

Making Metadata More FAIR Using Large Language Models

Sowmya S. Sundaram^[0000-0002-0086-7582] and Mark A. Musen^[0000-0003-3325-793X]

Stanford University

Abstract. With the global increase in experimental data artifacts, harnessing them in a unified fashion leads to a major stumbling block - bad metadata. To bridge this gap, this work presents a Natural Language Processing (NLP) informed application, called FAIRMetaText, that compares metadata. Specifically, FAIRMetaText analyzes the natural language descriptions of metadata and provides a mathematical similarity measure between two terms. This measure can then be utilized for analyzing varied metadata, by suggesting terms for compliance or grouping similar terms for identification of replaceable terms. The efficacy of the algorithm is presented qualitatively and quantitatively on publicly available research artifacts and demonstrates large gains across metadata related tasks through an in-depth study of a wide variety of Large Language Models (LLMs). This software can drastically reduce the human effort in sifting through various natural language metadata while employing several experimental datasets on the same topic.

Keywords: Metadata · NLP · FAIR data

1 Introduction

Metadata are the description of data that scientists, in turn, use for analysis and other complex downstream tasks. However, if the metadata are not of good quality, it can hamper retrieval, querying, analysing and harnessing the publicly available digital research artifacts. This problem has been often cited in the literature ([18, 4]) and was the basis of the keynote talk in a previous set of proceedings of this workshop [10]. This formidable bottleneck cripples the use of digital artifacts (including experimental datasets, scientific data descriptions, ontologies etc.) in an interoperable manner. Metadata often have representational heterogeneity - that is different terms refer to the same value, which often do not conform to specifications. The datasets are simply not ‘AI-ready’ if the metadata are not good.

In a bid to enhance the sharing of scientific data, the FAIR (Findable, Accessible, Interoperable and Reproducible) principles [26] were introduced. The purpose was to enable reuse and verify existing scientific data. Though such principles are exemplary for *future* scientific data artifacts, the mere presence

cell subtype	Fresh atypical memory B cells		
cell type	Primary cell		
disease	SLE	disease name	1.脑淀粉样血管病
disease stage	New-onset	Hereditary way	1.AD
ethnicity	Asian		...
health state	SLE	altitude	C
karyotype	46 chromosomes	Chr	chr20
population	Peripheral blood	Start	23618395
race	yellow race	End	23618395
sample type	leukocyte		
treatment	No treatment		
IndividualID	sle1		

isolate	missing'	sample name	MIBC-Pat_4
age	missing'	collected by	Jena University Hospital
biomaterial provider	Ying Hsiu Su, Blumberg Institute	collection date	2010-01
sex	female	sample type	MIBC
tissue	Liver	sex	w
disease	HCC	ENA first public	2021-01-06
		ENA last update	2020-11-13

Fig. 1. Real World Metadata - The red boxes bring attention to messy metadata in real records denoted by black boxes.

of such guidelines does little to unify *existing* artifacts across seemingly diverse but semantically similar resources

There is a long reported difficulty in finding datasets pertaining to a specific format or topic [18]. Our group has been consistently working towards making research artifacts FAIR, especially in the arena of processing metadata. First, an in-depth empirical analysis was done that conclusively determined that the quality of metadata is poor in digital research artefacts [8]. In this process, a robust software tool was developed - a FAIR Workbench ([9], [19]) that can recommend cleaning up a research artefact before submitting it.

FAIRMetaText was designed with principles deriving from an earlier research foray by our lab to incorporate NLP for automating metadata clean-up. The algorithm [8] used the NLP technology available then and performed metadata analysis in a semi-automated manner. The method handled spelling mistakes with difficulty and circumvented it by a syntactic similarity function. With the advent of Large Language Models (LLMs), NLP practitioners are able to make better strides and provide a completely automated approach, across many domains. Hence, in this work, we endeavor to upgrade our previous algorithm by (a) using improved NLP models and (b) reducing human intervention. In this work, an attempt to *automatically* address some of the metadata issues of web artifacts has been made with the use of the state-of-the-art Natural Language Processing (NLP) tools for the *existing* metadata.

