Outcomes Results for these experiments can be found in Table 3 for the two da-

tasets, respectively. First, we observe that both data augmentation and adversarial training substantially reduce the number of undersensitivity errors the model commits, consistently across adversarial search budgets, and consistently across the two datasets. This demonstrates that both training methods are effective defences and can mitigate – but not eliminate – the model’s undersensitivity problem. Notably the improved robustness – especially for data augmentation – is possible without sacrificing performance in the overall standard metrics EM and F_1 , even slight improvements are possible. Second, data augmentation is a more effective defence training strategy than adversarial training. This holds true both for standard and adversarial metrics and hints at potential adversarial overfitting.

Finally, a closer inspection of how performance changes on answerable (*HasAns*) vs. unanswerable (*NoAns*) samples of the datasets reveals that models with modified training objectives show improved performance on unanswerable samples, while sacrificing some performance on answerable samples.³ This suggests that the trained models – even though similar in standard metrics – evolve on different paths during training: the modified objectives prioritise unanswerable questions to a higher degree.

Evaluation on Held-Out Perturbation Spaces In Table 3 results are computed using the same perturbation spaces also used during training. These perturbation spaces are relatively large, and questions are about a disjoint set of articles at evaluation time. Nevertheless there is the potential of overfitting to the particular perturbations used during training. To measure the extent to which the defences generalise also to new, held out sets of perturbations, we assemble a new, disjoint perturbation space of identical size per NE tag as those used during training, and evaluate models on attacks with respect to these perturbations. Named entities are chosen from English Wikipedia using the same method as for the training perturbation spaces, and chosen such that they are disjoint from the training perturbation space. We then execute adversarial attacks using these new attack spaces on the previously trained models, and find that both vulnerability rates of the standard model, as well as relative defence success transfer to the new attack spaces. For example, with $\eta = 256$ we observe vulnerability ratios of 51.7%, 20.7%, and 23.8% on *SQuAD2.0* for standard training, data augmen-

³The *NoAns* threshold is tuned on the respective valid. sets.

	Person		Date		Numerical	
	EM	F_1	EM	F_1	EM	F_1
BERT w/ bias	55.9	63.1	48.9	58.2	38.7	48.0
+ Data Augm.	59.1	66.6	58.4	65.6	48.7	58.9
BERT w/o bias	69.2	78.1	73.2	81.7	69.6	80.5

Table 4: Robust training leads to improved generalisation with train/test distribution mismatch (w/ bias). Bottom: control experiment without train/test mismatch.

tation, and adversarial training, respectively. Detailed results for different values of η , as well as for *NewsQA* can be found in Appendix D.

Generalisation in a Biased Data Setting Datasets for high-level NLP tasks often come with annotation and selection biases; models then learn to exploit shortcut triggers which are dataset but not task-specific (Gururangan et al., 2018). For example, a model might be confronted with question/paragraph pairs which only ever contain one type-consistent answer span, e.g. mention one number in a text with a *How many...?* question. It is then sufficient to learn to pick out numbers from text to solve the task, irrespective of other information given in the question. Such a model might then have trouble generalising to articles that mention several numbers, as it never learned that it is necessary to take into account other relevant question information to determine the correct answer.

We test models in such a scenario: a model is trained on *SQuAD1.1* questions with paragraphs containing only a single type-consistent answer expression for either a person, date, or numerical answer. At test time, we present it with question/article pairs of the same respective question types, but now there are *multiple* possible type-consistent answers in the paragraph. We obtain such data from Lewis and Fan (2019), who first described this biased data scenario. Previous experiments on this dataset were conducted without dedicated development set, so while using the same training data, we split the test set with a 40/60% split into development and test data.⁴ We then test both a vanilla fine-tuned BERT BASE transformer model, and a model trained to be less vulnerable to undersensitivity attacks using data augmentation. Finally, we perform a control experiment, where we join and shuffle all data points from train/dev/test (of each question type, respectively), and split the

⁴We also include an experiment with the previous data setup used by Lewis and Fan (2019), see Appendix G.

