

UC San Diego

UC San Diego Previously Published Works

Title

FAIR for AI: An interdisciplinary and international community building perspective.

Permalink

<https://escholarship.org/uc/item/71s2x9j9>

Journal

Scientific Data, 10(1)

Authors

Huerta, E

Blaiszik, Ben

Brinson, L

et al.

Publication Date

2023-07-26

DOI

10.1038/s41597-023-02298-6

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

OPEN
COMMENT

FAIR for AI: An interdisciplinary and international community building perspective

E. A. Huerta^{1,2}✉, Ben Blaiszik^{1,3}, L. Catherine Brinson⁴, Kristofer E. Bouchard^{5,6,7}, Daniel Diaz⁸, Caterina Doglioni^{9,10}, Javier M. Duarte⁸, Murali Emami¹¹, Ian Foster^{1,2}, Geoffrey Fox¹², Philip Harris¹³, Lukas Heinrich¹⁴, Shantenu Jha^{15,16}, Daniel S. Katz^{17,18,19,20}, Volodymyr Kindratenko^{17,18,19}, Christine R. Kirkpatrick²¹, Kati Lassila-Perini²², Ravi K. Madduri¹, Mark S. Neubauer^{17,19,23}, Fotis E. Psoomopoulos²⁴, Avik Roy^{17,23}, Oliver Rübél⁵, Zhizhen Zhao^{17,19} & Ruike Zhu¹⁸

A foundational set of findable, accessible, interoperable, and reusable (FAIR) principles were proposed in 2016 as prerequisites for proper data management and stewardship, with the goal of enabling the reusability of scholarly data. The principles were also meant to apply to other digital assets, at a high level, and over time, the FAIR guiding principles have been re-interpreted or extended to include the software, tools, algorithms, and workflows that produce data. FAIR principles are now being adapted in the context of AI models and datasets. Here, we present the perspectives, vision, and experiences of researchers from different countries, disciplines, and backgrounds who are leading the definition and adoption of FAIR principles in their communities of practice, and discuss outcomes that may result from pursuing and incentivizing FAIR AI research. The material for this report builds on the FAIR for AI Workshop held at Argonne National Laboratory on June 7, 2022.

¹Data Science and Learning Division, Argonne National Laboratory, Lemont, Illinois, 60439, USA. ²Department of Computer Science, University of Chicago, Chicago, Illinois, 60637, USA. ³Globus, University of Chicago, Chicago, Illinois, 60637, USA. ⁴Department of Mechanical Engineering and Materials Science, Duke University, Durham, North Carolina, 27708, USA. ⁵Scientific Data Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA. ⁶Biological Systems & Engineering, Lawrence Berkeley National Laboratory, Berkeley, California, 94720, USA. ⁷Helen Wills Neuroscience Institute, University of California Berkeley, Berkeley, California, 94720, USA. ⁸Department of Physics, University of California San Diego, La Jolla, California, 92093, USA. ⁹Lund University, Department of Physics, Box 118, 221 00, Lund, Sweden. ¹⁰School of Physics & Astronomy, The University of Manchester, Manchester, M13 9PL, UK. ¹¹Leadership Computing Facility, Argonne National Laboratory, Lemont, Illinois, 60439, USA. ¹²Biocomplexity Institute and Department of Computer Science, University of Virginia, Charlottesville, Virginia, 22904, USA. ¹³Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts, 02139, USA. ¹⁴Technical University Munich, Arcisstraße 21, 80333, München, Germany. ¹⁵Computational Science Initiative Brookhaven National Laboratory Upton, New York, 11973, USA. ¹⁶Electrical and Computer Engineering, Rutgers, The State University of New Jersey, Piscataway, New Jersey, 08854, USA. ¹⁷National Center for Supercomputing Applications, University of Illinois, Urbana-Champaign, Urbana, Illinois, 61801, USA. ¹⁸Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois, 61801, USA. ¹⁹Department of Electrical & Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, 61801, USA. ²⁰School of Information Sciences, University of Illinois at Urbana-Champaign, Urbana, Illinois, 61801, USA. ²¹San Diego Supercomputer Center, University of California San Diego, La Jolla, California, 92093, USA. ²²Helsinki Institute of Physics, University of Helsinki, P.O. Box 64, Helsinki, 00014, Finland. ²³Department of Physics, University of Illinois at Urbana-Champaign, Urbana, Illinois, 61801, USA. ²⁴Institute of Applied Biosciences, Centre for Research and Technology Hellas, Thessaloniki, 57001, Greece. ✉e-mail: eliu@anl.gov