Experimental metadata are made of *reporting guidelines* that specify the name of metadata and *ontologies* that provide possible values. For example, while observing diseased samples, sex of the patient has been referred to variously as ‘F’, ‘female’ and ‘w’. Similarly, the disease is sometimes altogether missing, mentioned in abbreviations such as ‘SLE’ or presented in multi-lingual terms. In the previous examples, the field names are ‘sex’ and ‘disease’ and the field values have been described. As seen in Figure 1, ‘disease’ is variously described as ‘disease’ and ‘disease name’. FAIRMetaText can handle both field *names* and *values*.

Since the previous version of this algorithm differed significantly in its mechanism and was semi-automated, it has not been included in the empirical results. For the evaluation, both syntactic (spelling and phrase) and semantic similarities (meaning and context) of metadata are explored.

FAIRMetaText converts metadata into vectors based on the available large language models. These vectors can then be used in a wide variety of application. For our empirical study, there will be an exploration of two methods - retrieval and clustering. The experiments demonstrate that retrieval is useful for metadata compliance and that clustering is useful for metadata disambiguation through experiments. For this, both real and simulated datasets have been used. Through this tool, the authors propose that precious human effort applied to clean up metadata can be directed towards more useful endeavors.

2 Related Work

Metadata issues had been addressed time and again in the past ([18, 21, 2, 3]). The FAIR [26] principles suggest that metadata must be Findable, Accessible, Interoperable and Reusable. This work addresses a way to use LLMs [7] to make data FAIR.

There has been a lot of work in converting Knowledge Graphs, in their entirety, into vectorial embeddings for further processing [25]. This approach helps in question answering and other related tasks for structured data. This method is only useful where the metadata are specified perfectly. In this work, the focus is on developing embeddings for the NLP descriptions, rather than RDF triples or any other form of structural definitions. However, the presented algorithm is simply a text-based algorithm that can be applied across a wide variety of web artifacts.

Metadata alignment where manual mappings such as SSSOM [16] are used to bridge the gap between web resources, is a popular approach in the community. However, our work is designed to push the limits of automatically cleaning up metadata using state of the art NLP techniques.

Multiple research directions have been undertaken by our lab to enhance metadata authoring and recommendations ([6, 5, 15]). These methods harness advancements in machine learning and semantic technologies. To ensure the quality of metadata being submitted, [19], the FAIRware workbench highlights errors in submitted data with respect to the metadata specifications. The work-

bench ascertains and tests the quality of new scientific research being submitted. FAIRMetaText is designed with the *existing* metadata in mind by exploiting LLMs. There has been less work in this space of using NLP. As mentioned before, the earlier version [8] mapped possible candidates syntactically through edit distance and clustered semantic inputs using primitive word embeddings such as word2vec [17] and GloVe [22]. The method also involved a human in the loop. FAIRMetaText combines all the automated steps using the rich embeddings provided by the LLMs vis-a-vis the static word embeddings explored earlier. Hence, our current method is more robust, powerful and completely automated.

3 Large Language Models

Text is converted into *embeddings*, that is a piece of text is embedded in a vector space. Early naïve methods of embedding text into vector spaces included converting the word into a *one-hot vector*. Then, static word embeddings were in vogue where a single word is converted into a vector using a lookup [17]. This model fails when an unknown word is presented or when contexts change the meanings of words. The predecessor of this work [8], relied on a mix of syntactic analysis followed by the use of such static embeddings. Now, LLMs use complex text embeddings that incorporate context and can handle unknown words. The evolution of embeddings along with their corresponding visualizations is presented in Table 1.

Large Language Models (LLMs) have taken up the world of NLP by storm [7]. LLMs are trained over huge swathes of data (GPT3 uses a dataset of nearly 380 terabytes of data! [1]). They are instances of *self-supervised learning* where the data in itself is enough to train the systems without costly annotation. Leveraging the freely available text online, large language models identify patterns in text and have been proven to be successful in generating text. Under the hood of this generation process is the presence of word embeddings, or rather text embeddings, that are learned as a by-product. Text embeddings convert a piece of text to a vector such that *similar* pieces of text have vectors that are close to each other in distance. These models can be used on unseen data without further training for a particular application, which is known as zero-shot learning. Thus, it is a potent tool, in low resource data or data where annotation is hard, that has been harnessed for this application. Popular generic LLMs such as BERT, RoBERTa, CharBERT and PhraseBert have been used along with biologically specific LLMs such a BioMedLM and BioBERT have been tested. The last LLM used is GPT3.5. In this work, the embeddings¹ provided by OpenAI are plugged in directly, as they are more useful for our applications.

