

BatchEval: Towards Human-like Text Evaluation

Anonymous ACL submission

Abstract

Significant progress has been made in automatic text evaluation with the introduction of large language models (LLMs) as evaluators. However, current sample-wise evaluation paradigm suffers from the following issues: (1) Sensitive to prompt design; (2) Poor resistance to context noise; (3) Inferior ensemble performance with static reference. Inspired by the fact that humans treat both criterion definition and inter sample comparison as references for evaluation, we propose **BATCHEVAL**, a paradigm that conducts batch-wise evaluation iteratively to alleviate the above problems. We explore variants under this paradigm and confirm the optimal settings are two stage procedure with heterogeneous batch composition strategy and decimal scoring format. Comprehensive experiments across 3 LLMs on 4 text evaluation tasks demonstrate that **BATCHEVAL** outperforms state-of-the-art methods by 10.5% on Pearson correlations with only 64% API cost on average. Further analyses have verified the robustness, generalization, and working mechanism of **BATCHEVAL**¹.

1 Introduction

Accurately evaluating the text quality in specific criterion (e.g., coherence) can facilitate better understanding, application, and development of large language models (LLMs), which becomes more crucial with their recent rapid progress in text generation capabilities (OpenAI, 2023). Due to the labor-intensive and time-consuming nature of human evaluation, early works have explored automatic evaluation methods, which can be categorized into rule-based (Papineni et al., 2002; Lavie and Denkowski, 2009), embedding-based (Forgues et al., 2014; Zhang et al., 2020), and learning-based (Mehri and Eskénazi, 2020; Zhang et al., 2022) approaches. Continuous progress has been achieved




Figure 1: Both humans and LLM-based evaluators assess text based on criterion definition, but humans further conduct sample comparison for better evaluation.

through these methods, but there remains a significant gap in their consistency with human judgments (Sai et al., 2023).

Recently, the revolutionary power of LLMs has been applied across various fields, demonstrating performance that is even on par with humans (OpenAI, 2023; Guo et al., 2023a). In text evaluation field, LLM-based evaluators (Chiang and Lee, 2023a; Liu et al., 2023; Guo et al., 2023b; Chiang and Lee, 2023b) have also made significant progress compared to traditional methods, but they still lag behind human evaluators. We carefully compare their working procedures and find that the difference in evaluation references might be the reason for the performance disparity (Figure 1). Human evaluators analyze samples based on the criterion definition and provide discriminative scores through comparison between samples. However, LLM-based evaluators assess each sample individually, thus only having criterion as a reference.

We analyze that current sample-wise evaluation paradigm will face problems on three aspects: (1) *Robustness against prompt design?* Since criterion is the sole reference for evaluation, minor changes of prompt may significantly affect the evaluation results (See §4.4 for empirical validation). (2) *Robustness against noise?* Due to the absence of comparison between samples, the evaluation scores

¹Our code and data have been made public on the internet.

068	lack discrimination and exhibit a non-uniform distribution (See Figure 3), which can lead to reduced robustness against noise like random deletion or synonym substitution on samples (See Theorem 1).	119
069		120
070		
071	(3) <i>Better performance under ensemble?</i> Current LLM-based evaluators average scores from multiple generations as the final rating for given sample. However, generating multiple times from the static reference (criterion) induces a lack of diversity among scores (Figure 4), which can weaken the effect of ensemble according to Theorem 2.	
072		
073		
074		
075		
076		
077		
078		
079	To address the aforementioned problems, we propose BATCHEVAL, a new LLM-based text evaluation paradigm that assesses samples batch-wise, akin to the way of humans. Overall, BATCHEVAL iterates an allocation process where all samples are first split into batches, and then each batch is compiled into a prompt as the input of LLMs. By introducing in-batch samples as an additional reference apart from criterion, the orthogonal and complementary references can not only reduce the dependency on prompt design but also enhance the discrimination of scores between samples through in-batch comparison, leading to improved robustness against noise. Furthermore, the iteratively changing batch composition can provide LLMs with varying evaluation references, thereby enhancing diversity and the ensemble performance.	
080		
081		
082		
083		
084		
085		
086		
087		
088		
089		
090		
091		
092		
093		
094		
095		
096	While the idea of BATCHEVAL is simple, there are many ways it can be realized. We explored variants in evaluation procedure, format of scoring and composition of batch. Some of them work surprisingly well while some do not meet expectations. Experiments and analyses confirm that separate analyzing and scoring evaluation procedure, decimal scoring format, and quality-heterogeneous batch composition strategy yield the optimal results.	
097		
098		
099		
100		
101		
102		
103		
104		
105	We conduct extensive experiments on 4 text evaluation tasks primarily with GPT-4: turn-level response, dialogue, text summarization, and story generation. By allowing in-batch samples to share single prompt and applying a small iteration rounds, BATCHEVAL outperforms best performing LLM-based evaluators by a significant margin (10.5%) in terms of correlation with human evaluations, while incurring only 64% of API costs. We also validate the generalization of BATCHEVAL on more LLMs, robustness to prompt design and noise, and analyze the choice of hyperparameters through further experiments. Finally, we probe into the working mechanism of BATCHEVAL through attention analysis on Llama-2-70b-chat-hf. Our contributions are summarized as follows:	119
106		120
107		
108		
109		
110		
111		
112		
113		
114		
115		
116		
117		
118		
	1. We analyzed how the sample-wise evaluation paradigm of LLM-based evaluators, differing from human evaluators, limited their robustness and consistency with human judgment.	121
		122
		123
		124
	2. We proposed BATCHEVAL, a new paradigm that evaluates texts batch-wise, and experimentally validated its optimal settings.	125
		126
		127
	3. We validated through experiments on 4 tasks that BATCHEVAL outperforms public state-of-the-art methods by 10.5% while incurring only 64% of the API cost.	128
		129
		130
		131
	4. We analyzed the generalization, robustness, hyperparameter selection, and probed into the working mechanism of BATCHEVAL.	132
		133
		134
	2 Background	135
	2.1 Automatic Text Evaluation	136
	Automatic text evaluation method has been extensively studied as a supplement to labor-intensive and time-consuming human evaluation, with its correlation to human judgment as the criterion for assessment. Both <i>rule-based</i> (Papineni et al., 2002; Lavie and Denkowski, 2009) and <i>embedding-based</i> (Zhang et al., 2020; Forgues et al., 2014) evaluation methods rely on the assumption that high-quality generated texts should have a significant word overlap with reference texts. However, this assumption conflicts with the high entropy nature of text generation, restricting its consistency with humans. <i>Learning-based</i> methods consider directly assessing text quality through supervised (Lowe et al., 2017; Goyal and Durrett, 2021) and self-supervised (Mehri and Eskénazi, 2020; Zhang et al., 2022) approaches and achieve significant progress. Recently, <i>LLM-based</i> evaluators (Guo et al., 2023b; Chiang and Lee, 2023b; Liu et al., 2023) have demonstrated advanced consistency with humans leveraging their incredible knowledge and capabilities. However, typical sample-wise evaluation paradigm of the above methods leads to a lack of inter-sample comparison during scoring process, which serves as an important reference for human evaluators. Therefore, we propose BATCHEVAL to fill this gap for better alignment with humans.	137
		138
		139
		140
		141
		142
		143
		144
		145
		146
		147
		148
		149
		150
		151
		152
		153
		154
		155
		156
		157
		158
		159
		160
		161
		162
		163
	2.2 Supportive Theorems	164
	Theorem 1 <i>The robustness against noise correlates positively with the uniformity of evaluator</i>	165
		166

scoring distribution. (See Appendix A.1 for derivation in details)

Yuan et al. (2023) proposed this theorem and verified that learning-based evaluators, by adjusting the training loss function to uniformize the score distribution, can achieve better robustness against noise. We have experimentally proven that sample-wise LLM-based evaluators also exhibit an uneven score distribution (Figure 3), which can weaken their robustness against noise (Appendix C). Thus, we propose BATCHEVAL for a more uniform score distribution and better robustness against noise.

Theorem 2 Given scores from multiple generations of certain LLM $\mathcal{S} = \{s_i | i = 1, \dots, N\}$ and human evaluation score y for sample x , \bar{s} is the average of \mathcal{S} , the following equation holds:

$$Err(\bar{s}, y) = Err(\mathcal{S}, y) - Var(\mathcal{S}) \quad (1)$$

where:

$$\begin{aligned} Err(\bar{s}, y) &= (\bar{s} - y)^2 \\ Err(\mathcal{S}, y) &= \frac{1}{N} \sum_{i=1}^N (s_i - y)^2 \\ Var(\mathcal{S}) &= \frac{1}{N} \sum_{i=1}^N (s_i - \bar{s})^2 \end{aligned} \quad (2)$$

Eq. (1) (Zhou, 2012) (proof in Appendix A.2) implies that smaller average error in single prediction scores ($Err(\mathcal{S}, y)$) and larger variance among multiple prediction scores ($Var(\mathcal{S})$) induce smaller error in ensemble score ($Err(\bar{s}, y)$). However, current sample-wise LLM evaluators score multiple times based solely on static reference (criterion), resulting in smaller $Var(\mathcal{S})$ (Figure 5). To address this, we propose iterative quality-heterogenized batch composition strategy for LLMs to score with unbiased varying references, thus increasing $Var(\mathcal{S})$ for lower $Err(\bar{s}, y)$.

