

Fermilab

An MLCommons Scientific Benchmarks Ontology

FERMILAB-PUB-25-0701-CSAID

arXiv:2511.05614

This manuscript has been authored by Fermi Forward Discovery Group, LLC  
under Contract No. 89243024CSC000002 with the U.S. Department of Energy,  
Office of Science, Office of High Energy Physics.

# An MLCommons Scientific Benchmarks Ontology

1<sup>st</sup> Ben Hawks  
*CS and AI Directorate*  
*Fermilab*  
 Batavia, IL, USA  
 bhawks@fnal.gov

2<sup>nd</sup> Gregor von Laszewski  
*Biocomplexity Institute*  
*University of Virginia*  
 Charlottesville, VA, USA  
 laszewski@gmail.com

3<sup>rd</sup> Matthew D. Sinclair  
*CS Dept*  
*University of Wisconsin-Madison*  
 Madison, WI, USA  
 sinclair@cs.wisc.edu

4<sup>th</sup> Marco Colombo  
*Discovery Partners Institute*  
*University of Illinois Urbana-Champaign*  
 Chicago, IL, USA  
 mcolom4@illinois.edu

5<sup>th</sup> Shivararam Venkataraman  
*CS Dept*  
*University of Wisconsin-Madison*  
 Madison, WI, USA  
 shivaram@cs.wisc.edu

6<sup>th</sup> Rutwik Jain  
*CS Dept*  
*University of Wisconsin-Madison*  
 Madison, WI, USA  
 rnjain@wisc.edu

7<sup>th</sup> Yiwei Jiang  
*CS Dept*  
*University of Wisconsin-Madison*  
 Madison, WI, USA  
 jiang357@wisc.edu

8<sup>th</sup> Nhan Tran  
*CS and AI Directorate*  
*Fermilab*  
 Batavia, IL, USA  
 ntran@fnal.gov

9<sup>th</sup> Geoffrey Fox  
*Biocomplexity Institute and CS Dept*  
*University of Virginia*  
 Charlottesville, VA, USA  
 vxj6mb@virginia.edu

**Abstract**—Scientific machine learning research spans diverse domains and data modalities, yet existing benchmark efforts remain siloed and lack standardization. This makes novel and transformative applications of machine learning to critical scientific use-cases more fragmented and less clear in pathways to impact. This paper introduces an ontology for scientific benchmarking developed through a unified, community-driven effort that extends the MLCommons ecosystem to cover physics, chemistry, materials science, biology, climate science, and more. Building on prior initiatives such as XAI-BENCH, FastML Science Benchmarks, PDEBench, and the SciMLBench framework, our effort consolidates a large set of disparate benchmarks and frameworks into a single taxonomy of scientific, application, and system-level benchmarks. New benchmarks can be added through an open submission workflow coordinated by the MLCommons Science Working Group and evaluated against a six-category rating rubric that promotes and identifies high-quality benchmarks, enabling stakeholders to select benchmarks that meet their specific needs. The architecture is extensible, supporting future scientific and AI/ML motifs, and we discuss methods for identifying emerging computing patterns for unique scientific workloads. The MLCommons Science Benchmarks Ontology provides a standardized, scalable foundation for reproducible, cross-domain benchmarking in scientific machine learning. A companion webpage for this work has also been developed as the effort evolves: <https://mlcommons-science.github.io/benchmark/>

**Index Terms**—benchmark, mlcommons

## CONTENTS

<b>I</b>	<b>Introduction</b>	2
I-A	Related Work	2
FERMILAB-PUB-25-0701-CSAID		
<b>II</b>	<b>Scientific Benchmark Definition</b>	3
II-A	Problem Specification and Constraints	3
II-B	Dataset	3
II-C	Performance Metric(s)	3
II-D	Reference Solution	4
II-E	Documentation and Reproducible Protocol	4
<b>III</b>	<b>Existing Benchmarks, Rating and Endorsement System</b>	4
III-A	Software Environment	4
III-B	Problem Specification and Constraints	4
III-C	Dataset	4
III-D	Performance Metrics	4
III-E	Reference Solution	5
III-F	Documentation	5
III-G	Existing Benchmark Ontology	5
<b>IV</b>	<b>Motifs</b>	5
IV-A	Primary Users of the Benchmark Ontology	5
IV-B	Motifs within the ontology	5
IV-B1	Scientific Motifs	5
IV-B2	AI/ML Motifs	7
IV-C	Emerging Computing Patterns	8
IV-C1	Computing Motifs	8
IV-C2	Clustering Algorithm	8
<b>V</b>	<b>Summary</b>	9
<b>Appendix</b>		10
<b>References</b>		13

## I. INTRODUCTION

Benchmarking in scientific machine learning (ML) has emerged as a critical area to guide algorithm development, enable fair comparisons towards progress and innovation, and facilitate reproducibility. The development of ML benchmarks for science is especially critical because of the multi-disciplinary nature of the development, often including domain experts, computing hardware developers, and ML researchers. That, coupled with the variety of tasks and workloads, makes *high quality* benchmarking critical to making progress.

Our contribution, the MLCommons Science Benchmark Ontology, builds on prior scientific ML benchmarking efforts by integrating their strengths into a unified, community-driven paradigm. Unlike prior domain-specific benchmarking efforts, it provides a unifying principle through an ontology of curated benchmarks across multiple scientific domains. By situating this work under the MLCommons governance model, we ensure long-term community support, extensibility, and adoption, thereby moving from individual benchmark initiatives toward a standardized, widely applicable benchmark ontology for scientific machine learning.

Although each of these benchmarks can be valuable to their communities, there is no standard definition of a high-quality benchmark. This makes accessibility to the problem for different researchers with complementary expertise challenging. While there is the appearance that there are many “benchmarks” for scientific applications in the literature [1], [2], [3], [4], [5], [6], [7], [8], [9], our analysis finds that their quality varies significantly. Often, a dataset is provided without concrete tasks, the code is not reproducible, or the metrics are undefined.

to run all of these benchmarks simply due to the sheer volume.

A given stakeholder may also have different priorities in searching for benchmarks. For example, hardware vendors typically utilize benchmarks to understand the demands of state-of-the-art applications on potential future systems and ensure these systems provide high performance for important workloads. Researchers and domain scientists use these benchmarks to study how potential hardware and software optimizations will apply to current and/or future systems, and how different solutions to common domain-specific problems compare against others in a standardized format. Nevertheless, since different benchmarks from different domains value and stress different things, finding a way to identify relevant, representative subsets for different stakeholders is imperative.

Our goal is to develop a standard definition of benchmarking that can be used by the scientific community for ML benchmarking. They will be classified by different domains and ML motifs. We will design an accompanying web portal that is searchable based on benchmark quality, domain, and ML motif. Then we demonstrate an example of how these benchmark tasks can be used to understand emerging computational patterns for scientific HPC. This overall idea is illustrated in Fig. 1.

### A. Related Work

As shown in Table I, prior efforts have addressed aspects of this challenge. However, they have often done so within narrow domains or specialized contexts. Explainability-focused benchmarks such as XAI-BENCH [10] provide synthetic datasets with known ground-truth feature importance to systematically evaluate explainable AI methods. While highly valuable for interpretability studies, the scope

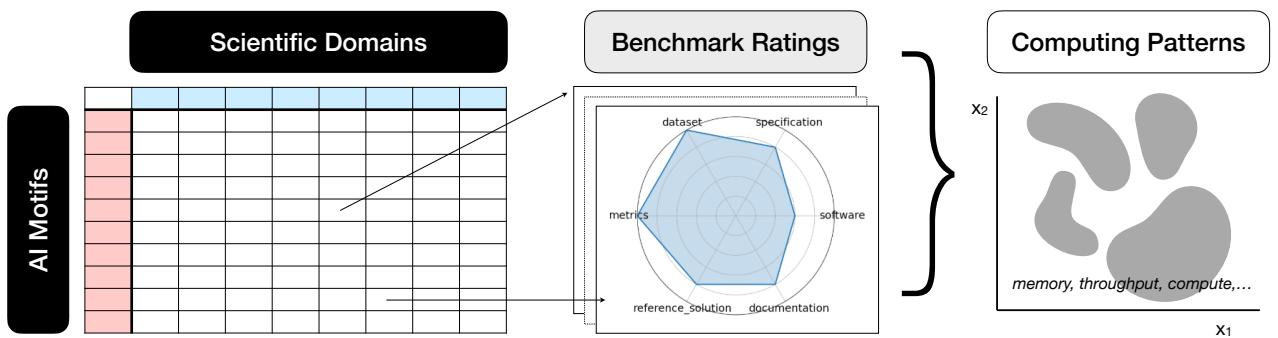


Fig. 1. Overview of the Scientific ML Benchmark ontology from the taxonomization of benchmarks by domain and ML motif to the qualification by benchmark through a standardized rating system to the use of the ontology to understand computing patterns for scientific workflows.

TABLE I  
COMPARISON OF PRIOR RELATED WORKS TO THIS WORK

Work	Focus Area	Scope	Dataset/Tasks	Key Contributions
XAI-Bench [10]	Explainable AI benchmarking	Feature attribution	Synthetic datasets with ground-truth feature importance	Provides controllable synthetic datasets for evaluating explainers; narrow focus on interpretability rather than general scientific ML
FastML Science Benchmarks [14]	Edge ML for physics	Real time, low latency ML workloads	Tasks in HEP jet tagging, sensor data compression, accelerator control	First to emphasize hardware-software co-design and real-time constraints in scientific ML; scope limited to latency-critical workloads
PDEBench [12]	PDE-based surrogate modeling	Scientific ML for PDE emulation	11 PDEs (1D–3D: Burgers, Navier–Stokes, shallow-water, Darcy flow, etc.)	Large-scale PDE datasets for forward + inverse problems; extensible APIs; baselines (FNO, U-Net, PINNs)
SciMLBench [13]	General benchmarking framework	Cross-domain ML benchmarks	Datasets from Astronomy, particle physics, materials, and life sciences	Proposes taxonomy: scientific, application, and system benchmarking; emphasizes community standards, curated datasets
<b>This Work</b>	End to end scientific ML benchmarking	Broad, community-driven suite integrating diverse workloads	Covers multiple scientific domains (e.g., physics, chemistry, materials, life sciences) with curated benchmarks emphasizing quality	Extends prior work by unifying domain-specific efforts into a standardized, extensible benchmark ontology under MLCommons; stresses reproducibility, FAIR data, community governance, and broad accessibility beyond single-domain focus

hardware/software codesign and real time inference, but they are domain specific and do not generalize across the wide range of scientific machine learning tasks.