4 Applications of Embeddings to Metadata


The following properties of embeddings have been exploited to process metadata (a) **Structural Similarity:** Text embeddings that have similar structure (such

¹ <https://platform.openai.com/docs/guides/embeddings/what-are-embeddings>


Table 1. Evolution of Word Embeddings - The figures are word embeddings representation. The axes are of large vector dimensions projected on to two dimensions, hence the dimensions hold no meaning and are unlabelled

Consider a dictionary of five words - ‘the’, ‘a’, ‘Greece’, ‘metadata’, ‘data’ and the corresponding vectors v_{the} , v_a , v_{Greece} , $v_{metadata}$, v_{data} .


One-hot Embeddings: Text is converted into unit vectors that are the length of the vocabulary. This leads to sparse vectors that are equidistant. All words are equidistant and hence does not capture semantics of similar words.




Frequency-based Embeddings: Instead of one-hot embeddings, the weight of the word is offset by their frequency in their document and the general frequency across documents. This captures some aspects of importance of words in a document but still does not capture semantics completely. The vectors are still high-dimensional and sparse.



Distributed Embeddings: Using neural networks and large corpus of data, in this method of building distributed embeddings, the context of words are used to learn dense vectorial representations. This method produces relatively denser embeddings of length 300. These embeddings can capture semantic relationships. In our example, the words ‘data’ and ‘metadata’ are closer. However, these embeddings cannot handle out-of-vocabulary words.



Transformer Embeddings: Transformer based models obtain denser representations based on larger models and more data. These embeddings can handle out-of-vocabulary words as the vector is constructed from representations of sub-words rather than words.



as spelling or character) are closer together in the vector space; and (b) **Semantic Similarity**: Text snippets that have similar meaning are closer together in vector space.

4.1 Retrieval for Metadata Compliance

Given a metadata term w and a permissible list of ontological terms $\mathcal{D} = \{d_1, d_2, \dots, d_k\}$, if $w \notin \mathcal{D}$, then w is said to be non-compliant. For compliance, the task is to retrieve $d \in \mathcal{D}$ such that the cosine similarity between the vector forms generated by model m to obtain $vec(d)$ and $vec(w)$, is maximum $\forall d \in \mathcal{D}$. For evaluation, a check against a specified ground truth of which ontological term $d_{true} \in \mathcal{D}$ matches is performed. Then, accuracy is calculated based on the number of matches across n data points $\sum d = d_{true}/n$.

4.2 Clustering for Metadata Unification

Given a set of metadata terms $\{t_1, t_2, \dots, t_n\} \in \mathcal{T}$, it is converted into a set of vectors $\{v_1, v_2, \dots, v_n\} \in \mathcal{V}$. Then, the popular clustering algorithm k-means [14] is employed to group them into clusters \mathcal{C} . For evaluation, the purity metric is used (by labeling each cluster with a majority vote) to measure against a set of cluster labels that were not used as part of the algorithm. Assume that the true cluster assignments $\mathcal{C}_{\lfloor \nabla \rfloor}$ are known beforehand for these terms. For each cluster $c \in \mathcal{C}$, determine the majority label of the data points according to $\mathcal{C}_{\lfloor \nabla \rfloor}$. Then, measure the fraction of data-points that adhere to this labelling. The experiments also involve qualitative examination of the clusters by projecting the vectors onto a two dimensional space using visualisation algorithms like t-SNE [13].

5 Methods

In this section, the methods of evaluating two tasks are presented - (i) Metadata compliance and (ii) Metadata Unification. The datasets and evaluation metrics are described along with the performance of FAIRMetaText on these tasks.

5.1 Retrieval for Metadata Compliance

For this task, three datasets have been chosen. To demonstrate the generalisability of our approach, two web artifacts - the Adult and Mushroom datasets are explored. These datasets were chosen from the UC Irvine Machine Learning Repository [4] as they had a large number of text categorical variables and can be easily extrapolated to *any* digital experimental data. The third dataset is the authors' own in-house expert-curated dataset on tissue samples. The dataset values are derived from HuBMAP specifications [23] and expert introduced errors. These datasets have a large number of categorical values. Also, these experiments involve only the dataset description in text and not the actual samples.

