

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60

PAPER

# Pre-Meta: Priors-augmented Retrieval for LLM-based Metadata Generation

Phil Tinn,<sup>1,\*,#</sup> Sondre Sørbo,<sup>1,\*</sup> Shanshan Jiang,<sup>1,\*</sup> Konstantinos Voutetakis,<sup>2</sup> Sotiris Moudouris Giounis,<sup>2</sup> Eleftherios Pilalis,<sup>2,3,‡</sup> Olga Papadodima<sup>2,‡</sup> and Dumitru Roman<sup>1,4</sup>

<sup>1</sup>SINTEF Digital, 4 Strindvegen, 7034, Trondheim, Norway, <sup>2</sup>Institute of Chemical Biology, National Hellenic Research Foundation, 48 Vassileos Constantinou Ave., 11635, Athens, Greece, <sup>3</sup>e-NIOS Applications PC, 196 Sygrou Ave., 17671, Kallithea, Greece and <sup>4</sup>Bucharest University of Economic Studies, Piața Romană nr. 6, Bucharest, Romania

<sup>#</sup>Corresponding author: phil.tinn@sintef.no

<sup>\*,‡</sup>Equal Contribution

FOR PUBLISHER ONLY Received on Date Month Year; revised on Date Month Year; accepted on Date Month Year

## Abstract

**Motivation:** While high-throughput sequencing technologies have dramatically accelerated genomic data generation, the manual processes required for dataset annotation and metadata creation impede the efficient discovery and publication of these resources across disparate public repositories. Large Language Models (LLMs) have the potential to streamline dataset profiling and discovery. However, their current limitations in generalizing across specialized knowledge domains, particularly in fields such as biomedical genomics, prevent them from fully realizing this potential. This paper presents Pre-Meta, an LLM-agnostic and domain-independent data annotation pipeline with an enriched retrieval procedure that leverages related priors—such as pre-generated metadata tags and ontologies—as auxiliary information to improve the accuracy of automated metadata generation.

**Results:** Validated using five selected metadata fields sampled across 1500 papers, the Pre-Meta assisted annotation experiment—without finetuning and prompt optimization—demonstrates a systemic improvement in the annotation task: shown through a 23%, 72%, and 75% accuracy gain from conventional RAG adoptions of GPT-4o mini, Llama 8B, and Mistral 7B respectively.

**Availability:** The code, data access, and scripts are available at: <https://github.com/SINTEF-SE/LLMDap>.

## 1 Introduction

The sharing of data in biomedical and genomic research is essential for accelerating scientific discovery. The main challenge lies in effectively discovering, sharing, integrating and analyzing biomedical datasets in various isolated portals and databases (Byrd et al. 2020). The Research Data Repository Registry<sup>1</sup>, a comprehensive global registry, lists more than 2,000 repositories across various disciplines, with a significant portion dedicated to biomedical sciences. For example, the National Center for Biotechnology Information (NCBI)<sup>2</sup> maintains several databases, including GenBank, which, as of 2023, contains more than 2 billion sequences. Similarly, NCBI's Gene Expression Omnibus (GEO) archives gene expression data from over 100,000 experiments. The European Bioinformatics Institute (EBI) hosts ArrayExpress<sup>3</sup>,

including data from more than 70,000 assays. Collectively, these repositories house millions of datasets, each varying in size and complexity.

Published datasets are accompanied by metadata that provide descriptive information about the data. To enhance data discovery, it is crucial to annotate datasets with rich, high-quality metadata. However, fragmentation of data repositories creates significant obstacles for both data submission and discovery, which are crucial tasks in biomedical research. In data discovery, searching across multiple platforms is time-consuming and often results in incomplete or inconsistent findings due to differences in data formats, metadata descriptions, and access policies. In data submission, researchers face difficulties navigating the diverse requirements and formats imposed by different repositories when submitting datasets, thus necessitating extensive reformatting and manual curation, which increases the overall workload and risk of errors.

Large Language Models (LLMs), an increasingly powerful auto-regressive text generation technology based on the

<sup>1</sup> <https://www.re3data.org>  
<sup>2</sup> <https://www.ncbi.nlm.nih.gov>  
<sup>3</sup> <https://www.ebi.ac.uk/biostudies/arrayexpress>

Transformer architecture (Vaswani et al. 2023), hold great potential in facilitating the processing and standardization of diverse metadata formats, thereby enhancing searchability and interoperability across disparate repositories. Retrieval-Augmented Generation (RAG) (Lewis et al. 2021) systems can further support discovery by retrieving relevant information from multiple sources and generating context-aware insights, minimizing the need for manual searches (Yang et al. 2025). In data submission, LLMs can assist researchers by automating metadata generation, ensuring adherence to repository requirements, and streamlining data formatting, ultimately reducing administrative burden.

Currently, several genomic data repositories are available, such as ArrayExpress<sup>4</sup>, GEO<sup>5</sup>, and cBioPortal<sup>6</sup>. Each portal has its own defined metadatasets, but there is no harmonization across portals. For example, the metadata field of "study type" in ArrayExpress maps directly to "experiment type" in GEO, whereas cBioPortal uses fields such as "type of cancer", "genetic alteration type" and "study description" to define the characteristics of the dataset. Furthermore, repositories such as ArrayExpress and GEO differ in their metadata annotation schemas, with ArrayExpress employing the MAGE-TAB format (Brazma et al. 2001) and GEO utilizing the MINiML schema (Barrett et al. 2006), reflecting different priorities in metadata structuring. This heterogeneity means that data providers must use portal-specific metadata formats and follow distinct submission procedures, while data consumers must use different methods (e.g., APIs, filters, or interfaces) that may return inconsistent or incomplete results due to inadequate dataset annotation. As a result, researchers often resort to searching the literature for dataset discovery rather than relying on data portals, highlighting the urgent need for more structured and harmonized metadata annotation.

A reliable approach to addressing these inconsistencies is the use of biomedical ontologies, which provide structured vocabularies for dataset annotation in a standardized way. Ontologies such as the Ontology for Biomedical Investigations (OBI)<sup>7</sup> and the Experimental Factor Ontology (EFO)<sup>8</sup> define hierarchical relationships between biological concepts, enabling better metadata annotation and semantic interoperability across repositories. ArrayExpress, for instance, integrates EFO to ensure metadata consistency, while GEO aligns with NCBI's ecosystem, incorporating structured metadata formats that facilitate machine-readability. Ontologies are also foundational in clinical and translational research, where resources such as SNOMED CT and MeSH support standardization in electronic health records (EHRs) and biomedical literature indexing (Bodenreider, Cornet, and Vreeman 2018). Despite their transformative potential, biomedical ontologies are often underutilized in dataset annotation, limiting their impact on improving metadata discoverability and usability.