A more domain specialized but broader dataset contribution is PDEBENCH (Takamoto et al., 2022) [12], which provides large-scale benchmarks for surrogate modeling of partial differential equations (PDEs). Covering forward and inverse problems with diverse PDEs, PDEBENCH offers extensible datasets and APIs along with physics aware evaluation metrics. However, its focus remains limited to PDE based simulations, without addressing broader scientific domains or end to end benchmarking pipelines.

In contrast, SciMLBench [13] introduces a conceptual framework for scientific ML benchmarking. It distinguishes between scientific, application, and system-level benchmarks, and emphasizes the importance of curated datasets, reproducibility, and community standards. While influential in defining the benchmarking landscape, SciMLBench functions more as a blueprint than an implemented suite of benchmarks.

## II. SCIENTIFIC BENCHMARK DEFINITION

A **benchmark** is a carefully defined *standardized* version of a scientific application that is used for making quantifiable comparisons of solutions. While many elements below are commonly identified in the literature, they provide the base on which we can structure our work.

### A. Problem Specification and Constraints

A succinct statement of the benchmark *task* describes the type of input data, the expected output, and any constraints on the task. An example of the input description could include its origin or representation, such as image data, time-series data, 3D point cloud, or natural language. The expected output describes the transformation of the

data, such as a regression task, anomaly detection, or generative model.

The system *constraints* are quantifiable elements of the benchmark (power, latency, etc.) or system specifications (hardware platform, technology node, etc.) that are not a part of the performance metric(s) being compared. System constraints are typically bounds with an upper or lower limit.

### B. Dataset

The dataset is the scaffolding which includes input data and potentially truth labels if a part of the benchmark. It adheres to the FAIR principles: each instance is uniquely identified and documented for *findability*; *access* is provided through persistent, open protocols; formats and metadata follow community standards to ensure *interoperability*; and versioning with associated preprocessing scripts enables full *reproducibility* of results. The dataset is bounded: no data augmentation, enrichment, or post hoc curation is permitted unless otherwise specified, ensuring a stable target for comparative evaluation. A canonical, non-overlapping split into training, validation, and test sets is provided to facilitate standardized model development and unbiased performance reporting.

### C. Performance Metric(s)

The performance metrics are quantifiable measures for comparison. There is an important distinction between constraints (fixed bounds) and performance metrics (measurable bases for comparison). A benchmark can have multiple *dimensions* such that more than one metric can be optimized simultaneously that forms a Pareto optimization.

For example, measures such as mean squared error (MSE) over a held-out test set, relative error in conserved quantities, computational cost (CPU/GPU time), or memory

footprint can all be considered. Sometimes benchmarks also track stability or robustness (e.g. performance when noise is added). The metrics must be well-defined, reproducible, and measurable when constructed from the validation dataset.

Multiple **benchmarks** may be defined for a given **task** based on its performance metrics. For example, for the *same* dataset, one benchmark may require a lower bound on accuracy constraint while testing hardware throughput and another benchmark may only test for the best algorithm accuracy.

#### D. Reference Solution

A reference solution satisfies the problem specification and constraints using the defined dataset and includes performance metrics measurements as a baseline to which other solutions can be compared.

#### E. Documentation and Reproducible Protocol

A benchmark must include a protocol to reproduce the reference solution. This will enable additional solutions to make direct comparisons. Additional proposed solutions must also be reproducible. This is typically through reference code within a well-defined software environment that is versioned. Hardware references must include thorough documentation and a bill of materials.

### III. EXISTING BENCHMARKS, RATING AND ENDORSEMENT SYSTEM

In curating the benchmark ontology, we identify numerous existing scientific benchmarks and attempt to categorize and rate them in order to compile them into a comprehensive benchmark ontology. Our goal in curating this ontology is not to create a ontology that is intended to be exhaustively and completely run, but instead to create a collaborative, high quality benchmark ontology that can be referenced and sampled from when a potential submitter might be searching for domain specific workloads, varying computing motifs, focus areas, and the next generation of challenging machine learning tasks.

However, modern workloads, especially in domains like ML, are evolving rapidly. Accordingly, our proposed benchmark ontology cannot remain static if we hope to continue representing the state-of-the-art. Thus, we allow new benchmarks to be proposed and potentially added to the benchmark ontology through an open submission process, reviewed by the MLCommons Science Working Group, and integrated into the ontology upon meeting evaluation standards. This model ensures that the benchmark remains both extensible and sustainable.

In order to determine quality benchmarks across different scientific domains, we have developed a domain agnostic rating system that can be used to identify quality benchmarks across six different categories, each aligning with a different aspect from the definition of the a benchmark as given in Section II.

Each category is given a score out of a possible 5 points. A benchmark may be scored a 0 in a category if none of the criteria for a given category are met, and a 5 if all the criteria for a given category are met. If a given benchmark obtains an average score across all 6 categories of at least 4.5 out of 5, it is given the "MLCommons Science Benchmark Endorsement", identifying it as a particularly high quality benchmark within the wider ontology.

The six categories that are used to evaluate a given benchmark, and their evaluation criteria, are as follows:

#### A. Software Environment

When evaluating the quality of a given benchmark's software environment, we award it a score based on the following criteria, giving it one point for each statement that is true:

- The code is available to reproduce the baseline reference solution.
- The provided code is complete.
- The code itself is well documented
- The code does not require any modifications to run.
- The environment is either containerized, or environment details and setup instructions are provided.

#### B. Problem Specification and Constraints

When evaluating the quality of a given benchmark's problem specification and system constraints, we award it points based on the following criteria, giving it one point for each statement that is true:

- System constraints, such as required power, latency, and/or throughput requirements, are provided.
- The benchmark task is clear.
- The dataset format is clearly specified.
- The task inputs are clearly specified.
- The task outputs are clearly specified.

#### C. Dataset

We evaluate the quality of a given benchmark's dataset by primarily judging how well it adheres to the FAIR Dataset Principles[15], giving a point for each of the 4 principles that are followed, and an additional point is given for if there are well defined train/test(/validation) splits present.

#### D. Performance Metrics

We evaluate the quality of a given benchmark's performance metrics by judging how well defined the metrics are, and how well they capture a given solutions performance. Points are awarded on the following scales, combining each scale's rating into a single rating for performance metrics overall:

Metric Definitions:

- 3 points if the metrics are fully defined
- 2 points if the metrics are clearly mentioned, but specific details/implementations (where applicable) are not defined.

- 1 point if some metrics are mentioned, but not clearly defined as what is being tracked for a given benchmark
- 0 points if no metrics are mentioned or defined.

Metric Quality:

- 2 Points if the metrics fully capture a given solution's performance
- 1 point if the metrics partially capture a given solution's performance
- 0 points if the metrics do not capture a given solution's performance

#### E. Reference Solution

When evaluating the quality of a given benchmark's reference solution, we award it points based on the following criteria, giving it one point for each statement that is true:

- A reference solution is publicly available
- The provided reference solution is well documented
- All hardware and/or software requirements to run the reference solution are listed
- All metrics defined as part of the benchmark are evaluated as part of the reference solution
- The baseline solution/model is openly available for study (and if a neural network, the architecture, hyperparameters, and code to train it is provided)

#### F. Documentation

When evaluating the quality of a given benchmark's documentation, we award it points based on the following criteria, giving it one point for each statement that is true:

- The task is explained and well documented.
- The task and benchmark background (scientific and otherwise) is clearly explained.
- The motivation for the benchmark is clearly explained.
- The evaluation criteria are clearly explained.
- An academic paper about the benchmark exists.

#### G. Existing Benchmark Ontology

We include a full list of benchmarks that are currently included in the ontology in Table II as a part of the Appendix, including the average score across the 6 evaluation categories. For full details on each benchmark, including individual scores for each category, please see von Laszewski, et al [16].

## IV. MOTIFS

### A. Primary Users of the Benchmark Ontology

We envision three primary types of users of the benchmark ontology, each having different needs for organizing and finding benchmarks within the ontology. To make sure that each user is best supported, we've made a website available along side the static report that allows for searching, filtering, and sorting entries within the benchmark ontology.

**User 1: Specific Benchmark:** In the first case, we expect a user that is searching for a specific benchmark that matches some given criteria, such as a specific type of

AI/ML task within a given scientific domain. This user is likely searching for a benchmark to test or evaluate a solution to a specific task. For this user, we've annotated each benchmark with metadata describing it, and have built in a search feature to the website that allows a user to search the metadata for all benchmarks.

**User 2: Specific Benchmark Category:** In the second case, we expect a user that is searching for a multiple benchmarks within given categories. This user is likely looking for a subset of benchmarks that are either entirely contained by a single category, such as all biology benchmarks, or some further narrowing of scope, such as all classification tasks within high energy physics. For this user, we provide a set of tools on the website to filter and sort the benchmark ontology.