Algorithm 1 Metadata Analysis

```

1: procedure GETVECTORS( $d, D, m$ )
2:    $\mathcal{V} \leftarrow \emptyset$ 
3:   for  $d \in \mathcal{D}$  do                                      $\triangleright$  For every word  $d$  in dataset  $\mathcal{D}$ 
4:      $v_d \leftarrow m(d)$                                 $\triangleright$  Get vector generated by  $m$  for  $d$ 
5:      $\mathcal{V} \leftarrow \mathcal{V} \cup v_d$ 
6:   return  $\mathcal{V}$ 
7: procedure METADATACOMPLIANCE(word  $w$ , accepted terms  $\mathcal{D}$ , model  $m$ )
8:    $sim \leftarrow 0$ 
9:    $\mathcal{V} \leftarrow$  GETVECTORS( $d, D, m$ )
10:  for  $v$  in  $\mathcal{V}$  do
11:     $sim_w \leftarrow$  COSINESIMILARITY( $w, v$ )
12:     $sim_d \leftarrow \emptyset$ 
13:    if then  $sim_w > sim$ 
14:       $sim \leftarrow sim_w$ 
15:       $sim_d \leftarrow d$                                 $\triangleright$  Storing the word  $d$  corresponding to  $v$ 
16:  return  $sim_d$ 
17: procedure METADATOUNIFICATION(metadata terms  $\mathcal{T}$ , model  $m, k$ )
18:   $\mathcal{V} \leftarrow$  GETVECTORS( $d, D, m$ )
19:   $clusters \leftarrow$  KMEANS( $V, k$ )
20:  return  $clusters_d$                                     $\triangleright$  Clusters of words corresponding to  $clusters$ 

```

With the Adult and Mushroom datasets, the dataset specifications are perturbed by single character substitutions using a publicly available NLP augementer [12]. The NLP augementer is a piece of software that introduces errors into words based on user requirement. Hence, this process generated the two new datasets after adding single character errors. These simulated datasets are used for evaluation. A single character perturbation can be easily rectified using Levenshtein distance [11]. However, with no training, this experiment demonstrates the efficacy of this method on many other facets through varied experiments. This experiment, as demonstrated in Table 2, measures the ability of the algorithm to identify syntactic similarities.

With the Tissue Sample dataset, users give a lot of semantic equivalences which are syntactically different as the wrong input. This phenomenon has been visualised in Figure 2. The input was two sets of both user input and metadata specification. The vectors are of size 784 and to visualise them, they were projected onto two axes using t-SNE [13]. The clustering algorithm correctly identified the diversity in text. In the figure, in the left bottom corner, the metadata specification ‘OCT embedded’ is closest to the user term optimal cutting temperature. These semantic details are hard to capture in traditional metadata tools.

The accuracy results for two settings - with and without definition for the expert curated dataset is presented in Table 3. Without definition, the specification terms \mathcal{D} are simply converted into vectors and compared against the query term w . In the other setting, the specification terms are augmented by con-

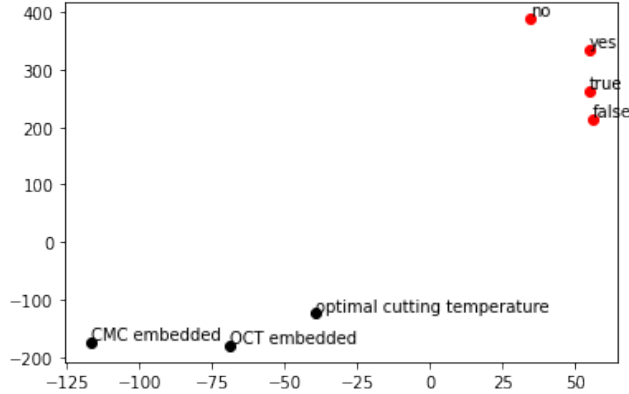


Fig. 2. Visualising Semantic Similarity of Embeddings: The vectorial representations of seven terms have been projected onto two dimensions and presented here. Given two types of terms - domain specific terms and a set of Boolean terms, the clustering algorithm correctly identifies the semantic categories. Furthermore, the full form ‘optimal cutting temperature’ is close to the phrase ‘OCT embedded’.

catenating their definitions as well. Hence, \mathcal{D} will become $\{d_1 + defn(d_1), d_2 + defn(d_2), \dots, d_k + defn(d_k)\}$. The NLP models will output a vector for these new chunks of text of size 784, that can then be used for the accuracy analysis reported in Table 3. The number of samples was small (62), again a highlight of the difficulty of low-resource datasets.

