

Exploring the Latest LLMs for Leaderboard Extraction

Notebook for the SimpleText Task4: SOTA? Lab at CLEF 2024

Salomon Kabongo^{1,*,\dagger}, Jennifer D'Souza^{1,2,\dagger} and Sören Auer^{1,2}

¹Leibniz University of Hannover, Hannover, Germany

²TIB Leibniz Information Centre for Science and Technology, Hannover, Germany

Abstract

The rapid advancements in Large Language Models (LLMs) have opened new avenues for automating complex tasks in AI research. This paper investigates the efficacy of different LLMs-Mistral 7B, Llama-2, GPT-4-Turbo and GPT-4.o in extracting leaderboard information from empirical AI research articles. We explore three types of contextual inputs to the models: DocTAET (Document Title, Abstract, Experimental Setup, and Tabular Information), DocREC (Results, Experiments, and Conclusions), and DocFULL (entire document). Our comprehensive study evaluates the performance of these models in generating (Task, Dataset, Metric, Score) quadruples from research papers. The findings reveal significant insights into the strengths and limitations of each model and context type, providing valuable guidance for future AI research automation efforts.

Keywords

Information Extraction, LLMs, Leaderboard, LLAMA-2, MISTRAL, GPT-4-Turbo, GPT-4.o

1. Introduction

The unprecedented expansion of scientific publications [1, 2] in the field of artificial intelligence (AI) has created a substantial challenge in systematically tracking and assessing advancements. Leaderboards, which rank AI models based on their performance across various tasks and datasets, serve as a crucial mechanism for monitoring progress and driving competition. The extraction of state-of-the-art information, represented by (Task, Dataset, Metric, Score) quadruples, is vital for the maintenance and accuracy of these leaderboards.

The effectiveness of LLMs in information extraction (IE) tasks is deeply influenced by the context provided during the input phase. In this context, “context” refers to the segments of the scholarly articles from which the (Task, Dataset, Metric, Score) data is extracted. Unlike the concept of in-context learning, where models are primed with examples, our focus is on selecting relevant sections of text to guide the LLMs toward accurate and relevant information extraction while minimizing the risk of generating irrelevant or false information (hallucinations).

Existing research has shown that longer contextual inputs can lead to decreased model accuracy and higher rates of hallucination due to the potential distraction of irrelevant information [3, 4]. Therefore, identifying the optimal length and specificity of context is crucial [3]. This study explores three distinct types of contextual inputs to empirically evaluate their impact on the task of extracting leaderboards: DocTAET (Document Title, Abstract, Experimental Setup, and Tabular Information), DocREC (Results, Experiments, and Conclusions), and DocFULL (the entire document).

Our primary contributions are twofold. Firstly, we evaluate the performance of the latest LLMs, both open-sourced Llama-2 and Mistral 7B [5, 6] and proprietary GPT-4-Turbo and GPT-4.o [7], in generating structured summaries and classifying papers with and without leaderboard data. Our first research question (RQ1) encapsulates this: Which LLM provides the most accurate performance for generating structured summaries and for leaderboard/no-leaderboard classification?

CLEF 2024: Conference and Labs of the Evaluation Forum, September 09–12, 2024, Grenoble, France

*Corresponding author.

^{\dagger}These authors contributed equally.

✉ kabenamalu@l3s.de (S. Kabongo); jennifer.dsouza@tib.eu (J. D'Souza); auer@tib.eu (S. Auer)

🌐 <https://skabongo.github.io/> (S. Kabongo)

🆔 0000-0002-0021-9729 (S. Kabongo); 0000-0002-6616-9509 (J. D'Souza); 0000-0002-0698-28646 (S. Auer)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

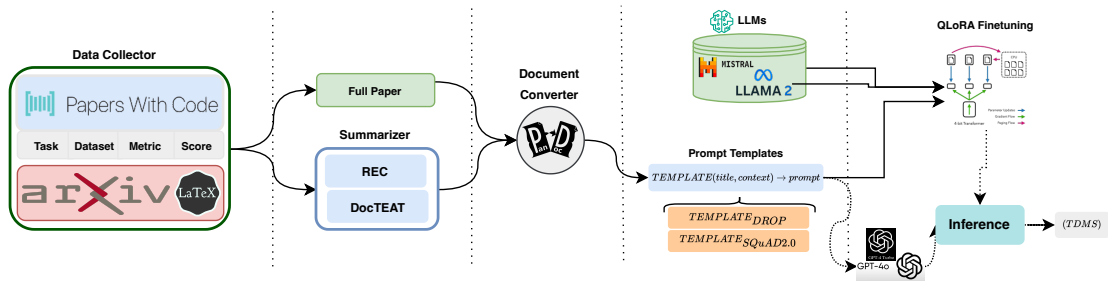


Figure 1: Main process overview

Secondly, we focus on the precise extraction of individual elements, addressing our second research question (RQ2): Which LLM offers the best trade-off between precision and other performance metrics in few-shot and zero-shot settings? This question is critical, given the importance of precision in scholarly communications and the broader implications for the reliability of model outputs.

Additionally, we participate in the “SOTA? *Tracking the State-of-the-Art in Scholarly Publications*” shared Task 4 in the SimpleText¹ track of CLEF 2024 [8, 9]. The goal of the SOTA? shared task is to develop systems that, given the full text of an AI paper, can recognize whether the paper reports model scores on benchmark datasets and, if so, extract all pertinent (Task, Dataset, Metric, Score) tuples presented within the paper.

2. Related Work

The aspiration for automatically generated leaderboards to monitor advancements in scientific research has been a notable ambition within the scientific community. This initiative first gained traction through the analysis of citation networks, employing then-cutting-edge methodologies such as Markov Random Fields and others [10, 11].

The advent of transformer-based models [12] marked a significant leap forward, setting new benchmarks across a myriad of machine learning tasks. In this vein, researchers at IBM [13] utilized the Bert pre-trained model [14] within a Natural Language Inference (NLI) framework to discern entailment from the complete texts of research papers to their corresponding leaderboards.

In the realm of information extraction, the application of LLMs such as GPT-3.5 and LLAMA has demonstrated significant promise. A study by [15] highlighted the capabilities of GPT-3.5 and GPT-4 in clinical named entity recognition (NER), showcasing their adeptness at processing intricate clinical datasets with limited prerequisite training. Through strategic prompt engineering, these models exhibited remarkable improvements in performance for extracting medical entities from clinical documentation, reinforcing the potential of LLMs in executing complex NER tasks within the healthcare sector [15].

The emergence of LLMs, including ChatGPT, has sparked a reconsideration of specialized versus general-purpose training approaches in the context of LLMs. Building upon the foundation of utilizing LLMs for specialized information extraction tasks, [16] delves into the utilization of these models in the domain of virology. The study showcases how LLMs, specifically tuned for scientific content, can efficiently parse and extract virology-related information from a plethora of scientific publications. This research underlines the importance of fine-tuning and prompt engineering in enhancing the model’s ability to discern relevant scientific facts, contributing to the development of domain-specific leaderboards. The approach exemplified in this paper demonstrates an effective strategy for context selection in LLM-based leaderboard generation, emphasizing the necessity for domain-specific adjustments to maximize the accuracy and relevance of the extracted information.