Recent advances in LLMs have demonstrated their ability to understand and generate structured and unstructured biomedical data, offering potential solutions to metadata fragmentation. However, LLMs are inherently biased due to their reliance on training data, which can be incomplete, imbalanced, or outdated (Thapa and Adhikari 2023). These

biases manifest as an over-representation of well-studied diseases, datasets, and populations, while under-representing rare conditions, minority groups, and non-English sources (Nazi and Peng 2024). Furthermore, discrepancies in terminology, ontologies, and regional clinical practices often lead LLMs to propagate inconsistencies in metadata generation (Ayoub et al. 2024). This underscores the importance of integrating curated biomedical ontologies and leveraging domain-specific priors to improve the reliability of LLM-generated metadata annotations and dataset discovery.

Relevant methodologies can be inferred from the adjacent area of the Extreme Multi-label Classification (XMC) problem, which refers to a supervised learning setting in which a data point may be associated with multiple labels drawn from an extremely large label space (Bhatia et al. 2016). Relevant XMC works thus far have focused mostly on fine-tuning (Remy, Demuynck, and Demeester 2022), prompt optimization (D'Oosterlinck et al. 2024), and multiplying the use of LLM calls (Zhu and Zamani 2023) to improve classification performance (further described in Section 4.4 Related Work); they have yet to be tested and extended to accommodate the diversity of data model standards seen across repositories and the heterogeneity of metadata fields inherent in genomic datasets.

To address these gaps, this research aims to leverage accessible structured data and Generative AI to provide rich, context-aware, and high-quality metadata annotation for biomedical datasets in an automated and efficient manner, supporting compliance with biomedical ontologies and repository/federating standards, such as Beacon V2 (Rueda et al. 2022a) defined by the Global Alliance for Genomics and Health (GA4GH).

In this paper, we introduce Pre-Meta, an LLM-based RAG pipeline for genomic metadata extraction that incorporates controlled vocabularies—in the form of pre-generated metadata and ontologies—as auxiliary data to improve annotation accuracy. Pre-Meta is model- and domain-agnostic, meaning that it provides a generic and adaptable metadata annotation procedure that is independent of both the LLM models used and the target metadata schema. We evaluated Pre-Meta using state-of-the-art open-weight and proprietary LLM models on real-world datasets curated from ArrayExpress and Europe PubMed Central (Europe PMC) and assessed its performance across five representative metadata fields. The results demonstrate that Pre-Meta significantly improves metadata annotation accuracy and highlight several systematic barriers that warrant further investigation. Furthermore, our work contributes to improving data stewardship promoted through FAIR principles (Findable, Accessible, Interoperable, Reusable) by reducing human intervention in the general process of data profiling, harmonization, and their downstream applications for data discovery.

## 2 Pre-Meta

### 2.1 Task Definition

We define the task of data annotation as a process that, for a dataset or an asset of interest for (re)annotation, fills out a given form with metadata about the asset, using values that conform to the schema. The process takes as input:

- A textual description  $x$  of the asset, such as a published paper, research log, etc.

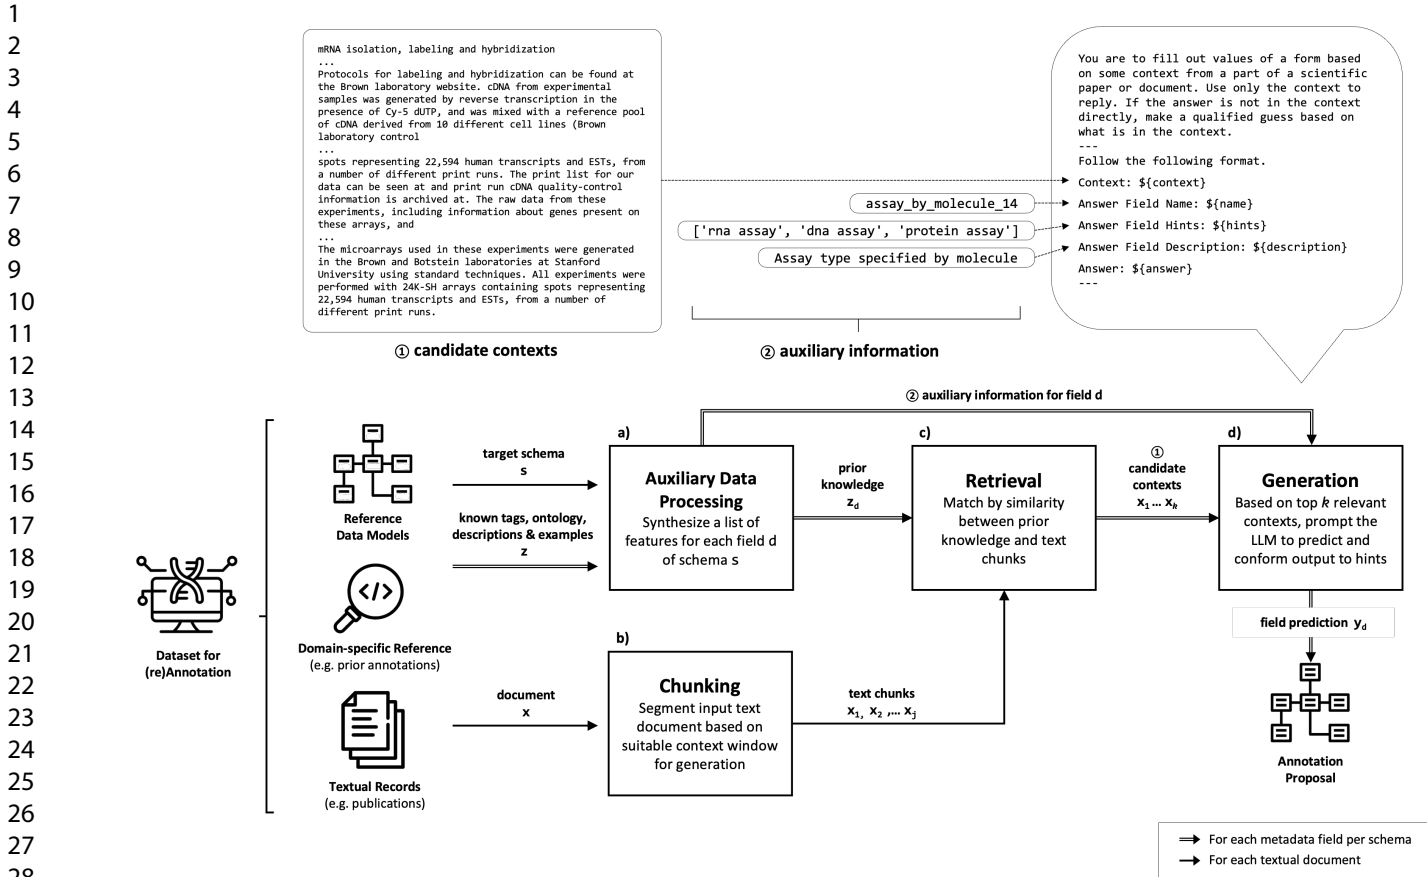
<sup>4</sup> <https://www.ebi.ac.uk/biostudies/arrayexpress>

