

Unveiling the Achilles' Heel of NLG Evaluators: A Unified Adversarial Framework Driven by Large Language Models

Yiming Chen[†] Chen Zhang^{†,*} Danqing Luo[†] Luis Fernando D'Haro^{*}
Robby T. Tan^{‡,†} Haizhou Li^{‡,†,§}

[†]National University of Singapore [‡]The Chinese University of Hong Kong, Shenzhen
^{*}Universidad Politécnica de Madrid [‡]ASUS Intelligent Cloud Services [§]Kriston AI Lab
{yiming.chen, chen_zhang}@u.nus.edu
danqing@nus.edu.sg luisfernando.dharo@upm.es
robby_tan@asus.com haizhouli@cuhk.edu.cn

Abstract

The automatic evaluation of natural language generation (NLG) systems presents a long-lasting challenge. Recent studies have highlighted various neural metrics that align well with human evaluations. Yet, the robustness of these evaluators against adversarial perturbations remains largely under-explored due to the unique challenges in obtaining adversarial data for different NLG evaluation tasks. To address the problem, we introduce AdvEval, a novel black-box adversarial framework against NLG evaluators. AdvEval is specially tailored to generate data that yield strong disagreements between human and victim evaluators. Specifically, inspired by the recent success of large language models (LLMs) in text generation and evaluation, we adopt strong LLMs as both the data generator and gold evaluator. Adversarial data are automatically optimized with feedback from the gold and victim evaluator. We conduct experiments on 12 victim evaluators and 11 NLG datasets, spanning tasks including dialogue, summarization, and question evaluation. The results show that AdvEval can lead to significant performance degradation of various victim metrics, thereby validating its efficacy.¹

1 Introduction

Recent advancements in generative language models, especially advanced large language models (LLMs), have led to their widespread application in creating human-like content across various domains. Concurrently, numerous automatic text evaluation metrics have been developed, aiming to identify unhelpful responses and understand the model's capacity. These metrics have exhibited good correlations with human judgment across diverse natural language generation (NLG) tasks, e.g., dialogue generation (Zhang et al., 2021a), summarization (Zhong et al., 2022; Yuan et al., 2021),

question generation (Mohammadshahi et al., 2023; Wang et al., 2022a). Despite their strong performance, the robustness of these metrics against adversarial attacks or misleading inputs remains critically under-explored. Therefore, in this paper, we aim to automatically identify "Achilles' Heel", critical weaknesses, within well-established automatic evaluation metrics.

Adversarial attack techniques are widely adopted to assess the robustness of neural models. Yet, their application in NLP has predominantly focused on general classification tasks, e.g., sentiment classification (Jin et al., 2020; Li et al., 2019a). These attack techniques all build upon the label-preserving assumption that the gold label of input text will remain unchanged after adding bounded perturbation on the input text (Wang et al., 2022b). This assumption, however, does not hold on the NLG evaluation tasks, where even slight modifications can drastically alter, often negatively, a text's perceived quality, which directly correlates with the regression label in the evaluation tasks. For example, changing a single word "movie" to "dish" in the dialogue response "That's my favorite movie" will perfectly preserve the positive sentiment classification label but significantly alter the response relevance evaluation label, given that the dialogue context is about a movie. Consequently, these traditional attacking techniques can easily obtain low-quality adversarial texts with high metric scores while struggling to produce high-quality adversarial texts that score low. Existing studies on NLG evaluators' robustness mainly rely on manual effort or heuristic perturbation rules to ensure the perturbed text accords with specified labels. These studies also fail to obtain high-quality adversarial texts with low metric scores and cannot scale up to generate large-scale diverse adversarial data (Khalid and Lee, 2022; Sai et al., 2021).

To overcome these limitations, we propose AdvEval, designed to obtain data that lead to two

* Corresponding author.

¹Code is available at github.com/MatthewCYM/AdvEval.

different types of disagreement between humans and victim evaluators: 1) text valued by humans but underrated by victim evaluators; 2) text underrated by humans but valued by victim evaluators. Inspired by recent works on using LLMs to iteratively improve their response based on current feedback (Madaan et al., 2023; Yang et al., 2023; Chen et al., 2023a), we adapt LLMs as the data generator to iteratively optimize the benign input text with the adversarial objective. As mentioned above, label preservation posts a unique challenge in attacking NLG tasks. Recognizing the recent trend of leveraging large language models (LLMs) as a substitute for human annotators to generate and annotate high-quality data (He et al., 2023; Ye et al., 2022a,b; Meng et al., 2022), we also utilize LLMs as the evaluator to give feedback during the optimization process, thereby imposing an implicit quality constraint on the generated text.

We evaluate AdvEval on 11 datasets covering three distinct NLG tasks against 12 automatic evaluation metrics that vary in their reliance on references, models, and LLMs. Experimental results demonstrate AdvEval’s ability to generate high-quality adversarial data effectively. Through manual inspection, we also confirm the label preservation property of AdvEval. Further experiments on adversarial learning, ablation studies, and case studies convincingly validate the effectiveness and generalization ability of our proposed AdvEval framework.

2 Background

2.1 NLG Evaluator

The evaluation of AI-generated content is increasingly important recently in various domains (Su et al., 2020; Gao et al., 2019; Jiang et al., 2023). Automatic NLG evaluators can be broadly divided into reference-based and reference-free metrics. Traditional reference-based metrics like BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) are favored for their simplicity and versatility across NLG tasks. Yet, their inability to capture complex textual semantics has led to the development of newer reference-based metrics utilizing contextualized embeddings for enhanced semantic understanding, such as BERTScore (Zhang et al., 2020), MoverScore (Zhao et al., 2019), and BLEURT (Selam et al., 2020).

In open-ended tasks like dialogue generation, reference-based metrics often fall short in mir-

roring human judgment (Liu et al., 2016). This has led to the development of trained reference-free metrics such as RUBER (Tao et al., 2018), DynaEval (Zhang et al., 2021a), UniEval (Zhong et al., 2022), TRUE (Honovich et al., 2022a), and PoE (Zhang et al., 2023). These reference free metrics offer closer alignment with human evaluations but still don’t reach the criteria for strong automatic evaluators² and remain task-specific.

The emergence of LLMs marks a significant step towards developing more general and effective NLG evaluators. Several recent works (Chen et al., 2023b; Liu et al., 2023a; Lin and Chen, 2023) examine the NLG evaluation capability of instruction-following LLMs via prompting and report strong correlations with human evaluation. Remarkably, GPT-4 (OpenAI, 2023) is found to exhibit human-level judgment while being more scalable and less expensive than crowd-sourced workers (Zheng et al., 2023; Gilardi et al., 2023; Chiang and Lee, 2023). Inspired by these findings, our AdvEval framework incorporates advanced proprietary LLMs like GPT-4 and PALM-2 to serve as proxies for human evaluators. Moreover, we conduct a comprehensive analysis of the robustness of metrics from the aforementioned types³ across various NLG tasks, including dialogue, summarization, and question generation.

2.2 Adversarial Attack

Adversarial attacks are common practice for robustness evaluation through creating adversarial examples to elicit incorrect predictions from neural networks (Papernot et al., 2016; Ebrahimi et al., 2018b; Li et al., 2019b; Wallace et al., 2019; Le et al., 2022; Cheng et al., 2020; Zhang et al., 2021b). These attacks manipulate benign inputs at different granularity levels. Mainstream perturbations include character-level (Belinkov and Bisk, 2018; Ebrahimi et al., 2018a), word-level (Ren et al., 2019; Jin et al., 2020; Li et al., 2021; Garg and Ramakrishnan, 2020), phrase-level (Lei et al., 2022; Chen et al., 2021a; Zheng et al., 2020), and sentence-level (Wang et al., 2020; Lin et al., 2021; Xu et al., 2021). The fundamental assumption of these works is label preservation under limited minor perturbation. The assumption does hold for various classification tasks, e.g., sentiment and intent classifi-

²Defined as having a correlation over 0.8 with human evaluations.