Introduction

The production, collection, and curation of data require painstaking planning and the use of sophisticated experimental and computational facilities. In order to maximize the impact of these investments and create best practices that lead to scientific discovery and innovation, a diverse set of stakeholders defined a set of findable, accessible, interoperable, and reusable (FAIR) principles in 2016^{1,2}. The original intent was that these principles would apply seamlessly to data and all scholarly digital objects, including research software³, workflows⁴, and even domain-specific custom digital objects⁵. However, because they were specifically written in the context of data, it became clear over time that the original set of FAIR principles would have to be translated or reinterpreted for digital assets beyond data^{6,7}. This realization has led to initiatives that have proposed and/or developed practical FAIR definitions for research software and workflows, and more recently, for artificial intelligence (AI) models^{8,9}.

In this document, we provide an inclusive and diverse perspective of FAIR initiatives in Europe and the US through the lens of researchers that are leading the definition, implementation, and adoption of FAIR principles in a variety of disciplines. This community was brought together at the [FAIR for AI Workshop \(https://indico.cern.ch/event/1152431/\)](https://indico.cern.ch/event/1152431/) at Argonne National Laboratory on June 7, 2022. We believe that this document provides a factual, straightforward, and inspiring description of what FAIR initiatives have accomplished, what is being done and planned at the time of writing this document, and describes the end goals of these disparate initiatives. Most importantly, we hope that the ideas presented in this document serve as a motivator to reach convergence on what FAIR means, in practice, for AI research and innovation.

FAIR Initiatives

We have identified the following non-exhaustive list of FAIR initiatives:

- **FAIR4HEP: Findable, Accessible, Interoperable, and Reusable Frameworks for Physics-Inspired Artificial Intelligence in High Energy Physics** (<https://fair4hep.github.io>). Funded by the US Department of Energy (DOE). In this project an interdisciplinary team of physicists, computer, and AI scientists use high energy physics as the science driver to develop a FAIR framework that advances our understanding of AI, provides new insights to apply AI techniques, and provides an environment where novel approaches to AI can be explored.
- **ENDURABLE: Benchmark Datasets and AI models with queryable metadata** (<https://sites.google.com/lbl.gov/endurable/home>). Funded by DOE. The goal of this project is to provide the scientific and machine learning (ML) communities with robust, scalable, and extensible tools to share and rigorously aggregate diverse scientific data sets for training state-of-the-art ML models.
- **The Common Fund Data Ecosystem** (<https://commonfund.nih.gov/dataecosystem>). Funded by the US National Institutes of Health (NIH). An [online discovery platform \(https://app.nih-cfde.org\)](https://app.nih-cfde.org) that enables researchers to create and search across FAIR datasets to ask scientific and clinical questions from a single access point.
- **BioDataCatalyst** (<https://biodatacatalyst.nhlbi.nih.gov>). Funded by NIH. Construct and enhance annotated metadata for heart, lung, and blood datasets that comply with FAIR data principles.
- **Garden: A FAIR Framework for Publishing and Applying AI Models for Translational Research in Science, Engineering, Education, and Industry** (<https://thegardens.ai>). Funded by the US National Science Foundation (NSF). This project will reduce barriers to the use of AI methods and promote the nucleation of communities around specific FAIR datasets, methods, and AI models. Model Gardens will provide a repository for models where they can be linked to papers, testing metrics, known model limitations, and code, plus computing and data storage resources through tools such as the [Data and Learning Hub for Science](#)¹⁰, [funcX](#)¹¹ and [Globus](#)¹².
- **Braid: Data Flow Automation for Scalable and FAIR Science** (<https://anl-braid.github.io/braid/>). Funded by DOE. This project aims to enable researchers to define sets of flows that individually and collectively implement application capabilities while satisfying requirements for rapid response, high reconstruction fidelity, data enhancement, data preservation, model training, etc.
- **HPC-FAIR: A Framework Managing Data and AI Models for Analyzing and Optimizing Scientific Applications** (<https://hpc-fair.github.io/>). Funded by DOE. This multi-institutional project aims to develop a generic High Performance Computing data management framework^{13,14} to make both training data and AI models of scientific applications FAIR.
- **The FAIR Surrogate Benchmarks Initiative** (<https://sbi-fair.github.io>). Funded by DOE. The research develops AI surrogates and studies their key features and the software environment to support their use¹⁵ in simulation based research. They collaborate with [MLCommons \(https://mlcommons.org/en/\)](https://mlcommons.org/en/), a consortium including 62 companies that host the MLPerf benchmarks, including those for science^{16–18}, and mirror their processes in the computational science domain. This involves rich metadata involving models, datasets, and the logging of their use with machine and power characteristics recorded, requiring multiple ontologies to be developed with FAIR approaches.
- **The Materials Data Facility (MDF)** (<https://www.materialsdatafacility.org>). Funded by the National Institute of Standards and Technology (NIST) and the Center for Hierarchical Materials Design, the MDF^{19,20} aims at making materials data easily publishable, discoverable, and reusable while following and building upon the FAIR principles. To date, MDF has collected over 80 TB of materials data in nearly 1000 datasets. In particular, this effort enables publication of datasets with millions of files or datasets comprising TB of data, and seeks to automatically index the contents in ways that provide unique queryable interfaces to the datasets. Recently, these capabilities have been augmented via the [Foundry \(https://github.com/MLMI2-CSSI/foundry\)](https://github.com/MLMI2-CSSI/foundry) to provide access to well-described ML-ready datasets with just a few lines of Python code.

- **Neurodata Without Borders (NWB)** (<https://www.nwb.org/>). Funded by the NIH BRAIN Initiative. NWB is an interdisciplinary project to create a FAIR data standard for neurophysiology, providing neuroscientists with a common standard to share, archive, use, and build common analysis tools for neurophysiology data. More than just a data standard, NWB is at the heart of a growing software ecosystem for neurophysiology data, including data from intracellular and extracellular electrophysiology experiments, data from optical physiology experiments, and tracking and stimulus data. A growing number of neurophysiology data generated by NIH BRAIN Initiative research projects and others are available on the **DANDI** neurophysiology data archive.
- **Materials Research Data Alliance (MaRDA)** (<https://www.marda-alliance.org/>). MaRDA is an organization dedicated to helping to build community to promote open, accessible and interoperable data in materials science. MaRDA has held two virtual workshops reaching 300 attendees last year, and has helped researchers form independent working groups. In August 2022, MaRDA leadership was funded via the NSF Research Coordination Network program to significantly expand efforts to build a sustainable community around these topics, to build consensus in metadata requirements, to train next generation workforce in ML/AI for materials, to develop shared community benchmark challenges, to host convening and coordination events, and more.
- **PUNCH4NFDI** (<https://www.punch4nfdi.de>) is the German National Research Data Infrastructure consortium of particle, astro, astroparticle, hadron, and nuclear physics, representing about 9,000 scientists with a PhD in Germany from universities, the Max Planck Society, the Leibniz Association, and the Helmholtz Association. The prime goal of PUNCH4NFDI is to set up a federated and FAIR science data platform, offering the infrastructures and interfaces necessary for access to and use of data and computing resources of the involved communities and beyond.
- **ESCAPE** (<https://www.projectescape.eu>) is the European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures, funded from the European Union's Horizon 2020 research and innovation programme. The goal of ESCAPE is to address the critical questions of open science and long term reuse of data for science and for innovation, many of the greatest European scientific facilities in physics and astronomy have combined forces to make their data and software interoperable and open, committing to make the European Science Cloud a reality. ESCAPE is delivering two Science Projects to aid the prototyping the European Open Science Cloud (EOSC) within the **EOSC-Future** (<https://eoscfuture.eu>), another Horizon 2020-funded project. These Science Projects will advance the science, the FAIR data and the software tools needed for dark matter searches and multi-messenger astronomy for extreme universe phenomena such as gravitational waves.
- **Awesome Materials Informatics** (<https://github.com/tilde-lab/awesome-materials-informatics>) is an interdisciplinary and community building effort to assemble a list of a holistic set of tools and best practices for Materials Science, encompassing software and products, cloud simulation platforms, and standardization initiatives.