<sup>5</sup> <https://www.ncbi.nlm.nih.gov/geo>

<sup>6</sup> <https://www.cbioportal.org>

<sup>7</sup> <https://obi-ontology.org>

<sup>8</sup> <https://www.ebi.ac.uk/efo/index.html>



**Fig. 1.** The Pre-Meta data annotation pipeline consists of four components: a) an auxiliary data processing step that takes a schema and references related to a domain (e.g. pre-generated metadata tags, ontologies, etc.) and outputs representations of the auxiliary data as useful prior knowledge to facilitate the retrieval process; b) a chunking step that segments the input document into suitable chunk sizes; c) a retrieval step that takes prior knowledge from the auxiliary data processor and text nodes from the chunking step, and outputs top  $k$  candidate contexts for prediction generation; d) a generation module taking  $k$  contexts and auxiliary information as hints for generation, and applies a user-defined LLM to predict a metadata field.

- A data model or schema  $s$  chosen among formats related to a data portal, sharing protocol, or newly defined metadata fields.
- Some auxiliary information or prior knowledge  $z_d$  related to each metadata field  $d$  in schema  $s$ , such as a description, pre-generated metadata, or similar example values e.g. from related ontologies.

The process can be represented as an RAG function  $f(\cdot)$  that predicts each metadata label  $y_d$  as an annotation to the target asset based on its textual description  $x$ :

$$\hat{y}_d = f(x, z_d), \quad \forall d \in s \quad (1)$$

## 2.2 Implementation

The Pre-Meta data annotation pipeline contains two main parts: first, the auxiliary data processing step shown in Figure 1-a; second, a simplified RAG defined as Eq. 1 consisting of three blocks of processes illustrated in Figure 1 b-c-d: chunking, retrieval, and generation. Given a dataset of interest requiring (re)annotation, the dataset's textual references are processed by a chunking step, which outputs chunks of text to the retrieval module. The retrieval module uses various representations of auxiliary information to identify  $k$  relevant contexts. The generation module, for each field of the target schema, proposes an annotation based on relevant contexts.

### 2.2.1 Auxiliary Data Processing

For the task of data annotation, given a data schema of interest, we considered three types of auxiliary information: 1) a set of pre-generated metadata tags associated with each schema field which is known to the user or accessible through related instances of data annotations on existing data portals; 2) an ontology of the associated knowledge domain, with nodes similar to the expected values of the field, and 3) example values provided in the given schema.

The processor (Figure 1-a) takes the data schema in the form of a `pydantic` object which characterizes each metadata field as a `Literal` type accompanied by a `Field` description. We also follow this convention to format pre-generated metadata tags as a `Literal` list (e.g. `Literal['rna assay', 'dna assay', 'protein assay']`) as shown in Figure 1) for use in both retrieval and generation steps. For ontologies, we selected a node and used all its descendant nodes (i.e. the nodes in the subtree/subontology with the selected node as root node). To do that, we used the `owlready2` package (Lamy 2017) to operate on each `.owl` file. As an example, for the field "assay by molecule", we used the Internationalized Resource Identifier (IRI) code `OBI:0000070` to locate the "assay" node object in the OBI ontology. Then we collated all descendant nodes for their descriptions without maintaining their tree structure. The resulting list of subontology information would then be passed onto the retriever for context ranking (Figure 1-c).

### 2.2.2 Chunking

The chunker (Figure 1-b) loads full text of papers and textual documents in .xml format using `UnstructuredXMLoader`<sup>9</sup>, which is then cleaned by removing noisy text headers and tags using `LangChain's MarkdownHeaderTextSplitter` and split into smaller chunks,  $\text{Chunks}(x) = \{x_1, x_2, \dots, x_j\}$ , using the `RecursiveCharacterTextSplitter`<sup>10</sup>. Finally, the resulting chunks are stored as `LlamaIndex's TextNode` data object (Liu 2022).

### 2.2.3 Retrieval

The retriever (Figure 1-c) embeds the document chunks and prior knowledge  $z_d$  for the field using a sentence transformer (Reimers and Gurevych 2019) `all-MiniLM-L6-v2` from Hugging Face, and compares these using cosine similarity to rank the chunks based on the relevance for the field. The retriever allows for a list of  $n$  queries,  $z_d = z_d^1, z_d^2, \dots, z_d^n$ , where each string  $z_d^i$  is embedded separately. The string  $z_d^i$  with maximum similarity to each chunk  $x_j$  is used for the ranking:

$$\text{relevance of chunk } x_j = \max_{i \leq n} (\cos.\text{sim}(\text{emb}(z_d^i), \text{emb}(x_j))) \quad (2)$$

Top  $k$  chunks with highest similarity scores are concatenated in text as one candidate context and provided to the generation module through a prompt template (the "candidate context" part as shown in Figure 1-1).

### 2.2.4 Generation

Generation (Figure 1-d) is independent of LLM models. Each target metadata field is annotated by feeding the generation model a prompt following the template in Figure 1, filled with the retrieved context and field property values (`name`, `hints` and `description`). Constrained generation is used to ensure an output format in line with the field `hints`. This works by masking the probabilities of the tokens not leading to allowed tags given as hints, before sampling the next token, thereby always ending up with a compliant output. The output can be an integer, a string with possible length constraint, or, as in our experiments, one of several listed possible strings. To do this, we use the `outlines` package (Willard and Louf 2023) for open-weight models, and OpenAI's native structured outputs for OpenAI models.

## 3 Experiments

### 3.1 Data Schema Selection

To experiment and evaluate a new dataset annotation method, a data schema with real-world metadata fields is preferred. In the experiment, metadata fields from the ArrayExpress portal were selected. ArrayExpress is one of the major public repositories for functional genomic datasets. For our experiment, we selected five fields shown in Table 1, based on field relevance and technical suitability. More specifically, the selected fields align strongly with a real-case scenario where a biologist or other biomedical researcher performs a dataset query, such as searching for a dataset concerning a particular tissue (ArrayExpress field: "organism part") or experiment design like compound treatment or time series (ArrayExpress field: "experiment design"). In addition, they are highly

descriptive and associable to existing ontologies, ensuring that they can be mapped to controlled values. Finally, selected fields from ArrayExpress are mapped to other public portals' similar schemas for dataset access (e.g., GEO), underlying their importance for dataset annotation.