3 Methodology

The core idea behind BATCHEVAL is to fully use in-batch sample comparison to enhance evaluation accuracy and robustness. Algorithm 1 illustrates the working process of BATCHEVAL, which involves N rounds of iteration: (1) B samples of each batch are compiled with pre-defined (task, criterion, evaluation procedure) into a single prompt for input to the LLM; (2) Based on the LLM’s assessment of the samples’ quality, we optimize batch allocation according to certain batch composition strategy. The core designs throughout the process

are **how to evaluate** (evaluation procedure), **what to input** (batch composition strategy), and **what to output** (scoring format). Below we discuss their potential variants in detail.

Algorithm 1 Workflow of BATCHEVAL.

Require: Samples $x^{1:|\mathcal{D}|}$, LLM \mathcal{M} , Evaluation procedure P
 Task and criterion T , Iteration rounds N , Batchsize B
 Batch composition strategy BATCHSTRATEGY

Ensure: Ensemble evaluation scores $\bar{s}^{1:|\mathcal{D}|}$

- 1: Randomly divide $x^{1:|\mathcal{D}|}$ into batches $b^{1:L}$, $L = \lceil \frac{|\mathcal{D}|}{B} \rceil$
- 2: $S_{all} \leftarrow \{i : [] \text{ for } i \in [1, |\mathcal{D}|]\}$
- 3: **for** $i \leftarrow 1, N$ **do**
- 4: $S_{current} \leftarrow \emptyset$
- 5: **for** $j \leftarrow 1, L$ **do:**
- 6: $S_{current} \leftarrow S_{current}.Append(\mathcal{M}(T, P, b^j))$
- 7: **end for**
- 8: $S_{all} \leftarrow S_{all}.Merge(S_{current})$
- 9: $b^{1:L} \leftarrow BATCHSTRATEGY(x^{1:|\mathcal{D}|}, S_{all}, B)$
- 10: **end for**
- 11: $\bar{s}^{1:|\mathcal{D}|} \leftarrow Average(S_{all})$

3.1 How to Evaluate

Sample-wise LLM evaluators work through a process of either analyzing followed by scoring (Guo et al., 2023b) or scoring followed by analyzing (Liu et al., 2023), where the former typically performs better (Chiang and Lee, 2023b) possibly due to the effect of chain-of-thought (Wei et al., 2022). On this basis, we further explore what procedures can better facilitate sample comparison for BATCHEVAL (Appendix J for prompts):

One stage As the most intuitive extension of sample-wise evaluation, LLM analyzes and scores each sample of the batch in order. This procedure enables adequate comparison between samples, but insufficient comparison between analyses (the analyses of subsequent samples cannot be referenced by the earlier samples for scoring).

Two stage To enhance the comparison among analyses, the LLM first analyzes all the samples. Based on the full comparisons among samples and analyses, the LLM further scores for each sample.

Three stage From human experience, it can be easier to first rank and then score the samples, as compared to directly scoring them. Therefore, we consider a procedure that sequentially performs analyzing, ranking, and scoring for all samples.

3.2 What to Input

The composition of the batch largely determines the efficacy of in-batch comparison as evaluation

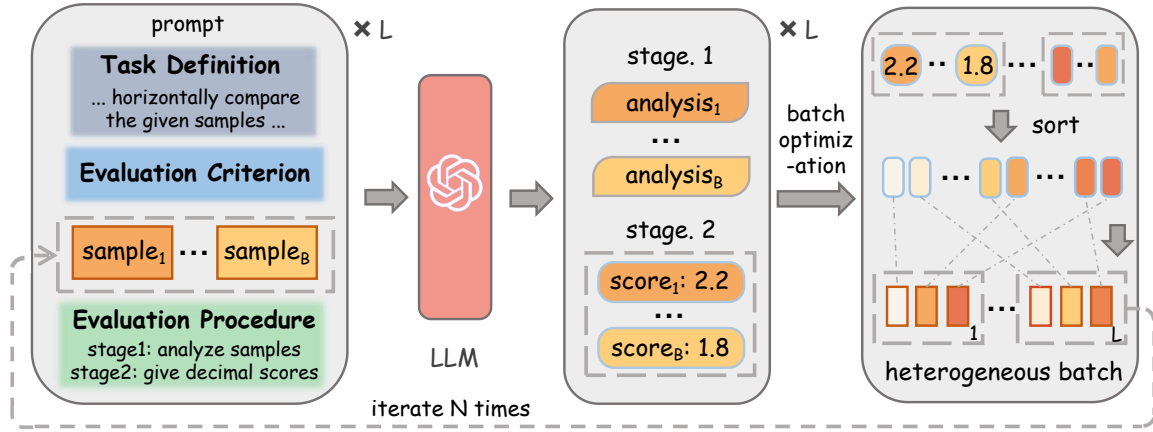


Figure 2: Overall illustration of BATCHEVAL.

reference. According to Theorem 2, we consider redrawing the batch divisions after each round of evaluation to provide the LLM with varying references when assessing a certain sample, thus can improve scoring diversity. Besides diversity, we are curious about what other characteristics the batch should possess to further enhance the effectiveness of BATCHEVAL, for which we explore the following strategies.²

Random Batch One base strategy is to reallocate batches randomly after each round of evaluation.

Homogeneous Batch Based on the idea of coarse-to-fine evaluation, we consider forming homogeneous batches in which samples have similar scores from the previous round of evaluation, in the hope that these samples can be further compared by LLM and ultimately attain discriminative scores.

Heterogeneous Batch A contrary idea is to select samples with diversified scores based on the previous round of evaluation results to form a new batch. In this way, LLM develops an unbiased perception of samples with different qualities through batch optimization, thus scoring more accurately.

3.3 What to Output

Sample-wise evaluation methods typically apply integers as the format for LLM scoring (Liu et al., 2023; Chiang and Lee, 2023b), and Lin and Chen (2023) proved that using more refined scoring format can not bring additional gains. *Will this trend be similar in BATCHEVAL?* Let us consider a concrete example: there are two samples with close but different quality, with human ratings of 2.2 and 1.8, respectively. Due to having only the criterion

²See Appendix F for strategies in detail.

as reference, sample-wise evaluators may consider them to be close to the 2-point standard and consequently assign a score of 2 regardless of whether decimal is allowed. However, if they appear in the same batch, on the basis of judging that they are all close to 2 points, LLM can further compare their quality directly. Thus, it is possible for LLM to give them differentiated decimal scores if it is allowed, thereby achieving more consistent judgments with humans. Based on the analysis above, we consider trying out two different scoring formats: **integer** and **decimal**.

Our default settings of BATCHEVAL include two stage evaluation procedure, heterogeneous batch composition strategy and decimal scoring format, as shown in Figure 2.

4 Experiments

Centered around BATCHEVAL, we will empirically explore the optimal variants in §4.2, demonstrate its performance on different LLMs and tasks in §4.3, validate the robustness in §4.4, and delve into its working mechanism in §4.5. We also investigate the choice of hyperparameters in Appendix §B.

4.1 Experimental settings

Benchmarks A brief introduction of benchmarks involved are listed as follows:

- **Topical-Chat** (Mehri and Eskénazi, 2020) is a benchmark for evaluating dialogue response generation. To save costs, we exclude knowledge as input to LLM and therefore choose criteria where knowledge is not necessary: Naturalness, Coherence, Engaging, Naturalness and Overall.

Type	Method	Scheme	Engaging		Understand		Naturalness		Coherence		Overall		Average		
			r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	\$/item
Human	Inter-annotator*	-	.575	.581	.510	.510	.486	.487	.558	.560	.710	.718	.568	.571	-
Rule	BLEU-4*	-	.232	.316	.201	.218	.180	.175	.131	.235	.216	.296	.192	.248	-
	METEOR*	-	.367	.439	.245	.225	.212	.191	.250	.302	.337	.391	.282	.310	-
Embedding	V-Extrema*	-	.210	.205	.156	.132	.101	.076	.184	.184	.203	.209	.171	.161	-
	BERTScore*	-	.317	.335	.256	.226	.226	.209	.214	.233	.298	.325	.262	.266	-
Learning	USR*	-	.456	.465	.293	.315	.276	.304	.416	.377	.422	.419	.373	.376	-
	BCR	-	.460	.463	.297	.325	.260	.298	.425	.391	.437	.421	.376	.380	-
LLM	G-Eval	-	.710	.719	.568	.593	.595	.605	.576	.584	.717	.705	.633	.641	.0614
	CloserLook	-	.651	.688	.649	.699	.656	.665	.675	.687	.778	.772	.682	.702	.0686
	CloserLook	+ ICL	.714	.743	.603	.685	.679	.693	.720	.733	.786	.783	.700	.727	.0856
		one stage	.780	.783	.642	.680	.706	.710	.727	.729	.785	.793	.728	.739	.0525
		three stage	.782	.778	.667	.725	.712	.704	.712	.714	.797	.798	.734	.744	.0541
	BATCHEVAL	random	.746	.743	.685	.724	.711	.700	.716	.720	.798	.799	.731	.737	.0528
	(Ours)	homogeneous	.654	.663	.639	.607	.671	.674	.669	.631	.722	.703	.671	.656	.0537
	integer	.771	.778	.686	.732	.726	.727	.722	.727	.790	.783	.739	.749	.0526	
	default	.792	.790	.694	.727	.730	.735	.740	.744	.805	.800	.752	.759	.0529	

Table 1: Turn-level Pearson (r_p) / Spearman (r_s) correlations and average API cost per sample (\$/item) of different metrics on Topical-Chat benchmark. The results of methods with * come from USR. We reproduced other methods with a unified API (the results were generally better than those reported in the original paper). All results of our replication are statistically significant (p-value < 0.05).

