

Metadata Representations for Queryable ML Model Zoos

Li, Z.; Hai, Rihan; Bozzon, Alessandro; Katsifodimos, Asterios

Publication date

2022

Document Version

Final published version

Citation (APA)

Li, Z., Hai, R., Bozzon, A., & Katsifodimos, A. (2022). *Metadata Representations for Queryable ML Model Zoos*. Paper presented at ICML 2022 Workshop: DataPerf Benchmarking Data for Data-Centric AI.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Metadata Representations for Queryable ML Model Zoos

Ziyu Li¹ Rihan Hai¹ Alessandro Bozzon¹ Asterios Katsifodimos¹

Abstract

Machine learning (ML) practitioners and organizations are building model zoos of pre-trained models, containing metadata describing properties of the ML models and datasets that are useful for reporting, auditing, reproducibility, and interpretability purposes. The metadata is currently not standardised; its expressivity is limited; and there is no interoperable way to store and query it. Consequently, model search, reuse, comparison, and composition are hindered. In this paper, we advocate for standardized ML model metadata representation and management, proposing a toolkit supported to help practitioners manage and query that metadata.

1. Introduction

To incentivize and enable the sharing and reuse of ML models, organizations are building repositories of pre-trained ML models, also known as *model zoos* or *model hubs*; known examples include HuggingFace ([hug](#)), Tensorflow Hub ([ten](#)), and PyTorch Hub ([pyt](#)). The metadata of a ML model can contain all the necessary information – from a model’s inference capabilities (e.g., identified object classes) to the architecture, execution time, training dataset, hyper-parameter configuration, and evaluation performance. Metadata management are developed in the context of an ongoing effort to promote Trustworthy and Responsible AI¹.

Metadata management for model zoos can facilitate MLOps and open up new opportunities for new and unexplored use cases. For instance: i) retrieving models from large repositories with complex filtering conditions; ii) continuous integration of models in production; iii) (semi-)automatic model composition; and iv) advanced data management system. We will further discuss use cases in Section 2.

¹Delft University of Technology, The Netherlands. Correspondence to: <z.li-14@tudelft.nl>.

The potential of model zoos is currently hindered by the lack of a structured, queryable metadata format. Current repositories include a wide range of information, in a form of a model card ([Mitchell et al., 2019](#)), but such information is mostly for human consumption, making it hard for automatic extension or management. At the same time, the level of the detail remains coarse-grained: for instance, Amazon SageMaker, AzureML, MLflow ([Zaharia et al., 2018](#)) do not require mandatory reporting of the related metadata, except for model name and version. Practitioners have to search on external websites for further metadata information such as the data instances, and they even have to evaluate the model at hand in order to assess its performance.

Table 1 lists some example queries that are practical for ML developers and users, and we will further discuss them in Section 2. These queries require more fine-grained model metadata that current model repositories such as HuggingFace do not support. This highlights the need for detailed metadata of trained ML models and datasets in a structured and queryable representation.

As a step toward this goal, we present a metadata model to represent the metadata of ML models and related datasets. Based on this metadata model, we are currently developing a ML model metadata management platform (prototype available at [modelsearch.io](#)) that can be used to query such metadata.

The contributions of this paper go as follows:

- We showcase two use cases abstracted from ML lifecycle, which advocates the need to define appropriate metadata representation for model zoos (Section 2).
- We propose a metadata model as the structured, queryable, and comprehensive metadata representation for model zoos (Section 3.1).
- We develop an early-stage tool to retrieve ML models and datasets based on our proposed metadata representation, which facilitates complex inference queries over model zoos (Section 3.2).