Table 1 shows the selected fields of ArrayExpress for evaluation, including the counterpart of each field in GEO and the respective sample size in the test dataset. Although cBioPortal does not offer direct equivalents for several ArrayExpress metadata fields, relevant information can often be included in the description field of the `meta_study.txt` file or within custom attributes (e.g., `Tissue.Site`) in the clinical data files (`data_clinical_sample.txt` or `data_clinical_patient.txt`).

### 3.2 Data Curation

To demonstrate the case of an automated data annotation pipeline, we need textual records associated with real-world biomedical datasets with the corresponding metadata that can be used as labels. We curated a dataset  $A = (X, Y)$  with 14,844 pairs of scientific papers and metadata files. Each pair  $(x, y) \in A$  consisted of  $x$ , an open access paper in .xml format from Europe PubMed Central (Europe PMC) and  $y$ , a set of labels that describe the associated dataset of the paper. The  $y$  is curated from the dataset's metadata in .json format downloaded from ArrayExpress.

#### 3.2.1 Label Standardization

Each metadata file  $y$  crawled from ArrayExpress contained a set of fields with varying quantities and consistencies of tag values. For example, some of the fields had several hundred or even thousands of unique tag values, many of which were semantically equivalent but syntactically different (e.g., "male and female" vs. "males and females") as a result of the lack of schema enforcement. We merged and discarded the low-occurrence metadata tags, standardizing to a dataset where each field had up to 25 different possible labels.

Each paper has values for a non-empty subset of fields and may have several values per field; but for the experiments in this paper we used only fields with a single value for simplicity. The resulting sample size for each field is shown in column 4 of Table 1.

The frequencies of the 25 possible labels were highly skewed, and for 3 of the fields, more than half of the papers had the same label. This leads to the potential for distributional "cheating", where a model with a high chance of predicting these values could get a high score without using the context as intended. To combat this, a subset of 1500 pairs of paper and labels was sampled for use in the experiments, providing 300 samples for each field. The papers with the least frequent of the 25 possible labels were picked first, ensuring a uniform distribution of the labels as much as possible. The resulting label distribution for each field is shown as the blue bars in Figure 2.

### 3.3 Design of Experiment

In order to easily and effectively evaluate and compare different versions of the pipeline, we focus on the case with limited output space by using schema fields with a predefined set of possible values—essentially as a multiple choice task. The alternative of using free-text fields would require a way of measuring the similarity of the output and the correct answer, for which embedding models are not necessarily accurate in representing specific biomedical terminologies. As the retrieval process is the part of the pipeline we focus on investigating

<sup>9</sup> <https://pypi.org/project/unstructured>

<sup>10</sup> <https://python.langchain.com>

**Table 1.** A mapping of the evaluation metadata fields (Col. 1) to their respective relevance (Col. 2), their GEO analogous fields (Col. 3) and viable sample sizes of paper per field after filtering, and finally to their associative ontologies (Col. 5) that can be pruned to subsets of node descriptions (Col. 7) using specific entity terms (Col. 6).

Metadata Field	Relevance	Pre-generated Metadata		Ontologies		
		GEO Mapping	Sample Size (papers)	Ontology	Subtree Top Node	No. of Nodes
hardware	Represents sequencing or/and microarray platform used to generate experimental data	“Platforms”	1526	OBI	“DNA sequencer”	37
organism part	Represents the specific biological sample origin, such as “liver”, “blood”, or “tumor tissue”	“Source name”	4773	BTO	“tissues, cell types, and enzyme sources”	6451
experimental design	Captures the structure and purpose of the study (e.g., “time-course”, “case-control”, or “dose-response”)	“Overall design”	5647	EFO	“Study design”	64
study type	Represents a more concise description of the experimental process for the dataset generation; defines the method or approach used in the experiment (e.g., “RNA-seq”, “ChIP-seq”, “DNA methylation”)	“Experiment Type”	14797	EFO	“Assay”	556
assay by molecule	Describes the molecular target of the assay (e.g., “RNA”, “protein”, “DNA”)	“Extracted molecule”	14344	OBI	“Assay”	1656

in this work, different versions of it were tested. These versions, presented in Section 3.3.1, are based on using varying knowledge to represent the field. In the case of using ontologies as auxiliary information, we retain the allowable metadata tags as `hints` (Figure 1-2) for the generation process in order to isolate the effects of the retrieval part, even though alternative retrieval methods are meant mainly for situations where this information is not known, such as working with open-string fields.

### 3.3.1 Retrieval Methods

The following list explains the different versions of the retrieval step of the pipeline that were used in the experiments. The first two are meant as baselines for comparison, while the rest are variants of Pre-Meta:

- **Fullpaper**: providing full text to GPT-4o mini without RAG and the use of auxiliary information.
- **Vanilla RAG**: using the original field description given by the target schema for ranking text chunks that represent candidate contexts.
- **Pre-Meta**: using individually encoded known metadata tags to rank chunks of candidate contexts using Eq. 2.
- **Pre-Meta Onto**: using individually encoded descriptions from the subtree of the related ontology to rank chunks of candidate contexts. The root nodes were chosen manually to include nodes similar to a substantial part of the known metadata tags, and are presented in Col. 5-7 of Table 1.
- **Pre-Meta (Retrieval-only)**: the retrieval-only case of Pre-Meta. Since a relevance score is calculated for each allowable metadata tag, it is possible to simply pick the one with the highest score as the prediction, skipping the generation step altogether.

### 3.3.2 Generation

We evaluated different generation setups, including open-weight models (Mistral-7B-v0.1-GPTQ and Meta-Llama-3.1-8B-Instruct-GPTQ-INT4), proprietary GPT-4o mini, and one without a generation model for the retrieval-only case of Pre-Meta.

### 3.3.3 Evaluation

For each paper, a score of 1 is given if the output of field prediction is the correct label and 0 otherwise. The reported scores are averages over the 300 papers for each field, and over five fields for each pipeline configuration. For fine-grain analysis and documentation, we recorded scores both field-by-field and as summary accuracy, together with all input RAG parameters and prompts using an AI developer platform, Weights & Biases<sup>11</sup>.

