

Evaluating the Adversarial Robustness of Retrieval-Based In-Context Learning for Large Language Models

Simon Yu^{*†}, Jie He^{*}, Pasquale Minervini & Jeff Z. Pan[†]

School of Informatics

University of Edinburgh

{s1967531, j.he, j.z.pan}@ed.ac.uk, pminervi@exseed.ed.ac.uk

Abstract

With the emergence of large language models, such as LLaMA and GPT, In-Context Learning (ICL) gained significant attention due to its effectiveness and efficiency. However, ICL is very sensitive to the choice, order, and verbaliser used to encode the demonstrations in the prompt. *Retrieval-Augmented ICL* methods try to address this problem by leveraging retrievers to extract semantically related examples as demonstrations. While this approach yields more accurate results, its robustness against various types of adversarial attacks, including perturbations on test samples, demonstrations, and retrieved data, remains under-explored. Our study reveals that retrieval-augmented models can enhance robustness against test sample attacks, outperforming vanilla ICL with a 4.87% reduction in Attack Success Rate (ASR); however, they exhibit overconfidence in the demonstrations, leading to a 2% increase in ASR for demonstration attacks. Adversarial training can help improve the robustness of ICL methods to adversarial attacks; however, such a training scheme can be too costly in the context of LLMs. As an alternative, we introduce an effective training-free adversarial defence method, *DARD*, which enriches the example pool with those attacked samples. We show that *DARD* yields improvements in performance and robustness, achieving a 15% reduction in ASR over the baselines. Code and data are released to encourage further research¹.

1 Introduction

Large Language models (LLMs) are revolutionising the field of NLP (Brown et al., 2020; Wei et al., 2022; Touvron et al., 2023). However, this also raises concerns about their robustness (Xu et al., 2023c). Although significant efforts are put into aligning LLMs for safety purposes, recent works still show LLMs can be vulnerable to adversarial inputs and jailbreaking (Zou et al., 2023; Xu et al., 2023b; Zeng et al., 2024; Longpre et al., 2024). As a result, it is essential to understand LLM security properties to guide the direction of LLMs that are secure and robust. In this paper, we focus on examining a common LLM inference method, namely In-Context Learning (ICL) and its variants, on how they might be susceptible to perturbations and their robustness against *adversarial attacks*.

ICL is a common few-shot training method for prompting LLM for downstream tasks (Brown et al., 2020). Its variants include retrieval-based ICL models, which utilise retrievers to identify similar examples as demonstrations (Wang et al., 2023c; Li et al., 2023); or *k*NN-ICL, which compute label demonstrations using a nearest neighbours algorithm; have been adopted and shown promising performance improvements when compared to vanilla ICL (Shi et al., 2022; Xu et al., 2023a). However, the robustness of these methods against adversarial attacks received less attention and remains unclear. Prior work—which we review in Section 2—shows that LLMs can still be vulnerable to different types of input

^{*}Equal Contributions

[†]Corresponding Author

¹<https://github.com/simonucl/adv-retrieval-icl>

perturbations, such as task instructions (Zhu et al., 2023; Sun et al., 2023) or inputs (Wang et al., 2023a); however, the robustness of ICL-based methods to such perturbations is still under-explored.

We address (cf. Section 4) the following question: *How sensitive are vanilla and retrieval-based ICL methods to perturbations to their test samples or demonstration instances*, which is particularly important given that the primary motivation of ICL is to facilitate few-shot learning without training; should any methods exhibit excessive sensitivity, this would limit their applicability. To investigate this problem, we employ commonly used perturbation attacks—namely *character-level*, *word-level*, and *sentence-level* attacks. Furthermore, we propose new attack methods that target the demonstration instances and datastore examples. Our findings indicate that retrieval-based ICL methods show greater robustness against test-sample attacks; however, they can be more vulnerable when their demonstrations are adversarially perturbed. Detailed ablation studies further reveal that this vulnerability to adversarial attacks persists in newer models, such as Mistral (Jiang et al., 2023) and Gemma (Gemma Team et al., 2024). Many of these attacks can be transferable, both from the same family (LLaMA-2-7B → LLaMA-2-13B, 70B) or different family (LLaMA-2-7B → Mistral-7B). We showed Mixture of Expert (MoE) models still posed a similar adversarial threat to its dense version. This encourages further investigation on enhancing model robustness.

We further explore the advantages of utilising retrieval-based models to improve robustness (cf. Section 5), considering that existing defence strategies often rely on an adversarial training process (Li et al., 2021; Si et al., 2021). However, this process is memory intensive in the context of LLMs. Instead, we show that retrieval-based ICL methods, combined with adversarially augmented examples, yield more robust results than no augmentation methods while maintaining their predictive accuracy on in-distribution evaluation sets. Therefore, we propose **DARD**, a training-free adversarial defence method which leverages adversarially perturbed samples to augment the retrieval pools, potentially improving the models’ robustness. Figure 1 presents an overview of the paper.

Our main contributions are summarized as follows.

- We first perform a comprehensive and in-depth evaluation on variants of ICL: vanilla, k NN-ICL, and retrieval-based ICL methods against Test-Sample, Demonstration and Datastore Adversarial Attacks.
- We show that Retrieval-based models can bring positive results on robustness when test samples are under adversarial attack, but it further reduces the robustness while demonstrations are perturbed. We show these attacks persist on larger models and are transferable.
- We propose **DARD**, a novel training-free adversarial defence method which leverages adversarially perturbed examples to improve the robustness.

2 Related Work

Adversarial attack and defences Adversarial attack has always been a topic in deep neural networks, where previous works showed neural networks can be vulnerable to adversarial examples that are crafted by adding imperceptible but adversarial noise on natural examples (Szegedy et al., 2013; Evtimov et al., 2017). In the real world, adversarial robustness is important for the application of a method (Tramèr et al., 2017).

Several works studied this problem and the main finding is that LMs struggle to answer correctly and factually under attacks and propose their defence via label smoothing or prompt tuning (Yang et al., 2022; Raman et al., 2023). It is important to note, however, that these experiments were mainly conducted using relatively small LMs such as BERT, RoBERTa. Recently, several works demonstrated LLM can still be vulnerable to these type of attacks, by showing the vulnerability of attacking the task instructions (PromptBench (Zhu et al., 2023)), or performing adversarial attacks on zero-shot settings and investigating the transferability of attacks from open-sourced models (i.e. Alpaca (Taori et al., 2023), Vicuna (Chiang et al., 2023)) to close-sourced (i.e. GPT-3.5 and GPT-4) (DecodingTrust (Wang et al., 2023a)). However, previous study (Gu et al., 2022) has shown that the LLMs are less sensitive to prompt / instruction variation when few-shot examples are provided

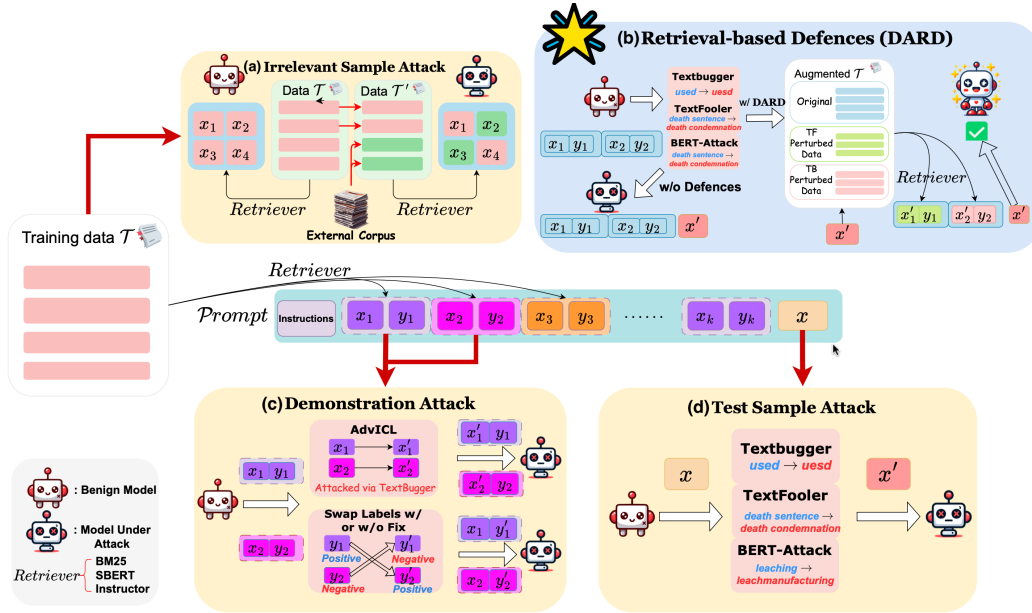


Figure 1: Overview of the paper. We visualize our seven adversarial attacks in (a), (c) and (d) (only 3 shots are used in the plot for display purposes). And our adversarial defence method, *DARD*, is showcased in the top right corner (Plot (b)).

in context (a.k.a ICL). However, there is still a lack of systematic analysis on this case which is the main contribution of this paper. Wang et al. (2023b) shows that attacking demonstrations can still be an effective way for adversarial attacks where LLMs are still vulnerable. Therefore, we aim to extend and evaluate the adversarial attack and defence performance on retrieval-based models.

ICL Sensitivity The existing literature on ICL analysis can be broadly divided into two streams, each focusing on different aspects. The first stream explores the influencing factors of ICL based on input perturbation, such as the order (Min et al., 2022), the formatting (Yoo et al., 2022; Lu et al., 2022) and the selection of the demonstration (Liu et al., 2022). Designing proper demonstration construction strategies (Ye et al., 2023; Li et al., 2023) could bring clear boosts to the ICL performance, while Chen et al. (2023) analyses the sensitivity of ICL via perturb on instructions and example ordering, and our work is a complement of their work on sensitivity. This paper provides a different view of when different ICL methods fail and when retrieval-based ICL can be a better approach for robustness.

Retrieval-based LLM Retrieval-based LLM has shown great success in lifting LLM performance and faithfulness and reducing hallucination (Pan et al., 2023), where the idea is to gather related articles or demonstrations from a collection of data (Huang et al., 2023). Specifically, we are interested in whether they can be used to improve the robustness against perturbations (Lin & Lee, 2024). Meanwhile, many papers also focus on analysing the robustness of RAG methods on irrelevant contexts (Shi et al., 2023; Asai et al., 2024), in which we adopt similar methods in our context for analysing how retrieval-based ICL being sensitive to irrelevant contexts.

3 Problem Formulation

3.1 Background

In this work, we consider several types of ICL methods, namely **Vanilla ICL**, and two types of retrieval-based ICL: *k*NN-ICL (Xu et al., 2023a) and **Retrieval-based ICL** (Wang

et al., 2023c; Li et al., 2023). Retrieval-based ICL methods demonstrated promising performance improvements over vanilla ICL ones; however, their robustness against adversarial perturbations remains unexplored.