- **FED** (Mehri and Eskenazi, 2020) includes human ratings on 11 criteria to evaluate the quality of dialogue. We choose the top 4 important criteria as claimed in the original paper for evaluation: Coherent, Understanding, Likeable and Overall.
- **HANNA** (Chhun et al., 2022) serves as a benchmark for meta-evaluating evaluation methods on story generation, with criteria including: Coherence, Relevance, Empathy, Surprise, Engagement and Complexity.
- **QAGS** (Chhun et al., 2022) is a benchmark for evaluating the Factual Consistency of summaries on CNN (Hermann et al., 2015) and XSUM (Narayan et al., 2018).

Baselines We introduce four types of baseline methods in the experiments. Among them, both rule-based and embedding-based methods need reference text, which is unavailable in FED and QAGS. Learning-based methods are typically task-specific. Below we briefly list their categories and snapshots of LLM-based methods. Refer to Appendix G for detailed introductions.

- **Rule-based:** BLEU (Papineni et al., 2002), METEOR (Lavie and Denkowski, 2009).

- **Embedding-based:** Vector Extrema (Forgues et al., 2014), BERTScore (Zhang et al., 2020)
- **Learning-based:** USR (Mehri and Eskénazi, 2020), BCR (Yuan et al., 2023), FED (Mehri and Eskenazi, 2020), DynaEval (Zhang et al., 2021), QAGS (Wang et al., 2020).
- **LLM-based**³: G-Eval (Liu et al., 2023) recommended using LLM to evaluate according to the procedures generated by itself. Chiang and Lee (2023b) tried various evaluation schemes and proved through experiments that *analyze-rate* led to the best performance, which we denote as CloserLook.

Details We explore variants of BATCHEVAL on Topical-Chat for its wide recognition. If not specified, FED serves as our default dataset for exploratory experiments as it only has 125 samples, thus can save API expenses. The other two benchmarks are used to confirm the generalization across tasks of BATCHEVAL. We primarily conduct experiments with GPT-4 (0613) and validate the generalization across models of BATCHEVAL with GPT-3.5-turbo (0613) and Llama-2-70b-chat-hf. We set iteration rounds as 5, batchsize as 10, decoding temperature as 0.2 for all the experiment. For other

³Two latest and well-known LLM evaluators are included. We are unable to reproduce some other methods due to incomplete disclosure of codes or prompts.




Figure 3: Score distribution and corresponding entropy ($-\sum_s p(s) \log_2 p(s)$) of different methods.




Figure 4: Comparisons between BATCHEVAL and CloserLook from the perspective of Theorem 2.




Figure 5: Average batch bias of different strategies.

LLM-based evaluators, we reproduced them according to their default settings (20 generations per sample) with the same API for a fair comparison. We choose Pearson and Spearman correlations to measure consistency with humans and also report API expenses for adequate comparison. We follow (Chiang and Lee, 2023b) to design prompts (See prompts in Appendix J).

4.2 Variants Exploration

As shown in Table 1, based on the default settings shown in Figure 2, we validate the effects of different variants (replacing the default setting with specific scheme) of BATCHEVAL.

Evaluation Procedure Compared to one stage procedure, the two stage procedure (default) achieves higher correlations by enhancing the comparison among analyses during scoring. Surprisingly, however, the three stage procedure does not

perform well as expected. We speculate this may be due to the LLM’s over-reliance on ranking results while neglecting the analyses and samples during scoring, and validate this in Appendix D.

Batch Composition Strategy As shown in Table 1, the performance of batch composition strategies ranks as follows: heterogeneous (default) > random > homogeneous. To investigate the reasons, we introduce batch bias as follows:

$$Bias(\mathcal{B}) = abs(\sum_{i \in \mathcal{B}} s_i^{\mathcal{B}} - \sum_{i \in \mathcal{B}} \bar{s}_i) / |\mathcal{B}| \quad (3)$$

where \mathcal{B} denotes the set of sample indexes of certain batch, $s_i^{\mathcal{B}}$ denotes score of sample x_i generated with batch \mathcal{B} , \bar{s}_i denotes average score of sample x_i across all the iterations. Ideally, we aspire for the batch bias to approach zero. This implies that LLM should not have the overall scores in a batch skewed either high or low compared to the ensemble scores. We evaluate the average $Bias(\mathcal{B})$ of different strategies and find that $Bias(\mathcal{B})$ correlates negatively with correlations r_s and r_p (Figure 5). This indicates that the more varied the quality of samples in a batch, the better they can simulate a real distribution as an unbiased reference to bring smaller batch bias for better correlations.

Scoring Format We observe from Table 1 that decimal scoring format brings around 1 point correlations improvement upon integer. As shown in Figure 3, the decimal scheme brings a more uniform scoring distribution. This implies that LLM indeed assigns more discriminative scores to different samples through in-batch comparison if decimal score is allowed, which verifies our hypothesis in §3.3 and accounts for the progress.

4.3 Overall Performance of BATCHEVAL

As shown in Table 1, 2, 3, 4, BATCHEVAL achieves an average of 6.5 points (10.5%) Pearson and 4.5 points (7.1%) Spearman correlations improvements with humans across four benchmarks compared to

Type	Method	Model	Likeable		Understand		Coherent		Overall		Average		
			r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	\$/item
Human	Inter-annotator	-	-	.838	-	.809	-	.809	-	.830	-	.822	-
	USR	-	.245	.226	.182	.178	.170	.185	.284	.302	.220	.223	-
Learning	FED	-	.248	.262	.295	.306	.262	.253	.460	.449	.316	.318	-
	DynalEval	-	.389	.393	.379	.368	.399	.409	.484	.490	.413	.415	-
LLM	CloserLook	Llama-2-70b	.525	.550	.574	.611	.640	.563	.634	.639	.593	.591	-
	BATCHEVAL	Llama-2-70b	.537	.563	.619	.597	.627	.648	.722	.732	.626	.635	-
	CloserLook	GPT-3.5-turbo	.681	.666	.691	.605	.726	.724	.687	.709	.696	.676	.0022
	BATCHEVAL	GPT-3.5-turbo	.682	.674	.704	.708	.733	.730	.705	.699	.706	.703	.0011
	G-Eval	GPT-4	.638	.692	.670	.625	.707	.721	.689	.652	.676	.673	.0667
	CloserLook w human prompt	GPT-4	.658	.680	.701	.614	.739	.751	.715	.684	.703	.682	.0785
	CloserLook w GPT-4 prompt	GPT-4	.632	.660	.678	.639	.725	.749	.723	.678	.690	.682	.0827
	BATCHEVAL w human prompt	GPT-4	.731	.741	.778	.696	.753	.753	.738	.729	.750	.730	.0314
	BATCHEVAL w GPT-4 prompt	GPT-4	.736	.741	.780	.700	.784	.749	.748	.727	.762	.729	.0314

Table 2: Dialog-level Pearson (r_p) / Spearman (r_s) correlations and average API cost per sample (\$/item) on FED-dialog benchmark. We implemented and tested all the methods with p-value < 0.05.

Method	Coherence		Relevance		Empathy		Surprise		Engagement		Complexity		Average		
	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	\$/item
BLEU-4	.220	.218	.135	.175	.242	.216	.178	.224	.242	.270	.362	.273	.230	.229	-
METEOR	.335	.273	.202	.190	.304	.282	.285	.283	.316	.338	.520	.482	.307	.307	-
BERTScore	.358	.293	.201	.188	.308	.303	.302	.290	.308	.331	.501	.472	.330	.313	-
G-Eval	.572	.578	.582	.584	.453	.461	.311	.347	.562	.591	.602	.557	.514	.520	.0772
CloserLook	.595	.591	.579	.597	.498	.478	.280	.339	.605	.607	.619	.568	.529	.530	.0835
BATCHEVAL	.678	.625	.702	.679	.546	.543	.368	.381	.617	.605	.625	.575	.589	.568	.0538

Table 3: Story-level Pearson (r_p) / Spearman (r_s) correlations and average API cost per sample (\$/item) of on HANNA benchmark. We implemented and tested all the methods with p-value < 0.05.

