

Calibrating LLM-Based Evaluator

Yuxuan Liu^{1*}, Tianchi Yang², Shaohan Huang², Zihan Zhang²
Haizhen Huang², Furu Wei², Weiwei Deng², Feng Sun², Qi Zhang²

¹ Peking University

² Microsoft Corporation
yx.liu@stu.pku.edu.cn

Abstract

Recent advancements in large language models (LLMs) and their emergent capabilities make LLM a promising reference-free evaluator on the quality of natural language generation, and a competent alternative to human evaluation. However, hindered by the closed-source or high computational demand to host and tune, there is a lack of practice to further calibrate an off-the-shelf LLM-based evaluator towards better human alignment. In this work, we propose `AUTO-CALIBRATE`, a multi-stage, gradient-free approach to automatically calibrate and align an LLM-based evaluator toward human preference. Instead of explicitly modeling human preferences, we first implicitly encompass them within a set of human labels. Then, an initial set of scoring criteria is drafted by the language model itself, leveraging in-context learning on different few-shot examples. To further calibrate this set of criteria, we select the best performers and re-draft them with self-refinement. Our experiments on multiple text quality evaluation datasets illustrate a significant improvement in correlation with expert evaluation through calibration. Our comprehensive qualitative analysis conveys insightful intuitions and observations on the essence of effective scoring criteria.

Keywords: Large Language Models (LLM), NLP Evaluation, Human Alignment

1. Introduction

The emergence of large language models is calling on a greater focus and importance on the quality of natural language generation evaluation. With the rapid improvement of language models, their goals have gone beyond simply fitting its output to a number of given samples to a broader human alignment. Traditional evaluation metrics like BLEU (Papineni et al., 2002), ROUGE (Lin, 2004) and CIDEr (Vedantam et al., 2015) often require curated reference outputs, whose application is limited when the output space is open and diversified, and show a low correlation with human judgments (Freitag et al., 2022). While sophisticated model-based evaluators like BERTScore (Zhang* et al., 2020) and COMET (Rei et al., 2020) yield correlation improvements, their performance is still limited by the quality of references. As a result, there is a surging demand for human-aligned, reference-free evaluators for NLG evaluations.

On this front, recent lines of research works explored leveraging state-of-the-art large language models (LLMs) as reference-free evaluators on various NLG tasks (Kocmi and Federmann, 2023; Fu et al., 2023; Wang et al., 2023a; Liu et al., 2023). Being optimized to follow human instructions (Ouyang et al., 2022) as well as attaining state-of-the-art performance on language modeling (OpenAI, 2023), modern LLMs could perform the task of evaluation when prompted accordingly. Multiple evidences show that LLMs are promising

competent in evaluating instruction-tuned models like Alpaca (Taori et al., 2023) and Vicuna (Zheng et al., 2023), and being a viable alternative to human expert evaluations (Zheng et al., 2023; Dubois et al., 2023).

Despite these promising results, emerging studies are raising concerns about the validity of LLM-based evaluators - whether LLM's underlying scoring mechanism aligns with human guidelines and preferences (Liu et al., 2023). Existing LLM-based evaluators enclose the candidate text together with the evaluation task into an instruction prompt. While this paradigm succeeds in presenting the task, it elicits several unaddressed issues, including the sensitivity and bias to output space (Wang et al., 2023a), sample ordering (Wang et al., 2023b), and prompt format (Zheng et al., 2023). Plus, as the scoring prompt is also human-written, it may also incorporate potential bias to the LLM.

To address this issue, we study calibrating an LLM-based evaluator towards better human alignment. We start from a retrospection into existing LLM-based evaluators and uncover they suffer from *insufficient prompting*, where the scoring guidelines are absent and only output spaces (e.g. 0-100) are provided, resulting in inconsistent and misaligned evaluations (Lu et al., 2023). We argue that such an issue could be mitigated by elucidating the scoring criteria. And by finalizing the scoring criteria, a *consensus* could be reached between humans and the LLM, as a means of alignment.

However, it is non-trivial to obtain adequate crite-

* Work done during internship at Microsoft.

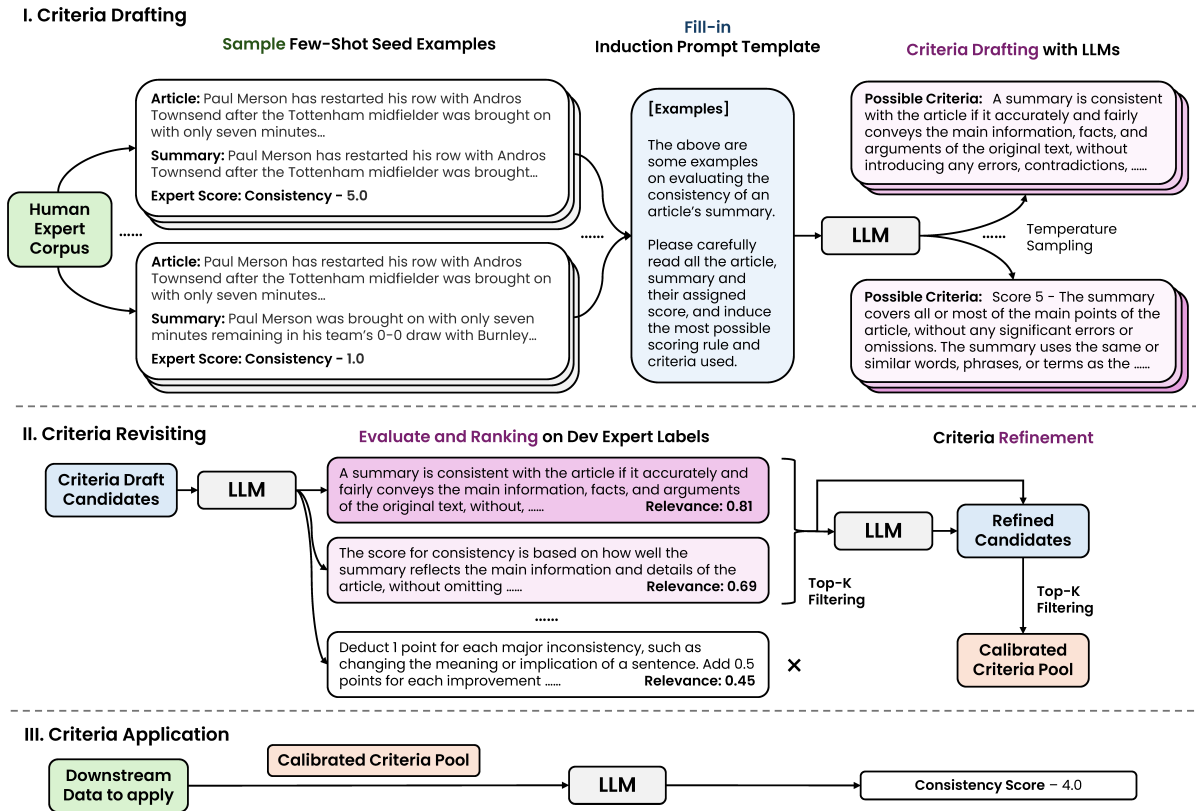


Figure 1: Overall framework of AUTO-CALIBRATE. To calibrate an LLM-based evaluator towards better human alignment, we propose a 3-stage procedure to draft, revisit, and apply high-quality scoring criteria.

ria¹, as it may require expert-level domain knowledge to assign rubrics and prevent personal bias. Drawing inspirations from the in-context learning capability (Dong et al., 2022) of LLMs, we propose AUTO-CALIBRATE, a framework to automatically align and calibrate an LLM-based evaluator through human alignment. To tackle the challenge of curating scoring criteria, we take a data-driven methodology to draft, filter, and refine rubrics using the LLM, based on human expert labels. By incorporating the mined and calibrated rubrics into scoring instructions, we obtained significant improvements in human alignment when evaluating text summarization, data-to-text generation, and hallucinations. Moreover, we release the optimal scoring criteria sets mined for the above tasks², and present detailed qualitative and quantitative analysis to uncover the essence that makes effective criteria.

¹ Results in (Chen et al., 2023) suggest that poorly curated criteria reduce relevance with human expert scoring. Uncalibrated random criteria would introduce extra bias as a misalignment between the standards used for human experts. And improperly assigned rubrics might reduce the difference between each score.

² Please refer to Appendix E.