Vanilla ICL is the most commonly used few-shot learning method in LLM. Given a training data set $\mathcal{T} = \{(x_i, y_i)\}$, where $x_i \in X$ denote the instances and $y_i \in Y$ represent their labels, the concatenation of the k -shot input-label pairs is used to generate the prompt P with respect to a test instance x_{test} :

$$P = \mathcal{I} \oplus \pi(x_1, y_1) \oplus \pi(x_2, y_2) \oplus \dots \oplus \pi(x_k, y_k) \oplus \pi(x_{\text{test}}, *) \quad (1)$$

where \mathcal{I} denotes the task instructions, π denotes the templates and verbaliser to turn inputs into prompts (details for each dataset are in Appendix A.3), \oplus is a concatenation operator, and $|Y|$ is the number of distinct labels. The label prediction y_{test} for the test instance x_{test} is then computed as:

$$y_{\text{test}} = \arg \max_{y \in Y} p(y|P; \theta_{\text{LLM}}) \quad (2)$$

k NN-Prompting was first introduced by Shi et al. (2022) to perform zero-shot inference and extended in later work by Xu et al. (2023a) to support few-shot inferences, simulating the ICL. Throughout the paper, we follow the implementation of Xu et al. (2023a) and call it k NN-ICL. The idea is to map each training instance into a latent representation using a LLM and cache the last token distribution into a datastore. At inference time, the prediction is based on the interpolation between the LLM prediction and the k -nearest neighbours label distribution in the datastore. Specifically, for each training instance x_i , Xu et al. (2023a) concatenate it into prompt P_i together with k in-context examples, where $k = |Y|$:

$$P_i = \pi(x_1, y_1) \oplus \dots \oplus \pi(x_{|Y|}, y_{|Y|}) \oplus \pi(x_i, *). \quad (3)$$

By querying the LLM using P_i , we obtain a distribution over tokens $p(v|P_i, \theta)$. Rather than mapping it back to label space Y , the distribution is cached as the key representation for the instance x_i :

$$\mathbf{k}_i = p(v|P_i; \theta_{\text{LLM}}), \quad (4)$$

which is mapped to the label y_i . The entire datastore thus consists of $\{(\mathbf{k}_i, y_i)\}$ pairs, where $K = \{\mathbf{k}_i\}_i$ denotes the set of keys. At inference time, we construct the same prompt and obtain the distribution $P_{\text{test}} = p(v|P_{\text{test}}; \theta_{\text{LLM}})$. We then match the distribution against cached K in the datastore, where standard KL -divergence is used to measure the distance:

$$D_{\text{KL}}(P_{\text{test}}||k_i) = \sum_v p(v|P_{\text{test}}; \theta_{\text{LLM}}) \log \frac{p(v|P_{\text{test}}; \theta_{\text{LLM}})}{p(v|P_i; \theta_{\text{LLM}})} \quad (5)$$

The predictions are then calculated by interpolating the LLM output and its m nearest neighbours distribution, where $\text{NN}_m(*, K)$ denotes the set of m nearest neighbours in K :

$$\hat{y}_{\text{pred}} = \arg \max_{y \in Y} \left[(1 - \alpha) \cdot p(y|P_{\text{test}}; \theta_{\text{LLM}}) + \alpha \cdot p \sum_{i \in \text{NN}_m(P_{\text{test}}, K)} \mathbb{I}(y_i^q = y) \right]. \quad (6)$$

Following Xu et al. (2023a), in this work, we use $\alpha = 0.2$ and $m = k/2$.

R-ICL Retrieval-augmented ICL (Lin et al., 2022) was proposed to mitigate the sensitivity of ICL methods to the choice and order of the in-context examples: rather than randomly sampling (x_i, y_i) , R-ICL use a retriever to identify the most similar examples to the test instance x_{test} from the training set \mathcal{T} . In this paper, we follow Li et al. (2023) and consider the following retrievers: **BM25** (Robertson & Zaragoza, 2009), **SBERT** (Reimers & Gurevych, 2019) and **Instructor** (Su et al., 2022); more details are available in Appendix A.1.

3.2 Attack Methods

This paper aims to study how perturbations can affect the performance of different ICL methods through adversarial attacks. We classify our 7 attacks into 3 categories: **Test Sample Attack**, **Demonstration Attack**, and **DataStore Attack**.

Test Sample Attacks We focus on the three different types of attack: typo-based perturbations — TextBugger (TB) (Li et al., 2018); embedding-similarity-based perturbations — TextFooler (TF) (Jin et al., 2019); and context-aware perturbations — BERT-Attack (BA) (Li et al., 2020).

Demonstration Attacks For demonstration attacks, we refer to adversarial attacks where the training demonstrations are adversarially perturbed. We consider this type of attack to investigate the sensitivity of LLMs to adversarial perturbations to the training demonstrations and their influence on the generalisation properties of the models.

Dataset Attacks One of our objectives is to analyse the sensitivity of retrieval-based ICL methods to adversarial perturbations. Therefore, we consider Irrelevant Context Attacks (Irr.), which work by contaminating the demonstration pools with irrelevant contexts. Our Irrelevant Context Attacks extended Min et al. (2022) — we replace 50% of the original demonstrations in the training data with Out-of-Distribution (OoD) examples.

Further details and analyses about attack methods are available in Appendix A.4.

4 Experiments

4.1 Settings

Datasets To assess the robustness, we selected six sentiment and multiple-choice text classification datasets: SST2 (Socher et al., 2013), RTE (Dagan et al., 2005), CR (Hu & Liu, 2004), MR (Pang & Lee, 2005), MNLI-matched (Williams et al., 2018), and TREC (Voorhees & Tice, 2000). We provide further details, including prompt templates and dataset statistics, in Appendix A.3. Unless explicitly mentioned, LLaMA-2-7B was mainly used as the base model for our experiments (Touvron et al., 2023). For experiments involving randomness (i.e., Vanilla ICL and k NN-ICL), we conducted trials with three different random seeds to ensure the reliability of our findings.

Evaluation To evaluate and compare the vulnerabilities across different types of attacks, in line with Wang et al. (2023a); Zhu et al. (2023), we assess the model’s performance in a benign (attack-free) context (**Clean**) and calculate the Attack Success Rate (**ASR**). The ASR is defined as $\frac{\text{Clean Accuracy} - \text{Attack Accuracy}}{\text{Clean Accuracy}}$, where Attack Accuracy denotes the model’s accuracy under attack.

4.2 Results

The 8-shot attack results are presented in Table 1 with 8-shot attacks, and detailed results are available in Appendix B.1.

Overview: Retrieval-based models exhibit higher Clean accuracy: 3.27% higher than vanilla ICL with the best performed SBERT retriever; similar finding with Rubin et al. (2021) that related demonstrations enhance the performance. For adversarial attacks, our proposed adversarial attack methods, Swap-Labels and Swap-Labels (Fix), are effective techniques for manipulating LLM predictions. On the other hand, irrelevant context attacks have a negligible effect on performance. We summarise two major findings about the adversarial robustness across different ICL methods.

Finding 1: Retrieval-based models are more robust against test-sample attacks; however, it could fall short under demonstration attack.

As shown in Table 1, all three R-ICL methods outperform vanilla ICL and k NN-ICL among all Text Sample Attacks (TB, TF and BA), with a drop in ASR by 4.87% and 2.47%, respectively.² In Figure 2, we observe an increased gap in the ASR between ICL and R-ICL

²Here we compare with the best performed $R_{\text{Instructor}}$ -ICL with ICL and k NN-ICL.

Tasks	Method	Clean [†]	TB [‡]	TF [‡]	BA [‡]	AdvICL [‡]	S-L [‡]	S-L (Fix) [‡]	Irr. [‡]	Avg [‡]
SST-2	ICL	92.66	37.29	43.69	68.40	33.64	92.33	62.99	1.69	48.58
	kNN-ICL	92.01	41.01	45.45	61.66	79.05	67.28	66.00	0.0	51.49
	R _{BM25} -ICL	94.61	43.39	45.94	57.33	25.63	99.88	63.88	0.85	46.7
	R _{SBERT} -ICL	95.18	33.01	33.01	56.26	27.21	99.64	66.84	6.09	46.01
	R _{Instructor} -ICL	<u>94.84</u>	38.33	32.62	58.65	28.22	99.54	69.32	6.73	47.63
RTE	ICL	<u>73.04</u>	82.54	75.62	97.86	14.32	91.59	58.31	11.04	<u>61.61</u>
	kNN-ICL	70.52	90.61	80.37	95.43	96.13	63.66	63.15	6.83	70.88
	R _{BM25} -ICL	71.48	68.69	63.14	90.4	18.77	92.42	60.60	15.15	58.45
	R _{SBERT} -ICL	72.20	83.50	83.01	98.12	27.49	95.50	64.50	15.50	66.80
	R _{Instructor} -ICL	73.29	78.33	72.41	96.67	28.08	90.64	66.00	10.36	63.21
MR	ICL	92.67	36.41	46.72	67.52	22.86	94.49	75.43	0.15	49.08
	kNN-ICL	91.77	39.60	47.18	65.95	78.99	71.81	72.90	2.78	54.17
	R _{BM25} -ICL	92.50	37.95	45.84	65.95	25.10	99.68	63.89	1.51	<u>48.56</u>
	R _{SBERT} -ICL	92.40	38.53	46.10	64.52	28.83	99.78	60.06	0.0	48.26
	R _{Instructor} -ICL	<u>92.6</u>	37.47	47.30	66.22	26.26	98.06	65.33	0.43	48.72
CR	ICL	91.31	35.54	52.43	71.16	22.8	99.90	91.87	1.22	51.56
	kNN-ICL	91.87	19.87	39.13	58.78	74.07	67.33	57.74	6.03	44.13
	R _{BM25} -ICL	93.09	29.14	47.43	65.72	33.20	100.00	56.29	7.30	48.44
	R _{SBERT} -ICL	<u>93.62</u>	31.25	46.30	71.59	33.81	100.00	62.32	4.76	50.00
	R _{Instructor} -ICL	93.88	34.55	39.38	53.82	35.81	100.00	64.88	4.02	<u>47.49</u>
MNLI-mm	ICL	53.63	68.49	57.92	40.85	32.48	88.81	43.84	1.73	47.73
	kNN-ICL	55.89	65.61	52.10	51.30	51.30	62.39	56.00	10.14	<u>49.83</u>
	R _{BM25} -ICL	<u>57.3</u>	68.41	59.16	40.98	38.45	95.99	63.89	0.35	52.46
	R _{SBERT} -ICL	57.8	69.03	55.07	43.27	37.01	99.31	63.55	2.77	52.86
	R _{Instructor} -ICL	54.30	66.48	51.93	36.06	37.83	100.00	69.67	2.58	52.08
TREC	ICL	76.07	61.09	57.50	60.65	18.69	65.56	40.58	8.19	44.61
	kNN-ICL	77.87	59.68	54.54	50.02	83.65	58.74	52.61	10.88	52.87
	R _{BM25} -ICL	79.20	50.00	50.76	57.07	28.91	55.56	9.60	3.79	36.53
	R _{SBERT} -ICL	87.8	43.74	45.56	47.15	34.61	53.08	10.02	9.11	34.75
	R _{Instructor} -ICL	<u>86.5</u>	45.43	70.87	47.51	31.31	44.97	3.35	7.28	<u>35.82</u>
Avg	ICL	79.90	53.56	55.65	67.74	24.13	88.78	59.84	4.00	50.53
	kNN-ICL	79.98	52.73	53.13	63.86	77.20	65.2	61.40	6.11	54.18
	R _{BM25} -ICL	81.36	49.6	<u>52.04</u>	<u>62.91</u>	<u>27.68</u>	90.59	53.02	4.82	48.52
	R _{SBERT} -ICL	83.17	<u>49.84</u>	51.51	63.48	31.49	91.22	54.55	6.37	49.78
	R _{Instructor} -ICL	<u>82.57</u>	50.10	52.42	59.82	31.25	88.87	56.42	5.23	<u>49.16</u>

Table 1: Adversarial Attack results with 8-shot ICL. **Clean** refers to the Benign accuracy, while all other columns refer to the Attack Success Rate (%) under corresponding Adversarial Attacks. The last column **Avg** refers to the mean ASR across 7 attacks. The higher the ASR, the lower the robustness of models against the attack.