2. Use Cases

To demonstrate the usefulness of a formalized metadata representation, in this section we explain two use cases with

Table 1. Example queries

Property	ID	Query
Dataset information	1	Retrieve text classification models trained on dataset crowdsourced by at least a group of 50 people
	2	Find a dataset collected from COCO and OpenImage with the all images containing "dog"
Model performance	3	Retrieve models trained on ImageNet with an accuracy higher than 90%
	4	Which model performs the best on COCO for person detection tasks?
Interpretability	5	Retrieve a person detection model with no gender bias
	6	Retrieve text generation models that do not generate hate speech
Hardware-related	7	Retrieve image classification models that are suitable to deploy on edge devices

different stakeholders, where a comprehensive and well-structured metadata representation is needed.

2.1. Metadata Retrieval Throughout the ML Life-cycle

Throughout the ML life-cycle, ML practitioners will require different metadata for tracking the ML model status, editing, comparing, or reporting. An ML practitioner often needs to query models in large repositories with complex filtering conditions, e.g., data instance, performance, and inherit mechanism. In Table 1, we list a few example queries revealing different properties of the metadata. For example, Query 1 and 2 require the metadata regarding the dataset, i.e., its attribute and source. Queries such as Query 3 to 6 require some other metadata properties. Query 7, on the other hand, requires more complex information regarding the inference performance with specified hardware settings. For example, an edge device may have constraints such as limited computation power and storage. To answer this query, practitioners will need to obtain the model performance of different objectives, e.g., inference speed and memory footprint.

2.2. (Semi-)Automatic Model Composition

Now we introduce a more advanced yet common use case. With metadata being captured and well-represented, ML practitioners can make good use of the models trained offline and apply them to answer complex, ad-hoc inference queries.

As shown in Figure 1, the ad-hoc query can consist of multiple ML inference tasks with different dependencies and relations. Moreover, the query can be composed of spec-

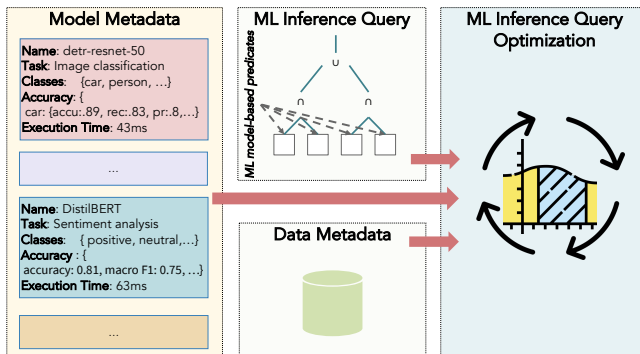


Figure 1. ML inference query optimization

ified constraints/requirements (e.g., latency and accuracy restrictions). Practitioners can select a composition of models from the model zoo to answer the ad-hoc queries. For example, an ML practitioner would like to design an application that can capture tweets with positive sentiment and a part-of-speech (PoS) tagger for data analysis. Since the data volume is significant and latency is also an essential factor to consider, the practitioner should select text classification model and PoS tagging model with fast inference speed. And the inference speed is greatly affected by the hardware being applied. If the application is deployed on the mobile phone, then the memory footprint is also a fundamental objective to be concerned with. To identify which set of models could best address the query and constraints, they may require information regarding the model performance with different objectives (e.g., accuracy, inference speed, memory footprint). The metadata of the dataset can also provide information to detect concept drift, for instance.

State-of-the-art model repositories would no be able to support these use cases, as they often maintain limited metadata, like the model name or the training dataset. A first step towards accessible repository is the specification of a rich and structured metadata format.

3. Methodology

This section defines a metadata model to represent different entities and relations. We also propose a practical tool to extract, store, and query the metadata. With our proposed comprehensive metadata in a structured representation and the tool, users can address the use cases discussed in the previous section by retrieving metadata in fine-grained details.

3.1. Metadata Model

Figure 2 presents a conceptual view of our proposed metadata model (from now on, *meta-model*), described using UML class diagrams. The *meta-model* composes of four packages: i) the ML Structure package, which defines

the ML models associated with architecture, hyperparameter settings, input and output; ii) the Dataset package, which represents information on datasets; iii) the Execution package, which describes the inference results of the model, possibly enriched with description from a knowledge graph; iv) the Evaluation package, which presents the run-time metrics obtained in specified hardware settings.