Table 2. Metadata Compliance Suggestion on Simulated Data: Accuracy of retrieval in percentage, dataset size in Parenthesis

Model	Adult	Mushroom
	[808]	[2222]
BERT	46.71	45.19
ROBERTA	56.14	68.09
CharBERT	43.00	47.14
PhraseBERT	47.71	51.09
GPT	87.78	86.95

Table 3. Metadata Compliance on Real Data: Accuracy of retrieval in percentage - with and without the use of definition of metadata

Model	No Defn	With Defn
BERT	58.70	45.65
ROBERTA	52.17	45.65
CharBERT	52.17	41.30
PhraseBERT	54.35	54.34
BioBERT	58.69	56.52
BioMedLM	47.82	45.65
GPT	63.04	67.39

Among all the LLMs, the GPT embeddings outperform with a large margin. The improved results with definitions suggests that with good representation of meanings of categorical specifications, text algorithms can do better. This observation is a pertinent point for the collection of metadata in future.

5.2 Clustering for MetaData Unification

For these experiments, the sources include the BioSample [24] resource which is a database maintained by the U.S. National Center for Biotechnology Information (NCBI) that provides descriptions and metadata for biological samples used in research. In BioSample, there is a list of attributes with their synonyms that can be used for simulating the clustering task. As already mentioned, clustering the vectors will lead us to discover groups of similarly meaning metadata terms. Purity will be used to measure the clustering. Since the knowledge of synonyms is known beforehand, one can evaluate quality of the clustering algorithm that had been presented with unlabelled metadata terms. In Table 4, the experiments have been presented by the outcome of the k-means clustering algorithm on BioSample synonym dataset at $k = 100, 200$ and 500 on a set of 1500 terms.

Table 4. Metadata Unification: Purity of clustering metadata terms and their synonyms using k means for varying k

Model	Purity (k=100)	Purity (k=200)	Purity (k=500)
BERT	59.36	59.40	57.18
ROBERTA	58.83	58.70	57.90
CharBERT	61.01	61.54	60.82
PhraseBERT	64.99	64.46	61.48
BioBERT	58.90	59.30	58.30
BioMedLM	59.03	59.03	58.10
GPT	79.48	75.51	72.99

Qualitatively, the NCBI’s GEO (U.S. National Center for Biotechnology Information’s Gene Expression Omnibus) database has been used to examine clusters more closely. GEO is a public database that provides access to a large collection of gene expression data. The GEO database includes data from a variety of platforms, such as micro-arrays, next-generation sequencing, and gene expression profiling assays. The text descriptions in the ‘characteristics’ column of all tuberculosis samples were extracted and clustered. This process was done using the R package GEOMetaDB [27]. With a similar approach of projection onto two dimensions using t-SNE, the findings are presented in Figure 3. As can be judged from the images, similar terms lie close to each other in the vector space. Given below in Figure 4 is another experiment where all ‘age’ related terms were clustered. These experiments describe how text information within metadata can be harnessed for practical applications of compliance and unification.

6 Discussion

Metadata are the often underappreciated aspects of experimental datasets. This approach is a roadblock to efficient usage of scientific artifacts. While one can postulate theories to improve the collection of such web artifacts in future, there is a pressing need to grapple with the existing huge mass of poor quality of metadata [21].




Fig. 3. Clustering of Personal vs Diagnostic Entities - Two dimensional projection of embedding vectors of metadata fields and values. Red terms correctly clustered diagnostic metadata and purple terms clustered patient metadata. The axes are a projection of large vector dimensions on to a 2D space and hence are unlabelled.

This work described FAIRMetaText which can help in this requirement. FAIRMetaText is an NLP based tool that can analyze the textual descriptions of metadata and help immensely in cleaning up the low quality data. Scientists spend a lot of time trying to look up relevant pieces of information on datasets that are simply hard to access. FAIRMetaText can cut their searching time drastically. Since LLMs encode complex semantic information, they can be leveraged to give text embeddings that can model metadata with no user/domain specific training. These embeddings were shown to capture both syntactic and semantic similarity of metadata. This study was done with an empirical analysis with different LLMs. Surprisingly, the generic GPT model was found to perform better than models trained especially for bio-medicine in the case of medical metadata. One possible surmise for this is that this is because of the much larger data these generic models have been trained on, that can account for vagaries in form, spelling and semantics.