**User 3: Similar Compute Patterns:** In the third case, we expect a user that is looking for multiple benchmarks, but not as defined by the given set of categories. This user might, for example, be searching for computationally similar workloads across multiple AI/ML tasks and scientific domains. For this user, we present a novel clustering algorithm as an example that can be used to find like benchmarks within the ontology based on the user's specific needs.

### B. Motifs within the ontology

Due to the broad range of benchmarks within the ontology, we attempt to organize the benchmarks within the ontology by different motifs. We attempt to identify motifs within 2 categories:

- Scientific Motifs - The scientific domain(s) that a given benchmark centers around (High energy physics, Chemistry, Biology, etc.)
- AI/ML Motifs - The type of AI/ML tasks that a given benchmark is trying to accomplish (Classification, Regression, Reasoning, etc.)

We list the broad motifs present in the current iteration of the benchmark ontology below, but intend to add more as it is expanded over time. The domains and motifs that are present in the ontology can additionally be seen in Fig. 2, which shows the number of tasks present in each domain/motif combination. Note that in the ontology, it is possible for a benchmark task to have multiple domains associated with it, but only a single AI/ML Motif. Building an exhaustive list fine grained domains and motifs is not possible as this is a constantly evolving ontology, but serves to broadly categorize each benchmark task within the ontology, as well as provide examples on which we can continue to add in order to create a more widely representative ontology of benchmarks.

1) *Scientific Motifs:* Within the benchmark ontology, we attempt to collect benchmarks that represent a wide variety of scientific domains.

a) *High-Energy Physics:* This domain includes benchmarks encompassing particle physics across all aspects of experiment and theory. Particular unique community

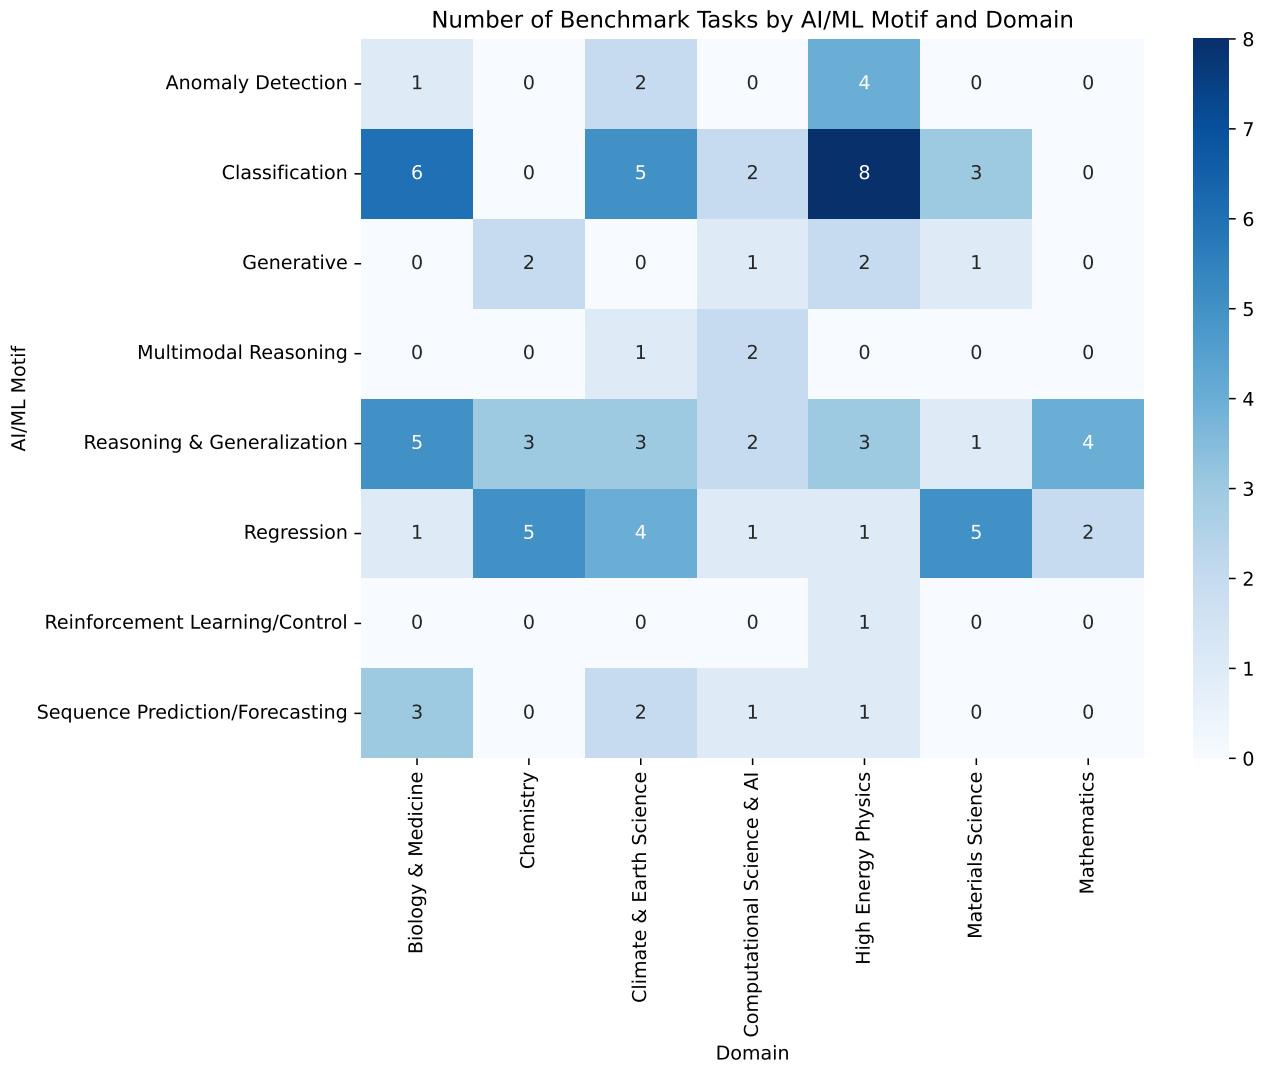


Fig. 2. A Heatmap showing the domain and AI/ML Motif for the tasks within the ontology. Note that each task can have multiple domains associated with it, but only a single AI/ML Motif

efforts in [11] are identified because they push the limits of inference latency, precision, and hardware-efficiency on custom accelerator back-ends. Jet-classification [11] and ultra-fast jet-tagging [17] benchmarks demand sub-microsecond inference on FPGA and quantify both accuracy and resource utilization. Beam-control is a reinforcement-learning problem formulated on a simulated accelerator that tests sample-efficient policy optimization and stability metrics [18]. Smart-pixels [19] and Neural Architecture Code-sign [20] further showcase on-chip inference and low-latency trigger generation for collider experiments.

*b) Chemistry:* Chemistry benchmarks span generative chemistry, catalytic modeling, and quantum-chemical corrections. MOLGEN asks a language model to generate chemically valid molecules in SELFIES and optimize log-P and docking scores, providing a stringent

distribution-learning test [21]. OCP (Open Catalyst Project) supplies a vast, DFT-derived adsorption-energy dataset that forces regression and graph-neural-network models to capture subtle electronic effects [22]. Delta-Squared-DFT adds a second layer of accuracy by training a correction model on CCSD(T) data, testing how well deep networks can interpolate high-level quantum chemistry from inexpensive DFT [23].

*c) Materials Science:* Materials-science benchmarks focus on high-throughput DFT databases and design-space exploration. Materials Project gives a global repository of >200 k inorganic crystal properties that support supervised learning for band-gap, formation energy, and elastic constants [24]. JARVIS-Leaderboard offers an open, community-driven leaderboard of design-method comparisons (CGCNN, ALIGNN, M3GNet) [25]. SuperCon3D is a

generative benchmark that asks models to predict high- $T_c$  superconductivity from 3-D crystal descriptors, combining regression and sampling [26].

*d) Biology & Medicine:* Benchmarks in biology emphasize data-driven discovery, multimodal understanding, and biomedical QA. BiasBench evaluates autonomous scientific discovery by requiring cell-type annotation and multiple-choice QA on scRNA-seq data [27]. SeafloorAI couples sonar imagery with natural-language questions, testing vision-language grounding in marine geology [28]. Vocal-Call Locator evaluates sound-source localization on multi-channel audio in a lab setting [29]. MedQA is a real-world medical board-exam dataset, measuring diagnostic inference and language grounding in clinical knowledge [30].

*e) Climate & Earth Sciences:* These benchmarks bring remote-sensing, time-series, and event-detection to the fore. ClimateLearn supplies ERA-5 based datasets for 3–5-day weather prediction, testing physics-aware deep learning [31]. HDR ML Anomaly Challenge (Sea-Level Rise) provides satellite-derived sea-level time-series for anomaly detection [32], while SatImgNet offers a multi-task satellite-image classification suite [33]. MLCommons Science – Earthquake forecasting is a simulation-based benchmark that evaluates inference speed and prediction accuracy for seismic events [34].

*f) Computational Science & AI Benchmarks:* These are testbeds that focus on evaluating general performance and throughput of models across multiple domains. SciCode is a scientist-curated coding benchmark across 16 scientific subfields, evaluating code synthesis for scientific computing tasks. [35]. MLPerf HPC benchmarks scientific-ML training on large-scale HPC clusters [36].

*g) Mathematics:* Mathematical benchmarks evaluate symbolic reasoning and long-horizon problem solving. FrontierMath pushes models on category theory, number theory, and algebraic geometry [37]. AIME [38] and PRM800k [39] provide structured, multiple-choice mathematics problems that evaluate generalization across difficulty levels.

## 2) *AI/ML Motifs:*

*a) Classification:* The classification benchmarks ask a model to assign a discrete label to a given input. Jet Classification evaluates a model’s ability to distinguish between different particle-jet types in high-energy physics data [11]. Smart Pixels for LHC tests on-chip inference by classifying pixel clusters in a 28nm CMOS detector [19]. BiasBench evaluates cell-type annotation from single-cell RNA-seq data, a multi-class classification problem [27]. SeafloorAI asks a model to segment and classify geological features in sonar imagery, while additionally evaluating natural-language questions [28]. SatImgNet is a multi-task satellite-image classification suite covering 27 remote-sensing datasets [33].

*b) Regression:* Regression benchmarks assess a model’s ability to predict continuous values. FEABench measures the accuracy of finite-element analysis solvers,

reporting solve-time and error norm [40]. CFDBench evaluates neural-operator regression on CFD data, reporting  $L^2$  error and MAE [41]. OCP (Open Catalyst Project) asks for regression of adsorption energies and forces [22]. Materials Project provides a large DFT database for regression of band-gap, formation energy, and elastic constants [24]. Delta Squared-DFT trains a correction network to map DFT outputs to CCSD(T) energies, evaluated on mean absolute error [23].

## *c) Sequence Prediction / Forecasting:*

Sequence-prediction tasks require a model to forecast future values of a time series. ClimateLearn provides standardized datasets and evaluation protocols for machine learning models in medium-range weather and climate forecasting using ERA5 reanalysis [31]. GB-Biology is a suite of large-scale biological network datasets (protein-protein interaction, drug-target, etc.) with standardized splits and evaluation protocols for node, link, and graph property prediction tasks. [42].

*d) Anomaly Detection:* Anomaly-detection benchmarks test a model’s ability to flag outliers or rare events. HDR ML Anomaly Challenge (Sea-Level Rise) uses satellite-derived sea-level time-series to detect flooding anomalies, evaluated with ROC-AUC and precision/recall [32]. HDR ML Anomaly Challenge (Gravitational Waves) detects anomalous gravitational-wave signals from LIGO/Virgo data, measured by ROC-AUC and precision/recall [43].

## *e) Reinforcement Learning / Control:*

Reinforcement-learning benchmarks involve learning a policy to control a system. Beam Control trains an RL agent to stabilize accelerator beam trajectories, evaluated on stability loss and control-latency [18]. Intelligent Experiments through Real-time AI uses FPGA-based trigger generation for sPHENIX/EIC, measured by accuracy and latency [44]. Quench Detection employs RL and autoencoders to detect superconducting magnet quenches in real time, evaluated by ROC-AUC and detection latency [45].

*f) Generative:* Generative benchmarks ask a model to produce novel data samples. MOLGEN generates chemically valid molecules in SELFIES, optimizing log-P, QED, and docking scores [21]. SuperCon3D generates high- $T_c$  superconductors from 3-D crystal descriptors, evaluated by MAE and validity [26]. Delta-Squared-DFT generates corrected DFT energies, evaluated by mean absolute error [23].

*g) Multimodal Reasoning:* Multimodal benchmarks combine vision, language, and sometimes audio. SPIQA (Scientific Paper Image QA) requires a model to answer questions about scientific figures, evaluated by accuracy and F1 [46]. SeafloorAI pairs sonar imagery with natural-language questions, measured by segmentation pixel accuracy and QA accuracy [28].

*h) Surrogate Modeling:* Surrogate-modeling benchmarks replace expensive traditional simulations with neural

surrogates. CFD Bench trains neural operators to emulate CFD, evaluated with  $L^2$  error and MAE [41]. The Well provides 16 different datasets across multiple scientific domains to evaluate physical surrogate models [47].

*i) Reasoning & Generalization:* Reasoning benchmarks probe a model’s ability to perform logical inference and extrapolate. MedQA presents multiple-choice medical board-exam questions, measuring diagnostic accuracy [30]. FrontierMath tests advanced mathematical reasoning across category theory, number theory, and algebraic geometry, evaluated by accuracy [37]. AIME evaluates high-school mathematics problem solving, measured by accuracy [38].

### C. Emerging Computing Patterns

*1) Computing Motifs:* Computational motifs across different benchmarks are somewhat more difficult to identify, as they tend to be a property of a given implementation rather than the benchmark itself. Nevertheless, some benchmarks are broadly categorizable based on specific system constraints, task architectures, and/or typical computational patterns that a benchmark exhibits. While we do attempt to do so in this work, the categories we identify can differentiate benchmarks by the typical computational area the task is bound by, such as latency, throughput, utilization, or memory. It is also possible that a benchmark can be bound by multiple of these categories.

- **Latency Bound** Latency-bound benchmarks are those where there is either a fixed latency constraint, or the objective is simply to produce an output as fast as possible.
- **Memory Bound** Memory-bound benchmarks typically involve tasks that require working with large amounts of data all at once.
- **Throughput Bound** Throughput-bound benchmarks have the need to be able to consume and produce data at very high rates, sometimes independent of the total system latency, typically due to bandwidth constraints present in a system.
- **Utilization Bound** Utilization-bound benchmarks are benchmarks that have a fixed limit on the amount of computational power they can use, or where the task is traditionally a computationally expensive task, a classical example of which is protein folding simulations.

*2) Clustering Algorithm:* Depending on the primary user’s area of expertise and goals when using MLCommons Science benchmarks, they will likely have different priorities in what they view as most important in a set of benchmarks. Some users will also have additional subgoals beyond a specific benchmark category. For example, ML practitioners typically prioritize the accuracy (e.g., error) of algorithms. Similarly, others care about how workloads apply to different compute scenarios (e.g., edge computing) and domain scientists (similar to **User 2**) prioritize having a

variety of representative workloads for their area. Alternatively, **User 2** may be interested in multiple different implementations for the same benchmark (*benchmark solutions*). Conversely, computer systems researchers and hardware vendors (similar to **User 3** and the **Computing Patterns**) prioritize ensuring that benchmarks cover a wide variety of architectural or system features (e.g., compute- and memory-intensive workloads, workloads that utilize different architectural features like AMD’s Matrix Core Engines [48], [49] or NVIDIA’s TensorCores [50], [51], [52], [53]. Notably, unlike some other users, systems researchers and vendors may be less interested in covering different domains – unless those domains have applications with different computational patterns they need to optimize their systems for. More broadly, these needs often can be categorized into motifs (Section IV-B).

Naturally, these conflicting goals and motifs make it challenging for a single benchmark ontology to be fully representative for all stakeholders. Thus, instead of attempting to design a single benchmark ontology that is representative for all stakeholders, we instead propose to design an algorithm that views the priorities of the different stakeholders as *axes* in a multi-dimensional space. Given input from the user about what their priorities are, our algorithm presents a representative subset of benchmarks or benchmark solutions of interest given those priorities.

Accordingly, we propose to create a workload classifier that clusters the benchmarks, including multiple benchmark solutions for the same application, in our ontology. By clustering the benchmarks solutions, we can identify workloads with similar behavioral characteristics without requiring expensive system-level profiling for every new application or input per application. Such a classification scheme can be then used to obtain information about a new workload’s behavior, including accuracy (error), compute platform, domain, power consumption, and performance. To achieve this, we will integrate profiling and clustering to construct a comprehensive view of a given workloads behavior in line with Figure 1.

For each axis (e.g., dataset, specification, software, documentation, reference solution, and performance metrics) we will use the rating system from Section III and user input about which axes are important to them to develop a clustering algorithm that emits the benchmark(s) most relevant to them. Given this information, we will leverage it to create a *selection bar* for the MLCommons website similar to Figure 1. Specifically, a user will input how important each of the five axes are. This information will be passed to our algorithm, which will generate clustering information given these preferences and output the appropriate, representative workloads of interest. This information will then be displayed on the MLCommons website for a user to use and identify their desired benchmark(s).

The data we need to gather for each axes will vary. For example, for performance and power metrics, we will

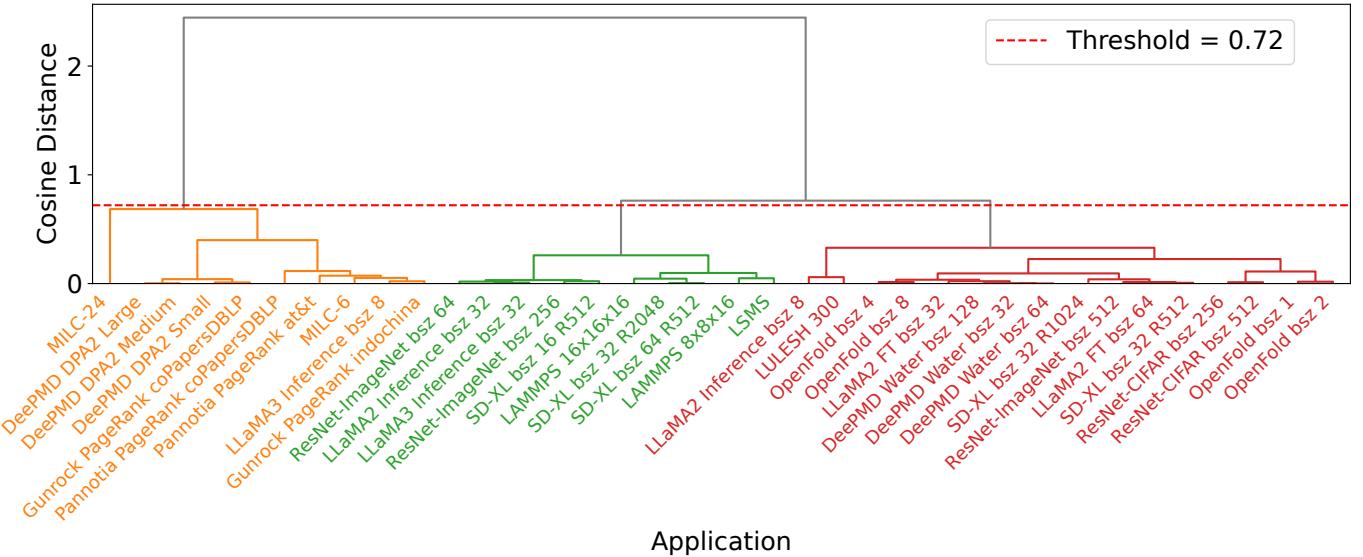


Fig. 3. Dendrogram showing hierarchical clustering based on power distributions of workloads. The clusters are labeled Low-power (orange), High-power (green), and Mixed (red), respectively, based on their power distribution.

leverage hardware profiler metrics (e.g., GPU resource utilization and instantaneous power metrics). Other metrics like documentation and reference solution will utilize static, offline information. Once we have gathered this information, we will effectively have an  $N$  dimensional vector per workload, where  $N$  represents the number of axes. A smaller  $N$  creates coarser grained bins which are easier to group but may group dissimilar applications together if their coarse, aggregate values look similar. Conversely, a larger  $N$  creates finer grained bins, enabling our approach to distinguish fine-grained power level variations, but also more aggressively separating workloads.

To identify similarities in behavior across workloads represented by these vectors, we apply **Hierarchical Clustering** to the collected feature vectors. Hierarchical clustering is an agglomerative classification, starting with leaf nodes where every application is its own cluster, and merging similar pairs of clusters together. To do this, the technique groups clusters using a distance metric such as Euclidean or cosine distance. We use cosine distance to compute pairwise distance between feature vectors since prior work [54], [55] suggests Euclidean distances are biased towards the magnitude of the feature vectors rather than the direction, and cosine similarity does not suffer from this bias. Like Fathom [56], we also propose to use hierarchical clustering to group workloads with similar dynamic power signatures.

Figure 3 shows an example of how we propose to do our clustering. Here, we analyze a variety of modern graph analytics, HPC, HPC+ML, and ML GPU workloads [3], [7], [9], [57], [58], [59], [60], [61], [62], [63], [64], [65], including different inputs for some of these workloads. For simplicity, here we refer to a given workload with different inputs as a

separate benchmark solution. We profiled each benchmark solution’s power consumption using NVIDIA’s `nvidia-smi` and `NVML` tools. Given these values, we then calculate the cosine distance between each pair of benchmark solutions. Figure 3’s dendrogram then hierarchically clusters the benchmark implementations based on these values. The dendrogram’s y-axis indicates the cosine distance between two workloads. A cosine distance of 0 indicates perfectly aligned power consumption while a larger cosine distance indicates that benchmark solutions are farther apart. Thus, we can slice the dendrogram at suitable cosine distances, to obtain  $K$  different groups or clusters. For example, slicing the dendrogram at a cosine distance of 0.72 yields three distinct power behavior groups ( $K=3$ ). While this example only uses a single type of information, and different dendograms can be generated by using different cosine distances, it nonetheless serves as a useful demonstration of how our clustering can be used to produce subsets of benchmark solutions for Figure 1.

## V. SUMMARY

This work presents the MLCommons Science Benchmark Ontology, a cohesive, community-driven benchmark ontology that extends the MLCommons ecosystem to scientific machine learning workloads. The ontology expands on existing domain-specific efforts, such as the XAI-BENCH synthetic explainability datasets, the FastML Science Benchmarks for low-latency physics workloads, the PDEBench surrogate modeling benchmarks, and the SciMLBench conceptual framework, into a single, extensible ontology.

In this work, we introduce a definition and ontology of scientific machine learning benchmarks, where benchmarks are classified and mapped to their scientific domain and

machine learning task type (see II). New benchmarks are added through an open submission workflow overseen by the MLCommons Science Working Group. Each submission is evaluated against a six category rating rubric (Software Environment, Problem Specification, Dataset, Performance Metrics, Reference Solution, Documentation) that assigns an overall rating and potential endorsement. The six category scoring framework enables stakeholders, researchers, domain scientists, and hardware vendors to identify representative subsets of benchmarks that align with their specific priorities. The ontology supports adding new scientific domains, AI/ML motifs, and computing motifs as it expands, and we additionally illustrate an example method for automated computing workload classification based on power and utilization characteristics.

Collectively, these elements transform a set of isolated domain-specific benchmark initiatives into a standardized, scalable ontology that promotes reproducibility, community participation, and broad applicability across the scientific machine learning landscape. This work positions MLCommons Scientific Benchmark Ontology as a reference point for future extensions and cross-domain benchmarking efforts.

An up-to-date, full report of all benchmarks in the ontology, as well as a searchable and filterable website, is available at <https://mlcommons-science.github.io/benchmark/>.

#### ACKNOWLEDGMENTS

This manuscript has been authored by FermiForward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

This work was supported by DOE ASCR Microelectronics Science Research Center Projects, BIA. This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357.

This work was partially supported by the NSF POSE Phase II Award 2303700.

This work has was supported by the DE-SC0023452; FAIR Surrogate Benchmarks and NSF - 2504401- OAC Core: Data I/O CyberInfrastructure grants.

The portion of this work done at UW-Madison is supported in part by NSF grant CNS-2312688; Advanced Micro Devices (AMD), Inc. under the AMD AI & HPC Cluster Program; and by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC-0026036.

#### APPENDIX

TABLE II: Table of all current benchmarks within the collection

Citation	Domain	AI/ML Motif	Average Rating
[31]	Climate & Earth Science	Sequence Prediction/Forecasting	<b>5.00</b>
[31]	Climate & Earth Science	Regression	<b>5.00</b>
[31]	Climate & Earth Science	Regression	<b>5.00</b>
[34]	Climate & Earth Science	Classification	<b>5.00</b>
[34]	Climate & Earth Science	Sequence Prediction/Forecasting	<b>5.00</b>
[34]	Biology & Medicine	Classification	<b>5.00</b>
[34]	Materials Science	Classification	<b>5.00</b>
[66]	Computational Science & AI	Reasoning & Generalization	<b>4.83</b>
[21]	Chemistry	Generative	<b>4.83</b>
[42]	Biology & Medicine	Sequence Prediction/Forecasting	<b>4.83</b>
[67]	Climate & Earth Science	Reasoning & Generalization	<b>4.67</b>
[35]	Computational Science & AI	Generative	<b>4.50</b>
[68]	High Energy Physics	Generative	<b>4.50</b>
[69]	Computational Science & AI, Climate & Earth Science, Mathematics	Regression	<b>4.50</b>
[70]	Climate & Earth Science	Regression	<b>4.50</b>
[70]	Climate & Earth Science	Classification	<b>4.50</b>
[70]	Climate & Earth Science	Classification	<b>4.50</b>
[70]	Climate & Earth Science	Anomaly Detection	<b>4.50</b>
[71]	Computational Science & AI	Multimodal Reasoning	4.42
[72]	Biology & Medicine	Classification	4.33
[72]	Biology & Medicine	Classification	4.33
[72]	Biology & Medicine	Classification	4.33
[28]	Climate & Earth Science	Classification	4.33
[28]	Climate & Earth Science	Reasoning & Generalization	4.33
[73]	High Energy Physics	Classification	4.33
[73]	High Energy Physics	Classification	4.33
[73]	Biology & Medicine	Classification	4.33
[73]	Materials Science	Regression	4.33
[22], [74], [75], [76]	Chemistry, Materials Science	Regression	4.17
[11]	High Energy Physics	Classification	4.17
[14]	High Energy Physics	Generative	4.17
[36]	High Energy Physics	Regression	4.17
[36]	Climate & Earth Science	Classification	4.17
[36]	Chemistry	Regression	4.17
[36]	Biology & Medicine	Sequence Prediction/Forecasting	4.17
[43]	High Energy Physics	Anomaly Detection	4.17
[26]	Materials Science	Regression	4.17
[26]	Materials Science	Generative	4.17
[27]	Biology & Medicine	Reasoning & Generalization	4.00
[27]	Biology & Medicine	Classification	4.00
[47]	Biology & Medicine, Computational Science & AI, High Energy Physics	Sequence Prediction/Forecasting	4.00
[77]	Computational Science & AI	Reasoning & Generalization	3.83
[33]	Climate & Earth Science	Multimodal Reasoning	3.83
[78]	Biology & Medicine, Chemistry, High Energy Physics	Reasoning & Generalization	3.83
[39]	Mathematics	Reasoning & Generalization	3.83
[40]	Mathematics	Reasoning & Generalization	3.83
[20]	High Energy Physics	Classification	3.83
[23]	Chemistry, Materials Science	Regression	3.83
[32]	Climate & Earth Science	Anomaly Detection	3.83
[29]	Biology & Medicine	Regression	3.83
[79]	Chemistry	Generative	3.75
[79]	Chemistry	Regression	3.75
[79]	Chemistry	Regression	3.75
[46]	Computational Science & AI	Multimodal Reasoning	3.67
[80]	Biology & Medicine, High Energy Physics, Chemistry	Reasoning & Generalization	3.67
[30]	Biology & Medicine	Reasoning & Generalization	3.50
[81]	Computational Science & AI	Classification	3.50
[41]	Mathematics	Regression	3.33
[82]	Materials Science, High Energy Physics, Biology & Medicine, Chemistry, Climate & Earth Science	Reasoning & Generalization	3.33
[19]	High Energy Physics	Classification	3.33
[83]	High Energy Physics	Anomaly Detection	3.33
[84]	Computational Science & AI	Classification	3.17

Continued on next page

<b>Citation</b>	<b>Domain</b>	<b>AI/ML Motif</b>	<b>Average Rating</b>
[17]	High Energy Physics	Classification	3.17
[85]	Materials Science	Classification	3.17
[86]	Materials Science	Classification	3.17
[18], [87]	High Energy Physics	Reinforcement Learning/Control	3.00
[44]	High Energy Physics	Classification	3.00
[88]	Biology & Medicine	Anomaly Detection	3.00
[89]	High Energy Physics	Anomaly Detection	2.83
[37]	Mathematics	Reasoning & Generalization	2.50
[38]	Mathematics	Reasoning & Generalization	2.33
[45]	High Energy Physics	Anomaly Detection	2.17
[24]	Materials Science	Regression	1.92
[90]	High Energy Physics	Classification	1.50

## REFERENCES

[1] S. Che, B. M. Beckmann, S. K. Reinhardt, and K. Skadron, "Pan-notia: Understanding Irregular GPGPU Graph Applications," in *IEEE International Symposium on Workload Characterization*, ser. IISWC, Sept 2013, pp. 185–195.

[2] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S.-H. Lee, and K. Skadron, "Rodinia: A Benchmark Suite for Heterogeneous Computing," in *2009 IEEE International Symposium on Workload Characterization*, ser. IISWC, 2009, pp. 44–54.

[3] Lawrence Livermore National Labs, "CORAL-2 Benchmarks," <https://asc.llnl.gov/coral-2-benchmarks>, 2020.

[4] P. Mattson, C. Cheng, C. Coleman, G. Diamos, P. Micikevicius, D. A. Patterson, H. Tang, G. Wei, P. Bailis, V. Bittorf, D. Brooks, D. Chen, D. Dutta, U. Gupta *et al.*, "MLPerf Training Benchmark," *CoRR*, vol. abs/1910.01500, 2019. [Online]. Available: <http://arxiv.org/abs/1910.01500>

[5] Oak Ridge National Labs, "OLCF-6 Benchmarks," <https://www.olcf.ornl.gov/draft-olcf-6-technical-requirements/benchmarks/>, 2024.

[6] NERSC, "NERSC-10 Benchmark Suite," <https://www.nersc.gov/systems/nersc-10/benchmarks/>, 3 2024.

[7] V. J. Reddi, C. Cheng, D. Kanter, P. Mattson, G. Schmuelling, C.-J. Wu, B. Anderson, M. Breughe, M. Charlebois, W. Chou, R. Chukka, C. Coleman, S. Davis, P. Deng *et al.*, "MLPerf Inference Benchmark," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture*, ser. ISCA, 2020, pp. 446–459.

[8] Standard Performance Evaluation Corporation, "Standard Performance Evaluation Corporation Benchmarks," <https://www.spec.org/>, 2025.

[9] Y. Wang, Y. Pan, A. Davidson, Y. Wu, C. Yang, L. Wang, M. Osama, C. Yuan, W. Liu, A. T. Riffel, and J. D. Owens, "Gunrock: GPU Graph Analytics," *ACM Trans. Parallel Comput.*, vol. 4, no. 1, Aug 2017. [Online]. Available: <https://doi.org/10.1145/3108140>

[10] Y. Liu, S. Khandagale, C. White, and W. Neiswanger, "Synthetic benchmarks for scientific research in explainable machine learning," in *Advances in Neural Information Processing Systems Datasets Track*, 2021.

[11] J. Duarte, N. Tran, B. Hawks, C. Herwig, J. Muhizi, S. Prakash, and V. J. Reddi, "Fastml science benchmarks: Accelerating real-time scientific edge machine learning," 2022. [Online]. Available: <https://arxiv.org/abs/2207.07958>

[12] M. Takamoto, T. Praditia, R. Leiteritz, D. MacKinlay, F. Alessiani, D. Pflüger, and M. Niepert, "Pdebenc: An extensive benchmark for scientific machine learning," <https://arxiv.org/abs/2210.07182>, 2022, arXiv preprint arXiv:2210.07182.

[13] J. Thiyyagalingam, M. Shankar, G. Fox, and T. Hey, "Scientific Machine Learning Benchmarks," *Nature Reviews Physics*, vol. 4, no. 6, pp. 413–420, 2022.

[14] J. Duarte, N. Tran, B. Hawks, C. Herwig, J. Muhizi, S. Prakash, and V. J. Reddi, "Fastml science benchmarks: Accelerating real-time scientific edge machine learning," 2022. [Online]. Available: <https://arxiv.org/abs/2207.07958>

[15] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne *et al.*, "The fair guiding principles for scientific data management and stewardship," *Scientific data*, vol. 3, no. 1, pp. 1–9, 2016.

[16] G. von Laszewski, R. Shiraishi, A. Krishnan, N. Tran, B. Hawks, and G. C. Fox, "Mlcommons science working group ai benchmarks collection," Github, Jul. 2025. [Online]. Available: <https://mlcommons-science.github.io/benchmark/benchmarks.pdf>

[17] P. Odagiu, Z. Que, J. Duarte, J. Haller, G. Kasieczka, A. Lobanov, V. Loncar, W. Luk, J. Ngadiuba, M. Pierini, P. Rincke, A. Seksaria, S. Summers, A. Sznajder *et al.*, "Ultrafast jet classification on fpgas for the hl-lhc," 2024. [Online]. Available: <https://arxiv.org/abs/2402.01876>

[18] J. Duarte, N. Tran, B. Hawks, C. Herwig, J. Muhizi, S. Prakash, and V. J. Reddi, "Fastml science benchmarks: Accelerating real-time scientific edge machine learning," 2022. [Online]. Available: <https://arxiv.org/abs/2207.07958>

[19] B. Parpillon, C. Syal, J. Yoo, J. Dickinson, M. Swartz, G. D. Guglielmo, A. Bean, D. Berry, M. B. Valentini, K. DiPetrillo, A. Badea, L. Gray, P. Maksimovic, C. Mills *et al.*, "Smart pixels: In-pixel ai for on-sensor data filtering," 2024. [Online]. Available: <https://arxiv.org/abs/2406.14860>

[20] J. Weitz, D. Demler, L. McDermott, N. Tran, and J. Duarte, "Neural architecture codesign for fast physics applications," 2025. [Online]. Available: <https://arxiv.org/abs/2501.05515>

[21] Y. Fang, N. Zhang, Z. Chen, L. Guo, X. Fan, and H. Chen, "Domain-agnostic molecular generation with chemical feedback," 2024. [Online]. Available: <https://arxiv.org/abs/2301.11259>

[22] L. Chanussot, A. Das, S. Goyal, T. Lavril, M. Shuaibi, M. Riviere, K. Tran, J. Heras-Domingo, C. Ho, W. Hu, A. Palizhati, A. Sriram, B. Wood, J. Yoon *et al.*, "The open catalyst 2020 (oc20) dataset and community challenges," *ACS Catalysis*, vol. 11, no. 10, pp. 6059–6072, 2021. [Online]. Available: <https://pubs.acs.org/doi/10.1021/acscatal.0c04525>

[23] K. Khrabrov, A. Ber, A. Tsybin, K. Ushenin, E. Rumiantsev, A. Telepov, D. Protasov, I. Shenbin, A. Alekseev, M. Shirokikh, S. Nikolenko, E. Tutubalina, and A. Kadurin, "Delta-squared dft: A universal quantum chemistry dataset of drug-like molecules and a benchmark for neural network potentials," 2024. [Online]. Available: <https://arxiv.org/abs/2406.14347>

[24] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, "The materials project: A materials genome approach," *APL Materials*, vol. 1, no. 1, 2013. [Online]. Available: <https://materialsproject.org/>

[25] K. Choudhary, "Jarvis-leaderboard," [Online; accessed 2025-06-02]. [Online]. Available: <https://pages.nist.gov/jarvis-leaderboard/>

[26] P. Chen, L. Peng, R. Jiao, Q. Mo, Z. Wang, W. Huang, Y. Liu, and Y. Lu, "Learning superconductivity from ordered and disordered material structures," in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 108 902–108 928. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/c4e3b55ed4ac9ba52d7df1f8bddbbf4-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/c4e3b55ed4ac9ba52d7df1f8bddbbf4-Paper-Datasets_and_Benchmarks_Track.pdf)

[27] E. Luo, J. Jia, Y. Xiong, X. Li, X. Guo, B. Yu, L. Wei, and X. Zhang, "Benchmarking ai scientists in omics data-driven biological research," 2025. [Online]. Available: <https://arxiv.org/abs/2505.08341>

[28] K. X. Nguyen, F. Qiao, A. Trembanis, and X. Peng, "Seafloorai: A large-scale vision-language dataset for seafloor geological survey," 2024. [Online]. Available: <https://arxiv.org/abs/2411.00172>

[29] R. E. Peterson, A. Tanelus, C. Ick, B. Mimica, N. Francis, V. J. Ivan, A. Choudhri, A. L. Falkner, M. Murthy, D. M. Schneider, D. H. Sanes, and A. H. Williams, "Vocal call locator benchmark (vcl) for localizing rodent vocalizations from multi-channel audio," in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 106 370–106 382. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/c00d37d6b04d73b870b963a4d70051c1-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/c00d37d6b04d73b870b963a4d70051c1-Paper-Datasets_and_Benchmarks_Track.pdf)

[30] D. Jin, E. Pan, N. Oufatolle, W.-H. Weng, H. Fang, and P. Szolovits, "What disease does this patient have? a large-scale open domain question answering dataset from medical exams," 2020. [Online]. Available: <https://arxiv.org/abs/2009.13081>

[31] T. Nguyen, J. Jewik, H. Bansal, P. Sharma, and A. Grover, "Climatelearn: Benchmarking machine learning for weather and climate modeling," 2023. [Online]. Available: <https://arxiv.org/abs/2307.01909>

[32] E. G. Campolongo, Y.-T. Chou, E. Govorkova, W. Bhimji, W.-L. Chao, C. Harris, S.-C. Hsu, H. Lapp, M. S. Neubauer, J. Namayanja, A. Subramanian, P. Harris, A. Anand, D. E. Carolyn *et al.*, "Building machine learning challenges for anomaly detection in science," 2025. [Online]. Available: <https://arxiv.org/abs/2503.02112>

[33] J. Roberts, K. Han, and S. Albanie, "Satin: A multi-task meta-dataset for classifying satellite imagery using vision-language models," *ICCV Workshop: Towards the Next Generation of Computer Vision Datasets*, 3 2023.

[34] J. Thiyagalingam, G. von Laszewski, J. Yin, M. Emani, J. Papay, G. Barrett, P. Luszczek, A. Tsaris, C. Kirkpatrick, F. Wang, T. Gibbs, V. Vishwanath, M. Shankar, G. Fox, and T. Hey, "Ai benchmarking for science: Efforts from the mlcommons science working group," in *High Performance Computing. ISC High Performance 2022 International Workshops*, H. Anzt, A. Bienz, P. Luszczek, and M. Baboulin, Eds. Cham: Springer International Publishing, 2022, pp. 47–64.

[35] M. Tian, L. Gao, S. D. Zhang, X. Chen, C. Fan, X. Guo, R. Haas, P. Ji, K. Krongchon, Y. Li, S. Liu, D. Luo, Y. Ma, H. Tong *et al.*, "Scicode: A research coding benchmark curated by scientists," 2024. [Online]. Available: <https://arxiv.org/abs/2407.13168>

[36] S. Farrell, M. Emani, J. Balma, L. Drescher, A. Drozd, A. Fink, G. Fox, D. Kanter, T. Kurth, P. Mattson, D. Mu, A. Ruhela, K. Sato, K. Shirahata *et al.*, "Mlperf hpc: A holistic benchmark suite for scientific machine learning on hpc systems," 2021. [Online]. Available: <https://arxiv.org/abs/2110.11466>

[37] E. Glazer, E. Erdil, T. Besiroglu, D. Chicharro, E. Chen, A. Gunning, C. F. Olsson, J.-S. Denain, A. Ho, E. de Oliveira Santos, O. Järvinimiemi, M. Barnett, R. Sandler, M. Vrzala *et al.*, "Frontiermath: A benchmark for evaluating advanced mathematical reasoning in ai," 2024. [Online]. Available: <https://arxiv.org/abs/2411.04872>

[38] TBD, "Aime," Mar. 2025, [Online accessed 2025-06-24]. [Online]. Available: <https://www.vals.ai/benchmarks/aime-2025-03-13>

[39] H. Lightman, V. Kosaraju, Y. Burda, H. Edwards, B. Baker, T. Lee, J. Leike, J. Schulman, I. Sutskever, and K. Cobbe, "Let's verify step by step," *arXiv preprint arXiv:2305.20050*, 2023.

[40] N. Mudur, H. Cui, S. Venugopalan, P. Raccuglia, M. P. Brenner, and P. Norgaard, "Feabench: Evaluating language models on multiphysics reasoning ability," 2025. [Online]. Available: <https://arxiv.org/abs/2504.06260>

[41] Y. Luo, Y. Chen, and Z. Zhang, "Cfdbench: A large-scale benchmark for machine learning methods in fluid dynamics," 2024. [Online]. Available: <https://arxiv.org/abs/2310.05963>

[42] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec, "Open graph benchmark: Datasets for machine learning on graphs," 2021. [Online]. Available: <https://arxiv.org/abs/2005.00687>

[43] E. G. Campolongo, Y.-T. Chou, E. Govorkova, W. Bhimji, W.-L. Chao, C. Harris, S.-C. Hsu, H. Lapp, M. S. Neubauer, J. Namayanya, A. Subramanian, P. Harris, A. Anand, D. E. Carlyn *et al.*, "Building machine learning challenges for anomaly detection in science," 2025. [Online]. Available: <https://arxiv.org/abs/2503.02112>

[44] J. Kvapil, G. Borca-Tasciuc, H. Bossi, K. Chen, Y. Chen, Y. C. Morales, H. D. Costa, C. D. Silva, C. Dean, J. Durham, S. Fu, C. Hao, P. Harris, O. Hen *et al.*, "Intelligent experiments through real-time ai: Fast data processing and autonomous detector control for sphenix and future eic detectors," 2025. [Online]. Available: <https://arxiv.org/abs/2501.04845>

[45] M. Khan, S. Krave, V. Marinozzi, J. Ngadiuba, S. Stoynev, and N. Tran, "Benchmarking and interpreting real time quench detection algorithms," in *Fast Machine Learning for Science Conference 2024*. Purdue University, IN: indico.cern.ch, Oct. 2024. [Online]. Available: [https://indico.cern.ch/event/1387540/contributions/6153618/attachments/2948441/5182077/fast\\_ml\\_magets\\_2024\\_final.pdf](https://indico.cern.ch/event/1387540/contributions/6153618/attachments/2948441/5182077/fast_ml_magets_2024_final.pdf)

[46] X. Zhong, Y. Gao, and S. Gururangan, "Spiqa: Scientific paper image question answering," 2024. [Online]. Available: <https://arxiv.org/abs/2407.09413>

[47] R. Ohana, M. McCabe, L. Meyer, R. Morel, F. J. Agocs, M. Beneitez, M. Berger, B. Burkhardt, S. B. Dalziel, D. B. Fielding, D. Fortunato, J. A. Goldberg, K. Hirashima, Y.-F. Jiang *et al.*, "The well: a large-scale collection of diverse physics simulations for machine learning," in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 44 989–45 037. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/4f9a5acd91ac76569f2fe291b1f4772b-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/4f9a5acd91ac76569f2fe291b1f4772b-Paper-Datasets_and_Benchmarks_Track.pdf)

[48] G. H. Loh, M. J. Schulte, M. Ignatowski, V. Adhinarayanan, S. Aga, D. Aguren, V. Agrawal, A. M. Aji, J. Alsop, P. Bauman, B. M. Beckmann, M. V. Beigi, S. Blagodurov, T. Boraten *et al.*, "A Research Retrospective on AMD's Exascale Computing Journey," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, ser. ISCA. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: <https://doi.org/10.1145/3579371.3589349>

[49] A. Smith, G. H. Loh, M. J. Schulte, M. Ignatowski, S. Naffziger, M. Mantor, N. Kalyanasundaram, V. Alla, N. Malaya, J. L. Greathouse, E. Chapman, and R. Swaminathan, "Realizing the AMD Exascale Heterogeneous Processor Vision : Industry Product," in *51st ACM/IEEE Annual International Symposium on Computer Architecture*, ser. ISCA. Piscataway, NJ, USA: IEEE, 2024, pp. 876–889. [Online]. Available: <https://doi.org/10.1109/ISCA59077.2024.00068>

[50] J. Choquette, O. Giroux, and D. Foley, "Volta: Performance and Programmability," *IEEE Micro*, vol. 38, no. 2, pp. 42–52, 2018.

[51] R. Krashinsky, O. Giroux, S. Jones, N. Stam, and S. Ramaswamy, "NVIDIA Ampere Architecture In-Depth," <https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth/>, 2020.

[52] NVIDIA, "NVIDIA H100 Tensor Core GPU Architecture," in *Proceedings of GPU Technology Conference*, ser. GTC, 2022.

[53] J. Choquette, "NVIDIA Hopper H100 GPU: Scaling Performance," *IEEE Micro*, vol. 43, no. 03, pp. 9–17, May 2023. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/MM.2023.3256796>

[54] Scikit Learn, "API Reference: sklearn.metrics: cosine distances," 2007. [Online]. Available: [https://scikit-learn.org/stable/modules/generated/sklearn.metrics.pairwise.cosine\\_distances.html](https://scikit-learn.org/stable/modules/generated/sklearn.metrics.pairwise.cosine_distances.html)

[55] P. Xia, L. Zhang, and F. Li, "Learning Similarity with Cosine Similarity Ensemble," *Information Sciences*, vol. 307, pp. 39–52, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0020025515001243>

[56] Adolf, Robert and Rama, Saketh and Reagen, Brandon and Wei, Gu-Yeon and Brooks, David, "Fathom: Reference workloads for modern deep learning methods," in *IEEE International Symposium on Workload Characterization*, ser. IISWC, IEEE. Los Alamitos, CA, USA: IEEE Computer Society, 9 2016, pp. 1–10. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/IISWC.2016.7581275>

[57] P. Mattson, C. Cheng, C. Coleman, G. Diamos, P. Micikevicius, D. Patterson, H. Tang, G.-Y. Wei, P. Bailis, V. Bittorf, D. Brooks, D. Chen, D. Dutta, U. Gupta *et al.*, "MLPerf Training Benchmark," *arXiv*, 2019. [Online]. Available: <https://arxiv.org/abs/1910.01500>

[58] G. Derevyanko, G. Lamoureux, C. Outeiral, T. Oda, F. Fuchs, S. P. Mahajan, J. Moult, J. Haas, P. Maragakis, T. Ruizmetov, and M. AlQuraishi, "OpenFold2: Replicating AlphaFold2 in the Dark," <https://lupoglaz.github.io/OpenFold2/>, 2023.

[59] O. R. L. C. Facility, *OLCF-6 Technical Requirements: Benchmarks*, 2025, accessed 2025. [Online]. Available: <https://www.olcf.ornl.gov/draft-olcf-6-technical-requirements/benchmarks/>

[60] H. Wang, L. Zhang, J. Han, and W. E, "DeePMD-kit: A deep learning package for many-body potential energy representation and molecular dynamics," *Computer Physics Communications*, vol. 228, pp. 178–184, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010465518300882>

[61] S. Che, B. M. Beckmann, S. K. Reinhardt, and K. Skadron, "Pannotta: Understanding Irregular GPGPU Graph Applications," in *IEEE International Symposium on Workload Characterization*, ser. IISWC, 9 2013, pp. 185–195.