	AddSent		AddOneSent		Dev 2.0	
	EM	F ₁	EM	F ₁	EM	F ₁
BERT Large	61.3	66.0	70.1	74.9	78.3	81.4
+ Data Augm.	64.0	70.3	70.2	76.5	78.9	82.1

Table 5: Comparison between BERT LARGE and BERT LARGE + data augmentation using NE perturbations, on two sets of adversarial examples: ADDSENT and ADDONESENT from Jia and Liang (2017).

dataset into new parts of the same size as before, which now follow the same data distribution (w/o data bias setting). Table 4 shows the results. In this biased data scenario we observe a marked improvement across metrics and answer type categories when a model is trained with unanswerable samples. This demonstrates that the negative training signal stemming from related – but unanswerable – questions counterbalances the signal from answerable questions in such a way, that the model learns to better take into account relevant information in the question, which allows it to correctly distinguish among several type-consistent answer possibilities in the text, which the standard BERT BASE model does not learn well.

Evaluation on Adversarial SQuAD We next evaluate BERT LARGE and BERT LARGE + Augmentation Training on ADDSENT and ADDONESENT, which contain adversarially composed samples (Jia and Liang, 2017). Our results, summarised in Table 5, show that including altered samples during the training of BERT LARGE improves EM/F₁ scores by 2.7/4.3 and 0.1/1.6 points on the two datasets, respectively.

Transferability of Attacks We train a RoBERTa model (Liu et al., 2019) on SQuAD2.0, and conduct undersensitivity attacks ($\rho = 6$; $\eta = 256$). Notably the resulting undersensitivity error rates are considerably lower for RoBERTa (34.5%) than for BERT (54.7%). Interestingly then, when considering only those samples where RoBERTa was found vulnerable, BERT has an undersensitivity error rate of 90.7%. That is, the same samples tend to be vulnerable to undersensitivity attacks in *both* models. Even more, we find that concrete adversarial inputs x' selected with RoBERTa transfer when evaluating them on BERT for 17.5% of samples (i.e. satisfy Inequality 1).

7 Conclusion

We have investigated undersensitivity: a problematic behaviour of RC models, where they are overly

stable in their predictions when given semantically altered questions. A model’s robustness to undersensitivity attacks can be drastically improved with appropriate defences without sacrificing nominal performance, and the resulting models become more robust also in other adversarial data settings.

Acknowledgements

This research was supported by an Engineering and Physical Sciences Research Council (EPSRC) scholarship, and the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 875160.

References

- Chris Alberti, Daniel Andor, Emily Pitler, Jacob Devlin, and Michael Collins. 2019. Synthetic QA corpora generation with roundtrip consistency. In *ACL (1)*, pages 6168–6173. Association for Computational Linguistics.
- 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.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *International Conference on Learning Representations*.
- 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.
- 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.
- Allyson Ettinger, Sudha Rao, Hal Daumé III, and Emily M. Bender. 2017. Towards linguistically generalizable NLP systems: A workshop and shared task. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018.

- Pathologies of neural models make interpretations difficult. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3719–3728, Brussels, Belgium. Association for Computational Linguistics.
- Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hanna Hajishirzi, Gabriel Ilharco, Daniel Khoshabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. Evaluating nlp models via contrast sets.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. In *NAACL-HLT (2)*, pages 107–112. Association for Computational Linguistics.
- Hossein Hosseini, Baicen Xiao, and Radha Poovendran. 2017. Deceiving google’s cloud video intelligence API built for summarizing videos. In *CVPR Workshops*, pages 1305–1309. IEEE Computer Society.
- Minghao Hu, Furu Wei, Yuxing Peng, Zhen Huang, Nan Yang, and Dongsheng Li. 2019. Read + verify: Machine reading comprehension with unanswerable questions. In *AAAI*, pages 6529–6537. AAAI Press.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018a. 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.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018b. Adversarial example generation with syntactically controlled paraphrase networks. In *NAACL-HLT*, pages 1875–1885. Association for Computational Linguistics.
- Joern-Henrik Jacobsen, Jens Behrmann, Richard Zemel, and Matthias Bethge. 2019. Excessive invariance causes adversarial vulnerability. In *International Conference on Learning Representations*.
- 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.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*.
- Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*.
- Mike Lewis and Angela Fan. 2019. Generative question answering: Learning to answer the whole question. In *International Conference on Learning Representations*.
- 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.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2018. Gender bias in neural natural language processing. *ArXiv*, abs/1807.11714.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Tong Niu and Mohit Bansal. 2018. Adversarial oversensitivity and over-stability strategies for dialogue models. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 486–496, Brussels, Belgium. 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.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018a. Anchors: High-precision model-agnostic explanations. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018b. 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. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of NLP models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. *ICLR*.

