

Pay Attention to Real World Perturbations! Natural Robustness Evaluation in Machine Reading Comprehension

Anonymous ACL submission

Abstract

As neural language models achieve human-comparable performance on Machine Reading Comprehension (MRC) and see widespread adoption, ensuring their robustness in real-world scenarios has become increasingly important. Current robustness evaluation research, though, primarily develops synthetic perturbation methods, leaving unclear how well they reflect real life scenarios. Considering this, we present a framework to automatically examine MRC models on naturally occurring textual perturbations, by replacing paragraph in MRC benchmarks with their counterparts based on available Wikipedia edit history. Such perturbation type is *natural* as its design does not stem from an artificial generative process, inherently distinct from the previously investigated synthetic approaches. In a large-scale study encompassing SQUAD datasets and various model architectures we observe that natural perturbations result in performance degradation in pre-trained encoder language models. More worryingly, these state-of-the-art Flan-T5 and Large Language Models (LLMs) inherit these errors, with the largest observed drop reaching 28.28%. Further experiments demonstrate that our findings generalise to natural perturbations found in other more challenging MRC benchmarks such as DROP and HOTPOTQA. In an effort to mitigate these errors, we show that robustness to natural perturbations can be improved through adversarial training for encoder-only models or through in-context demonstrations of perturbed instances for LLMs, although a more generalisable and effective defence strategy remains to be developed.

1 Introduction

Transformer-based pre-trained language models demonstrate remarkable efficacy in addressing questions based on a given passage of text, a task commonly referred to as Machine Reading Comprehension (MRC) (Devlin et al., 2019; Brown

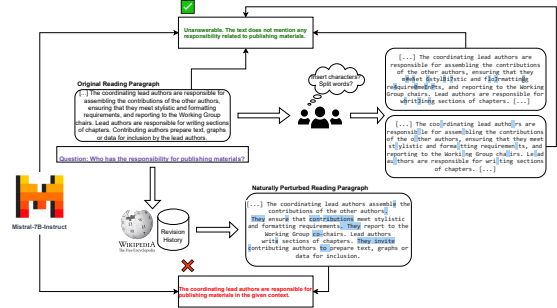


Figure 1: Given a reading paragraph, we extract and use Wikipedia revision history to construct its naturally perturbed version for a more realistic robustness evaluation (Bottom), rather than relying on a set of synthetic methods (Top). While Mistral-7B-Instruct-v0.2 generates the correct answer for both the original and synthetically perturbed passages, it fails under natural perturbation.

et al., 2020; He et al., 2021; Wei et al., 2022; Touvron et al., 2023; OpenAI et al., 2024b). Despite these advancements, high-performing MRC systems are known to succeed by exploiting shortcuts in existing benchmark datasets rather than truly demonstrating understanding of the passage and question, thereby exhibiting a lack of robustness to various types of test-time perturbations. Such phenomenon is observed not only in encoder-only and encoder-decoder models (Ho et al., 2023; Schlegel et al., 2023), but also in state-of-the-art (SOTA) Large Language Models (LLMs) (Levy et al., 2023; Fang et al., 2023; Gupta et al., 2024).

Extensive textual perturbation methods have been developed to reveal the vulnerabilities of MRC models to various linguistic challenges (Ribeiro et al., 2018; Jiang and Bansal, 2019; Welbl et al., 2020; Tan et al., 2020; Tan and Joty, 2021; Schlegel et al., 2021; Cao et al., 2022; Tran et al., 2023). Despite the insights they provide, all of them craft perturbations in a synthetic manner—that is, based on hypothesised manipulation

strategies (Figure 1 Top)—which may not accurately reflect the types of challenges encountered in real-world scenarios. As a result, there is a risk of overlooking the genuine weaknesses of reading comprehension systems when deployed in practical settings, potentially hindering efforts to improve their reliability in real applications.

To counteract this issue, in this paper, we develop a framework to inject textual changes that arise in real-world conditions into MRC datasets and audit how well contemporary language models perform under such perturbations. We deem them as *natural* because the perturbation process does not involve any artificial manipulation, in line with the definitions by Belinkov and Bisk (2018); Hendrycks et al. (2021); Pedraza et al. (2022); Agarwal et al. (2022); Le et al. (2022) (Figure 1 Bottom). Results of robustness evaluation are therefore more representative of real-world applications. Similar to Belinkov and Bisk (2018), our approach utilises Wikipedia revision histories as the source of natural perturbations, given that the differences between revisions authentically capture the textual modifications made by human editors in the real world. Despite this, significant differences exist in the perturbation construction methodology between us. Perturbation in (Belinkov and Bisk, 2018) is restricted to single word replacements and applied on non-English source-side sentences in machine translation. In detail, they build a look-up table of possible lexical replacements by harvesting naturally occurring errors (typos, misspellings, etc.) from available corpora of French/German Wikipedia edits (Max and Wisniewski, 2010; Zesch, 2012). Afterwards, they replace every word in the source-side sentences with an error if one exists in the look-up table. Different from (Belinkov and Bisk, 2018), our approach does not restrict the perturbation level and utilise English Wikipedia. By comparing the variances between each adjacent revision, we identify perturbed versions for each Wikipedia reading passage in the original MRC benchmarks (if it exists). This enables us to capture more comprehensive and critical natural perturbation patterns (see Section 5.2) that can not be possible to capture in (Belinkov and Bisk, 2018). Our perturbation method only alter the reading context, while the questions and ground truth answers remain unchanged.

With the established framework, we conduct extensive experiments on six datasets, evaluating forty-two models, including recently proposed

LLMs. Experimental results on Stanford Question Answering Dataset (SQUAD) (Rajpurkar et al., 2016, 2018) indicate that natural perturbations encompass rich linguistic variations and can lead to failures in the encoder-only models, while humans are almost undeterred by their presence. Crucially, these errors also transfer to larger and more powerful models, such as FLan-T5 and SOTA LLMs, with performance drops ranging from 4.4% to 28.28%. These findings also generalise to other and more challenging MRC benchmarks (e.g., Mistral-7B-Instruct-v0.2’s 5.25% decrease on BOOLQ (Clark et al., 2019) and Llama-3.2-1B-Instruct’s 9.98% decline on DROP (Dua et al., 2019)), emphasising the harmful effects of natural perturbations. Adversarial re-training/in-context demonstration with either naturally or synthetically perturbed MRC instances can enhance the robustness against natural perturbations, with the latter sometimes providing greater benefits. However, there is still ample room for improvement, calling for better defense strategies.

The contributions of this paper are as follows:

- A framework—based on Wikipedia revision history—for studying model robustness under real-world natural perturbations. This is relevant, even in the LLM era, as our framework can be applied to any other tasks with input from Wikipedia and also to any types of models.
- Perturbed datasets for six diverse MRC tasks. Two SQUAD challenge sets derived from error analysis of encoder-only models, on which SOTA LLMs struggle, even without being involved in the creation in any capacity.
- Empirical demonstration of the validity of natural perturbations across both encoder-only models and LLMs, their characterisation by different linguistic phenomena and their harmful effects on diverse model architectures across benchmarks generated with the proposed framework.
- Showcasing adversarial re-training with natural or, especially, synthetic perturbations, as well as adversarial in-context demonstrations as a way to enhance the robustness of encoder-only models and LLMs, respectively, against natural perturbations.

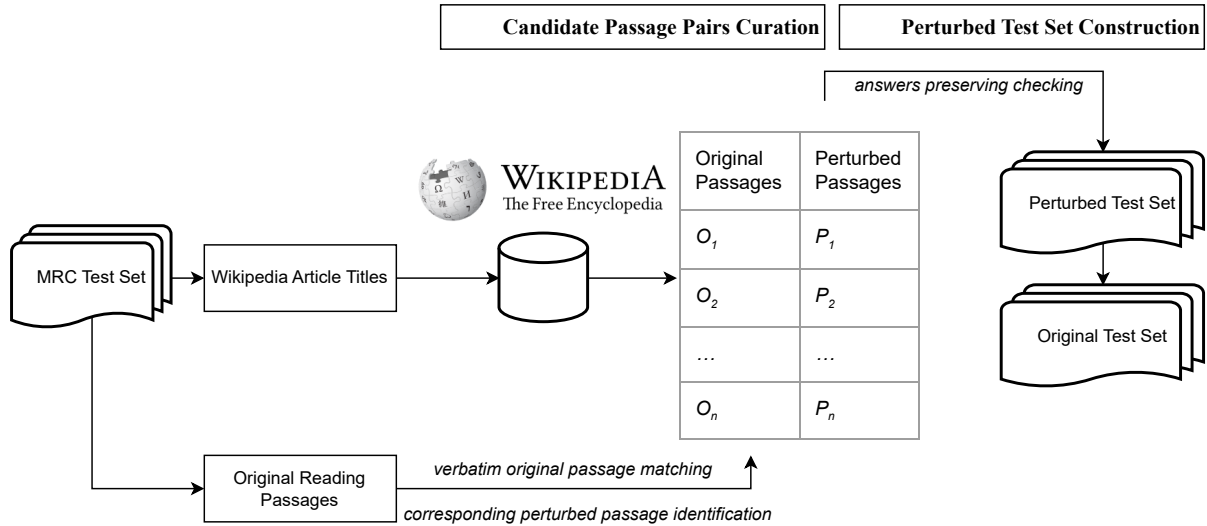


Figure 2: Process of generating naturally perturbed MRC test sets.

2 Related Work

Robustness Evaluation in MRC A typical approach to evaluate the robustness of MRC models is via test-time perturbation. This line of research develops different perturbation methods as attacks, such as adversarial distracting sentence addition (Jia and Liang, 2017; Tran et al., 2023), low-level attacks (Eger and Benz, 2020), word substitution (Wu et al., 2021), character swap (Si et al., 2021), entity renaming (Yan et al., 2022) and paraphrasing (Gan and Ng, 2019; Lai et al., 2021; Wu et al., 2023). Our work also fits within the category of test-time perturbation, but differs from previous works in that we introduce perturbations that naturally occur in real-world scenarios, therefore contributing to a more practical robustness test.