A natural question at this juncture is whether doing away with FAIRMetaText and directly employing ChatGPT can help. ChatGPT has a strict text-to-text interface (Fig 5). FAIRMetaText goes under the hood and employs embeddings in a way suitable for metadata analysis. For example, clustering over 20000 terms would be tedious. Hence, FAIRMetaText builds on top of GPT embeddings and is tailored for metadata clean-up. Using FAIRMetaText, (a) one can monitor the quality of data and metadata of research artefacts being submitted online and (b) analyze and compare existing metadata to make artefacts inter-operable. This tool can be used on *any* metadata term - both field names and field values. The prototype will be released soon. The requirement for FAIRMetaText to enforce metadata compliance is a set of *accessible* metadata specifications that can be used to model the metadata terms. For metadata unification, the text descriptions of metadata must be machine accessible - easily

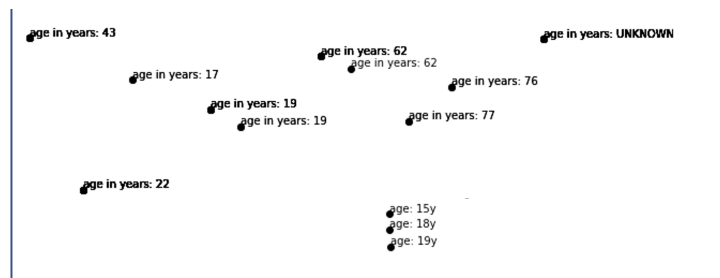


Fig. 4. Clustering of ‘Age’ Related Terms - Two dimensional projection of embedding vectors of metadata fields and values relating to ‘age’. Different configurations of the term ‘age’ are clustered into the same bin - whether they are ‘age = [X]y’ or ‘age in years: [X]’. The axes are a projection of large vector dimensions on to a 2D space and hence are unlabelled.

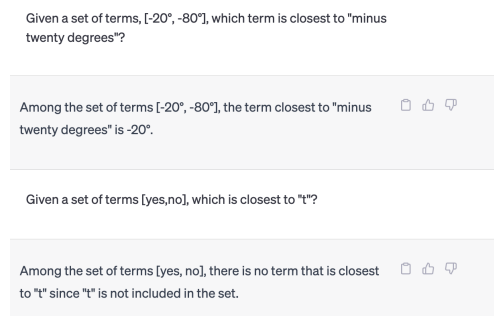


Fig. 5. ChatGPT and Metadata Compliance

downloadable and easy to parse. In future iterations of this product, FAIRMeta-Text will involve parsing various formats such as ‘json’ and ‘xml’ and ‘text’ files for generalised processing.

When this model was tested on the real data requirements, the GPT model, though best, could only provide a retrieval accuracy of 60 percent. This suggests that further training or *finetuning* of these LLMs for the purpose of metadata analysis could potentially improve the performance. User data and information of metadata non-compliance can also be logged in future to enhance the model. The experiments show that, despite the shortcomings, there is an enormous potential to be harnessed as the formidable results from generic models suggest. Unlike traditional deep learning models that require large amounts of data, just a fraction of that size is enough for *fine-tuning* or enhancing the embeddings with domain-specific information. This investigation leads to a conclusion that indeed, FAIRMetaText is a significant and promising direction for *automatically* making varied *existing* messy metadata FAIR and for specifying FAIR compliant metadata suggestions for future experimental datasets.

7 Conclusion

To address concerns of the quality of *existing* metadata, the authors presented an automatic metadata analysis and improvement tool. In accordance with the FAIR principles, the power of Large Language Models was used for *automatically* analyzing metadata. The current text embedding methods fare better than the erstwhile word embeddings [8] due to the following properties, namely identifying both syntactic and semantic properties congruently and the ability to handle out of vocabulary texts. The efficacy of our algorithms were demonstrated for metadata compliance and unification. These experiments were performed on both simulated and real datasets. In future work, the authors look to analyse actual user logs for metadata compliance and develop annotated datasets for evaluating the algorithms better. Fine-tuning LLMs with domain-specific data has been shown to improve performance in the literature [7]. In this work, commercial GPT embeddings have been used. An exploration of open-source versions of the same forms the basis of the next increment of this tool.

8 Acknowledgments

This work was supported in part by grant R01 LM013498 from the U.S. National Library of Medicine and by Award OT2 OD033759 from the U.S. National Institutes of Health Common Fund. The authors thank Josef Hardi for providing the expert-curated dataset and extensively working on the metadata of HuBMAP [23]. This project would not be possible without the accesible data available on BioPortal [20] and NCBI's GEO [27].