2. Methodology

2.1. Overview of AUTO-CALIBRATE

Figure 1 illustrates the overall framework of AUTO-CALIBRATE. To calibrate an LLM-based evaluator, we focus on optimizing the evaluation prompt template \mathcal{T} applied to improve the correlation and alignment between LLM's scores and human preference. Specifically, we mine and tune the **scoring criteria** in pursuing such alignment. To express human preference, we first construct a golden set D^* , containing ground-truth sample-label pairs from human expert labeler. We then follow a novel two-stage procedure to optimize candidate scoring criteria, including drafting and revisiting. Initial criteria drafts are first inferred from in-context labels and an induction prompt, filtered on expert labels, and then refined to mitigate erroneous evaluations.

2.2. Problem formulation

In this section, we elaborate on the calibration medium and objective of AUTO-CALIBRATE - the scoring criteria. Denote D the dataset which contains multiple samples to evaluate. Based on different tasks, a sample $d_i \in D$ can contain various components: *single text*, for tasks like evaluating grammat-

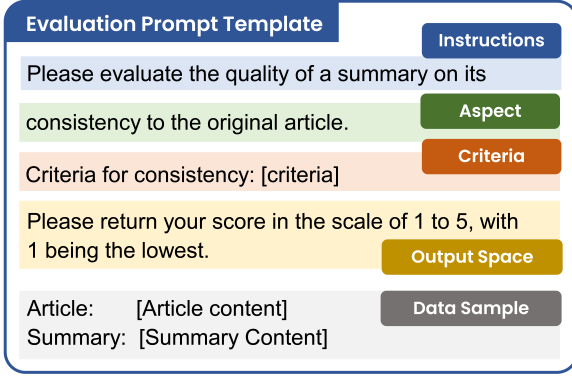


Figure 2: Example of an evaluation prompt template applied by a LLM-based evaluator.

ical correctness; *source-target*, for the vast majority of conditional generations, and *multi-turn*, like assessing multi-turn conversations.

To guide the LLM to evaluate the quality of sample d_i , prompts are applied to provide sufficient instructions and clarifications of the task. To calibrate the prompt template \mathcal{T} applied during evaluation, we regulate it by decomposing it into the following building blocks: *instructions*, *criteria*, *aspect*, *output format*, and *data sample* to evaluate, as illustrated in Figure 2. For an arbitrary sample $d_i \in D$, given a prompt template \mathcal{T} (guides the LLM to perform evaluation on NLG quality), scoring criteria \mathcal{C} , evaluation aspect a (e.g., fluency, coherence, consistency) and a large language model $LLM(\cdot)$, the NLG quality of d_i could be evaluated as

$$\hat{s}_{i,a} = LLM(\mathcal{T}(d_i, \mathcal{C}, a)). \quad (1)$$

Denote D^* a golden set consists of curated sample-label pairs $(d_i^*, s_{i,a})$ from human experts, and $f(\cdot)$ a correlation metric. In AUTO-CALIBRATE, we focus on calibrating the scoring criteria \mathcal{C} to maximize the correlation between predicted labels and human expert labels, as

$$\mathcal{C} = \arg \max_{\mathcal{C}} f \left[\cup_{d_i^* \sim D^*} (\hat{s}_{i,a}, s_{i,a}) \right]. \quad (2)$$

2.3. AUTO-CALIBRATE

Data Labeling as Human Preference To calibrate an LLM-based evaluator, one primary question is: *how* to represent and model the preference of human experts. On existing approaches, sophisticated model-based evaluators like COMET (Rei et al., 2020) directly train on human labels, while ranking-based human labels are widely adopted in RLHF to model human preference (Ouyang et al., 2022). However, these methods require extra fine-tuning, which makes them computationally intensive and impracticable to API-based LLMs. To mitigate these limitations, We implicitly encode human expert preference to a set of sample-label pairs

and form a golden set D^* . Compared with curating scoring criteria and guidelines with joint human expert works, it is more feasible to collect labels leveraging crowdsourcing dispatch, and easier to validate and merge opinions from different experts.

Criteria Drafting After constructing the expert label set D^* , we utilize the instruction following and in-context learning capability of LLMs to independently infer scoring criteria \mathcal{C} from few-shot exemplars. One crucial part here is to ensure the diversity of recalled criteria. To mitigate the label bias and position bias of in-context learning (Zhao et al., 2021), we construct various Monte-Carlo samples from D^* to obtain few-shot in-context exemplars. Given drafting prompt template \mathcal{T}_D and a few-shot exemplar set $D_s = \cup(d_i^*, s_{i,a}) \subset D^*$, an corresponding criteria is inferred as

$$\hat{\mathcal{C}} = \arg \max_{\mathcal{C}} P_{\theta}(\mathcal{C} | \mathcal{T}_D(D_s, a)), \quad (3)$$

where a denotes the evaluation aspect. Temperature sampling is also applied to draw scoring criteria in diversified presentations from the LLM. Following this procedure, we obtain the initial set of scoring criteria for evaluation and refinement. Example prompt templates are provided in Appendix D.1.

Criteria Revisiting Inferred from various few-shot exemplars, criteria within the initial draft set are diversified, but may be sub-optimal or contain potential bias (e.g., to particular scoring labels). To filter out high-quality candidates, we first revisit them leveraging D^* and select the top performing candidates w.r.t their human relevance.

To mitigate disagreements between human experts and the drafted criteria, we prompt LLMs to refine (Madaan et al., 2023) the previously generated criteria by providing them samples with strong disagreement in their scores. When refining the criteria, we suggest the following atomic editing operations via prompting to the LLM³:

- *Modification*: Adjust parts of the criteria to increase its correlation with human experts.
- *Paraphrase*: If a criteria is good enough, paraphrase it to make it clearer and more concise.
- *Adding Aspects or Details*: When LLM discovers new underlying scoring rules that are not covered by the current criteria, consider appending them to the current criteria, but make sure not to make the criteria redundant.
- *Calibrate*: Any other modifications that the LLM considers beneficial in improving alignment.

³ Detailed prompt examples and templates are provided in Appendix D.

Algorithm 1 Overall Procedure of AUTO-CALIBRATE

Require: LLM θ , human expert labels D^* , meta-correlation metric $f(\cdot)$, Monte-Carlo trial count N , in-context exemplar size $L = \{l_1, \dots, l_m\}$, aspect a , target criteria candidate pool size K

- 1: **for** Few-shot exemplar size l_i in L **do**
- 2: **for** Monte-Carlo trial j in $1 : N$ **do**
- 3: Sample few-shot examples of human expert labels $D_s = \cup(d_i^*, s_{i,a})$ from D^*
- 4: Draft candidate criteria via LLM according to Eq.(3) using temperature sampling
- 5: Add obtained criteria \mathcal{C}_i to candidate set \mathcal{C}
- 6: **end for**
- 7: **end for**
- 8: Revisit \mathcal{C} and retain top-K candidates with highest correlation: $\mathcal{C} \leftarrow \arg\text{TopK}_{c_i \in \mathcal{C}} f(c_i, D^*)$
- 9: Collect mis-aligned examples D_i^R for c_i in \mathcal{C}
- 10: **for** Candidate criteria c_i in \mathcal{C} **do**
- 11: **for** Monte-Carlo trial j in $1 : N$ **do**
- 12: Sample few-shot examples of misaligned labels $D_s^R = \cup(d_i^R, s_{i,a}^R)$ from D_i^R
- 13: Refine candidate criteria with LLM, then add obtained criteria to candidate set \mathcal{C}
- 14: **end for**
- 15: **end for**

Return $\mathcal{C}_{fin} \leftarrow \arg\max_{c_i \in \mathcal{C}} f(c_i, D^*)$

As illustrated in Figure 1, after obtaining refined candidate criteria, we first filter them with D^* and then combine them with the pre-filtered draft criteria to obtain a calibrated set of scoring rules.

Conclusion and Discussion Combining the above, we obtain AUTO-CALIBRATE, an automatic pipeline in calibrating LLM-based evaluators. The overall procedure is summarized in Algorithm 1.

The benefits of choosing criteria as a medium for calibration are multitudinous. First, we do not require gradients or access to model parameters, which makes AUTO-CALIBRATE applicable to API-based LLMs. Second, since criteria remain in a natural language form (compared with soft prompt-tuning), calibrating the criteria is essential to reaching an agreement between humans and the LLM. Therefore, the process is more explainable and controllable (e.g., one may perform *human-in-the-loop* adjustments to scoring criteria in case of preference changes, or to avoid corner cases).

3. Experimental Setup

3.1. Tasks and Datasets

We evaluate AUTO-CALIBRATE on three text quality evaluation tasks, including *text summarization*, *data-to-text generation*, and *evaluating hallucinations*. We select tasks following previous re-

searches (Zhong et al., 2022; Fu et al., 2023). We select two datasets for each task, consisting of 6 datasets in total, each containing human expert labels as golden samples. Specifically, we select NewsRoom (Grusky et al., 2018) and SummEval (Fabbri et al., 2021) for evaluating summarization; SFRES (Wen et al., 2015) and SFHOT (Wen et al., 2015) for data-to-text, QAGS-XSUM and QAGS-CNN (Wang et al., 2020a) for evaluating hallucinations. To evaluate the alignment between the scoring from LLM and human experts, we perform a meta-evaluation following (Zhong et al., 2022). Details on evaluation are listed in Appendix A.