## 3.4 Results

Based on experimenting with 300 sampled paper-metadata pairs for each of five chosen fields, we observe considerable advantages in incorporating domain-specific auxiliary information for the task of predicting annotations for metadata fields. The results are illustrated as a summary of the accuracy scores in Table 2 and as distributions of predictions versus labels in Figure 2. We observe that the Pre-Meta retrieval methods by far outperform conventional retrieval (**Vanilla RAG**), for pairing with both GPT-4o mini and open-weight models (Llama 8B and Mistral 7B). The accuracy gain is seen in all fields and Pre-Meta versions, with the exception of the “experimental design” field for GPT-4o mini, where conventional retrieval scored slightly

<sup>11</sup> <https://wandb.ai>



**Table 2.** A ranking of combinations of retrieval methods and generation models based on their accuracy scores of annotation predictions averaged across five evaluation fields. The rows in bold text illustrate that the combination of Pre-Meta retrieval (both variants) with GPT-4o mini as a generation model outperforms running the **Fullpaper** through GPT without RAG and running the GPT with **Vanilla RAG**. Notably, both variants of Pre-Meta also add significant accuracy gain to Llama 8B and Mistral 7B compared to using conventional RAG. The pairings of both Pre-Meta methods with open-weight Llama 8B achieve comparable result to that of GPT-4o mini running on conventional RAG. Among GPT-4o mini pairings, **Pre-Meta** contributes the highest accuracy to the fields of **hardware**, **organism part**, and **study type**, **Pre-Meta Onto** to the field of **assay by molecule**, while **Vanilla RAG** retains the highest on **experimental design**. (n=300)

Retrieval Method	Generation Model	label: <b>hardware</b>	label: <b>organism part</b>	label: <b>experimental design</b>	label: <b>assay by molecule</b>	label: <b>study type</b>	Average Accuracy
<b>Pre-Meta</b>	<b>GPT-4o mini</b>	<b>0.540</b>	<b>0.787</b>	<b>0.330</b>	<b>0.770</b>	<b>0.513</b>	<b>0.588</b>
<b>Pre-Meta Onto</b>	<b>GPT-4o mini</b>	<b>0.513</b>	<b>0.743</b>	<b>0.323</b>	<b>0.817</b>	<b>0.480</b>	<b>0.575</b>
Pre-Meta	Llama 8B	0.390	0.710	0.220	0.723	0.370	0.483
<b>Vanilla RAG</b>	<b>GPT-4o mini</b>	<b>0.323</b>	<b>0.673</b>	<b>0.343</b>	<b>0.627</b>	<b>0.367</b>	<b>0.467</b>
Pre-Meta Onto	Llama 8B	0.410	0.593	0.180	0.740	0.390	0.463
<b>Fullpaper</b>	<b>GPT-4o mini</b>	<b>0.340</b>	<b>0.730</b>	<b>0.160</b>	<b>0.683</b>	<b>0.270</b>	<b>0.437</b>
Pre-Meta	---	0.110	0.547	0.170	0.740	0.303	0.374
Pre-Meta	Mistral 7B	0.190	0.560	0.157	0.627	0.140	0.335
Pre-Meta Onto	Mistral 7B	0.167	0.413	0.093	0.643	0.123	0.288
Vanilla RAG	Llama 8B	0.170	0.483	0.123	0.490	0.140	0.281
Vanilla RAG	Mistral 7B	0.063	0.260	0.077	0.490	0.063	0.191

better. In combination with GPT-4o mini, **Pre-Meta** and **Pre-Meta Onto** add 26% and 23% performance gain compared to the baseline of **Vanilla RAG**. The gain is especially pronounced for the pairing with Llama 8B and Mistral 7B, where **Pre-Meta** adds 72% and 75% to each respective model's baseline performance. The **Pre-Meta (Retrieval-only)** setup, which involves directly choosing the best matching allowable metadata tags derived from auxiliary information without using a generation model, performs better than the baselines for the field "assay by molecule", and is better on average than Mistral 7B with **Pre-Meta**. As this is a considerably cheaper method than the others, it shows that generation models alone do not always provide a better solution, at least for the tasks explored in this experiment.

Through improving the context ranking in retrieval alone, Pre-Meta methods add considerable gain in the prediction quality to our annotation tasks for the five selected fields, demonstrating the value of using pre-generated structured data as auxiliary information.

## 4 Discussion

### 4.1 Error Analysis

#### 4.1.1 Domain-related

We note the high variance across five test fields. For example, all scores in "assay by molecule", and almost all scores in "organism part", are higher than every score in "experimental design". This variability reflects the inherent heterogeneity of biomedical data and the uneven adoption of ontologies in existing repositories. For example, nuanced fields like "experimental design" remain challenging to annotate, as they require more comprehension of domain-specific contexts.

Figure 2 highlights the distribution of predictions versus labels across five field types for three different pipeline configurations. The distributions of predictions tend to match the labels better when using Pre-Meta, although certain field labels are systematically over- or under-represented.

Some labels of the **organism part** field are not mutually exclusive, e.g., "peripheral blood" and "whole blood" could be regarded as subcategories of "blood". Thus, some errors can

be attributable to imperfections in the dataset and evaluation method, not to the models. Theoretically, a merging of semantically similar tags in the dataset or evaluation process would result in a fairer score, but this would require a systematic analysis of the metadata tags that was considered outside the scope of this work.

#### 4.1.2 Model-related

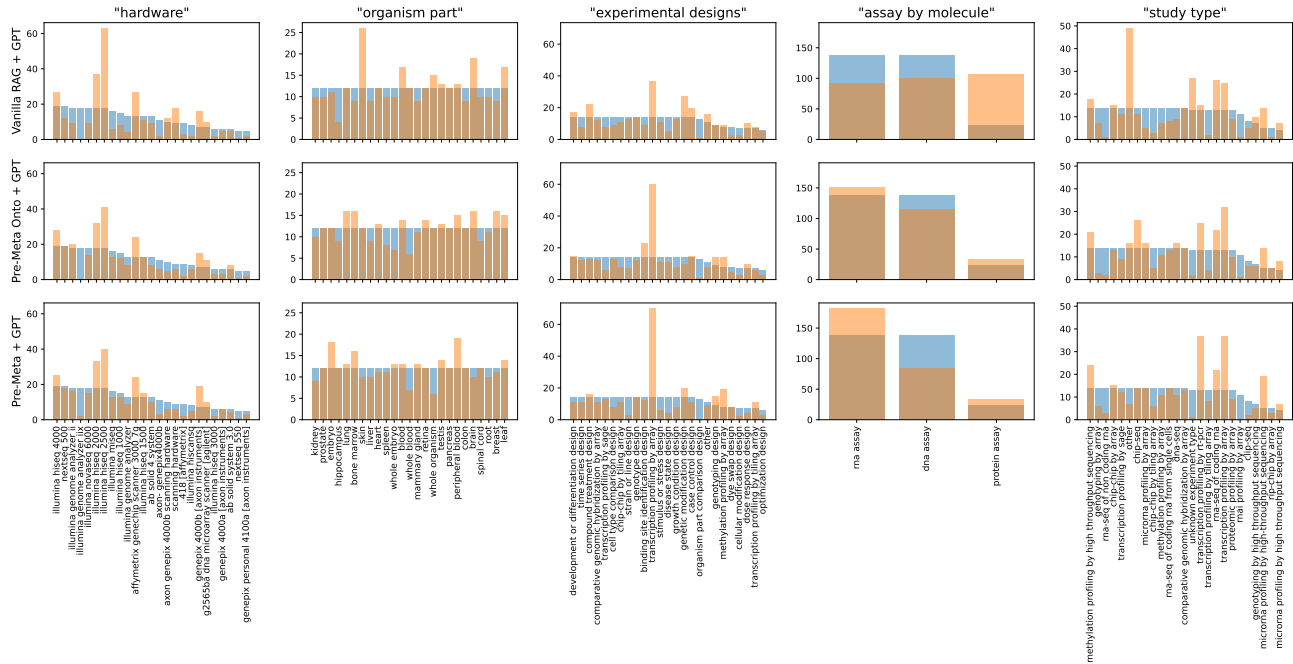
Although our work focuses mainly on retrieval, we observe that the choice of generation models affects the resulting performance more than the choice of retrieval methods. From Table 2, we see that, regardless of retrieval methods, the GPT-4o-mini model's overall scores are better than those of two open-weight models. Generation models are known to be sensitive to how the task is presented. The factor that turned out to play a considerable role in our experiment was the order of the list of **hints** given to the prompt. As a result, we shuffled the list each time to avoid systematic errors. Figure 4 shows that all generation models used in the experiment systematically focus more on answers that appear early among the list of **hints** provided to generation, even with a fairly short list of 25 answers. The bias is similar to the observed tendency that LLMs have in attending to information in the early and late parts of long context windows<sup>12</sup>. This order bias in token selection also provides us with a perspective beyond general tasks benchmarks<sup>13</sup> that reflect the difference between models (e.g. Llama 8B vs GPT-4o mini) used in the experiment.