In our previous study [17], we expanded on these findings by empirically investigating different

¹<https://simpletext-project.com/2024/en/>

ways to select context in creating leaderboards. We analyze the impact of tailored context cues and the integration of domain-specific knowledge on the precision and utility of automatically generated leaderboards. Our contributions included a thorough assessment of the latest LLMs, both open-source and proprietary, for assessing the effectiveness of different context selection strategies and the development of a novel methodology that significantly enhances the accuracy and relevance of LLM-based leaderboard generation. This study not only corroborates the pivotal role of context in leveraging LLMs for information extraction but also introduces innovative techniques that refine the process of generating insightful and reliable leaderboards in the scientific community.

Table 1

Our DocREC (Documents Result[s], Experimentation[s] and Conclusion) corpora statistics. The “papers w/o leaderboard” refers to papers that do not report leaderboard.

	Our Corpus		
	Train	Test-Few-shot	Test Zero-shot
Papers w/ leaderboards	7,987	753	241
Papers w/o leaderboards	4,401	648	548
Total TDM-triples	415,788	34,799	14,800
Distinct TDM-triples	11,998	1,917	1,267
Distinct <i>Tasks</i>	1,374	322	236
Distinct <i>Datasets</i>	4,816	947	647
Distinct <i>Metrics</i>	2,876	654	412
Avg. no. of TDM per paper	5.12	4.81	6.11
Avg. no. of TDMS per paper	6.95	5.81	7.86

3. Methodology

This section details the methodology employed in our study, encompassing data collection, preprocessing, model selection, and evaluation metrics. The goal is to systematically evaluate the performance of four state-of-the-art LLMs—Mistral 7B, Llama-2, GPT-4-Turbo and GPT-4.o—across different context types for the task of leaderboard extraction.

3.1. Data Collection and Preprocessing

We utilized data that was previously published by [17] and re-released as the SOTA? shared task training corpus. The dataset consists of (T, D, M, S) annotations for thousands of AI articles available on PwC (CC BY-SA). These articles cover various AI domains such as Natural Language Processing, Computer Vision, Robotics, Graphs, Reasoning, etc., making them representative for empirical AI research. The specific PwC source was downloaded on December 09, 2023. The corpus comprised over 8,000 articles, with 7,987 used for training and 994 for testing, including 751 in the few-shot setting and 241 in the zero-shot setting. These articles, originally sourced from arXiv under CC-BY licenses, are available as latex source code, each accompanied by one or more (T, D, M, S) annotations from PwC. The articles’ metadata was directly obtained from the PwC data release, and the articles collection had to be reconstructed by downloading them from arXiv under CC-BY licenses.

After downloading the article’s source code (.tex), we needed to preprocess it to convert it to plain text. Sometimes, articles written in LaTeX are split into multiple files. To address this, we first created and executed a custom script to merge the project source code into a single LaTeX file corresponding to the arXiv ID of the paper. Next, we used another custom script to extract specific sections of the paper (DocTEAT or DocREC) from the ``arxiv_id.tex`` file, ensuring that the file remained compilable by LaTeX, which is necessary for our ``tex_to_text`` parsing process.

To convert the resulting all-in-one `arxiv_id.tex` file to plain text, we used the command `pandoc -to = plain`. Subsequently, each article’s parsed text was annotated with (T, D, M, S) quadruples via distant labeling. The overall corpus statistics are reported in Table 1.

Another important subset of our data, in addition to our base dataset reported in Table 1, was the “no leaderboards papers”. We included a set of approximately 4,401 and 648 articles that do not report leaderboards into the train and test sets, respectively. These articles were randomly selected by leveraging the arxiv category feature, then filtering it to papers belonging to domains unrelated to AI/ML/Stats. These articles were annotated with the *unanswerable* label to finetune our language model in recognizing papers without (T,D,M,S) mentions in them.

We phrased the following question to formulate our task objective w.r.t. the (T, D, M, S) extraction target: *What are the values for the following properties to construct a Leaderboard for the model introduced in this article: task, dataset, metric, and score?* In essence, it encapsulates an IE task.

Instruction tuning [18, 19, 20, 21, 22, 23] boosts LLMs by providing specific finetuning instructions, improving adaptability and performance on new tasks [24, 25]. This method offers a more efficient approach than traditional unlabeled data methods [26, 27, 28, 29], allowing for versatile task prompting with single instructions.

In this vein, the “Flan 2022 Collection” [23] was a large-scale open-sourced collection of 62 prior publicly released datasets in the NLP community clustered as 12 task types, such as reading comprehension (RC), sentiment, natural language inference (NLI), struct to text, etc. It is the most comprehensive resource facilitating open-sourced LLM development as generic multi-task models. Importantly, and of relevance to this work, FLAN was not just a super-amalgamation of datasets encapsulating different learning objectives, but also included at least 10 human-curated natural instructions per dataset that described the task for that dataset. As such, we select a set of instructions to guide the LLM for our complex IE task from the FLAN collection. Specifically, we identified the applicable instructions to our task were those designed for the SQuAD_v2 [30, 31] and DROP [32] datasets. Specifically, 8 SQuAD and 7 DROP instructions were found suitable. The general characteristic of the selected instructions is that they encode a context and the SOTA task objective, and instruct the model to fulfill the objective. The context, in our case, is a selection from specific sections from the full-text of an article where (T, D, M, S) information is most likely shared. This is discussed next.

As introduced in our prior work [17], we compared the performances of the four state-of-the-art LLMs on the following three contexts:

3.2. DocTAET

Delineated in prior work [13], the context to the LLM comprises text selected from the (T)-title, (A)-abstract, (E)-experimental setup, and (T)-tabular information parts of the full-text. It yields an average context length of 493 words, ranging from a minimum of 26 words to a maximum of 7,361 words. These specific selections targeted the areas of the paper where the (T, D, M, S) were most likely to be found. An example of this context selection is illustrated in Figure 2.

3.3. DocREC

Introduced for the first time in this work, the DocREC context comprises text selected from the sections named (R)-results, (E)-experiments, and (C)-conclusions with allowances for variations in the three names. Complementary but still unique to DocTAET, the DocREC context representation aims to distill the essence of the research findings and conclusions into a succinct format. This context, ended by being much longer than DocTAET, yielded an average length of 1,586 words, with a minimum length of 27 words and a maximum length of 127,689 words. An example of this context selection is illustrated in Figure 3.



{'id': '2306.00834v1.pdf', 'prompt': 'Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network for Motion Deblurring' Event cameras differ from conventional RGB cameras in that they produce asynchronous data sequences. While RGB cameras capture every frame at a fixed rate, event cameras only capture changes in the scene, resulting in sparse and asynchronous data output. Despite the fact that event data carries useful information that can be utilized in motion deblurring of RGB cameras, integrating event and image information remains a challenge. Recent state-of-the-art CNN-based deblurring solutions produce multiple 2-D event frames based on the accumulation of event data over a time period. In most of these techniques, however, the number of event frames is fixed and predefined, which reduces temporal resolution drastically, particularly for scenarios when fast-moving objects are present or when longer exposure times are required. It is also important to note that recent modern cameras (e.g., cameras in mobile phones) dynamically set the exposure time of the image, which presents an additional problem for networks developed for a fixed number of event frames. A Long Short-Term Memory (LSTM)-based event feature extraction module has been developed for addressing these challenges, which enables us to use a dynamically varying number of event frames. Using these modules, we constructed a state-of-the-art deblurring network, Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network (DLEFNet). It is particularly useful for scenarios in which exposure time varies depending on factors such as lighting conditions or the presence of fast-moving objects in the scene. It has been demonstrated through evaluation results that the proposed method can outperform the existing state-of-the-art networks for deblurring task in synthetic and real-world data sets. Visual comparison of the state-of-the-art algorithms on REBLur dataset The Real Event Blur (REBLur) dataset is a recently proposed collection of event data and real blurry images. As done in, we fine-tuned our network trained on the GoPro dataset with the training portion of the REBLur dataset. The evaluation was then conducted on the test set. Visual comparison of the state-of-the-art algorithms on REBLur dataset The Real Event Blur (REBLur) dataset is a recently proposed collection of event data and real blurry images. As done in, we fine-tuned our network trained on the GoPro dataset with the training portion of the REBLur dataset. The evaluation was then conducted on the test set. Please answer a question about this article. What are the values for the following properties to construct a Leaderboard for the mode