3.2. Models and Baselines

To implement AUTO-CALIBRATE, we select OpenAI’s GPT-4 model (GPT-4-32K) as the LLM for the evaluator. We release prompt templates for criteria drafting, evaluation, and refinement for each tasks in Appendix Section D. We set the temperature to 0 during evaluation, and 1 when obtaining initial criteria drafts and their refined versions. Please refer to Appendix Section C for detailed configurations of each task.

We compare AUTO-CALIBRATE with various state-of-the-art and/or widely applied evaluators. We first include ROUGE (Lin, 2004), a widely-applied n-gram-based evaluation metric for text summarization. We then select various evaluators based on smaller neural (language) models, including BERTScore (Zhang* et al., 2020), MoverScore (Zhao et al., 2019), PRISM (Thompson and Post, 2020), BartScore (Yuan et al., 2021), CTC (Deng et al., 2021), and UniEval (Zhong et al., 2022). Finally, we compare evaluators based on state-of-the-art LLMs (e.g. GPT-3.5 and GPT-4), including GPTScore (Fu et al., 2023), ChatGPT (Wang et al., 2023a), and GPT-Eval (Liu et al., 2023).

4. Experimental Results

4.1. Evaluating Summarization

We conduct meta-correlation analysis on NewsRoom and SummEval to evaluate AUTO-CALIBRATE’s performance to calibrate an LLM-based evaluator on text summarization. Following (Liu et al., 2021), we perform summary-level Spearman and Kendall correlation analysis on each of the 4 aspects. To represent the performance of *un-calibrated* LLM, we add a GPT-4 baseline, whose evaluations are obtained through a single call using an evaluation prompt where scoring criteria is omitted⁴.

Results on NewsRoom and SummEval task are listed in Table 1 and 2, respectively. On News-

⁴ For a fair comparison, the only difference is the removal of criteria from prompt. We keep the rest identical.

Metrics	Coherence		Relevance		Informative		Fluency		Average	
	ρ	τ	ρ	τ	ρ	τ	ρ	τ	ρ	τ
ROUGE-1	0.095	0.076	0.147	0.112	0.130	0.099	0.104	0.082	0.119	0.092
ROUGE-2	0.026	0.009	0.091	0.065	0.079	0.052	0.048	0.032	0.061	0.092
ROUGE-L	0.064	0.051	0.106	0.083	0.089	0.064	0.072	0.061	0.083	0.065
BERTSCORE	0.147	0.116	0.162	0.126	0.130	0.105	0.171	0.128	0.152	0.119
MOVERSCORE	0.161	0.127	0.195	0.157	0.188	0.151	0.120	0.086	0.166	0.130
PRISM	0.573	0.478	0.553	0.460	0.561	0.472	0.532	0.443	0.555	0.463
BARTSCORE (CNN)	0.653	0.547	0.567	0.478	0.616	0.510	0.640	0.540	0.619	0.519
CHATGPT (DA)	0.469	0.405	0.461	0.392	0.578	0.498	0.507	0.427	0.504	0.430
CHATGPT (Stars)	0.428	0.375	0.402	0.348	0.557	0.487	0.451	0.385	0.460	0.399
GPT-4	0.557	0.498	0.574	0.511	0.581	0.521	0.601	0.535	0.578	0.516
AUTOCALIBRATE	0.602	0.540	0.656	0.585	0.654	0.590	0.640	0.575	0.638	0.573

Table 1: Summary-level Spearman (ρ) and Kendall (τ) correlations of aspects on NewsRoom.

Metrics	Coherence		Consistency		Fluency		Relevance		Average	
	ρ	τ	ρ	τ	ρ	τ	ρ	τ	ρ	τ
ROUGE-1	0.167	0.126	0.160	0.130	0.115	0.094	0.326	0.252	0.192	0.150
ROUGE-2	0.184	0.139	0.187	0.155	0.159	0.128	0.290	0.219	0.205	0.161
ROUGE-L	0.128	0.099	0.115	0.092	0.105	0.084	0.311	0.237	0.165	0.128
BERTSCORE	0.284	0.211	0.110	0.090	0.193	0.158	0.312	0.243	0.225	0.175
MOVERSCORE	0.159	0.118	0.157	0.127	0.129	0.105	0.318	0.244	0.191	0.148
PRISM	0.249	0.196	0.212	0.163	0.345	0.285	0.254	0.205	0.265	0.212
CTC (Consistency)	0.223	0.172	0.415	0.345	0.335	0.276	0.166	0.124	0.285	0.229
CTC (Relevance)	0.402	0.310	0.366	0.301	0.299	0.245	0.428	0.336	0.374	0.298
BARTSCORE (CNN)	0.448	0.342	0.382	0.315	0.356	0.292	0.356	0.273	0.385	0.305
UNI-EVAL (Multi-task)	0.495	0.374	0.435	0.365	0.419	0.346	0.424	0.327	0.443	0.353
UNI-EVAL (Continual)	0.575	0.442	0.446	0.371	0.449	0.371	0.426	0.325	0.474	0.377
CHATGPT (DA)	0.451	0.383	0.432	0.399	0.380	0.351	0.439	0.379	0.425	0.378
G-EVAL-3.5	0.440	0.335	0.386	0.318	0.424	0.347	0.385	0.293	0.401	0.320
G-EVAL-4	0.582	0.457	0.507	0.425	0.455	0.378	0.547	0.433	0.514	0.418
GPT-4	0.535	0.464	0.466	0.432	0.440	0.413	0.532	0.465	0.493	0.443
AUTOCALIBRATE	0.570	0.493	0.500	0.467	0.487	0.452	0.560	0.483	0.529	0.474

Table 2: Summary-level Spearman (ρ) and Kendall (τ) correlations of aspects on SummEval.

Room (Table 1), our AUTOCALIBRATE significantly outperforms the LLM-based ChatGPT evaluator. It also surpasses the vanilla GPT-4-based evaluator by a large margin (with a 10.4% improvement on Spearman and 11% on Kendall correlation), demonstrating the significance of the calibration procedure. While BartScore obtained a competent performance on NewsRoom, it falls short on SummEval. We conjecture that since it utilizes a smaller model, its consistency scoring capability might be hindered by the distribution of its fine-tuning corpus.

In contrast, our AUTOCALIBRATE demonstrated a consistent human relevance uplift on both datasets, since the pre-training knowledge in LLM is more in-depth and generalizable. On SummEval, AUTOCALIBRATE improves human correlation of GPT-4 evaluations by 7.3%, and superior to a strong baseline G-EVAL-4 that also utilizes GPT-4. Noteworthy, G-EVAL-4 requires 20 calls from LLM to obtain an average score to mitigate replicated evaluations. While this improves Spearman correlation by creating a more continuous distribution, it reduces the Kendall coefficient. In contrast, by *elucidating*

the scoring rule with calibrated criteria, AUTOCALIBRATE improves both Spearman (2.9%) and Kendall (13.4%) coefficients with only one forward call.

4.2. Evaluating Data-to-Text

We consider SFRES and SFHOT for evaluation of data-to-text generation task and follow (Fu et al., 2023) to conduct dataset-level meta-evaluation. Results are listed in Table 3. AUTOCALIBRATE significantly outperforms the most competent trained evaluator (UNI-EVAL) over 30%, and yields an over 20% and 10% improvement on Spearman correlation over GPT-SCORE (based on 175B-LLM GPT-3.5) and uncalibrated GPT-4 evaluator, respectively. These results suggest that the proposed procedures within AUTOCALIBRATE could promptly curate adequate scoring criteria for different NLG tasks and sample distributions.