- 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.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Alon Talmor and Jonathan Berant. 2019. MultiQA: An empirical investigation of generalization and transfer in reading comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4911–4921, Florence, Italy. Association for Computational Linguistics.
- 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.
- Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. 2019. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for NLP. *CoRR*, abs/1908.07125.
- 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.
- Johannes Welbl, Po-Sen Huang, Robert Stanforth, Sven Gowal, Krishnamurthy (Dj) Dvijotham, Martin Szummer, and Pushmeet Kohli. 2020. Towards verified robustness under text deletion interventions. In *International Conference on Learning Representations*.
- Adams Wei Yu, David Dohan, Quoc Le, Thang Luong, Rui Zhao, and Kai Chen. 2018. Fast and accurate reading comprehension by combining self-attention and convolution. In *International Conference on Learning Representations*.
- Wei Emma Zhang, Quan Z. Sheng, and Ahoud Abdulrahmn F. Alhazmi. 2019. Generating textual adversarial examples for deep learning models: A survey. *CoRR*, abs/1901.06796.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. Gender bias in coreference resolution: Evaluation and debiasing methods. 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 15–20, New Orleans, Louisiana. Association for Computational Linguistics.
- Haichao Zhu, Li Dong, Furu Wei, Wenhui Wang, Bing Qin, and Ting Liu. 2019. Learning to ask unanswerable questions for machine reading comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4238–4248, Florence, Italy. Association for Computational Linguistics.

SQuAD2.0 Adv. budget η	Undersensitivity Error Rate			
	@32	@64	@128	@256
BERT LARGE	40.7	45.2	48.6	51.7
+ Data Augment.	4.8	7.9	11.9	20.7
+ Adv. Training	9.2	12.2	16.5	23.8

Table 6: Breakdown of undersensitivity error rate on *SQuAD2.0* with a held-out attack space (lower is better).

NewsQA Adv. budget η	Undersensitivity Error Rate			
	@32	@64	@128	@256
BERT BASE	32.8	33.9	35.0	36.2
+ Data Augment.	3.9	6.5	11.9	17.5
+ Adv. Training	17.6	20.7	25.4	28.5

Table 7: Breakdown of undersensitivity error rate on *NewsQA* with a held-out attack space (lower is better).

A Appendix: PoS Perturbation Details.

We exclude these PoS-tags when computing perturbations: ‘IN’, ‘DT’, ‘:’, ‘VBD’, ‘VBZ’, ‘WP’, ‘WRB’, ‘WDT’, ‘CC’, ‘MD’, ‘TO’.

B Appendix: BERT Training Details - *SQuAD2.0*

We first train a BERT LARGE model on the full training set for 2 epochs, where it reaches 78.32%EM and 81.44%F₁, largely comparable to results (78.7%EM and 81.9%F₁) reported by Devlin et al. (2019). We then however choose a different training setup as we would like to conduct adversarial attacks on data inaccessible during training: we split off 5% from the original training set for development purposes and retain the remaining 95% for training, stratified by articles. We use this development data to tune hyperparameters and perform early stopping, evaluated every 5,000 steps with batch size 16 and patience 5, and tune hyperparameters for defence on it.