Natural Perturbation for Robustness Assessment Compared with deliberately crafting the perturbed instances, the study of natural perturbation is quite under-explored. In the computer vision domain, researchers find that real-world clean images without intentional modifications can confuse deep learning models as well, terming them as natural adversarial examples (Hendrycks et al., 2021; Pedraza et al., 2022). Similarly, in the field of Natural Language Processing (NLP), naturally occurring perturbations extracted from human-written texts can also degrade model performance in tasks such as machine translation (Blinkov and Bisk, 2018) and toxic comments detection (Le et al., 2022). Motivated by these, we attempt to harvest natural perturbations from available Wikipedia revision histories and utilise them to

modify the original MRC instances. To the best of our knowledge, we are the first to investigate MRC model robustness under real natural perturbations.

3 Natural Perturbation Pipeline

We design a pipeline to automatically construct label-preserving stress MRC test sets with noises that occur in real-world settings by leveraging Wikipedia revision histories (Figure 2). Our approach comprises two modules: *candidate passage pairs curation* and *perturbed test set construction*.

Candidate passage pairs curation. For each English Wikipedia article within the development set¹ of MRC datasets, we systematically extract its entire revision histories and preprocess them, including the removal of markups and the segmentation of content. Subsequently, we obtain the content differences between each current revision and the previous adjacent one, identifying three distinct editing patterns: addition, deletion, and modification. In the case of an edit falling within the modification pattern, we retain the paragraph from the prior version as the *original* and the corresponding one from the current version as the *perturbed*, provided both paragraphs exceed 500 characters².

Perturbed test set construction. To generate the naturally perturbed test set, we begin by acquiring all reading passages from the development set

¹Since not all test sets are public, we apply natural perturbations to the development sets. For simplicity, we use the term “test set” throughout.

²This threshold setting adheres to the methodology employed in the collection of SQuAD 1.1 (Rajpurkar et al., 2016).

of each MRC dataset and identifying their entries in the collection of previously extracted candidate original passages, along with the corresponding perturbed counterparts. Subsequently, for the matched original passages with a single occurrence, we keep them and the corresponding perturbed passages; whereas for those with multiple occurrences, we randomly select one instance for each and extract its perturbed version. After obtaining the perturbed reading passages, we retain only those with at least one question where all annotated ground truth answers (or all plausible answers for the unanswerable question) can still be located within the perturbed context, resulting in the *Perturbed* test set. For the sake of comparison, we also construct an *Original* version of the test set keeping only the original passages and questions corresponding to those that were included in the *Perturbed* version.

4 Experiment Setup

4.1 Datasets

We use six English MRC datasets: SQUAD 1.1 (Rajpurkar et al., 2016), SQUAD 2.0 (Rajpurkar et al., 2018), BOOLQ (Clark et al., 2019), DROP (Dua et al., 2019), HOTPOTQA (distractor) (Yang et al., 2018) and TYDI QA (gold passage task in English) (Clark et al., 2020). These are chosen as their reading passages are sourced from Wikipedia, thereby enabling the utilisation of Wikipedia editing histories to generate the naturally perturbed test set.

4.2 Models

Our evaluation study involves MRC models across three different types: encoder-only, encoder-decoder, and decoder-only. Under the encoder-decoder and decoder-only model evaluation settings, we reframe MRC as the text generation task based on the given context and question. Access to and experimentation with all models are possible via the use of the HuggingFace’s *Transformers* library (Wolf et al., 2020), the vLLM library (Kwon et al., 2023), two 80GB Nvidia A100 GPUs and the OpenAI ChatGPT API.

Encoder-only: We select BERT (Devlin et al., 2019) and its various variants for evaluation, including DistilBERT (Sanh et al., 2019), SpanBERT (Joshi et al., 2020), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2020) and DeBERTa (He et al., 2021). Some of these model types also come with different variations, such as size (e.g., *base* and

large for RoBERTa), versions (e.g., *v1* and *v2* for ALBERT) and whether the input text is cased or not (e.g., *cased* and *uncased* for BERT), all of which are included in the evaluation. We fine-tune these encoder-only pre-trained language models on the training set of the two SQUAD datasets (Rajpurkar et al., 2016, 2018) and evaluate them on the constructed original and perturbed test sets. Model details and the hyperparameters used in model fine-tuning are shown in Appendix A.

Encoder-Decoder: Instruction finetuning has been demonstrated to be effective in enhancing zero-shot performance of pretrained language models, resulting in the development of Finetuned Language Net (FLAN) (Wei et al., 2022). In this work, we use the instruction-finetuned version of T5 model class, specifically the Flan-T5 (Chung et al., 2022), available in sizes ranging from *small* (80M), *base* (250M), *large* (780M) to *xl* (3B). During evaluation, we utilise the instruction templates from MRC task collection in open-sourced FLAN repository and report the model performance as the average of those obtained across the employed templates. Refer to Appendix B for various instruction templates used for the evaluation on the test sets with the format as the two SQUAD datasets.

Decoder-only: There is an exponential increase of pre-trained generative LLMs and their fine-tuned chat versions, inspired by the remarkable success of ChatGPT (Bang et al., 2023). Therefore, our experiments incorporate a broad range of recently proposed language model families, including GPT 3.5 Turbo, GPT-4o (OpenAI et al., 2024a), Gemma (Mesnard et al., 2024), Gemma 2 (Riviere et al., 2024), Llama 2 (Touvron et al., 2023), Llama 3 and Llama 3.1 (Dubey et al., 2024), Llama 3.2, Mistral (Jiang et al., 2023), OLMo (Groeneveld et al., 2024), Qwen2.5 (Qwen et al., 2025), Falcon (Almazrouei et al., 2023), Falcon3 (Team, 2024), and DeepSeek LLM (DeepSeek-AI et al., 2024). The zero-shot prompts designed for soliciting their responses are presented in Appendix C.

4.3 Evaluation Metrics

In line with existing literature, we choose the (instance-averaged) Token-F1 score to assess the performance of both encoder-only and encoder-decoder models (Rajpurkar et al., 2016), as on SQUAD-style test sets, they are optimised to output the shortest continuous span from the context as the answer (or predict the question as unanswer-

able) during inference. However, the outputs of the decoder-only models do not consistently adhere to the instruction due to their conversational style, rendering F1 unsuitable for evaluation. Consequently, we employ a more lenient metric, namely Inclusion Match (IM), which measures whether the response of the model contains any of the ground truth answers (Bhuiya et al., 2024). Furthermore, if the model’s output includes phrases such as “I cannot answer this/the question” or “unanswerable”³, we deem that the model believes the question is not answerable. Model robustness is quantified by measuring the relative variation in performance (as reflected in the F1 or IM) under natural perturbations.

5 MRC under Natural Perturbation

In this section, we present and discuss the results of our experiments. We first evaluate encoder-only models on SQUAD to establish a baseline evaluation of model behaviour under natural perturbations. While neither represents the current SOTA, SQUAD’s simplicity, the stable, super-human performance of encoder-only models, and its advantage of being free from benchmark leakage impact enable a focused and controlled examination of perturbation effects (Section 5.1), error sources (Section 5.2), and adversarial instance validity (Section 5.3). Then, we investigate the transferability of errors from encoder-only models to other architectures, showing both FLAN-T5 and LLMs carry these errors significantly (Section 5.4)⁴. We finally generalise the findings from the baseline evaluation to SOTA LLMs and other more complex datasets (Section 5.5).

5.1 Are Encoder-only Models Resilient to Natural Perturbation?

Table 1 presents the relative F1 change for encoder-only MRC models on the naturally perturbed SQUAD test set. It can be seen that overall, the performance of all the examined models decreases, indicating that *encoder-only MRC models suffer from natural perturbation*. Nonetheless, the performance decline of all models is not at a significant level, with the biggest drop being only

³We collate a collection of such phrases by manually examining the decoder-only models’ outputs (Check Appendix D for the full set).

⁴To supplement, we further evaluate the full test set on FLAN-T5 and several LLMs, and measure the transferability of adversarial examples across all model architectures (Appendix E).

3.06%. This suggests that these models also exhibit considerable robustness to natural perturbation.

Victim	SQUAD 1.1	SQUAD 2.0	
		Overall	(Ans./Unans.)
distilbert-base	−0.6	−0.71	(−2.76/1.71)
bert-base-cased	−0.21	−0.63	(−1.84/0.6)
bert-base-uncased	−0.87	−0.49	(−1.88/0.94)
bert-large-cased	−0.63	−0.53	(−1.61/0.55)
bert-large-uncased	−0.35	−1.38	(−2.51/−0.24)
spanbert-base-cased	−0.26	−1.24	(−2.66/0.15)
spanbert-large-cased	−0.51	−1.20	(−1.9/−0.56)
roberta-base	−0.61	−0.60	(−2.09/0.81)
roberta-large	−0.29	−1.52	(−2.6/−0.54)
albert-base-v1	−1.0	−1.07	(−2.02/−0.22)
albert-base-v2	−0.34	−1.08	(−2.03/−0.22)
albert-large-v1	−0.42	−0.41	(−1.42/0.52)
albert-large-v2	−0.8	−0.69	(−1.66/0.22)
albert-xxlarge-v1	−0.75	−1.23	(−3.06/0.49)
albert-xxlarge-v2	−0.46	−1.28	(−3.02/0.36)
deberta-large	−0.52	−1.05	(−2.2/0.0)

Table 1: Relative F1 change (%) for encoder-only MRC systems subjecting to natural perturbations. For SQUAD 2.0, the overall values are broken down to answerable and unanswerable questions, respectively.

5.2 Which Categories of Natural Perturbation Lead to Model Failure?

To investigate the sources of natural perturbation and reveal the perturbation phenomena contributing to encoder-only models’ error, we manually label linguistic features between passages where models succeed and fail, to identify how they differ.