#### 4.1.3 Retrieval by Metadata Tags vs. Ontologies


In instances where not all potential metadata tags are present, certain tags may still be available. To emulate this for **Pre-Meta**, we experimented with giving only a sampled subset of allowable metadata tags to the retriever as auxiliary information. The results of this sub-sampling are shown in Figure 3 for comparison with **Pre-Meta Onto** pairings that

<sup>12</sup> [https://github.com/gkamradt/LLMTest\\_NeedleInAHaystack](https://github.com/gkamradt/LLMTest_NeedleInAHaystack)

<sup>13</sup> <https://aimlapi.com/comparisons/llama-3-1-8b-vs-chatgpt-4o-mini>



**Fig. 2.** A comparison of the distributions of predicted metadata (orange bar) against actual annotations (blue bar) across five field types based on results from three GPT-4o mini pipelines (**Vanilla RAG**, **Pre-Meta Onto**, and **Pre-Meta**). Certain field labels, such as "illumina hiseq 2500" in hardware, "skin" in organism part, and "other" in study type show observable attenuation, while "transcription profiling by array" in experimental design shows further biased attention. (n=300)



**Fig. 3.** A comparison of the impact on annotation accuracy given different sampling densities applied to pre-generated metadata tags for ranking text chunks at the retrieval stage. For all paired models, we see comparable performance between using halved pre-generated metadata tags (12) and using ontology subtrees as auxiliary information. Notably, with full sampling at 25 metadata tags, the Llama 8B-paired accuracy (red solid) slightly exceeds that of GPT-4o mini on Vanilla RAG (blue mixed dash).

use all available subontology field descriptions. We observe that even having a single metadata example given as a hint can yield superior retrieval quality compared to using just the field description itself in the **Vanilla RAG** baseline for Llama 8B, Mistral 7B; additionally, only two metadata examples are needed as a hint to bring improvement over the baseline for GPT-4o mini. Meanwhile, for all three paired models, the use of subontologies gives better results than using half or less than half of the available metadata tags.

#### 4.1.4 Full Context Length vs Chunks

Overall, **Vanilla RAG** performs better than **Fullpaper** by a slight margin (0.467 vs 0.437), although for only two of the five metadata fields. This shows some of the reason for using retrieval in the first place—having all the information in one large context window is not necessarily better than having a selection of small chunks of it, even when this selection is suboptimal (seeing as the Pre-Meta chunk selections are better).

#### 4.2 Cost and computation

Table 3 itemizes the resource requirements across different usages of LLMs in the study. Compared to running on full paper, the use of GPT-4o mini on retrieved text chunks incurred significantly lower dollar costs due to lower token counts. For open-weight models running on a desktop GPU (Nvidia RTX 3090), the compute time of a few seconds per field also offers a promising option for deployment settings without access to non-industrial grade compute resources.

#### 4.3 Addressing FAIR principles

The experimental results demonstrate a promising step toward an LLM-driven automation of dataset (re)annotation. This effort provides the community with a useful building block to address the principle of machine-actionability (i.e., the capacity of computational systems to find, access, interoperate and reuse data with none or minimal human intervention) emphasized in FAIR, in particular the characteristics stipulated in the subprinciples F2, F4, I1, and R1.3.<sup>14</sup>

<sup>14</sup> Relevant FAIR subprinciples are: F4. (Meta)data are registered or indexed in a searchable resource; F2. Data

**Table 3.** Computational/monetary cost of experiments for each generation model. Computation times for Llama 8B and Mistral 7B includes just the generation step, not the document loading, retrieval and evaluation. Computation times are based on a single unit Nvidia RTX 3090 GPU.

Generation Model	Cost per field	Cost of 300 papers with 5 fields
GPT-4o mini with retrieved chunks	0.0002 USD	0.3 USD / 1,878,594 input tokens
GPT-4o mini with full paper	0.0027 USD	4.1 USD / 26,992,061 input tokens
Llama 8B	3.5s	88 minutes
Mistral 7B	3.9s	98 minutes
Retrieval only	2.2s	55 minutes

#### 4.4 Related Work

Efforts to increase the utility of LLMs in the biomedical domain include various Question Answer (QA) benchmarks such as BioASQ (Tsatsaronis et al. 2015), MedQA (D Jin et al. 2021), MMLU (Hendrycks et al. 2021), MedMCQA (Pal, Umapathi, and Sankarasubbu 2022). There are also specialized pre-trained models like BioMedLM (Bolton et al. 2024) along with various fine-tuned, instruction-tuned, and Reinforcement Learning-tuned models seen on Open Medical-LLM Leaderboard (Pal, Minervini, et al. 2024), and more recently, pipelines driven by tool augmentations such as GeneGPT (Q Jin et al. 2024) that incorporate in-context learning through Web APIs of NCBI. These efforts, however, have not been extended to performing short-form answer generation tasks at the precision level suitable for metadata standardization.