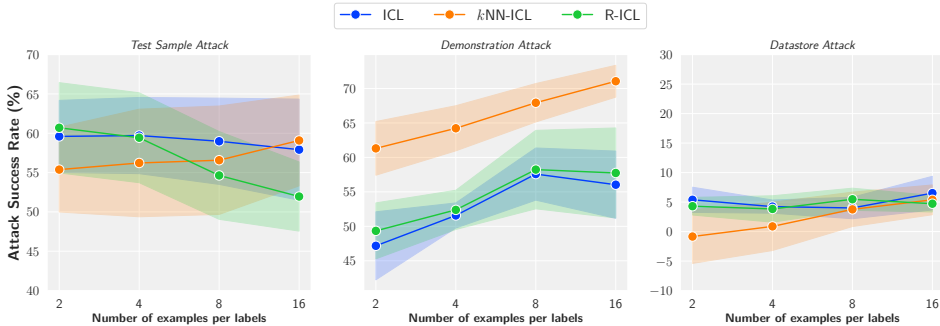


Figure 2: Attack Success Rate (ASR) across shots among different ICL methods. We aggregated the results for ICL and kNN-ICL across 3 seeds and R-ICL results across 3 retrievers.

when the number of shots goes up. Therefore, we believe in-context learning benefits from drawing semantically related demonstrations against adversarial perturbations on both performance and robustness. Recall that these textual attacks are designed to deceive LLMs by replacing words with their synonyms or introducing typos in the input sentences. This strategy can be viewed as exploiting LLM’s *parametric knowledge* (Szegedy et al., 2013; Goyal et al., 2022). In this case, R-ICL can provide supported demonstrations by retrieving related




Figure 3: Attack Success Rate (%) for adversarial attacks across various models, based on experiments conducted on the RTE dataset with 8-shot demonstrations. The results are based on R-ICL and the mean attack success rate among the three retrievers we used: BM25, SBERT, and Instructor.

examples to enhance the *contextual knowledge* to the LLM, therefore reducing the noise of the LLM prediction.

On the other side, we see an increase in ASR for Retrieval-based ICL when demonstrations are under attack, demonstrating a 1%-6% increase in ASR compared with vanilla ICL. This demonstrates the deficit of retrieval-based models that can sometimes overly rely on the retrieved demonstrations and are more vulnerable when demonstrations are perturbed.

Finding 2: *kNN-ICL is highly sensitive to adversarial perturbations on test samples and demonstration.*

In Table 1, *kNN-ICL* shows the highest ASR in *Demonstration Attack*, specifically showing the highest vulnerability to AdvICL, 40%-50% higher than ICL and R-ICL. Recall that for *kNN-ICL*, demonstrations and test samples are mapped into the hidden space to compute the label distribution with nearest neighbours. Any perturbations on the demonstration might notably affect their position in the hidden space and alter the label distribution significantly. Also, in Figure 2, *kNN-ICL* exhibits a consistently ascending trend in ASR as the number of examples increases, with 15% higher than the other methods. This further illustrates that incorporating nearest neighbours algorithms inevitably includes more noise, thereby increasing vulnerability.

Overall, our experiment results show that few-shot methods in LLM could still be vulnerable to adversarial perturbations, leading to a 50% drop in performance on average. Pairing with previous works on jailbreaking leads to the need for a continuous effort to improve the robustness of LLMs. We include quantitative results in the Appendix B.2.

4.3 Attack on Model Variants

In this section, we extend our experiments to four other models: **LLaMA-2-13B**, larger version of 7B model; **Vicuna-7B** (Chiang et al., 2023), instruction-tuned from LLaMA-2-7B; **Mistral-7B** (Jiang et al., 2023) and **Gemma-7B** (Gemma Team et al., 2024), aiming to provide a view for models from different families to prevent potential bias posed within the pretraining data for LLaMA models. These results allowing us to analyze the vulnerability of modern models with better coverage. We conducted the experiments on RTE datasets with 8-shot demonstrations. Even larger models such as LLaMA-2-70B or Mixtral-8x7B (Jiang et al., 2024) are ignored due to attacking them is expensive and slow. Instead, we leave them in the following section for attack transferability (§4.4).

The results are shown in Figure 3 for R-ICL (corresponding results for ICL are presented in Figure 5). It can be observed that most adversarial attacks remain effective across all models, including the Swap-Labels attack we proposed. Thus, it still posed a major threat on modern LLMs. Specifically, Vicuna-7B tends to be more sensitive when attack on text sample (*TextBugger*) or when demonstrations are under attack (*AdvICL*); while Gemma-7B—most recently released model—shows a higher vulnerability when the context or label in demonstrations is adversarial perturbed, compared to other base models. Notably, the




Figure 4: Analysis of the attack transferability for Retrieval-ICL on larger variant models within the same family: (Left) the LLaMA family; (Mid) the Mistral family; and (Right) across models from different families. The models’ orders are sorted by their parameter sizes and release date. The models highlighted in bold are the models being attacked.

irrelevant context attacks become less vulnerable to more recent release models, especially Vicuna, which shows nearly no effect after being tuned with instructions, showing improved capability in the latest LLMs for filtering OoD contexts (Shi et al., 2023).

4.4 Attack Transferability

In this section, we would like to answer two questions: Are these attacks **transferable** across different models, and can larger models **mitigate** the vulnerability to these attacks? To this end, we conducted two sets of experiments: (1) transfer attacks within the same model family from small models to their large variants, by testing on LLaMA-2 and Mistral families; (2) whether attacks still be effective across models from different family, by transferring attacks to Vicuna-7B, Mistral-7B and Gemma-7B, following the same settings as Section 4.3. Our experiments focused on the RTE dataset, chosen for its pronounced drop in robustness when subjected to these attacks. Due to memory constraints, both Mixtral-8×7B and LLaMA-2-70B were run with 8-bit quantization (Frantar et al., 2022), while the remaining models were in full precision. Hong et al. (2024) tested on AdvGLUE++ and showed that quantized LLMs have negligible drops in adversarial robustness to their source model.

The results of the transfer attack with R-ICL are shown in Figure 4 (corresponding results for ICL are presented in Figure 6). Most attacks are transferable and achieve an ASR of 15%-40% on similar scale models; except for AdvICL, which cannot transfer between models. Also, Swap-Labels attacks are shown to be less effective when targeting instruction-tuned models (e.g. Vicuna and Mistral-7B-Instruct), pointing to the instruction-tuning process leading models to be less sensitive to its demonstration (Sun et al., 2023). Specifically, the results shown in Figure 4(L) confirm that larger models tend to be more robust, as indicated by a reduction of 6% on average in ASR between the 13B and 70B models. However, the MoE variants of Mistral-7B-Instruct exhibit a similar or even higher ASR compared to its dense version, as illustrated in Figure 4(M). This finding deviates from previous work by Puigcerver et al. (2022), which investigates adversarial robustness in MoE, showing Vision MoE models are less vulnerable to adversarial attacks. We believe this discrepancy might be subject to the difference between on how MoE is applied in NLP (token) and Vision (image-patch) models. Remember that in the training of MoE models, experts are specialised to handle different tasks (Xue et al., 2024); perturbations in the input texts might change the router output and be routed to experts that are less capable for the targeted tasks, leading to more vulnerability and changes in the output. We hope our findings encourage future work to study the robustness of MoE models, both from the training and router perspectives, given their widespread use nowadays.

Defence	Shots	Clean [↑]	TB [↓]	TF [↓]	BA [↓]	Avg [↓]
↔ No Defence	8	71.48	68.69	63.14	90.4	74.08
Augmentation						
↔ Random Addition	8	72.01	67.75	67.76	91.6	75.70
↔ Random Deletion	8	69.18	65.72	62.78	85.96	71.49
DARD						
↔ R-ICL (BM25)	8	<u>74.39</u>	58.30	51.43	79.32	63.01
↔ R-ICL (SBERT)	8	72.16	<u>58.53</u>	<u>44.46</u>	72.60	58.53
↔ R-ICL (Instructor)	8	71.26	<u>69.00</u>	61.81	84.31	71.70
Adversarial Training	8	77.22	62.22	41.96	<u>77.22</u>	<u>60.47</u>

Table 2: *DARD* for adversarial defences. We show the Clean Accuracy and the ASR. For Random Addition and Deletion baselines, BM25 is the retriever as it performs the best.

5 Defence Methods: DARD

In this section, we propose an easy yet effective method to enhance the robustness of retrieval-based models against adversarial attacks on test samples. Specifically, we introduce an adversarial defence method named **Demonstration Augmentation Retrieval Defences (DARD)**. This approach performs adversarial augmentation on the training data and mixes them into the retrieval bases for R-ICL. The initiative is that adversarial training in LLM can be costly in memory, and Xie et al. (2021); Dai et al. (2022) showed that in-context learning can be seen as an implicit finetuning method.

5.1 Methodology

The methods begin by adversarially perturbing the training examples, similar to previous approaches. Our experiments use the RTE dataset, focusing on an 8-shot setting for this section. For each of the 872 test samples in RTE, the retriever selects the 8-shot most similar examples for demonstrations. After deduplication, we are left with 7,293 distinct examples. These are concatenated with one-shot examples, and adversarial attacks (TF, TB, and BA) are performed on the extracted training samples. Only the successful examples are retained, resulting in 16,206 perturbed examples. In the adversarial training paradigm, such perturbed examples would typically be used to fine-tune models. However, our approach, *DARD*, reintegrates these examples into the training demonstration for retrieval during the inference stage. It is important to note that we impose a constraint whereby each example and its perturbed examples can be retrieved no more than once in same prompt for ICL. This is because we observed that most retrieved samples were variants of the same examples with perturbations without this constraint, which reduced both performance and robustness.

5.2 Results

For baselines, we compared with **No Defence**, **Random Augmentation** (Addition, Deletion) inspired by Robey et al. (2023), which are smoothing methods used to defend against attacks. We also compared with adversarially trained models following vanilla ICL setup. We show the results in Table 2 (detailed results are in Table 19). Although the adversarially trained models show a higher clean accuracy as it’s directly fine-tuned on RTE data, our methods with SBERT outperform it by 2% reduction in ASR; both BM25 and SBERT variants of **DARD** have significant improvement on robustness against **No Defence** baseline.

6 Summary

While ICL sensitivity is a widely known issue, our paper provides an alternative view of the problem through adversarial attacks. In summary, our paper conducts a comprehensive analysis of the effects of perturbations on modern few-shot methods for LLMs inference and also proposes new adversarial attacks targeted at ICL models. Besides their perfor-

mance, the findings show the other side regarding retrieval-based models related to their robustness. Furthermore, we also demonstrate the potential of retrieval-based models in data augmentation strategies for adversarial defense.

There are many open questions for future work. One of the most profound findings is that our study on attack transferability demonstrates that larger models generally exhibit better robustness; however, this does not necessarily hold true for Mixture of Experts (MoE) models. Given the recent surge in their popularity, it is crucial to investigate the conditions under which MoE models may fail and to understand the reasons behind their vulnerabilities.

7 Limitations

We view our work as an early exploration of the adversarial robustness of Retrieval-based in-context learning. Here, we only show *how* different types of attack can be applied to, but not *why*. Given the standard use of retrieval and in-context learning as inference methods, it is crucial to understand its deficiency. Future work should focus on *certified robustness* (Zhao et al., 2022; Xiang et al., 2024), which analyzes with theoretical proof the extent to which the model is unaffected by adversarial attacks.

Acknowledgement

Computations described in this research were supported by the Edinburgh International Data Facility (EIDF) and the Data-Driven Innovation Programme at the University of Edinburgh. We thank Giwon Hong and Yu Zhao for the insightful discussion at the beginning of this research work.

References

- Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hanna Hajishirzi, and Wen tau Yih. Reliable, adaptable, and attributable language models with retrieval. 2024. URL <https://api.semanticscholar.org/CorpusID:268248911>.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020. URL <https://api.semanticscholar.org/CorpusID:218971783>.
- Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown, and He He. On the relation between sensitivity and accuracy in in-context learning. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 155–167, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.12. URL <https://aclanthology.org/2023.findings-emnlp.12>.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, 2005. URL <https://api.semanticscholar.org/CorpusID:8587959>.

- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. Why can gpt learn in-context? language models implicitly perform gradient descent as meta-optimizers. 2022. URL <https://api.semanticscholar.org/CorpusID:258686544>.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *ArXiv*, abs/2307.08691, 2023. URL <https://api.semanticscholar.org/CorpusID:259936734>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. URL <https://api.semanticscholar.org/CorpusID:52967399>.
- I. Evtimov, Kevin Eykholt, Earlene Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Xiaodong Song. Robust physical-world attacks on deep learning models. *arXiv: Cryptography and Security*, 2017. URL <https://api.semanticscholar.org/CorpusID:4730292>.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *ArXiv*, abs/2210.17323, 2022. URL <https://api.semanticscholar.org/CorpusID:253237200>.
- Thomas Mesnard Gemma Team, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, and et al. Gemma. 2024. doi: 10.34740/KAGGLE/M/3301. URL <https://www.kaggle.com/m/3301>.
- Shreyansh Goyal, Sumanth Doddapaneni, Mitesh M.Khapra, and Balaraman Ravindran. A survey of adversarial defences and robustness in nlp. 2022. URL <https://api.semanticscholar.org/CorpusID:247447518>.
- Jiasheng Gu, Hanzi Xu, Liang Nie, and Wenpeng Yin. Robustness of learning from task instructions. In *Annual Meeting of the Association for Computational Linguistics*, 2022. URL <https://api.semanticscholar.org/CorpusID:254366640>.
- Felix Hamborg, Norman Meuschke, Corinna Breiting, and Bela Gipp. news-please: A generic news crawler and extractor. In *Proceedings of the 15th International Symposium of Information Science*, pp. 218–223, March 2017. doi: 10.5281/zenodo.4120316.
- Junyuan Hong, Jinhao Duan, Chenhui Zhang, Zhangheng Li, Chulin Xie, Kelsey Lieberman, James Diffenderfer, Brian Bartoldson, Ajay Jaiswal, Kaidi Xu, Bhavya Kailkhura, Dan Hendrycks, Dawn Song, Zhangyang Wang, and Bo Li. Decoding compressed trust: Scrutinizing the trustworthiness of efficient llms under compression. 2024. URL <https://api.semanticscholar.org/CorpusID:268680727>.
- Minqing Hu and Bing Liu. Mining and summarizing customer reviews. *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2004. URL <https://api.semanticscholar.org/CorpusID:207155218>.
- Wenyu Huang, Mirella Lapata, Pavlos Vougiouklis, Nikos Papasrantopoulos, and Jeff Z. Pan. Retrieval Augmented Generation with Rich Answer Encoding. 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 (IJCNLP 2023)*, 2023.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L’elio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mixtral of experts. *ArXiv*, abs/2401.04088, 2024. URL <https://api.semanticscholar.org/CorpusID:266844877>.

- Albert Qiaochu Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Deendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L'elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b. *ArXiv*, abs/2310.06825, 2023. URL <https://api.semanticscholar.org/CorpusID:263830494>.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *AAAI Conference on Artificial Intelligence*, 2019. URL <https://api.semanticscholar.org/CorpusID:202539059>.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. Textbugger: Generating adversarial text against real-world applications. *ArXiv*, abs/1812.05271, 2018. URL <https://api.semanticscholar.org/CorpusID:54815878>.
- Linyang Li, Ruotian Ma, Qipeng Guo, X. Xue, and Xipeng Qiu. Bert-attack: Adversarial attack against bert using bert. *ArXiv*, abs/2004.09984, 2020. URL <https://api.semanticscholar.org/CorpusID:216036179>.
- Xiaonan Li, Kai Lv, Hang Yan, Tianya Lin, Wei Zhu, Yuan Ni, Guo Tong Xie, Xiaoling Wang, and Xipeng Qiu. Unified demonstration retriever for in-context learning. *ArXiv*, abs/2305.04320, 2023. URL <https://api.semanticscholar.org/CorpusID:258557751>.
- Zongyi Li, Jianhan Xu, Jiehang Zeng, Linyang Li, Xiaoqing Zheng, Qi Zhang, Kai-Wei Chang, and Cho-Jui Hsieh. Searching for an effective defender: Benchmarking defense against adversarial word substitution. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3137–3147, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.251. URL <https://aclanthology.org/2021.emnlp-main.251>.
- Bill Yuchen Lin, Kangmin Tan, Chris Miller, Beiwen Tian, and Xiang Ren. Un-supervised cross-task generalization via retrieval augmentation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 22003–22017. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/8a0d3ae989a382ce6e50312bc35bf7e1-Paper-Conference.pdf.
- Ziqian Lin and Kangwook Lee. Dual operating modes of in-context learning. 2024. URL <https://api.semanticscholar.org/CorpusID:268063278>.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. What makes good in-context examples for GPT-3? In Eneko Agirre, Marianna Apidianaki, and Ivan Vulić (eds.), *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pp. 100–114, Dublin, Ireland and Online, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.deelio-1.10. URL <https://aclanthology.org/2022.deelio-1.10>.
- Shayne Longpre, Sayash Kapoor, Kevin Klyman, Ashwin Ramaswami, Rishi Bommasani, Borhane Blili-Hamelin, Yangsibo Huang, Aviya Skowron, Zheng-Xin Yong, Suhas Kotha, Yi Zeng, Weiyan Shi, Xianjun Yang, Reid Southen, Alexander Robey, Patrick Chao, Diyi Yang, Ruoxi Jia, Daniel Kang, Sandy Pentland, Arvind Narayanan, Percy Liang, and Peter Henderson. A safe harbor for ai evaluation and red teaming. *ArXiv*, abs/2403.04893, 2024.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8086–8098, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.556. URL <https://aclanthology.org/2022.acl-long.556>.

- Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaitė, and Vincent Zhao. Dr.icl: Demonstration-retrieved in-context learning. *ArXiv*, abs/2305.14128, 2023. URL <https://api.semanticscholar.org/CorpusID:258841276>.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 11048–11064, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.759. URL <https://aclanthology.org/2022.emnlp-main.759>.
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 119–126, 2020.
- Jeff Z. Pan, Simon Razniewski, Jan-Christoph Kalo, Sneha Singhania, Jiaoyan Chen, Stefan Dietze, Hajira Jabeen, Janna Omeliyanenko, Wen Zhang, Matteo Lissandrini, ussa Biswas, Gerard de Melo, Angela Bonifati, Edlira Vakaj, Mauro Dragoni, and amien Graux. Large language models and knowledge graphs: Opportunities and challenges. *Transactions on Graph Data and Knowledge*, 2023.
- Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In Kevin Knight, Hwee Tou Ng, and Kemal Oflazer (eds.), *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pp. 115–124, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics. doi: 10.3115/1219840.1219855. URL <https://aclanthology.org/P05-1015>.
- Joan Puigcerver, Rodolphe Jenatton, Carlos Riquelme, Pranjal Awasthi, and Srinadh Bhojanapalli. On the adversarial robustness of mixture of experts. *ArXiv*, abs/2210.10253, 2022. URL <https://api.semanticscholar.org/CorpusID:252992497>.
- Mrigank Raman, Pratyush Maini, J. Zico Kolter, Zachary Chase Lipton, and Danish Pruthi. Model-tuning via prompts makes nlp models adversarially robust. *ArXiv*, abs/2303.07320, 2023. URL <https://api.semanticscholar.org/CorpusID:257495746>.
- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.
- Stephen E. Robertson and Hugo Zaragoza. The probabilistic relevance framework: Bm25 and beyond. *Found. Trends Inf. Retr.*, 3:333–389, 2009. URL <https://api.semanticscholar.org/CorpusID:207178704>.
- Alexander Robey, Eric Wong, Hamed Hassani, and George J. Pappas. Smoothllm: Defending large language models against jailbreaking attacks. *ArXiv*, abs/2310.03684, 2023. URL <https://api.semanticscholar.org/CorpusID:263671542>.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. *ArXiv*, abs/2112.08633, 2021. URL <https://api.semanticscholar.org/CorpusID:245218561>.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 31210–31227. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/shi23a.html>.

- Weijia Shi, Julian Michael, Suchin Gururangan, and Luke Zettlemoyer. Nearest neighbor zero-shot inference. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 3254–3265, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.214. URL <https://aclanthology.org/2022.emnlp-main.214>.
- Chenglei Si, Zhengyan Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. Better robustness by more coverage: Adversarial and mixup data augmentation for robust finetuning. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 1569–1576, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.137. URL <https://aclanthology.org/2021.findings-acl.137>.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/D13-1170>.
- Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. One embedder, any task: Instruction-finetuned text embeddings. 2022. URL <https://arxiv.org/abs/2212.09741>.
- Jiu Sun, Chantal Shaib, and Byron Wallace. Evaluating the zero-shot robustness of instruction-tuned language models. *ArXiv*, abs/2306.11270, 2023. URL <https://api.semanticscholar.org/CorpusID:259203613>.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, D. Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013. URL <https://api.semanticscholar.org/CorpusID:604334>.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *ArXiv*, abs/2307.09288, 2023. URL <https://api.semanticscholar.org/CorpusID:259950998>.
- Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Dan Boneh, and Patrick Mcdaniel. Ensemble adversarial training: Attacks and defenses. *ArXiv*, abs/1705.07204, 2017. URL <https://api.semanticscholar.org/CorpusID:21946795>.
- Ellen M. Voorhees and Dawn M. Tice. The TREC-8 question answering track. In M. Gavriliadou, G. Carayannis, S. Markantonatou, S. Piperidis, and G. Stainhauer (eds.), *Proceedings of the Second International Conference on Language Resources and Evaluation (LREC'00)*, Athens, Greece, May 2000. European Language Resources Association (ELRA). URL <http://www.lrec-conf.org/proceedings/lrec2000/pdf/26.pdf>.

- Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, Sang Truong, Simran Arora, Mantas Mazeika, Dan Hendrycks, Zi-Han Lin, Yuk-Kit Cheng, Sanmi Koyejo, Dawn Xiaodong Song, and Bo Li. Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. *ArXiv*, abs/2306.11698, 2023a. URL <https://api.semanticscholar.org/CorpusID:259202782>.
- Jiong Wang, Zi yang Liu, Keun Hee Park, Muhao Chen, and Chaowei Xiao. Adversarial demonstration attacks on large language models. *ArXiv*, abs/2305.14950, 2023b. URL <https://api.semanticscholar.org/CorpusID:258865399>.
- Liang Wang, Nan Yang, and Furu Wei. Learning to retrieve in-context examples for large language models. *ArXiv*, abs/2307.07164, 2023c. URL <https://api.semanticscholar.org/CorpusID:259924840>.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=gEZrGCozdqR>.
- Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122. Association for Computational Linguistics, 2018. URL <http://aclweb.org/anthology/N18-1101>.
- Chong Xiang, Tong Wu, Zexuan Zhong, David Wagner, Danqi Chen, and Prateek Mittal. Certifiably robust rag against retrieval corruption, 2024. URL <https://arxiv.org/abs/2405.15556>.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. *ArXiv*, abs/2111.02080, 2021. URL <https://api.semanticscholar.org/CorpusID:241035330>.
- Benfeng Xu, Quan Wang, Zhendong Mao, Yajuan Lyu, Qiaoqiao She, and Yongdong Zhang. knn prompting: Beyond-context learning with calibration-free nearest neighbor inference. *ArXiv*, abs/2303.13824, 2023a. URL <https://api.semanticscholar.org/CorpusID:257756989>.
- Nan Xu, Fei Wang, Ben Zhou, Bangzheng Li, Chaowei Xiao, and Muhao Chen. Cognitive overload: Jailbreaking large language models with overloaded logical thinking. *ArXiv*, abs/2311.09827, 2023b. URL <https://api.semanticscholar.org/CorpusID:265221395>.
- Rongwu Xu, Brian S. Lin, Shujian Yang, Tianqi Zhang, Weiyan Shi, Tianwei Zhang, Zhixuan Fang, Wei Xu, and Han Qiu. The earth is flat because...: Investigating llms’ belief towards misinformation via persuasive conversation. *ArXiv*, abs/2312.09085, 2023c.
- Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. Openmoe: An early effort on open mixture-of-experts language models. *arXiv preprint arXiv:2402.01739*, 2024.
- Yahan Yang, Soham Dan, Dan Roth, and Insup Lee. In and out-of-domain text adversarial robustness via label smoothing. *ArXiv*, abs/2212.10258, 2022. URL <https://api.semanticscholar.org/CorpusID:254877265>.
- Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. Compositional exemplars for in-context learning. In *International Conference on Machine Learning*, 2023. URL <https://api.semanticscholar.org/CorpusID:256826793>.
- Kang Min Yoo, Junyeob Kim, Hyuhng Joon Kim, Hyunsoo Cho, Hwiyeol Jo, Sang-Woo Lee, Sang-goo Lee, and Taeuk Kim. Ground-truth labels matter: A deeper look into input-label demonstrations. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang

(eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 2422–2437, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.155. URL <https://aclanthology.org/2022.emnlp-main.155>.

Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyang Shi. How johnny can persuade llms to jailbreak them: Rethinking persuasion to challenge ai safety by humanizing llms. *ArXiv*, abs/2401.06373, 2024. URL <https://api.semanticscholar.org/CorpusID:266977395>.

Haiteng Zhao, Chang Ma, Xinshuai Dong, Anh Tuan Luu, Zhi-Hong Deng, and Hanwang Zhang. Certified robustness against natural language attacks by causal intervention, 2022. URL <https://arxiv.org/abs/2205.12331>.

Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Weirong Ye, Neil Zhenqiang Gong, Yue Zhang, and Xingxu Xie. Promptbench: Towards evaluating the robustness of large language models on adversarial prompts. *ArXiv*, abs/2306.04528, 2023. URL <https://api.semanticscholar.org/CorpusID:259095572>.

Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *ArXiv*, abs/2307.15043, 2023. URL <https://api.semanticscholar.org/CorpusID:260202961>.

A Experiment Settings

This section discusses the details of the ICL methods used, including their hyperparameters, and also the verbalizer and prompt templates for classification tasks. The details on each attack are also provided for reproducibility. Code are included as supplementary materials and will be publicly available upon the paper is accepted.

A.1 Details about the ICL methods

kNN-ICL the exact formulation follows the original paper (Xu et al., 2023a).

R-ICL We give more details about the retriever used throughout the paper here:

1. **BM25** (Robertson & Zaragoza, 2009): A prevailing sparse retriever. For BM25, we follow Luo et al. (2023) and use uncased BERT wordpiece tokenization (Devlin et al., 2019).
2. **SBERT** (Reimers & Gurevych, 2019): We use the Sentence-BERT as the dense demonstration retriever. Specifically, we used all-MiniLM-L6-v2 to encode the test input and training set’s inputs, and retrieve the examples with the most similar input as demonstrations.
3. **Instructor** (Su et al., 2022): Instructor is a recently proposed competitive text embedding model trained on 330 tasks with instructions. By providing specialized instruction, it can serve for demonstration retrieval. We conduct experiments on its released large-size model ³.

A.2 The statistics of datasets

The table 3 shows the number of categories for each classification dataset, the number of examples in the training and test sets, and the average sentence length.

A.3 Templates and verbaliser for datasets

See Table 4 for the used templates and verbaliser for mapping discrete labels into label space (Adopted from Min et al. (2022) and Lu et al. (2022)).

³[hkunlp/instructor-large](https://github.com/hkunlp/instructor-large)

Datasets	No. of Labels	Train/Test Size	Avg. Length
SST-2	2	67348/872	23.07
RTE	2	2489/277	96.03
MR	2	8530/1066	37.76
CR	2	3394/376	30.80
MNLI-mm	3	8832/982	55.90
TREC	6	5452/500	19.59

Table 3: Datasets Information

Task	Template	Verbaliser
SST2	Review: contains no wit , only labored gags Sentiment: negative Review: the film is powerful , accessible and funny . Sentiment:	0: negative, 1: positive
RTE	A man is due in court later charged with the murder 26 years ago of a teenager whose case was the first to be featured on BBC One’s Crimewatch. Colette Aram, 16, was walking to her boyfriend’s house in Keyworth, Nottinghamshire, on 30 October 1983 when she disappeared. Her body was later found in a field close to her home. Paul Stewart Hutchinson, 50, has been charged with murder and is due before Nottingham magistrates later.” The question is Paul Stewart Hutchinson is accused of having stabbed a girl. True or False? The Answer is: false For women earning 22,000 a year, the total pay accumulated after six months maternity leave would be just 5,300 in the UK and 5,850 in Ireland. Entitlements in Germany would also be relatively low, at 5,900, along with those in France, Spain and the Netherlands, all at 6,750. At the other end of the scale, pay received after six months leave in Italy would be 9,150 while in Denmark and Norway it would be as much as 11,000. The question is Maternity leave varies in Europe. True or False? The Answer is:	0: false, 1: true
MR	Review: “you might say tykwer has done all that heaven allows , if you wanted to make as anti-kieslowski a pun as possible . suffice to say its total promise is left slightly unfulfilled .” Sentiment: negative Review: an alternately raucous and sappy ethnic sitcom . . . you’d be wise to send your regrets . Sentiment:	0: negative, 1: positive
CR	Review: it ‘s not as stylized as a sony or samsung . Sentiment: negative Review: i went out and got the canon today . Sentiment:	0: negative, 1: positive
MNLI	Instruction: Please identify whether the premise entails the hypothesis. The answer should be exactly ‘yes’, ‘no’ or ‘maybe’. Premise: We serve a classic Tuscan meal that includes a Florentine terrine made with dick and chicken livers . Hypothesis: We serve a meal of Florentine terrine . Prediction: yes Premise: After a lifetime of trials, Donna not only earned her GED at Goodwill, she earned a job here. Hypothesis: Donna went through Goodwill to get her GED, but was still unemployed. Prediction:	0: yes, 1: maybe, 2: no
TREC	Question: How did serfdom develop in and then leave Russia ? Type: description Question: What is Shakespeare ‘s nickname? Type:	0: expression, 1: entity, 2: description, 3: human, 4: location, 5: number,

Table 4: Templates and verbaliser used across the experiments. These are minimum cases with only one demonstration example for illustration.

A.4 Attack Details

We implemented our proposed attack methods and conducted all our experiments using the TextAttack framework (Morris et al., 2020), which is model-agnostic and can be easily tested on any model. Here are the details regarding to each attack method:

1. **TextBugger**: We employ the TextBugger (Li et al., 2018), which determines the importance of words and replaces key words to conduct the attack. We use GloVe embeddings to retrieve nearby words of the words to be replaced for the attack.
2. **TextFooler**: We also evaluated another popular word-level method, TextFooler (Jin et al., 2019), which is similar principle to TextBugger. However, TextFooler greedily searches for a large number of nearby words in the embedding space for each word, as long as they meet some constraints on embedding similarity and sentence quality. Additional constraints require the replacement words to match the Part-Of-Speech (POS) of the original words.
3. **BERT-Attack**: We used the popular character-level attack, BERT-Attack (Li et al., 2020), which utilizes BERT to generate candidate replacement words based on the importance of words. Then, by calculating the impact of each replacement word on the classification outcome, the replacement word that has the most significant impact on the classification result is chosen as the final replacement.
4. **AdvICL**: We mostly follow the implementations by Wang et al. (2023b), except for limiting the attack percentage to be 15%. Also, for RTE and MNLI datasets, which consist of both premise and hypothesis, we only allow perturbations on the hypothesis as it's closer to the realistic cases, and the premise is usually much longer than the hypothesis.
5. **Swap-Labels/Swap-Labels (Fix)**: For both of the attacks, we reused the *Greedy-WordSwapWIR* implementation in textattack framework. The attacks start by measuring the importance of each label in the demonstration. It is done by first replacing each of the labels with a meaningless word, fetching it into the LLM to get the label distribution, and computing the difference to the original label distribution. Based on this, the Swap-Labels attack focuses on the most important labels first, y_i , by replacing it with other labels from the label set, namely $y'_i \in \{Y/y_i\}$. If that reduces the LLM output probability on the correct label, the attack models continue to attack the subsequent most important labels with the perturbed examples with y_i replaced by y'_i ; otherwise, it keeps y_i and attack another labels. This process iterates until the attack is successful or the maximum perturbations allowed are reached. The only difference between Swap-Labels (Fix) and Swap-Labels is that an additional constraint is imposed: the label distribution on the attacked samples should be the same as the original distribution.
6. **Irrelevant Context** Following Min et al. (2022), CC-news (Hamborg et al., 2017) is used as the corpus providing OoD sentences. Sentence length is considered to select the most relevant sentences to the original demonstration. The contamination rate is 50%.

Examples of each attack are demonstrated in Table 6-12. Code will be released to facilitate further research in this field.

A.5 Reproducibility

We conducted all attack experiments using either two A100-40GB GPUs or one A100-80GB GPU. The batch size is 8. For ICL and k NN-ICL that consists of randomness, we conducted experiments with seed $\{1, 13, 42\}$. Most attacks took between 1 and 2.5 hours to run. The exception was the 16-shot demonstration attack, which usually lasted 5 hours or longer. We used Flash-Attention 2 (Dao, 2023) to speed up the experiments, which gave us a 2 to 2.5 times speed increase. The speedup significantly helped our comprehensive experiments.

B Complete Results

B.1 Main results

We include the complete results breakdown in Table 13-18. The results shown in those tables are the Benign accuracy and Attack accuracy.




Figure 5: Attack Success Rate (%) for adversarial attacks across various models, based on experiments conducted on the RTE dataset with 8-shot demonstrations. The results are based on ICL and involve aggregating results from 3 shots.




Figure 6: Transferable attack for ICL on larger variant models within the same family: (L) the LLaMA family; (M) the Mistral family; and (R) across models from different families. The models' orders are sorted by their parameter sizes and release date. The models highlighted in bold are the adversarial models under attack.

B.2 Quantitative Results

We conducted a human study to evaluate the feasibility of adversarial attacks and estimate human performance against them. Specifically, we provided 100 input examples to students and asked them to determine the correct label for each given sentence (answer options: true or false). We used the RTE dataset and instructed participants to give their answers solely based on the provided context. We assessed various types of attacks.

In this set of 100 sentences, the baseline model we used, $R_{\text{Instructor-ICL}}$, achieved an accuracy of 88% on the original, unaltered sentences. First, we observed that word-level attacks on test examples had a minor impact on human judgment, with an average decrease of only 3% (Table 5). Attacks such as advicl and irrelevant attacks resulted in decreases of only 7% and 0%, respectively. However, these attacks could significantly affect retrieval models, causing attack successful accuracies (ASR) ranging from 15% to 95%, indicating that the models are more fragile and sensitive than anticipated.

For swap-label and swap-label (fix) attacks, human performance decreased by 25% and 13%, respectively. This suggests that attacks manipulating the label of demonstrations indeed caused confusion for humans, demonstrating their effectiveness.

	TB	TF	BA	AdvICL	S-L	S-L (Fix)	Irr.
Perturbed Examples	84%	87%	82%	81%	63%	75%	88%

Table 5: Summary of the responses from the human evaluation. The total number of evaluated examples is 100. The accuracy under unperturbed conditions on these 100 examples is 88%.

Unperturbed	<p>only a few mag-lev trains have been used commercially such as at the birmingham airport in the uk., The question is: maglev is commercially used. True or False? The answer is: ture.</p> <p>they, too, have not produced a practical, commercially acceptable maglev. The question is: maglev is commercially used. True or False? The answer is: false</p> <p>two weeks ago, china became the first nation to operate a maglev railway commercially, when officials inaugurated a 30-kilometer-long line between downtown shanghai and the city’s airport. The question is: maglev is commercially used., True or False? The answer is: true.</p> <p>the m-2000 uses a commercially manufactured nbti superconductor, similar to that used in the m-2000 maglev magnets., The question is: maglev is commercially used. True or False? The answer is: false</p> <p>it appears that the super-conducting maglev system is technically ready to be used commercially as a very high-speed, large-capacity transportation system. The question is: maglev is commercially used. True or False? The answer is false.</p>
Perturbed	<p>only a few mag-lev trains have been used commercially such as at the birmingham airport in the uk., The question is: maglev is commercially used. True or False? The answer is: ture.</p> <p>they, too, have not produced a practical, commercially acceptable maglev., The question is: maglev is commercially used. True or False? The answer is: false</p> <p>two weeks ago, china became the first nation to operate a maglev railway commercially, when officials inaugurated a 30-kilometer-long line between downtown shanghai and the city’s airport. The question is: maglev is commercially used., True or False? The answer is: true.</p> <p>the m-2000 uses a commercially manufactured nbti superconductor, similar to that used in the m-2000 maglev magnets., The question is: maglev is commercially used. True or False? The answer is: false</p> <p>it appears that the super-conducting maglev system is technically ready to be used commercially as a very high-speed, large-capacity transportation system. The question is: maglev is commercially utilize. True or False? The answer is true.</p>

Table 6: This example illustrates the TextBugger attack instance. The portion surrounded by a gray background represents our test example, and the **text** in blue font indicates the words that were altered in the TextBugger attack question.

