

Verif.ai: Towards an Open-Source Scientific Generative Question-Answering System with Referenced and Verifiable Answers

Miloš Košprdić

*Institute for Artificial Intelligence Research
and Development of Serbia*
Fruškogorska 1, Novi Sad, Serbia
email: milos.kosprdic@ivi.ac.rs

Adela Ljajić

*Institute for Artificial Intelligence Research
and Development of Serbia*
Fruškogorska 1, Novi Sad, Serbia
email: adela.ljajic@ivi.ac.rs

Bojana Bašaragin

*Institute for Artificial Intelligence Research
and Development of Serbia*
Fruškogorska 1, Novi Sad, Serbia
email: bojana.basaragin@ivi.ac.rs

Darija Medvecki

*Institute for Artificial Intelligence Research
and Development of Serbia*
Fruškogorska 1, Novi Sad, Serbia
email: darija.medvecki@ivi.ac.rs

Nikola Milošević

*R&D Data Sciences and AI
Bayer A.G.*
Müllerstraße 178, Berlin, Germany
email: nikola.milosevic@bayer.com

Abstract—In this paper, we present the current progress of the project Verif.ai, an open-source scientific generative question-answering system with referenced and verified answers. The components of the system are (1) an information retrieval system combining semantic and lexical search techniques over scientific papers (PubMed), (2) a fine-tuned generative model (Mistral 7B) taking top answers and generating answers with references to the papers from which the claim was derived, and (3) a verification engine that cross-checks the generated claim and the abstract or paper from which the claim was derived, verifying whether there may have been any hallucinations in generating the claim. We are reinforcing the generative model by providing the abstract in context, but in addition, an independent set of methods and models are verifying the answer and checking for hallucinations. Therefore, we believe that by using our method, we can make scientists more productive, while building trust in the use of generative language models in scientific environments, where hallucinations and misinformation cannot be tolerated.

Keywords—question-answering; automatic referencing; generative search; large language models; natural language inference.

I. INTRODUCTION

In recent years, the advent of large language models has revolutionized various domains, offering unprecedented capabilities in natural language understanding, generation, and interaction [1]–[6]. Particularly within the scientific community, these models hold tremendous potential for accelerating research processes, automating information retrieval [7], and enhancing the generation of complex scientific content. However, as these models become integral to scientific workflows, a critical challenge emerges – the issue of hallucinations, or the inadvertent generation of false or misleading information [8]–[10].

In scientific domains where accuracy and reliability are paramount, the occurrence of hallucinations poses a significant impediment to the widespread adoption of Large Language Models (LLMs) [11]. The potential for misinformation introduces an inherent trust deficit, hindering scientists from fully embracing generative language models. It is imperative to address this challenge comprehensively to ensure that the benefits of these models are harnessed without compromising the integrity of scientific knowledge.

In response to this pressing concern, we introduce the Verif.ai project, an open-source initiative aimed at mitigating the risk of hallucinations in scientific generative question-answering systems. Our approach relies on information retrieval, leveraging both semantic and lexical techniques over a vast repository of scientific papers such as PubMed [12], complemented with Retrieval-Augmented Generation (RAG) using a fine-tuned generative model, Mistral 7B, for answer generation with traceable references. Notably, the system goes beyond mere answer generation by incorporating a verification engine that cross-checks the generated claims against the abstracts or papers from which they are derived. We believe that the system, which makes the best effort to indicate possible hallucinations to the user, coupled with hallucination reduction techniques, its open-source nature, and community support, will instill trust in the scientific community in the use of LLM-based scientific systems.

The rest of the paper is structured as follows. In Section 2, we present the methodology overview. In Section 3, we present the preliminary results of our evaluations. We conclude in Section 4 and provide the information about the availability of code and the models in Section 5.