These initiatives suggest that researchers are developing methods, approaches, and tools from scratch to address specific needs in their communities of practice. Thus, it is timely and important to identify common needs and gaps in disparate disciplines, abstract them and then create commodity, generic tools that address similar challenges across fields. Interdisciplinary efforts of this nature may leverage work led by several research data consortia which tend to be more general, e.g., the **Research Data Alliance (RDA)** (<https://www.rd-alliance.org/>), the International Science Council's Committee on Data (**CODATA**) (<https://codata.org/>), and **GO FAIR** (<https://www.go-fair.org/>). This translational approach has been showcased in the context of scientific datasets²¹ and for AI models and datasets⁸. These recent efforts pose an important question: what is the optimal composition of interdisciplinary teams that may work together to create sufficiently generic solutions that may then be specialized down to specific disciplines and projects? As these interdisciplinary teams are assembled, and they work to define, implement, and then showcase how to adopt FAIR principles, it is critical to keep in mind that FAIR is not the goal *per se*, rather the science and innovation that such principles and best practices will enable. As well, FAIR is not a goal as much as a continual process.

In high energy physics (HEP), the experiments at the Large Hadron Collider at CERN are committed to bringing their data into the public domain²² through the **CERN Open Data portal** (<http://opendata.cern.ch/>). The CMS experiment has led the effort and, since 2014, made close to 3PB of research-level data public. Their availability opens unprecedented opportunities to process samples from original HEP experiment data for different AI studies. While the experiment data distribution follows FAIR principles, they remain complex, and their practical reusability has required further thoughts on the FAIR principles concretely applicable to software and workflows. Furthermore, the application of FAIR principles to data and AI models is important for the sustainability of HEP science and enhancing collaborative efforts with others, both inside and outside of the HEP domain. Ensuring that data and AI models are FAIR facilitates a better understanding of their content and context, enabling more transparent provenance and reproducibility^{23,24}. There is a strong connection between FAIRness and interpretability, as FAIR models facilitate comparisons of benchmark results across models²⁵ and applications of *post-hoc* explainable AI methods²⁶. As described in ref. ²⁷ data and AI models preserved in accordance with FAIR principles can facilitate education in data science and machine learning in several ways, such as interpretability of AI models, uncertainty quantification, and ease of access of data and models for key HEP use cases. In this way, they can be reliably reused to reproduce benchmark results for both research and pedagogical purposes. For instance, the detailed analysis of FAIR and AI-readiness of the CMS $H(b\bar{b})$ dataset in ref. ²¹ has explained how the FAIR readiness of this dataset has been useful in building ML exercises for open source courses on AI for HEP²⁸.

In the materials science domain, the importance of broad accessibility of research data on all materials and the transformative potential impact of FAIR data and use of data driven and AI approaches was recognized with the advent of the Materials Genome Initiative (MGI) in 2011²⁹, and with a recently released MGI Strategic Plan in late 2021³⁰. In the decade since the launch of MGI, the power of integrating data science with materials science has unleashed an explosion of productivity^{31,32}. Early adopters were computational materials scientists who launched a number of accessible data portals for hard materials and who have begun working together across the world on interoperability standards³³. Subsequently, significant efforts have been launched towards capturing FAIR experimental data and tackling the complexities of soft materials³⁴. In the last several years, **MaRDA** has developed and flourished with multiple workshops and working groups addressing issues of FAIR data and models across all aspects of materials science.