Method	QAGS-C		QAGS-X		Average		
	r_p	r_s	r_p	r_s	r_p	r_s	\$/item
BERTScore*	.576	.505	.024	.008	.300	.256	-
QAGS*	.545	-	.175	-	.375	-	-
G-Eval*	.631	.685	.558	.537	.599	.611	-
CloserLook	.581	.602	.549	.573	.498	.478	.0691
BATCHEVAL	.785	.643	.618	.634	.682	.639	.0521

Table 4: Results on QAGS benchmark (QAGS with -C and -X denote subset CNN and XSUM respectively). Results with * come from G-EVAL. we present the original results of G-Eval here as our replication is not good as those reported in the original paper.

the best performing methods. From the perspective of Theorem 2, as shown in Figure 4, we found that the reason BATCHEVAL outperforms CloserLook under score ensemble ($Err(\bar{s}, y)$) is twofold. First, BATCHEVAL attains more accurate single predictions ($Err(\mathcal{S}, y)$) through thorough in-batch comparison. Second, the scoring diversity ($Var(\mathcal{S})$) of BATCHEVAL is significantly improved. This validates that iterative heterogeneous batch composition strategy can provide LLM with unbiased

varying evaluation references, thus stably enhancing diversity and ensemble performance.

In terms of cost, BATCHEVAL only consumes 64% API expenses of the best performing baselines. This is because we only use the average scores from 5 iterations and allow in-batch samples to share single prompt, while the LLM-based baselines average scores from 20 generations.⁴ Considering that baselines reach ensemble saturation at about 20 generations, BATCHEVAL has broad potential for performance improvement with more iterations.

4.4 Robustness of BATCHEVAL

Robustness against Prompt Design We test BATCHEVAL and CloserLook respectively on prompts written by human and rewritten by GPT-4, with results as shown in Table 2. We calculate the average difference in correlations across metrics under two types of prompts. The standard deviation

⁴Due to changes in the prompt during iteration, the prompt expense needs to be billed 5 times for our method, whereas baselines require only once. Therefore the expenditure ratio (64%) is higher than the proportion of generations (5:20).




Figure 6: Normalized attention matrices of the first (top figure) and last (bottom figure) transformer layer with Llama-2-70b-chat-hf. We set batchsize as 3 for clear demonstration. See Appendix H for the normalizing process.

of r_p and r_s are 0.009 and 0.007 for CloserLook, while only 0.006 and 0.002 for BATCHEVAL. This verifies that BATCHEVAL attains better robustness against prompt design by introducing in-batch samples as additional references.

Robustness against Noise As shown in Figure 3, the score distribution of BATCHEVAL is more uniform and has lower entropy compared with CloserLook due to in-batch comparison with decimal scoring format, which can theoretically enhance robustness against noise according to Theorem 1. We further experimentally validate this in Appendix C.

4.5 Further Discussion and Analysis

Relationship with In-context-learning ICL (Brown et al., 2020) can also provide sample-side references by incorporating samples and corresponding answers into the prompt. The main differences between ICL and BATCHEVAL are: (1) BATCHEVAL can provide LLM with varying and comprehensive references through iterative heterogeneous batch, while the references provided by ICL are relatively fixed and may bring bias (sensitive to prompt design). (2) BATCHEVAL uses in-batch samples as references to each other, thus saving the costs of demonstrations in ICL prompts. Thanks to the aforementioned advancements, BATCHEVAL outperforms CloserLook with ICL by more than 5 points Pearson correlations while only incurs 61.8% expense (Table 1).

Working Mechanism of BATCHEVAL To further understand how BATCHEVAL benefits from

in-batch comparison, we visualized the normalized attention matrices of the first and last layers of Llama-2-70b-chat-hf (Figure 6). The value at (X,Y) represents the average normalized attention of tokens corresponding to X towards tokens corresponding to Y. We observe that in the final scoring phase (red box), LLM first perceives samples with varied qualities based on the already generated scores and analyses at the shallower layers. Afterwards, LLM completes scoring based on criterion and comparison between samples at the deeper layers. This process demonstrates the in-batch comparison mechanism of BATCHEVAL, which we hope can inspire future research.

5 Conclusions

In this paper, we propose BATCHEVAL, a new text evaluation paradigm that evaluate samples batch-wise to alleviate the limitations of sample-wise evaluation paradigm. We explore variants of BATCHEVAL on multiple dimensions and figure out the optimal settings. Following the human evaluation method, BATCHEVAL treats in-batch samples and criterion as complementary references and optimizes the batch composition through iteration to eliminate batch bias. Comprehensive experiments have confirmed that BATCHEVAL can achieve higher consistency with humans at a lower cost, while also demonstrating better robustness to prompt design and noise. We further analyze and reveal the working mechanism of BATCHEVAL, shedding lights on future work.

507 Limitations

508 From an objective perspective , we think there are
509 two main limitations of this paper:

- 510 1. BATCHEVAL requires LLMs to have a certain
511 capability to handle longer contexts. From
512 Appendix B, we found that as the batchsize
513 increases, LLMs struggle to handle too many
514 samples, leading to a performance decline.
515 We also attempted to test BATCHEVAL’s per-
516 formance on Llama-2-13b-chat-hf and found
517 that the batchsize must be set to 2 or 3 to
518 see any benefits. Therefore, when setting the
519 batchsize, we cannot exceed the limit of how
520 many samples an LLM can process in a single
521 context. Fortunately, we discovered that
522 a batchsize of 10 is suitable for current main-
523 stream LLMs. Additionally, as LLMs continue
524 to advance, they can handle increasingly
525 larger contexts. Thus, from this perspective,
526 BATCHEVAL is a scalable method that im-
527 proves alongside the capabilities of LLMs (in-
528 creasing the batchsize within the capabilities
529 of the LLM can enhance the evaluation effec-
530 tiveness of the LLM).
- 531 2. We only explored a limited number of
532 schemes of BATCHEVAL. We leave exploring
533 possible schemes of BATCHEVAL for future
534 research.

535 Ethics Statement

536 All of the datasets used in this study were publicly
537 available, and no annotators were employed for our
538 data collection. We confirm that the datasets we
539 used did not contain any harmful content and was
540 consistent with their intended use (research). We
541 have cited the datasets and relevant works used in
542 this study.

543 References

544 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie
545 Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind
546 Neelakantan, Pranav Shyam, Girish Sastry, Amanda
547 Askell, Sandhini Agarwal, Ariel Herbert-Voss,
548 Gretchen Krueger, Tom Henighan, Rewon Child,
549 Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,
550 Clemens Winter, Christopher Hesse, Mark Chen, Eric
551 Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess,
552 Jack Clark, Christopher Berner, Sam McCandlish,
553 Alec Radford, Ilya Sutskever, and Dario Amodei.
554 2020. [Language models are few-shot learners](#). In *Ad-
555 vances in Neural Information Processing Systems 33:*

*Annual Conference on Neural Information Process-
ing Systems 2020, NeurIPS 2020, December 6-12,
2020, virtual*. 556
557
558

Cyril Chhun, Pierre Colombo, Fabian M Suchanek, and
Chloé Clavel. 2022. [Of human criteria and auto-
matic metrics: A benchmark of the evaluation of
story generation](#). In *29th International Conference
on Computational Linguistics (COLING 2022)*. 559
560
561
562
563

David Cheng-Han Chiang and Hung-yi Lee. 2023a. [Can
large language models be an alternative to human
evaluations?](#) In *Proceedings of the 61st Annual
Meeting of the Association for Computational Lin-
guistics (Volume 1: Long Papers), ACL 2023, Toronto,
Canada, July 9-14, 2023*, pages 15607–15631. Asso-
ciation for Computational Linguistics. 564
565
566
567
568
569
570

David Cheng-Han Chiang and Hung-yi Lee. 2023b. [A
closer look into using large language models for auto-
matic evaluation](#). In *Findings of the Association for
Computational Linguistics: EMNLP 2023, Singapore,
December 6-10, 2023*, pages 8928–8942. Association
for Computational Linguistics. 571
572
573
574
575
576

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and
Kristina Toutanova. 2019. [BERT: pre-training of
deep bidirectional transformers for language under-
standing](#). In *Proceedings of the 2019 Conference of
the North American Chapter of the Association for
Computational Linguistics: Human Language Tech-
nologies, NAACL-HLT 2019, Minneapolis, MN, USA,
June 2-7, 2019, Volume 1 (Long and Short Papers)*,
pages 4171–4186. Association for Computational
Linguistics. 577
578
579
580
581
582
583
584
585
586

Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang,
Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy
Liang, and Tatsunori B Hashimoto. 2023. [Al-
pacafarm: A simulation framework for methods
that learn from human feedback](#). *arXiv preprint
arXiv:2305.14387*. 587
588
589
590
591
592

Gabriel Forgues, Joelle Pineau, Jean-Marie
Larchevêque, and Réal Tremblay. 2014. [Bootstrapping
dialog systems with word embeddings](#). In *Nips,
modern machine learning and natural language
processing workshop*, volume 2, page 168. 593
594
595
596
597