4.3. Evaluating Hallucinations

Hallucination is an important issue in NLG where the output is based on fabricated facts, and it

Metrics	SFRES-INF		SFRES-NAT		SFHOT-INF		SFHOT-NAT		Average	
	ρ	τ	ρ	τ	ρ	τ	ρ	τ	ρ	τ
ROUGE-1	0.129	0.098	0.109	0.081	0.116	0.089	0.113	0.084	0.117	0.088
ROUGE-2	0.124	0.094	0.094	0.069	0.080	0.061	0.086	0.064	0.096	0.072
ROUGE-L	0.097	0.073	0.097	0.071	0.088	0.067	0.102	0.076	0.096	0.072
BERTSCORE	0.156	0.119	0.138	0.102	0.135	0.104	0.126	0.094	0.172	0.105
MOVERSCORE	0.021	-0.016	0.075	0.056	0.042	0.033	0.038	0.029	0.044	0.026
BARTSCORE (CNN)	0.154	0.117	0.138	0.101	0.164	0.126	0.075	0.055	0.133	0.100
UNIEVAL (Multi-task)	0.225	0.169	0.333	0.247	0.249	0.191	0.320	0.238	0.282	0.211
GPT-SCORE (D01)	0.270	-	0.317	-	-	-	-	-	0.294	-
GPT-SCORE (D03)	0.296	-	0.270	-	-	-	-	-	0.283	-
GPT-4	0.283	0.247	0.389	0.329	0.315	0.277	0.389	0.331	0.344	0.296
AUTOCALIBRATE	0.315	0.272	0.416	0.351	0.357	0.313	0.440	0.383	0.382	0.330

Table 3: Dataset-level Spearman (ρ) and Kendall (τ) correlations of different evaluation aspects on SFRES and SFHOT. -INF and -NAT denote informativeness and naturalness, respectively.

Metrics	QAGS-CNN		QAGS-XSUM		Average	
	ρ	τ	ρ	τ	ρ	τ
ROUGE-1	0.318	0.248	-0.049	-0.040	0.134	0.104
ROUGE-2	0.418	0.333	0.083	0.068	0.250	0.200
ROUGE-L	0.324	0.254	-0.011	-0.009	0.156	0.122
BERTSCORE	0.505	0.399	0.008	0.006	0.256	0.202
MOVERSCORE	0.347	0.271	0.044	0.036	0.195	0.153
FACTCC	0.484	0.376	0.259	0.212	0.371	0.294
BARTSCORE	0.680	0.557	0.159	0.130	0.420	0.343
CTC	0.564	0.450	0.295	0.242	0.430	0.346
UNIEVAL	0.662	0.532	0.488	0.399	0.575	0.465
G-EVAL-3.5	0.516	0.410	0.406	0.343	0.461	0.377
G-EVAL-4	0.685	0.591	0.537	0.472	0.611	0.525
GPT-4	0.649	0.606	0.637	0.637	0.643	0.622
AUTOCALIBRATE	0.744	0.663	0.662	0.662	0.703	0.663

Table 4: Dataset-level Spearman (ρ) and Kendall-Tau (τ) correlations on QAGS-CNN/XSUM.

is becoming an increasingly important topic for trustworthy LLMs (Ji et al., 2023). To test AUTOCALIBRATE on evaluating hallucinations, we select QAGS-CNN and QAGS-XSUM and perform dataset-level meta-analysis following (Liu et al., 2023). As presented in Table 4, AUTOCALIBRATE uplift the Spearman correlation by 15% over G-EVAL-4. Noteworthy, since fine-tuned on CNN data, BartScore achieves promising human relevance on QAGS-CNN, but significantly falls short on QAGS-XSUM, while LLM-based AUTOCALIBRATE performs consistently on both datasets. This further indicates that LLMs, given their immense knowledge gained during pre-training, are strong candidates for unified evaluators, and their performance could be further boosted with calibration.

4.4. Ablation Experiments

We conduct ablation studies on the procedure of AUTOCALIBRATE to better investigate the contribution of each process in calibrating LLM-based evaluator. The main ablation experiments are listed in Table 5. As illustrated in the table, removing criteria in

Dataset		ρ			τ		
		OG	-Crit	-Rfi	OG	-Crit	-Rfi
News	COH	0.602	0.557	0.593	0.540	0.498	0.531
	REL	0.656	0.574	0.619	0.585	0.511	0.550
	Room	INF	0.654	0.581	0.617	0.590	0.521
SFRES	FLU	0.640	0.601	0.628	0.575	0.535	0.563
	INF	0.315	0.283	0.300	0.272	0.247	0.264
SFHOT	NAT	0.416	0.389	0.405	0.351	0.329	0.346
	INF	0.357	0.315	0.345	0.313	0.277	0.303
QAGS	NAT	0.440	0.389	0.425	0.383	0.331	0.368
	CNN	0.744	0.649	0.724	0.663	0.606	0.642
	XSUM	0.662	0.637	0.651	0.662	0.637	0.651

Table 5: Ablations on each proposed module. We report Spearman (ρ) and Kendall (τ) correlations. ‘OG’ denotes original method, ‘-Crit’ and ‘-Rfi’ denote removing criteria and refine, respectively.

the prompt significantly reduces the human correlation of GPT-4. This corroborates our argument that previously LLMs suffered from a vaguely defined scoring principle, and this could be calibrated to increase the human alignment of LLM evaluators. The self-refine process also positively contributed to the improvements in human alignment. This indicates that LLMs could accordingly adjust the scoring criteria towards better human alignment.

5. Analysis

5.1. Essence of Effective Criteria

In this chapter, we present statistical analysis on the pool of draft candidates of scoring criteria, and mine for possible essence that contributes to effective scoring criteria for LLM with high human relevance. The main results are presented in Figure 3.

Effect of Few-Shot Example Size We study the sensitivity of AUTOCALIBRATE to the sample size of few-shot in-context samplers. As illustrated in Figure 3(A), the size of in-context few-shot exemplars

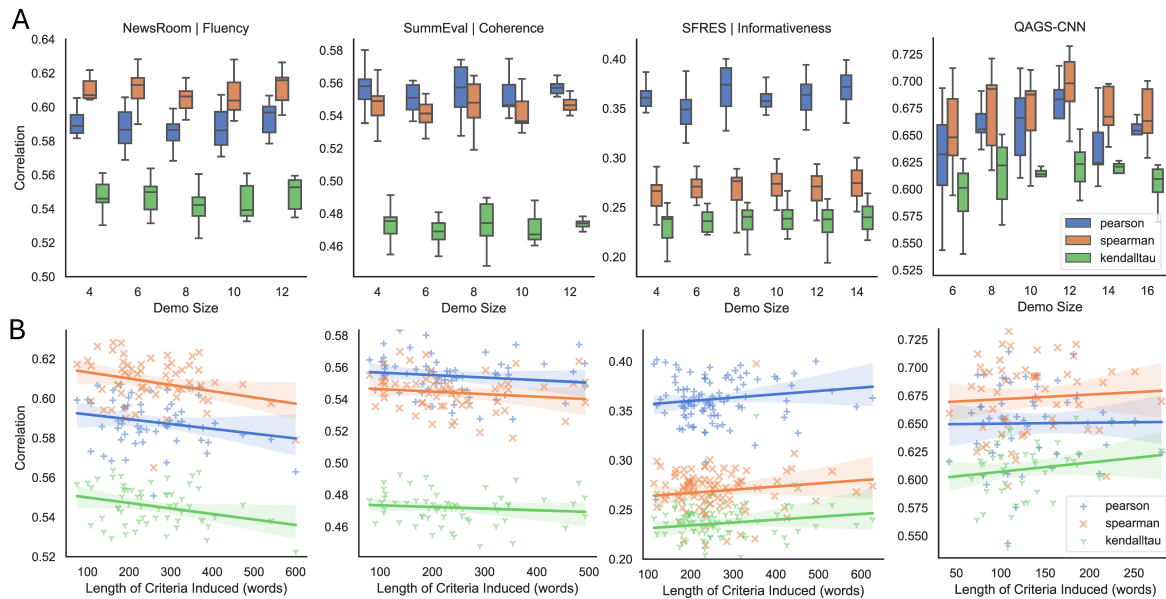


Figure 3: Statistics of criteria. A) Human relevance of criteria drafted with various in-context sizes. B) Correlation between human relevance and criteria length. Shaded areas denote 95% confidence interval.

Evaluation Criteria Induced	Relevance	
	ρ	τ
A) A summary should capture the main idea and key details of the article, without introducing any new or misleading information. A summary should use clear and concise language, avoiding unnecessary repetition or filler words. A summary should be proportionate to the length and complexity of the article, reflecting the most important aspects and leaving out less relevant details.	0.56	0.49
B) Possible scoring rule: A score of 5 means the summary is very relevant, covering all the essential elements of the article and omitting any unnecessary or misleading information. A score of 4 means the summary is mostly relevant, covering most of the essential elements of the article and omitting or including only minor or trivial information. A score of 3 means the summary is somewhat relevant, covering some of the essential elements of the article, but omitting or including some important or relevant information. A score of 2 means the summary is slightly relevant, covering only a few of the essential elements of the article, and omitting or including a lot of important or relevant information. A score of 1 means the summary is irrelevant, covering none or almost none of the essential elements of the article, and omitting or including a lot of inaccurate or irrelevant information.	0.52	0.45

Table 6: Case study on patterns of criteria induced from SummEval-REL. Criteria mined tend to follow two major patterns of its form: *holistic* (A) and *specific* (B). The former commonly describe what makes a good or bad sample, while the latter generate specific rubrics for each of the scores.

yields no significant impact except for QAGS-CNN. The results indicate that AUTO CALIBRATE is mostly robust to the size of in-context samples. Thanks to the sufficient prior knowledge obtained during pretraining by the LLM, AUTO CALIBRATE is capable of inferring the underlying criteria using only a few examples in context. As illustrated in the figure, a few-shot size of 8 to 12 is sufficient in mining effective criteria across all tasks. This intriguing feature enables a reduction in search space for cost reductions upon deployment.