In the life sciences, AI is becoming increasingly popular as an efficient mechanism to extract knowledge and new insights from the vast amounts of data that are constantly generated. AI has the potential for transformative impact on the life sciences since almost half of global life sciences professionals are either using, or are interested in using, AI in some area of their work³⁵. This transition is clearly shown in the explosion of ML articles in life sciences over the past decade: from around 500 such publications in 2010 to approximately 14k publications in 2020, an exponential increase that does not show any signs of slowing down in the short-term³⁶. However, AI is not a one-solution-fits-all, nor a magic wand that can address any challenge in the life sciences and beyond. In this context, scientists pursuing domain aware AI applications may benefit from defining community-backed standards, such as the DOME recommendations³⁶, which were spearheaded by the **ELIXIR infrastructure** (<https://elixir-europe.org/>). As scientists adopt these guidelines, and prioritize openness in all aspects of their work processes, FAIR AI research will streamline the creation of AI applications that are trustworthy, high quality, reliable, and reproducible.

Towards a practical definition of FAIR for AI models. There are several efforts that aim to define, at a practical level, what FAIR means for scientific datasets and AI models. As a starting point, researchers have created platforms that provide, in an integrated and centralized manner, access to popular AI models and standardized datasets, e.g., the **Hugging Face** (<https://huggingface.co>) platform, and the **Data and Learning Hub for Science**¹⁰.

While these efforts are necessary and valuable, additional work is needed to leverage these AI models and datasets, and translate them for AI R&D in scientific applications. This is because state-of-the-art AI models become valuable tools for scientific discovery when they encode domain knowledge, and are capable of learning complex features and patterns in experimental datasets, which differ vastly from standardized datasets (ImageNet, Google's Open Images, xView, etc.). Creating scientific AI tools requires significant investments to produce, collect and curate experimental datasets, and then incorporate domain knowledge in the design, training and optimization of AI models. Often, this requires the development and deployment of distributed training algorithms in high performance computing (HPC) platforms to reduce time-to-insight^{37,38}, and the optimization of fully trained AI models for accelerated inference on HPC platforms and/or at the edge^{39,40}. How can this wealth of knowledge be leveraged, extended or seamlessly used by other researchers that face similar challenges in similar or disparate disciplines?

While peer-reviewed publications continue to be the main avenue to communicate advances in AI for science, researchers increasingly recognize that articles should also be linked to data, AI models, and scientific software needed to reproduce and validate data-driven scientific discovery. Doing so is in line with the norm in scientific machine learning, which is characterized by open access to state-of-the-art AI models and standardized datasets. This is one of the central aims in the creation of FAIR datasets and AI models, namely, to share knowledge, resources, and tools following best practices to accelerate and sustain discovery and innovation.

Several challenges, however, need to be addressed when researchers try to define, implement and adopt FAIR principles in practice. This is because there is a dearth of simple-to-follow guidelines and examples, and of lack of consistent metrics that indicate when the FAIRification of datasets and AI models has been done well or not, and how to improve. Furthermore, while the FAIR principles are simple to read, they can be difficult to implement, and work is needed to build consensus about what they mean in specific cases, how they can be met, and how implementation can be measured, not only for data but also for other types of digital objects, such as AI models and software. The need to integrate FAIR mechanisms throughout the research lifecycle has been noted⁴¹. Researchers are actively trying to address these gaps and needs in the context of datasets²¹ and AI models.