Due to space limitation, and for the sake of clarity, we present a compact version of the *meta-model*, which is amenable to implementation in different formats, including relational databases and Linked Data.

3.1.1. ML STRUCTURE PACKAGE

The package composes of three classes: ML model, Hyperparameter, and Architecture. The ML Model class encodes information about the *input* and *output* of the model, the prediction *task*, and the data *transformation* processes steps (e.g. data transformation and post-processing). For traditional ML model, the transformation pipelines can include steps for feature engineering. Such a recording of the ML model details enables exploration tasks such as model selection, data lineage, and visualization. It also helps reproduce the ML model and compare it to others from different perspectives. The Hyperparameter, as the name suggests, capture the type and values of hyperparameters used to train the specific model. Finally, the Architecture class contains information about the specific structural characteristic of a model, e.g. if it is a Feed-Forward Neural Network.

3.1.2. DATASET PACKAGE

The behaviour of a ML model heavily rely on the data it used. Thus the *meta-model* includes a Dataset packaged, representing both Datasets and their Data Instances. With the Dataset element, we present the metadata of the datasets that is significant for data management and reporting, e.g., data source, data version. With Data Instance we facilitate data provenance, for which it is of utmost importance to know which instances the model is trained on and how the data is utilized. Moreover, it is also significant to know if the data has revealed sensitive information, for example, violating the GDPR regulations.

3.1.3. EXECUTION PACKAGE

The Execution package includes two important classes. The Model Prediction class allows the description of inference metrics on one or more prediction output. Optionally a Model Prediction can be linked with a *Semantic Concept* from a knowledge base, thus allowing complex reasoning. Consider the example that an ML practitioner is building an ML model for image classification of cars, and she tries to conduct model diagnosis. She may have

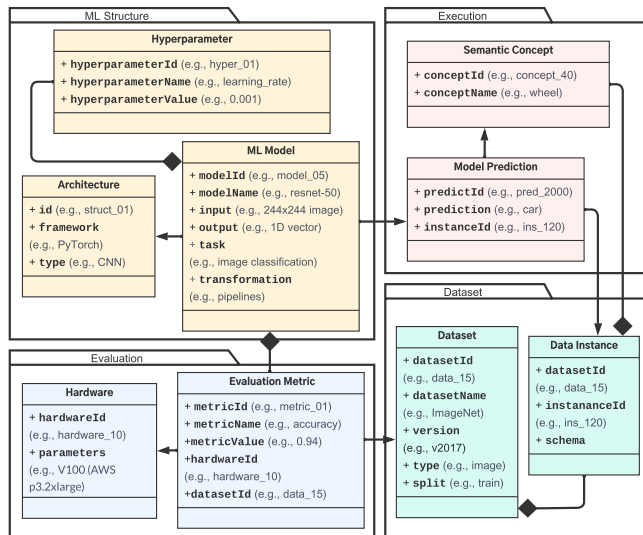


Figure 2. Modeling the metadata throughout ML life-cycle

several questions: what makes the model identify a car as a car? What are the semantic concepts that the model is capable to identify? Are the wheels that make it believe that it is a car? To support such use-cases, the *meta-model* allows to store information about inference performance on specific data instances, which can be used to describe the the behavior of a model, i.e., in what circumstances a model can perform well and why. Such information and awareness of the model prediction, may significantly improve ML model interpretability in various applications such as health care, law enforcement and finance.

3.1.4. EVALUATION PACKAGE

The performance evaluation of an ML model is critical during both the training and deployment phases. The practitioners need to deploy a suitable ML model for the task given a specified environment, e.g., Hardware parameters. A mismatch of the deployment will lead to latency issues or reliability concerns, which results in user dissatisfaction. Thus a well-rounded evaluation of the model performance is significant. Hence, we have the Hardware preserving the hardware settings and Evaluation Metric with metric-related metadata. Different models will need different evaluation metrics, for instance, a regression model would have a *Mean Square Error - MSE*, while a multi-class classification model would have an *accuracy* value for each class.