³Including reference-based and reference-free, trained and untrained, as well as BERT-based and LLM-based metrics.

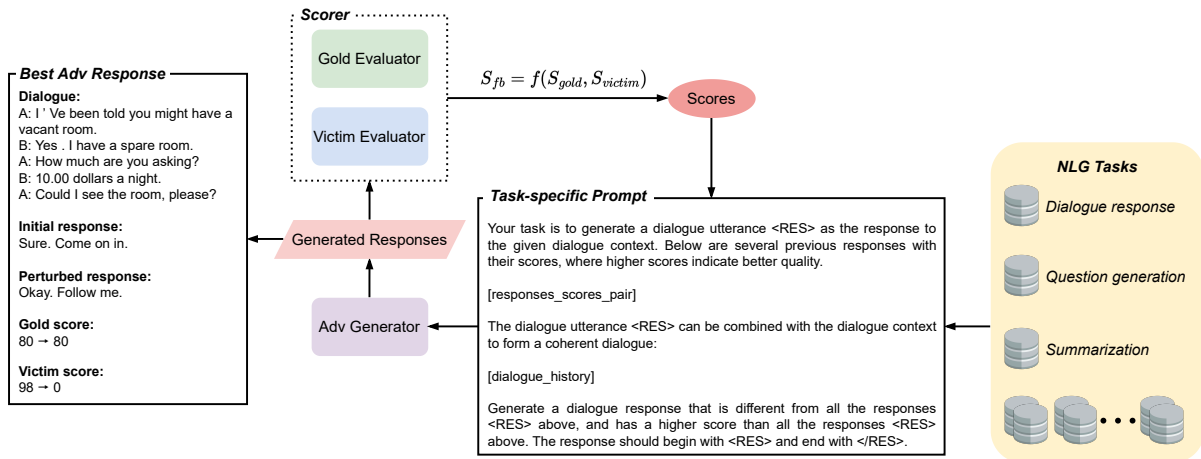


Figure 1: Overall framework of AdvEval. Optimization prompt shown above is for dialogue response generation.

ation. Yet, these perturbations are highly likely to affect the text quality, which is critical in NLG evaluation tasks. This discrepancy renders traditional adversarial techniques unsuitable for attacking NLG evaluators.

While there are also adversarial works targeting NLG evaluators (Khalid and Lee, 2022; Sai et al., 2021), they bypass the label concern by using only heuristic perturbation rules. These rules include invariant rules, e.g., phrase expansion/contraction, and deterioration rules, e.g., dropping words and reordering sentences. Yet, these rules are relatively simple and follow a certain pattern, rendering them infeasible to scale up and create diverse, high-quality adversarial data. In contrast, AdvEval utilizes a powerful generative model to craft adversarial samples with automatic quality control, which can generate large-scale, diverse natural adversarial samples.

3 Methodology

3.1 Problem Formulation

Our study focuses on conditional generation tasks, where a single data point consists of input context C , output response R , and reference ground-truth response G . Given the context and response, the evaluator E will assign a score S , indicating the quality of the response w.r.t. context:

$$S = E(C, R, G). \quad (1)$$

We aim to design a generalizable framework to generate adversarial samples for various NLG evaluation tasks. Given a victim evaluator E_{victim} , we target to obtain a perturbed response from initial response R_0 that leads to the disagreement between

score S_{victim} from E_{victim} and human judgment S_{human} . Specifically, the perturbed responses include two types: 1) R^+ where S_{human} is high and S_{victim} is low; 2) R^- where S_{human} is low and S_{victim} is high. For all E_{victim} , we assume black-box access to all E_{victim} , which means only the evaluation scores are accessible. To tackle this problem, we employ advanced LLMs as both adversarial generator G and evaluator E_{gold} . For these two models, we only have black-box access, where only output text is available.

3.2 Overview

Fig. 1 illustrates the overall framework of AdvEval. The generator G iteratively optimizes the adversarial response R . In each step, G generates new candidate responses with the task-specific prompt as the input. These new candidates are then scored by both E_{gold} and E_{victim} , generating a feedback score S_{fb} . Afterwards, these candidate responses, along with their feedback scores, will be included in the generation prompt for the next iteration to guide G towards generating better adversarial responses. The optimization process ends if the score reaches a pre-defined threshold or the iteration count hits the maximum limit. Finally, the response with the highest score will be returned as the adversarial response. The following sections will elaborate on the task-specific prompt and scorer designs.

3.3 Generation Prompt

We design task-specific prompts to iteratively improve the generated adversarial samples. Fig. 1 shows a dialogue task example prompt. The task-specific prompt includes below components.

Evaluator	Reference-free	Model-based	LLM-based	Studied task
SacreBLEU (Post, 2018)				Dialog, Summ, QA
ROUGE (Lin, 2004)				Summ
BLEURT (Sellam et al., 2020)		✓		Dialog, QA
BERTScore (Zhang et al., 2020)		✓		Dialog, Summ
PoE (Zhang et al., 2023)	✓	✓		Dialog
UniEval (Zhong et al., 2022)	✓	✓		Dialog, Summ
BARTScore (Yuan et al., 2021)	✓	✓		Summ
TRUE (Honovich et al., 2022b)	✓	✓		Summ
RQUGE (Mohammadshahi et al., 2023)	✓	✓		QA
Baichuan (Baichuan, 2023)	✓	✓	✓	Dialog, QA
Vicuna (Zheng et al., 2023)	✓	✓	✓	Dialog, QA
InstructGPT (Ouyang et al., 2022)	✓	✓	✓	Dialog

Table 1: Summary of victim models studied in AdvEval.

Task instruction: We give instructions to the adversarial generator about the optimization goal ("generate dialogue response with higher score"), the output format for easier parsing ("begin with <RES> and end with <RES>"), and also NLG task definition ("dialogue utterance as the response to the given dialogue context").

Evaluation criteria: In addition, we find that the viable response space is considerably limited, particularly for more constrained tasks like summarization. During the optimization process, the generator is prone to generate bad response, which reduces the R^+ probability. Therefore, when generating R^+ , we optionally append the evaluation criteria (e.g., "The summary must be factually consistent with the article." for factual consistent summarization) to the task instruction. This emphasis makes the generator better understand the desired output, and implicitly constrain the output space.

Optimization trajectory: Rather than direct instructions on how to update the responses, we add optimization trajectory in the generation prompt. The optimization trajectory includes a history of generated responses and their associated scores, organized in ascending score order. The score serves as the proxy of disagreement between human and victim evaluator (Sec. 3.4). We only keep top-10 responses with highest scores to reduce the context length and inference cost. We aim to use the optimization trajectory to guide G towards improved responses by combining existing good responses, or perturbing the responses along a potential desired direction.

3.4 Human Judgement Proxy

As mentioned previously, our optimization goal here is to maximize disagreement between S_{victim} and S_{human} . The disagreement serves as both the criterion for successful attacks and crucial feedback to the adversarial generator. Yet, unlike previous

work on prompt optimization, where the goal function is simply accuracy and easily accessible, it's costly and infeasible to ask human to participate in the optimization process, complicating the assignment of scores that accurately reflect human judgment. To address this problem, we propose to use a strong evaluator E_{gold} to produce S_{gold} as a proxy of S_{human} . Specifically, the E_{gold} can be a single LLM-based evaluator or an ensemble of several powerful evaluators. The E_{gold} is asked to rate the response on a 100-point scale using the manual/automatic prompt shown in Appx. B. Since we observe that LLM-based evaluators sometimes produce unstable scores, we utilize self-agreement to produce a consistent score. We first obtain 8 ratings from the evaluator through repeatedly sampling. Then, we take the arithmetic mean of these ratings as the final gold score S_{gold} . After getting the S_{gold} , it is combined with S_{victim} to get the final score S_{fb} as the feedback to generator:

$$S_{fb} = \begin{cases} \alpha \times S_{gold} - S_{victim}, & \text{if } R^+ \\ S_{victim} - \alpha \times S_{gold}, & \text{if } R^- \end{cases}, \quad (2)$$

where α is hyper-parameter balancing the influence of S_{gold} and S_{victim} . A higher α encourages the generator to exploit the response space adhere to judgement of gold evaluator, while a lower α allows the generator to more aggressively explore diverse responses that can lead to misjudgement of victim evaluator.