Effect of Criteria Length Distribution of lengths of generated criteria and their human relevance is illustrated in Figure 3(B). Most criteria curated

with AUTO CALIBRATE lie in the range of 60 to 600 words. We discover different trends on the preference of AUTO CALIBRATE to different lengths of criteria. While fluency and coherence metrics on text summarization lean towards shorter criteria, lengthier versions are favored by the informativeness metric on data-to-text and evaluating hallucinations. Despite this difference, AUTO CALIBRATE enjoys the capability to generate effective criteria at each length. We conjecture this nuance is caused by the intrinsic complexity of the aspect to evaluate.

Patterns of Criteria We observe two significant patterns on the criteria drafted by AUTO CALIBRATE: *holistic* and *specific*. The former typically character-

Evaluation Criteria Induced	Relevance	
	ρ	τ
Before ... It should use appropriate vocabulary and punctuation, and avoid repetition or redundancy. It should also capture the tone and style of the original article. A summary with a medium score (3) should have few or minor errors that do not interfere with the overall meaning and readability of the summary. It should use mostly appropriate vocabulary and punctuation, and avoid repetition or redundancy. It should also capture the tone and style of the article. - A summary with a low score (1 or 2) should have frequent or major errors that affect the overall meaning and readability of the summary ... also fail to capture the tone and style of the original article.	0.62	0.56
After ... It should also capture the tone and style of the original article and use the correct genre and format (e.g., not writing a summary as a list of bullet points). A summary with a medium score (3) should have few or minor errors that do not interfere with the overall meaning and readability of the summary. It should use mostly appropriate vocabulary and punctuation, and minimize repetition or redundancy. It should also attempt to capture the tone and style of the original article and use the correct genre and format, but may have some inconsistencies or inaccuracies . - A summary with a low score (1 or 2) should have frequent or major errors that affect the overall meaning and readability of the summary ... also fail to capture the tone and style of the original article and use the wrong genre or format .	0.64	0.58

Table 7: Case study of criteria refinement on NewsRoom-FLU. To refine a criteria, the model automatically infer new patterns from bad cases and promptly adjust the criteria to incorporate them. Modifications are highlighted in blue, and some parts of generated criteria are omitted for space.




Figure 4: Performance of different patterns of scoring criteria induced by AUTO-CALIBRATE.

izes the common features possessed by high and low-quality samples, while the latter generates a segment of the corresponding rubric for each evaluation score (e.g., 1 to 5). A random example of these patterns of criteria is listed in Table 6. These two patterns emerge across all sets of experiments on different benchmarks. The performance distribution of these two patterns across different datasets is illustrated in Figure 4. As illustrated in the figure, there is no significant difference in human expert correlation between holistic and specific patterns, indicating that both patterns generated from AUTO-CALIBRATE are of high quality. Therefore, the performance of AUTO-CALIBRATE is robust to the patterns of criteria generated.

5.2. Case Study

To investigate the effect of criteria refinement, we present a case study in Table 7. When prompted with previous misaligned evaluation cases and instructions on modifications (Section 2.3), AUTO-CALIBRATE automatically infers new patterns of underlying scoring principles, then promptly adapts the existing criteria to accommodate them. As illus-

trated in the table, AUTO-CALIBRATE discovers that the *genre and format* is crucial to the *fluency* of summary from in-context examples provided, adjusts the criteria accordingly, and achieves higher human relevance. These findings corroborate with (Madaan et al., 2023) that LLM is capable of self-refine, and opens a future research direction on the multi-turn, iterative calibration of LLM-based evaluators.

6. Related Work

Automatic NLG Evaluation This paragraph outlines automatic evaluation metrics before the era of LLM. (1) *N-gram-based metrics*: as the most widely adopted method, n-gram-based metrics like BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) measure the quality of a candidate text by the overlap of its lexical fraction between references. Despite being widely applied, their human relevance is undesired (Freitag et al., 2022). (2) *Embedding-based metrics*: this line of method leverages a pre-trained language model (e.g. BERT (Devlin et al., 2019)) to measure the similarity between word embedding of the candidate and reference text (Zhang*

et al., 2020; Zhao et al., 2019). Their major limitation lies in the similarity-based paradigm and high dependency on the quality and diversity of references. (3) *Trained neural evaluators*: more recent research focus on specializing the PLMs by either fine-tuning on human (Rei et al., 2020) or synthetic (Zhong et al., 2022) labels, or pretraining on domain-relevant documents (Yuan et al., 2021). However, these metrics either focus on a single dimension (Wang et al., 2020b; Huang et al., 2020) or are limited in human relevance (Mehri and Eskenazi, 2020; Zhong et al., 2022).

LLM-Based NLG Evaluation With the emergence of LLM, recent research works focus on LLM-based evaluators given their promising instruction-following and generalization capability. A first line of work goes through preliminary explorations on LLM-based evaluators, including prompting methods and model variants (Fu et al., 2023; Kocmi and Federmann, 2023; Wang et al., 2023a; Chen et al., 2023; Liu et al., 2023). Successor research focuses on various aspects of improving LLM-based evaluators, including factuality (Min et al., 2023), interpretability (Lu et al., 2023), mitigating position bias (Wang et al., 2023b), and agreement to human evaluation (Zheng et al., 2023). Different from the above approaches, we focus on a general method to calibrate an off-the-shelf LLM with gradient-free approaches, to improve its alignment with human preferences on a desired NLG evaluation task.

7. Conclusion

We propose AUTO-CALIBRATE, a novel pipeline to calibrate an LLM-Based evaluator towards human alignment in a gradient-free fashion. Inferred from human expert labels and refined according to previous misalignment samples by the LLM, the criteria curated by AUTO-CALIBRATE demonstrate significant improvements in human correlation across various NLG evaluation tasks. Our qualitative analysis conveys insightful intuitions and observations on the essence of effective scoring criteria.

8. Ethical Statements

In this chapter, we discuss the limitations and potential societal impacts of AUTO-CALIBRATE.

Broader Impacts This work study on calibrating a strong LLM-based evaluator towards better human alignment. Beyond manual prompt engineering, AUTO-CALIBRATE automates the calibration process of LLM-based evaluators and provides a first experimental study on how further LLM-based evaluators could be strengthened with better prompting. We envision AUTO-CALIBRATE being potentially

applied to a wider spectrum of tasks in NLG and beyond, including LLM safety, content moderation, harm or toxic content detection, and etc.

Limitations The primary limitation is that only criteria are mined to improve alignment. After carefully analysing prompts, we conclude that the criteria are most crucial, as they are most causal to the scores given, and can be regarded as a shared consensus between humans and LLMs due to their natural language form. Plus, the criteria section is the hardest to curate compared with other parts of the prompt template (e.g., scoring scale, task definition), on which we primarily focus. Besides, A more comprehensive research on advancing and assessing other components of prompts to calibrate a LLM-based evaluator, and adapting it to wider tasks and languages is open to future work.

9. Bibliographical References

- Yi Chen, Rui Wang, Haiyun Jiang, Shuming Shi, and Ruifeng Xu. 2023. Exploring the use of large language models for reference-free text quality evaluation: A preliminary empirical study. *arXiv preprint arXiv:2304.00723*.
- Mingkai Deng, Bowen Tan, Zhengzhong Liu, Eric Xing, and Zhiting Hu. 2021. Compression, transduction, and creation: A unified framework for evaluating natural language generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7580–7605.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. AlpacaFarm: A simulation framework for methods that learn from human feedback. *arXiv preprint arXiv:2305.14387*.

- Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André FT Martins. 2022. Results of wmt22 metrics shared task: Stop using bleu—neural metrics are better and more robust. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies. In *Proceedings of NAACL-HLT*, pages 708–719.
- 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.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Tom Kocmi and Christian Federmann. 2023. Large language models are state-of-the-art evaluators of translation quality. *arXiv preprint arXiv:2302.14520*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Pengfei Liu, Jinlan Fu, Yang Xiao, Weizhe Yuan, Shuaichen Chang, Junqi Dai, Yixin Liu, Zihuiwen Ye, and Graham Neubig. 2021. Explain-able: An explainable leaderboard for nlp. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 280–289.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Qingyu Lu, Baopu Qiu, Liang Ding, Liping Xie, and Dacheng Tao. 2023. Error analysis prompting enables human-like translation evaluation in large language models: A case study on chatgpt. *arXiv preprint arXiv:2303.13809*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*.
- Shikib Mehri and Maxine Eskenazi. 2020. Ustr: 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.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. *arXiv preprint arXiv:2305.14251*.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- 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.
- 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.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. Comet: A neural framework for mt evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Brian Thompson and Matt Post. 2020. Automatic machine translation evaluation in many languages via zero-shot paraphrasing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121.

- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020a. Asking and answering questions to evaluate the factual consistency of summaries. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020b. Asking and answering questions to evaluate the factual consistency of summaries. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023a. Is chatgpt a good nlg evaluator? a preliminary study. *arXiv preprint arXiv:2303.04048*.
- Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023b. Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277.
- 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*.
- 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.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*.
- 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.

A. Evaluation Strategy

In this section, we introduce meta-evaluation strategies for assessing human alignment that are applied in this work. We select evaluation strategies primarily following previous works (Zhong et al., 2022; Fu et al., 2023; Liu et al., 2023). Given a dataset \mathcal{D} consisting of NLG samples from M diverse systems and J source text samples, evaluation metric $f(\cdot)$ (e.g., BLEU (Papineni et al., 2002)) and correlation metric $g(\cdot)$, we could perform meta-evaluation at either sample or dataset level.

Sample Level For sample-level meta-evaluation, we first compute correlation values on multiple candidate response (from each system) to a individual sample, then average across all samples:

$$f_{sample} = \frac{1}{J} \sum_{i=1}^J (g([\hat{s}_{i,1}, \dots, \hat{s}_{i,M}], [s_{i,1}, \dots, s_{i,M}])), \quad (4)$$

where $\hat{s}_{u,v}$ and $s_{u,v}$ denote the evaluation results (if not, converted to a numeric value) for the v -th response to u -th sample from evaluator $f(\cdot)$ and human experts, respectively.

Dataset Level For dataset-level meta-evaluation, we evaluate the correlations on all samples in the dataset (with a total of $M \times J$ samples), as follows:

$$f_{dataset} = g([\hat{s}_{i,1}, \dots, \hat{s}_{J,M}], [s_{i,1}, \dots, s_{J,M}]). \quad (5)$$

B. On Performance of Adding Chain-of-Thoughts

Chain-of-thought (CoT) (Wei et al., 2022) prompting elicits reasoning in large language models by encouraging models to generate their rationales before obtaining an answer. As studied in recent research (Liu et al., 2023), chain-of-thoughts are beneficial to improving human alignment in NLG evaluation, if incorporated in the scoring prompt template. Therefore, we study whether AUTO-CALIBRATE could further benefit from adding a CoT into our calibrated scoring prompts.

To obtain the CoT for each scoring aspect, we follow (Liu et al., 2023), and results are illustrated in Table 8. As shown in the figure, adding CoTs to our calibrated prompts yields negligible difference. We conjecture the effectiveness of ‘CoT’ is marginalized by providing informative and instructive scoring criteria. In contrast to math, the assessment of text quality is not a strictly chained reasoning process, so providing a CoT is essentially clarifying the evaluation rubrics, which is consistent with the meaning of the criteria in this paper, and thus obtained no additional benefit. Plausibly, the ‘CoT’s here act to

elucidate the scoring rules, rather than providing reasoning paths to follow.

C. Configuration Details

In this section, we list the configuration details of AUTO-CALIBRATE for each experiments. Detailed configurations for AUTO-CALIBRATE are listed in Table 9. We apply the same set of configurations to each of the two datasets within a task.

D. List of Prompt Templates

In this section, we list prompt templates applied throughout this study, including induction templates for criteria drafting, evaluation templates that utilize the generated scoring criteria, and templates for self-refinement of criteria.

D.1. Criteria Drafting Templates

Prompt templates for criteria drafting are listed in Figure 5, 6 and 7. The [Aspect] denote placeholders for aspects to evaluate (e.g. coherence, consistency, etc.), and sampled few-shot in-context exemplars are placed at [In-Context Few-Shot Samples], including samples and their expert scores.

D.2. Evaluation Templates

Prompt templates for evaluation are listed in Figure 8, 9 and 10. The [Aspect] denotes placeholders for aspects to evaluate (e.g. coherence, consistency, etc.). Evaluation samples and calibrated scoring criteria for each aspect are filled into corresponding placeholders during evaluation.

D.3. Criteria Refinement Templates

An example prompt template for criteria refinement can be found in Figure 11. As illustrated in the figure, we first fill in the aspect and tasks to the instructions, then prompt the LLM with the previous criteria, few-shot in-context samples of misaligned evaluations, together with suggested means of modifications to obtain a modified version of scoring criteria for this task.

E. Extended Case Study

E.1. List of Criteria

In this section, we present a case study on scoring criteria generated by AUTO-CALIBRATE for each evaluation aspect of each benchmark throughout this study in Table 10, 11, 12 and 13. As illustrated in the tables, scoring criteria generated with AUTO-CALIBRATE are informative, covering significant rubrics to evaluate a given aspect of the target NLG task.

Metrics	Coherence		Consistency		Fluency		Relevance		Average	
	ρ	τ	ρ	τ	ρ	τ	ρ	τ	ρ	τ
w/ Auto CoT	0.550	0.477	0.495	0.461	0.482	0.447	0.564	0.492	0.523	0.469
w/o Auto CoT	0.550	0.478	0.500	0.467	0.480	0.447	0.560	0.483	0.523	0.467

Table 8: Performance comparison of w/ and w/o CoT in base prompt template on SummEval.

Task	Summarization	Data-to-text	Hallucination
Model	GPT-4-32K	GPT-4-32K	GPT-4-32K
Evaluation Temperature	0.0	0.0	0.0
Max Tokens	20	20	20
Criteria Drafting Temperature	1.0	1.0	1.0
In-context Sample Size	4,6,8,10,12	4,6,8,10,12,14	6,8,10,12,14,16
Monte-Carlo Trials	4	4	3
Temperature Sampling Count	3	3	3
Max Tokens	768	768	768
Criteria Refining Temperature	1.0	1.0	1.0
In-context Sample Size	1,2,4	1,2,4	1,2,4
Monte-Carlo Trials	4	4	4
Temperature Sampling Count	2	2	2
Max Tokens	768	768	768

Table 9: Detailed configurations of AUTO-CALIBRATE for different experiments.

```

## Instructions
Please infer the scoring criteria for the following task:
[Score the following summary of a news article on its [Aspect]. Please return your score on how the summary
is consistent with the article in the scale of 1 to 5, with 1 being the lowest.]

- The following is some examples on evaluation scores of [Aspect] of summary (in the range of 1 to 5, where 1
being the lowest).
- Please carefully read all the article, summary and their assigned score, and induce the most possible scoring
rule and criteria used.
- It is optimal that, by using the induced criteria, you are very likely to assign a same score on [Aspect] to the
provided reference scores.

## Criteria for [Aspect]
- The scoring criteria been used. Now it is not explicitly provided, and you should induce it from the following
samples.
- The induced criteria should be able to explain the scores of all the samples provided, being generic and
concise.

## Examples
[In-Context Few-Shot Samples]

## Induced Criteria
Criteria for [Aspect]:

```

Figure 5: Prompt template for criteria drafting on text summarization (SummEval, NewsRoom).

Instructions

Please infer the scoring criteria for the following task:

[Task data-to-text is to generate natural language sentences from structured data sources. This can be useful for creating chatbots, voice assistants, or text summarizers. Please score the following natural language sentence generated according to a structured data expression. Please return your score on **[Aspect]** of the sentence, in the scale of 1 to 6, with 1 being the lowest.]

- The following is some examples on evaluation of **[Aspect]** of the natural language sentence generated from structured data expression (in the range of 1 to 6, where 1 being the lowest).
- Please carefully read all expressions, generated sentence and its assigned score, and induce the most possible scoring rule and criteria used.
- It is optimal that, by using the same criteria, you are very likely to assign a same score to the provided reference scores.

Criteria for **[Aspect]**

- The scoring criteria been used. Now it is not explicitly provided, and you should induce it from the following samples.
- The induced criteria should be able to explain the scores of all the samples provided, being generic and concise.