Tanya Goyal and Greg Durrett. 2021. [Annotating and
modeling fine-grained factuality in summarization](#).
In *Proceedings of the 2021 Conference of the North
American Chapter of the Association for Computa-
tional Linguistics: Human Language Technologies,
NAACL-HLT 2021, Online, June 6-11, 2021*, pages
1449–1462. Association for Computational Linguis-
tics. 598
599
600
601
602
603
604
605

Zishan Guo, Renren Jin, Chuang Liu, Yufei Huang, Dan
Shi, Supryadi, Linhao Yu, Yan Liu, Jiaxuan Li, Bo-
jian Xiong, and Deyi Xiong. 2023a. [Evaluating large
language models: A comprehensive survey](#). *CoRR*,
abs/2310.19736. 606
607
608
609
610

611	Zishan Guo, Renren Jin, Chuang Liu, Yufei Huang, Dan Shi, Supryadi, Linhao Yu, Yan Liu, Jiaxuan Li, Bojian Xiong, and Deyi Xiong. 2023b. Evaluating large language models: A comprehensive survey . <i>CoRR</i> , abs/2310.19736.	667
612		668
613		669
614		
615		
616	Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. <i>Advances in neural information processing systems</i> , 28.	
617		
618		
619		
620		
621	Alon Lavie and Michael J. Denkowski. 2009. The meteor metric for automatic evaluation of machine translation . <i>Mach. Transl.</i> , 23(2-3):105–115.	
622		
623		
624	Yen-Ting Lin and Yun-Nung Chen. 2023. Llm-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models . <i>CoRR</i> , abs/2305.13711.	
625		
626		
627		
628	Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: NLG evaluation using gpt-4 with better human alignment . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023</i> , pages 2511–2522. Association for Computational Linguistics.	
629		
630		
631		
632		
633		
634		
635		
636	Ryan Lowe, Michael D. Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses . In <i>Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers</i> , pages 1116–1126. Association for Computational Linguistics.	
637		
638		
639		
640		
641		
642		
643		
644		
645	Shikib Mehri and Maxine Eskenazi. 2020. Unsupervised evaluation of interactive dialog with dialogpt. In <i>21th Annual Meeting of the Special Interest Group on Discourse and Dialogue</i> , page 225.	
646		
647		
648		
649	Shikib Mehri and Maxine Eskénazi. 2020. USR: an unsupervised and reference free evaluation metric for dialog generation . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020</i> , pages 681–707. Association for Computational Linguistics.	
650		
651		
652		
653		
654		
655	Shashi Narayan, Shay Cohen, and Maria Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In <i>2018 Conference on Empirical Methods in Natural Language Processing</i> , pages 1797–1807. Association for Computational Linguistics.	
656		
657		
658		
659		
660		
661		
662	OpenAI. 2023. GPT-4 technical report . <i>CoRR</i> , abs/2303.08774.	
663		
664	Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation . In <i>Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA</i> , pages 311–318. ACL.	670
665		671
666		672
		673
	Ananya B. Sai, Akash Kumar Mohankumar, and Mitesh M. Khapra. 2023. A survey of evaluation metrics used for NLG systems . <i>ACM Comput. Surv.</i> , 55(2):26:1–26:39.	674
		675
		676
	Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and answering questions to evaluate the factual consistency of summaries. In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 5008–5020.	677
		678
		679
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models . In <i>NeurIPS</i> .	680
		681
		682
		683
	Peiwen Yuan, Xinglin Wang, Jiayi Shi, Bin Sun, Yiwei Li, and Kan Li. 2023. Better correlation and robustness: A distribution-balanced self-supervised learning framework for automatic dialogue evaluation. In <i>Thirty-seventh Conference on Neural Information Processing Systems</i> .	684
		685
		686
		687
		688
		689
	Chen Zhang, Yiming Chen, Luis Fernando D’Haro, Yan Zhang, Thomas Friedrichs, Grandee Lee, and Haizhou Li. 2021. Dynaeval: Unifying turn and dialogue level evaluation . In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021</i> , pages 5676–5689. Association for Computational Linguistics.	690
		691
		692
		693
		694
		695
		696
		697
		698
		699
	Pengfei Zhang, Xiaohui Hu, Kaidong Yu, Jian Wang, Song Han, Cao Liu, and Chunyang Yuan. 2022. MME-CRS: multi-metric evaluation based on correlation re-scaling for evaluating open-domain dialogue . <i>CoRR</i> , abs/2206.09403.	700
		701
		702
		703
		704
	Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT . In <i>8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020</i> . OpenReview.net.	705
		706
		707
		708
		709
		710
	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. <i>arXiv e-prints</i> , pages arXiv–2306.	711
		712
		713
		714
		715
	Zhi-Hua Zhou. 2012. <i>Ensemble methods: foundations and algorithms</i> . CRC press.	716
		717

A Proof of Theorems Involved

A.1 Theorem 1

For any $f(x)$, the probability density function of score distribution, the Spearman correlation $\mathbb{E}(r_s)$ between the original scores and scores adding a small disturbance has an upper bound:

$$\mathbb{E}(r_s) \leq 1 - \frac{6\mathbb{E}(\lambda)^2}{n^2 - 1}, \quad (4)$$

and the equality condition is $f(x) \equiv 1, \forall x \in [0, 1]$.

Proof 1 The ranking difference $d(x)$ before and after disturbance is :

$$d(x) = \int_x^{x+\lambda} f(x) dx \quad (5)$$

According to the definition of Spearman correlations, $\mathbb{E}(r_s)$ can be written as:

$$\mathbb{E}(r_s) = \mathbb{E}\left(1 - \frac{6 \sum_{i=1}^n d(x_i)^2}{n(n^2 - 1)}\right), \quad (6)$$

we derive the lower bound of $\mathbb{E}(d(x)^2)$ as follows:

$$\begin{aligned} & \mathbb{E}(d(x)^2) \\ &= \int_0^1 \left(\int_x^{x+\mathbb{E}(\lambda)} f(u) du \right)^2 f(x) dx \\ &= \int_0^1 \left(\int_x^{x+\mathbb{E}(\lambda)} f(u) du \sqrt{f(x)} \right)^2 dx \\ &= \int_0^1 \left(\int_x^{x+\mathbb{E}(\lambda)} f(u) du \sqrt{f(x)} \right)^2 dx \\ & \quad \cdot \int_0^1 f(x) dx \\ &\geq \left(\int_0^1 \int_x^{x+\mathbb{E}(\lambda)} f(u) du f(x) dx \right)^2 \quad (7) \\ & \quad (\text{Cauchy's Inequality}) \\ &= \left(\int_0^1 \mathbb{E}(\lambda) \cdot f(x) \cdot f(x) dx \right)^2 \quad (\mathbb{E}(\lambda) \rightarrow 0) \\ &= \mathbb{E}(\lambda)^2 \left(\int_0^1 f(x) \cdot f(x) dx \right)^2 \\ &= \mathbb{E}(\lambda)^2 \left(\int_0^1 f(x)^2 dx \cdot \int_0^1 1^2 dx \right)^2 \\ &\geq \mathbb{E}(\lambda)^2 \left(\left(\int_0^1 f(x) dx \right)^2 \right)^2 \\ & \quad (\text{Cauchy's Inequality}) \\ &= \mathbb{E}(\lambda)^2 \end{aligned}$$

The equality condition is $f(x) \equiv 1$ for $x \in [0, 1]$. Taking the lower bound of $\mathbb{E}(d(x)^2)$ into

Eq. (6), we conclude the proof. Note that higher $\mathbb{E}(r_s)$ denotes better robustness against noise. Hence, we can derive that the robustness against noise correlates positively with the uniformity of score distribution.

A.2 Theorem 2

Given scores from multiple generations of certain LLM $\mathcal{S} = \{s_i | i = 1, \dots, N\}$ and human evaluation score y for sample x , \bar{s} is the average of \mathcal{S} , the following equation holds:

$$Err(\bar{s}, y) = Err(\mathcal{S}, y) - Var(\mathcal{S}) \quad (8)$$

where:

$$\begin{aligned} Err(\bar{s}, y) &= (\bar{s} - y)^2 \\ Err(\mathcal{S}, y) &= \frac{1}{N} \sum_{i=1}^N (s_i - y)^2 \\ Var(\mathcal{S}) &= \frac{1}{N} \sum_{i=1}^N (s_i - \bar{s})^2 \end{aligned} \quad (9)$$

Proof 2

$$\begin{aligned} & Err(\mathcal{S}, y) - Var(\mathcal{S}) \\ &= \frac{1}{N} \sum_{i=1}^N (s_i - y)^2 - \frac{1}{N} \sum_{i=1}^N (s_i - \bar{s})^2 \\ &= \frac{1}{N} \left(\sum_{i=1}^N (s_i^2 + y^2 - 2s_i y - s_i^2 - \bar{s}^2 + 2s_i \bar{s}) \right) \\ &= y^2 - \bar{s}^2 - \frac{1}{N} \left(\sum_{i=1}^N 2s_i y \right) + \frac{1}{N} \left(\sum_{i=1}^N 2s_i \bar{s} \right) \\ &= y^2 - \bar{s}^2 - 2\bar{s}y + 2\bar{s}^2 \\ &= y^2 + \bar{s}^2 - 2\bar{s}y \\ &= (\bar{s} - y)^2 \\ &= Err(\bar{s}, y) \end{aligned} \quad (10)$$

B Hyperparameter Analysis

In the experiments of the main text, we set the batch size to 10 and the temperature to 0.2. In this section, we explore the impact of different hyperparameter choices on performance.