4 Experiments

4.1 Experiment Setup

Due to space limitations, for more details about model implementation and hyper-parameters, please refer to Appx. A.

Datasets: We conduct our experiments on three popular NLG evaluation tasks: dialogue response,

Model	DailyDialog							DREAM						
	SR	EX	JB	NU	BA	TF	Ours	SY	EX	JB	NU	BA	TF	Ours
SacreBLEU	3/-	4/-	-/-	-/-	20/-	29/-	100/-	5/-	2/-	-/-	-/-	10/-	17/-	100/-
BLEURT	0/-	0/-	-/-	-/-	1/-	6/-	79/-	0/-	0/-	-/-	-/-	0/-	3/-	92/-
BERTScore	0/-	0/-	-/-	-/-	4/-	16/-	93/-	0/-	0/-	-/-	-/-	6/-	8/-	99/-
PoE-Base	10/-	13/-	-/39	-/66	53/35	63/13	99/95	18/-	8/-	-/2	-/78	39/58	40/35	90/96
PoE-Large	6/-	10/-	-/49	-/72	54/46	52/25	95/98	8/-	4/-	-/2	-/77	28/63	39/34	81/97
UniEval	0/-	3/-	-/22	-/76	55/69	27/74	95/96	5/-	4/-	-/4	-/75	31/74	34/41	98/96
Baichuan-7B	6/-	12/-	-/34	-/24	56/68	43/51	95/94	16/-	19/-	-/34	-/30	32/69	31/59	90/88
Baichuan-13B	0/-	3/-	-/27	-/41	57/76	16/98	76/97	2/-	8/-	-/27	-/45	19/83	21/100	73/98
Vicuna-7B	6/-	9/-	-/25	-/14	58/63	43/36	98/88	9/-	14/-	-/22	-/27	36/59	33/37	92/84
Vicuna-13B	7/-	6/-	-/20	-/12	59/42	44/31	92/88	8/-	16/-	-/17	-/10	34/50	40/32	90/87
Avg.	3.8/-	6/-	-/30.9	-/43.6	41.7/57.0	33.9/46.9	92.2/93.7	7.1/-	7.5/-	-/12.3	-/48.9	23.5/65.1	26.6/48.3	90.5/92.3

	MuTual							PersonaChat						
	SY	EX	JB	NU	BA	TF	Ours	SY	EX	JB	NU	BA	TF	Ours
SacreBLEU	3/-	4/-	-/-	-/-	8/-	10/-	100/-	2/-	2/-	-/-	-/-	2/-	2/-	100/-
BLEURT	0/-	0/-	-/-	-/-	0/-	4/-	81/-	0/-	0/-	-/-	-/-	1/-	1/-	69/-
BERTScore	0/-	0/-	-/-	-/-	4/-	3/-	85/-	0/-	0/-	-/-	-/-	1/-	1/-	85/-
PoE-Base	12/-	7/-	-/0	-/84	33/51	39/23	92/96	20/-	9/-	-/0	-/80	3/33	7/32	77/100
PoE-Large	5/-	4/-	-/1	-/84	16/60	39/26	74/99	17/-	8/-	-/1	-/87	3/36	8/44	72/97
UniEval	2/-	5/-	-/5	-/74	17/74	13/39	96/97	13/-	9/-	-/9	-/84	3/40	8/45	92/94
Baichuan-7B	7/-	10/-	-/26	-/41	24/73	25/75	89/91	17/-	12/-	-/14	-/23	9/41	11/26	80/93
Baichuan-13B	2/-	4/-	-/32	-/48	13/84	20/100	73/99	11/-	6/-	-/25	-/49	3/46	0/100	52/97
Vicuna-7B	4/-	9/-	-/11	-/26	22/62	26/23	92/88	23/-	13/-	-/4	-/5	4/26	4/9	81/73
Vicuna-13B	6/-	14/-	-/8	-/12	19/42	28/18	88/88	25/-	14/-	-/5	-/2	5/12	10/8	85/75
Avg.	4.1/-	5.7/-	-/11.9	-/51.7	15.6/63.7	20.7/43.4	87/94	12.8/-	7.3/-	-/8.3	-/47.1	3.4/33.4	5.2/37.7	79.3/89.9

Table 2: Experiment results on dialogue response evaluation. Each row refers to different victim evaluators, and each column refers to different baseline methods: synonym replacement (SR), extension (EX), jumble (JB), negate previous utterance (NU), BERT-Attack (BA), TextFooler (TF). Each cell gives the R^+ ASR (left) and R^- ASR (right) respectively. For those methods that are designed to generate only single type of adversarial data, we report only one ASR score. We also omit R^- ASR for reference-based metrics, since adding any perturbation will lead to deteriorated score. Results with highest ASR are bold.

summarization, and question evaluation. We assess proposed AdvEval with 100 samples from each dataset. For dialogue response, we utilize DailyDialog (Li et al., 2017), DREAM (Sun et al., 2019), PersonaChat (Zhang et al., 2018), and MuTual (Cui et al., 2020). For summarization, data come from CNN/DailyMail (See et al., 2017), XSum (Narayan et al., 2018), and DialogSum (Chen et al., 2021b). For question generation, we use SQuAD (Rajpurkar et al., 2016), NewsQA (Trischler et al., 2017), HotpotQA (Yang et al., 2018) and Natural Questions (Kwiatkowski et al., 2019).

Evaluation dimensions: In our study, each NLG task focuses on a key evaluation dimension. For dialogue response, we emphasize relevance, which strongly correlates with human judgment. For the summarization task, we prioritize factual consistency, a vital aspect of this task. Lastly, for question evaluation, we assess overall quality, encompassing both the answer-ability of the questions and their relevance to the provided answers.

Adversarial generators & gold evaluators: For consistency and simplicity, we utilize the same LLM as gold evaluator and adversarial generator. For dialogue response generation, we use PALM-2. For summarization and question generation, PALM-2 experiences severe hallucination,

and fails to generate factual consistent summaries or answerable questions in our preliminary experiments. Therefore, we use GPT-4 as the R^+ generator/evaluator, and PALM-2 as the R^- generator/evaluator.

Victim evaluators: Our study involves a thorough examination of a diverse array of evaluators, which vary in mechanism and model size. Tab. 1 lists these evaluators.⁴

Metrics: Our primary metric for evaluating adversarial methods is the attack success rate (ASR). We generate two types of adversarial data, and a successful attack must satisfy specific criteria:

$$S_{gold} > \tau_1 \wedge S_{gold} - S_{victim} > \tau_2, \text{ if } R^+ \quad (3)$$

$$S_{gold} < \tau_1 \wedge S_{victim} - S_{gold} > \tau_2, \text{ if } R^-$$

where τ_1 and τ_2 are pre-defined thresholds. The first part of each condition ensures the adversarial data has a confident label (positive/negative), while the second part ensures the incorrect prediction by victim evaluators. Note that our ASR criteria involve E_{gold} , which is a neural metric. Therefore, on top of ASR, we also manually check the label validity of the adversarial data for trustworthy scoring. We report the rate of data with correct labels.

⁴Experiments on InstructGPT is presented in APPX. C.