References

1. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. CoRR **abs/2005.14165** (2020), <https://arxiv.org/abs/2005.14165>
2. Bruce, T.R., Hillmann, D.: The continuum of metadata quality: Defining, expressing, exploiting (2004)
3. Bui, Y., Park, J.r.: An assessment of metadata quality: A case study of the national science digital library metadata repository. Proceedings of the Annual Conference of CAIS (Oct 2013). <https://doi.org/10.29173/cais166>, <https://journals.library.ualberta.ca/ojs.cais-acsi.ca/index.php/cais-ascii/article/view/166>
4. Dua, D., Graff, C.: UCI machine learning repository (2017), <http://archive.ics.uci.edu/ml>
5. Egyedi, A.L., O'Connor, M.J., Martínez-Romero, M., Willrett, D., Hardi, J., Graybeal, J., A Musen, M.: Using semantic technologies to enhance metadata submissions to public repositories in biomedicine (2018)

6. Egyedi, A.L., Willrett, D., Hardi, J., Graybeal, J., Musen, M.A.: Embracing semantic technology for better metadata authoring in biomedicine.
7. Eloundou, T., Manning, S., Mishkin, P., Rock, D.: Gpts are gpts: An early look at the labor market impact potential of large language models (2023)
8. 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. Springer (2019)
9. Gonsalves, R.S., O'Connor, M.J., nez Romero, M., Egyedi, A.L., Willrett, D., Graybeal, J., Musen, M.A.: The CEDAR Workbench: An Ontology-Assisted Environment for Authoring Metadata that Describe Scientific Experiments. *Semant Web ISWC* **10588**, 103–110 (Oct 2017)
10. Kramer, B.: Open metadata-where do we stand?(keynote). In: DaMaLOS (2021), <https://dx.doi.org/10.4126/FRL01-006429411>
11. Levenshtein, V.I.: Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady* **10**, 707 (Feb 1966)
12. Ma, E.: Nlp augmentation. <https://github.com/makcedward/nlpaug> (2019)
13. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(11) (2008)
14. MacQueen, J.: Classification and analysis of multivariate observations. In: 5th Berkeley Symp. Math. Statist. Probability. pp. 281–297. University of California Los Angeles LA USA (1967)
15. Martínez-Romero, M., O'Connor, M.J., Egyedi, A.L., Willrett, D., Hardi, J., Graybeal, J., Musen, M.A.: Using association rule mining and ontologies to generate metadata recommendations from multiple biomedical databases. *Database* **2019** (2019)
16. Matentzoglou, N., Balhoff, J.P., Bello, S.M., Bizon, C., Brush, M., Callahan, T.J., Chute, C.G., Duncan, W.D., Evelo, C.T., Gabriel, D., Graybeal, J., Gray, A., Gyori, B.M., Haendel, M., Harmse, H., Harris, N.L., Harrow, I., Hegde, H.B., Hoyt, A.L., Hoyt, C.T., Jiao, D., Jiménez-Ruiz, E., Jupp, S., Kim, H., Koehler, S., Liener, T., Long, Q., Malone, J., McLaughlin, J.A., McMurry, J.A., Moxon, S., Munoz-Torres, M.C., Osumi-Sutherland, D., Overton, J.A., Peters, B., Putman, T., Queralt-Rosinach, N., Shefchek, K., Solbrig, H., Thessen, A., Tudorache, T., Vasilevsky, N., Wagner, A.H., Mungall, C.J.: A Simple Standard for Sharing Ontological Mappings (SS-SOM). *Database* **2022** (05 2022). <https://doi.org/10.1093/database/baac035>, <https://doi.org/10.1093/database/baac035>, [baac035](https://doi.org/10.1093/database/baac035)
17. 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)
18. Musen, M.A.: Without appropriate metadata, data-sharing mandates are pointless. *Nature* **609**(7926), 222–222 (2022)
19. 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 (Nov 2022). <https://doi.org/10.1038/s41597-022-01815-3>, <https://doi.org/10.1038/s41597-022-01815-3>
20. Noy, N.F., Shah, N.H., Whetzel, P.L., Dai, B., Dorf, M., Griffith, N., Jonquet, C., Rubin, D.L., Storey, M.A., Chute, C.G., Musen, M.A.: BioPortal: ontologies and integrated data resources at the click of a mouse. *Nucleic Acids Res* **37**(Web Server issue), W170–173 (Jul 2009)