Within the original and the naturally perturbed test set pair generated based on SQUAD 2.0 development set, we first identify 384 instances where at least one encoder-only model succeeds on the original but fails⁵ on the perturbed (i.e., being adversarial), and then randomly select the same number of instances on which all encoder-only models succeed on both the original and perturbed versions (Naik et al., 2018). We refer to these two types of instances as C2W (correct to wrong) and C2C (correct to correct) instances, respectively. Among the identified C2W and C2C instances, we further remove duplicates, resulting in 210 and 244 unique original and perturbed paragraph pairs, respectively. Furthermore, as natural perturbation can occasionally help the model to get the answer correct, we also filter 85 unique W2C (wrong to correct) in-

⁵For answerable questions, a model’s prediction is considered correct if Exact Match (EM) score equals 1, and incorrect if F1 score is 0 or it determines the question is unanswerable. For unanswerable questions, a model’s prediction is correct if it predicts the question is unanswerable, and wrong if it provides an answer span.

stances on which at least two encoder-only models fail on the original but succeed on the perturbed. Finally, utilising an 8-category taxonomy of the semantic edit intentions in Wikipedia revisions derived from Yang et al. (2017), the chosen 210 samples of C2W and C2C, as well as the 85 W2C were annotated, with 20% of the annotated C2W and C2C examples presented to a second annotator for additional validation. See Appendix F for the instruction provided to the annotators, along with detailed explanations of each edit intention. We calculate the (micro-averaged) F1 score to evaluate the inter-annotator agreement, which is 0.82. This suggests that the annotators’ annotations align closely. Figure 3 reports the annotation results.




Figure 3: The percentage (%) of samples annotated with each edit intention in the C2W, C2C and W2C categories. The percentages do not add up to 100% because a single revision may fall into multiple intentions.

Distribution of perturbation types shown in Figure 3 generally aligns with the edit intentions distribution annotated in (Yang et al., 2017), with *Copy Editing* and *Elaboration* appearing more frequently than others, such as *Clarification*, *Fact Update*, and *Refactoring*. This reflects the inherent characteristics of Wikipedia revisions. From Figure 3, we observe that there is no significant difference in the distribution of annotated edit intentions between C2W and C2C examples, suggesting that *though these types of natural perturbations confuse the encoder-only MRC models, there seems no correlation with human-perceivable features*. A roughly similar distribution is also observed in the W2C examples, which indicates that these natural perturbation types can also facilitate correct answers by the models, i.e., being beneficial. These demonstrate that on SQUAD 2.0, there might be no correlation between the quality of the naturally per-

turbed passage and its potential for being adversarial⁶. Certain text edits aimed at improving the passage quality, such as *Copy Editing* and *Elaboration*, do render the perturbation adversarial, whereas edits intended to damage the article may not consistently result in adversarial instances; in fact, *vandalism* can even assist models in providing correct answers. Instead, we infer that whether an edit to the passage can render the MRC instance adversarial or not depends on the location of the edits in relation to the question. Among the 384 C2W and C2C examples, we measure the proportion of answerable questions with the answer sentence(s) in the original passage remaining unmodified in the naturally perturbed version, which is 34.5% and 71.5%, respectively. This confirms our hypothesis that if the edits affect the answer sentence(s), there is a higher likelihood of the perturbed example becoming adversarial; otherwise, it might not. *Copy Editing* appears to alter the answer sentences in the reading passage more frequently, making it the most impactful category that confuses models (contributing to more than 40% of error cases), while other types have a lesser effect. Appendix G presents one perturbed example for each of the C2W, C2C, and W2C categories, respectively, along with the annotated natural perturbation type(s).

5.3 To What Extent Do Natural Adversarial Examples Preserve Validity?

To accurately assess a model’s robustness under perturbation, it is vital to examine the validity of adversarial example, i.e., whether humans can still find the correct answer under the perturbation (Dymishi et al., 2023). Two human annotators are recruited to verify the validity of the 210 C2W examples in Section 5.2 and the inter-annotator agreement is measured by computing the Cohen’s κ coefficient (Cohen, 1960). We then involve a third human annotator to annotate the adversarial examples on which the first two annotators disagree and take the majority label as ground truth. This validity verification process is detailed in Appendix H. Out of 210 C2W examples, we find that 86% of them are valid (0.77 Cohen’s κ), indicating that *a substantial proportion of natural adversarial examples for encoder-only MRC model(s) are valid*.

⁶We also find little or no significant correlation between the perturbation magnitude (measured as byte-level changes between the original and perturbed passages) and model failure, with point biserial correlation coefficient close to 0.

5.4 Can Errors from Encoder-only Models Affect Other Architectures?

We further investigate whether the errors identified in encoder-only models carry over to other more recent models and architectures, as SOTA advancements in NLP would suggest otherwise. Therefore, we propose an exhaustive search algorithm (Appendix I) to zoom in on the errors of encoder-only models as much as possible, curate the challenging natural perturbed test set, and finally examine the performance of Flan-T5 and LLMs. With the development set of SQUAD 1.1 and SQUAD 2.0 as the source, the algorithm results in two challenge perturbed test sets: NAT_V1_CHALLENGE (184 contexts, 234 questions) and NAT_V2_CHALLENGE (214 contexts, 442 questions (226 unanswerable)).

Table 2 shows the evaluation results on the newly generated challenge test sets. From the table, we observe that *the errors caused by natural perturbation in encoder-only MRC models transfer to both Flan-T5 and LLMs*. On the NAT_V1_CHALLENGE, flan-t5-small demonstrates the greatest susceptibility to natural perturbations, experiencing a 14.27% decrease in F1; while among LLMs, Gemma-7B-IT emerges as the least robust, with a 16.66% IM drop, followed by Gemma-2B-IT (−15.83%) and Llama-3.1-8B-Instruct (−15.61%). Transitioning to the NAT_V2_CHALLENGE, the base version of flan-t5 exhibits the largest performance decline at 13.83% and Falcon-7B-Instruct stands out as the LLM with the lowest robustness (−28.28%). Other LLMs such as Qwen2.5-7B-Instruct and deepseek-llm-7b-chat also show severe robustness loss, with drops of 12.21% and 11.29%, respectively. Further, we observe that the robustness of models under natural perturbations does not necessarily size-dependent. While larger models tend to exhibit greater robustness in some cases (e.g., Qwen2.5-14B-Instruct vs. Qwen2.5-3B-Instruct), exceptions within the Falcon and Llama model series suggest that factors beyond model size—such as training corpora, training and fine-tuning methodology, and architectural differences may also significantly affect susceptibility to natural perturbations. In Appendix J, we showcase two adversarial examples targeting LLMs sourced from our generated challenge sets.

Model	Performance			
	original vs. perturbed			
	NAT_V1_CHALLENGE		NAT_V2_CHALLENGE	
flan-t5-small	58.76/64.76	48.58/55.52 _{−14.27}	42.57/44.57	39.71/41.81 _{−6.19}
flan-t5-base	79.49/85.01	66.17/73.42 _{−13.63}	70.66/72.85	61.16/62.78 _{−13.83}
flan-t5-large	88.1/92.53	76.57/82.31 _{−11.05}	79.11/81.01	70.14/72.13 _{−10.96}
flan-t5-xl	86.25/91.57	75.0/81.45 _{−11.05}	83.71/85.84	73.19/74.86 _{−12.79}
GPT-3.5-turbo-0125	91.03	83.33 _{−8.46}	51.58	47.06 _{−8.76}
gpt-4o-2024-11-20	93.16	85.9 _{−7.79}	80.09	75.11 _{−6.22}
Gemma-2B-IT	51.28	43.16 _{−15.83}	55.66	50.23 _{−9.76}
Gemma-7B-IT	82.05	68.38 _{−16.66}	59.95	57.01 _{−4.9}
Gemma-2-2b-IT	85.47	78.21 _{−8.49}	48.87	43.44 _{−11.11}
Gemma-2-9b-IT	89.32	81.62 _{−8.62}	64.93	59.95 _{−7.67}
Llama-2-chat-7B	82.91	73.93 _{−10.83}	41.63	38.69 _{−7.06}
Llama-2-chat-13B	80.77	73.93 _{−8.47}	46.83	41.18 _{−12.06}
Llama-3-8B-Instruct	88.89	77.35 _{−12.98}	51.81	46.61 _{−10.04}
Llama-3.1-8B-Instruct	87.61	73.93 _{−15.61}	61.31	55.43 _{−9.59}
Llama-3.2-1B-Instruct	54.27	47.86 _{−11.81}	35.29	32.13 _{−8.95}
Llama-3.2-3B-Instruct	81.2	71.37 _{−12.11}	48.42	43.44 _{−10.29}
Mistral-7B-Instruct-v0.2	84.19	73.08 _{−13.2}	54.98	51.36 _{−6.58}
OLMo-7B-0724-Instruct	90.17	82.91 _{−8.05}	51.36	49.1 _{−4.4}
Qwen2.5-3B-Instruct	78.63	68.38 _{−13.04}	61.31	54.07 _{−11.81}
Qwen2.5-7B-Instruct	88.03	81.2 _{−7.76}	76.02	66.74 _{−12.21}
Qwen2.5-14B-Instruct	92.31	81.62 _{−11.58}	80.54	74.21 _{−7.86}
Falcon-7B-Instruct	53.42	50.00 _{−6.4}	32.81	23.53 _{−28.28}
Falcon-40B-Instruct	69.66	62.82 _{−9.82}	38.69	36.88 _{−4.68}
Falcon3-7B-Instruct	88.03	79.49 _{−9.7}	59.28	55.43 _{−6.49}
Falcon3-10B-Instruct	90.6	82.91 _{−8.49}	64.48	59.73 _{−7.37}
deepseek-llm-7b-chat	70.51	64.1 _{−9.09}	42.08	37.33 _{−11.29}

Table 2: Performance (%) of Flan-T5 and SOTA LLMs on NAT_V1_CHALLENGE and NAT_V2_CHALLENGE. Values in smaller font are changes (%) relative to the original performance of the model.

5.5 Do Our Findings Generalise to Other Challenging MRC Datasets?

The two SQUAD datasets investigated previously are relatively simple, as they lack challenging features (Schlegel et al., 2020), leading to super-human performance of MRC models (Lan et al., 2020). To generalise our findings to more challenging MRC benchmarks, we apply the natural perturbation methodology (Section 3) to the development set of four more datasets and assess the performance changes of multiple LLMs, as shown in Table 3. For DROP (Dua et al., 2019), we first use the GPT-4o mini to infer the likely Wikipedia article title from which each passage is retrieved⁷ and extract the revision histories for those articles. For HOTPOTQA (Yang et al., 2018), we only perturb the paragraphs containing the supporting facts, while the distracting passages remain unchanged. Furthermore, using the validity verification method described in Section 5.3, we manually verify the validity percentage of all adversarial examples in DROP and TyDi QA, as well as 50 randomly selected adversarial examples from BOOLQ and HOTPOTQA, as reported in Table 3.