Figure 2: DocTAET representation of the paper title "Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network for Motion Deblurring" With Dashed lines representing Task, Dataset, Metrics, and Score present in the paper but not captured by paper with codes



{'id': '2306.00834v1.pdf', 'prompt': 'Title: Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network for Motion Deblurring Abstract: Event cameras differ from conventional RGB cameras in that they produce asynchronous data sequences. While RGB cameras capture every frame at a fixed rate, event cameras only capture changes in the scene, resulting in sparse and asynchronous data output. Despite the fact that event data carries useful information that can be utilized in motion deblurring of RGB cameras, integrating event and image information remains a challenge. Recent state-of-the-art CNN-based deblurring solutions produce multiple 2-D event frames based on the accumulation of event data over a time period. In most of these techniques, however, the number of event frames is fixed and predefined, which reduces temporal resolution drastically, particularly for scenarios when fast-moving objects are present or when longer exposure times are required. It is also important to note that recent modern cameras (e.g., cameras in mobile phones) dynamically set the exposure time of the image, which presents an additional problem for networks developed for a fixed number of event frames. A Long Short-Term Memory (LSTM)-based event feature extraction module has been developed for addressing these challenges, which enables us to use a dynamically varying number of event frames. Using these modules, we constructed a state-of-the-art deblurring network, Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network (DLEFNet). It is particularly useful for scenarios in which exposure times vary depending on factors such as lighting conditions or the presence of fast-moving objects in the scene. It has been demonstrated through evaluation results that the proposed method can outperform the existing state-of-the-art networks for deblurring task in synthetic and real-world data sets. Experimental Evaluation Method PSNR SSIM Params BHA SRN 29.06 0.943 30.26 0.934 DeepDeblur 29.23 0.916 11.7M SVDN 11.32M PSS-NSC 30.92 0.942 2.8M SRN 30.26 0.934 6.8M DGN 30.49 0.938 11.32M PSS-NSC 30.92 0.942 2.8M MT-RNN 31.15 0.945 2.6M DMPHN 31.20 0.945 11.32M PSS-NSC 30.92 0.942 2.8M N/A LEBMD 31.79 0.949 31.20 0.960 31.98 0.928 23.4M SAPHN 31.76 0.953 32.02 0.953 32.44 0.957 85.6M MPRNET 32.66 0.959 20.1M MIMO-UNet++ 32.68 0.959 5.42M BANET 32.44 0.957 85.6M MPRNET 32.66 0.959 20.1M MIMO-UNet++ 32.68 0.959 16.1M HINet 32.71 0.967 67.9M MADANet+ 33.84 0.964 16.9M Vitoria et al. 34.33 0.944 N/A NAFNet 33.71 0.967 67.9M MADANet+ 33.84 0.964 16.9M Vitoria et al. 34.33 0.944 N/A EFNNet 35.46 0.972 8.5M DLEFNet (ours) 35.61 0.973 12M Deblurring results on GoPro dataset GoPro Dataset We adhere to the recommended protocol for training and testing by averaging a varying number of nearby frames (ranging from 7 to 13) to generate the blurry image. In that manner, the synthetic events corresponding to the averaged frames are accumulated into varying number of event frames ranging from 7 to 13. ESIM simulator was used to generate the corresponding events. We use a total of 2103 samples for training, consisting of blurry image, event frame stack, and ground-truth frames, and 1111 samples for testing. Table 1 compares the performance of DLEFNet with other state-of-the-art deblurring methods. BHA, LEBMD, ERDN, Vitoria et al., and EFNNet are event-aided solutions, while DGN provides depth-aware deblurring, MBRNN, GSTA, and PVDNet are video-based (multi-frame) methods. The proposed solution significantly outperforms the SoTA deblurring techniques as can be seen in the Table 1. Visual comparison of the algorithms are given in Figure [fig:result]. REBLur Dataset Visual comparison of the state-of-the-art algorithms on REBLur dataset The Real Event Blur (REBLur) dataset is a recently proposed collection of event data and real blurry images. REBLur encompasses 12 different types of linear and nonlinear motions or 3 distinct moving patterns and for the camera itself. It includes 36 sequences and 1469 pairs of blurry-sharp images with associated events. Of these pairs, 486 are used for training, and 983 for testing. As done in, we fine-tuned our network trained on the GoPro dataset with the training portion of the REBLur dataset. The evaluation was then conducted on the test set. The performance comparison of algorithms is presented in Table 2. The results of the algorithms with "*" are from, in which the blurry images and event frames produced in are fed and trained together. A visual comparison of EFNNet and DLEFNet is provided in Figure 1. Our proposed solution outperforms the state-of-the-art network EFNNet both quantitatively and visually. PSNR Params SRN 35.58 88.67M BHA+ 36.62 0.51M SRN+ 36.87 10.43M HINet+ 37.68 88.85M EFNNet 38.12 8.47M DLEFNet (ours) 38.40 12M Deblurring results on REBLur dataset Conclusions Event cameras differ from RGB cameras in that they produce asynchronous data sequences by capturing only changes in the scene. This unique property makes it challenging to utilize existing computer vision algorithms with event data. Recent state-of-the-art solutions generate event frames by dividing the exposure time, assuming it remains constant between frames, into a fixed number of smaller time intervals. However, this approach is unrealistic as modern cameras automatically adjust exposure time based on scene information obtained from multiple sensors. Additionally, fixing the number of frames results in large time interval bins to compensate for sparse event data, causing reduced temporal resolution in scenarios with fast-moving objects or longer exposure times. To address these issues, the Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network (DLEFNet) has been proposed, demonstrating superior performance compared to existing deblurring solutions on both synthetic and real-world datasets such as GoPro and REBLur. Please answer a question about this article. If the question is unanswerable, say "unanswerable". What are the values

Figure 3: DocREC representation of the paper title "Deformable Convolutions and LSTM-based Flexible Event Frame Fusion Network for Motion Deblurring". With Dashed lines representing Task, Dataset, Metrics, and Score present in the paper but not captured by paper with codes

3.4. DocFULL

As a last representation and to test the hypothesis that longer contexts which are not selectively tailored to suit the desired task objective tends to distract and thereby hinder the LLM performance, we used the full paper text as context. This approach entailed compiling the LaTeX source code of the document and translating its entirety into a plain text file. DocFULL ended producing the longest contexts compared to DocTAET and DocREC, in an average length of 5,948 words.