B.1 Effect of Batchsize

On FED dataset, we test BATCHEVAL with batchsize among [1, 2, 5, 10]. As shown in Figure 7, we found that as the batch size increases, the performance generally undergoes a process of initial improvement followed by a decline. Similar observations were made on other datasets as well. We further discovered that the performance turning point of the ensemble results from five iterations is slightly delayed compared to a single prediction. Considering that increasing the batchsize will make




Figure 7: Dialog-level Pearson correlations on FED-dialog dataset of BATCHEVAL with different batchsize.




Figure 8: Dialog-level Pearson correlations on FED-dialog dataset of BATCHEVAL with different temperature.

766 the combination of in-batch samples more diverse, 767 thereby increasing scoring diversity, we have the 768 following conjecture about Figure 7: When the 769 batchsize starts to increase from 1, due to the effect 770 of in-batch comparison and the increase in diversity, 771 the performance of both 1-round score and ensemble 772 score increase a lot. However, as the batchsize 773 continues to increase, LLM finds it difficult to handle 774 too many samples simultaneously, resulting in a decrease 775 in 1-round score performance. When the rate of decrease 776 in 1-round score performance gets greater than the rate 777 of increase in diversity, ensemble score performance 778 also begins to decrease according to Theorem 2. Therefore, 779 the batchsize should not be too large or too small. We 780 found that setting the batchsize to 10 can achieve superior 781 performance on different tasks. We also believe that 782 for LLMs with weaker ability to handle longer context, 783 the batchsize should be set to be smaller. Fortunately, 784 we have noticed that current LLMs are continually 785 improving in processing long contextual texts, which 786 illuminates further development prospects for BATCHEVAL 787 in the future. 788

789 B.2 Effect of Temperature

790 We also test BATCHEVAL with temperature among 791 [0, 0.2, 0.5, 1]. We found that as the temperature 792 rises in Figure 8, the performance of BATCHEVAL 793 does not exhibit a uniform trend of change. Overall, 794 the performance of 5 iterations is relatively stable 795 along the temperature dimension, suggesting that

BATCHEVAL is quite robust to temperature variations. 796 797

798 C Robustness against Noise

799 To test the robustness against noise of BATCHEVAL, 800 we use an external tool⁵ to add noise to the input 801 and calculate the changes in performance before and 802 after the noise is added. For the sake of noise balance, 803 we randomly replace 5% of tokens with synonyms and 804 randomly delete 5% of tokens. As shown in Table 5, 805 CloserLook experiences a decrease of 0.109 in Pearson 806 correlation and 0.081 in Spearman correlation, respectively. 807 In contrast, BATCHEVAL only shows a decrease of 808 0.003 and 0.009, respectively. This indicates that 809 BATCHEVAL has much better robustness to noise. 810

811 D Inferior Performance of Three Stage Procedure

812 As shown in Table 1, we observe a performance 813 drop of BATCHEVAL with three stage procedure, 814 though it may be closer to human evaluation procedure. 815 We speculate this may be due to the LLM’s over-reliance 816 on ranking results while neglecting the analyses and 817 samples during scoring. To validate this, we delete the 818 ranking and scoring contents of LLM’s three stage 819 procedure response and ask LLM to score based on the 820 remaining con- 821

⁵nlpaug(<https://github.com/makcedward/nlpaug>)

Method	Likeable		Understand		Coherent		Overall		Average		
	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	\$/item
CloserLook w/o noise	.658	.680	.701	.614	.739	.755	.715	.684	.703	.683	.0785
CloserLook w noise	.509	.580	.626	.606	.608	.605	.632	.616	.594 (-.109)	.602 (-.081)	.0866
BATCHEVAL w/o noise	.731	.741	.778	.696	.753	.757	.738	.729	.750	.731	.0314
BATCHEVAL w noise	.729	.718	.775	.700	.764	.754	.720	.724	.747 (-.003)	.724 (-.007)	.0344

Table 5: Story-level Pearson (r_p) / Spearman (r_s) correlations and average API cost per sample (\$/item) of on HANNA benchmark. We tested all the methods for a fair comparison with p-value < 0.05.

Method	Scheme	Engaging		Understand		Naturalness		Coherence		Overall		Average		
		r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	\$/item
BATCHEVAL	default	.792	.790	.694	.727	.730	.735	.740	.744	.805	.800	.752	.759	.0529
	3 stage	.782	.778	.667	.725	.712	.704	.712	.714	.797	.798	.734	.744	.0541
	3 stage w/o rank results	.789	.785	.701	.733	.721	.727	.735	.747	.810	.808	.751	.760	-

Table 6: Comparison of BATCHEVAL with different scheme. *3 stage w/o ranking results* means results of deleting the ranking and scoring contents of LLM’s three stage procedure response and asking LLM to score based on the remaining contents (samples and analyses)

tents (samples and analyses). If the new scoring results perform similarly to BATCHEVAL with two stage procedure, the inferior performance of BATCHEVAL with three stage procedure can be attributed to its excessive focus on ranking results. Otherwise, the reason lies in the decrease in the quality of analyses. As shown in Table 6, the performance of three stage w/o rank results is on par with that of two stage procedure. This validates our conjecture that the over-reliance on ranking results causes the performance drop of BATCHEVAL with three stage procedure.

E Relationship with Pair-wise Evaluation

The current mainstream text evaluation approach adopts sample-wise assessment. Alternatively, an LLM evaluator is presented with a question and two answers, and is tasked with determining which one is better or declaring a tie (Zheng et al., 2023; Dubois et al., 2023). However, as the number of models to be evaluated grows, the scalability of pairwise comparison becomes a challenge, due to the quadratic increase in the potential number of pairs. Therefore, this pair-wise paradigm has not been as extensively studied as sample-wise evaluation. Zheng et al. (2023) validates that this method performs slightly better than a sample-wise evaluator, potentially due to its ability to discern subtle differences between specific pairs.

Similarly, we have enhanced the evaluation capabilities of the LLM evaluator through in-batch sample comparison. The main difference lies in the composition of our batches, which consist of dif-

ferent samples rather than responses from different models to the same sample, thereby offering good scalability.

F Batch Composition Strategies

F.1 Homogenized Batch

Given scores $s^{1:|\mathcal{D}|}$ for samples $x^{1:|\mathcal{D}|}$ predicted by LLM in the previous round, we first sort the scores and attain the corresponding indexes $index^{1:|\mathcal{D}|}$. Based on this, we get indexes of homogenized batch $b^i = index^{1+(i-1)*10:i*10}$.

F.2 Heterogenized Batch

Given scores $s^{1:|\mathcal{D}|}$ for samples $x^{1:|\mathcal{D}|}$ predicted by LLM in the previous round, we first sort the scores and attain the corresponding indexes $index^{1:|\mathcal{D}|}$. Considering that our default batchsize is 10, we group the indexes into 10 splits $split^{1:10}$, where $split^i = index^{1+(i-1) \times \lceil \frac{|\mathcal{D}|}{10} \rceil : i \times \lceil \frac{|\mathcal{D}|}{10} \rceil}$. Based on this, we get indexes of heterogenized batch $b^i = \{split^{j,i} | j \in [1, 10]\}$.

G Introduction of Baselines

G.1 Rule-based Methods

BLEU (Papineni et al., 2002) BLEU is a renowned metric for measuring word overlap, which evaluates n-gram precision in a generated sequence against a reference. It includes a brevity penalty to counteract its inherent preference for shorter sentences, ensuring a more comprehensive assessment.