Model	HotpotQA							NaturalQuestion						
	SR	CO	CA	CW	BA	TF	Ours	SR	CO	CA	CW	BA	TF	Ours
SacreBLEU	2/-	0/-	-/-	-/-	21/-	35/-	100/-	2/-	0/-	-/-	-/-	29/-	35/-	100/-
BLEURT	0/-	0/-	-/-	-/-	1/-	9/-	95/-	1/-	0/-	-/-	-/-	7/-	10/-	99/-
RQUGE	5/-	4/-	-/33	-/59	63/28	64/5	77/89	6/-	8/-	-/34	-/44	47/47	51/9	96/73
Baichuan-7B	0/-	0/-	-/65	-/84	11/92	13/92	11/100	1/-	2/-	-/50	-/66	9/94	14/95	14/100
Baichuan-13B	3/-	0/-	-/56	-/67	25/72	34/13	52/99	0/-	1/-	-/45	-/52	27/78	35/21	25/100
Vicuna-7B	6/-	3/-	-/13	-/37	28/35	34/100	61/96	7/-	10/-	-/5	-/17	26/31	41/100	53/98
Vicuna-13B	5/-	2/-	-/47	-/54	39/47	40/5	60/98	3/-	5/-	-/41	-/59	25/33	44/1	35/98
Avg.	3/-	1.29/-	-/42.8	-/60.2	26.9/54.8	32.7/43	65.1/96.4	2.86/-	3.71/-	-/35	-/47.6	24.3/56.6	32.9/45.2	60.3/93.8

Model	NewsQA							SQUAD						
	SR	CO	CA	CW	BA	TF	Ours	SR	CO	CA	CW	BA	TF	Ours
SacreBLEU	2/-	4/-	-/-	-/-	44/-	55/-	100/-	1/-	2/-	-/-	-/-	49/-	50/-	100/-
BLEURT	0/-	0/-	-/-	-/-	7/-	13/-	95/-	0/-	0/-	-/-	-/-	3/-	16/-	98/-
RQUGE	0/-	3/-	-/63	-/64	29/50	44/15	80/68	2/-	0/-	-/67	-/79	64/55	57/10	94/81
Baichuan-7B	1/-	1/-	-/86	-/91	5/92	16/48	12/100	1/-	1/-	-/50	-/80	20/88	34/98	25/99
Baichuan-13B	0/-	0/-	-/86	-/61	15/85	27/88	14/98	0/-	0/-	-/82	-/78	44	37/18	33/98
Vicuna-7B	1/-	7/-	-/15	-/27	45/34	50/95	67/85	3/-	2/-	-/21	-/46	43/37	48/100	26/92
Vicuna-13B	1/-	1/-	-/53	-/69	29/41	41/12	42/87	0/-	0/-	-/65	-/79	41/44	46/5	32/94
Avg.	0.7/-	2.3/-	-/60.6	-/62.4	24.9/60.4	35.1/51.6	58.6/87.6	1/-	0.71/-	-/57	-/72.4	37.7/61.2	41.1/46.2	58.3/92.8

Table 3: Experiment results on question evaluation. Baselines include: synonym replacement (SR), contraction (CO), change question to assertion (CA), change question word (CW), BERT-Attack (BA), TextFooler (TF).

Baselines: We compare AdvEval against two types of baselines. The first type of baselines are specially designed for NLG evaluators, and utilize sophisticated rule-based perturbations to ensure the label validity (Sai et al., 2021; Khalid and Lee, 2022). The perturbations include invariant perturbation (synonym adjective, expansion, contraction) for R^+ generation, and deterioration perturbation (jumble, change name, negate previous utterance, change question to assertion, change question word) for R^- generation. The first type of baselines serve as our major baseline for their high relevance. For the second type of baselines, we adapt existing attack methods originally designed for general classification tasks, including TextFooler (Jin et al., 2020) and BERT-Attack (Li et al., 2020), for targeted attack with target label 0 to generate R^+ and target label 1 to generate R^- . We sample a random response as the initial response during generating R^- to align with the label preservation assumption in these methods. For fair comparison, we allocate the same query budget 300 to both AdvEval and the baseline methods.

4.2 Main Results

Dialogue response evaluation: The comparison of different attacking methods against dialogue evaluation metrics are shown in Tab. 2. Overall, we have the following key observations. First, all evaluated metrics, regardless of their working mechanisms, are vulnerable to adversarial perturbations. Especially for R^- case, even simple heuristic rule-based perturbations can deceive some models effectively. For example, negating previous utterance

achieves 47.8% ASR on average. Second, despite the promising performance of rule-based methods on generating R^- , their effectiveness in creating R^+ is limited, with only ASR 6.9% with synonym replacement, and 6.6% with phrase extension. These methods maintain label accuracy but lack data diversity and are easily detectable. Third, compared to rule-based methods, optimization-based baselines are more effective for R^+ generation, but fall short in generating R^- . Additionally, it struggles to produce adversarial samples for reference-based methods, which often assign high scores to perturbed responses due to the high vocabulary overlap with the reference caused by local perturbations. Lastly, AdvEval significantly outperforms all baselines, achieving an ASR of 84.8% for R^+ and 92.5% for R^- . This underscores the efficiency of AdvEval in navigating both R^+ and R^- scenarios. In addition, we observe that it’s easier to craft R^- than R^+ for all attacking methods, likely due to limited good response space compared to bad response space.

Question evaluation: The results on question evaluation are shown in Tab. 3. Findings from the question evaluation largely mirror those from the dialogue evaluation tasks. AdvEval achieves average R^+ success rate 60.6% and R^- success rate 92.7%. However, we observe a lower R^+ ASR compared to dialogue tasks. The lower R^+ ASR might be attributed to the stringent constraint, i.e., questions to be answerable from the article and aligned with the provided answer. This leads to a limited range of appropriate responses. Despite this, AdvEval still outperforms baseline methods

Model	CNN/Dailymail					
	SR	JB	CN	BA	TF	Ours
SacreBLEU	11/-	-/-	-/-	0/-	6/-	100/-
ROUGE	0/-	-/-	-/-	0/-	20/-	100/-
BERTScore	3/-	-/-	-/-	1/-	1/-	100/-
BARTScore	15/-	-/0	-/18	69/3	76/1	84/40
UniEval	4/-	-/11	-/9	18/30	29/7	68/80
TRUE	6/-	-/62	-/2	42/8	53/2	72/41
Avg.	6.5/-	-/24.3	-/9.7	21.7/13.7	30.8/3.3	87.3/53.7
	XSum					
	SR	JB	CN	BA	TF	Ours
SacreBLEU	7/-	-/-	-/-	6/-	20/-	100/-
ROUGE	0/-	-/-	-/-	1/-	4/-	100/-
BERTScore	0/-	-/-	-/-	1/-	4/-	100/-
BARTScore	23/-	-/0	-/20	58/25	68/3	81/35
UniEval	10/-	-/0	-/12	60/45	23/6	63/60
TRUE	17/-	-/11	-/1	71/1	57/1	85/43
Avg.	9.5/-	-/3.7	-/11	32.8/23.7	29.3/3.3	88.2/46
	DialogSum					
	SR	JB	CN	BA	TF	Ours
SacreBLEU	6/-	-/-	-/-	11/-	17/-	100/-
ROUGE	2/-	-/-	-/-	5/-	3/-	100/-
BERTScore	0/-	-/-	-/-	3/-	6/-	97/-
BARTScore	22/-	-/0	-/40	64/23	69/1	80/43
UniEval	0/-	-/21	-/21	45/79	48/13	33/99
TRUE	5/-	-/53	-/15	49/1	65/0	67/88
Avg.	5.8/-	-/24.7	-/25.3	29.5/34.3	34.7/4.7	79.5/76.7

Table 4: Experiment results on summarization evaluation. Baselines include: synonym replacement (SR), jumble (JB), change name (CN), BERT-Attack (BA), TextFooler (TF).

in average performance.

Summarization evaluation: The results on summarization evaluation are shown in Tab. 4. AdvEval surpasses baseline methods with an 85% R^+ ASR, and 58.8% R^- ASR. In addition, these summarization metrics are more robust against R^- perturbation compared to dialogue and question metrics. The highest rule-based average ASR is only 15.3%, in stark contrast to 60.2% for question metric, and 47.8% for dialogue metric.