3.2. A Metadata Management Tool

We are currently developing a tool² to manage the metadata for model zoos, whose structure is shown in Figure 3.

²<https://modelsearch.io>

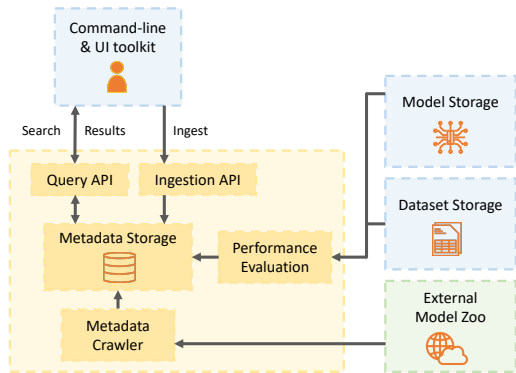


Figure 3. The structure of our metadata management tool

We collect metadata in the following three ways. (i) a user can add the metadata regarding a model or a dataset by filling in specified fields. Then such information is translated by the *Ingestion API* into structured representation according to the above-mentioned metadata model, and stored in the *Metadata Storage*. (ii) The tool can also automatically extract information from external model zoos, e.g., HuggingFace and PyTorch Hub. We can extract metadata from their API or information on the web page and record the source of the external metadata in our metadata field. (iii) To gather the information regarding the model performance, we apply a third way to obtain the metadata. We obtain the performance evaluation by executing the model on a dataset with specified hardware settings. This process can be done by utilizing cloud resources.

By the time we are writing the paper, the tool is available to crawl information from the external model zoos and support querying the metadata stored in structured and queryable representations. Besides ingesting, extracting, and storing the metadata, our proposed tool allows a user to (i) retrieve the models that help them identify a model, (ii) compare multiple models, (iii) or explore the properties of models/data by composing queries on the metadata.

4. Related work

Recent studies focus on different aspects of management during ML life-cycle, from model versioning, model reporting to model evaluation. Each is important for practitioners to manage and understand the models. We observe a gap among the profound works, a comprehensive and queryable metadata representation. With the metadata representation, we can thus better manage the ML models and data, including the interactions between them.

Metadata Captured in ML Versioning and Data Versioning

Recent works developed tools/systems to manage ML models and datasets used in ML life-cycle. Works such as Mod-

eldb (Vartak et al., 2016), ModelHub (Miao et al., 2017) and MLflow (Zaharia et al., 2018) developed tools to manage ML models with different versions and captured the metadata of the models in different levels of details. However, they focused on the abstractions on the model, they lack information on the model performance under different hardware settings (e.g., inference speed, accuracy, memory footprint), and it paid little attention on the datasets that the models consume. For ML dataset management and versioning, research work such as Mldp (Agrawal et al., 2019) and DataLab (Zhang et al., 2016) can be complements to the above-mentioned model management system by supporting data management.

Model Cards and Data Cards

Recent research also focus on the reporting of models and dataset, covering aspects not only limited to basic informative components, but also including ethical, inclusive and fair considerations. Model cards (Mitchell et al., 2019), for example, proposed to include information regarding model intended use cases, potential pitfalls and other contexts that can improve model understanding. Similar idea also lies in data cards/sheets. Examples include (Boyd, 2021; Gebru et al., 2021; Miceli et al., 2021). Though model cards and data cards contain rich information, the Q&A format is nonetheless unfriendly to machine to process and thus cannot be easily managed and retrieved.

Model Performance Benchmarking

A growing body of published work also focus into the benchmarking of ML model performance, such as MLperf (Mattson et al., 2020; Reddi et al., 2020), fathom (Adolf et al., 2016), Dawnbench (Coleman et al., 2017). These platforms covered a set of metadata including metrics, training/inference configurations with specified hardware/software setting. Their focus is the report of the model performance at different ML life-cycle stages (training or inference). They paid little attention on the dataset the model used, whose path is provided as an argument filled by the user. The model process pipeline is also not covered besides the model scripts.