882	METEOR (Lavie and Denkowski, 2009) is an	attention between token pairs at different positions	928
883	advancement over BLEU, utilizing a harmonic	vary. If we use $Att(x, y)$ to represent the attention	929
884	mean of precision and recall, and also incorporat-	of the x^{th} token to the y^{th} token, then its expected	930
885	ing stemming and synonym use in its evaluation.	value is $\frac{1}{x}$. Since tokens at different positions will	931
886	G.2 Embedding-based Methods	be visualized into the same graph, we first multiply	932
887	Vector Extrema (Forgues et al., 2014) is a scor-	each $Att(x, y)$ by x to make its expected value	933
888	ing method that uses cosine similarity between sen-	1. On this basis, we determine the token intervals	934
889	tence embeddings, identifying the highest value in	corresponding to different strings through word	935
890	each dimension of the word embedding for evalua-	matching, and calculate $Att(string1, string2)$ as	936
891	tion.	follows:	937
892	BERTScore (Zhang et al., 2020) is a method that	$Att(s1, s2) = Avg(\{Att(x, y) x \in s1, y \in s2\})$	938
893	utilizes a pretrained BERT (Devlin et al., 2019)	(11)	939
894	model to optimally align each word in a reference	according to which we plot our attention matrices.	
895	response with a single word in the generated se-	I Case Study	940
896	quence. By doing so, BERTScore computes the	Below is an example response from GPT-4 regard-	941
897	recall of the generated sequence.	ing FED-coherence:	942
898	G.3 Learning-based Methods	<i>I will do my best to provide individual</i>	943
899	USR (Mehri and Eskénazi, 2020) is a dialogue	<i>analysis for each sample. Analysis:</i>	944
900	response evaluation method that uses one masked		945
901	language model and two dialogue retrieval models	<i>Sample1: The system maintains a</i>	946
902	to assess various sub-qualities of a sample and then	<i>coherent conversation with the user, but</i>	947
903	integrates these evaluations into a comprehensive	<i>there are some instances where the sys-</i>	948
904	overall score.	<i>tem's responses are not directly related</i>	949
905	BCR (Yuan et al., 2023) is a dialogue response	<i>to the user's questions. For example,</i>	950
906	evaluation method that use a dynamic loss func-	<i>when asked about the technology that</i>	951
907	tion to train a BERT model with uniform score	<i>powers it, the system fails to provide a</i>	952
908	distribution.	<i>clear answer.</i>	953
909	FED (Mehri and Eskenazi, 2020) is a unified		954
910	dialogue evaluation method that uses pretrained	<i>Sample2: The system maintains a</i>	955
911	language models to calculate scores based on the	<i>good flow of conversation and responds</i>	956
912	difference in the probability of generating positive	<i>appropriately to the user's questions.</i>	957
913	and negative evaluation words for a certain crite-	<i>However, there is a slight misunderstand-</i>	958
914	rion.	<i>ing about the release date of a movie.</i>	959
915	DynaEval (Zhang et al., 2021) is also a unified di-		960
916	alogue evaluation method that leverages graph con-	<i>Sample3: The system maintains a</i>	961
917	volutional network to model the sentences among	<i>coherent conversation with the user,</i>	962
918	a dialogue for accurate evaluation.	<i>responding appropriately to the user's</i>	963
919	QAGS (Wang et al., 2020) is a method that based	<i>statements and questions.</i>	964
920	on question-answering, which creates questions		965
921	from a summary and then verifies whether their an-	<i>Sample4: The system's responses</i>	966
922	swers are present in the original source document.	<i>are not coherent with the user's state-</i>	967
923	H Details of Normalizing Process	<i>ments. For example, the user talks about</i>	968
924	We will introduce how to normalize the attention	<i>a ski trip and the system responds with a</i>	969
925	matrix to make it more visually appealing like in	<i>question about French.</i>	970
926	Figure 6. Due to the autoregressive generation		971
927	mode of mainstream LLMs, the expected values of	<i>Sample5: The system maintains a</i>	972
		<i>coherent conversation with the user,</i>	973
		<i>but there are some instances where</i>	974
		<i>the system's responses are not directly</i>	975

976	<i>related to the user's questions. For</i>	<i>You should carefully horizontally com-</i>	1025
977	<i>example, when asked about its age, the</i>	<i>pare the given samples in order to assign</i>	1026
978	<i>system provides a confusing response.</i>	<i>a suitable float score to each sample.</i>	1027
979			
980	<i>Sample6: The system maintains a</i>	<i>Please make sure you read and under-</i>	1028
981	<i>coherent conversation with the user,</i>	<i>stand these instructions carefully. Please</i>	1029
982	<i>responding appropriately to the user's</i>	<i>keep this document open while reviewing,</i>	1030
983	<i>statements and questions.</i>	<i>and refer to it as needed.</i>	1031
984			1032
985	<i>Sample7: The system maintains a</i>	<i>Evaluation Criteria:</i>	1033
986	<i>coherent conversation with the user,</i>		1034
987	<i>but there are some instances where</i>	<i>Coherence (floating point numbers</i>	1035
988	<i>the system's responses are not directly</i>	<i>within the interval [1,3]): Does the</i>	1036
989	<i>related to the user's statements. For</i>	<i>response serve as a valid continuation of</i>	1037
990	<i>example, the user talks about Taylor</i>	<i>the conversation history?</i>	1038
991	<i>Swift and the system responds with a</i>	<i>- A float score near 1 (no) means that</i>	1039
992	<i>statement about Leonardo da Vinci.</i>	<i>the response drastically changes topic or</i>	1040
993		<i>ignores the conversation history.</i>	1041
994	<i>Sample8: The system's responses</i>	<i>- A float score near 2 (somewhat) means</i>	1042
995	<i>are not coherent with the user's state-</i>	<i>the response refers to the conversation</i>	1043
996	<i>ments. For example, the user asks to</i>	<i>history in a limited capacity (e.g., in a</i>	1044
997	<i>see the system's humanlike part and</i>	<i>generic way) and shifts the conversation</i>	1045
998	<i>the system responds with unrelated</i>	<i>topic.</i>	1046
999	<i>statements.</i>	<i>- A float score near 3 (yes) means</i>	1047
1000		<i>the response is on topic and strongly</i>	1048
1001	<i>Sample9: The system maintains a</i>	<i>acknowledges the conversation history.</i>	1049
1002	<i>coherent conversation with the user,</i>		1050
1003	<i>responding appropriately to the user's</i>	<i>Conversations and corresponding</i>	1051
1004	<i>statements and questions.</i>	<i>potential response to be evaluated:</i>	1052
1005			1053
1006	<i>Sample10: The system maintains</i>	<i>{{Data}}</i>	1054
1007	<i>a coherent conversation with the user,</i>	<i>Evaluation Form (Answer by starting</i>	1055
1008	<i>responding appropriately to the user's</i>	<i>with "I will do my best to provide indi-</i>	1056
1009	<i>statements and questions.</i>	<i>vidual analysis for each sample. Analy-</i>	1057
1010		<i>sis:" to analyze the given samples regard-</i>	1058
1011	<i>Float Scores: [Sample1:2.5, Sam-</i>	<i>ing the evaluation criteria as concise</i>	1059
1012	<i>ple2:2.8, Sample3:3, Sample4:1.5,</i>	<i>as possible (Attention: Don't give your</i>	1060
1013	<i>Sample5:2.5, Sample6:3, Sample7:2.5,</i>	<i>scores during this step). After analysing</i>	1061
1014	<i>Sample8:1.5, Sample9:3, Sample10:3]</i>	<i>all the samples, please give all the float</i>	1062
1015		<i>scores in order following the template</i>	1063
1016	J Example Prompts	<i>"Float Scores: [Sample1:score of Sam-</i>	1064
1017	J.1 Evaluate Coherence for Topical-Chat	<i>ple1,...,Sample{{number}}:score of Sam-</i>	1065
1018	default prompt	<i>ple{{number}}]".</i>	1066
1019	<i>You will be given a batch of {{number}}</i>	<i>- Coherence:</i>	1067
1020	<i>samples. Each sample contains a conver-</i>	one stage prompt	1068
1021	<i>sation between Speaker A and Speaker B</i>	<i>You will be given a batch of {{number}}</i>	1069
1022	<i>and one potential response for the next</i>	<i>samples. Each sample contains a conver-</i>	1070
1023	<i>turn.</i>	<i>sation between Speaker A and Speaker B</i>	1071
1024	<i>Your task is to assign a float score to the</i>	<i>and one potential response for the next</i>	1072
	<i>response on one metric.</i>	<i>turn.</i>	1073

1074	<i>Your task is to assign a float score to the response on one metric.</i>		1122
1075			1123
1076	<i>You should carefully horizontally compare the given samples in order to assign a suitable float score to the given samples one by one.</i>		1124
1077			1125
1078			1126
1079			1127
1080	<i>Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.</i>		1128
1081			1129
1082			1130
1083			1131
1084			1132
1085	<i>Evaluation Criteria:</i>		1133
1086	<i>Coherence (floating point numbers within the interval [1,3]): Does the response serve as a valid continuation of the conversation history?</i>		1134
1087			1135
1088	<i>- A float score near 1 (no) means that the response drastically changes topic or ignores the conversation history.</i>		1136
1089			1137
1090			1138
1091	<i>- A float score near 2 (somewhat) means the response refers to the conversation history in a limited capacity (e.g., in a generic way) and shifts the conversation topic.</i>		1139
1092			1140
1093			1141
1094			1142
1095			1143
1096	<i>- A float score near 3 (yes) means the response is on topic and strongly acknowledges the conversation history.</i>		1144
1097			1145
1098			1146
1099			1147
1100	<i>Conversations and corresponding potential response to be evaluated:</i>		1148
1101			1149
1102	<i>Conversations and corresponding potential response to be evaluated:</i>	<i>{{Data}}</i>	1150
1103			1151
1104	<i>{{Data}}</i>		1152
1105			1153
1106	<i>Evaluation Form (Answer by starting with "I will do my best to provide individual analysis and give a suitable float score for each sample in order". When rating for each sample, please follow the template "Score of SampleX:[float score]").</i>	<i>Answer by starting with "I will do my best to provide individual analysis for each sample. Analysis:" to analyze the given samples regarding the evaluation criteria as concise as possible (Attention: Don't give your scores during this step). After analysing all the samples, please horizontally compare the given samples, rank all the samples according to the analysis of the response and give the corresponding reasons. After ranking, according to the analysis and rank, please give all the float scores in order following the template "Float Scores: [Sample1:score of Sample1,...,Sample{{number}}:score of Sample{{number}}]".</i>	1154
1107			1155
1108			1156
1109			1157
1110			1158
1111			1159
1112			1160
1113	<i>- Coherence:</i>		1161
1114	three stage prompt		1162
1115	<i>You will be given a batch of {{number}} samples. Each sample contains a conversation between Speaker A and Speaker B and one potential response for the next turn.</i>		1163
1116			1164
1117			1165
1118			1166
1119			1167
1120	<i>You will be introduced to a metric to be evaluated.</i>	<i>- Coherence:</i>	1168
1121			1169
		Integer prompt	

