

Knockout LLM Assessment: Using Large Language Models for Evaluations through Iterative Pairwise Comparisons

Isik Baran Sandan, Tu Anh Dinh, Jan Niehues
Karlsruhe Institute of Technology

Abstract

Large Language Models (LLMs) have shown to be effective evaluators across various domains such as machine translations or the scientific domain. Current LLM-as-a-Judge approaches rely mostly on individual assessments or a single round of pairwise assessments, preventing the judge LLM from developing a global ranking perspective. To address this, we present Knockout Assessment, an LLM-as-a-Judge method using a knockout tournament system with iterative pairwise comparisons. Experiments across three LLMs on two datasets show that knockout assessment improves scoring accuracy, increasing Pearson correlation with expert evaluations by 0.07 on average for university-level exam scoring and machine translation evaluations, aligning LLM assessments more closely with human scoring.

1 Introduction

Across various domains, and especially for scientific research, accurate and consistent evaluations are very crucial for informed decision-making. However, the inherent scale and subjectivity make this task very challenging and time-consuming. In recent years, the methodology of “LLM-as-a-Judge” (Zheng et al., 2023) has emerged to tackle this challenge, where instead of humans, Large Language Models (LLMs) take the role of the expert to evaluate complex tasks. Using LLMs as evaluators allows us to mimic the abilities of human experts, making evaluations cost-effective and scalable.

Although many approaches to LLM-as-a-Judge exist, the most common is individual assessment, in which the evaluation prompt consists of only the question and the corresponding answer, which is to be evaluated (Chiang and Lee, 2023). While this approach has already shown to yield good evaluation results next to providing scalability (Chiang and Lee, 2023; Dinh et al., 2024), it does not consider

the relative strength of answers in a set to a given question. The more recent approach of pairwise assessment tries to address this issue by providing two responses to the judge LLM each time, however, it still fails to account for a global ranking perspective, as pairwise comparisons do not analyze how all responses compare to each other in the broader sense.

In this paper we present an LLM-as-a-judge method called **Knockout Assessment** to address this challenge, which can be seen as a variation of the tournament system used by Zheng et al. (2023), differing in that it makes use of iterative pairwise comparisons. Instead of isolating responses individually or in pairs for evaluation, Knockout Assessment focuses on an iterative process where responses are compared against one another multiple times in a tournament manner. In each round, stronger responses advance to compete against each other in later rounds, allowing us to refine the scores progressively throughout the tournament. This approach allows the judge LLMs to develop a global perspective on responses without requiring all replies to be included in a single prompt, which would otherwise result in an impractically long context length.

To summarize, our contributions are as follows:

- Knockout Assessment, an LLM-as-a-Judge methodology which makes use of iterative pairwise comparisons in a knockout tournament system for more accurate evaluations
- Analysis of Knockout Assessment’s performance compared to individual assessment’s and naive pairwise assessment’s performance on two different datasets concerning scientific evaluation and machine translation evaluation.

2 Related Work

Individual Assessment One approach to LLM-as-a-Judge is individual assessment, where the

judge LLM is provided with a prompt or a question, the corresponding answer, the scoring criteria and is asked to provide an evaluation such as a grade. Various studies have used this method for tasks such as evaluating story generation, scoring quality of different texts according to different criteria, or grading university-level exams (Chiang and Lee, 2023; Wang et al., 2023; Dinh et al., 2024).

Pairwise Assessment Another more recent LLM-as-a-Judge approach is pairwise prompting, in which the LLM judge is provided with two responses to the prompt instead of one. The judge LLM is then asked to evaluate both responses. This has shown to be an effective LLM-as-a-Judge method for ranking documents, as it gives the judge LLM direct comparison points while making evaluations (Liusie et al., 2024). However, this method still does not make the Judge LLM develop a global ranking perspective.

Chatbot Arena Approach Zheng et al. (2023) made use of an ELO system in which all possible answer pairs are evaluated against one another. This approach is thus able to make implicit use of a global view over the dataset while assigning scores. However, pairing all possible answers results in a computational time of $O(N^2)$.

Sorting Based Approaches To address this inefficiency, Qin et al. (2024) introduced two new methods. First approach uses Heapsort with pairwise comparisons to sort out the possible answers ($O(N \log N)$). Second is a sliding window approach, making use of individual passes in the Bubble Sort algorithm for a constant number K times ($O(N)$).