Beacon V2 has been introduced as open source software to facilitate the transformation of genomic metadata to conform to a hierarchical data model with seven entities: analysis, biosamples, cohorts, datasets, genomicVariations, individuals, and runs (Rueda et al. 2022b). The input, notably, is an Excel file manually completed by the end user transcribing the metadata from the original storage format (e.g. text files, CSV, database, PEF, Electronic Health Records, etc.) of the metadata. Regardless of the format, the origination of the metadata remains a user-defined process.

A potential prospect of automating the data profiling process can be seen in Extreme Multi-label Classification (XMC), a class of problem where a classifier selects relevant subsets of labels from an extremely large label set (Bhatia et al. 2016). XMC typically involves either fine-tuning a retriever (Decorte, Verlinden, et al. 2023) over the label space or training a binary classifier per class (Clavié and Soulié 2023), both of which require substantial labeled data. To reduce manual labeling, researchers use distant supervision (Decorte, Haute, et al. 2022), generate synthetic data with LLMs (Clavié and Soulié 2023; De Raedt et al. 2023), or fine-tune on related tasks with available data (Remy, Demuyne, and Demeester 2022). At inference time, performance can be boosted by reranking candidate labels using LLMs. (Zhu and Zamani 2023) avoids fine-tuning but relies on multiple LLM and retrieval calls per input to generate synthetic prompts and rerank labels. The Infer-Retrieve-Rank method introduced by (D'Oosterlinck et al. 2024) defined a multi-step, in-context interactions between LLMs and the retriever to predict relevant queries and interpret retrieved documents for better ranking. However, this requires dedicated prompt optimization for different

datasets. Furthermore, the method focuses on the single task of predicting adverse drug reactions from an input document, which does not reflect the heterogeneity of metadata inherent in the rich description of genomic datasets. Pre-Meta, on the other hand, focuses on the setting of labeling diverse metadata fields through a complementary, light-weight procedure of leveraging available priors for candidate context ranking. Its contribution is also independent of performance gains that can be derived from fine-tuning, prompt optimization, and multiplying the use of LLM calls.

#### 4.5 Limitations

Our current implementation focused on metadata that has string/literal values in order to leverage the LLM's capabilities on natural languages. For other types of metadata, for example, metadata with numerical values and alpha-numeric strings, additional experiments and iteration of the pipeline will be required, given the incoherent nature of the way numbers and decimals are tokenized by language models (Marjeh et al. 2025).

Due to the lack of existing reference or benchmark datasets for evaluating the tasks of biomedical data annotation, our choices of the five metadata fields were naturally constrained by the quality and availability of the metadata labels found on ArrayExpress. To further systematize the evaluation of Pre-Meta on a wider set and diversity of metadata fields that may appear under a different or newly introduced schema, additional data sources beyond ArrayExpress will need to be identified first and evaluated. Similarly, this applies to potential adoption of Pre-Meta to other domains. As part of further research, we aim to facilitate the automatic identification of relevant priors, i.e. from ontologies, that relate to specific new fields appearing through a new schema.


From the deployment perspective, the curation of data labels from pre-generated metadata tags in ArrayExpress required extensive cleaning due to the lack of schema enforcement. As a result, the raw state of metadata tags found in platforms like ArrayExpress would not be considered ready-to-go auxiliary information that can effectively support the implementation of Pre-Meta on the platform. But this gap can be addressed, portal by portal, through tailoring the auxiliary information preprocessor to reduce the data noise accordingly.

#### Conclusion

This study presents Pre-Meta, an LLM-agnostic and domain-independent data annotation pipeline designed to facilitate the automation of data profiling and downstream discovery. By leveraging auxiliary data—such as pre-generated metadata tags and ontologies—for retrieval to support the pairing with LLMs, Pre-Meta addresses key challenges in metadata standardization

are described with rich metadata; II. (Meta)data use a formal, accessible, shared, and broadly applicable language for knowledge representation; and R1.3. (Meta)data meet domain-relevant community standards. <https://www.go-fair.org/fair-principles>





**Fig. 4.** The graphs show how often an answer was generated, based on its position in the list of possible answers given in the generation prompt, across the fields with 25 possible answers. The list was shuffled at each LLM call. This shows that LLM models, especially Mistral 7B, have a systematic bias towards choosing an answer appearing early in the list. (n=10800)

and dataset interoperability, as seen across repositories like GEO, ArrayExpress, and cBioPortal in the biomedical domain.

The scalability offered by LLMs in processing large datasets across repositories is both an opportunity and a challenge. While it enables the rapid annotation of diverse datasets, the reliance on unstructured or inconsistent metadata introduces variability in quality and interpretation (Higashi et al. 2024). Pre-Meta addresses this by promoting harmonization through auxiliary information and ontology-driven annotation. However, achieving full interoperability will require a broader community effort to establish harmonized metadata frameworks and robust cross-repository standards.

This study also emphasizes the need for benchmarkability in the evaluation of LLMs for data annotation in the biomedical domain. Transparent benchmarks and robust metrics are critical for assessing model performance across domain-relevant tasks, ensuring both accuracy and generalizability. Pre-Meta contributes to this goal by providing a scalable and flexible approach to metadata annotation, which lays the groundwork for improving data sharing and discovery in biomedical research.

### Code and Data Availability

The code, data access, and pre-processing scripts are available in the GitHub repository: <https://github.com/SINTEF-SE/LLMDap>.

### Acknowledgments

The work was supported through the projects UPGAST (HE 101093216), enRichMyData (HE 101070284), DataPACT (HE 101189771), and CauseFinder (PNRR 760049).

### References

Ayoub, Noel F. et al. (2024) “Inherent Bias in Large Language Models: A Random Sampling Analysis”. In: *Mayo Clinic Proceedings: Digital Health* 2.2, pp. 186–191. ISSN: 2949-7612. DOI: 10.1016/j.mcpdig.2024.03.003.

Barrett, Tanya et al. (Nov. 2006) “NCBI GEO: mining tens of millions of expression profiles—database and tools update”. In: *Nucleic Acids Research* 35.suppl1, pp. D760–D765. DOI: 10.1093/nar/gkl887.

Bhatia, K. et al. (2016) *The extreme classification repository: Multi-label datasets and code*. URL: <http://manikvarma.org/downloads/XC/XMLRepository.html>.

Bodenreider, Oliver, Ronald Cornet, and Daniel J Vreeman (2018) “Recent developments in clinical terminologies — SNOMED CT, LOINC, and RxNorm”. In: *Yearbook of medical informatics* 27.01, pp. 129–139.

Bolton, Elliot et al. (2024) “Biomedlm: A 2.7 b parameter language model trained on biomedical text”. In: *arXiv preprint arXiv:2403.18421*.