3.5. Models

In the period following the introduction of Flan-T5 [25], the community has witnessed the emergence of multiple advanced LLMs that have outperformed it. In this paper, we adopt the Flan-T5 fine-tuning paradigm and implement it in a comparative experimental setup on two of the latest state-of-the-art LLMs featured on the public LLM leaderboard.

In this work, we selected two open-sourced models as well as two closed-source models.

Mistral-7B. As the first open-sourced model, we selected Mistral-7B [6]. This model as the name suggested is a 7-billion-parameter language model optimized for performance and efficiency. It introduces the Grouped-Query Attention (GQA) for rapid inference and reduced memory requirements, and Sliding Window Attention (SWA) for handling long sequences with lower computational costs. The model surpasses existing models in benchmarks, including reasoning, mathematics, and code generation tasks. It also features fine-tuning capabilities for instruction following, achieving superior performance in human and automated benchmarks. Mistral 7B is designed for real-time applications, supports high throughput, and its architecture enables effective sequence generation with optimized memory usage. The model is released under the Apache 2.0 license, with its source code on Github, facilitating broad accessibility and application in various tasks.

Llama-2 7B. As the second open-sourced model, we selected the Llama-2 model [5]. The Llama-2 model is a collection of LLMs that range from 7 billion to 70 billion parameters, designed for both general and dialogue-specific applications. From the three available Llama-2 model checkpoints, i.e. 7B, 13B, and 70B, for comparability with our first select model, i.e. Mistral-7B, we choose the Llama-2 7B model. The Llama-2 models are fine-tuned for enhanced dialogue use cases and exhibit improved performance over existing open-source models in terms of helpfulness and safety, based on benchmarks and human evaluations. The Llama-2 family includes models optimized for different scales and introduces safety and fine-tuning methodologies to advance the responsible development of LLMs.

GPT-4-Turbo. As the first proprietary model, we leveraged GPT-4-Turbo, developed by OpenAI. GPT-4-Turbo is designed to offer enhanced performance and efficiency, providing faster response times and reduced computational costs. This model excels in various natural language processing tasks, including text generation, translation, summarization, and question answering. Optimized for high-throughput applications, GPT-4-Turbo maintains a high level of accuracy and consistency, making it suitable for real-time AI applications. Additionally, the model incorporates improvements in handling extensive contextual inputs, which enhances its capability in generating coherent and contextually relevant outputs.

GPT-4.o. For the proprietary model, we leveraged GPT-4.o, developed by OpenAI. GPT-4.o (“o” for “omni”) is designed to enable more natural human-computer interactions, accepting a wide range of input types including text, audio, image, and video, and generating outputs in text, audio, and image formats. The model is optimized for multimodal tasks, offering a significant advancement in vision and audio understanding compared to previous models.

We employed the QLORA (Quantum-enhanced Learning Optimization for Robust AI) [33] framework for fine-tuning our open-sourced models, leveraging its advanced optimization capabilities to enhance model performance. QLORA has been recognized for its innovative approach to integrating quantum computing principles with machine learning, offering a novel pathway to overcoming traditional optimization challenges.

4. Evaluations

Experimental setup. For training, we had one main experimental setting based on the 15 instructions. As elicited earlier in subsection 3.1, each of the 15 instruction were instantiated with the 7,987 (T, D, M, S) data instances and the SOTA question resulting in a total of 119,805 instances to instruction finetune out LLMs. In this scenario, we hypothesized that this repetition in the data instances across the instructions would cause the resulting model to overfit the training dataset. Thus to control for this, we applied the following experimental setup. Each instruction was instantiated with a random

selection of only half the 7,987 (T, D, M, S) data instances resulting in a finetuning dataset of a sizeable 41,340 instances that had leaderboards. The papers w/o leaderboards were also similarly halved. As shown in Table 1, the few-shot test included 1,401 paper instances, comprising 648 papers without leaderboards and 735 papers with leaderboards. The zero-shot test set consisted of 789 papers, with 548 papers without leaderboards and 241 papers with leaderboards. Model hyperparameter details are in Appendix C. In terms of compute, all experiments including inference were run on an NVIDIA 3090 GPU. Training took 20 hours on the 50% sampled dataset, while inference lasted 10 minutes for 2k test instances.

Table 2

Evaluation results of Llama-2, Mistral, GPT-4-Turbo, and GPT-4.0 for the shared task, reported using the metrics proposed for the task. The output evaluations are conducted as a structured summary generation task (reported with ROUGE metrics) and as a binary classification task to distinguish between papers with and without leaderboards (reported as General Accuracy).

Model	Few-shot					Zero-shot					
	Rouge1	Rouge2	RougeL	RougeLsum	General -Accuracy	Rouge1	Rouge2	RougeL	RougeLsum	General -Accuracy	Context
Llama-2 7B	49.68	10.18	48.91	49.02	83.51	68.15	4.81	67.59	67.78	86.82	DocREC
	49.70	17.62	48.81	48.81	83.62	62.75	10.88	62.07	62.18	86.22	DocTAET
	5.38	0.79	4.96	5.13	57.54	7.55	0.71	7.24	7.35	37.80	DocFULL
Mistral 7B	55.46	14.11	54.54	54.64	88.44	72.98	6.87	72.42	72.35	92.40	DocREC
	57.24	19.67	56.28	56.19	89.68	73.54	12.23	73.01	72.95	95.97	DocTAET
	6.73	0.77	6.36	6.49	58.45	9.38	0.59	9.11	9.23	39.28	DocFULL
GPT-4-Turbo	52.64	5.82	51.99	51.76	60.89	72.80	2.66	72.35	72.09	77.06	DocREC
	43.14	2.41	42.97	42.91	47.33	59.98	0.48	59.89	59.74	61.18	DocTAET
	48.50	3.21	48.06	47.96	52.87	70.10	1.8	69.75	69.73	72.65	DocFULL
GPT-4.0	58.59	16.81	56.37	55.45	83.21	74.94	9.02	73.65	73.02	87.94	DocREC
	52.10	13.72	50.77	49.26	80.63	69.59	8.81	68.57	67.43	87.56	DocTAET
	55.41	17.82	53.01	51.79	79.56	70.05	10.89	68.42	67.51	78.95	DocFULL

Metrics. We evaluated our models in two main settings. In the first setting, we applied standard summarization ROUGE metrics [34]. Furthermore, we also tested the models ability to identify papers with leaderboards and those without. This task was simple. For the papers with leaderboards, the model replied with a structured summary and for those it identified as without it replied as “unanswerable.” For these evaluations we applied simple accuracy measure. In the second setting, we evaluated the model JSON output in a fine-grained manner w.r.t. each of the individual (T, D, M, S) elements and overall for which we reported the results in terms of the standard F1 score and Precision score.

4.1. Results and Discussion

This section analyzes the comparative performance of different contexts provided to LLMs, examining their impact on model precision and reliability across various tasks and settings. Additionally, we utilized a matching algorithm, fuzz.ratio, with a threshold of 50% for partial metrics to account for variations in how tasks, datasets, and metrics are reported in research papers compared to the pwc code dumps.

Results for RQ1: Performance in Generating Structured Summaries and Classification

1. Few-shot Performance (Table 2):

- **Llama-2 7B:** Demonstrates consistent performance across different contexts (DocREC, DocTAET) with high General Accuracy in DocREC (83.51%) and DocTAET (83.62%). It shows substantial differences in ROUGE scores, especially in the ROUGE-2 and ROUGE-Lsum metrics, indicating its ability to capture more nuanced information.
- **Mistral 7B:** Exhibits the highest General Accuracy in DocTAET (89.68%) and competitive ROUGE scores across all contexts. This model shows significant improvement over Llama-2 7B and GPT-4-Turbo in most metrics, showcasing the strengths of domain-specific fine-tuned open-source models.