Unperturbed	<p>actual statistics about the deterrent value of capital punishment are not available because it is impossible to know who may have been deterred from committing a crime. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>capital punishment is a catalyst for more crime. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>rather than deterring crime, capital punishment actually increases the level of brutality in society., The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>, the death penalty is not a deterrent. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>a closely divided u.s. supreme court said on thursday its 2002 ruling that juries and not judges must impose a death sentence applies only to future cases, a decision that may affect more than 100 death row inmates. The question is: the supreme court decided that only judges can impose the death sentence. True or False? The answer is false.</p>
Perturbed	<p>actual statistics about the deterrent value of capital punishment are not available because it is impossible to know who may have been deterred from committing a crime. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>capital punishment is a catalyst for more crime. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>rather than deterring crime, capital punishment actually increases the level of brutality in society. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>, the death penalty is not a deterrent. The question is: capital punishment is a deterrent to crime. True or False? The answer is false.</p> <p>a closely divided u.s. supreme court said on thursday its 2002 ruling that juries and not judges must impose a death sentence applies only to future cases, a decision that may affect more than 100 death row inmates. The question is: the high court decision that only magistrates can impose the death condemnation. True or False? The answer is true.</p>

Table 7: This example illustrates the TextFooler attack instance. The portion surrounded by a gray background represents our test example, and the **text** in blue font indicates the words that were altered in the TextFooler attack question.

Unperturbed	<p>gold mining operations in california and nevada use cyanide to extract the precious metal. The question is: cyanide is used in gold mining. True or False? The answer is: true</p> <p>access to the underground workings at the la camorra mine is via a ramp from the surface, excavated at a -15% grade and connecting numerous levels. The question is: la camorra is a mine. True or False? The answer is: true</p> <p>the mine would operate nonstop seven days a week and use tons of cyanide each day to leach the gold from crushed ore. The question is: a weak cyanide solution is poured over it to pull the gold from the rock. True or False? The answer is: false</p> <p>the dam covering what used to be its football pitch fill up with millions of tons of poisonous mining waste. The question is: mine waste-water poses an environmental hazard. True or False? The answer is: false</p> <p>known as “heap leach” mining, the method has become popular in the last decade because it enables microscopic bits of gold to be economically extracted from low-grade ore. The question is: the mining industry uses a method known as heap leaching. True or False? The answer is true.</p>
Perturbed	<p>gold mining operations in california and nevada use cyanide to extract the precious metal. The question is: cyanide is used in gold mining. True or False? The answer is: true</p> <p>access to the underground workings at the la camorra mine is via a ramp from the surface, excavated at a -15% grade and connecting numerous levels. The question is: la camorra is a mine. True or False? The answer is: true</p> <p>the mine would operate nonstop seven days a week and use tons of cyanide each day to leach the gold from crushed ore. The question is: a weak cyanide solution is poured over it to pull the gold from the rock. True or False? The answer is: false</p> <p>the dam covering what used to be its football pitch fill up with millions of tons of poisonous mining waste. The question is: mine waste-water poses an environmental hazard. True or False? The answer is: false</p> <p>known as “heap leach” mining, the method has become popular in the last decade because it enables microscopic bits of gold to be economically extracted from low-grade ore. The question is: the mining industry uses a method known as heap leachmanufacturing. True or False? The answer is false.</p>

Table 8: This example illustrates the BERT-Attack attack instance. The portion surrounded by a gray background represents our test example, and the **text** in blue font indicates the words that were altered in the BERT-Attack attack question.

Unperturbed	<p>in other words, with its 2 million inhabitants, slovenia has only 5.5 thousand professional soldiers. The question is: slovenia has 5.5 million inhabitants. True or False? The answer is: false nicholas burns stated in an exclusive interview for "utrinski vesnik" from pristina that it would be disgraceful if athens puts a veto on macedonia's application for membership in the european union and nato. The question is: greece and macedonia are in dispute over name. True or False? The answer is: false greece objects to the neighbouring skopje government using the name macedonia, saying it implies claims on a greek province of the same name. The question is: greece and macedonia are in dispute over name. True or False? The answer is: true in other words, with its 2 million inhabitants, slovenia has only 5.5 thousand professional soldiers. The question is: slovenia has 2 million inhabitants. True or False? The answer is: true the croatian intent is even more problematic because the border between slovenian and croatian territorial waters has not yet been established. the dispute about this border began in 1991 when both countries became independent. The question is: there is a territorial waters dispute. True or False? The answer is true.</p>
Perturbed	<p>in other words, with its 2 million inhabitants, slovenia has only 5.5 thousand professional soldiers. The question is: slovenia has 5.5 billion inhabitants. True or False? The answer is: false nicholas burns stated in an exclusive interview for "utrinski vesnik" from pristina that it would be disgraceful if athens puts a veto on macedonia's application for membership in the european union and nato. The question is: greece and macedonia are in dispute over name. True or False? The answer is: false greece objects to the neighbouring skopje government using the name macedonia, saying it implies claims on a greek province of the same name. The question is: greece and macedonia are in dispute over name. True or False? The answer is: true in other words, with its 2 million inhabitants, slovenia has only 5.5 thousand professional soldiers. The question is: slovenia has 2 million inhabitants. True or False? The answer is: true the croatian intent is even more problematic because the border between slovenian and croatian territorial waters has not yet been established. the dispute about this border began in 1991 when both countries became independent. The question is: there is a territorial waters dispute. True or False? The answer is false.</p>

Table 9: This example illustrates the AdvICL attack instance. The portion surrounded by a gray background represents our test example, and the text in blue font indicates the words that were altered in the AdvICL attack question.

Unperturbed	<p>money raised from the sale will go into a trust for hepburn’s family. The question is:proceeds go to hepburn’s family. True or False? The answer is:true</p> <p>treasures belonging to hollywood legend katharine hepburn have raised £3.2m at a two-day auction in america. The question is: a two-day auction of property belonging to actress katharine hepburn brought in 3.2 million pounds. True or False? The answer is: true</p> <p>a two-day auction of property belonging to actress katharine hepburn brought in 3.2 million pounds. The question is: a two-day auction of property belonging to actress katharine hepburn brought in £3.2m. True or False? The answer is: true</p> <p>hepburn, a four-time academy award winner, died last june in connecticut at age 96. The question is: hepburn, who won four oscars, died last june aged 96. True or False? The answer is: true</p> <p>hepburn’s family will receive the proceeds from the sale. The question is: proceeds go to hepburn’s family. True or False? The answer is true.</p>
Perturbed	<p>money raised from the sale will go into a trust for hepburn’s family. The question is:proceeds go to hepburn’s family. True or False? The answer is: false</p> <p>treasures belonging to hollywood legend katharine hepburn have raised £3.2m at a two-day auction in america. The question is: a two-day auction of property belonging to actress katharine hepburn brought in 3.2 million pounds. True or False? The answer is: true</p> <p>a two-day auction of property belonging to actress katharine hepburn brought in 3.2 million pounds. The question is: a two-day auction of property belonging to actress katharine hepburn brought in £3.2m. True or False? The answer is: false</p> <p>hepburn, a four-time academy award winner, died last june in connecticut at age 96. The question is: hepburn, who won four oscars, died last june aged 96. True or False? The answer is: true</p> <p>hepburn’s family will receive the proceeds from the sale. The question is: proceeds go to hepburn’s family. True or False? The answer is false.</p>

Table 10: This example illustrates the Swap-Labels attack instance. The portion surrounded by a gray background represents our test example, and the **text** in blue font indicates the words that were altered in the Swap-Labels attack question.

<p>Unperturbed</p>	<p>herceptin was already approved to treat the sickest breast cancer patients, and the company said, monday, it will discuss with federal regulators the possibility of prescribing the drug for more breast cancer patients. The question is: herceptin can be used to treat breast cancer. True or False? The answer is: true only 14 percent of u.s. mothers exclusively breast-feed their babies for the minimum recommended six months. The question is: there are many benefits from breast-feeding. True or False? The answer is: false new research shows there has been a sharp increase in disfiguring skin cancers, particularly in women under the age of 40, providing more evidence that young people are not heeding warnings about the dangers of tanning. The question is: tanning may cause skin cancers. True or False? The answer is: true weinstock painstakingly reviewed dozens of studies for evidence of any link between sunscreen use and either an increase or decrease in melanoma. The question is: skin cancer numbers increase. True or False? The answer is: false a compound in breast milk has been found to destroy many skin warts, raising hopes it also might prove effective against cervical cancer and other lethal diseases caused by the same virus The question is: breast milk may help fight cervical cancer. True or False? The answer is true.</p>
<p>Perturbed</p>	<p>herceptin was already approved to treat the sickest breast cancer patients, and the company said, monday, it will discuss with federal regulators the possibility of prescribing the drug for more breast cancer patients. The question is: herceptin can be used to treat breast cancer. True or False? The answer is: false only 14 percent of u.s. mothers exclusively breast-feed their babies for the minimum recommended six months. The question is: there are many benefits from breast-feeding. True or False? The answer is: false new research shows there has been a sharp increase in disfiguring skin cancers, particularly in women under the age of 40, providing more evidence that young people are not heeding warnings about the dangers of tanning. The question is: tanning may cause skin cancers. True or False? The answer is: true weinstock painstakingly reviewed dozens of studies for evidence of any link between sunscreen use and either an increase or decrease in melanoma. The question is: skin cancer numbers increase. True or False? The answer is: true a compound in breast milk has been found to destroy many skin warts, raising hopes it also might prove effective against cervical cancer and other lethal diseases caused by the same virus The question is: breast milk may help fight cervical cancer. True or False? The answer is false.</p>

Table 11: This example illustrates the Swap-Labels (Fix) attack instance. The portion surrounded by a gray background represents our test example, and the **text** in blue font indicates the words that were altered in the Swap-Labels (Fix) attack question.

Unperturbed	<p>rabies virus infects the central nervous system, causing encephalopathy and ultimately death. early symptoms of rabies in humans are nonspecific, consisting of fever, headache, and general malaise. The question is: rabies is fatal in humans. True or False? The answer is: true</p> <p>rabies is a viral disease of mammals and is transmitted primarily through bites. annually, 7,000 to 8,000 rabid animals are detected in the united states, with more than 90 percent of the cases in wild animals. The question is: rabies is fatal in humans. True or False? The answer is: false</p> <p>former tennis star vitas gerulaitis died in such fashion in september, from a lethal carbon monoxide buildup related to the faulty installation of a propane heater. The question is: vitas gerulaitis died of carbon monoxide poisoning. True or False? The answer is: true</p> <p>more than 6,400 migratory birds and other animals were killed in nevada by drinking water in the cyanide-laced ponds produced by gold mining operations. The question is: animals have died by the thousands from drinking at cyanide-laced holding ponds. True or False? The answer is: true</p> <p>a farmer who was in contact with cows suffering from bse – the so-called mad cow disease – has died from what is regarded as the human form of the disease. The question is: bovine spongiform encephalopathy is another name for the “mad cow disease”. True or False? The answer is true.</p>
Perturbed	<p>The further into summer we go, the more California’s winter swells seem like a distant memory. Do you remember? Tons of rain, a few massive days up and down the coast, and more rain. ... The latter half’s roping barrels will make you want to throw on your booties, crank the heat, and grab your step-up. The question is: A motorcyclist struck a cow and suffered fatal injuries in Pierce County on Saturday morning, authorities said. True or False? The answer is: true</p> <p>most of the mines are in arid areas and animals searching for water are attracted to the cyanide-laced holding ponds that are an integral part of the mining operations. The question is: animals have died by the thousands from drinking at cyanide-laced holding ponds. True or False? The answer is: false</p> <p>the unconfirmed case in dundee concerns a rabies-like virus known only in bats. The question is: a case of rabies was confirmed. True or False? The answer is: false</p> <p>this case of rabies in western newfoundland is the first case confirmed on the island since 1989. The question is: a case of rabies was confirmed. True or False? The answer is: true.</p> <p>a farmer who was in contact with cows suffering from bse – the so-called mad cow disease – has died from what is regarded as the human form of the disease. The question is: bovine spongiform encephalopathy is another name for the “mad cow disease”. True or False? The answer is false.</p>

Table 12: This example illustrates the Irrelevant attack instance. The portion surrounded by a gray background represents our test example, and all demonstrations have been replaced by irrelevant attacks.