5. Conclusion and outlook

In this paper, we advocate for the need of a structured, queryable, and comprehensive metadata representation for model zoos. We propose a metadata model for such metadata representation to tackle different use cases. We also develop a tool that helps practitioners to manage and query the metadata. Future work can develop tools/frameworks to extract more useful information from the textual description in the model/data cards by applying natural language processing techniques and mapping it to the predefined metadata representation.

References

- Huggingface. <https://huggingface.co/>. Accessed: 2022-04-10.
- Pytorch hub. <https://pytorch.org/hub/>. Accessed: 2022-04-10.
- Tensorflow hub. <https://www.tensorflow.org/hub>. Accessed: 2022-04-10.
- Adolf, R., Rama, S., Reagen, B., Wei, G.-Y., and Brooks, D. Fathom: Reference workloads for modern deep learning methods. In *2016 IEEE International Symposium on Workload Characterization (IISWC)*, pp. 1–10. IEEE, 2016.
- Agrawal, P., Arya, R., Bindal, A., Bhatia, S., Gagneja, A., Godlewski, J., Low, Y., Muss, T., Paliwal, M. M., Raman, S., et al. Data platform for machine learning. In *Proceedings of the 2019 International Conference on Management of Data*, pp. 1803–1816, 2019.
- Boyd, K. L. Datasheets for datasets help ml engineers notice and understand ethical issues in training data. *Proceedings of the ACM on Human-Computer Interaction*, 5 (CSCW2):1–27, 2021.
- Coleman, C., Narayanan, D., Kang, D., Zhao, T., Zhang, J., Nardi, L., Bailis, P., Olukotun, K., Ré, C., and Zaharia, M. Dawnbench: An end-to-end deep learning benchmark and competition. *Training*, 100(101):102, 2017.
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., and Crawford, K. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.
- Mattson, P., Cheng, C., Damos, G., Coleman, C., Micikevicius, P., Patterson, D., Tang, H., Wei, G.-Y., Bailis, P., Bittorf, V., et al. Mlperf training benchmark. *Proceedings of Machine Learning and Systems*, 2:336–349, 2020.
- Miao, H., Li, A., Davis, L. S., and Deshpande, A. Towards unified data and lifecycle management for deep learning. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pp. 571–582. IEEE, 2017.
- Miceli, M., Yang, T., Naudts, L., Schuessler, M., Serbanescu, D., and Hanna, A. Documenting computer vision datasets: an invitation to reflexive data practices. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 161–172, 2021.
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., and Gebru, T. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pp. 220–229, 2019.
- Reddi, V. J., Cheng, C., Kanter, D., Mattson, P., Schmuelling, G., Wu, C.-J., Anderson, B., Breughe, M., Charlebois, M., Chou, W., et al. Mlperf inference benchmark. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 446–459. IEEE, 2020.
- Vartak, M., Subramanyam, H., Lee, W.-E., Viswanathan, S., Husnoo, S., Madden, S., and Zaharia, M. Modeldb: a system for machine learning model management. In *Proceedings of the Workshop on Human-In-the-Loop Data Analytics*, pp. 1–3, 2016.
- Zaharia, M., Chen, A., Davidson, A., Ghodsi, A., Hong, S. A., Konwinski, A., Murching, S., Nykodym, T., Ogilvie, P., Parkhe, M., et al. Accelerating the machine learning lifecycle with mlflow. *IEEE Data Eng. Bull.*, 41 (4):39–45, 2018.
- Zhang, Y., Xu, F., Frise, E., Wu, S., Yu, B., and Xu, W. Datalab: a version data management and analytics system. In *Proceedings of the 2nd International Workshop on BIG Data Software Engineering*, pp. 12–18, 2016.