21. Park, J.R.: Metadata quality in digital repositories: A survey of the current state of the art. *Cataloging & Classification Quarterly* **47**(3-4), 213–228 (2009). <https://doi.org/10.1080/01639370902737240>, <https://doi.org/10.1080/01639370902737240>
22. 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)
23. Snyder, M.P., Lin, S., Posgai, A., Atkinson, M., Regev, A., Rood, J., Rozenblatt-Rosen, O., Gaffney, L., Hupalowska, A., Satija, R., Gehlenborg, N., Shendure, J., Laskin, J., Harbury, P., Nystrom, N.A., Silverstein, J.C., Bar-Joseph, Z., Zhang, K., Börner, K., Lin, Y., Conroy, R., Procaccini, D., Roy, A.L., Pillai, A., Brown, M., Galis, Z.S., Cai, L., Trapnell, C., Jackson, D., Nolan, G., Greenleaf, W.J., Plevritis, S., Ahadi, S., Nevins, S.A., Lee, H., Schuerch, C.M., Black, S., Venkataraman, V.G., Esplin, E., Horning, A., Bahmani, A., Sun, X., Jain, S., Hagood, J., Pryhuber, G., Kharchenko, P., Bodenmiller, B., Brusko, T., Clare-Salzler, M., Nick, H., Otto, K., Wasserfall, C., Jorgensen, M., Brusko, M., Maffioletti, S., Caprioli, R.M., Spraggins, J.M., Gutierrez, D., Patterson, N.H., Neumann, E.K., Harris, R., deCaestecker, M., Fogo, A.B., van de Plas, R., Lau, K., Yuan, G.C., Zhu, Q., Dries, R., Yin, P., Saka, S.K., Kishi, J.Y., Wang, Y., Goldaracena, I., Ye, D., Burnum-Johnson, K.E., Piehowski, P.D., Ansong, C., Zhu, Y., Desai, T., Mulye, J., Chou, P., Nagendran, M., Teichmann, S.A., Paten, B., Murphy, R.F., Ma, J., Kiselev, V.Y., Kingsford, C., Ricarte, A., Keays, M., Akoju, S.A., Ruffalo, M., Vella, M., McCallum, C., Cross, L.E., Friedman, S.H., Heiland, R., Herr, B., Macklin, P., Quardokus, E.M., Record, L., Sluka, J.P., Weber, G.M., Blood, P.D., Ropelewski, A.J., Shirey, W.E., Scibek, R.M., Mabee, P., Lenhardt, W.C., Robasky, K., Michailidis, S., Marioni, J., Butler, A., Stuart, T., Fisher, E., Ghazanfar, S., Eraslan, G., Biancalani, T., Vaishnav, E.D., Roy, A., Galis, Z., Srinivas, P., Pawlyk, A., Sechi, S., Wilder, E., Anderson, J., Consortium, H., Group, W., TMC, C.U., TMC, S.W., TMC, U., of Florida TMC, U., TMC, V.U., of Technology TTD, C.I., TTD, H., TTD, P., TTD, S., HuBMAP Integration, V., Engagement (HIVE) Collaboratory: Carnegie Mellon, T.C., Harvard Medical School, T.C., Indiana University Bloomington, M.C., Center, P.S., University of Pittsburgh, I., Component, E., University of South Dakota, C.C., New York Genome Center, M.C., Group, N.H.W.: The human body at cellular resolution: the nih human biomolecular atlas program. *Nature* **574**(7777), 187–192 (Oct 2019). <https://doi.org/10.1038/s41586-019-1629-x>, <https://doi.org/10.1038/s41586-019-1629-x>
24. T., B.: Biosample (2013 Nov 14), <https://www.ncbi.nlm.nih.gov/books/NBK169436/>
25. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering* **29**(12), 2724–2743 (2017)
26. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.W., da 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)
27. Zhu, Y., Davis, S., Stephens, R., Meltzer, P.S., Chen, Y.: GEOmetadb: powerful alternative search engine for the Gene Expression Omnibus. *Bioinformatics* **24**(23), 2798–2800 (11 2008). <https://doi.org/10.1093/bioinformatics/btn520>, <https://doi.org/10.1093/bioinformatics/btn520>