Shots	Method	Clean	Test Sample Attack			Demonstration Attack			Datstore Attack Irr.
			TB	TF	BA	AdvICL	S-L	S-L (Fix)	
$m = 2$	ICL	94.65 \pm 0.85	57.57 \pm 2.14	49.81 \pm 0.96	25.46 \pm 0.35	68.33 \pm 6.72	25.11 \pm 14.34	74.29 \pm 2.24	92.01 \pm 0.89
	kNN-ICL	88.3 \pm 7.85	49.08 \pm 2.88	47.17 \pm 1.15	33.29 \pm 6.29	24.15 \pm 2.17	30.11 \pm 4.47	30.28 \pm 1.89	89.22 \pm 6.47
	R _{BM25} -ICL	90.48	43.12	44.04	36.35	73.62	6.88	55.3	89.91
	R _{SBERT} -ICL	92.43	56.42	51.61	43.46	74.33	7.43	59.20	89.45
	R _{Instructor} -ICL	94.95	58.49	53.56	41.06	71.29	16.9	57.39	92.66
$m = 4$	ICL	94.49 \pm 1.1	58.53 \pm 1.89	51.34 \pm 1.27	25.92 \pm 0.3	69.33 \pm	16.11 \pm 9.49	61.29 \pm 3.24	90.48 \pm 0.64
	kNN-ICL	90.25 \pm 5.47	52.03 \pm 4.04	50.3 \pm 1.95	34.48 \pm 2.99	25.28 \pm 4.18	30.11 \pm 4.47	34.28 \pm 4.90	91.01 \pm 2.12
	R _{BM25} -ICL	93.0	53.67	49.66	42.32	76.83	1.49	57.8	91.86
	R _{SBERT} -ICL	87.84	52.64	46.67	43.12	71.28	3.27	52.73	83.69
	R _{Instructor} -ICL	92.78	59.86	53.56	42.32	72.09	4.77	51.84	89.89
$m = 8$	ICL	92.66 \pm 3.18	58.11 \pm 5.01	52.18 \pm 2.29	29.28 \pm 1.56	61.49 \pm 5.23	7.11 \pm 8.49	34.29 \pm 8.24	91.09 \pm 3.6
	kNN-ICL	92.01 \pm 2.13	54.28 \pm 2.57	50.19 \pm 1.26	35.28 \pm 2.38	19.28 \pm 4.28	30.11 \pm 4.47	31.28 \pm 3.92	92.01 \pm 1.53
	R _{BM25} -ICL	94.61	53.56	51.15	40.37	79.82	0.11	34.17	93.81
	R _{SBERT} -ICL	95.18	63.76	63.76	41.63	69.28	0.34	31.56	89.38
	R _{Instructor} -ICL	94.84	58.49	63.90	39.22	68.08	0.44	29.10	88.46
$m = 16$	ICL	92.12 \pm 1.06	56.69 \pm 5.07	49.39 \pm 4.7	28.29 \pm 3.46	60.67 \pm 5.2	4.11 \pm 4.27	18.30	78.02 \pm 21.45
	kNN-ICL	93.65 \pm 1.49	53.97 \pm 1.79	52.37 \pm 1.74	36.57 \pm 1.28	15.79 \pm 6.22	30.11 \pm 4.47	31.99 \pm 3.20	92.54 \pm 2.56
	R _{BM25} -ICL	94.61	56.77	50.25	38.76	70.5	0.33	15.14	91.72
	R _{SBERT} -ICL	95.07	60.32	53.29	41.51	64.08	0.00	12.47	92.72
	R _{Instructor} -ICL	94.38	59.06	53.18	41.28	62.08	0.00	11.99	92.61

Table 13: Complete results for the SST-2 dataset. The term **Clean** refers to the benign accuracy before attacks; attack accuracy is provided under each attack column.

Shots	Method	Clean	Test Sample Attack			Demonstration Attack			Datstore Attack Irr.
			TB	TF	BA	AdvICL	S-L	S-L (Fix)	
$m = 2$	ICL	68.47 \pm 5.98	18.65 \pm 7.33	20.46 \pm 8.66	3.01 \pm 2.41	60.73 \pm 3.87	20.22 \pm 9.07	48.26 \pm 11.07	67.27 \pm 4.41
	kNN-ICL	61.25 \pm 6.47	23.47 \pm 18.96	25.39 \pm 16.97	4.28 \pm 3.11	3.28 \pm 2.77	29.96 \pm 2.96	31.05 \pm 2.73	64.86 \pm 7.09
	R _{BM25} -ICL	70.76	14.08	20.58	3.25	60.38	13.78	56.68	62.82
	R _{SBERT} -ICL	70.04	11.91	21.30	1.81	59.39	11.49	36.29	59.57
	R _{Instructor} -ICL	71.48	14.83	18.05	1.81	59.40	9.09	37.02	68.59
$m = 4$	ICL	70.88 \pm 2.12	17.09 \pm 6.32	20.7 \pm 6.55	2.65 \pm 1.99	61.27 \pm 5.39	12.03 \pm 3.36	43.68 \pm 10.4	66.19 \pm 3.51
	kNN-ICL	64.26 \pm 6.25	9.93 \pm 5.87	15.73 \pm 3.33	3.18 \pm 2.62	4.27 \pm 2.19	25.15 \pm 3.99	26.35 \pm 4.34	59.33 \pm 6.09
	R _{BM25} -ICL	68.59	15.16	27.80	3.18	61.22	8.66	34.77	59.21
	R _{SBERT} -ICL	71.84	14.44	16.97	0.72	58.27	9.20	31.29	60.29
	R _{Instructor} -ICL	71.48	19.86	18.05	2.76	55.38	10.23	37.78	61.48
$m = 8$	ICL	73.04 \pm 1.27	12.75 \pm 2.4	17.81 \pm 3.76	1.56 \pm 0.56	62.58 \pm 1.37	6.14 \pm 1.81	30.45 \pm 3.36	64.98 \pm 7.19
	kNN-ICL	70.52 \pm 3.90	6.62 \pm 2.56	13.84 \pm 5.28	3.22 \pm 1.41	2.73 \pm 1.19	25.63 \pm 3.82	25.99 \pm 3.77	65.7 \pm 4.7
	R _{BM25} -ICL	71.48	22.38	26.35	6.86	58.06	5.42	28.16	60.65
	R _{SBERT} -ICL	72.20	11.91	12.27	1.36	52.35	3.25	25.63	61.01
	R _{Instructor} -ICL	73.29	15.88	20.22	2.44	52.71	6.86	24.92	65.70
$m = 16$	ICL	73.53 \pm 0.75	8.9 \pm 2.18	11.91 \pm 1.88	1.2 \pm 0.21	59.21 \pm 8.90	4.17 \pm 1.62	24.26 \pm 3.26	68.95 \pm 3.37
	kNN-ICL	70.88 \pm 3.51	14.44 \pm 7.05	16.97 \pm 6.03	1.32 \pm 0.98	1.73 \pm 0.22	24.89 \pm 2.91	25.12 \pm 3.02	69.67 \pm 3.82
	R _{BM25} -ICL	64.15	20.94	41.52	18.41	38.27	3.83	27.87	58.84
	R _{SBERT} -ICL	73.98	15.78	26.71	4.33	54.29	4.87	26.80	64.26
	R _{Instructor} -ICL	69.33	15.16	23.83	1.44	56.86	7.12	28.12	64.26

Table 14: Complete results for the RTE dataset. The term **Clean** refers to the benign accuracy before attacks; attack accuracy is provided under each attack column.

Shots	Method	Clean	Test Sample Attack		Demonstration Attack			Datastore Attack		Irr.
			TB	TF	BA	AdvICL	S-L	S-L (Fix)		
$m = 2$	ICL	91.97 \pm 0.21	57.4 \pm 1.65	49.67 \pm 0.32	24.67 \pm 0.32	75.33 \pm 3.72	23.11 \pm 12.36	48.57 \pm 0.64	86.93 \pm 7.05	
	kNN-ICL	90.9 \pm 0.42	55.95 \pm 0.92	49.67 \pm 1.86	33.67 \pm 4.11	23.77 \pm 1.17	25.38 \pm 8.47	26.88 \pm 7.47	87.12 \pm 12.01	
	R _{BM25} -ICL	89.6	44.38	47.92	26.83	73.26	16.1	75.3	88.9	
	R _{SBERT} -ICL	90.4	57.4	53.6	29.22	70.33	9.7	78.8	89.7	
	R _{Instructor} -ICL	91.3	42.9	32.6	28.45	71.43	8.4	86.9	89.8	
$m = 4$	ICL	92.07 \pm 0.32	58.0 \pm 2.6	50.87 \pm 1.21	28.97 \pm 0.5	73.29 \pm 4.02	14.11 \pm 8.36	45.37 \pm 2.66	92.13 \pm 0.8	
	kNN-ICL	92.3 \pm 1.04	56.93 \pm 1.5	49.67 \pm 1.86	31.67 \pm 3.20	21.33 \pm 4.18	24.76 \pm 3.40	26.12 \pm 3.51	91.22 \pm 4.01	
	R _{BM25} -ICL	92.0	48.2	49.8	29.01	71.38	3.3	61.0	90.3	
	R _{SBERT} -ICL	91.6	56.6	54.2	33.28	70.24	4.5	51.8	91.4	
	R _{Instructor} -ICL	91.6	49.5	40.7	31.47	69.11	2.8	78.4	91.7	
$m = 8$	ICL	92.67 \pm 0.35	58.93 \pm 2.63	49.37 \pm 1.07	30.1 \pm 0.5	71.49 \pm 6.23	5.11 \pm 4.28	22.77 \pm 14.96	92.53 \pm 0.49	
	kNN-ICL	91.77 \pm 0.32	55.43 \pm 3.52	48.47 \pm 2.08	31.25 \pm 1.20	19.28 \pm 3.21	25.87 \pm 2.43	24.87 \pm 5.43	89.22 \pm 3.01	
	R _{BM25} -ICL	92.5	57.4	50.1	31.5	69.28	0.3	33.4	91.1	
	R _{SBERT} -ICL	92.4	56.8	49.8	32.78	65.76	0.2	36.9	92.4	
	R _{Instructor} -ICL	92.6	57.9	48.8	31.28	68.28	1.8	32.1	92.2	
$m = 16$	ICL	92.77 \pm 0.45	57.8 \pm 0.53	50.53 \pm 0.59	29.57 \pm 0.5	69.67 \pm 3.1	4.11 \pm 2.27	21.77 \pm 6.96	93.0 \pm 0.4	
	kNN-ICL	92.13 \pm 0.29	55.63 \pm 1.86	48.9 \pm 1.35	32.25 \pm 0.34	17.79 \pm 2.12	22.87 \pm 3.61	22.87 \pm 2.43	87.22 \pm 2.01	
	R _{BM25} -ICL	93.5	59.6	51.0	29.01	65.5	4.4	20.34	92.3	
	R _{SBERT} -ICL	93.0	55.7	48.8	31.47	61.12	6.6	19.85	92.5	
	R _{Instructor} -ICL	93.1	57.2	49.3	33.28	63.33	5.3	16.38	92.5	

Table 15: Complete results for the MR dataset. The term **Clean** refers to the benign accuracy before attacks; attack accuracy is provided under each attack column.