Overall, our proposed AdvEval significantly outperforms existing adversarial techniques across all examined tasks. The superior performance of AdvEval can be attributed to two key advantages. Firstly, AdvEval benefits from a larger search space. While all baseline methods are restricted to adding only local perturbations to the benign input, which significantly limits their search space, AdvEval leverages the remarkable generative capabilities of LLMs to perform both local and global perturbations. This results in a substantially larger search space, facilitating the discovery of novel and effective responses. Secondly, AdvEval includes improved search constraints. Baseline methods typically rely on either rule-based or sentence embedding constraints to maintain label integrity. However, as discussed in our paper, these constraints

fall short in text evaluation scenarios due to their inability to capture the nuances of natural language. To overcome this limitation, AdvEval adopts an LLM evaluator as a human proxy, which ensures superior label integrity by more accurately assessing the quality of generated responses.

4.3 Label Validity

In AdvEval, gold evaluator E_{gold} is crucial as it acts as a human proxy for assessing response quality and guiding the adversarial generator. To evaluate its reliability, we conduct a manual examination on 780 adversarial data points, spanning various evaluation tasks and generated through different attack methods, with 260 data points per method. Three annotators manually label each data point as 'good' or 'not good' based on context and scoring criteria, with majority voting determining the final label. Then, we convert the S_{gold} to binary high-quality/low-quality label with threshold 50. Finally, We compute the accuracy score between S_{gold} binary labels and human annotated ones to measure the reliability of E_{gold} on adversarial data. Due to space limitations, detailed annotation schema can be found in APPX. E.

As shown in Tab. 5, E_{gold} achieves an exceptionally high average accuracy over 80%, aligning closely with the highest inter-annotator agreement rates of 77%. This accuracy indicates a high degree of reliability of our gold evaluator, affirming its utility as a viable proxy for human judgment in guiding the optimization process. Additionally, we claim that a key issue with earlier adversarial methods is their inability to consistently preserve the quality label of R^+ . This is supported by our manual labeling results that 65% of R^+ samples from TextFooler are labelled as bad response by human annotators.

Method	Dialog	Summ	QA
TextFooler	85.0	78.3	99.0
BERT-Attack	78.0	83.3	90.0
AdvEval	75.0	78.3	72.0

Table 5: Agreement between human and gold evaluator on adversarial data.

4.4 Adversarial Learning

To assess potential countermeasures against adversarial attacks, we experiment with the adversarial learning on PoE-Large. We generate 3,649 adversarial data points from mixed data source of

DREAM, DailyDialog, and MuTual with AdvEval, and also 3,634 adversarial data points with TextFooler. Then, we further finetune the model on these two types of adversarial data mixed with original vanilla data respectively. As shown in Tab. 6, TextFooler can be easily defended with naive adversarial training, while our proposed AdvEval is more robust against this defense technique. After adversarial training, the model can hardly achieve any improvement, which further validates the capacity of AdvEval. We hypothesize that AdvEval can craft more diverse and complex data points that challenge the model’s adaptability.

Method	DailyDialog	DREAM	MuTual
AdvEval	95 / 98	81 / 97	74 / 99
+Adv. Train	93 / 98	81 / 97	74 / 99
TextFooler	52 / 27	39 / 44	39 / 32
+Adv. Train	9 / 6	6 / 3	5 / 8

Table 6: Adversarial learning on PoE-Large using data crafted by different methods.

4.5 Case Study

Tab. 7 presents adversarial data from different methods. More samples can be found in Appx. D. These samples are sampled from DailyDialog and SQUAD. Compared to previous methods, AdvEval is able to generate more valuable hard adversarial data with high quality and diversity. Regarding R^+ generation, BERT-Attack and TextFooler exhibit similar tendencies, notably their failure to preserve the positive label after perturbation due to the inadequacy of their perturbation techniques. For example, in the dialogue sample, A asks whether the dish is good, but the perturbed response answers that the flat is popular, which is not related to the dialogue context at all. Therefore, it should be assigned a negative label. On the contrary, AdvEval generate a good response "You'll never eat better" that directly answer the question from A. In addition, the AdvEval adversarial data are notably more complex. For instance, one need to infer "Ming dynasty succeeded Yuan dynasty" from "Yuan dynasty preceded Ming dynasty" to check the question’s answer-ability. For R^- generation, all methods can generate samples with correct label. Compared to random responses from TextFooler and BERT-Attack, AdvEval samples are also more challenging to validate. The case study convincingly demonstrates the limitations of current attack

<i>Dialog Response Evaluation</i>	
A:	I'd like to taste some local dishes . What would you recommend ?
B:	That's fine . You must try this dish .
A:	Could you tell me how this thing is cooked ?
B:	It's fish steamed and served with our special sauce .
A:	Is it good ?
[Benign R]	Sure, It's a most popular dish.
[AdvEval R^+]	You'll never eat better.
[AdvEval R^-]	I have never had it.
[BERT-Attack R^+]	Sure. It's a most popular day.
[BERT-Attack R^-]	So, what's the bad?
[TextFooler R^+]	Sure. It's a most popular flat.
[TextFooler R^-]	So, what's the question?
<i>Question Evaluation</i>	
Article: In official Chinese histories, the Yuan dynasty bore the Mandate of Heaven, following the Song dynasty and preceding the Ming dynasty.
Answer:	Ming dynasty
[Benign R]	What dynasty came after the Song?
[AdvEval R^+]	Which dynasty succeeded the Yuan dynasty according to the official Chinese historical records?
[AdvEval R^-]	What is the only dynasty mentioned to have received the Mandate of Heaven in the text?
[BERT-Attack R^+]	What dynasty came after the qing?
[BERT-Attack R^-]	which ranking did Harvard find in war?
[TextFooler R^+]	What dynasty came after the Dollar?
[TextFooler R^-]	What structured has Amherst unearth in 1900?

Table 7: Adversarial data from AdvEval, BERT-Attack, and TextFooler.

approaches and highlights the effectiveness and necessity of our proposed AdvEval framework.

4.6 Ablation Study

Gold evaluator: To understand the impact of the gold evaluator in AdvEval, we conduct a study with an AdvEval variant that only uses the victim score for feedback. For R^- generation, we use random response as the initial input to avoid direct termination with initial input. As shown in Tab. 8, excluding the gold evaluator results in a significant drop in performance. The generation process tends to end prematurely because the score threshold is quickly met without the guidance of the gold evaluator.

Evaluation criteria: We optionally include eval-

Model	DailyDialog	Dream
UniEval	95 / 96	98 / 96
- E_{gold}	34 / 6	53 / 9
Vicuna-7B	98 / 88	92 / 84
- E_{gold}	33 / 4	34 / 2

Table 8: Ablation study on the effect of gold evaluator for dialogue response generation.

uation criteria into the generation prompt, aiming to generate responses following the criteria. As shown in Tab. 9, excluding criteria leads to lower ASR. The results validate the effect of criteria in the generation prompt for guiding the LLM to generate response more adhere to the criteria.

Method	CNN/Dailymail	DialogSum
AdvEval	68	33
-criteria	61	28

Table 9: Results on UniEval with(out) criteria included in the generation prompt. R^+ ASR is reported.

5 Conclusion and Future Works

In this paper, we systematically evaluate the robustness of NLG evaluators. We also propose AdvEval, a novel black-box attack framework that effectively craft high-quality diverse adversarial samples that lead to performance degradation of various NLG evaluators. Specifically, AdvEval generates adversarial samples using LLM generator with feedback from LLM evaluator. Extensive experimental demonstrate the superior effectiveness of AdvEval across variety of NLG tasks. Future works include the development of defense methods against AdvEval, and the extension of AdvEval to other NLP tasks and NLG metrics.

Limitations

The paper proposes the first optimization-based framework to examine the adversarial robustness of NLG evaluators. Yet, there are several limitations. Firstly, the framework is only able to find absolutely bad text with score around 0 or absolutely good text with score around 100. Techniques that can generate medium-quality text with score around 50 remains unexplored. We leave it to future work. Secondly, our research primarily targets the robustness of general NLG evaluators rather than specific evaluation dimensions. The dimensions we selected for study are either overall quality

or dimension that is highly correlated to the overall quality across various tasks. We leave dimension-oriented attacks as an important direction for future research. Thirdly, adversarial attack on strong proprietary LLM evaluators, e.g., InstructGPT, demands ensemble of several powerful evaluators, which requires more computational resources. We leave a more computational efficient framework to future exploration.