Overall, *when natural perturbations are applied to other more challenging benchmarks, SOTA LLMs also exhibit a lack of robustness*, with the largest 9.98 performance decrease observed

⁷This is because the raw Wikipedia title information cannot be found in the original development set of DROP. We use the prompt: “Given a reading paragraph, return the Wikipedia page title from which it is likely retrieved.”

LLM	IM Relative Change (%)			
	BOOLQ	DROP	HOTPOTQA	TYDI QA
<i>adversarial validity (%) (Cohen's κ)</i>	72 (0.54)	85.7 (0.46)	88 (0.6)	87.5 (0.52)
Gemma 2-2b-IT	-3.91	-2.22	—	-1.61
Gemma 2-9b-IT	-3.92	-1.69	-2.21	-1.51
Llama-3.1-8B-Instruct	-3.05	-7.13	-0.91	3.17
Llama-3.2-1B-Instruct	-3.81	-9.98	-1.73	-9.1
Llama-3.2-3B-Instruct	-3.74	—	-2.05	—
Mistral-7B-Instruct-v0.2	-5.25	-1.85	-1.16	—
OLMo-7B-0724-Instruct	-4.49	-7.9	-2.36	2.94
Qwen2.5-7B-Instruct	-4.24	-3.85	-1.18	-2.78
Qwen2.5-14B-Instruct	-3.22	—	-1.84	1.38
Falcon3-7B-Instruct	-5.1	2.04	-0.06	-8.82
Falcon3-10B-Instruct	-3.58	1.8	-1.81	-4.22

Table 3: IM changes (%) of SOTA LLMs on naturally perturbed test set of other more challenging MRC datasets.

for Llama-3.2-1B-Instruct on DROP. This further demonstrates the broad and severe impact of natural perturbations on diverse MRC tasks, especially in light of on three out of four benchmarks, our human annotators are still able to correctly answer over 85% of the adversarial examples. BOOLQ exhibits a lower adversarial validity rate (72%). However, in most cases, this is not due to the perturbation degrading the passage, but rather the poor quality of the original instance (Tedeschi et al., 2023), i.e., even with the corresponding original passage, humans are unable to assign the correct label. For instance, from the annotators’ perspective, the question itself may be ambiguous, such as “do you need a visa to visit oman”, or entirely unanswerable due to missing information in the original passage. Our human annotations also observe diverse error patterns in LLMs caused by natural perturbations (e.g., *Copy Editing*, *Elaboration* and *Vandalism*), as what are presented in Figure 3. Finally, we note that for a very small proportion of adversarial examples, the limitations of the IM metric mean that the model’s answer under perturbation does not necessarily indicate an error, emphasising the need for future efforts to address this concern.

6 Dealing With Natural Perturbations

In this section, we provide an initial exploration of methods to defend against natural perturbations. Due to page constraints, we present only the approaches and general findings here; additional details and results are provided in Appendix K.

To enhance encoder-only model robustness, we conduct adversarial training by presenting them with both original training data and the generated naturally or synthetically perturbed training examples. In general, we find that retraining with nat-

ural perturbations enhances the performance on naturally perturbed test sets and improves the robustness to such perturbations as well, though this can lead to varying reductions in performance on the clean test set. Encouragingly, adversarial training with synthetically perturbed examples benefits the model’s capability to handle natural perturbations as well, a phenomenon differs from what is reported in machine translation task (Belinkov and Bisk, 2018). In some cases, the improvement even exceeds what achieved by retraining the model on natural perturbations alone.

Similarly, for the LLMs, we adopt a few-shot prompting approach by including both the original MRC instance and its naturally or synthetically perturbed counterpart as demonstrations, and assess how model performance and robustness change compared to the zero-shot setting. Although not widely observed, in certain cases, in-context demonstrations can improve an LLM’s resilience to natural perturbations, regardless of whether natural or synthetic perturbed examples are demonstrated. This phenomenon is particularly evident in models such as Llama-3.2-3B-Instruct, OLMo-7B-0724-Instruct and Falcon3-10B-Instruct. However, it can also have detrimental effects, further decreasing LLM robustness and resulting in a performance decline on both the clean and naturally perturbed test sets.

7 Conclusion

In this paper, we first study the robustness of MRC models to *natural* perturbations, which occur under real-world conditions without intentional human intervention. Using the proposed evaluation framework, we show that certain naturally perturbed examples can indeed be adversarial, i.e., lead to model failure, even when the modifications aim to improve the overall passage quality. Natural perturbations also appear to differ significantly from synthetic ones, exhibiting a wide range of rich linguistic phenomena and may be more effective in generating valid adversarial instances. Adversarial training via augmentation with either naturally or synthetically perturbed samples is generally beneficial for enhancing the model’s robustness to natural perturbations; yet, it can decrease performance on clean test set. Future work includes the exploration of alternative natural perturbation approaches and the design of more effective defensive strategies.

Limitations

We acknowledge several limitations in our work: (1) Our perturbation framework constructs natural perturbations from Wikipedia edit history and therefore only works with Wikipedia-based benchmarks. Since the phenomenon of natural perturbations is by no means limited to Wikipedia and can occur in any kind of text that evolves over time, future work should explore alternative methods to generate natural perturbations for non-Wikipedia MRC datasets. (2) As training data augmentation and in-context demonstration have a relatively limited impact, further research is needed to develop better techniques for improving the robustness of both encoder-only models and LLMs to natural perturbations, and to investigate the relationship between robustness to natural and synthetic perturbations. (3) Potential benchmark contamination may affect our findings on LLM evaluation. Investigating its extent and impact on LLM performance and robustness evaluation will be a focus of our future research efforts.

Ethical Considerations

All datasets, extracted natural perturbations, and models used in this work are publicly available, used consistently with their intended purpose and under the permitted license. A very small proportion of natural perturbations may contain offensive content, as they come from reverted Wikipedia revisions intended to damage the articles. We include these to raise awareness within the community about their potential impact on MRC models and to call for methods to improve the safety of MRC models—especially those LLMs operating under such adversarial conditions. While our ultimate goal is to enhance model robustness, the findings from this work may carry the risk of being misused by malicious attackers to refine adversarial attack strategies and craft attacks against similar systems. Before starting the annotation task, we provide all annotators with clear instructions and inform the intended use of their annotations, obtaining their explicit consent. No private or sensitive information was collected, other than their annotations.

References

Akshay Agarwal, Nalini Ratha, Mayank Vatsa, and Richa Singh. 2022. Exploring robustness connection between artificial and natural adversarial examples. In *Proceedings of the IEEE/CVF Conference*

on Computer Vision and Pattern Recognition (CVPR) Workshops, pages 179–186.

Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. *The falcon series of open language models*. Preprint, arXiv:2311.16867.

Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenhao Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. *A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity*. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. 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*.

Neeladri Bhuiya, Viktor Schlegel, and Stefan Winkler. 2024. *Seemingly plausible distractors in multi-hop reasoning: Are large language models attentive readers?* In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2514–2528, Miami, Florida, USA. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Yu Cao, Dianqi Li, Meng Fang, Tianyi Zhou, Jun Gao, Yibing Zhan, and Dacheng Tao. 2022. *TASA: Deceiving question answering models by twin answer sentences attack*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11975–11992, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, and 16 others. 2022. *Scaling instruction-finetuned language models*. Preprint, arXiv:2210.11416.

- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). 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 2924–2936, Minneapolis, Minnesota. 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 Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- DeepSeek-AI, :, Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiusi Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong Guo, and 69 others. 2024. [Deepseek llm: Scaling open-source language models with longtermism](#). *Preprint*, arXiv:2401.02954.
- 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.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). 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 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 516 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Salijona Dyrnishi, Salah Ghamizi, and Maxime Cordy. 2023. [How do humans perceive adversarial text? a reality check on the validity and naturalness of word-based adversarial attacks](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8822–8836, Toronto, Canada. Association for Computational Linguistics.
- Steffen Eger and Yannik Benz. 2020. [From hero to zéro: A benchmark of low-level adversarial attacks](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 786–803, Suzhou, China. Association for Computational Linguistics.
- Jingliang Fang, Hua Xu, Zhijing Wu, Kai Gao, Xiaoyin Che, and Haotian Hui. 2023. [Robustness-eva-mrc: Assessing and analyzing the robustness of neural models in extractive machine reading comprehension](#). *Intelligent Systems with Applications*, 20:200287.
- Wee Chung Gan and Hwee Tou Ng. 2019. [Improving the robustness of question answering systems to question paraphrasing](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6065–6075, Florence, Italy. Association for Computational Linguistics.
- Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, and 24 others. 2024. [OLMo: Accelerating the science of language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand. Association for Computational Linguistics.
- Ashim Gupta, Rishanth Rajendhran, Nathan Stringham, Vivek Srikumar, and Ana Marasovic. 2024. [Whispers of doubt amidst echoes of triumph in NLP robustness](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5533–5590, Mexico City, Mexico. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [{DEBERTA}: {DECODING}-{enhanced} {bert} {with} {disentangled} {attention}](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. 2021. [Natural adversarial examples](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15262–15271.
- Xanh Ho, Johannes Mario Meissner, Saku Sugawara, and Akiko Aizawa. 2023. [A survey on measuring and mitigating reasoning shortcuts in machine reading comprehension](#). *Preprint*, arXiv:2209.01824.
- 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.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. *Mistral 7b*. *Preprint*, arXiv:2310.06825.
- Yichen Jiang and Mohit Bansal. 2019. *Avoiding reasoning shortcuts: Adversarial evaluation, training, and model development for multi-hop QA*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2726–2736, Florence, Italy. 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.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. *Efficient memory management for large language model serving with pagedattention*. In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP ’23*, page 611–626, New York, NY, USA. Association for Computing Machinery.
- 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.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. *Albert: A lite bert for self-supervised learning of language representations*. In *International Conference on Learning Representations*.
- Thai Le, Jooyoung Lee, Kevin Yen, Yifan Hu, and Dongwon Lee. 2022. *Perturbations in the wild: Leveraging human-written text perturbations for realistic adversarial attack and defense*. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2953–2965, Dublin, Ireland. Association for Computational Linguistics.
- Mosh Levy, Shauli Ravfogel, and Yoav Goldberg. 2023. *Guiding LLM to fool itself: Automatically manipulating machine reading comprehension shortcut triggers*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8495–8505, Singapore. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. *Roberta: A robustly optimized BERT pretraining approach*. *CoRR*, abs/1907.11692.
- Edward Ma. 2019. *Nlp augmentation*. <https://github.com/makcedward/nlpaug>.
- Aur  lien Max and Guillaume Wisniewski. 2010. *Mining naturally-occurring corrections and paraphrases from Wikipedia’s revision history*. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Gemma Team Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, L. Sifre, Morgane Riviere, Mihir Kale, J Christopher Love, Pouya Dehghani Tafti, L  onard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Am  lie H  liou, Andrea Tacchetti, and 88 others. 2024. *Gemma: Open models based on gemini research and technology*. *ArXiv*, abs/2403.08295.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. *Stress test evaluation for natural language inference*. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander M  dry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024a. *Gpt-4o system card*. *Preprint*, arXiv:2410.21276.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024b. *Gpt-4 technical report*. *Preprint*, arXiv:2303.08774.
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. *PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification*. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics.
- Anibal Pedraza, Oscar Deniz, and Gloria Bueno. 2022. *Really natural adversarial examples*. *Internationa*