Examples

[\[In-Context Few-Shot Sampls\]](#)

Induced Criteria

Criteria for **[Aspect]**:

Figure 6: Prompt template for criteria drafting on data-to-text (SFRES, SFHOT).

Instructions

Please infer the scoring criteria for the following task:

[Evaluate the **factual consistency** of the summary to the article. Check how well the summary is supported by the article and whether it contains untruthful or misleading facts. Score 1 if the summary is factually consistent with the article, 0 otherwise.]

- The following is some examples on evaluation of **factual consistency** of generated summary to the article.
- Please carefully read all summary - article pairs and its assigned score, and induce the most possible scoring rule and criteria used.
- It is optimal that, by using the same criteria, you are very likely to assign a same score to the provided reference scores.

Criteria for **factual consistency**

- The scoring criteria been used. Now it is not explicitly provided, and you should induce it from the following samples.
- The induced criteria should be able to explain the scores of all the samples provided, being generic and concise.

Examples

[\[In-Context Few-Shot Samples\]](#)

Induced Criteria

Criteria for **factual consistency**:

Figure 7: Prompt template for criteria drafting on evaluating hallucinations (QAGS-XSUM/CNN).

Instructions
Score the following summary of a news article on its [Aspect].
Please return your score on how the summary is [Aspect] with the article in the scale of 1 to 5, with 1 being the lowest.

Example
[Article and Summary to be evaluated]

Criteria for [Aspect]
[Calibrated criteria for evaluating this aspect]

Evaluation
Now, please evaluate how [Aspect] is the summary to the article (on a scale of 1-5, with 1 being the lowest). Please carefully read the article and summary, and follow the scoring criteria above to score the [Aspect] of the summary to the article.
Please first return your score, and then provide your reasoning for the score.

[Aspect] Score (1-5):

Figure 8: Prompt template for evaluation on text summarization (SummEval, NewsRoom).

Instructions
Please score on the [Aspect] of a following natural language sentence generated according to a structured data expression.
Please return your score on how [Aspect] is the sentence, in the scale of 1 to 6, with 1 being the lowest.

Example
[Data expression and text to be evaluated]

Criteria for [Aspect]
[Calibrated criteria for evaluating this aspect]

Evaluation
Now, please evaluate how [Aspect] is the generated sentence. (on a scale of 1-6, with 1 being the lowest) Please carefully read the sentence and the structured data expression, and follow the scoring criteria to score the [Aspect] of the sentence.
Please first return your score, and then provide your reasoning for the score.

[Aspect] Score (1-5):

Figure 9: Prompt template for evaluation on data-to-text (SFRES, SFHOT).

Is the sentence supported by the article?
Answer 1 if the summary is factually consistent with the article, 0 otherwise.

[Article and Summary to be evaluated]

[Calibrated criteria for evaluating hallucination]

Answer:

Figure 10: Prompt template for evaluation on evaluating hallucinations (QAGS-XSUM/CNN).

Please refine and improve a scoring criteria used by a large language model in evaluating the [Aspect] of [Task].

Large language models (LLMs) are powerful neural models that can evaluate the quality of [Task]. However, LLMs may not always agree with human judgments. Please refine the scoring criteria used by LLMs to improve its correlation with human expert scores.

To refine the scoring criteria used by the LLM in evaluating the [Aspect] of [Task], please follow the following instructions step-by-step:

1. Carefully read each example, understand each [Source acronym (e.g. article)] and its corresponding [Target acronym (e.g. summary)], and get your initial assessment of its quality on [Aspect].

2. Compare the test score obtained by the LLM according to the criteria and the ground-truth score from human experts. Please think why the correlation is limited by using the current criteria, and how can you improve the criteria to increase the correlation between LLM's score and human expert score. If there is a small gap or no gap, this means the criteria work well in this case.

3. Read all of the test cases and rethink how you could refine the current criteria based on your observations and analysis. Then, refine the criteria to make it concise, accurate, and consistent with human judgments. When refining the criteria, you can do the following: 1) modification: adjust some parts of the criteria to increase its correlation with the scoring criteria that you think might used by human experts; 2) paraphrase: if the criteria is good enough, you can consider paraphrasing it to make more concise and easy to understand; 3) adding aspects or details: if you find some new underlying scoring rules not covered by the current criteria, consider adding them as a new line of injecting to current criteria, but make sure not to make the criteria too long and redundant; 4) calibrate: you can take other methods you think being helpful to improve the correlation with human experts.

Please return only your refined criteria without any additional sentences.

Old criteria: [Previous Criteria Drafts]

Examples: [In-Context Few-Shot Samples]

Figure 11: Prompt template for criteria refinement with GPT-4.

Aspect	Example Scoring Criteria
COH	<p>- The summary should be a concise and accurate representation of the main points and ideas of the article. It should avoid including irrelevant or minor details that are not essential to the article's purpose or message. - The summary should have a clear and logical structure that follows the article's original order of information, or provides a coherent alternative order if it improves the summarization. The summary should use transitions and connectors to link the sentences and paragraphs coherently. - The summary should use the same or similar terminology and tone as the article, unless there is a need to simplify or clarify some terms for the intended audience. The summary should avoid introducing new or unfamiliar words or concepts that are not in the article or relevant to the summary. - The summary should maintain the same perspective and point of view as the article, unless there is a reason to shift or contrast it. The summary should not express the summarizer's own opinion, interpretation, or evaluation of the article, unless it is explicitly stated as such. - The summary should be grammatically correct and free of spelling, punctuation, and capitalization errors. The summary should use direct or indirect quotations and citations appropriately to acknowledge the source of the article. - The summary should be coherent and consistent with the article's topic and genre. The summary should avoid introducing information or claims that contradict or deviate from the article's main message. The summary should also avoid repeating information or using unnecessary filler words.</p>
INF	<p>- A summary should convey the main idea and the most important details of the original article in a concise and clear way. - A summary should avoid repeating information that has already been mentioned or is irrelevant to the article's main point. - A summary should use accurate and specific words to describe the article's content, and avoid vague or ambiguous expressions. - A summary should maintain the same tone and perspective as the original article, and avoid adding personal opinions or interpretations. - A summary should follow the logical order and structure of the original article, and use transition words or phrases to connect sentences if needed. - Possible scoring scheme: - 5: The summary meets all the criteria and has no significant flaws or errors. - 4: The summary meets most of the criteria and has minor flaws or errors that do not affect the overall comprehension. - 3: The summary meets some of the criteria and has moderate flaws or errors that affect the comprehension of some parts. - 2: The summary meets few of the criteria and has major flaws or errors that affect the comprehension of most parts. - 1: The summary meets none of the criteria and has severe flaws or errors that make it incomprehensible.</p>
FLU	<p>- The score for fluency should reflect how well the summary is written in clear, natural and grammatical language, without unnecessary or confusing repetitions, contradictions or omissions. - A summary with a score of 5 should be fluent, coherent and engaging, with no errors or awkward expressions. It should use appropriate vocabulary, syntax and punctuation, and convey the main information and meaning of the article. - A summary with a score of 4 should be mostly fluent, coherent and engaging, with minor errors or awkward expressions that do not affect the overall meaning. It should use mostly appropriate vocabulary, syntax and punctuation, and convey most of the main information and meaning of the article. - A summary with a score of 3 should be somewhat fluent, coherent and engaging, but with some errors or awkward expressions that may affect the overall meaning or readability of the summary. It should use some appropriate vocabulary, syntax and punctuation, and convey some of the main information and meaning of the article, but may have some gaps or inaccuracies. - A summary with a score of 2 should be poorly fluent, coherent and engaging, with frequent errors or awkward expressions that significantly affect the overall meaning or readability of the summary. It should use limited or inappropriate vocabulary, syntax and punctuation, and convey little of the main information and meaning of the article, or may have some major distortions or misunderstandings. - A summary with a score of 1 should be very poorly fluent, coherent and engaging, with severe errors or awkward expressions that make the summary incomprehensible or unintelligible. It should use very limited or inappropriate vocabulary, syntax and punctuation, and convey none of the main information and meaning of the article, or may have some nonsensical or irrelevant content.</p>
REL	<p>- The summary should capture the main topic, events, and outcomes of the article in a concise and accurate way. - The summary should not omit any essential information that is necessary to understand the article's purpose and significance. - The summary should not include any irrelevant or redundant details that distract from the article's main points or introduce confusion. - The summary should use the same or similar terminology and tone as the article, unless the article uses obscure or jargon words that need to be simplified. - The summary should reflect the article's structure and organization, presenting the information in a logical and coherent order. Examples of scoring: - Score 5: The summary meets all the criteria for relevance and provides a clear and comprehensive overview of the article, without any errors or gaps. - Score 4: The summary meets most of the criteria for relevance and provides a mostly clear and comprehensive overview of the article, but may have some minor errors or gaps, such as missing a minor detail, using a slightly different word, or omitting a transition. - Score 3: The summary meets some of the criteria for relevance and provides a partially clear and comprehensive overview of the article, but has some noticeable errors or gaps, such as missing a key detail, using a vague or inaccurate word, or skipping a logical connection. - Score 2: The summary meets few of the criteria for relevance and provides a vaguely clear and comprehensive overview of the article, but has many errors or gaps, such as missing several important details, using inappropriate or misleading words, or presenting the information in a confusing or contradictory order. - Score 1: The summary meets none or almost none of the criteria for relevance and provides a unclear and incomplete overview of the article, with severe errors or gaps, such as missing the main topic, using incorrect or irrelevant words, or omitting the entire conclusion.</p>

Table 10: **Case study on criteria** on each aspect for NewsRoom generated by AUTO CALIBRATE.