3 Knockout Assessment

We propose using multiple iterative pairwise comparisons instead of individual assessment with Knockout Assessment. In each pairwise assessment, one pairwise ranking prompt similar to the comparative prompt introduced by Liusie et al. (2024) is used. In each prompt, one question and two answers to that question are provided to the judge LLM, which is asked to evaluate both of those answers. We call this a “question-level-match”. The exact prompts we used for our experiments can be found in Appendix A.2.

From the response generated by the judge LLM, the score each individual answer got is parsed and saved to the list of scores for that answer, which

keeps track of all the scores an answer got throughout all its question-level-matches. The answer which got the higher score advances to the next round to be matched up against another answer.

The order of texts in pairwise rankings has shown to be an influential factor in the LLMs decision making (Resnik, 2024), thus we also collected the results with using a debiasing methodology similar to the one introduced by Liusie et al. (2024), averaging scores from both possible orderings of each answer pair. Debiasing thus results in double the compute-time compared to a regular question-level-match.

The main methodology behind our approach is a knockout tournament system that iteratively uses the question-level matches. In each tournament round, the N available answers to a question are randomly assigned to pairs. Each pair then enters a question-level match, and the higher scoring response advances to the next round. In the case when N is odd, one answer directly advances to the next round. This continues until we reach a tournament round with a single response.

Once the tournament ends, the final evaluation score for each answer is computed as the average of all the scores it received throughout the tournament. An example tournament with $N = 4$ answers is depicted in Figure 1. The full algorithm is given in A.1.

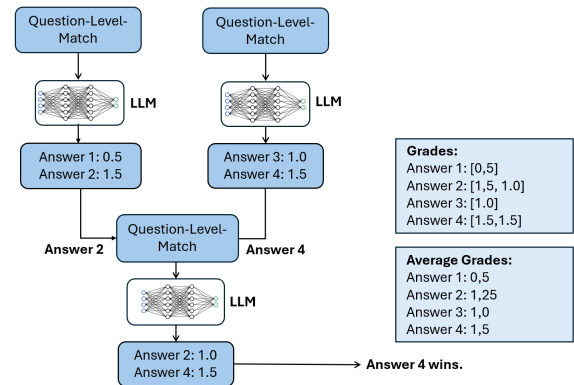


Figure 1: An example Knockout-Tournament with 4 answers for a question.

4 Experiments

Datasets Overall, the datasets we used include the task input, machine outputs, and human-assigned scores on the machine output. The first dataset is SciEx (Dinh et al., 2024), which consists of university exam questions, LLM answers,

human grades and LLM grades. Each question is labeled by difficulty and language. SciEx contains 1120 question-answer pairs in total. The second dataset is taken from the WMT Metrics Shared Task (Kocmi et al., 2024), which includes a list of source sentences and machine translations, accompanied by a human evaluation on a scale of 0 to 100. We filtered the dataset by languages supported by the Llama models (English, German, French, Italian, Portuguese, Hindi, Spanish, and Thai). We also remove identical translations of the same source sentence, as sometimes the human evaluations differed for the same translation, causing inconsistencies. The final processed WMT dataset includes 2100 source-translation pairs.

Baselines We compare our approach to 2 baselines: (1) individual assessment, where each answer is scored individually and (2) naive pairwise assessment, where answers are paired up and scored only one time, without multiple knockout rounds.

Evaluation Metrics We use Pearson correlation as our primary evaluation metric, which measures the linear relationship between the scores provided by our method and the human-provided scores. Additionally, we use pairwise ranking accuracy (Kocmi et al., 2021), which indicates how often the LLM judges select the correct winning answer given a pair of answers. The results for pairwise ranking accuracy can be found in Appendix A.3 since they mirror our Pearson correlation findings.

Models The models we used as judges for our experiments are Meta’s Llama 3.2 1B parameter model, Llama 3.2 3B parameter model and Llama 3.1 70B parameter model. All the model checkpoints for our experiments were obtained from the HuggingFace model hub. For consistency with similar work (Dinh et al., 2024) the temperature parameter of the models was set to 0.1 for our experiments.

Hardware The experiments on the 70B parameter model were conducted on 4 NVIDIA Tesla V100 GPUs with 32GB VRAM each. The experiments for the smaller models were conducted on 1 NVIDIA Tesla V100 GPU with 32GB VRAM.