II. METHODOLOGY

Our methodology employs a toolbox to discover relevant information and provide context to the question-answering system. Currently, the primary component of this toolbox is the information retrieval engine (PubMed). The question-answering system utilizes a fine-tuned LLM to generate answers based on the information from the toolbox. A fact-checking or verification engine examines the generated answer within the toolbox, identifying any potential hallucinations in the system. The final component of the system is a user interface, enabling users to ask a question, review answers and offer a feedback functionality, so they can contribute to the improvement of the **Verif.ai** project. The overview of the methodology is depicted in Figure 1. In the following subsections, we provide details of the methods envisioned for each of the components.

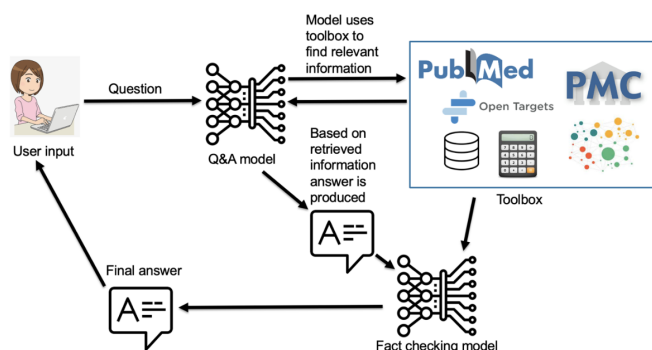


Fig. 1. Methodology overview of the **Verif.ai** project

A. Toolbox and Information Retrieval

The major component that has been implemented so far in our toolbox is the information retrieval engine. Our information retrieval engine is based on OpenSearch [13], an open-source engine that was forked from Elasticsearch and is under the Apache 2 license. We have indexed PubMed articles using lexical indexing provided by OpenSearch. Additionally, we have created an index storing embeddings of documents using the MSMARCO model for semantic search. This model was selected because it can handle asymmetric searches (e.g., different lengths of queries compared to the searched texts) [14]. Embeddings were stored in the OpenSearch field, allowing for the combination of lexical and semantic search. This approach emphasizes direct matches while also finding semantically similar phrases and parts of the text where the text does not match. The user question is first transformed into a query, and the most relevant documents are retrieved before being passed to the LLM that generates the answer.

B. Question-answering with references

For generating the answers, we have used Mistral 7B parameter model with instruction fine-tuning [15]. This model was further fine-tuned using questions from PubMedQA dataset [16] and generated answers using GPT3.5 with the most

relevant documents from PubMed passed as context. The following prompt was used to generate answers:

Please carefully read the question and use the provided research papers to support your answers. When making a statement, indicate the corresponding abstract number in square brackets (e.g., [1][2]). Note that some abstracts may appear to be strictly related to the instructions, while others may not be relevant at all.

We have selected 10,000 random PubMedQA questions to generate this dataset. The dataset was then used to fine-tune the Mistral 7B model using the QLoRA methodology [17]. The training was performed using a rescaled loss, a rank of 64, an alpha of 16, and a LoRA dropout of 0.1, resulting in 27,262,976 trainable parameters. The input to the training contained the question, retrieved documents (as many as can fit into the context), and the answer. We made this preliminary generated QLoRA adapter available on Hugging Face [18].

We then used the fine-tuned model for answer generation. Using the exactly same input as in training did not produce the expected results, and therefore, we added an instruction at the beginning of the prompt:

[INST] Answer the question using the given abstracts. Reference claims with the relevant abstract id in brackets (e.g. (PUBMED:123456) at the end of the sentence). Answer may contain references to many abstracts. Be as factual as possible and always use references in brackets. Use exclusively provided abstracts and their ids. Make answer look similar to the following: Several genes play role in breast cancer. For example BRAC1, BRAC2 are well studied targets (PUBMED:554433). The other targets involve IRAK4, CAS2 and HMPA (PUBMED:665544).

The instruction was followed by the set of relevant documents obtained by querying OpenSearch and the question asked by the user. To prompt Mistral-7B-Instruct-v0.1-pqa, we use the mentioned template and default parameters with only two differences: we set `max_new_tokens` to 1000 and `repetition_penalty` to 1.1.

C. Verifying claims from the generated answer

The aim of the verification engine is to parse sentences and references from the answer generation engine and verify that there are no hallucinations in the answer. Our assumption is that each statement is supported by one or more references. For verification, we compare the XLM-RoBERTa-large model [19] and DeBERTa model [20], treating it as a natural language inference problem. The selected model has a significantly different architecture than the generation model and is fine-tuned using the SciFact dataset [8]. The dataset is additionally cleaned (e.g., claims were deduplicated, and instances with multiple citations in no-evidence examples were split into multiple samples, one for each reference). The input to the model contains the CLS token (class token), the statement, a separator token, and the joined referenced article title and

abstract, followed by another separation token. The output of the model falls into one of three classes:

- **Supports** - in case statement is supported by the content of the article
- **Contradicts** - in case the statement contradicts the article
- **No Evidence** - in case there is no evidence in the article for the given claim

The fine-tuned model serves as the primary method for flagging contradictions or unsupported claims. However, additional methods for establishing user trust in the system will be implemented, including presenting to the user the sentences from the abstracts that are most similar to the claim.

D. User feedback integration

The envisioned user interface would present the answer to the user’s query, referencing documents containing the answer and flagging sentences that contain potential hallucinations. However, users are asked to critically evaluate answers, and they can provide feedback either by changing a class of the natural language inference model or even by modifying generated answers. These modifications are recorded and used in future model fine-tuning, thereby improving the system.

III. PRELIMINARY EVALUATION

In this section, we present the results based on our preliminary evaluation. At the time of writing of this article, the project was in the 3rd month of implementation, and we are working on improving our methodology and creating a web application that integrates all the described components.

A. Information retrieval

We have qualitatively evaluated OpenSearch’s results on a small set of indexed PubMed articles. We compared lexical search, semantic search, and a hybrid combination of both lexical and semantic search. We observed that lexical search may perform better when the search terms can be exactly matched in the documents, while semantic search works well with paraphrased text or synonymous terms. Hybrid search managed to find documents containing terms that could be exactly matched, as well as ones that were paraphrased or contained synonyms. While semantic search would also find documents that contained an exact match of the terms, it often happened that they were not prioritized. Hybrid search helped in putting such documents at the top of the search results. Based on several user discussions, we have concluded that users expect the top results to be based on exact matches and later to find relevant documents that do not contain the searched terms.

B. Answer generation

As we previously mentioned, we have fine-tuned Mistral 7B for question answering on questions coming from PubMedQA and answers generated using PubMed searches for relevant abstracts and GPT-3.5 for actual answer generation. The evaluation loss for the fine-tuning process can be seen in Figure 2.




Fig. 2. Evaluation loss for fine-tuning of Mistral 7B model on PubMedQA questions with generated and referenced answers

The fine-tuning of the Mistral 7B model improved the model’s performance, making the generated answers comparable to those of much larger GPT-3.5 and GPT-4 models for the referenced question-answering task.

After manually comparing answers from GPT-3.5, GPT-4, and Mistral-7B-Instruct-v0.1-pqa to a test set of 50 questions and extracted abstracts, no model showed a clear advantage over the others. The quality, referenced abstracts, and length of the answers varied within each model and among the models. In terms of referenced abstracts, most of the time all three models referenced the same abstracts as relevant.

C. Verification and hallucination detection

The evaluation of the fine-tuned XLM-RoBERTa and DeBERTa model on the SciFact dataset that can be used for hallucination detection can be seen in Table I. The model used 10% of the data for validation and 10% of the dataset for evaluation (test set). All three sets have homogenous distribution of the classes (36%:42%:22% for NO_EVIDENCE, SUPPORT and CONTRADICT classes respectively).

TABLE I
THE EVALUATION OF THE ENTAILMENT MODEL FINE-TUNED FROM XLM-ROBERTA-LARGE AND DeBERTA-LARGE MODEL USING SCIFACT DATASET

	XLM-RoBERTa		
	Precision	Recall	F1-score
NO_EVIDENCE	0.91	0.96	0.95
SUPPORT	0.91	0.75	0.82
CONTRADICT	0.59	0.81	0.68
Weighted Avg	0.87	0.85	0.85
	DeBERTa		
NO_EVIDENCE	0.88	0.86	0.87
SUPPORT	0.87	0.92	0.90
CONTRADICT	0.88	0.81	0.85
Weighted Avg	0.88	0.88	0.88

As can be seen from the table, the models exhibited state-of-the-art performance, surpassing the reported scores in [8] for the label prediction task, and DeBERTa-large model showed superior performance compared to the RoBERTa-large. We use fine-tuned DeBERTa-large model for verification and hallucination detection. We also evaluated the SciFact label prediction task using the GPT-4 model, resulting in a precision

of 0.81, recall of 0.80, and an F-1 score of 0.79. Therefore, our models outperformed GPT-4 model in zero-shot regime with carefully designed prompt for label prediction for the claims and abstracts in the SciFact dataset. It is important to note that the SciFact dataset contains challenging claim/abstract pairs, demanding a significant amount of reasoning for accurate labeling. Thus, in a real-use case where answers are generated by Mistral or another generative model, the task becomes easier. We believe that this model provides a good starting point for hallucination detection, as supported by our qualitative analysis of several pairs of generated claims and abstracts, which demonstrated good performance.

However, this model has some limitations. While it is capable of reasoning around negations, detecting contradicting claims, differing in just few words switching the context of the claim compared to the text of the abstract, proves to be a challenge. Additionally, we observe that neither model handle well situations where numerical values in claims are slightly different from the ones in the abstract.

IV. CONCLUSION

In this short paper, we present the current progress on the **Verif.ai** project, an open-source generative search with referenced and verifiable answers based on PubMed. We describe our use of OpenSearch to create a hybrid search based on both semantic and lexical search methods, an answer generation method based on fine-tuning the Mistral 7B model, and our first hallucination detection and answer verification model based on fine-tuned DeBERTa-large model. However, there are still a number of challenges to be addressed and work to be done.

LLMs are rapidly developing, and performant, smaller LLMs, with larger context size are becoming more available. We aim to follow this development and use the best available open-source model for the task of referenced question-answering. We also aim to release early and collect user feedback. Based on this feedback, we aim to design an active learning method and incorporate user feedback into the iterative training process for both answer generation and answer verification and hallucination detection.

The model for hallucination detection and answer verification exhibits some limitations when it needs to deal with numerical values or perform complex reasoning and inference on abstracts. We believe that a single model may not be sufficient to verify the abstract well, but it may be the case that a solution based on a mixture of experts may be required [21][22]. To build user trust, we aim to offer several answer verification methods, some of which should be based on explainable AI and be easy for users to understand. In the future, this may include, for example, verification based on sentence similarity scores.

Currently, the system is designed for use in the biomedical domain and provides answers based on scientific articles indexed in PubMed. However, we believe that the system can be easily extended to other document formats and become a base for a personal, organizational, or corporate generative

search engine with trustworthy answers. In the future, our version may incorporate additional sources, contributing to the trust and safety of the next generation internet.

V. AVAILABILITY

Code created so far in this project is available on GitHub [23] under AGPLv3 license. Our fine-tuned qLoRA adapter model for referenced question answering based on Mistral 7B [15] is available on HuggingFace [18]. The verification models are available on HuggingFace [24][25]. More information on the project can be found on the project website: [26].

ACKNOWLEDGMENT

The project **Verif.ai** is a collaborative effort of Bayer A.G. and the Institute for Artificial Intelligence Research and Development of Serbia, funded within the framework of the NGI Search project under Horizon Europe grant agreement No 101069364.

REFERENCES

- [1] OpenAI, "Gpt-4 technical report," 2023.
- [2] A. Q. Jiang *et al.*, "Mistral 7b," *arXiv preprint arXiv:2310.06825*, 2023.
- [3] S. Bubeck *et al.*, "Sparks of artificial general intelligence: Early experiments with gpt-4," *arXiv preprint arXiv:2303.12712*, 2023.
- [4] J. S. Park, J. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative agents: Interactive simulacra of human behavior," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 2023, pp. 1–22.
- [5] H. Touvron *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [6] D. M. Katz, M. J. Bommarito, S. Gao, and P. Arredondo, "Gpt-4 passes the bar exam," *Available at SSRN 4389233*, 2023.
- [7] P. Lewis *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [8] D. Wadden, S. Lin, K. Lo, L. L. Wang, M. van Zuylen, A. Cohan, and H. Hajishirzi, "Fact or fiction: Verifying scientific claims," *arXiv preprint arXiv:2004.14974*, 2020.
- [9] L. Huang *et al.*, "A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions," *arXiv preprint arXiv:2311.05232*, 2023.
- [10] C. Malaviya, S. Lee, S. Chen, E. Sieber, M. Yatskar, and D. Roth, "Expertqa: Expert-curated questions and attributed answers," *arXiv preprint arXiv:2309.07852*, 2023.
- [11] J. Boyko *et al.*, "An interdisciplinary outlook on large language models for scientific research," *arXiv preprint arXiv:2311.04929*, 2023.
- [12] NIH. (2024) Pubmed. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/>
- [13] Amazon. (2024) Opensearch.org. [Online]. Available: <https://opensearch.org/>
- [14] N. Craswell, B. Mitra, E. Yilmaz, D. Campos, and J. Lin, "Ms marco: Benchmarking ranking models in the large-data regime," in *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 1566–1576.
- [15] F. Almeida. (2024) Mistral-7b-instruct-v0.1-sharded. [Online]. Available: <https://huggingface.co/filipealmeida/Mistral-7B-Instruct-v0.1-sharded>
- [16] Q. Jin, B. Dhingra, Z. Liu, W. W. Cohen, and X. Lu, "Pubmedqa: A dataset for biomedical research question answering," *arXiv preprint arXiv:1909.06146*, 2019.
- [17] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "Qlora: Efficient finetuning of quantized llms," *arXiv preprint arXiv:2305.14314*, 2023.
- [18] B. Basaragin. (2024) Mistral-7b-instruct-v0.1-pqa. [Online]. Available: <https://huggingface.co/BojanaBas/Mistral-7B-Instruct-v0.1-pqa>
- [19] A. Conneau and K. Khandelwal. (2024) xlm-roberta-large. [Online]. Available: <https://huggingface.co/xlm-roberta-large>

- [20] Microsoft. (2024) deberta-v3-large. [Online]. Available: <https://huggingface.co/microsoft/deberta-v3-large>
- [21] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, pp. 79–87, 1991.
- [22] A. Q. Jiang *et al.*, "Mixtral of experts," *arXiv preprint arXiv:2401.04088*, 2024.
- [23] N. Milosevic, A. Ljajic, M. Kosprdic, B. Basaragin, and D. Medvecki. (2024) Verif.ai github repository. [Online]. Available: <https://github.com/nikolamilosevic86/verif.ai>
- [24] N. Milosevic. (2024) Xlm roberta large based model fine-tuned on scifact dataset. [Online]. Available: https://huggingface.co/nikolamilosevic/SCIFACT_xlm_roberta_large
- [25] M. Kosprdic. (2024) Deberta large based model fine-tuned on scifact dataset. [Online]. Available: <https://huggingface.co/MilosKosRad/DeBERTa-v3-large-SciFact>
- [26] Verif.ai. (2024) Verif.ai website. [Online]. Available: <https://verifai-project.com>