976	<i>tional Journal of Machine Learning and Cybernetics</i> ,	language understanding. <i>Natural Language Engi-</i>	1033
977	13(4):1065–1077.	neering, 29(1):1–31.	1034
978	Jeffrey Pennington, Richard Socher, and Christopher	Viktor Schlegel, Marco Valentino, Andre Freitas, Goran	1035
979	Manning. 2014. GloVe: Global vectors for word	Nenadic, and Riza Batista-Navarro. 2020. A frame-	1036
980	representation . In <i>Proceedings of the 2014 Confer-</i>	work for evaluation of machine reading comprehen-	1037
981	<i>ence on Empirical Methods in Natural Language Pro-</i>	sion gold standards . In <i>Proceedings of the Twelfth</i>	1038
982	<i>cessing (EMNLP)</i> , pages 1532–1543, Doha, Qatar.	<i>Language Resources and Evaluation Conference</i> ,	1039
983	Association for Computational Linguistics.	pages 5359–5369, Marseille, France. European Lan-	1040
984	Qwen, :, An Yang, Baosong Yang, Beichen Zhang,	guage Resources Association.	1041
985	Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan	Chenglei Si, Ziqing Yang, Yiming Cui, Wentao Ma,	1042
986	Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan	Ting Liu, and Shijin Wang. 2021. Benchmarking ro-	1043
987	Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin	bustness of machine reading comprehension models .	1044
988	Yang, Jiaxi Yang, Jingren Zhou, and 25 oth-	In <i>Findings of the Association for Computational Lin-</i>	1045
989	ers. 2025. Qwen2.5 technical report . <i>Preprint</i> ,	<i>guistics: ACL-IJCNLP 2021</i> , pages 634–644, Online.	1046
990	arXiv:2412.15115.	Association for Computational Linguistics.	1047
991	Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018.	Samson Tan and Shafiq Joty. 2021. Code-mixing on	1048
992	Know what you don’t know: Unanswerable ques-	sesame street: Dawn of the adversarial polyglots . In	1049
993	tions for SQuAD . In <i>Proceedings of the 56th Annual</i>	<i>Proceedings of the 2021 Conference of the North</i>	1050
994	<i>Meeting of the Association for Computational Lin-</i>	<i>American Chapter of the Association for Computa-</i>	1051
995	<i>guistics (Volume 2: Short Papers)</i> , pages 784–789,	<i>tional Linguistics: Human Language Technologies</i> ,	1052
996	Melbourne, Australia. Association for Computational	pages 3596–3616, Online. Association for Computa-	1053
997	Linguistics.	tional Linguistics.	1054
998	Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and	Samson Tan, Shafiq Joty, Min-Yen Kan, and Richard	1055
999	Percy Liang. 2016. SQuAD: 100,000+ questions for	Socher. 2020. It’s morphin’ time! Combating lin-	1056
1000	machine comprehension of text . In <i>Proceedings of</i>	guistic discrimination with inflectional perturbations .	1057
1001	<i>the 2016 Conference on Empirical Methods in Natu-</i>	In <i>Proceedings of the 58th Annual Meeting of the As-</i>	1058
1002	<i>ral Language Processing</i> , pages 2383–2392, Austin,	<i>sociation for Computational Linguistics</i> , pages 2920–	1059
1003	Texas. Association for Computational Linguistics.	2935, Online. Association for Computational Lin-	1060
1004	Marco Tulio Ribeiro, Sameer Singh, and Carlos	guistics.	1061
1005	Guestrin. 2018. Semantically equivalent adversar-	Falcon-LLM Team. 2024. The falcon 3 family of open	1062
1006	ial rules for debugging NLP models . In <i>Proceedings</i>	models .	1063
1007	<i>of the 56th Annual Meeting of the Association for</i>	Simone Tedeschi, Johan Bos, Thierry Declerck, Jan	1064
1008	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	Hajič, Daniel Hershcovich, Eduard Hovy, Alexan-	1065
1009	pages 856–865, Melbourne, Australia. Association	der Koller, Simon Krek, Steven Schockaert, Rico	1066
1010	for Computational Linguistics.	Sennrich, Ekaterina Shutova, and Roberto Navigli.	1067
1011	Gemma Team Morgane Riviere, Shreya Pathak,	2023. What’s the meaning of superhuman perfor-	1068
1012	Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupati-	mance in today’s NLU? In <i>Proceedings of the 61st</i>	1069
1013	raju, L’eonard Hussenot, Thomas Mesnard, Bobak	<i>Annual Meeting of the Association for Computational</i>	1070
1014	Shahriari, Alexandre Ram’e, Johan Ferret, Peter Liu,	<i>Linguistics (Volume 1: Long Papers)</i> , pages 12471–	1071
1015	Pouya Dehghani Tafti, Abe Friesen, Michelle Cas-	12491, Toronto, Canada. Association for Computa-	1072
1016	bon, Sabela Ramos, Ravin Kumar, Charline Le Lan,	tional Linguistics.	1073
1017	Sammy Jerome, Anton Tsitsulin, and 176 others.	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	1074
1018	2024. Gemma 2: Improving open language mod-	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	1075
1019	els at a practical size . <i>ArXiv</i> , abs/2408.00118.	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	1076
1020	Victor Sanh, Lysandre Debut, Julien Chaumond, and	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton	1077
1021	Thomas Wolf. 2019. DistilBERT, a distilled version	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	1078
1022	of BERT: smaller, faster, cheaper and lighter . In <i>5th</i>	Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 oth-	1079
1023	<i>Workshop on Energy Efficient Machine Learning and</i>	ers. 2023. Llama 2: Open foundation and fine-tuned	1080
1024	<i>Cognitive Computing @ NeurIPS 2019</i> .	chat models . <i>Preprint</i> , arXiv:2307.09288.	1081
1025	Viktor Schlegel, Goran Nenadic, and Riza Batista-	Son Quoc Tran, Phong Nguyen-Thuan Do, Uyen Le,	1082
1026	Navarro. 2021. Semantics altering modifications	and Matt Kretchmar. 2023. The impacts of unanswer-	1083
1027	for evaluating comprehension in machine reading .	able questions on the robustness of machine reading	1084
1028	<i>Proceedings of the AAAI Conference on Artificial</i>	comprehension models . In <i>Proceedings of the 17th</i>	1085
1029	<i>Intelligence</i> , 35(15):13762–13770.	<i>Conference of the European Chapter of the Asso-</i>	1086
1030	Viktor Schlegel, Goran Nenadic, and Riza Batista-	<i>ciation for Computational Linguistics</i> , pages 1543–	1087
1031	Navarro. 2023. A survey of methods for revealing	1557, Dubrovnik, Croatia. Association for Computa-	1088
1032	and overcoming weaknesses of data-driven natural	tional Linguistics.	1089

Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.

Johannes Welbl, Pasquale Minervini, Max Bartolo, Pontus Stenetorp, and Sebastian Riedel. 2020. [Under-sensitivity in neural reading comprehension](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1152–1165, Online. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Winston Wu, Dustin Arendt, and Svitlana Volkova. 2021. [Evaluating neural model robustness for machine comprehension](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2470–2481, Online. Association for Computational Linguistics.

Yulong Wu, Viktor Schlegel, and Riza Batista-Navarro. 2023. [Are machine reading comprehension systems robust to context paraphrasing?](#) In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 184–196, Nusa Dua, Bali. Association for Computational Linguistics.

Jun Yan, Yang Xiao, Sagnik Mukherjee, Bill Yuchen Lin, Robin Jia, and Xiang Ren. 2022. [On the robustness of reading comprehension models to entity renaming](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 508–520, Seattle, United States. Association for Computational Linguistics.

Diyi Yang, Aaron Halfaker, Robert Kraut, and Eduard Hovy. 2017. [Identifying semantic edit intentions from revisions in Wikipedia](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2000–2010, Copenhagen, Denmark. 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.

Torsten Zesch. 2012. [Measuring contextual fitness using error contexts extracted from the Wikipedia revision history](#). In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 529–538, Avignon, France. Association for Computational Linguistics.

A Encoder-only Model Parameters and Hyperparameters for Fine-tuning

Table 4 shows the hyperparameters used to fine-tune the pre-trained encoder-only MRC models in this work and their number of parameters contained.

Model _{Parameters(M)}	d	b	lr	ep
DistilBERT ₍₆₆₎	384	8	3e − 5	3
BERT _(110/340)	384	8	3e − 5	2
SpanBERT _(110/340)	512	4	2e − 5	4
RoBERTa _(125/355)	384	8	3e − 5	2
ALBERT _(11/17/223)	384	4	3e − 5	2
DeBERTa ₍₃₅₀₎	384	4	3e − 6	3

Table 4: Number of parameters in each type of pre-trained encoder-only MRC model and the hyperparameters used to fine-tune them. For BERT, SpanBERT, RoBERTa and ALBERT, we show the number of model parameters in the order of *base*, *large* and *xxlarge* (if applicable) version. d is the size of the token sequence fed into the model, b is the training batch size, lr is the learning rate, and ep is the number of training epochs. We used stride = 128 for documents longer than d tokens.