4.1 General Results

From Table 1, it can be seen that our knockout assessment method improves performance on all datasets and model judges compared to individual

assessment. Generally, the scores sampled using debiasing have improved performance even further.

One failing case is on the WMT subset, where the best performing assessment method is to use individual assessment with Llama 3.1 70B. In this case, knockout assessment has decreased performance for the largest model. However, smaller models still saw an increase in performance with knockout assessment compared to individual assessment. Overall, we see an average increase of 0.07 Pearson correlation over individual assessment across all our experiments, when debiased knockout assessment is used.

One possible reason for this is that, Knockout Assessment is more useful when the evaluation task is more complex. Therefore, it consistently helps improving the assessment performance on the SciEx dataset, which contains difficult university-level scientific question-answer pairs. On the other hand, evaluating the machine translation task is more simple, thus a large model like Llama 3.1 70B can have good performance with just individual assessment. In this case, introducing other answers with Knockout Tournament could introduce noise, thus lower the performance.

4.2 Comparison Against Naive Pairwise Assessment

In this section, we check whether knockout assessment results in any additional performance increase compared to regular pairwise comparisons with only one round. We report the Pearson correlations of the sets of scores the answers got, based on their round of elimination.

For both datasets, the answers/translations which got eliminated on the first round of the knockout tournament, got only one pairwise comparison, compared to the multiple pairwise comparisons the answers/translations which advanced to later rounds got. As can be seen in Table 2, for SciEx, the answers which got eliminated in later rounds have an overall higher Pearson correlation in the grades they got, across the three models we used. This shows that more pairwise comparisons result in more accurate grades from the judge LLM.

However, as can be seen in Table 2, the responses which advanced further in the tournament showed lower alignment with human experts in the WMT dataset. This suggests that the iterative comparisons do not increase the scoring accuracy for the task of machine translation, but rather introduce noises.

Method	SciEx Question-Level			SciEx Exam-Level			WMT Dataset			Overall
	3.2 1B	3.2 3B	3.1 70B	3.2 1B	3.2 3B	3.1 70B	3.2 1B	3.2 3B	3.1 70B	
Ind. Assessment	0.400	0.365	0.615	0.504	0.465	0.667	0.050	0.187	0.397	0.405
KO. Assessment (No Debiasing)	0.434	0.541	0.622	0.627	0.555	0.652	0.113	0.207	0.222	0.441
KO. Assessment (Debiased)	0.443	0.558	0.648	0.672	0.540	0.697	0.087	0.259	0.268	0.475

Table 1: LLM scores’ Pearson correlation to expert scores for different datasets, subdivided by models

Dataset	Elimination	Knockout Assessment							
		1B		3B		70B		Overall	
		Biased	Debiased	Biased	Debiased	Biased	Debiased	Biased	Debiased
SciEx	First Round	0.3737	0.3223	0.5400	0.5400	0.5264	0.5801	0.4800	0.4808
	Later Rounds	0.4816	0.4428	0.5393	0.5692	0.6602	0.6782	0.5604	0.5634
	Difference	+0.1079	+0.1205	-0.0007	+0.0292	+0.1338	+0.0981	+0.0804	+0.0826
WMT	First Round	0.1245	0.0836	0.2222	0.2917	0.2688	0.3173	0.2052	0.2309
	Later Rounds	0.0777	0.0875	0.1603	0.2013	0.1150	0.1581	0.1177	0.1490
	Difference	-0.0468	+0.0039	-0.0619	-0.0904	-0.1538	-0.1592	-0.0875	-0.0819

Table 2: Comparison of LLM Grader’s performance on SciEx and WMT datasets for answers graded once versus multiple times.

4.3 Effect of Difficulty Levels

We investigate how the difficulty level of the task effect the performance of our assessment method. We use the question-level difficulty labels from SciEx, and report the performance splitted by the labels of "Easy", "Medium" and "Hard". The results are shown in Figure 2. As can be seen, for individual assessment, the models perform better on scoring answers of difficult questions than easier questions, which is rather counter-intuitive. However, this aligns with the finding by Dinh et al. (2024) that LLMs can perform worse on easy questions, since they may lack specific course knowledge compared to the students.




Figure 2: Performance by Model and Difficulty Level: Knockout (debiased) vs. Individual Assessment

With our Knockout Assessment method, the performance of the judges on easy questions significantly increases. This shows that, by including

information of multiple candidate answers, we give more global view to the LLMs, thus help them to better provide assessment scores, even when their own internal knowledge for the question is lacking.

5 Conclusion

We address a key limitation of many existing LLM-as-a-Judge methods: not having a global view over the responses while evaluating them. To address this, we proposed an alternative LLM-as-a-Judge method called knockout assessment and tested it with three different LLMs on two different datasets. Knockout Assessment improves Pearson correlation to human evaluations by 0.07 over individual assessment on average. The performance increase was more significant in scientific evaluation compared to machine translation evaluation, especially for the larger LLM. Furthermore, for scientific evaluation, the responses which progressed further in the knockout assessment process had 0.08 better accuracy compared to the responses which got eliminated on the first round, which indicates that knockout assessment results in a performance increase from regular pairwise assessments.

Acknowledgments

This work was supported by the Helmholtz Programme-oriented Funding, with project number

46.24.01, project name AI for Language Technologies. We acknowledge the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the Federal Ministry of Education and Research.

References

- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Tu Anh Dinh, Carlos Mullov, Leonard Bärmann, Zhaolin Li, Danni Liu, Simon Reiß, Jueun Lee, Nathan Lerzer, Fabian Ternava, Jianfeng Gao, Tobias Röddiger, Alexander Waibel, Tamim Asfour, Michael Beigl, Rainer Stiefelhagen, Carsten Dachsbacher, Klemens Böhm, and Jan Niehues. 2024. [Sciex: Benchmarking large language models on scientific exams with human expert grading and automatic grading.](#) *Preprint*, arXiv:2406.10421.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Christof Monz, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popović, and 3 others. 2024. [Findings of the WMT24 general machine translation shared task: The LLM era is here but MT is not solved yet.](#) In *Proceedings of the Ninth Conference on Machine Translation*, pages 1–46, Miami, Florida, USA. Association for Computational Linguistics.
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To ship or not to ship: An extensive evaluation of automatic metrics for machine translation.](#) In *Proceedings of the Sixth Conference on Machine Translation*, pages 478–494, Online. Association for Computational Linguistics.
- Adian Liusie, Potsawee Manakul, and Mark J. F. Gales. 2024. [Llm comparative assessment: Zero-shot nlg evaluation through pairwise comparisons using large language models.](#) *Preprint*, arXiv:2307.07889.
- Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky. 2024. [Large language models are effective text rankers with pairwise ranking prompting.](#) *Preprint*, arXiv:2306.17563.
- Philip Resnik. 2024. [Large language models are biased because they are large language models.](#) *Preprint*, arXiv:2406.13138.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023. [Is chatgpt a good nlg evaluator? a preliminary study.](#) *Preprint*, arXiv:2303.04048.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena.](#) *Preprint*, arXiv:2306.05685.

A Appendix

A.1 Algorithm for Knockout Assessment

Algorithm 1 Knockout Assessment

Input: prompt P , Set of responses \mathcal{R} , scoring function G (evaluates response quality) **Output:** Final Champion Response, All Scores, Final Average Scores

```
1: while  $|\mathcal{R}| > 1$  do
2:   Form consecutive pairs of responses from  $\mathcal{R}$ 
3:   for  $k = 1, 3, 5, \dots, |\mathcal{R}| - 1$  do  $\triangleright$  Match responses in consecutive pairs
4:     (Score1, Score2)  $\leftarrow$  Question_Level_Match( $P, G, r_k, r_{k+1}$ )
5:     Update All Scores with Score1 and Score2
6:     if Score1 > Score2 then
7:       Advance  $r_k$  to the next round
8:     else
9:       Advance  $r_{k+1}$  to the next round
10:    end if
11:  end for
12:  if  $|\mathcal{R}|$  is odd then
13:    Advance the unmatched response to the next round
14:  end if
15: end while
16: Compute Final Average Scores for each response across rounds
17: answer: Final Champion  $r^*$  (last remaining response), Final Average Scores, All Scores
```

A.2 Prompts

This section includes all the prompts we used to generate scores for our experiments.

Individual Grading Prompt in English: You are a university professor exam grader. Grade the following answer on a scale of 0 to <maxpoints> (allowing half points) based on its correctness and relevancy given the following question.