Ethics Statement

We propose an adversarial attack framework against NLG evaluators on various NLG tasks in this work. We aim to study the robustness of NLG evaluators and provide insight to inspire future works on robust NLG evaluators. Our proposed framework may be used to attack online evaluation services. However, we believe that exploring this vulnerability and robustness of evaluator is more important than the above risks. Research studying attacks on evaluator will motivate improvements to the system security to defend against the attacks.

Acknowledgement

This work is partly supported by AS-TOUND (101071191- HORIZON- EIC -2021PATHFINDERCHALLENGES-01) funded by the European Commission, BEWORD (PID2021-126061OB-C43) funded by MCIN/AEI/10.13039/501100011033 and, as appropriate, by “ERDF A way of making Europe”, by the “European Union”, and the Research Grants for Young Investigators from Universidad Politécnic de Madrid (GENIUS:APOYO-JOVENES-21-TAXTYC-32K61X37) funded by Comunidad de Madrid, Shenzhen Key Laboratory of Cross-Modal Cognitive Computing funded by Shenzhen Science and Technology Innovation Program (2024) 02, Shenzhen, China.

References

- Jacopo Amidei, Paul Piwek, and Alistair Willis. 2018. [Rethinking the agreement in human evaluation tasks](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3318–3329, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Baichuan. 2023. [Baichuan 2: Open large-scale language models](#). *arXiv preprint arXiv:2309.10305*.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine transla-

- tion. In *International Conference on Learning Representations*.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023a. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*.
- Yangyi Chen, Jin Su, and Wei Wei. 2021a. [Multi-granularity textual adversarial attack with behavior cloning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4511–4526, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yi Chen, Rui Wang, Haiyun Jiang, Shuming Shi, and Ruifeng Xu. 2023b. Exploring the use of large language models for reference-free text quality evaluation: A preliminary empirical study. *arXiv preprint arXiv: Arxiv-2304.00723*.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021b. [DialogSum: A real-life scenario dialogue summarization dataset](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5062–5074, Online. Association for Computational Linguistics.
- Minhao Cheng, Jinfeng Yi, Pin-Yu Chen, Huan Zhang, and Cho-Jui Hsieh. 2020. Seq2sick: Evaluating the robustness of sequence-to-sequence models with adversarial examples. In *AAAI Conference on Artificial Intelligence*, volume 34, pages 3601–3608.
- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. 2020. [MuTual: A dataset for multi-turn dialogue reasoning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1406–1416, Online. Association for Computational Linguistics.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Javid Ebrahimi, Daniel Lowd, and Dejing Dou. 2018a. [On adversarial examples for character-level neural machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 653–663, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018b. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.
- Xiaoxue Gao, Xiaohai Tian, Rohan Kumar Das, Yi Zhou, and Haizhou Li. 2019. [Speaker-independent spectral mapping for speech-to-singing conversion](#). In *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 159–164.
- Siddhant Garg and Goutham Ramakrishnan. 2020. [BAE: BERT-based adversarial examples for text classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6174–6181, Online. Association for Computational Linguistics.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv: 2303.15056*.
- Prakhar Gupta, Cathy Jiao, Yi-Ting Yeh, Shikib Mehri, Maxine Eskenazi, and Jeffrey Bigham. 2022. [InstructDial: Improving zero and few-shot generalization in dialogue through instruction tuning](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 505–525, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xingwei He, Zhenghao Lin, Yeyun Gong, Alex Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, Weizhu Chen, et al. 2023. Annollm: Making large language models to be better crowdsourced annotators. *arXiv preprint arXiv:2303.16854*.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022a. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3905–3920, Seattle, United States. Association for Computational Linguistics.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022b. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the Second DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 161–175, Dublin, Ireland. Association for Computational Linguistics.
- Lishan Huang, Zheng Ye, Jinghui Qin, Liang Lin, and Xiaodan Liang. 2020. [GRADE: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems](#). In *Proceedings of the*

- 2020 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9230–9240, Online. Association for Computational Linguistics.
- Dongfu Jiang, Yishan Li, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, and Wenhao Chen. 2023. Tigerscore: Towards building explainable metric for all text generation tasks. *arXiv preprint arXiv:2310.00752*.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. [Is bert really robust? a strong baseline for natural language attack on text classification and entailment](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8018–8025.
- Baber Khalid and Sungjin Lee. 2022. [Explaining dialogue evaluation metrics using adversarial behavioral analysis](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5871–5883, Seattle, United States. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Thai Le, Jooyoung Lee, Kevin Yen, Yifan Hu, and Dongwon Lee. 2022. [Perturbations in the wild: Leveraging human-written text perturbations for realistic adversarial attack and defense](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2953–2965, Dublin, Ireland. Association for Computational Linguistics.
- Yibin Lei, Yu Cao, Dianqi Li, Tianyi Zhou, Meng Fang, and Mykola Pechenizkiy. 2022. [Phrase-level textual adversarial attack with label preservation](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1095–1112, Seattle, United States. Association for Computational Linguistics.
- Dianqi Li, Yizhe Zhang, Hao Peng, Liquan Chen, Chris Brockett, Ming-Ting Sun, and Bill Dolan. 2021. [Contextualized perturbation for textual adversarial attack](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5053–5069, Online. Association for Computational Linguistics.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019a. [Textbugger: Generating adversarial text against real-world applications](#). In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019b. [Textbugger: Generating adversarial text against real-world applications](#). In *Annual Network and Distributed System Security Symposium*. The Internet Society.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. [BERT-ATTACK: Adversarial attack against BERT using BERT](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6193–6202, Online. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, Jia-Chen Gu, and Chongyang Tao. 2024. Leveraging large language models for nlg evaluation: A survey. *arXiv preprint arXiv:2401.07103*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jieyu Lin, Jiajie Zou, and Nai Ding. 2021. [Using adversarial attacks to reveal the statistical bias in machine reading comprehension models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 333–342, Online. Association for Computational Linguistics.
- Yen-Ting Lin and Yun-Nung Chen. 2023. [LLM-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models](#). In *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*, pages 47–58, Toronto, Canada. Association for Computational Linguistics.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruochen Xu, and Chenguang Zhu. 2023a. [G-Eval: NLG evaluation using GPT-4 with better human alignment](#). *arXiv preprint arXiv: Arxiv-2303.16634*.
- Yuxuan Liu, Tianchi Yang, Shaohan Huang, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng,

- Feng Sun, and Qi Zhang. 2023b. Calibrating llm-based evaluator. *arXiv preprint arXiv:2309.13308*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*.
- Shikib Mehri and Maxine Eskenazi. 2020a. **Unsupervised evaluation of interactive dialog with DialoGPT**. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 225–235, 1st virtual meeting. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 2020b. **USR: An unsupervised and reference free evaluation metric for dialog generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. Generating training data with language models: Towards zero-shot language understanding. In *Advances in Neural Information Processing Systems*.
- Alireza Mohammadshahi, Thomas Scialom, Majid Yazdani, Pouya Yanki, Angela Fan, James Henderson, and Marzieh Saeidi. 2023. **RQUGE: Reference-free metric for evaluating question generation by answering the question**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6845–6867, Toronto, Canada. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. **Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- OpenAI. 2023. **GPT-4 technical report**.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Nicolas Papernot, Patrick D. McDaniel, Ananthram Swami, and Richard E. Harang. 2016. **Crafting adversarial input sequences for recurrent neural networks**. In *IEEE Military Communications Conference*, pages 49–54. IEEE.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. **A call for clarity in reporting BLEU scores**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. **SQuAD: 100,000+ questions for machine comprehension of text**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. **Generating natural language adversarial examples through probability weighted word saliency**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097, Florence, Italy. Association for Computational Linguistics.
- Ananya B. Sai, Tanay Dixit, Dev Yashpal Sheth, Sreyas Mohan, and Mitesh M. Khapra. 2021. **Perturbation CheckLists for evaluating NLG evaluation metrics**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7219–7234, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ananya B Sai, Akash Kumar Mohankumar, and Mitesh M Khapra. 2022. A survey of evaluation metrics used for nlg systems. *ACM Computing Surveys (CSUR)*, 55(2):1–39.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. **Get to the point: Summarization with pointer-generator networks**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. **BLEURT: Learning robust metrics for text generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. 2020. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3667–3676.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. **DREAM: A challenge data set and models for dialogue-based reading comprehension**. *Transactions of the Association for Computational Linguistics*, 7:217–231.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. **RUBER: An unsupervised method for automatic evaluation of open-domain dialog systems**.

- Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. **NewsQA: A machine comprehension dataset**. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200, Vancouver, Canada. Association for Computational Linguistics.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. **Universal adversarial triggers for attacking and analyzing NLP**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Boxin Wang, Hengzhi Pei, Boyuan Pan, Qian Chen, Shuohang Wang, and Bo Li. 2020. **T3: Tree-autoencoder constrained adversarial text generation for targeted attack**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6134–6150, Online. Association for Computational Linguistics.
- Xiaoqiang Wang, Bang Liu, Siliang Tang, and Lingfei Wu. 2022a. **QRelScore: Better evaluating generated questions with deeper understanding of context-aware relevance**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 562–581, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xuezhi Wang, Haohan Wang, and Diyi Yang. 2022b. **Measure and improve robustness in NLP models: A survey**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4569–4586, Seattle, United States. Association for Computational Linguistics.
- Ying Xu, Xu Zhong, Antonio Jimeno Yepes, and Jey Han Lau. 2021. **Grey-box adversarial attack and defence for sentiment classification**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4078–4087, Online. Association for Computational Linguistics.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2023. Large language models as optimizers. *arXiv preprint arXiv:2309.03409*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. **HotpotQA: A dataset for diverse, explainable multi-hop question answering**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022a. **ZeroGen: Efficient zero-shot learning via dataset generation**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11653–11669, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2022b. **ProGen: Progressive zero-shot dataset generation via in-context feedback**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3671–3683, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yi-Ting Yeh, Maxine Eskenazi, and Shikib Mehri. 2021. **A comprehensive assessment of dialog evaluation metrics**. In *The First Workshop on Evaluations and Assessments of Neural Conversation Systems*, pages 15–33, Online. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. **BARTScore: Evaluating generated text as text generation**. In *Advances in Neural Information Processing Systems*.
- Chen Zhang, Yiming Chen, Luis Fernando D’Haro, Yan Zhang, Thomas Friedrichs, Grandee Lee, and Haizhou Li. 2021a. **DynaEval: Unifying turn and dialogue level evaluation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5676–5689, Online. Association for Computational Linguistics.
- Chen Zhang, Luis Fernando D’Haro, Qiquan Zhang, Thomas Friedrichs, and Haizhou Li. 2023. **Poe: A panel of experts for generalized automatic dialogue assessment**. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:1234–1250.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. **Personalizing dialogue agents: I have a dog, do you have pets too?** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. **Bertscore: Evaluating text generation with bert**. In *International Conference on Learning Representations*.
- Xinze Zhang, Junzhe Zhang, Zhenhua Chen, and Kun He. 2021b. **Crafting adversarial examples for neural machine translation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1:*

Long Papers), pages 1967–1977, Online. Association for Computational Linguistics.

Tianyu Zhao, Divesh Lala, and Tatsuya Kawahara. 2020. [Designing precise and robust dialogue response evaluators](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 26–33, Online. Association for Computational Linguistics.

Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. [MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 563–578, Hong Kong, China. Association for Computational Linguistics.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-bench and chatbot arena](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Xiaoqing Zheng, Jiehang Zeng, Yi Zhou, Cho-Jui Hsieh, Minhao Cheng, and Xuanjing Huang. 2020. [Evaluating and enhancing the robustness of neural network-based dependency parsing models with adversarial examples](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6600–6610, Online. Association for Computational Linguistics.

Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. [Towards a unified multi-dimensional evaluator for text generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2023–2038, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A Detailed Experiment Setup

A.1 Evaluator & Generator Implementation

We select victim evaluators that are across diverse categories, and commonly employed in latest studies (Sai et al., 2022). For standard metrics like SacreBLEU, ROUGE, BLEURT, and BERTScore, we use the implementation from huggingface evaluate library⁵. For PoE, UniEval, BARTScore, TURE, RQUGE, we use the official public repositories. For open-source LLM-based evaluators like Baichuan and Vicuna, we follow the methodologies in prior studies (Gupta et al., 2022; Fu et al.,

⁵github.com/huggingface/evaluate

2023) to compute the evaluation score directly with output logits. We use latest version of these LLMs, i.e., Baichuan2 and Vicuna-V1.5. For InstructGPT, we use latest gpt-3.5-turbo-instruct-0914. Regarding the gold evaluators and adversarial generator, the PALM-2 version we use is text-bison-001. The GPT-4 version is gpt-4-1106-preview. For all evaluation score S , we normalize it to the same scale of 0 to 100.

A.2 Hyper-parameters

Note that our objective is to get responses of distinctly high quality (with scores exceeding 70) and those of distinctly low quality (with scores below 30). Accordingly, we set τ_1 as 70 for R^+ and 30 for R^- . τ_2 is set at 40 across evaluations. These thresholds are chosen to ensure that adversarial samples meet two requirements: 1) distinctly representing either extremely positive or negative labels; 2) having a clear gap in scoring between the victim and gold evaluators. The second requirement implicitly ensure incorrect classification by S_{victim} . For simplicity, α is fixed at 1 across all tasks without further tuning. We observe that this simple hyper-parameter setting has demonstrated strong performance.

B Evaluation Prompt

In this section, we list the prompt used as input to our gold evaluator (Fig. 2, 3, 4). To boost the accuracy of our evaluator, the prompt can be manually crafted or automatically crafted using existing prompt optimization methods (Yang et al., 2023; Liu et al., 2023b) to align with human judgement. In this work, the dialogue response evaluation prompt is optimized from publicly available annotated data (Zhao et al., 2020; Mehri and Eskenazi, 2020b; Huang et al., 2020). The initial instruction provided is very straightforward and consistent with our evaluation dimension mentioned in Sec. 4.1: “Given the above context and response, rate the relevance of the response to the context on a scale of 1 to 100, providing only the numerical score as your output.” The extra learned criteria can be considered as auxiliary information to support the evaluation of our targeted relevance dimension (Li et al., 2024). Similarly, We use automatically optimized summarization evaluation prompt from previous work (Liu et al., 2023b). For question evaluation, we use manually crafted prompt. Note that all gold evaluation prompts ex-

Context:
[dialogue_history]

Response:
[candidate_response]

Question:
Evaluate the context and response provided, determining how well the response sustains the dialogue, its relevance, and coherence. Using a scale of 0 to 100, with higher scores indicating a more suitable and contextually appropriate response, give a numerical score. This score should be your sole output.

Criteria:
Assign a score considering the following dimensions:

1. Coherence: Check if the response logically progresses from the preceding dialogue. It should align with both the immediate previous statement and the broader context of the conversation.
2. Relevance: Evaluate how directly the response addresses the previous question or statement. It should either answer the question or build upon the preceding statement.
3. Dialogue advancement: Determine how the response contributes to maintaining the conversation flow and promotes further interaction.
4. Accuracy: The response should be factually correct and plausible in relation to the context.
5. Use of appropriate language: The response should use language that is consistent with the dialogue, ensuring it is neither overly formal nor overly casual.

For each dimension, assign a score between 0 to 20, with a higher score indicating better performance on that dimension. Sum the scores for the final evaluation. A final score near 100 denotes a highly coherent, relevant, dialogue-advancing, accurate, and appropriately worded response, while a score near 0 implies a lack of these qualities. For responses that present a mixed performance, assign a score around 50.

Answer:

Figure 2: Example prompt for dialogue response evaluation.

Article:
[article]

Summary:
[candidate_summary]

Question:
Given the above article and summary, rate whether the summary is factually consistent with the article on a scale of 1 to 100 and just output the score.