B Instruction Templates for Flan-T5 Evaluation

In Table 5, we present the instruction templates employed in constructing the inputs to the Flan-T5 model for the SQUAD 1.1 format and SQUAD 2.0 format test sets, respectively.

C MRC Prompts

We use the following zero-shot prompts to instruct the decoder-only models to generate responses in the task of MRC.

SQUAD 1.1 & TyDi QA: *Use the provided article delimited by triple quotes to answer question. Provide only the shortest continuous span from the context without any additional explanation.* \n\n“““{context}”””\n\nQuestion: {question}

SQUAD 1.1	
1	“Read this and answer the question\n\n{context}\n\n{question}”
2	“{context}\n\n{question}”
3	“Answer a question about this article:\n\n{context}\n\n{question}”
4	“Here is a question about this article: {context}\n\nWhat is the answer to this question: {question}”
5	“Article: {context}\n\nQuestion: {question}”
6	“Article: {context}\n\nNow answer this question: {question}”
SQUAD 2.0	
1	“Read this and answer the question. If the question is unanswerable, say \“unanswerable\”. \n\n{context}\n\n{question}”
2	“{context}\n\n{question} (If the question is unanswerable, say \“unanswerable\”)”
3	“{context}\n\nTry to answer this question if possible (otherwise reply \“unanswerable\”): {question}”
4	“{context}\n\nIf it is possible to answer this question, answer it for me (else, reply \“unanswerable\”): {question}”
5	“{context}\n\nAnswer this question, if possible (if impossible, reply \“unanswerable\”): {question}”
6	“Read this: {context}\n\nNow answer this question, if there is an answer (If it cannot be answered, return \“unanswerable\”): {question}”

Table 5: Various instruction templates for Flan-T5 model evaluation.

SQUAD 2.0: *Use the provided article delimited by triple quotes to answer question. Provide only the shortest continuous span from the context without any additional explanation. If the question is unanswerable, return “unanswerable”.*\n\n“““{context}”””\n\nQuestion: {question}

DROP & HOTPOTQA: *Use the provided article delimited by triple quotes to answer question. Provide only the answer without any additional explanation.*\n\n“““{context}”””\n\nQuestion: {question}

BOOLQ: *Use the provided article delimited by triple quotes to answer question. Return only TRUE or FALSE.*\n\n“““{context}”””\n\nQuestion:

{question}

D Indicators of Unanswerable

We manually identify a set of phrases contained in the output of LLMs that indicate the unanswerability of the question, including “*I cannot answer this/the question*”, “*unanswerable*”, “*There is no indication in the provided article*”, “*The context provided does not provide enough information*”, “*There is no reference in the given article*”, “*The answer to the question is not provided in the given article*”, “*it is not possible*”, “*question cannot be answered*” and “*context/question/article/text/article provided/passage does not*”.

E Impact of the Complete Set of Perturbed Instances on Encoder-Decoder and Decoder-Only Architectures

We supplement Table 1 in Section 5.1 with additional experiments on Flan-T5 and some SOTA LLMs such as Gemma 2 (Riviere et al., 2024) and Llama 3.2, to study the effect of all perturbed instances on these two architecture types (in addition to the encoder-only one). The results are presented in Table 6. From Table 6, we can see that similar to encoder-only models, Flan-T5 and LLMs generally exhibit varying degrees of performance degradation under natural perturbations, but also exhibit considerable robustness.

Victim	SQUAD 1.1	SQUAD 2.0
flan-t5-small	−0.69	−0.64
flan-t5-base	−0.91	−1.32
flan-t5-large	−0.77	−1.13
flan-t5-xl	−0.98	−1.37
Gemma 2-2b-IT	—	−0.76
Gemma 2-9b-IT	−0.89	−0.92
Llama-3.1-8B-Instruct	−0.38	0.39
Llama-3.2-3B-Instruct	−0.96	−0.37
Mistral-7B-Instruct-v0.2	0.39	−1.28
Falcon-7B-Instruct	−0.88	−5.38
Falcon-40B-Instruct	−0.80	—

Table 6: Performance change (%) for Flan-T5 and LLMs subjecting to natural perturbations.

We then measure the transferability of adversarial examples across all the evaluated model architectures and observe that these models share similar error patterns, with LLMs (especially Falcon) showing moderate differences. However, the lowest transferability metric is still as high as 0.86.

F Human Annotation Instructions

In Figure 4, we show the instructions given to human annotators for error analysis (Section 5.2) and adversarial validity checking (Section 5.3), respectively. All our human annotators are university students in the United Kingdom and China. Before commencing each task, we ask the annotators to annotate some examples and report the average time spent on each. As compensation, annotators receive 40 pence for each annotated example.

G Demonstration of Perturbed MRC Examples for Encoder-only Models

Figure 5 illustrates a naturally perturbed MRC instance each for categories C2W, C2C, and W2C, with the annotated perturbation type(s).

H Process of Adversarial Validity Verification

We first present two human annotators with the same collection of adversarial instances, which includes only perturbed contexts and their corresponding questions, and then ask them to answer the question based on the perturbed context. The annotators are required to select the shortest continuous span in the perturbed context that answers the question and are allowed to leave the answer blank if they are confident that the question is not answerable. Full instructions given to the annotators can be seen in Appendix F. Subsequently, for both annotators, we measure the correctness (1 or 0) of their provided answers by comparing each of them with the corresponding ground truth answers⁸. The inter-annotator agreement is then measured by computing the Cohen’s κ coefficient (Cohen, 1960). We then involve a third human annotator to annotate the adversarial examples on which the first two annotators disagree and then take the majority label as ground truth.

I Exhaustive Search Algorithm for Challenging Test Set Construction

We propose an exhaustive search algorithm that leverages the predictions of all encoder-only models to create the challenging natural perturbed test set. In detailed terms, for each matched reading passage from the prior version and its counterpart from

⁸Here, as long as one of the ground truth answers is included in the human-provided answer span, we consider the prediction to be correct.

Error Analysis

You will be presented with pairs of reading contexts and their modified versions. The task is to compare each context and its modified version, observe the changes made and classify them into one or more of the semantic edit intention categories detailed below:

- *Copy Editing*: Rephrase; improve grammar, spelling, tone, or punctuation
- *Clarification*: Specify or explain an existing fact or meaning by example or discussion without adding new information
- *Elaboration*: Extend/add new content; insert a fact or new meaningful assertion
- *Fact Update*: Update numbers, dates, scores, episodes, status, etc. based on newly available information
- *Refactoring*: Restructure the article; move and rewrite content, without changing the meaning of it
- *Simplification*: Reduce the complexity or breadth of discussion; may remove information
- *Vandalism*: Deliberately attempt to damage the article
- *Other*: None of the above

We will use your annotation to calculate the percentage of each edit category.

Adversarial Validity Checking

Please read each provided context carefully and answer a corresponding question. Select the shortest continuous span from the context as your answer. If you believe a question cannot be answered, leave the answer blank. Your answer will be compared with the ground truth answers, and the result will only be used to decide the human answerability of the question.

Figure 4: Instructions for the two distinct human annotation tasks. In the error analysis task, the eight semantic edit intentions are adopted from (Yang et al., 2017).

the current version, we determine which should be designated as the *original* and which as the *per-*

<p>Category: C2W</p> <p>Original Paragraph: <i>Jacksonville, like most large cities in the United States, suffered from negative effects of rapid urban sprawl after World War II. The construction of highways led residents to move to newer housing in the suburbs. After World War II, the government of the city of Jacksonville began to increase spending to fund new public building projects in the boom that occurred after the war. [...]</i></p> <p>Perturbed Paragraph: <i>Jacksonville, like most large cities in the United States, suffered from negative effects of rapid urban sprawl after World War V. The construction of highways led residents to move to newer housing in the suburbs. After World War II, the government of the city of Jacksonville began to increase spending to fund new public building projects in the boom that occurred after the war. [...]</i></p> <p>Question: What did Jacksonville suffer from following World War I?</p> <p>Prediction of distilbert-base and spanbert-large-cased: unanswerable→rapid urban sprawl</p> <p>Annotated Natural Perturbation Type: Vandalism</p>
<p>Category: C2C</p> <p>Original Paragraph: <i>Construction projects can suffer from preventable financial problems. Underbids happen when builders ask for too little money to complete the project. Cash flow problems exist when the present amount of funding cannot cover the current costs for labour and materials, and because they are a matter of having sufficient funds at a specific time, can arise even when the overall total is enough. Fraud is a problem in many fields, but is notoriously prevalent in the construction field. Financial planning for the project is intended to ensure that a solid plan with adequate safeguards and contingency plans are in place before the project is started and is required to ensure that the plan is properly executed over the life of the project.</i></p> <p>Perturbed Paragraph: <i>Financial planning ensures adequate safeguards and contingency plans are in place before the project is started, and ensures that the plan is properly executed over the life of the project. Construction projects can suffer from preventable financial problems. Underbids happen when builders ask for too little money to complete the project. Cash flow problems exist when the present amount of funding cannot cover the current costs for labour and materials; such problems may arise even when the overall budget is adequate, presenting a temporary issue. Fraud is also an occasional construction issue.</i></p> <p>Question: What can construction projects suffer from?</p> <p>Prediction of all encoder-only models: preventable financial problems→preventable financial problems</p> <p>Annotated Natural Perturbation Type: Copy Editing; Refactoring; Simplification</p>
<p>Category: W2C</p> <p>Original Paragraph: <i>[...] The antigens expressed by tumors have several sources; some are derived from oncogenic viruses like human papillomavirus, which causes cervical cancer, while others are the organism's own proteins that occur at low levels in normal cells but reach high levels in tumor cells. [...] A third possible source of tumor antigens are proteins normally important for regulating cell growth and survival, that commonly mutate into cancer inducing molecules called oncogenes.</i></p> <p>Perturbed Paragraph: <i>[...] The antigens expressed by tumors have several sources; some are derived from oncogenic viruses like human papillomavirus, which causes cancer of the cervix, vulva, vagina, penis, anus, mouth, and throat, while others are the organism's own proteins that occur at low levels in normal cells but reach high levels in tumor cells. [...] A third possible source of tumor antigens are proteins normally important for regulating cell growth and survival, that commonly mutate into cancer inducing molecules called oncogenes.</i></p> <p>Question: What is a fourth possible source for tumor antigens?</p> <p>Prediction of bert-base-uncased: proteins normally important for regulating cell growth and survival→unanswerable</p> <p>Annotated Natural Perturbation Type: Elaboration</p>

Figure 5: Natural perturbed MRC example in C2W, C2C and W2C categories.

turbed based on which scenario can yield the questions on which the maximum sum of the number of encoder-only models demonstrates the lack of robustness phenomenon⁹. To be specific:

Given a matched reading passage (P) from the prior version, its counterpart (P') from the current version, and the associated questions:

First Scenario: We treat (P) as the original passage and (P') as the perturbed one. We then evaluate, for each associated question, how many encoder-only models demonstrate the lack of robustness phenomenon, i.e., succeed on (P) but fail on (P'). We finally obtain the total number of models that demonstrate the lack of robustness phenomenon across all questions, denoted as (N). Questions on which none of the models demonstrate the lack of robustness phenomenon are removed, leaving (Q) questions.