Shots	Method	Clean	Test Sample Attack		Demonstration Attack			Datastore Attack		Irr.
			TB	TF	BA	AdvICL	S-L	S-L (Fix)		
$m = 2$	ICL	93.35 \pm 0.96	66.04 \pm 2.17	47.61 \pm 2.27	27.48 \pm 2.47	76.21 \pm 4.27	16.04 \pm 13.28	63.03 \pm 4.52	88.12 \pm 1.26	
	kNN-ICL	71.87 \pm 11.2	29.4 \pm 2.38	29.4 \pm 2.38	29.4 \pm 2.38	27.15 \pm 6.27	15.28 \pm 2.38	17.90 \pm 1.09	64.4 \pm 6.38	
	R _{BM25} -ICL	92.29	63.03	51.06	0	69.62	24.20	80.59	91.29	
	R _{SBERT} -ICL	92.02	62.23	51.60	32.71	71.34	8.24	82.18	92.0	
	R _{Instructor} -ICL	93.09	63.56	58.51	39.36	72.29	1.86	90.69	92.1	
$m = 4$	ICL	93.26 \pm 0.85	61.79 \pm 7.05	45.92 \pm 2.47	28.28 \pm 0.67	71.33 \pm 2.01	1.68 \pm 2.47	47.87 \pm 1.88	89.2 \pm 1.44	
	kNN-ICL	73.87 \pm 8.40	29.4 \pm 2.38	29.4 \pm 2.38	29.4 \pm 2.38	26.28 \pm 3.18	17.28 \pm 1.29	19.4 \pm 4.29	67.4 \pm 1.38	
	R _{BM25} -ICL	91.49	65.65	45.48	0	67.24	1.86	61.70	90.34	
	R _{SBERT} -ICL	93.09	56.38	48.67	32.18	65.28	0.53	69.68	91.2	
	R _{Instructor} -ICL	93.27	62.50	60.37	42.82	64.09	0.00	88.03	90.47	
$m = 8$	ICL	91.31 \pm 3.31	58.86 \pm 7.22	43.44 \pm 4.25	26.33 \pm 1.88	70.49 \pm 3.25	0.09 \pm 0.16	20.21 \pm 16.51	90.2 \pm 0.96	
	kNN-ICL	91.87 \pm 4.29	29.4 \pm 2.38	29.4 \pm 2.38	29.4 \pm 2.38	20.19 \pm 2.11	18.28 \pm 2.66	18.9 \pm 2.37	69.4 \pm 1.22	
	R _{BM25} -ICL	93.09	65.96	48.94	31.91	62.18	0.00	40.69	86.29	
	R _{SBERT} -ICL	93.62	64.36	50.27	26.60	61.97	0.00	35.28	89.16	
	R _{Instructor} -ICL	93.88	61.44	56.91	43.35	60.26	0.00	32.97	90.11	
$m = 16$	ICL	84.31 \pm 12.73	53.37 \pm 10.37	41.58 \pm 5.53	28.9 \pm 2.03	66.67 \pm 2.5	3.1 \pm 0.15	61.88 \pm 21.06	84.1 \pm 0.96	
	kNN-ICL	83.29 \pm 2.1	29.4 \pm 2.38	29.4 \pm 2.38	29.4 \pm 2.38	18.79 \pm 1.99	19.28 \pm 1.83	20.2 \pm 1.97	74.4 \pm 0.27	
	R _{BM25} -ICL	93.09	68.62	60.11	38.30	59.34	9.31	85.90	87.28	
	R _{SBERT} -ICL	93.88	69.95	54.52	30.05	61.27	6.12	74.47	90.32	
	R _{Instructor} -ICL	94.00	64.10	60.90	45.21	63.11	4.21	77.13	87.10	

Table 16: Complete results for the CR dataset. The term **Clean** refers to the benign accuracy before attacks; attack accuracy is provided under each attack column.

Shots	Method	Clean	Test Sample Attack			Demonstration Attack			Datstore Attack Irr.
			TB	TF	BA	AdvICL	S-L	S-L (Fix)	
$m = 2$	ICL	51.07 \pm 4.14	19.07 \pm 4.92	21.83 \pm 5.79	24.22 \pm 4.79	42.10 \pm 6.54	10.43 \pm 0.32	35.10 \pm 4.72	50.4 \pm 3.38
	kNN-ICL	53.07 \pm 5.28	17.03 \pm 5.29	22.03 \pm 2.26	21.38 \pm 5.22	39.14 \pm 3.98	19.03 \pm 8.22	28.48 \pm 3.18	49.07 \pm 4.67
	R _{BM25} -ICL	58.4	13.6	16.5	30.18	28.4	2.8	28.67	49.6
	R _{SBERT} -ICL	55.7	12.6	20.8	26.43	30.42	0.8	26.1	53.7
	R _{Instructor} -ICL	53.2	14.6	20.4	25.71	31.74	1.6	22.5	54.1
$m = 4$	ICL	54.17 \pm 2.19	17.33 \pm 2.61	22.47 \pm 2.0	27.68 \pm 3.25	38.23 \pm 4.66	8.03 \pm 0.06	32.10 \pm 6.15	52.7 \pm 1.59
	kNN-ICL	55.07 \pm 2.72	16.12 \pm 3.87	25.13 \pm 3.26	23.4 \pm 3.19	30.14 \pm 1.57	22.03 \pm 6.31	26.7 \pm 3.75	51.46 \pm 3.67
	R _{BM25} -ICL	54.8	17.4	19.9	29.75	34.57	1.1	24.36	55.4
	R _{SBERT} -ICL	57.5	13.0	16.8	28.84	33.81	0.5	24.08	55.8
	R _{Instructor} -ICL	54.3	16.6	14.2	27.79	32.68	0.3	20.47	56.5
$m = 8$	ICL	53.63 \pm 2.99	16.9 \pm 4.33	22.57 \pm 4.94	31.72 \pm 2.17	36.21 \pm 3.27	6.0 \pm 0.0	30.12 \pm 1.48	52.7 \pm 1.75
	kNN-ICL	55.89 \pm 3.61	19.22 \pm 2.62	26.77 \pm 2.52	27.22 \pm 3.01	27.22 \pm 1.25	21.02 \pm 4.18	24.59 \pm 3.11	50.22 \pm 1.90
	R _{BM25} -ICL	57.3	18.1	23.4	33.82	35.27	2.3	20.69	57.1
	R _{SBERT} -ICL	57.8	17.9	25.97	32.79	36.41	0.4	21.07	56.2
	R _{Instructor} -ICL	54.3	18.2	26.1	34.72	33.76	0.0	16.47	52.9
$m = 16$	ICL	53.3 \pm 3.13	17.73 \pm 4.9	22.63 \pm 2.87	35.72 \pm 1.88	38.32 \pm 5.81	4.3 \pm 2.0	29.48 \pm 2.46	51.8 \pm 0.82
	kNN-ICL	57.89 \pm 1.34	20.79 \pm 1.93	25.68 \pm 1.94	36.17 \pm 1.05	18.22 \pm 6.31	23.48 \pm 1.23	25.12 \pm 4.15	49.37 \pm 0.63
	R _{BM25} -ICL	59.1	20.2	21.78	38.09	38.02	1.1	19.72	52.3
	R _{SBERT} -ICL	58.9	21.5	23.4	36.4	37.45	0.2	20.78	54.3
	R _{Instructor} -ICL	55.5	22.4	26.9	33.21	33.47	0.4	18.93	53.8

Table 17: Complete results for the MNLI dataset. The term **Clean** refers to the benign accuracy before attacks; attack accuracy is provided under each attack column.

Shots	Method	Clean	Test Sample Attack			Demonstration Attack			Datstore Attack Irr.
			TB	TF	BA	AdvICL	S-L	S-L (Fix)	
$m = 2$	ICL	69.87 \pm 9.93	19.6 \pm 10.05	27.53 \pm 7.76	25.13 \pm 5.77	60.73 \pm 4.87	0.13 \pm 0.23	0.07 \pm 0.12	59.12 \pm 4.93
	kNN-ICL	71.87 \pm 11.2	18.4 \pm 7.38	26.4 \pm 8.92	31.48 \pm 5.89	33.28 \pm 6.77	30.28 \pm 2.38	31.90 \pm 4.09	64.4 \pm 6.38
	R _{BM25} -ICL	70.2	26	10.6	14.4	60.28	0.2	30.6	65.3
	R _{SBERT} -ICL	84.0	35.6	11.6	41.8	61.34	2.0	41.0	78.12
	R _{Instructor} -ICL	75.6	30.8	11.8	39.0	62.44	1.8	46.0	72.10
$m = 4$	ICL	70.53 \pm 9.56	22.67 \pm 8.62	26.27 \pm 11.44	24.67 \pm 8.77	59.27 \pm 5.39	21.87 \pm 2.53	12.53 \pm 9.1	65.22 \pm 3.78
	kNN-ICL	73.87 \pm 8.40	27.4 \pm 4.38	31.2 \pm 5.22	36.1 \pm 4.29	24.27 \pm 5.2	31.98 \pm 1.29	34.41 \pm 3.29	67.4 \pm 1.38
	R _{BM25} -ICL	82.8	5.0	13.6	20.2	55.32	27	52.2	79.2
	R _{SBERT} -ICL	84.8	43.2	18.4	44.0	54.76	28.2	54.6	81.4
	R _{Instructor} -ICL	82.0	34.6	22.0	38.4	51.29	30.4	58.4	80.6
$m = 8$	ICL	76.07 \pm 5.88	29.6 \pm 6.55	32.33 \pm 6.93	29.93 \pm 6.64	61.85 \pm 3.17	26.2 \pm 3.68	45.2 \pm 6.32	69.84 \pm 2.11
	kNN-ICL	77.87 \pm 4.29	31.4 \pm 5.22	35.4 \pm 3.92	38.92 \pm 2.76	12.73 \pm 2.18	32.13 \pm 3.66	36.9 \pm 2.37	69.4 \pm 1.22
	R _{BM25} -ICL	79.2	39.6	39.0	34.0	56.3	35.2	71.6	76.2
	R _{SBERT} -ICL	87.8	49.4	47.8	46.4	57.41	41.2	79.0	79.8
	R _{Instructor} -ICL	86.5	47.2	25.2	45.4	59.42	47.6	83.6	80.2
$m = 16$	ICL	83.53 \pm 3.41	38.07 \pm 3.23	42.2 \pm 5.03	39.53 \pm 3.82	59.21 \pm 8.90	32.2 \pm 1.78	65.2 \pm 3.32	71.21 \pm 1.33
	kNN-ICL	83.29 \pm 2.1	39.4 \pm 3.28	41.39 \pm 3.21	42.1 \pm 3.12	10.73 \pm 2.23	33.28 \pm 4.83	36.2 \pm 1.97	74.4 \pm 0.27
	R _{BM25} -ICL	82.0	42.4	40.4	37.0	58.27	41.2	77.2	80.2
	R _{SBERT} -ICL	90.0	54.8	56.8	48.4	56.13	43.8	84.2	87.4
	R _{Instructor} -ICL	89.6	52.0	35.0	47.4	56.68	42.4	86.6	88.2

Table 18: Complete results for the TREC dataset. The term **Clean** refers to the benign accuracy before attacks; attack accuracy is provided under each attack column.

Defence	Shots	Clean Accuracy	TB	TF	BA
↔ No Defence	8	71.48	22.38	26.35	6.86
Augmentation					
↔ Random Addition ¹	8	72.01	23.22	23.21	7.29
↔ Random Deletion ¹	8	69.18	23.71	25.75	9.71
DARD					
↔ R-ICL (BM25)	8	<u>74.39</u>	31.02	36.13	15.38
↔ R-ICL (SBERT)	8	<u>74.16</u>	<u>30.53</u>	<u>41.19</u>	20.32
↔ R-ICL (Instructor)	8	71.38	22.13	27.26	11.20
Adversarial Training	8	77.22	29.17	44.82	<u>17.59</u>

Table 19: *DARD* for adversarial defences. We show the Clean Accuracy and the Attack Accuracy.