Criteria:

- A summary is consistent with the article if it accurately and faithfully reflects the main points, and facts of the article without changing, adding, or omitting any significant information.
- A summary should avoid introducing any errors, contradictions, or distortions of the original article, unless they are explicitly marked as the summary writer’s opinions or interpretations.
- A summary should only contain information from the article itself. Any external information, even if correct, should not be included in the summary.
- All information in the summary should be clearly supported by the article.

Answer:

Figure 3: Example prompt for summarization evaluation.

licitly specify the output format as a single score number only. We take that score number as the gold score.

C Attacking Proprietary LLM

To study the possibility of using AdvEval against powerful proprietary LLMs, we further extend AdvEval to attack InstructGPT evaluator on 50 data points from DailyDialog dataset. Given the similar strong performance of available proprietary LLMs, we employ an ensemble of GPT-4 and PALM-2 as our gold evaluator. The AdvEval R^+ and R^-

ASR are 24% and 84% respectively. Compared to smaller victim models, InstructGPT shows superior robustness against R^+ perturbations, but remains vulnerable to R^- perturbations. The results suggest the potential of using our proposed AdvEval framework to study the robustness issue of cutting-edge NLG evaluators, potentially guiding future improvements in evaluator robustness.

D Additional Case Study

In Tab. 10, we present example adversarial data point from DialogSum with different attack tech-

<p>Article: [article]</p> <p>Candidate question: [candidate question]</p> <p>Provided answer: [provided_answer]</p> <p>Question: Given the above article, candidate question and provided answer, rate the quality of the candidate question on a scale of 1 to 100 based on the below criteria and just output the score number only.</p> <p>Criteria: The rating of candidate question should follow the below steps: 1. If the candidate question cannot be answered using only the information from the article, then directly return a low score. 2. If the candidate question is answerable, then verify whether the provided answer adequately addresses the candidate question. 3. In cases where the provided answer fully and accurately addresses the candidate question, award a high score. If it does not, a low score should be given. 4. If the provided answer can only partially address the candidate question, then assign a score around 50.</p> <p>Answer:</p>

Figure 4: Example prompt for question evaluation.

<i>Summarization Evaluation</i>
<p>Jo: I worked hard for a whole year. I really need a break. Phyllis: That's true. You need to take some time off to relax. Jo: You said it. I'm looking forward to my annual vacation. Phyllis: When are you going to take your vacation? Jo: Later this month. I can't wait! Phyllis: I really envy you. You know I'm not taking my vacation until December.</p>
<p>[Benign R] Jo will take an annual vacation, but Phyllis cannot take it until December.</p>
<p>[AdvEval R^+] Jo plans to take their well-deserved annual vacation later this month while Phyllis is scheduled to have hers in December. [AdvEval R^-] Jo is looking forward to her annual vacation in December but Phyllis feels like her time will be wasted since she is taking her vacation then.</p>
<p>[BERT-Attack R^+] Jo will take an biennial vacation, but Phyllis cannot take it until December. [BERT-Attack R^-] Nolan didn't fail to college only. not want to skip class tomorrow to go to the movies.</p>
<p>[TextFooler R^+] Qiu will take an annual vacation, but Phyllis cannot take it until December. [TextFooler R^-] Wes didn't gonna to school fri. Brooks wished to skip echelon tommorrow to proceed to the movies.</p>

Table 10: Adversarial data from AdvEval, BERT-Attack, and TextFooler.

niques. Similar to the findings on dialogue and question evaluation, previous attack methods struggles in the NLG evaluation context, especially for the R^+ generation, which reinforces the necessity and impact of our approach for enhancing evaluator robustness.

E Human Validation of Gold Evaluator Labels

To verify the validity of gold evaluators, we instruct PhD students who are proficient in English to annotate adversarial data produced by our AdvEval framework, as well as those from TextFooler and BERT-Attack baselines. Drawing upon manual validation scale from previous studies (Mehri and Eskenazi, 2020b; Yeh et al., 2021; Mehri and Eskenazi, 2020a), a total of 780 data points are evaluated by 9 student annotators (each data point is annotated by three students). Regarding the data to be annotated, we sample adversarial data produced by AdvEval, BERT-Attack, and TextFooler. The data sources include DailyDialog, MuTual, NewsQA, SQUAD, CNN/Dailymail, and DialogSum. The victim models include BLEURT, PoE-Large, and Vicuna-13B for dialogue response; ROUGE, and UniEval for summarization; BLEURT, RQUGE, and Baichuan-7B for question generation. For each combination, we sample 10 R^+ and 10 R^- (if reference-free).

Instead of using 100-scale annotations, We design the annotation task as a binary classification that determines whether the output text is "good" or "bad" conditioned on the corresponding contexts and the scoring criteria. The instruction manual is presented in Tab. 11. This binary classification design aligns with Eq. 3 and the hyper-parameter choices detailed in Sec. A.2. With the constraints from S_{gold} part of Eq. 3, we only get responses of distinctly high-quality or low-quality, which can

===== DIALOGUE =====
Evaluate the context of the dialogue and the response provided, and judge the overall quality of the response in maintaining the dialogue, relevance, and coherence. It's a binary classification task, please annotate in "Yes or No."
Consider the following criteria comprehensively:
Coherence: Check if the response smoothly develops from the logic of the previous conversation. It should be consistent with the previous statement and the broader context of the dialogue.
Relevance: Evaluate how the response directly addresses the previous question or statement. It should answer the question or build upon the previous statement.
Dialogue Advancement: Determine how the response helps maintain the flow of conversation and encourages further interaction.
Accuracy: The response should be factually correct and credible in the context.
Use of Appropriate Language: The response should use language consistent with the conversation, ensuring it is neither too formal nor too casual.
===== SUMMARIZATION =====
Based on the provided text and summary, evaluate whether the summary is factually consistent with the text. It's a binary classification task, please annotate in "Yes or No".
Consider the following criteria comprehensively:
If the summary accurately and faithfully reflects the main points and facts of the article, without changing, adding, or omitting any important information, then the summary is consistent with the article. The summary should avoid introducing any errors, contradictions, or distortions of the original text unless these are clearly marked as the views or interpretations of the author of the summary.
===== QG =====
Based on the given criteria, the article, and the provided answer, evaluate the question. It's a binary classification task, please annotate in "Yes or No".
Consider the following criteria comprehensively:
If the candidate question cannot be answered using only the information in the article, then judge it as "No."
If the candidate question has an answer in the article, verify whether the provided answer addresses the candidate question.
If the provided answer accurately addresses the candidate question, judge it as "Yes."
Otherwise, judge it as "No."

Table 11: Instruction manual provided to the students

be validated with binary manual annotations. Regarding the margin part of Eq. 3, we find that the

S_{victim} is much more extreme than bounds constrained by Eq. 3 on our final AdvEval adversarial data, which have an average S_{victim} of 18.1 for R^+ , and 85.4 for R^- . That extremely low/high value of S_{victim} ensures that E_{victim} makes wrong prediction on our AdvEval data as long as the binary label computed from S_{gold} is correct. Therefore, our human verification process is naturally a binary classification problem instead of a regression problem. The rationale for generating binary good/bad adversarial responses is further supported by various evaluators' training datasets, which are often constructed as positive and negative instances (binary labels) (Yeh et al., 2021).

The inter-annotator agreements are found to be 0.77, 0.77, and 0.68 for the dialogue, summarization, and question generation tasks respectively. Since achieving perfect accuracy is exceptionally challenging due to the inherent subjectivity of the task (Amidei et al., 2018; Zheng et al., 2023; Dubois et al., 2023), this inter-annotator agreements suggest good quality of the annotations. Note that the highest inter-annotator agreement in our study is 77%, which also shows the challenging nature of achieving complete alignment even among knowledgeable human annotators. Critically, the agreement rate between E_{gold} and human annotators closely mirrors this highest inter-annotator score, indicating the reliability of our gold evaluation as human proxy to guide the optimization process.