Second Scenario: We treat (P') as the original passage and (P) as the perturbed one. We then repeat the same evaluation process as described in the first scenario and obtain the total number of models demonstrating the lack of robustness phenomenon across all questions, denoted as (N'). Questions on which none of the models demonstrate the lack of robustness phenomenon are removed as well, leaving (Q') questions.

If ($N > N'$), we consider (P) as the original passage and (P') as the perturbed version.

If ($N < N'$), we consider (P') as the original and (P) as the perturbed.

If ($N = N'$), we compare (Q) and (Q'):

- If ($Q > Q'$), we consider (P) as the original passage and (P') as the perturbed version.
- If ($Q < Q'$), we consider (P') as the original and (P) as the perturbed.
- If ($Q = Q'$), the order does not matter, and we randomly decide which one should be the original and which should be the perturbed.

We finally process the identified original and perturbed passage pairs to ensure that the original passages are within the original SQUAD 1.1 development set. For those original passages with multiple occurrences, we select the one with the maximum number of questions reserved.

⁹We define A model as lacking robustness to the perturbation if it achieves 1 EM on the original question but attains less than 0.4 F1 on the perturbed one (for answerable questions).

J Natural Adversarial Samples for LLMs

We demonstrate two naturally perturbed reading comprehension examples that pose challenges for LLMs in Figure 6.

K Dealing With Natural Perturbations

In this section, we provide an initial exploration of methods to defend against natural perturbations. To enhance encoder-only model robustness, we first conduct adversarial training by identifying six encoder-only model architectures that already exhibit the highest robustness to natural perturbations in their respective categories (except albert-xxlarge-v2 on NAT_V2_CHALLENGE), and presenting them with both original training data and the generated naturally perturbed training examples. We extract the entire Wikipedia revision histories for the 392 articles in the original SQUAD training set, and then obtain 5,262 (with 22,033 questions) and 5,311 (with 32,993 questions) perturbed contexts to augment the original SQUAD 1.1 and SQUAD 2.0 training set, respectively, using the methodology described in Section 3. Table 7 compares the performance of these models on NAT_V1_CHALLENGE and NAT_V2_CHALLENGE, before and after retraining.

Model	Performance (EM/F1)			
	original vs. perturbed			
	NAT_V1_CHALLENGE		NAT_V2_CHALLENGE	
distilbert-base	64.53/70.45	41.03/47.6-32.43	56.56/59.08	41.18/43.3-26.71
	57.26/63.44	43.59/51.87-18.24	53.17/55.4	43.89/45.51-17.85
bert-large-cased	79.06/83.66	63.68/70.23-16.05	66.29/68.35	53.17/55.04-19.47
	74.79/80.14	59.83/67.5-15.77	67.87/69.31	58.37/59.53-14.11
spanbert-large-cased	84.19/88.2	67.95/74.77-15.23	78.73/80.68	62.44/64.99-19.45
	82.48/86.6	69.66/76.05-12.18	78.28/80.0	65.61/67.12-16.1
roberta-large	86.75/90.21	73.93/79.47-11.91	82.13/84.27	66.29/68.52-18.69
	83.33/87.15	70.94/76.53-12.19	81.22/82.67	70.59/71.84-13.1
albert-xxlarge-v2	84.62/89.64	73.93/78.77-12.13	84.62/86.07	68.1/69.61-19.12
	86.32/90.93	75.64/81.07-10.84	82.58/84.08	70.59/72.78-13.44
deberta-large	88.46/92.5	73.5/78.48-15.16	85.07/86.65	71.49/73.0-15.75
	88.03/91.84	76.92/81.53-11.23	83.03/85.1	72.62/74.48-12.48

Table 7: Comparison of the performance of several encoder-only MRC systems on NAT_V1_CHALLENGE and NAT_V2_CHALLENGE, before and after retraining. The results shown in the shaded areas represent the performance of the model retrained on the augmented training set with naturally perturbed instances.

Apart from re-training with the same type of noise, we also ask whether exposing models to synthetic perturbations can help them confront natural ones. Therefore, we incorporate thirteen synthetic perturbation techniques spanning character and word levels (see Appendix L). Afterwards, we first retrain deberta-large with perturbed training samples generated by each syn-

NAT_V1_CHALLENGE
<p>Original Paragraph: <i>In business, notable alumni include Microsoft CEO Satya Nadella, Oracle Corporation founder and the third richest man in America Larry Ellison, Goldman Sachs and MF Global CEO as well as former Governor of New Jersey Jon Corzine, McKinsey & Company founder and author of the first management accounting textbook James O. McKinsey, Arley D. Cathey, Bloomberg L.P. CEO Daniel Doctoroff, Credit Suisse CEO Brady Dougan, Morningstar, Inc. founder and CEO Joe Mansueto, Chicago Cubs owner and chairman Thomas S. Ricketts, and NBA commissioner Adam Silver.</i></p> <p>Perturbed Paragraph: <i>In business, notable alumni include Microsoft CEO Satya Nadella, Oracle Corporation founder and the third richest man in America Larry Ellison, Goldman Sachs and MF Global CEO as well as former Governor of New Jersey Jon Corzine, McKinsey & Company founder and author of the first management accounting textbook James O. McKinsey, co-founder of the Blackstone Group Peter G. Peterson, co-founder of AQR Capital Management Cliff Asness, founder of Dimensional Fund Advisors David Booth, founder of The Carlyle Group David Rubenstein, Lazard CEO Ken Jacobs, entrepreneur David O. Sacks, CEO of TPG Group and former COO of Goldman Sachs Jon Winkelreid, former COO of Goldman Sachs Andrew Alper, billionaire investor and founder of Oaktree Capital Management Howard Marks, Bloomberg L.P. CEO Daniel Doctoroff, Credit Suisse CEO Brady Dougan, Morningstar, Inc. founder and CEO Joe Mansueto, Chicago Cubs owner and chairman Thomas S. Ricketts, and NBA commissioner Adam Silver.</i></p> <p>Question: What Goldman Sachs CEO is also an alumni of the University of Chicago?</p> <p>Prediction of GPT-3.5-turbo-0125 and Llama-3-8B-Instruct: Jon Corzine→Jon Winkelreid</p> <p>Prediction of Falcon-40B-Instruct: Jon Corzine→David Rubenstein, co-founder of The Carlyle Group, is also an alumnus of the University of Chicago.</p>
NAT_V2_CHALLENGE
<p>Original Paragraph: <i>Each chapter has a number of authors who are responsible for writing and editing the material. A chapter typically has two "coordinating lead authors", ten to fifteen "lead authors", and a somewhat larger number of "contributing authors". The coordinating lead authors are responsible for assembling the contributions of the other authors, ensuring that they meet stylistic and formatting requirements, and reporting to the Working Group chairs. Lead authors are responsible for writing sections of chapters. Contributing authors prepare text, graphs or data for inclusion by the lead authors.</i></p> <p>Perturbed Paragraph: <i>Each chapter has a number of authors to write and edit the material. A typical chapter has two coordinating lead authors, ten to fifteen lead authors and a larger number of contributing authors. The coordinating lead authors assemble the contributions of the other authors. They ensure that contributions meet stylistic and formatting requirements. They report to the Working Group co-chairs. Lead authors write sections of chapters. They invite contributing authors to prepare text, graphs or data for inclusion.</i></p> <p>Question: Who has the responsibility for publishing materials?</p> <p>Prediction of Mistral-7B-Instruct-v0.2: Unanswerable. The text does not mention any responsibility related to publishing materials.→The coordinating lead authors are responsible for publishing materials in the given context.</p>

Figure 6: Natural perturbed MRC examples that confuse LLMs.

thetic perturbation method, respectively, and assess the performance changes compared to the vanilla version on both NAT_V1_CHALLENGE and NAT_V2_CHALLENGE (Figure 8 in Appendix M). As we observe that synthetic adversarial training can assist deberta-large in handling natural perturbations, we further retrain five other models in

the same manner and quantify the performance difference on NAT_V1_CHALLENGE compared to the vanilla version, as shown in Figure 7.

In general, for encoder-only MRC models, retraining with natural perturbations enhances the performance on naturally perturbed test sets and improves the robustness to such perturbations as




Figure 7: Absolute changes in original and perturbed performance (F1), as well as the robustness of five encoder-only models under natural perturbations (on NAT_V1_CHALLENGE), following retraining with each synthetic perturbation.

well, though this can lead to varying reductions in performance on the clean test set. Encouragingly, adversarial training with synthetically perturbed examples benefits the model’s capability to handle natural perturbations as well, a phenomenon differs from what is reported in machine translation task (Belinkov and Bisk, 2018). In some cases, the improvement even exceeds what achieved by retraining the model on natural perturbations alone. We also observe that the effectiveness of adversarial training varies with model size and architecture. Generally, adversarial training brings the most significant benefits for the weakest distilbert-base, with the benefits diminishing in larger and more complex model architectures.