Table 3

Evaluation results of Llama-2, Mistral, GPT-4-Turbo, and GPT-4.0 for the shared task, reported using the metrics proposed for the task. The evaluation considers the individual (Task, Dataset, Metric, Score) elements and Overall in the model JSON generated output, reported in terms of **F1 score**.

Model	Mode	Few-shot					Zero-shot					Context
		Task	Dataset	Metric	Score	Overall	Task	Dataset	Metric	Score	Overall	
Llama-2 7B	Exact	20.93	13.06	13.96	3.04	12.75	13.97	6.83	11.72	2.61	8.78	DocREC
	Partial	31.37	22.50	21.99	3.46	19.83	24.05	16.6	18.28	3.10	15.51	
	Exact	29.53	16.68	20.02	1.14	16.84	21.75	11.26	16.99	0.77	12.69	DocTAET
	Partial	43.37	30.36	30.51	1.38	26.40	38.48	23.10	27.09	0.96	22.41	
	Exact	1.59	1.36	0.94	0.23	1.03	2.06	1.30	1.52	0.33	1.30	DocFULL
	Partial	2.29	1.82	1.68	0.37	1.54	3.36	2.49	2.49	0.54	2.22	
Mistral 7B	Exact	26.77	15.68	18.70	6.36	16.88	17.99	11.80	15.55	5.04	12.60	DocREC
	Partial	39.75	27.28	28.49	7.08	25.65	29.88	21.05	23.16	5.75	19.96	
	Exact	33.38	18.51	24.23	1.87	19.50	26.99	14.32	22.04	1.20	16.14	DocTAET
	Partial	46.35	32.75	34.16	2.25	28.88	44.90	27.29	32.23	1.41	26.46	
	Exact	0.81	0.57	0.57	0.56	0.63	0.22	0.33	0.33	0.76	0.42	DocFULL
	Partial	1.19	0.85	0.81	0.84	0.92	0.56	0.67	0.78	0.87	0.72	
GPT-4-Turbo	Exact	7.61	6.19	4.92	4.25	5.74	4.26	5.35	3.86	3.28	4.18	DocREC
	Partial	16.48	13.96	11.03	7.03	12.13	13.76	11.09	10.19	5.46	10.13	
	Exact	2.99	2.69	0.95	0.75	1.84	1.13	0.79	0.34	0.11	0.59	DocTAET
	Partial	6.22	5.42	3.03	1.63	4.08	2.72	1.59	1.59	0.11	1.5	
	Exact	3.38	3.16	1.98	2.48	2.75	2.45	2.98	1.81	2.77	2.5	DocFULL
	Partial	7.03	6.41	4.96	4.15	5.64	6.49	5.85	4.47	3.56	5.09	
GPT-4.0	Exact	16.14	16.11	15.50	10.76	14.63	16.04	15.05	17.43	10.38	14.72	DocREC
	Partial	38.40	32.63	29.35	15.20	28.90	37.23	31.16	29.97	14.96	28.33	
	Exact	14.10	12.76	9.91	2.11	9.72	13.78	10.25	11.01	2.36	9.35	DocTAET
	Partial	31.84	26.65	20.83	4.22	20.92	29.33	23.87	19.50	3.71	19.12	
	Exact	16.72	14.53	14.67	11.25	14.29	13.08	14.94	16.09	11.17	13.82	DocFULL
	Partial	36.56	31.0	27.61	16.50	27.93	35.59	28.28	27.38	14.80	26.52	

Table 4

Evaluation results of Llama-2, Mistral, GPT-4-Turbo, and GPT-4.0 w.r.t. the individual (Task, Dataset, Metric, Score) elements and Overall in the model JSON generated output in terms of **Precision score**.

Model	Mode	Few-shot					Zero-shot					Context
		Task	Dataset	Metric	Score	Overall	Task	Dataset	Metric	Score	Overall	
Llama-2 7B	Exact	34.10	21.27	22.74	4.99	20.78	31.89	15.77	26.77	6.06	20.12	DocREC
	Partial	51.13	36.66	35.82	5.59	32.32	54.92	38.32	41.73	7.27	35.56	
	Exact	30.61	17.29	20.78	1.18	17.46	24.56	12.72	19.19	0.87	14.34	DocTAET
	Partial	44.96	31.48	31.66	1.43	27.38	43.46	26.09	30.60	1.09	25.31	
	Exact	34.69	29.59	20.41	5.10	22.45	32.20	20.34	23.73	5.08	20.34	DocFULL
	Partial	50.00	39.80	36.73	8.16	33.67	52.54	38.98	38.98	8.47	34.75	
Mistral 7B	Exact	37.65	22.15	26.38	8.94	23.78	35.68	23.40	31.02	9.98	25.02	DocREC
	Partial	55.90	38.52	40.18	9.95	36.14	59.25	41.73	46.20	11.46	39.66	
	Exact	39.48	21.89	28.66	2.21	23.06	38.46	20.41	31.41	1.71	23.00	DocTAET
	Partial	54.82	38.73	40.41	2.65	34.15	64.00	38.89	45.94	2.03	37.71	
	Exact	32.43	32.43	32.43	9.6	30.76	25.00	37.50	37.50	14.00	28.50	DocFULL
	Partial	71.43	48.65	45.95	14.52	45.13	62.50	75.00	87.50	21.62	61.66	
GPT-4-Turbo	Exact	30.46	24.80	19.67	17.30	23.06	20.57	25.96	18.66	15.94	20.28	DocREC
	Partial	65.96	55.95	44.07	23.89	47.47	66.51	53.85	49.28	21.67	47.83	
	Exact	32.79	29.51	10.27	8.20	20.19	29.41	20.59	8.82	2.94	15.44	DocTAET
	Partial	68.31	59.56	32.97	15.64	44.12	70.59	41.18	41.18	4.55	39.37	
	Exact	28.52	27.56	17.05	21.65	23.70	24.73	30.11	18.28	27.96	25.73	DocFULL
	Partial	59.32	55.91	42.64	30.69	47.14	65.59	59.14	45.16	30.91	50.20	
GPT-4.0	Exact	17.99	17.96	17.28	12.00	16.31	19.5	18.72	21.67	12.89	18.31	DocREC
	Partial	42.82	36.38	32.72	16.09	32.00	46.31	38.5	37.27	18.18	35.13	
	Exact	19.70	17.83	13.85	2.95	13.58	25.31	18.83	20.22	4.37	17.18	DocTAET
	Partial	44.50	37.25	29.11	7.78	29.66	53.86	43.83	35.80	9.07	35.64	
	Exact	19.18	16.67	16.84	12.90	16.40	15.30	17.48	18.83	12.05	16.17	DocFULL
	Partial	41.95	35.56	31.70	18.15	31.84	41.64	33.08	32.03	16.66	30.85	

- **GPT-4-Turbo:** Shows moderate performance with lower General Accuracy in DocTAET (47.33%). However, it demonstrates notable improvements in ROUGE-L and ROUGE-Lsum, indicating its potential in capturing detailed summaries.
- **GPT-4.0:** Shows strong performance in ROUGE metrics, especially in DocREC and DocTAET

Table 5

Evaluation results of Llama-2, Mistral, GPT-4-Turbo, and GPT-4.o w.r.t. the individual (Task, Dataset, Metric, Score) elements and Overall in the model JSON generated output in terms of **Recall score**.