Aspect	Example Scoring Criteria
COH	Coherence is the quality of being consistent, logical, and well-organized in the summary. A summary is coherent if it accurately captures the main ideas and key information from the article, and presents them in a clear and concise manner. A summary is not coherent if it omits important details, contradicts the article, or introduces irrelevant or confusing information. The score for coherence is based on the following scale: - 5: The summary is very coherent, with no errors or flaws. - 4: The summary is mostly coherent, with only minor errors or gaps. - 3: The summary is somewhat coherent, but has some significant errors or omissions. - 2: The summary is poorly coherent, with many errors, inconsistencies, or redundancies. - 1: The summary is not coherent at all, with little or no relation to the article.
CON	- A summary is consistent with the article if it accurately and faithfully reflects the main points, facts, and tone 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 use clear and precise language that matches the style and genre of the article, and avoid any vague or ambiguous expressions that could mislead the reader or obscure the meaning of the article. - A summary should maintain the logical structure and coherence of the article, and present the information in a well-organized and easy-to-follow manner. - A summary should be concise and avoid any unnecessary or redundant details that do not contribute to the main purpose or message of the article.
FLU	- A fluent summary should reflect the main content and structure of the original article, using clear and coherent language that avoids redundancy and errors. - A fluent summary should retain the key information and details from the article, without introducing any irrelevant or inaccurate information that distorts the meaning of the original text. - A fluent summary should use appropriate transition words, connectors, and referents to ensure the logical flow and cohesion of the summary, and avoid abrupt or confusing shifts in topic or perspective. - A fluent summary should use varied and precise vocabulary and grammar that suits the tone and style of the article, and avoid repetition or ambiguity. - A fluent summary should use correct spelling, punctuation, and capitalization throughout the summary, and follow the conventions of standard written English. A possible scoring rubric based on these criteria is: - 5: The summary is fluent and meets all the criteria listed above. It captures the main points and details of the article accurately and effectively, using clear and coherent language that follows the logical structure of the article. The summary uses appropriate transition words, connectors, and referents to ensure cohesion, and varied and precise vocabulary and grammar that suits the tone and style of the article. The summary has no or minimal errors in spelling, punctuation, and capitalization. - 4: The summary is mostly fluent and meets most of the criteria listed above. It captures the main points and details of the article fairly well, using mostly clear and coherent language that follows the logical structure of the article. The summary uses mostly appropriate transition words, connectors, and referents to ensure cohesion, and mostly varied and precise vocabulary and grammar that suits the tone and style of the article. The summary has few errors in spelling, punctuation, and capitalization. - 3: The summary is somewhat fluent and meets some of the criteria listed above. It captures some of the main points and details of the article, but may omit or misrepresent some important information. The summary uses somewhat clear and coherent language, but may deviate from the logical structure of the article or have some lapses in cohesion. The summary uses some appropriate transition words, connectors, and referents, but may also have some inappropriate or confusing ones. The summary uses some varied and precise vocabulary and grammar, but may also have some repetition or ambiguity. The summary has several errors in spelling, punctuation, and capitalization. - 2: The summary is not very fluent and meets few of the criteria listed above. It captures few of the main points and details of the article, and may omit or misrepresent many important information. The summary uses unclear or incoherent language, and does not follow the logical structure of the article or have much cohesion. The summary uses few or no appropriate transition words, connectors, and referents, and may have many inappropriate or confusing ones. The summary uses limited or imprecise vocabulary and grammar, and may have many repetition or ambiguity. The summary has many errors in spelling, punctuation, and capitalization. - 1: The summary is not fluent and meets none of the criteria listed above. It captures none or almost none of the main points and details of the article, and may omit or misrepresent most or all of the important information. The summary uses incomprehensible or irrelevant language, and does not follow the logical structure of the article or have any cohesion. The summary uses no or almost no appropriate transition words, connectors, and referents, and may have only inappropriate or confusing ones. The summary uses very limited or inaccurate vocabulary and grammar, and may have only repetition or ambiguity. The summary has numerous and severe errors in spelling, punctuation, and capitalization.
REL	- A summary is relevant if it captures the main points or the most important information from the article, without leaving out any crucial details or adding any unnecessary or inaccurate ones. - A summary is more relevant if it uses the same or similar terms and expressions as the article, as long as they are clear and concise. - A summary is less relevant if it omits or misrepresents some of the key facts or arguments from the article, or if it introduces irrelevant or erroneous information that is not supported by the article. - A summary is irrelevant if it does not correspond to the article at all, or if it only mentions a minor or peripheral aspect of the article.

Table 11: Case study on criteria on each aspect for SummEval generated by AUTOCALIBRATE.

Aspect	Example Scoring Criteria
INF	- A natural language sentence is informative if it conveys all the relevant information from the data expression, without omitting, adding, or distorting any facts. - A sentence is more informative if it uses clear and natural language, without grammatical errors, ambiguity, or redundancy. - A sentence is less informative if it leaves out some information from the data expression, or if it uses vague, unnatural, or incorrect language. - A possible scoring rule for informativeness is as follows: - 6: The sentence conveys all the information from the data expression, using clear and natural language. - 5.5: The sentence conveys all the information from the data expression, using mostly clear and natural language, but with minor issues (e.g., word choice, punctuation, etc.) - 5: The sentence conveys all the information from the data expression, but with some issues in language clarity or naturalness. - 4.5: The sentence conveys most of the information from the data expression, using clear and natural language, but omitting one detail. - 4: The sentence conveys most of the information from the data expression, but with some issues in language clarity or naturalness, or omitting more than one detail. - 3: The sentence conveys some of the information from the data expression, but with significant issues in language clarity or naturalness, or omitting several details. - 2: The sentence conveys little of the information from the data expression, or with major issues in language clarity or naturalness, or adding or distorting facts. - 1: The sentence conveys none of the information from the data expression, or with unintelligible or irrelevant language.
NAT	- A natural language sentence is natural if it is fluent, coherent, grammatical, and conveys the meaning of the data expression accurately and concisely. - The score of naturalness ranges from 1 to 6, where 1 is the lowest and 6 is the highest. - The score is assigned based on the following criteria: - A sentence that is completely natural, without any errors or awkwardness, and expresses the data expression fully and succinctly, gets a 6. - A sentence that is mostly natural, with minor errors or redundancy, and expresses the data expression adequately, gets a 5 or 5.5. - A sentence that is somewhat natural, with noticeable errors or incompleteness, and expresses the data expression partially or vaguely, gets a 4 or 4.5. - A sentence that is barely natural, with serious errors or confusion, and expresses the data expression incorrectly or irrelevantly, gets a 3 or 3.5. - A sentence that is not natural at all, with unacceptable errors or nonsense, and does not express the data expression at all, gets a 1 or 2.

Table 12: Case study on criteria on each aspect for SFRES generated by AUTOCALIBRATE.

Aspect	Example Scoring Criteria
FACT	- Score 1 if the sentence accurately and concisely summarizes the main facts and information from the article, without omitting, distorting, or adding any significant details. - Score 0.75 if the sentence summarizes the main facts and information from the article, but has some minor issues such as: omitting, distorting, or adding some less important details; using vague or imprecise language; or being too long or verbose. - Score 0.5 if the sentence captures some of the facts and information from the article, but has some major issues such as: omitting, distorting, or adding some important details; using incorrect or misleading language; or being too short or incomplete. - Score 0.25 if the sentence only captures a few of the facts and information from the article, and has many issues such as: omitting, distorting, or adding most of the details; using irrelevant or contradictory language; or being too general or specific.- Score 0 if the sentence does not capture any of the facts and information from the article, or contradicts or misrepresents the article entirely.

Table 13: **Case study on criteria** on each aspect for QAGS-CNN generated by AUTOCALIBRATE.