Similarly, for the LLMs, we adopt a few-shot prompting approach by including both the original MRC instance and its naturally or synthetically perturbed counterpart as demonstrations, and assess how model performance and robustness change compared to the zero-shot setting (see Table 8 and Table 9). A total of two original-perturbed instance pairs are used, with the original samples taken from the SQUAD training set. Although not widely observed, in certain cases, in-context demonstrations can improve an LLM’s resilience to natural perturbations, regardless of whether natural or synthetic perturbed examples are demonstrated. This phenomenon is particularly evident in models such as Llama-3.2-3B-Instruct, OLMo-7B-0724-Instruct and Falcon3-10B-Instruct. However, it can also have detrimental effects, further decreasing LLM robustness and resulting in a performance decline on both the clean and naturally perturbed test sets.

L Synthetic Perturbation Methods

Table 10 presents the synthetic perturbation methods used in this study.

LLM	NAT_V1_CHALLENGE			NAT_V2_CHALLENGE		
	Orig./Pert.	IM Drop	zero-shot	Orig./Pert.	IM Drop	zero-shot
Gemma 2-2b-IT	80.77/72.22 ↓	-10.59	-8.49	45.93/42.53 ↓	-7.4	-11.11
Gemma 2-9b-IT	↑91.88/79.91 ↓	-13.03	-8.62	62.9/57.24 ↓	-9.0	-7.67
Llama-3.1-8B-Instruct	↓85.04/74.79 ↑	-12.05	-15.61	45.25/41.86 ↓	-7.49	-9.59
Llama-3.2-3B-Instruct	70.09/62.82 ↓	-10.37	-12.11	43.67/43.44 ↓	-0.53	-10.29
Mistral-7B-Instruct-v0.2	↓83.33/77.78 ↑	-6.66	-13.2	50.45/46.61 ↓	-7.61	-6.58
OLMo-7B-0724-Instruct	76.5/76.07 ↓	-0.56	-8.05	53.62/51.58 ↑	-3.8	-4.4
Qwen2.5-3B-Instruct	61.11/47.44 ↓	-22.37	-13.04	67.42/61.54 ↑	-8.72	-11.81
Qwen2.5-7B-Instruct	86.32/73.08 ↓	-15.34	-7.76	76.24/69.46 ↑	-8.89	-12.21
Qwen2.5-14B-Instruct	85.47/72.65 ↓	-15.0	-11.58	80.09/73.76 ↓	-7.9	-7.86
Falcon3-7B-Instruct	86.32/75.21 ↓	-12.87	-9.7	↑60.86/54.98 ↓	-9.66	-6.49
Falcon3-10B-Instruct	85.9/79.06 ↓	-7.96	-8.49	61.54/59.5 ↓	-3.31	-7.37
deepseek-11m-7b-chat	↑70.94/55.98 ↓	-21.09	-9.09	59.73/51.58 ↑	-13.64	-11.29

Table 8: Performance and IM drop of LLMs in the few-shot setting with both original and *naturally* perturbed MRC instances demonstrated. *zero-shot* represents the IM drop in the zero-shot setting, adopted from Table 2. Results that evidence robustness improvement in the few-shot setting are underlined.

LLM	NAT_V1_CHALLENGE			NAT_V2_CHALLENGE		
	Orig./Pert.	IM Drop	zero-shot	Orig./Pert.	IM Drop	zero-shot
Gemma 2-2b-IT	82.48/72.22 ↓	-12.44	-8.49	47.29/42.53 ↓	-10.07	-11.11
Gemma 2-9b-IT	↑91.45/81.62 ↓	-10.75	-8.62	61.54/56.33 ↓	-8.47	-7.67
Llama-3.1-8B-Instruct	83.33/72.65 ↓	-12.82	-15.61	49.32/44.57 ↓	-9.63	-9.59
Llama-3.2-3B-Instruct	66.24/62.82 ↓	-5.16	-12.11	42.53/41.86 ↓	-1.58	-10.29
Mistral-7B-Instruct-v0.2	↓81.62/76.92 ↑	-5.76	-13.2	50.68/46.61 ↓	-8.03	-6.58
OLMo-7B-0724-Instruct	73.84/75.53 ↓	2.29	-8.05	54.98/53.39 ↑	-2.89	-4.4
Qwen2.5-3B-Instruct	59.4/51.71 ↓	-12.95	-13.04	66.52/61.09 ↑	-8.16	-11.81
Qwen2.5-7B-Instruct	85.04/73.5 ↓	-13.57	-7.76	76.47/69.0 ↓	-9.77	-12.21
Qwen2.5-14B-Instruct	84.19/73.5 ↓	-12.7	-11.58	↑81.97/74.21	-9.39	-7.86
Falcon3-7B-Instruct	85.47/74.36 ↓	-13.0	-9.7	61.09/56.56 ↓	-7.42	-6.49
Falcon3-10B-Instruct	85.9/79.91 ↓	-6.97	-8.49	61.31/57.92 ↓	-5.53	-7.37
deepseek-11m-7b-chat	↑72.22/56.41 ↓	-21.89	-9.09	61.09/53.62 ↑	-12.23	-11.29

Table 9: Performance and IM drop of LLMs in the few-shot setting with both original and *synthetically* perturbed MRC instances demonstrated. *zero-shot* represents the IM drop in the zero-shot setting, adopted from Table 2. Results that evidence robustness improvement in the few-shot setting are underlined.

We employ methods including WSplit, WSynSub and WInsert (WE) to each sentence in the original reading passage, and then recombine the modified sentences to generate the perturbed version. Conversely, other perturbation approaches are directly executed on the entire paragraph, as implementing them at the sentence-level might result in perturbed text that is even difficult for humans to read and comprehend (Si et al., 2021). The implementation of all character-level and word-level methods is carried out using the NLPAug library (Ma, 2019). Moreover, we set the perturbation rate to 30%, in line with the default settings within the NLPAug library.

M Impact of Synthetic Adversarial Training

Figure 8 describes the impact of synthetic adversarial training (for deberta-large) on handling natural and synthetic perturbations.

Method	Description
<i>character-level</i>	
CharOCR	Replace characters with Optical Character Recognition (OCR) errors.
CharInsert	Inject new characters randomly.
CharSubstitute	Substitute original characters randomly.
CharSwapMid	Swap adjacent characters within words randomly, excluding the first and last character.
CharSwapRand	Swap characters randomly without constraint.
<i>word-level</i>	
WInsert (CWE)	Insert new words to random position according to contextual word embeddings calculation from RoBERTa-base (Liu et al., 2019).
WSubstitute (CWE)	Substitute words according to contextual word embeddings calculation from RoBERTa-base (Liu et al., 2019).
WSplit	Split words to two tokens randomly.
WSwap	Swap adjacent words randomly.
WDelete	Delete words randomly.
WCrop	Remove a set of continuous word randomly.
Word Synonym Substitution (WSynSub)	Substitute words with synonyms from large size English PPDB (Pavlick et al., 2015).
WInsert (WE)	Insert new words to random position according to GloVe (Pennington et al., 2014) word embeddings calculation (we use <i>glove.6B.300d.txt</i>).

Table 10: Various synthetic perturbation approaches.

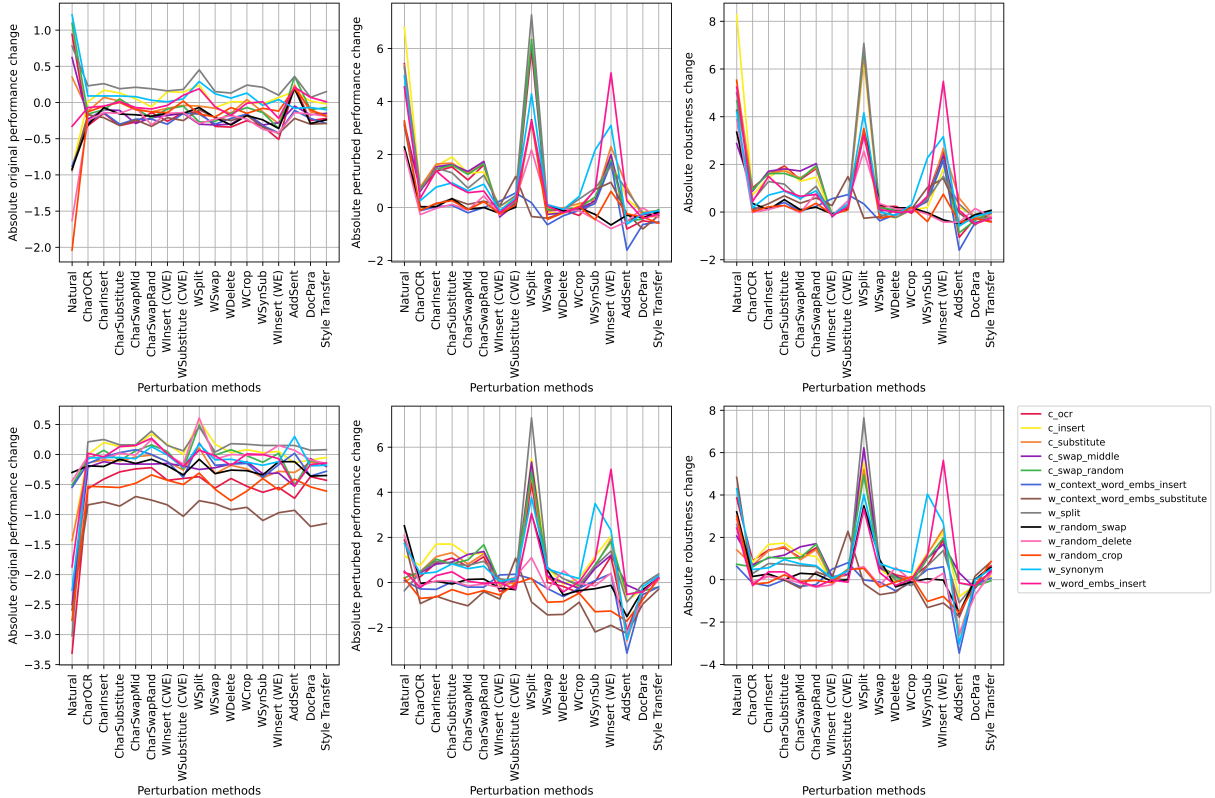


Figure 8: Absolute changes in original and perturbed performance (F1), as well as the robustness of deberta-large under natural and various synthetic noises, following retraining with each synthetic perturbation. The upper row and the bottom row illustrate the results on the SQUAD 1.1 and SQUAD 2.0 format test sets, respectively.