C Appendix: Adversarial Defence Experiments

SQuAD2.0 We train the BERT LARGE model, tuning the hyperparameter $\lambda \in \{0.0, 0.01, 0.1, 0.25, 0.5, 0.75, 1.0, 2.0\}$, and find $\lambda = 0.25$ to work best for both defences. We tune the threshold for predicting *NoAnswer* based on validation data and report results on the test set (the original *SQuAD2.0* Dev set). All experiments are executed with batch size 16, NE




Figure 4: Vulnerability to undersensitivity attacks on *NewsQA*.

perturbations are used for the defence methods, and adversarial attacks with $\eta = 32$ and $\rho = 1$ in adversarial training. Where no attack is found for a given question sampled during SGD training, we instead consider a different sample from the original training data. We evaluate the model on its validation data every 5,000 steps and perform early stopping with a patience of 5.

NewsQA Following the experimental protocol for *SQuAD*, we further test a BERT BASE model on *NewsQA*, which – like *SQuAD2.0* – contains unanswerable questions. As annotators often do not fully agree on a single annotation in *NewsQA*, we opt for a conservative choice and filter the dataset, such that only samples with the same majority annotation are retained, following the preprocessing pipeline of Talmor and Berant (2019).

D Appendix: Generalisation to Held-out Perturbations

Vulnerability results for new, held-out perturbation spaces, disjoint from those used during training, can be found in Table 6 for *SQuAD2.0*, and in Table 7 for *NewsQA*.

E Appendix: Adversarial Example from a Question Collection

Searching in a large collection of (mostly unrelated) natural language questions, e.g. among all questions in the *SQuAD2.0* training set, yields several cases where the prediction of the model increases, compared to the original question, see Table 8 for one such example. These cases are however relatively rare, and we found the yield of this type of search to be very low.

F Appendix: Attack Examples

Table 9 shows more examples of successful adversarial attacks on *SQuAD2.0*.

G Appendix: Biased Data Setup

For completeness and direct comparability, we also include an experiment with the data setup of Lewis and Fan (2019) (not holding aside a dedicated validation set). Results can be found in Table 10. We again observe improvements in the biased data setting. Furthermore, the robust model outperforms GQA (Lewis and Fan, 2019) in two of the three subtasks.

H Appendix: Vulnerability Analysis on NewsQA

Fig. 4 depicts the vulnerability of a BERT LARGE model on *NewsQA* under attacks using NE perturbations.

Given Text	[...] The Normans were famed for their martial spirit and eventually for their Christian piety, becoming exponents of the <i>Catholic orthodoxy</i> [...]
Q (orig)	What religion were the Normans? (0.78)
Q (adv.)	IP and AM are most commonly defined by what type of proof system? (0.84)

Table 8: Drastic example for lack of specificity: unrelated questions can trigger the same prediction (here: *Catholic orthodoxy*), and with higher probability.

Original / Modified Question	Prediction	Annotation	Scores
What ethnic neighborhood in (Fresno) (Kilbride) had primarily Japanese residents in 1940?	Chinatown	valid	(0.998) (0.999)
The (Mitchell Tower) (MIT) is designed to look like what Oxford tower?	Magdalen Tower	valid	(0.96) (0.97)
What are some of the accepted general principles of (European Union) (Al-Andalus) law?	fundamental rights [...]	valid	(0.59) (0.61)
What does the EU’s (legitimacy) (digimon) rest on?	the ultimate authority of [...]	valid	(0.38) (0.40)
What is Jacksonville’s hottest recorded (temperature) (atm)?	104°F	valid	(0.60) (0.62)
Who leads the (Student) (commissioning) Government?	an Executive Committee	same answer	(0.61) (0.65)

Table 9: Example adversarial questions ((original), (attack)), together with their annotation as either a valid counterexample or other type. Top 3 rows: Named entity (NE) perturbations. Bottom 3 rows: PoS perturbations.

	Person		Date		Numerical	
	EM	F ₁	EM	F ₁	EM	F ₁
GQA (Lewis and Fan, 2019)	53.1	61.9	64.7	72.5	58.5	67.6
BERT BASE - w/ data bias	66.0	72.5	67.1	72.0	46.6	54.5
+ Robust Training	67.4	72.8	68.1	74.4	56.3	64.5

Table 10: Robust training leads to improved generalisation under train/test distribution mismatch (data bias).