[62] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen *et al.*, "Llama 2: Open Foundation and Fine-Tuned Chat Models," 2023. [Online]. Available: <https://arxiv.org/abs/2307.09288>

[63] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan, A. Yang, A. Fan, A. Goyal, A. Hartshorn *et al.*, "The

Llama 3 Herd of Models,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.21783>

[64] Markus Eisenbach and Jeff Larkin and Justin Lutjens and Steven Rennich and James H. Rogers, “GPU acceleration of the Locally Selfconsistent Multiple Scattering code for first principles calculation of the ground state and statistical physics of materials,” *Computer Physics Communications*, vol. 211, pp. 2–7, 2017, high Performance Computing for Advanced Modeling and Simulation of Materials. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010465516301953>

[65] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in’t Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen *et al.*, “Lammps-a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales,” *Computer Physics Communications*, vol. 271, p. 108171, 2022.

[66] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord, “Think you have solved question answering? try arc, the ai2 reasoning challenge,” *arXiv:1803.05457v1*, 2018.

[67] H. Zhang, J. Sun, R. Chen, W. Liu, Z. Yuan, X. Zheng, Z. Wang, Z. Yang, H. Yan, H.-S. Zhong, X. Wang, W. Ouyang, F. Yang, and N. Dong, “Empowering and assessing the utility of large language models in crop science,” in *The Thirty-eighth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. [Online]. Available: <https://openreview.net/forum?id=hMj6jZ6JWU>

[68] C. Krause, M. F. Giannelli, G. Kasieczka, B. Nachman, D. Salamani, D. Shih, A. Zaborowska, O. Amram, K. Borras, M. R. Buckley, E. Buhmann, T. Buss, R. P. D. C. Cardoso, A. L. Caterini *et al.*, “Calochallenge 2022: A community challenge for fast calorimeter simulation,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.21611>

[69] M. Takamoto, T. Praditia, R. Leiteritz, D. MacKinlay, F. Alesiani, D. Pflüger, and M. Niepert, “Pdebench: An extensive benchmark for scientific machine learning,” 2024. [Online]. Available: <https://arxiv.org/abs/2210.07182>

[70] Y. Wang, T. Wang, Y. Zhang, H. Zhang, H. Zheng, G. Zheng, and L. Kong, “Urbanbatalayer: A unified data pipeline for urban science,” in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 7296–7310. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/0db7f135f6991e8cec5e516ecc66bfba-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/0db7f135f6991e8cec5e516ecc66bfba-Paper-Datasets_and_Benchmarks_Track.pdf)

[71] S. Pramanick, R. Chellappa, and S. Venugopalan, “Spiqa: A dataset for multimodal question answering on scientific papers,” 2025. [Online]. Available: <https://arxiv.org/abs/2407.09413>

[72] A. Karargyris, R. Umeton, M. J. Sheller, A. Aristizabal, J. George, A. Wuest, S. Pati, H. Kassem, M. Zenk, U. Baid, P. Narayana Moorthy, A. Chowdhury, J. Guo, S. Nalawade *et al.*, “Federated benchmarking of medical artificial intelligence with medperf,” *Nature Machine Intelligence*, vol. 5, no. 7, pp. 799–810, Jul. 2023. [Online]. Available: <https://doi.org/10.1038/s42256-023-00652-2>

[73] D. Zou, S. Liu, S. Miao, V. Fung, S. Chang, and P. Li, “Gess: Benchmarking geometric deep learning under scientific applications with distribution shifts,” in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 92499–92528. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/a8063075b00168dc39bc81683619f1a8-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/a8063075b00168dc39bc81683619f1a8-Paper-Datasets_and_Benchmarks_Track.pdf)

[74] R. Tran, J. Lan, M. Shuaibi, B. M. Wood, S. Goyal, A. Das, J. Heras-Domingo, A. Kolluru, A. Rizvi, N. Shoghi, A. Sriram, F. Therrien, J. Abed, O. Voznyy *et al.*, “The open catalyst 2022 (oc22) dataset and challenges for oxide electrocatalysts,” *ACS Catalysis*, vol. 13, no. 5, pp. 3066–3084, 2023. [Online]. Available: <https://pubs.acs.org/doi/10.1021/acscatal.2c05426>

[75] L. Chanussot, A. Das, S. Goyal, T. Lavril, M. Shuaibi, M. Riviere, K. Tran, J. Heras-Domingo, C. Ho, W. Hu, A. Palizhati, A. Sriram, B. Wood, J. Yoon *et al.*, “Open catalyst 2020 (oc20) dataset and community challenges,” *ACS Catalysis*, vol. 11, no. 10, pp. 6059–6072, 2021. [Online]. Available: <https://doi.org/10.1021/acscatal.0c04525>

[76] R. Tran, J. Lan, M. Shuaibi, B. M. Wood, S. Goyal, A. Das, J. Heras-Domingo, A. Kolluru, A. Rizvi, N. Shoghi, A. Sriram, F. Therrien, J. Abed, O. Voznyy *et al.*, “The open catalyst 2022 (oc22) dataset and challenges for oxide electrocatalysts,” *ACS Catalysis*, vol. 13, no. 5, pp. 3066–3084, Feb. 2023. [Online]. Available: <http://dx.doi.org/10.1021/acscatal.2c05426>

[77] D. Hendrycks, C. Burns, and S. Kadavath, “Measuring Massive Multitask Language Understanding,” 2021. [Online]. Available: <https://arxiv.org/abs/2009.03300>

[78] D. Rein, B. L. Hou, and A. C. Stickland, “Gpqa: A graduate-level google-proof q and a benchmark,” 2023. [Online]. Available: <https://arxiv.org/abs/2311.12022>

[79] R. Bushuiev, A. Bushuiev, N. F. de Jonge, A. Young, F. Kretschmer, R. Samusevich, J. Heirman, F. Wang, L. Zhang, K. Dührkop, M. Ludwig, N. A. Haupt, A. Kalia, C. Brungs *et al.*, “Massspecgym: A benchmark for the discovery and identification of molecules,” in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 110010–110027. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/c6c31413d5c53b7d1c343c1498734b0f-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/c6c31413d5c53b7d1c343c1498734b0f-Paper-Datasets_and_Benchmarks_Track.pdf)

[80] D. Rein, B. L. Hou, A. C. Stickland, J. Petty, R. Y. Pang, J. Dirani, J. Michael, and S. R. Bowman, “Gpqa: A graduate-level google-proof q and a benchmark,” 2023. [Online]. Available: <https://arxiv.org/abs/2311.12022>

[81] G. D. Guglielmo, B. Du, J. Campos, A. Boltasseva, A. V. Dixit, F. Fahim, Z. Kudyshev, S. Lopez, R. Ma, G. N. Perdue, N. Tran, O. Yesilyurt, and D. Bowring, “End-to-end workflow for machine learning-based qubit readout with qick and hls4ml,” 2025. [Online]. Available: <https://arxiv.org/abs/2501.14663>

[82] H. Cui, Z. Shamsi, G. Cheon, X. Ma, S. Li, M. Tikhonovskaya, P. Norgaard, N. Mudur, M. Plomecka, P. Raccuglia, Y. Bahri, V. V. Albert, P. Srinivasan, H. Pan *et al.*, “Curie: Evaluating llms on multitask scientific long context understanding and reasoning,” 2025. [Online]. Available: <https://arxiv.org/abs/2503.13517>

[83] T. Aarrestad, E. Govorkova, J. Ngadiuba, E. Puljak, M. Pierini, and K. A. Wozniak, “Unsupervised new physics detection at 40 mhz: Training dataset,” 2021. [Online]. Available: <https://zenodo.org/record/5046389>

[84] J. Bowles, S. Ahmed, and M. Schuld, “Better than classical? the subtle art of benchmarking quantum machine learning models,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.07059>

[85] Z. Liu, H. Sharma, J.-S. Park, P. Kenesei, A. Miceli, J. Almer, R. Kettimuthu, and I. Foster, “Bragggnn: Fast x-ray bragg peak analysis using deep learning,” 2021. [Online]. Available: <https://arxiv.org/abs/2008.08198>

[86] S. Qin, J. Agar, and N. Tran, “Extremely noisy 4d-tem strain mapping using cycle consistent spatial transforming autoencoders,” in *AI for Accelerated Materials Design - NeurIPS 2023 Workshop*, 2023. [Online]. Available: <https://openreview.net/forum?id=7yt3N0o0W9>

[87] D. Kafkes and J. S. John, “Bostr: A dataset for accelerator control systems,” 2021. [Online]. Available: <https://arxiv.org/abs/2101.08359>

[88] E. G. Campolongo, Y.-T. Chou, E. Govorkova, W. Bhimji, W.-L. Chao, C. Harris, S.-C. Hsu, H. Lapp, M. S. Neubauer, J. Namayanja, A. Subramanian, P. Harris, A. Anand, D. E. Carolyn *et al.*, “Building machine learning challenges for anomaly detection in science,” 2025. [Online]. Available: <https://arxiv.org/abs/2503.02112>

[89] A. A. Abud, B. Abi, R. Acciarri, M. A. Acero, G. Adamov, D. Adams, M. Adinolfi, A. Aduszkiewicz, Z. Ahmad, J. Ahmed, T. Alion, S. A. Monsalve, M. Alrashed, C. Alt *et al.*, “Deep underground neutrino experiment (dune) near detector conceptual design report,” 2021. [Online]. Available: <https://arxiv.org/abs/2103.13910>

[90] Y. Wei, R. F. Forelli, C. Hansen, J. P. Levesque, N. Tran, J. C. Agar, G. D. Guglielmo, M. E. Mauel, and G. A. Navratil,

“Low latency optical-based mode tracking with machine learning deployed on fpgas on a tokamak,” 2024. [Online]. Available: <https://arxiv.org/abs/2312.00128>