1170	<i>You will be given a batch of {{number}} samples. Each sample contains a conversation between Speaker A and Speaker B and one potential response for the next turn.</i>	J.2 Evaluate Coherent for FED-Dialogue	1217
1171		default prompt	1218
1172		<i>You will be given a batch of {{number}} samples. Each sample contains a conversation between User and a dialogue System.</i>	1219
1173			1220
1174			1221
1175	<i>Your task is to rate the responses on one metric.</i>		1222
1176		<i>Your task is to assign a float score to the sample on one metric.</i>	1223
1177	<i>You should carefully horizontally compare the given samples in order to assign a score to each sample.</i>		1224
1178		<i>You should carefully horizontally compare the given samples in order to assign a suitable float score to each sample.</i>	1225
1179			1226
1180	<i>Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.</i>		1227
1181		<i>Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.</i>	1228
1182			1229
1183			1230
1184			1231
1185	<i>Evaluation Criteria:</i>		1232
1186	<i>Coherence (1-3): Does the response serve as a valid continuation of the conversation history?</i>	<i>Evaluation Criteria:</i>	1233
1187		<i>Coherent (floating point numbers within the interval [1,3]): Does System maintain coherence and a good flow of conversation throughout the dialogue?</i>	1234
1188			1235
1189	<i>- A score of 1 (no) means that the response drastically changes topic or ignores the conversation history.</i>		1236
1190		<i>- A float score near 1 (not coherent) means that System's responses are unrelated to the conversation topic and may disrupt or confuse the flow of the dialogue.</i>	1237
1191			1238
1192	<i>- A score of 2 (somewhat) means the response refers to the conversation history in a limited capacity (e.g., in a generic way) and shifts the conversation topic.</i>		1239
1193		<i>- A float score near 2 (somewhat coherent) means that System's responses are partially related to the conversation topic but may not be clear or direct.</i>	1240
1194			1241
1195		<i>- A float score near 3 (very coherent) means that System's responses are closely related to the conversation topic and contribute to maintaining a smooth dialogue.</i>	1242
1196	<i>- A score of 3 (yes) means the response is on topic and strongly acknowledges the conversation history.</i>		1243
1197			1244
1198			1245
1199			1246
1200	<i>Conversations and corresponding potential response to be evaluated:</i>	<i>- A float score near 3 (very coherent) means that System's responses are closely related to the conversation topic and contribute to maintaining a smooth dialogue.</i>	1247
1201			1248
1202	<i>{{Data}}</i>	<i>Conversations to be evaluated:</i>	1249
1203		<i>{{Data}}</i>	1250
1204	<i>Evaluation Form (Answer by starting with "I will do my best to provide individual analysis for each sample. Analysis:" to analyze the given samples regarding the evaluation criteria as concise as possible (Attention: Don't give your scores during this step). After analysing all the samples, please give all the scores in order following the template "Scores: [Sample1:score of Sample1,...,Sample{{number}}:score of Sample{{number}}]").</i>		1251
1205			1252
1206		<i>Evaluation Form (Answer by starting with "I will do my best to provide individual analysis for each sample. Analysis:" to analyze the given samples regarding the evaluation criteria as concise as possible (Attention: Don't give your scores during this step). After analysing all the samples, please give all the float scores in order following the</i>	1253
1207			1254
1208			1255
1209			1256
1210			1257
1211			1258
1212			1259
1213			1260
1214			1261
1215			1262
1216	<i>- Coherence:</i>		1263
			1264

1265	<i>template "Float Scores: [Sample1:score of Sample1,...,Sample{{number}}:score of Sample{{number}}]".</i>	<i>individual analysis for each sample.</i>	1312
1266		<i>Analysis:" to analyze the given samples</i>	1313
1267		<i>regarding the evaluation criteria as</i>	1314
1268	<i>- Coherent:</i>	<i>concise as possible (Attention: Don't</i>	1315
		<i>give your scores during this step). After</i>	1316
1269	J.3 Evaluate Coherence for HANNA	<i>analysing all the samples, please give</i>	1317
1270	default prompt	<i>all the float scores in order following the</i>	1318
1271	<i>You will be given a batch of {{num-</i>	<i>template "Float Scores: [Sample1:score</i>	1319
1272	<i>ber}} samples. Each sample contains a</i>	<i>of Sample1,...,Sample{{number}}:score</i>	1320
1273	<i>prompt and a story generated following</i>	<i>of Sample{{number}}]".</i>	1321
1274	<i>the prompt.</i>	<i>- Coherence:</i>	1322
1275	<i>Your task is to assign a float score to</i>		
1276	<i>the story according to the prompt on one</i>	J.4 Evaluate Factual Consistency for QAGS	1323
1277	<i>metric.</i>	default prompt	1324
1278	<i>You should carefully horizontally com-</i>	<i>You will be given a batch of {{number}}</i>	1325
1279	<i>pare the given samples in order to assign</i>	<i>samples. Each sample contains an arti-</i>	1326
1280	<i>a suitable float score to each sample.</i>	<i>cle and a sentence.</i>	1327
1281	<i>Please make sure you read and under-</i>	<i>Your task is to determine if the sentence</i>	1328
1282	<i>stand these instructions carefully. Please</i>	<i>is factually correct given the contents of</i>	1329
1283	<i>keep this document open while reviewing,</i>	<i>the article.</i>	1330
1284	<i>and refer to it as needed.</i>	<i>You should carefully horizontally com-</i>	1331
1285		<i>pare the given samples in order to assign</i>	1332
1286	<i>Evaluation Criteria:</i>	<i>a suitable float score to each sample.</i>	1333
1287	<i>Coherence (floating point numbers</i>	<i>Please make sure you read and under-</i>	1334
1288	<i>within the interval [1,5]) Measures</i>	<i>stand these instructions carefully. Please</i>	1335
1289	<i>whether the story makes sense?</i>	<i>keep this document open while reviewing,</i>	1336
1290	<i>- A float score near 1 means the story</i>	<i>and refer to it as needed.</i>	1337
1291	<i>does not make sense at all. For instance,</i>		1338
1292	<i>the setting and/or characters keep chang-</i>	<i>Evaluation Criteria:</i>	1339
1293	<i>ing, and/or there is no understandable</i>	<i>Consistency ([1,3]) - Is the sentence sup-</i>	1340
1294	<i>plot.</i>	<i>ported by the article? (consistent with</i>	1341
1295	<i>- A float score near 2 means most of the</i>	<i>the article)</i>	1342
1296	<i>story does not make sense.</i>	<i>- A float score near 1 (not) means that the</i>	1343
1297	<i>- A float score near 3 means the story</i>	<i>sentence is totally not supported by the</i>	1344
1298	<i>mostly makes sense but has some inco-</i>	<i>article.</i>	1345
1299	<i>herences.</i>	<i>- A float score near 2 (somewhat) means</i>	1346
1300	<i>- A float score near 4 means the story</i>	<i>that the sentence is partially supported</i>	1347
1301	<i>almost makes sense overall, except for</i>	<i>by the article.</i>	1348
1302	<i>one or two small incoherences.</i>	<i>- A float score near 3 (very) means that</i>	1349
1303	<i>- A float score near 5 means the story</i>	<i>the sentence is completely supported by</i>	1350
1304	<i>makes sense from beginning to end.</i>	<i>the article.</i>	1351
1305			1352
1306	<i>Prompts and corresponding stories to be</i>	<i>Articles and corresponding sentences to</i>	1353
1307	<i>evaluated:</i>	<i>be evaluated:</i>	1354
1308	<i>{{Data}}</i>	<i>{{Data}}</i>	1355
1309			1356
1310	<i>Evaluation Form (Answer by start-</i>	<i>Evaluation Form (Answer by start-</i>	1357
1311	<i>ing with "I will do my best to provide</i>	<i>ing with "I will do my best to provide</i>	1358

1359 *individual analysis for each sample.*
1360 *Analysis:" to analyze the given samples*
1361 *regarding the evaluation criteria as*
1362 *concise as possible (Attention: Don't*
1363 *give your scores during this step). After*
1364 *analysing all the samples, please give*
1365 *all the float scores in order following the*
1366 *template "Float Scores: [Sample1:score*
1367 *of Sample1,...,Sample{{number}}:score*
1368 *of Sample{{number}}]".*

1369 - Consistency: