arXiv:2504.05307v1 [cs.IR] 13 Feb 2025

Toward Total Recall: Enhancing FAIRness through AI-Driven Metadata Standardization

Sowmya S. Sundaram and Mark A. Musen

Stanford Center for Biomedical Informatics Research (BMIR), Stanford University, California, USA.

*Corresponding author(s). E-mail(s): sowmyasm@stanford.edu;
Contributing authors: musen@stanford.edu;

Abstract

Current metadata often suffer from incompleteness, inconsistency, and incorrect formatting, hindering effective data reuse and discovery. Using GPT-4 and a metadata knowledge base (CEDAR), we devised a method that standardizes metadata in scientific data sets, ensuring the adherence to community standards. The standardization process involves correcting and refining metadata entries to conform to established guidelines, significantly improving search performance and recall metrics. The investigation uses BioSample and GEO repositories to demonstrate the impact of these enhancements, showcasing how standardized metadata lead to better retrieval outcomes. The average recall improves significantly, rising from 17.65% with the baseline raw datasets of BioSample and GEO to 62.87% with our proposed metadata standardization pipeline. This finding highlights the transformative impact of integrating advanced AI models with structured metadata curation tools in achieving more effective and reliable data retrieval.

Keywords: metadata, FAIR, NLP

Introduction

Effective data sharing can be improved by adhering to FAIR standards[1]—ensuring data is Findable, Accessible, Interoperable, and Reusable. An important aspect of making data FAIR is high quality metadata[2], which refers to descriptive information about the data. Metadata play a pivotal role in organizing, categorizing, and enhancing the discoverability of data by including details such as keywords, formats, and

contextual information about the data's origin, purpose, and usage. However, manually improving metadata quality is a complex and time-consuming process due to lack of consensus, variability in community standards, etc. To address this challenge of metadata quality, we propose an automated method for standardizing metadata. In this research endeavor, we investigate the use of Natural Language Processing (NLP) techniques for automated metadata correction and assess their impact on recall performance, a quantitative metric for data retrieval.

The focal point of our investigation is on improving the metadata, as high-quality metadata enhance the searchability of public datasets, ultimately benefiting scientific research (Figure 1).

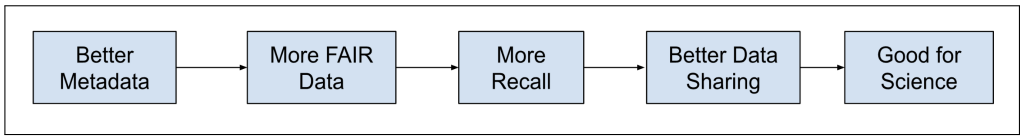

Fig. 1 The Big Picture

Metadata comprise lists of field name–value pairs that may be augmented reporting guidelines. In Figure 2—a snapshot of a record sourced from the BioSample repository of the National Center for Biotechnology Information (NCBI)[3]—the black box highlights an example of metadata. In this example, the field name is tissue and the field value is lung cancer, which is inaccurate as lung cancer is not a type of tissue. Searches for scientific datasets primarily involve querying metadata. Consequently, a researcher querying for the name–value pair of tissue would overlook this potentially valuable record during their search. As important as metadata are, current metadata are usually incomplete, inconsistent, and incorrectly formatted[4]. Standardizing metadata is therefore a necessary process for ensuring that datasets are properly described and accessible, so that they they can be more easily reused and integrated with other datasets, thus facilitating secondary use and analysis of scientific data..

The process of metadata standardization involves the analysis of the language used to specify name–value pairs. To the best of our knowledge, there have been a lack of concerted efforts in this research area, underscoring the relevance of our lab's consistent work in automated metadata standardization. In a previous attempt to enhance BioSample records using NLP techniques[5], our team employed the earliest available embeddings (vectorial representations of words learned by training over text)—specifically word2vec[6] and GloVe[7]—to represent metadata terms. These embeddings are powerful tools for capturing semantic relationships among words. However, we encountered significant scalability challenges because these embeddings were limited to terms present in the training set, thereby neglecting the specialized and niche terms often found in medical and scientific data.

To address these limitations, we expanded our approach by incorporating Generative Pre-trained Transformer 4 (GPT-4)[8] embeddings[9], which provide a more comprehensive and flexible representation of language. Unlike traditional embeddings, GPT-4 can accommodate a broader range of terms, including those unseen during

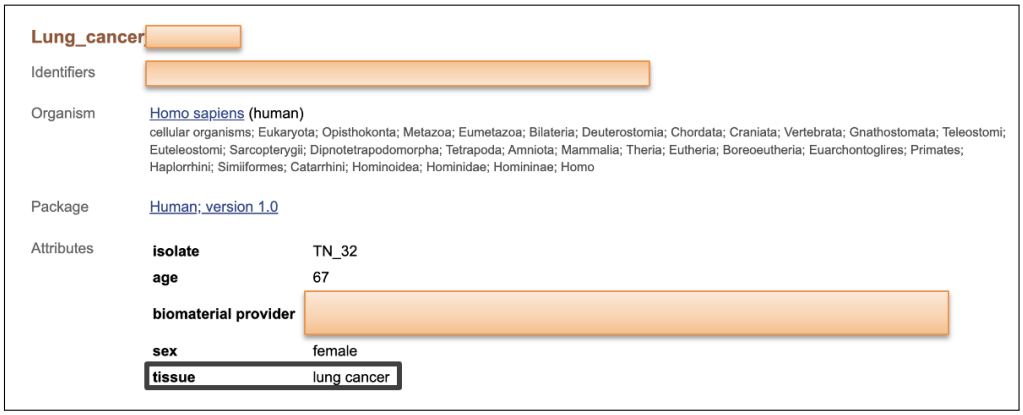

Fig. 2 Record from BioSample where the black box highlights a field name–value pair. In this example, the field name *tissue* is wrongly associated with field value *lung cancer*. Orange boxes mask identifying information

training or those that are less common or domain-specific. We subsequently analyzed how GPT-4 can improve the metadata of a small dataset of 200 metadata records from BioSample[10] through a peer review board. Our results indicated that the domain experts preferred GPT-4 augmented with cues from a metadata knowledge source in its prompts.

In our current study, we expand our investigation to a large dataset of 2,400 samples each from BioSample and the Gene Expression Omnibus (GEO)[11], evaluating whether the necessity of a metadata knowledge source remains valid across different LLMs, including GPT-4, Large Language Model Meta AI (LLaMA-3) [12], and MedLLaMA-2 [13]. Instead of a peer review board, we measure the impact on search outcomes. While many authors highlight the theoretical benefits of the FAIR principles, the absence of quantitative analysis complicates the assessment of the tangible effects of FAIR data. A key contribution of our work is the quantification of the retrieval of BioSample records following metadata standardization and its effect on findability of scientific data.

The Search Experiment

Our experiment is to improve metadata standardization of BioSample records and measure the impact of metadata improvements on search metrics. We begin by detailing the dataset utilized in our experiment.

Dataset

For the experiment, we utilized records from two separate datasets: one from BioSample and another from GEO. Each dataset contains 2,400 samples related to three types of cancer: lung cancer, liver cancer, and ovarian cancer. The queries used to retrieve the records were as follows: lung cancer *(query: lung cancer[All Fields] AND "human*

1 0"[filter])[1], liver cancer *(query: liver cancer[All Fields] AND "human 1 0"[filter])*, and ovarian cancer *(query: ovarian cancer[All Fields] AND "human 1 0"[filter])*. For each query, we randomly sampled 800 records from both datasets. We initially sampled 1,000 records for each query, removed those with formatting errors, and selected the maximum uniform number of well-formatted records across all datasets, which happened to be 800. Therefore, our total dataset comprises 4,800 records: 2,400 from BioSample and 2,400 from GEO (Table 1).

Dataset Description			
BioSample (2400)	LungCancer (800)	LiverCancer (800)	OvarianCancer (800)
GEO (2400)	LungCancer (800)	LiverCancer (800)	OvarianCancer (800)
Total - 4800 records			

Table 1 Dataset Composition

Method

Manual standardization of metadata involves a human who typically interprets applicable metadata standards and examines the metadata record to assess adherence, subsequently suggesting corrections. Similarly, for a given record, we sought to replicate this process by directing the LLM to utilize the data dictionary[14], which includes the allowed list of metadata name–value pairs, their formats, and descriptions, to correct the record. This process is illustrated in Figure 3 for GPT-4.

Building upon our previous efforts, we also measured the impact of using a metadata knowledge source on search metrics. Specifically, we used metadata templates sourced from the Center for Expanded Data Annotation and Retrieval (CEDAR)[15]. CEDAR houses metadata templates that ease data sharing. CEDAR templates are machine-readable data structures that describe how metadata should be structured, providing guidelines for field names and corresponding values. The guidelines for field values include data format restrictions, value set restrictions, and ontological restrictions to ensure that metadata entries are consistent with established standards (Table 2). Given its comprehensive repository of metadata templates[16], CEDAR is a natural choice for supporting metadata-related endeavors. For example, as shown in Figure 4, CEDAR recommends values from the UBERON ontology for the attribute name "tissue," restricting the user to select standardized terms.

We use the described framework across a set of LLMs (GPT-4, LLaMA-3 and MedLLaMA2) and tested it on our dataset. In this manner, we have two versions of our each dataset - one that is augmented by using the official data dictionary (LLM+DD) and another which is augmented by the CEDAR templates (LLM+CEDAR). We then perform search on a few queries for the field name *tissue*. The retrieval from the dataset is done based on exact match. We then measure commonly used search metrics - precision, recall and F1-score.

[1] This query translates to searching all the fields of a record to match with *lung cancer* and filtering the result set to only contain records from human samples

Input to GPT-4: Metadata Record with Data Dictionary

Convert the record: "age:67 — sex:female — tissue:lung cancer" to the format given by the BioSample data dictionary:

"Name — Description — Value format

age — age at the time of sampling; relevant scale depends on species and study, e.g. could be seconds for amoebae or centuries for trees — {float} {unit}

...

...

tissue — Type of tissue the sample was taken from — {term}

disease — list of diseases diagnosed; can include multiple diagnoses. the value of the field depends on host; for humans the terms should be chosen from DO (Disease Ontology), free text for non-human. For DO terms, please see http://gemina.svn.sourceforge.net/viewvc/gemina/trunk/Gemina/ontologies/gemina_symptom.obo?view=log — {term}

Output:

biosample_provider: NA

age: 67

sex: female

disease: lung cancer

Fig. 3 GPT-4 Correction of BioSample - The prompt given to GPT-4 includes both the metadata to be corrected and the entire data dictionary to which BioSample records should adhere.

Feature	**Description**
Field Definitions	Field names and clear definitions of what each field is intended to represent
Data Format Restrictions	Specifications on the format of data entries, such as date formats or string patterns
Value Restrictions	Predefined lists or ontological restrictions

Table 2 Key Features of CEDAR Templates

Evaluation Metrics

In our search experiment, we lack access to a gold standard. To derive the gold standard values for our records, we manually examine the 'tissue' values in our BioSample dataset. We then develop simplistic rules to assign tissue values based on the subcohort (Table 3). This approach allows us to devise an *approximate* gold standard for our experiment.

We employed traditional search metrics—precision, recall, and F1-score to evaluate the datasets. The description of these metrics is provided in Table 4. We performed experiments on two queries: *tissue:(major organ)* and *tissue:blood*.

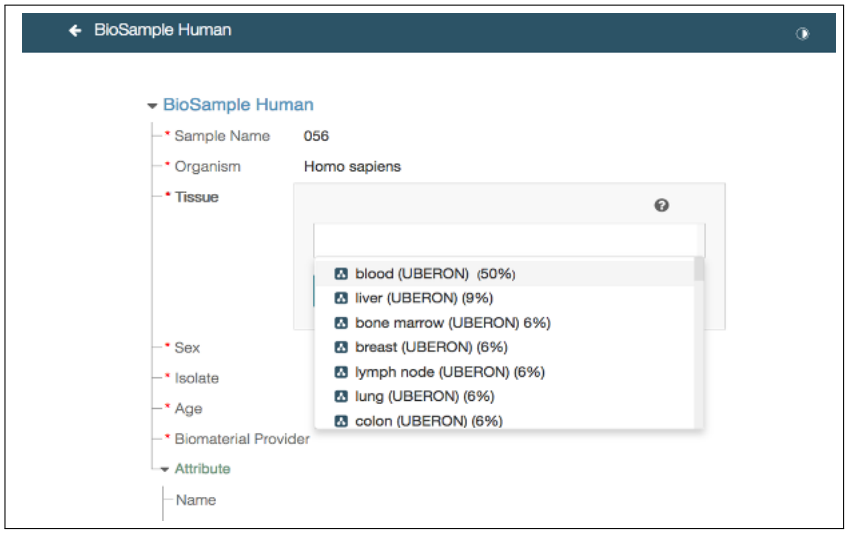

Fig. 4 A screen capture depicting the CEDAR metadata entry form derived from the template for BioSample. Field names specified in the data dictionary such as *tissue* and *organism* are listed. CEDAR recommends options for the "tissue" field from the UBERON ontology[17].

Ovarian Cancer Tissue	**Liver Cancer Tissue**	**Lung Cancer Tissue**
a) Initialize label as 'unknown'.	a) Initialize label as 'unknown'.	a) Initialize label as 'unknown'.
b) If the tissue contains the word 'ovary' or 'ovarian', set label to 'ovary'.	b) If the tissue contains the word 'liver' or 'HCC', set label to 'liver'.	b) If the tissue contains the word 'lung', set label to 'lung'.
c) Else, if the tissue contains the word 'plasma', set label to 'plasma'.	c) Else, if the tissue contains the word 'PBMC' or 'blood', set label to 'blood'.	c) Else, if the tissue contains the word 'PBMC' or 'blood', set label to 'blood'.
d) Else, if the tissue contains the word 'PBMC' or 'blood', set label to 'blood'.	d) Else, if the tissue contains the word 'lymph', set label to 'lymph'.	d) Else, if the tissue contains the word 'lymph', set label to 'lymph'.
	e) Else, if the tissue contains the word 'plasma', set label to 'plasma'.	e) Else, if the tissue contains the word 'plasma', set label to 'plasma'.

Table 3 Annotation Rules for Gold Standard

Results

We first present the results of our search experiment on overall recall metrics with GPT-4. The average recall value rises from 17.65% to 62.87% from the baseline raw datasets of BioSample and GEO, to the GPT4+CEDAR versions.

Input to GPT-4: Record with CEDAR template
Convert the record: "age:67 — sex:female — tissue:lung cancer" to the format given by the CEDAR template
"Name — Description — Comments
age — age at the time of sampling; relevant scale depends on species and study, e.g. could be seconds for amoebae or centuries for trees — {float} {unit}
tissue — type of tissue sample — **Must be from UBERON ontology**
...
...
disease — Name of the disease — **Must be from Disease Ontolology (DO) ontology** "
GPT-4 Output:
biosample_accession: NA
organism: Homo sapiens
age: 67
sex: female
tissue: lung
disease: lung cancer
...
...
population: NA
race: NA
sample_type: tissue

Fig. 5 GPT-4 Correction of BioSample with CEDAR template

Metric	**Explanation**	**Formula**
Precision	Precision is the ratio of correctly retrieved relevant instances to the total retrieved instances.	$\frac{\text{True Positives}}{\text{True Positives}+\text{False Positives}}$
Recall	Recall is the ratio of correctly retrieved relevant instances to the total relevant instances.	$\frac{\text{True Positives}}{\text{True Positives}+\text{False Negatives}}$
F1-Score	F1-Score is the harmonic mean of precision and recall.	$2 \times \frac{\text{Precision}\times\text{Recall}}{\text{Precision}+\text{Recall}}$

Table 4 Explanation of Search Metrics: Precision, Recall, and F1-Score

Liver Cancer

First, we evaluated the methods on the BioSample and GEO Liver Cancer datasets.

Ovarian Cancer

For the domain of ovarian cancer , we compared the search efficacy of different datasets on two queries: 'tissue:ovary' and 'tissue:blood'. The results are presented below in Figures 9 and 10, with precision, recall, and F1-score calculated for each method.

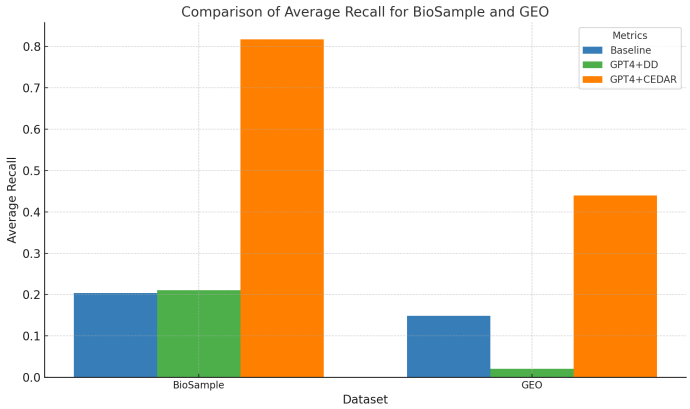

Fig. 6 Comparison of average recall values for BioSample and GEO datasets across three result sets: Raw Samples, GPT4+DD, and GPT4+CEDAR

Lung Cancer

Finally, we tested the lung cancer related datasets for the 'tissue:lung' and 'tissue:blood' queries. Figures 11 and 12 summarize the results, showing a strong performance of GPT4 + CEDAR in both cases.

Discussion

Standardized metadata are essential for making data FAIR[15]. Effective metadata design ensures that data can be easily discovered, understood, and consumed by computer algorithms, thus fostering collaboration and advancing scientific discovery. Our results demonstrate that data's improved conformance to FAIR principles through more standards-adherent metadata, leads to enhanced recall during dataset searches.

On closer examination of our results, we found only minor enhancements after using GPT-4 with the respective repository's data dictionary. We examined some of the errors and surmised that correction of existing metadata is difficult due to variations in describing a concept (such as *lung afflicted with cancer*, *lung cancer sample*). In some cases, this variability even caused a slight dip in recall. The experiments suggested a mechanism is required for providing guard rails to the generation process. This suggestion is also backed by recent studies in LLM alignment to real-world problems that incorporate knowledge to reduce hallucination and enhance reasoning[18, 19]. In our method, we utilized CEDAR for the providing the knowledge in the form of a metadata template. Notably, our LLM prompted method, GPT4+CEDAR, consistently excels in retrieval and recall, although there is a slight decrease in precision. In scientific data discovery, researchers typically aim to retrieve as many datasets as possible and can identify the wrongly retrieved ones, if they are small in number. This observation highlights the significance of a human-in-the-loop approach.

Metadata correction with LLMs presents its own challenges. We experimented with LlaMA-3, Mistral, and MedLLaMA-2. While GPT-4 successfully generated accurate corrections for 4,800 samples, the other LLMs produced formatting errors. Hence, we

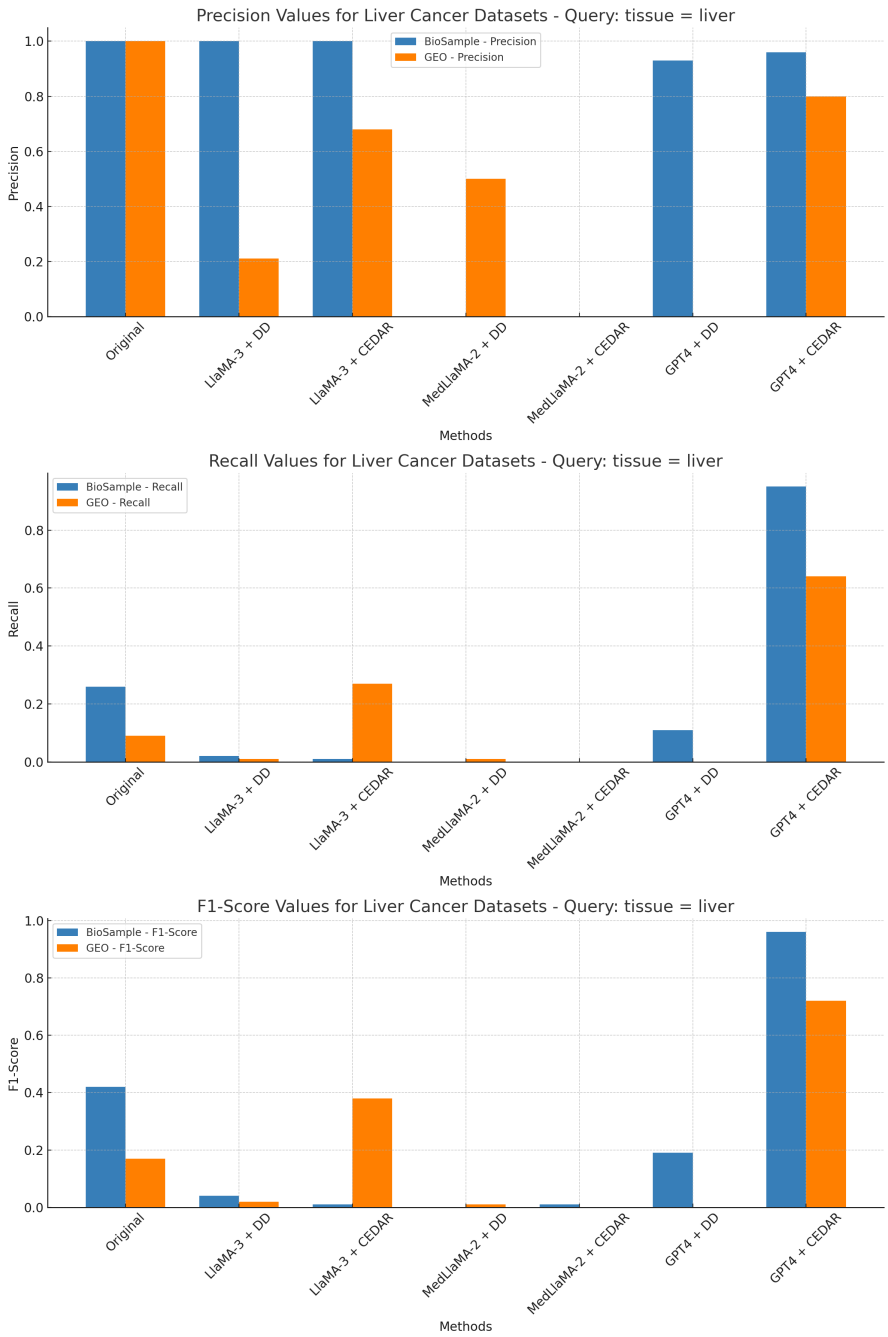

Fig. 7 Liver Cancer Datasets derived from BioSample and GEO: Query - tissue:liver

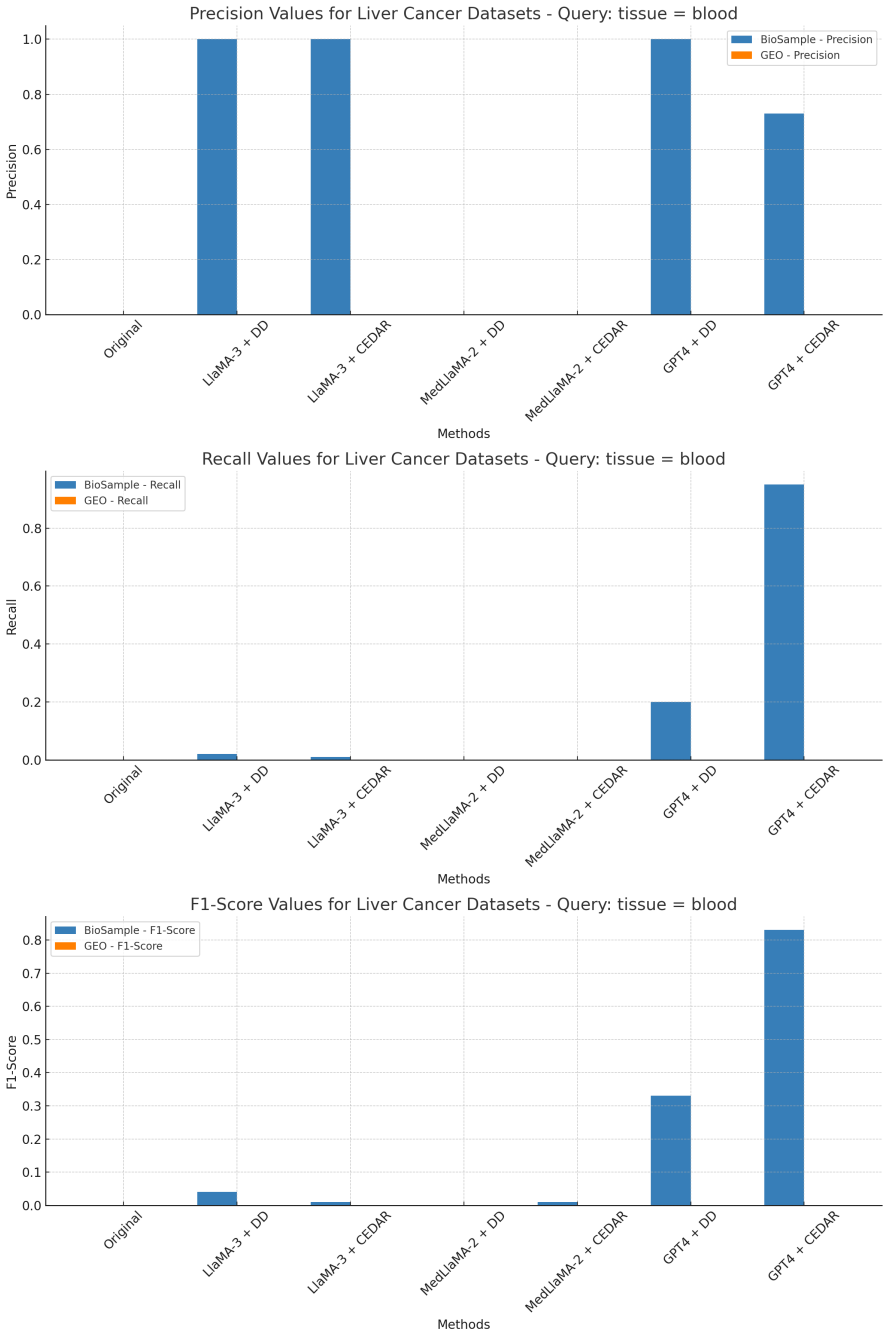

Fig. 8 Liver Cancer Datasets derived from BioSample and GEO: Query - tissue:blood

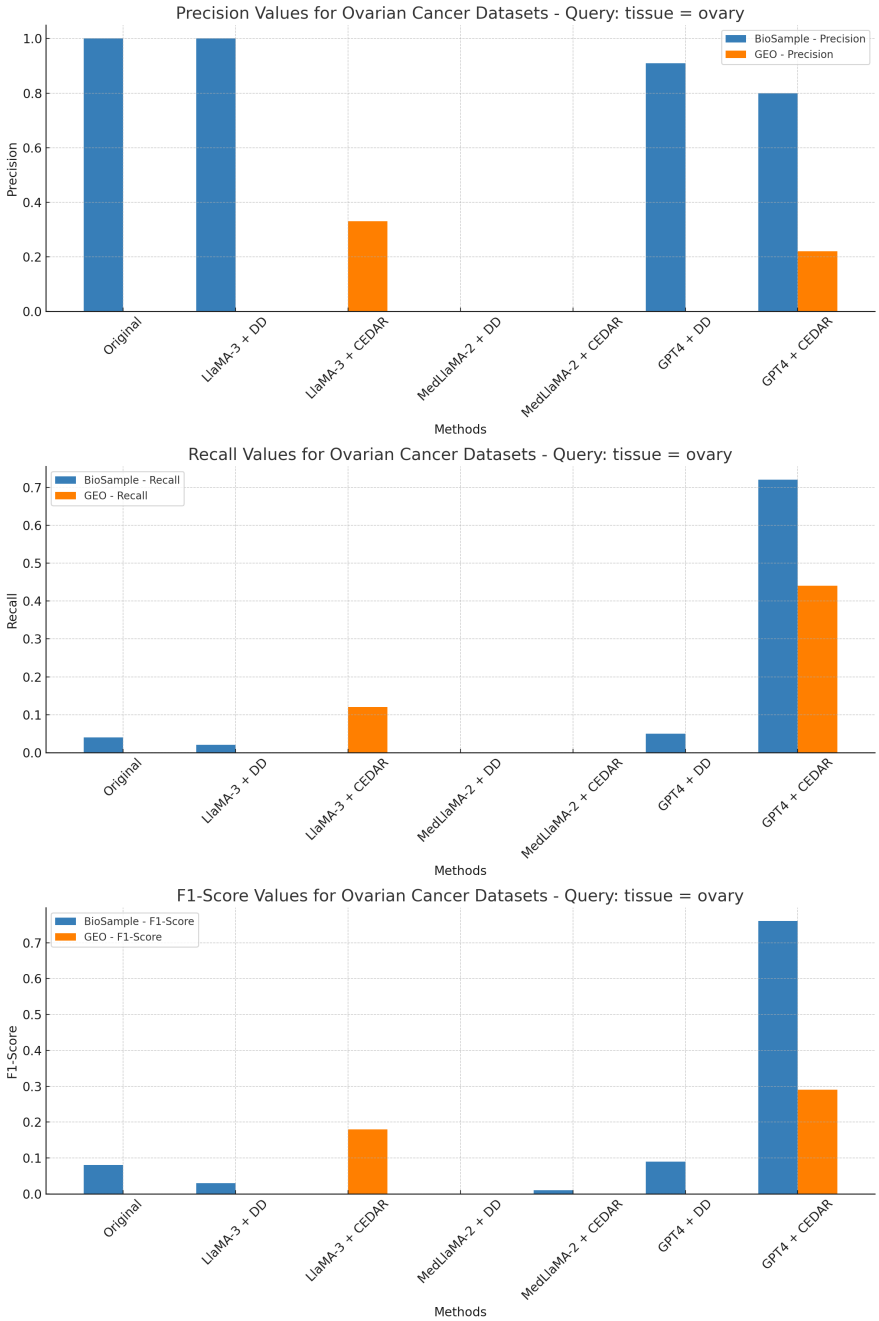

Fig. 9 Ovarian Cancer Datasets derived from BioSample and GEO: Query - tissue:ovary

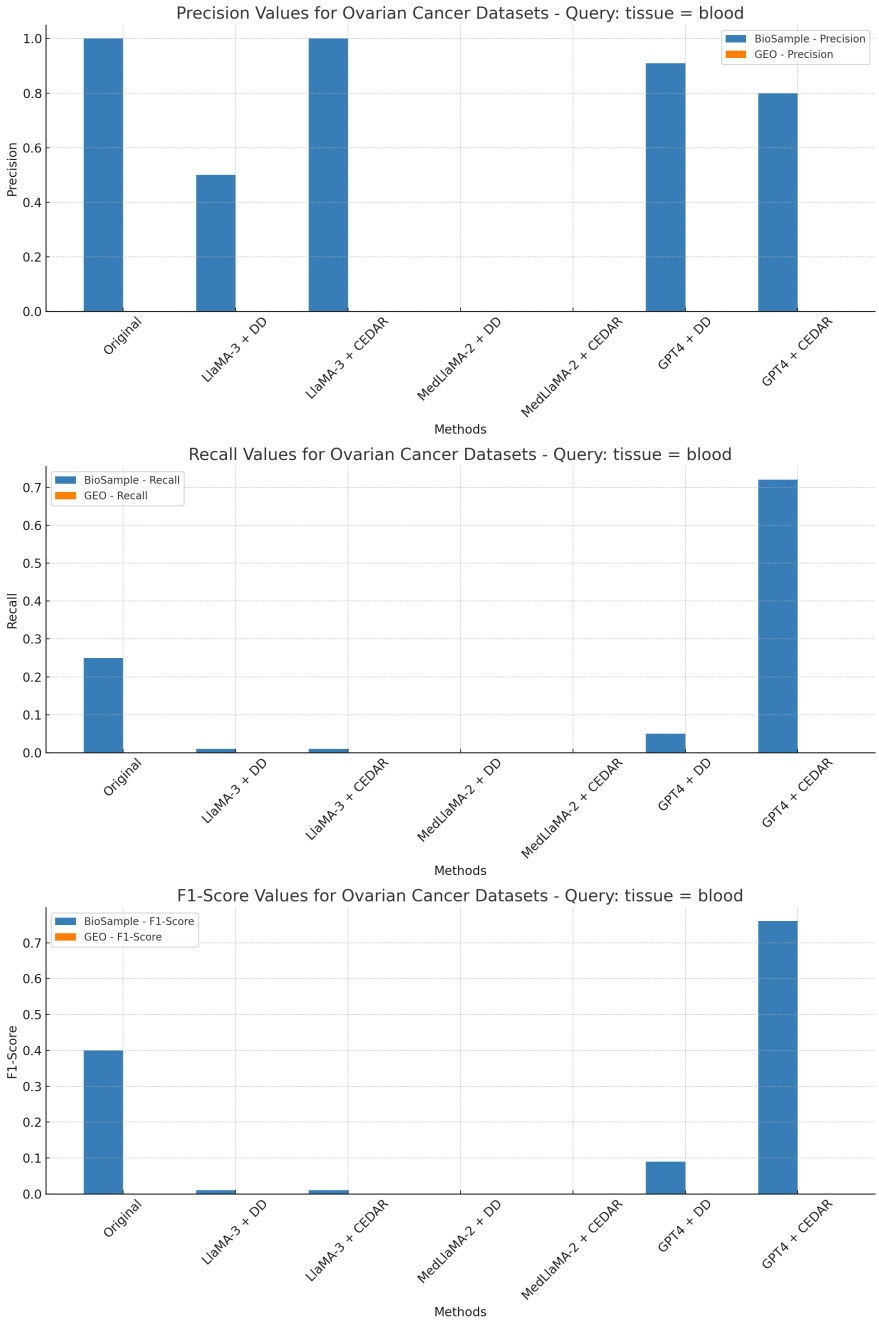

Fig. 10 Ovarian Cancer Datasets derived from BioSample and GEO: Query - tissue:blood

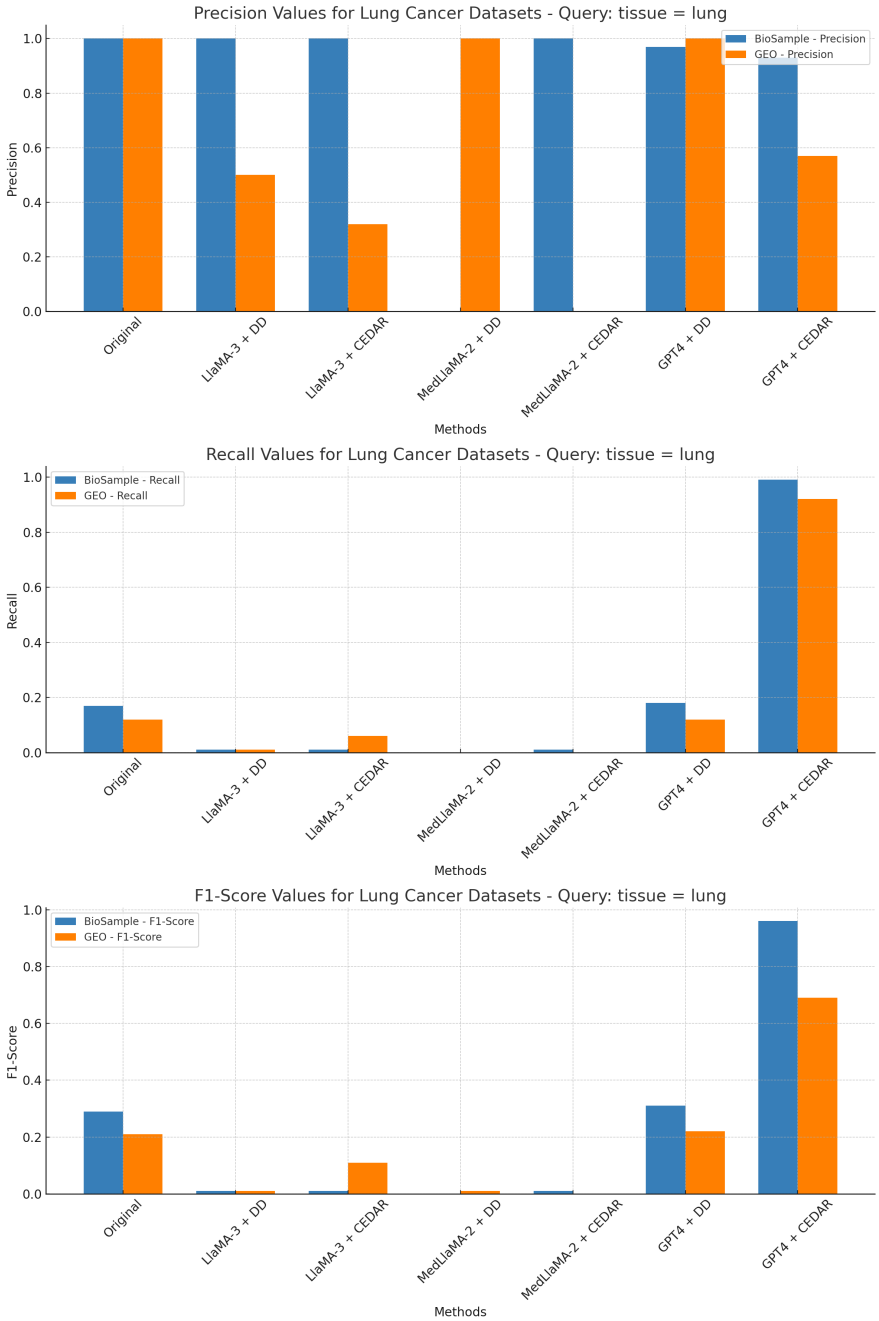

Fig. 11 Lung Cancer Datasets derived from BioSample and GEO: Query - tissue:lung

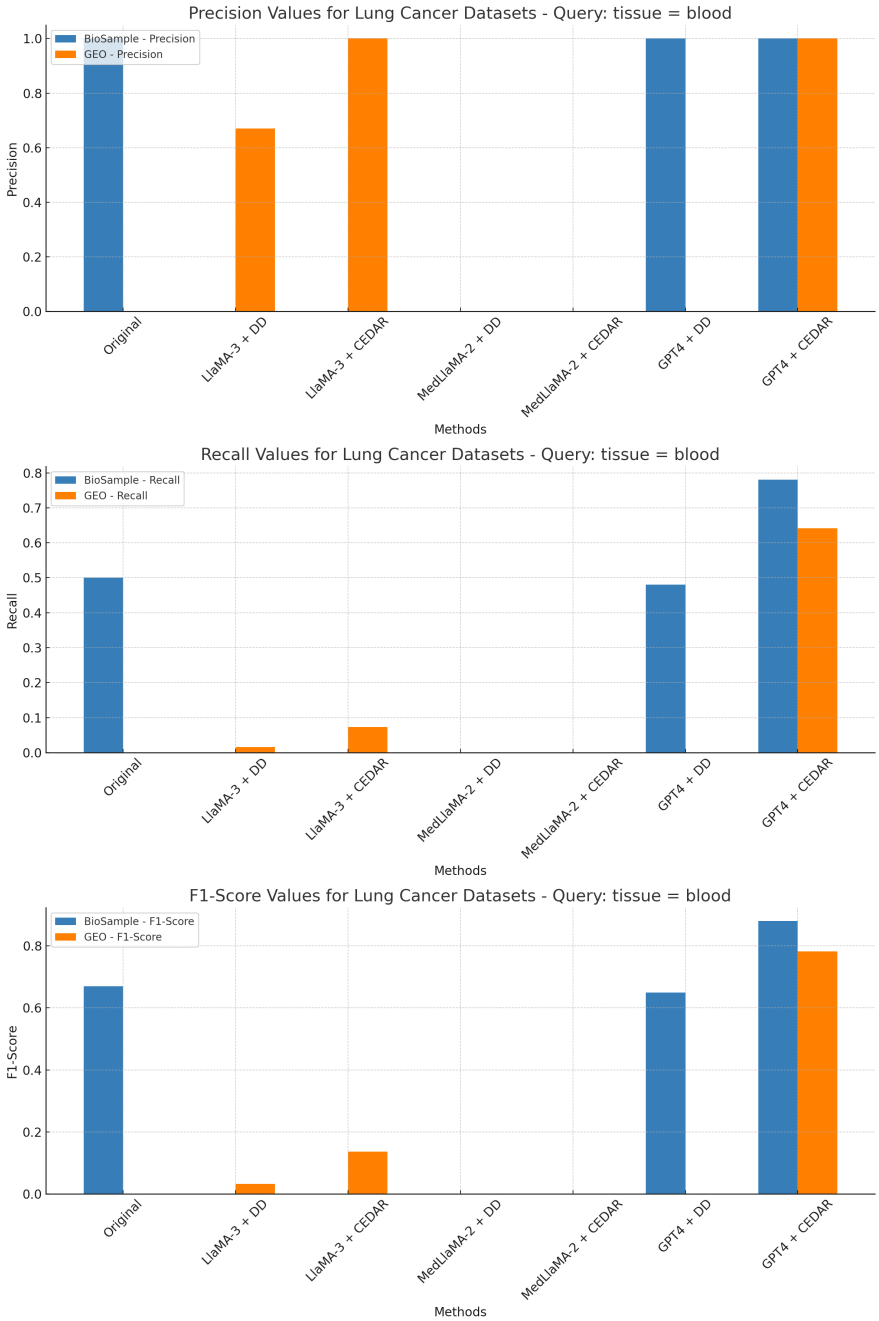

Fig. 12 Lung Cancer Datasets derived from BioSample and GEO: Query - tissue:blood

had to employ complex post-processing of the output, hand-crafted for each LLM, thereby affecting deployment at scale. Consistently, GPT-4+CEDAR achieves the best recall values.

Another insight from the qualitative analysis of precision values is that the drop in precision can be attributed to errors introduced by GPT-4. For example, we observed instances where GPT-4 incorrectly changes the clearly stated *tissue* value, such as *blood*, to *lung*, even after augmentation with CEDAR. This typically occurs when a BioSample record is longer than average, and the presence of other metadata entries mentioning *lung* seem to interfere with the generation of the correct metadata. We also observed that the rule-based labeling approach is conservative and labels many tissue values as *unknown* and GPT-4 actually extracts the correct tissue value but is penalized because the gold standard is not accurate and is an approximation.

The analyses above highlight the need for additional safeguards following the integration of CEDAR. To address this, we implemented a method that retrieves all potential samples based on the new records while simultaneously displaying both the old and new records. This allows users to make informed decisions about which record to use, effectively reducing the cognitive load for researchers engaged in secondary data analysis. We will release this software for BioSample as a part of our data for this work.

We plan further enhancements by integrating ontology values into a retrieval framework using Retrieval-Augmented Generation (RAG) pipelines, ensuring consistent field naming. An ambitious goal involves developing a shadow database for the entirety of BioSample—reported to contain approximately 5 million samples—leveraging the CEDAR-augmented approach.

Conclusion

Efforts to achieve FAIR data have become increasingly prominent in the scientific community, yet mere assertions of data FAIRness without the implementation of rich, standardized metadata fail to render datasets truly findable. Our study underscores that the cornerstone of data FAIRness is the meticulous standardization of metadata. By employing advanced NLP techniques and a structured metadata knowledge source (CEDAR) to correct metadata, we demonstrated significant improvements in search recall and precision. The pronounced gains in retrieval performance presented in our results emphasize that through rigorous adherence to metadata standards, the advantages of data FAIRness can be actualized, thereby enhancing data accessibility and usability in scientific research.

Data Availability

will add the link to visualization code here

Acknowledgements

This work was support in part by grant R01 LM013498 from the National Library of Medicine.

References

[1] Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., Silva Santos, L.B., Bourne, P.E., *et al.*: The fair guiding principles for scientific data management and stewardship. Scientific data **3**(1), 1–9 (2016)

[2] Musen, M.A.: Without appropriate metadata, data-sharing mandates are pointless. Nature **609**(7926), 222–222 (2022)

[3] Barrett, T., Clark, K., Gevorgyan, R., Gorelenkov, V., Gribov, E., Karsch-Mizrachi, I., Kimelman, M., Pruitt, K.D., Resenchuk, S., Tatusova, T., *et al.*: Bioproject and biosample databases at ncbi: Facilitating capture and organization of metadata. Nucleic acids research **40**(D1), 57–63 (2012)

[4] Gonçalves, R.S., Musen, M.A.: The variable quality of metadata about biological samples used in biomedical experiments. Scientific data **6**(1), 1–15 (2019)

[5] Gonçalves, R.S., Kamdar, M.R., Musen, M.A.: Aligning biomedical metadata with ontologies using clustering and embeddings. In: The Semantic Web: 16th International Conference, ESWC 2019, Portorož, Slovenia, June 2–6, 2019, Proceedings 16, pp. 146–161 (2019). Springer

[6] Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. Advances in neural information processing systems **26** (2013)

[7] Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1532–1543 (2014)

[8] Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)

[9] S Sundaram, S., Musen, M.: Making metadata more fair using large language models (2023) https://doi.org/10.4126/FRL01-006444995

[10] Sundaram, S.S., Solomon, B., Khatri, A., Laumas, A., Khatri, P., Musen, M.A.: Use of a Structured Knowledge Base Enhances Metadata Curation by Large Language Models (2024)

[11] Edgar, R., Domrachev, M., Lash, A.E.: Gene expression omnibus: Ncbi gene expression and hybridization array data repository. Nucleic acids research **30**(1), 207–210 (2002)

[12] Touvron, H., al.: Llama 3: Open and efficient foundation language models. arXiv

preprint arXiv:2308.12001 (2023)

[13] Alvi, O., al.: Medllama: A large language model for medicine. arXiv preprint arXiv:2308.09109 (2023)

[14] NCBI: BioSample Data Dictionary. https://submit.ncbi.nlm.nih.gov/biosample/template/?organism-organism_name=&organism-taxonomy_id=&package-0=Human.1.0&action=definition

[15] Musen, M.A., Bean, C.A., Cheung, K.-H., Dumontier, M., Durante, K.A., Gevaert, O., Gonzalez-Beltran, A., Khatri, P., Kleinstein, S.H., O'Connor, M.J., *et al.*: The center for expanded data annotation and retrieval. Journal of the American Medical Informatics Association **22**(6), 1148–1152 (2015)

[16] Musen, M.A., O'Connor, M.J., Schultes, E., Martínez-Romero, M., Hardi, J., Graybeal, J.: Modeling community standards for metadata as templates makes data fair. Scientific Data **9**(1), 696 (2022) https://doi.org/10.1038/s41597-022-01815-3

[17] Mungall, C.J., Torniai, C., Gkoutos, G.V., Lewis, S.E., Haendel, M.A.: Uberon, an integrative multi-species anatomy ontology. Genome biology **13**(1), 5 (2012)

[18] Dessimoz, C., Thomas, P.D.: Ai and the democratization of knowledge. Scientific Data **11**(1), 268 (2024)

[19] Lu, P., Peng, B., Cheng, H., Galley, M., Chang, K.-W., Wu, Y.N., Zhu, S.-C., Gao, J.: Chameleon: Plug-and-play compositional reasoning with large language models. Advances in Neural Information Processing Systems **36** (2024)