Brazma, Alvis et al. (2001) “Minimum information about a microarray experiment (MIAME)—toward standards for microarray data”. In: *Nature genetics* 29.4, pp. 365–371.

Byrd, James Brian et al. (2020) “Responsible, practical genomic data sharing that accelerates research”. In: *Nature Reviews Genetics* 21.10, pp. 615–629.

Clavié, Benjamin and Guillaume Soulié (2023) *Large Language Models as Batteries-Included Zero-Shot ESCO Skills Matchers*. arXiv: 2307.03539 [cs.CL] URL: <https://arxiv.org/abs/2307.03539>.

D’Oosterlinck, Karel et al. (2024) *In-Context Learning for Extreme Multi-Label Classification*. arXiv: 2401.12178 [cs.CL] URL: <https://arxiv.org/abs/2401.12178>.

De Raedt, Maarten et al. (2023) “IDAS: Intent discovery with abstractive summarization”. In: *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*. Toronto, Canada: Association for Computational Linguistics, pp. 71–88.

Decorte, Jens-Joris, Jeroen Van Haute, et al. (2022) *Design of Negative Sampling Strategies for Distantly Supervised Skill Extraction*. arXiv: 2209.05987 [cs.CL] URL: <https://arxiv.org/abs/2209.05987>.

Decorte, Jens-Joris, Severine Verlinden, et al. (2023) “Extreme multi-label skill extraction training using large language models”. In: *AI4HR & PES, the International Workshop on AI for Human Resources and Public Employment Services, ECML-PKDD 2023 Workshop*, pp. 1–10.

Hendrycks, Dan et al. (2021) *Measuring Massive Multitask Language Understanding*. arXiv: 2009.03300 [cs.CY]

Higashi, Koichi et al. (2024) “Automated Harmonization and Large-Scale Integration of Heterogeneous Biomedical Sample Metadata Using Large Language Models”. In: *bioRxiv*, pp. 2024–10.

Jin, Di et al. (2021) “What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams”. In: *Applied Sciences* 11.14. ISSN: 2076-3417. DOI: 10.3390/app11146421.

Jin, Qiao et al. (2024) “Genegpt: Augmenting large language models with domain tools for improved access to biomedical information”. In: *Bioinformatics* 40.2, btac075.

Lamy, Jean-Baptiste (2017) “Owlready: Ontology-oriented programming in Python with automatic classification and high level constructs for biomedical ontologies”. In: *Artificial intelligence in medicine* 80, pp. 11–28. DOI: 10.1016/j.artmed.2017.07.002.

Lewis, Patrick et al. (2021) *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. arXiv: 2005.11401 [cs.CL] URL: <https://arxiv.org/abs/2005.11401>.

Liu, Jerry (Nov. 2022) *LlamaIndex*. DOI: 10.5281/zenodo.1234.

Marjeh, Raja et al. (2025) *What is a Number, That a Large Language Model May Know It?* arXiv: 2502.01540 [cs.CL]

Nazi, Zabir Al and Wei Peng (2024) “Large Language Models in Healthcare and Medical Domain: A Review”. In: *Informatics* 11.3. ISSN: 2227-9709. DOI: 10.3390/informatics11030057.

- 1
  - 2 Pal, Ankit, Pasquale Minervini, et al. (2024) *openlifescienceai*
  - 3 */open\_medical\_llm\_leaderboard*. [https://huggingface.co/spaces/openlifescienceai/open\\_medical\\_llm\\_leaderboard](https://huggingface.co/spaces/openlifescienceai/open_medical_llm_leaderboard).
  - 4
  - 5 Pal, Ankit, Logesh Kumar Umapathi, and Malaikannan
  - 6 Sankarasubbu (2022) *MedMCQA : A Large-scale*
  - 7 *Multi-Subject Multi-Choice Dataset for Medical domain*
  - 8 *Question Answering*. arXiv: 2203.14371 [cs.CL].
  - 9 Reimers, Nils and Iryna Gurevych (Nov. 2019) “Sentence-
  - 10 BERT: Sentence Embeddings using Siamese BERT-
  - 11 Networks”. In: *Proceedings of the 2019 Conference on*
  - 12 *Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. URL: <https://arxiv.org/abs/1908.10084>.
  - 13 Remy, François, Kris Demuynck, and Thomas Demeester (2022)
  - 14 *BioLORD: Learning Ontological Representations from*
  - 15 *Definitions (for Biomedical Concepts and their Textual*
  - 16 *Descriptions)*. arXiv: 2210.11892 [cs.CL]. URL: <https://arxiv.org/abs/2210.11892>.
  - 17 Rueda, Manuel et al. (Aug. 2022a) “Beacon v2 Reference
  - 18 Implementation: a toolkit to enable federated sharing
  - 19 of genomic and phenotypic data”. In: *Bioinformatics*
  - 20 38.19, pp. 4656–4657. ISSN: 1367-4803. DOI: 10.1093/bioinformatics/btac568.
  - 21 — (2022b) “Beacon v2 Reference Implementation: a toolkit to
  - 22 enable federated sharing of genomic and phenotypic data”. In: *Bioinformatics* 38.19, pp. 4656–4657.
  - 23
  - 24
  - 25
  - 26
  - 27
  - 28
  - 29
  - 30
  - 31
  - 32
  - 33
  - 34
  - 35
  - 36
  - 37
  - 38
  - 39
  - 40
  - 41
  - 42
  - 43
  - 44
  - 45
  - 46
  - 47
  - 48
  - 49
  - 50
  - 51
  - 52
  - 53
  - 54
  - 55
  - 56
  - 57
  - 58
  - 59
  - 60
- Thapa, Surendrabikram and Surabhi Adhikari (2023) “ChatGPT, bard, and large language models for biomedical research: opportunities and pitfalls”. In: *Annals of biomedical engineering* 51.12, pp. 2647–2651.
- Tsatsaronis, George et al. (2015) “An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition”. In: *BMC bioinformatics* 16, pp. 1–28.
- Vaswani, Ashish et al. (2023) *Attention Is All You Need*. arXiv: 1706.03762 [cs.CL].
- Willard, Brandon T and Rémi Louf (2023) “Efficient Guided Generation for LLMs”. In: *arXiv preprint arXiv:2307.09702*.
- Yang, Rui et al. (2025) “Retrieval-augmented generation for generative artificial intelligence in health care”. In: *npj Health Systems* 2.1, pp. 1–5. DOI: 10.1038/s44401-024-00004-1.
- Zhu, Yaxin and Hamed Zamani (2023) “ICXML: An In-Context Learning Framework for Zero-Shot Extreme Multi-Label Classification”. In: *arXiv preprint arXiv:2311.09649*. DOI: 10.48550/arXiv.2311.09649. URL: <https://arxiv.org/abs/2311.09649>.