Model	Mode	Few-shot					Zero-shot					Context
		Task	Dataset	Metric	Score	Overall	Task	Dataset	Metric	Score	Overall	
Llama-2 7B	Exact	15.09	9.42	10.07	2.19	9.19	8.94	4.36	7.51	1.66	5.62	DocREC
	Partial	2.63	16.23	15.86	2.49	14.30	15.40	10.60	11.70	1.97	9.92	
	Exact	28.52	16.11	19.32	1.10	16.26	19.52	10.10	15.24	0.69	11.39	DocTAET
	Partial	41.89	29.33	29.43	1.33	25.50	34.53	20.73	24.31	0.85	20.10	
	Exact	0.81	0.69	0.48	0.12	0.53	1.06	0.67	0.78	0.17	0.67	DocFULL
Partial	1.17	0.93	0.86	0.19	0.79	1.74	1.29	1.29	0.28	1.15		
Mistral 7B	Exact	20.77	12.14	14.49	4.93	13.08	12.03	7.89	10.38	3.37	8.42	DocREC
	Partial	30.84	21.12	22.07	5.49	19.88	19.98	14.07	15.45	3.84	13.33	
	Exact	28.91	16.03	20.99	1.62	16.89	20.79	11.03	16.97	0.93	12.43	DocTAET
	Partial	40.15	28.36	29.59	1.95	25.01	34.58	21.02	24.83	1.08	20.38	
	Exact	0.41	0.29	0.29	0.29	0.32	0.11	0.17	0.17	0.39	0.21	DocFULL
Partial	0.6	0.43	0.41	0.43	0.47	0.28	0.34	0.39	0.44	0.36		
GPT-4-Turbo	Exact	4.35	3.53	2.81	2.42	3.28	2.37	2.98	2.15	1.83	2.33	DocREC
	Partial	9.42	7.98	6.30	4.12	6.95	7.67	6.18	5.68	3.12	5.66	
	Exact	1.56	1.41	0.50	0.39	0.96	0.58	0.40	0.17	0.06	0.30	DocTAET
	Partial	3.26	2.84	1.59	0.86	2.14	1.39	0.81	0.81	0.06	0.76	
	Exact	1.8	1.68	1.05	1.32	1.46	1.29	1.57	0.95	1.46	1.32	DocFULL
Partial	3.73	3.40	2.63	2.23	3.00	3.42	3.08	2.35	1.89	2.68		
GPT-4.o	Exact	14.63	14.60	14.05	9.75	13.26	13.41	12.58	14.57	8.68	12.31	DocREC
	Partial	34.81	29.58	26.60	14.41	26.35	31.13	26.05	25.06	12.71	23.74	
	Exact	10.97	9.93	7.72	1.64	7.57	9.47	7.04	7.56	1.62	6.42	DocTAET
	Partial	24.79	20.75	16.21	2.89	16.16	20.15	16.40	13.39	2.33	13.07	
	Exact	14.82	12.88	13.00	9.97	12.66	11.42	13.05	14.05	9.76	12.07	DocFULL
Partial	32.41	27.48	24.46	15.13	24.87	31.08	24.69	23.91	13.32	23.25		

contexts, with high General Accuracy in DocREC (83.21%) and DocTAET (80.63%). However, the DocFULL context results of 79.56% which indicate the robustness of the DocREC context.

2. Zero-shot Performance (Table 2):

- **Llama-2 7B:** Performs well with notable General Accuracy in DocREC (86.82%) and DocTAET (86.22%). However, there is a significant drop in performance in the DocFULL context (37.8% General Accuracy), likely due to the long and complex nature of the full document context which introduces distractions and challenges the model's ability to maintain focus and coherence over extended text. This suggests that while the models perform well in more concise contexts, the added complexity and length in DocFULL can detract from their overall accuracy and effectiveness.
- **Mistral 7B:** Outperforms other models in General Accuracy, especially in DocTAET (95.97%) and DocREC (92.40%). This indicates its robustness in zero-shot settings, making it a reliable choice for generating structured summaries without prior examples.
- **GPT-4-Turbo:** Shows lower General Accuracy in zero-shot settings, particularly in DocREC (77.06%) and DocTAET (61.18%), highlighting the need for further optimization.
- **GPT-4.o:** Continues to show competitive performance in DocREC (87.94%) and DocTAET (87.56%) contexts, reinforcing its capability in zero-shot scenarios.

Results for RQ2: Precision and Performance Trade-offs

1. Few-shot F1 Score analysis (Table 3):

- **Llama-2 7B:** Shows a balanced performance between exact and partial metrics. For example, it achieves an Overall score of 26.40 in the partial mode under DocTAET context, indicating its ability to handle partial matches effectively.
- **Mistral 7B:** Excels in both exact and partial metrics, with the highest Overall score of 28.88 in the partial mode under DocTAET context. This model provides the best trade-off between precision and other performance metrics.

- **GPT-4-Turbo:** Shows moderate performance with an Overall score of 12.13 in the partial mode under DocREC context. It requires further optimization to compete with other models.
- **GPT-4.o:** Demonstrates strong partial mode performance, with an Overall score of 28.90 under DocREC context.

2. Zero-shot F1 Score analysis (Table 3):

- **Llama-2 7B:** Shows variability in performance with a notable Overall score of 22.41 in the partial mode under DocTAET context.
- **Mistral 7B:** Continues to lead with the highest Overall score of 16.14 in the exact mode under DocTAET context, reinforcing its capability in zero-shot scenarios.
- **GPT-4-Turbo:** Provides moderate performance in partial mode with an Overall score of 10.13 under DocFULL context, highlighting areas for improvement.
- **GPT-4.o:** Provides solid performance in partial mode with an Overall score of 23.87 under DocTAET context.

3. Precision Scores (Table 4):

- **Llama-2 7B:** Achieves moderate precision scores, particularly in the partial mode with an Overall score of 35.56 in zero-shot settings under DocREC context.
- **Mistral 7B:** Dominates precision metrics with the highest Overall scores in both exact and partial modes across different contexts, demonstrating its precision and reliability.
- **GPT-4-Turbo:** Shows competitive precision scores in zero-shot settings, particularly in the partial mode with an Overall score of 50.20 under DocFULL context.
- **GPT-4.o:** Achieves moderate precision scores in few-shot settings, particularly in the exact mode, with an Overall score of 16.40 under DocFULL context.

4. Recall Scores (Table 5):

- **Llama-2 7B:** Demonstrates notable recall scores in few-shot settings, particularly in the partial mode with an Overall score of 25.50 under DocTAET context. However, its performance varies significantly across different contexts and modes.
- **Mistral 7B:** Leads in recall scores, achieving the highest Overall scores in both exact and partial modes across various contexts. Notably, it achieves an Overall score of 16.89 in the exact mode in few-shot setting and 12.43 in few-shot under DocTAET context.
- **GPT-4-Turbo:** Shows moderate recall performance, with room for improvement, particularly in the zero-shot settings. Its best recall score is in the exact mode under DocREC context with an Overall score of 3.28.
- **GPT-4.o:** Excels in recall scores, especially in few-shot settings, achieving an Overall score of 26.35 in the partial mode under DocREC context. It maintains competitive recall scores across different contexts and modes.