The Question: <Question>

The Answer: <Answer 1>

Give your grade in the format: Explanation: [explanation] Score: [score]/<maxpoints>

Pairwise Grading Prompt in English: You are a university professor exam grader. Grader the following answers on a scale of 0 to <maxpoints>

(allowing half points) based on how well they answer the question.

The Question: <Question>

Answer 1: <Answer 1>

Answer 2: <Answer 2>

Grade the 2 answers on a scale of 0 to <maxpoints> (Half points such as 0.5 or 1.5 are allowed.) in the format: Explanation: [explanation] Answer 1: X/<maxpoints> Answer 2: Y/<maxpoints>

Individual Grading Prompt in German: Sie sind ein Universitätsprofessor. Bewerten Sie die folgende Antwort auf die unten stehende Frage. Geben Sie eine Punktzahl von 0 bis maxpoints basierend auf Korrektheit und Relevanz an.

Die Frage: <question>

Die Antwort: <Answer 1>

Bewerten Sie die Antworten auf einer Skala von 0 bis maxpoints (halbe Punkte wie 0,5 oder 1,5 sind erlaubt) im Format: Begründung:[begründung] Punktzahl: X/<maxpoints>

Pairwise Grading Prompt in German: Sie sind ein Universitätsprofessor und bewerten Prüfungsantworten. Bewerten Sie die folgenden Antworten auf einer Skala von 0 bis maxpoints (halbe Punkte sind erlaubt) basierend darauf, wie gut sie die Frage beantworten.

Die Frage: question

Antwort 1: answer1

Antwort 2: answer2

Bewerten Sie die beiden Antworten auf einer Skala von 0 bis maxpoints (halbe Punkte wie 0,5 oder 1,5 sind erlaubt) im Format: Begründung: [begründung] Antwort 1: X/maxpoints Antwort 2: Y/maxpoints.

Individual Scoring Prompt for MT: You are a translation evaluator. Evaluate the quality of the translation provided. Give a score from 0 to 100 based on clarity, accuracy and grammar.

Source: <source>

Translation: <tgt>

Output only: : Explanation: [explanation] Score: [score]/100

Pairwise Scoring Prompt for MT: You are a translation evaluator. Your task is to evaluate the quality of two translations for a given source sentence. You will provide a score from 0 to 100, based solely on clarity, accuracy and grammar of the translations.

Source: <source>

Translation 1: <tgt1>

Translation 2: <tgt2>

Output only: Explanation: [explanation] Translation 1: [score]/100 Translation 2: [score]/100

A.3 Pairwise Ranking Accuracy Results

The performance of each evaluation method with pairwise ranking accuracy as an evaluation metric can be seen in Table 3, divided by the grading LLM.

Method	SciEx Exam Level			Overall
	3.2 1B	3.2 3B	3.1 70B	
Ind. Assessment	0.529	0.533	0.695	0.586
KO. Assessment	0.557	0.543	0.676	0.592
Debiased KO. Assessment	0.610	0.591	0.767	0.656

Table 3: LLM scores’ pairwise ranking accuracy to expert scores for different methods

A.4 Influential Factors for SciEx

Our findings for the impact of knockout assessment on different examinees and different languages can be seen in Tables 4 and 5

Examinee	LLama Models			Overall
	1B	3B	70B	
Llava	+0.0150	-0.0040	+0.0724	+0.0283
Mistral	+0.2452	+0.1043	-0.0247	+0.1082
Mixtral	-0.1247	+0.0407	-0.0776	-0.0539
Qwen	-0.1263	+0.0184	-0.0106	-0.0395
Claude	-0.2975	+0.0760	-0.0018	-0.0744
GPT-3.5	-0.0310	+0.1943	+0.0228	+0.0620
GPT-4V	+0.1903	+0.3529	+0.0046	+0.1826

Table 4: Performance difference with knockout assessment vs. individual assessment, by examinee and model.

Language	Individual			Knockout					
	1B	3B	70B	1B		3B		70B	
				Biased	Debiased	Biased	Debiased	Biased	Debiased
English	0.2348	0.176	0.6759	0.5695	0.5691	0.6365	0.6892	0.6429	0.6952
German	0.5474	0.5451	0.6263	0.3706	0.3824	0.5456	0.5628	0.6497	0.6790

Table 5: Pearson correlations for LLM graders’ performance across languages (English and German), for individual and knockout assessment.