Score Extraction Analysis

When focusing on the extraction of score entities, GPT models, particularly GPT-4.o, demonstrate a clear advantage. As shown in Table 3, GPT-4.o outperforms other models in extracting scores with higher accuracy and precision. For example, in few-shot settings, GPT-4.o achieves an Overall score of 11.25 in the exact mode and 16.50 in the partial mode under the DocFULL context, significantly higher than Llama-2 7B and Mistral 7B. Similarly, in zero-shot settings, GPT-4.o maintains a competitive edge with an Overall score of 14.96 in the partial mode under DocREC context, surpassing the performance of the other models.

This superior performance in score extraction can be attributed to the advanced capabilities of GPT-4.o and its optimized architecture for handling diverse and long contexts. This makes it particularly effective in accurately recognizing and extracting numerical and textual data, which is

crucial for generating accurate leaderboard metrics and scores.

Discussion

Open-source models like Mistral 7B demonstrate competitive and sometimes superior performance compared to proprietary models like GPT-4.o and GPT-4-Turbo. This is evident in both structured summary generation and classification tasks, particularly in the DocTAET context, where Mistral 7B consistently outperforms others.

The trade-off between precision and performance is well-balanced in DocREC on Mistral 7B, making it a reliable choice for applications requiring high precision. This is crucial for scholarly communications where accuracy and reliability are paramount.

Few-shot settings show that all models, including open-source ones, perform robustly, but Mistral 7B often leads in various metrics. This highlights the potential of fine-tuning open-source models to achieve high performance that surpasses closed-source models.

5. Conclusions

Our participation in the shared task has demonstrated that fine-tuning open-source models like Mistral 7B and Llama-2 7B can yield competitive, and in some cases superior, results compared to proprietary models such as GPT-4.o and GPT-4-Turbo. Throughout our experiments, the DocTAET context typically delivered dependable and accurate performance, while the DocREC context excelled in scenarios where precision is paramount. The interaction among different context types and their consequent impact on model performance provided valuable insights into ongoing research and practical deployment of LLMs.

The implications of our findings within the context of the shared task are substantial. They indicate that with meticulous context design and implementation, fine-tuned open-source LLMs are well-suited for tracking and synthesizing scientific progress. This capability enables the provision of current and nuanced leaderboards for any given academic field. The potential of this technology to support and augment the efforts of academics and policymakers is significant, heralding novel opportunities for automated or semi-automated leaderboard construction.

While we have made significant progress in elucidating and enhancing context selection for LLMs, several avenues for future exploration have emerged. Investigating hybrid context selection methods, domain-specific adaptations, and the integration of structured data could lead to even more sophisticated leaderboard generation. As LLMs continue to advance, optimizing their potential remains a dynamic and impactful field of research.

In conclusion, our involvement in the shared task has not only highlighted the effectiveness of fine-tuned open-source models but also emphasized the importance of strategic context selection in maximizing model performance. These insights contribute to the broader understanding of LLM capabilities and pave the way for future advancements in the field.

Acknowledgements

This work was jointly supported by the German BMBF project SCINEXT (01IS22070) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number: NFDI4DataScience (460234259).

References

- [1] S. Fortunato, C. T. Bergstrom, K. Börner, J. A. Evans, D. Helbing, S. Milojević, A. M. Petersen, F. Radicchi, R. Sinatra, B. Uzzi, et al., Science of science, *Science* 359 (2018) eaao0185.

- [2] L. Bornmann, R. Haunschild, R. Mutz, Growth rates of modern science: a latent piecewise growth curve approach to model publication numbers from established and new literature databases, *Humanities and Social Sciences Communications* 8 (2021) 1–15.
- [3] F. Shi, X. Chen, K. Misra, N. Scales, D. Dohan, E. H. Chi, N. Schärli, D. Zhou, Large language models can be easily distracted by irrelevant context, in: *International Conference on Machine Learning*, PMLR, 2023, pp. 31210–31227.
- [4] N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, P. Liang, Lost in the middle: How language models use long contexts, *ArXiv abs/2307.03172* (2023). doi:10.48550/arXiv.2307.03172.
- [5] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al., Llama 2: Open foundation and fine-tuned chat models, *arXiv preprint arXiv:2307.09288* (2023).
- [6] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mistral 7b, 2023. arXiv:2310.06825.
- [7] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al., Gpt-4 technical report, *arXiv preprint arXiv:2303.08774* (2023).
- [8] J. D’Souza, S. Kabongo, H. B. Giglou, Y. Zhang, Overview of the CLEF 2024 simpletext task 4: SOTA? tracking the state-of-the-art in scholarly publications, in: G. Faggioli, N. Ferro, P. Galuščáková, A. G. S. de Herrera (Eds.), *Working Notes of CLEF 2024 - Conference and Labs of the Evaluation Forum*, CEUR Workshop Proceedings, CEUR-WS, Online, 2024.
- [9] L. Ermakova, E. SanJuan, S. Huet, H. Azarbyonad, G. M. D. Nunzio, F. Vezzani, J. D’Souza, J. Kamps, Overview of the CLEF 2024 simpletext track — improving access to scientific texts for everyone, in: L. Goeuriot, P. Mulhem, G. Quénot, D. Schwab, L. Soulier, G. M. D. Nunzio, P. Galuščáková, A. G. S. de Herrera, G. Faggioli, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the Fifteenth International Conference of the CLEF Association (CLEF 2024)*, Lecture Notes in Computer Science (LNCS), Springer, Heidelberg, Germany, 2024.
- [10] V. Qazvinian, D. R. Radev, Identifying non-explicit citing sentences for citation-based summarization., in: J. Hajič, S. Carberry, S. Clark, J. Nivre (Eds.), *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Uppsala, Sweden, 2010, pp. 555–564. URL: <https://aclanthology.org/P10-1057>.
- [11] A. Athar, S. Teufel, Detection of implicit citations for sentiment detection, in: A. Van Den Bosch, H. Shatkay (Eds.), *Proceedings of the Workshop on Detecting Structure in Scholarly Discourse*, Association for Computational Linguistics, Jeju Island, Korea, 2012, pp. 18–26. URL: <https://aclanthology.org/W12-4303>.
- [12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [13] Y. Hou, C. Jochim, M. Gleize, F. Bonin, D. Ganguly, Identification of tasks, datasets, evaluation metrics, and numeric scores for scientific leaderboards construction, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019. URL: <https://aclanthology.org/P19-1513>. doi:10.18653/v1/P19-1513.
- [14] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [15] Y. Hu, Q. Chen, J. Du, X. Peng, V. K. Keloth, X. Zuo, Y. Zhou, Z. Li, X. Jiang, Z. Lu, et al., Improving large language models for clinical named entity recognition via prompt engineering, *Journal of the American Medical Informatics Association* (2024) ocad259.
- [16] M. Shamsabadi, J. D’Souza, S. Auer, Large language models for scientific information extraction: An empirical study for virology, in: Y. Graham, M. Purver (Eds.), *Findings of the Association for Computational Linguistics: EACL 2024*, Association for Computational Linguistics, St. Julian’s, Malta, 2024, pp. 374–392. URL: <https://aclanthology.org/2024.findings-eacl.26>.
- [17] S. Kabongo, J. D’Souza, S. Auer, Effective context selection in llm-based leaderboard generation: An empirical study, 2024. arXiv:2407.02409.

- [18] D. Khashabi, S. Min, T. Khot, A. Sabharwal, O. Tafjord, P. Clark, H. Hajishirzi, Unifiedqa: Crossing format boundaries with a single qa system, in: Findings of the Association for Computational Linguistics: EMNLP 2020, 2020, pp. 1896–1907.
- [19] T. Xie, C. H. Wu, P. Shi, R. Zhong, T. Scholak, M. Yasunaga, C.-S. Wu, M. Zhong, P. Yin, S. I. Wang, et al., Unifiedskg: Unifying and multi-tasking structured knowledge grounding with text-to-text language models, in: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, 2022, pp. 602–631.
- [20] Y. Wang, S. Mishra, P. Alipoormolabashi, Y. Kordi, A. Mirzaei, A. Arunkumar, A. Ashok, A. S. Dhanasekaran, A. Naik, D. Stap, et al., Super-naturalinstructions: generalization via declarative instructions on 1600+ tasks, in: EMNLP, 2022.
- [21] V. Sanh, A. Webson, C. Raffel, S. H. Bach, L. Sutawika, Z. Alyafeai, A. Chaffin, A. Stiegler, T. Le Scao, A. Raja, et al., Multitask prompted training enables zero-shot task generalization, in: ICLR 2022-Tenth International Conference on Learning Representations, 2022.
- [22] O. Honovich, T. Scialom, O. Levy, T. Schick, Unnatural instructions: Tuning language models with (almost) no human labor, arXiv preprint arXiv:2212.09689 (2022).
- [23] S. Longpre, L. Hou, T. Vu, A. Webson, H. W. Chung, Y. Tay, D. Zhou, Q. V. Le, B. Zoph, J. Wei, et al., The flan collection: Designing data and methods for effective instruction tuning, arXiv preprint arXiv:2301.13688 (2023).
- [24] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al., Training language models to follow instructions with human feedback, Advances in Neural Information Processing Systems 35 (2022) 27730–27744.
- [25] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma, et al., Scaling instruction-finetuned language models, arXiv preprint arXiv:2210.11416 (2022).
- [26] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, The Journal of Machine Learning Research 21 (2020) 5485–5551.
- [27] X. Liu, P. He, W. Chen, J. Gao, Multi-task deep neural networks for natural language understanding, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 4487–4496.
- [28] A. Aghajanyan, A. Gupta, A. Shrivastava, X. Chen, L. Zettlemoyer, S. Gupta, Muppet: Massive multi-task representations with pre-finetuning, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 5799–5811. URL: <https://aclanthology.org/2021.emnlp-main.468>. doi:10.18653/v1/2021.emnlp-main.468.
- [29] V. Aribandi, Y. Tay, T. Schuster, J. Rao, H. S. Zheng, S. V. Mehta, H. Zhuang, V. Q. Tran, D. Bahri, J. Ni, et al., Ext5: Towards extreme multi-task scaling for transfer learning, arXiv preprint arXiv:2111.10952 (2021).
- [30] P. Rajpurkar, J. Zhang, K. Lopyrev, P. Liang, Squad: 100,000+ questions for machine comprehension of text, in: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 2016, pp. 2383–2392.
- [31] P. Rajpurkar, R. Jia, P. Liang, Know what you don’t know: Unanswerable questions for squad, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2018, pp. 784–789.
- [32] D. Dua, Y. Wang, P. Dasigi, G. Stanovsky, S. Singh, M. Gardner, Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), 2019, pp. 2368–2378.
- [33] T. Dettmers, A. Pagnoni, A. Holtzman, L. Zettlemoyer, Qlora: Efficient finetuning of quantized llms, Advances in Neural Information Processing Systems 36 (2024).
- [34] C.-Y. Lin, Rouge: A package for automatic evaluation of summaries, in: Text summarization branches out, 2004, pp. 74–81.
- [35] N. Shazeer, M. Stern, Adafactor: Adaptive learning rates with sublinear memory cost, in:

A. Instructions: Qualitative Examples

In this section, we elicit each of the instructions that were considered in this work as formulated in the FLAN 2022 Collection for the SQuAD_v2 and DROP datasets.

Table 6
Comparative Instructions for the SQuAD_v2 and DROP datasets.

ID	SQuAD_v2 Instructions	DROP Instructions
1	Please answer a question about this article. If unanswerable, say "unanswerable".	Answer based on context.
2	{Context} {Question} If unanswerable, say "unanswerable".	Answer this question based on the article.
3	Try to answer this question if possible (otherwise reply "unanswerable").	{Context} {Question}
4	Please answer a question about this article, or say "unanswerable" if not possible.	Answer this question: {Question}
5	If possible to answer this question, do so (else, reply "unanswerable").	Read this article and answer this question.
6	Answer this question, if possible (if impossible, reply "unanswerable").	Based on the above article, answer a question.
7	What is the answer? (If it cannot be answered, return "unanswerable").	Context: {Context} Question: {Question} Answer:
8	Now answer this question, if there is an answer (else, "unanswerable").	

B. ROUGE Evaluation Metrics

The ROUGE metrics [34] are commonly used for evaluating the quality of text summarization systems. ROUGE-1 measures the overlap of unigram (single word) units between the generated summary and the reference summary. ROUGE-2 extends this to measure the overlap of bigram (two consecutive word) units. ROUGE-L calculates the longest common subsequence between the generated and reference summaries, which takes into account the order of words. ROUGE-LSum is an extension of ROUGE-L that considers multiple reference summaries by treating them as a single summary.

C. Additional Data statistics and Hyperparameters

Table 7
Full Paper dataset statistics. The "papers w/o leaderboard" refers to papers that do not report leaderboard.

	Our Corpus		
	Train	Test-Few-shot	Test Zero-shot
Papers w/ leaderboards	7,744	961	630
Papers w/o leaderboards	4,063	604	507
Total TDM-triples	612,709	74,475	76,936
Distinct TDM-triples	62,629	8,748	8,434
Distinct <i>Tasks</i>	1,365	432	339
Distinct <i>Datasets</i>	4,733	1,379	1,077
Distinct <i>Metrics</i>	2,845	850	825
Avg. no. of TDM per paper	5.1	5.1	6.7
Avg. no. of TDMS per paper	6.9	6.3	9.0

We used a context length of 2400 and based on GPU availability, a batch size of 2 and gradient_accumulation_steps of 4 were used, leading to a final batch size of 8. All experiments were run on five epochs and we used AdafactorSchedule and Adafactor optimizer [35] with scale_parameter=True, relative_step=True, warmup_init=True, lr=1e-4.

Table 8

DOCTEAT dataset statistics. The “papers w/o leaderboard” refers to papers that do not report leaderboard.

	Our Corpus		
	Train	Test-Few-shot	Test Zero-shot
Papers w/ leaderboards	7,025	903	573
Papers w/o leaderboards	3,033	444	353
Total TDM-triples	515,203	63,041	61,474
Distinct TDM-triples	56,486	8,241	7,557
Distinct <i>Tasks</i>	1,232	417	323
Distinct <i>Datasets</i>	4,473	1,317	1,010
Distinct <i>Metrics</i>	2,687	812	778
Avg. no. of TDM per paper	5.0	5.0	6.7
Avg. no. of TDMS per paper	6.8	6.2	8.9