On the latter point, two recent studies^{8,9} have presented practical FAIR guidelines for AI models. Common themes in these studies encompass: 1) the need to define the realm of applicability of these principles in the AI R&D cycle, i.e., they consider AI models that have been fully trained and whose FAIRness is quantified for AI-driven inference; 2) the use of common software templates to develop and publish AI models, e.g., the template generator `cookiecutter4fair`⁴²; and 3) the use of modern computing environments and scientific data infrastructure to transcend barriers in hardware architectures and software to speak a common AI language. To ground these ideas, refs. ^{8,9} proposed definitions of a (FAIR) AI model, which we have slightly modified as follows: "an AI model comprises a computational graph and a set of parameters that can be expressed as scientific software that, combined with modern computing environments, may be used to extract knowledge or insights from experimental or synthetic datasets that describe processes, systems, etc. An AI model is Findable when a digital object identifier (DOI) can direct a human or machine to a digital resource that contains the model, its metadata, instructions to run the model on a data sample, and uncertainty quantification metrics to evaluate the soundness of AI predictions; it is Accessible when it and its metadata may be readily downloaded or invoked by humans or machines via standardized protocols to run inference on data samples; it is Interoperable when it can seamlessly interact with other models, data, software, and hardware architectures; and it is Reusable

when it can be used by humans, machines and other models to reproduce its expected inference capabilities, and provide reliable uncertainty quantification metrics when processing datasets that differ from those originally used to create it and quantify its performance”.

Furthermore, the work presented by Ravi *et al.*⁸, emphasizes the need to create computational frameworks that link FAIR and AI-ready datasets (produced by scientific facilities or large scale simulations and either hosted at data facilities or broadcast to supercomputing centers) with FAIR AI models (hosted at model hubs), and that can leverage computing environments (e.g., supercomputers, AI-accelerator machines, edge computing devices, and the cloud) to automate data management and scientific discovery. All these elements may be orchestrated and steered by *Globus* workflows.

Rationale to invest in FAIR research. There are many compelling reasons to create and share FAIR AI models and datasets. Recent studies argue that FAIR data practices are not only part of good research practices, but will save research teams time by decreasing the need for data cleanup and preparation⁴³. It is easy to dismiss anything that sounds new as an “unfunded mandate”. However, FAIR directly relates to many compatible initiatives and goals of most scientifically focused organizations. For instance, FAIRness is closely connected to, and perhaps a prerequisite of, reproducibility. It is also needed for data exploration and is closely connected to ethics issues. FAIR principles can contribute to transparency and other tenets of Open Science.

On the other hand, Supercomputing resources (e.g., Argonne Leadership Computing Facility, Oak Ridge Leadership Computing Facility, National Center for Supercomputing Applications, Texas Advanced Computing Center, etc.) and scientific data facilities (e.g., Advanced Photon Source at Argonne, National Synchrotron Light Source II at Brookhaven National Laboratory, etc.) produce valuable data that may only be effectively shared and reused through the adoption of practical, easy to follow FAIR principles, and the design and deployment of smart software infrastructure. In brief, FAIR is an important step towards an optimal use of taxpayer dollars, it maximizes the science reach of large scale scientific and cyberinfrastructure facilities to power automated AI-driven discovery.

Needs and gaps in AI research that may be addressed by adopting FAIR principles. In the article that established the FAIR Principles, it is emphasized these principles should enable machine actionable data⁴⁴. This is synergistic with the rapid adoption and increased use of AI in research. The more data is easy to locate (Findable), easy to access (A), well described with good and interoperable metadata (I), and available for reuse (R) the easier it will be to use existing data as training or validation sets for AI models. Specific benefits of FAIR AI research throughout the entire discovery cycle include:

- Rapid discovery of data via search and visualization tools, ability to download data for benchmarking and meta-analyses using AI for further scientific discovery.
- Reproducibility of papers and AI models published with them.
- Easy-to-follow guides for how to make data and AI models FAIR are needed, as this process can be difficult, particularly for researchers to whom it is new.
- Establish and promote tools and data infrastructures that accept, store, and offer FAIR and AI-ready data.
- In biomedicine and healthcare, AI models could improve generalization by exposing them to diverse, FAIR datasets.
- Engagement from industry partners is vital to this effort, since they are a major force in AI innovation.
- Get publishers involved and committed to using FAIR, both for data and for other objects such as AI models and software, as they are where research results are shared.
- Adopting the FAIR principles in AI research will also facilitate more effective reporting. Inadequate explanations on the main parts of AI methods not only lead to distrust of the results, but also act as blockers in transferring them to an applied context, such as the clinic and patient care.
- Making FAIR datasets available in HEP is crucial to obtaining benchmark performances of AI models that make AI-driven discovery possible. While a large number of models have been developed for targeted tasks like classification of jets in collider experiments⁴⁵, their performances vary with the choice of training datasets, their preprocessing, and training conditions. Developing FAIR datasets and FAIRifying AI models with well defined hyperstructure and training conditions will allow uniform comparison of these models.
- Establishing seamless and interoperable data e-infrastructures. As these infrastructures mature, a new *AI services layer* will emerge; defining the FAIR principles in advance is thus important in order to accelerate this process.
- Computer science and AI research on efficient generic surrogate architectures and methods to derive reliable surrogate performance for a given accuracy (i.e., towards general surrogate performance models) will benefit extensively from FAIR data and processes.
- One element that has been often debated in AI solutions is of fair (or unbiased) models. This issue is one of the most critical in life sciences, and especially when considering applications that have a direct consequence to human health. FAIR AI and data can facilitate the overall process of identifying potential biases in the involved process.
- Where reproducibility cannot be guaranteed, FAIR data and processes can help establish at a minimum scientific correctness.

Agreed-upon approaches/best practices to identify foundational connections between scientific (meta)datasets, AI models, and hardware. Since this work is in its infancy, there is an urgent need to create incentive structures to impel researchers to invest time and effort to adopt FAIR principles in their

research, since these activities will lower the barrier to adopting AI methodologies. Adopting FAIR best practices will bring about immediate benefits. For instance, FAIR AI models can be constantly reviewed and improved by researchers. Furthermore, software can be optimized for performance or expanded in functionality, rather than standing still and stagnant. In materials science and chemistry, and many other disciplines, thousands of AI models are published each year. Thus, it is critical to rank best AI models, FAIRly share them, and develop APIs to streamline their use within minutes or seconds. Specific initiatives to address these needs encompass:

- **GOFAIRUS** (<https://www.gofair.us>). FAIR papers that efficiently link publications, AI models, and benchmarks to produce figures of merit that quantify performance of AI models and sanity of datasets.
- **MLCommons** (<https://mlcommons.org/en/>). A consortia that brings industry and academic partners together in a pre-competitive space to compare performance of specific tasks and datasets using different hardware architectures and software/hardware combinations.
- **Garden** (<https://thegardens.ai>). A platform for publishing, discovering and resuing FAIR AI models, linked to FAIR and AI-ready datasets, in physics, chemistry and materials science.
- **Bridge2AI** (<https://commonfund.nih.gov/bridge2ai>). FAIR principles can enable ethics inquiries in datasets, easing their use by communities of practice.

While these approaches aim to ease the adoption and development of AI models for scientific discovery and to develop methods to quantify the statistical validity, reliability and reproducibility of AI for inference, there are other lines of research that explore the interplay between datasets, AI models, optimization methods, hardware architectures, and computing approaches from training through to inference. It is expected that FAIR and AI-ready datasets may facilitate these studies. For instance, scientific visualization and accelerated computing have been combined to quantify the impact of multi-modal datasets to optimize the performance of AI models for healthcare⁴⁶, cosmology^{47,48}, high energy physics^{9,49}, and observational astronomy^{50,51}, to mention a few exemplars. These studies shed new light into the features and patterns that AI extracts from data to make reliable predictions. Similarly, recent studies^{52–54} have demonstrated that incorporating domain knowledge in the architecture of AI models, and optimization methods (through geometric deep learning and domain aware loss functions) leads to faster (even zero shot) learning and convergence, and optimal performance with smaller training and validation datasets.

It is also worth mentioning that publishing a FAIR AI model including all relevant (meta)data, e.g., set of initial weights for training, all relevant hyperparameters, libraries, dependencies, and the software needed for training and optimization may not suffice to attain full reproducibility. This is because users may use different hardware to train and optimize AI models, and thus the selection of batchsize and learning rate may have to be adjusted if only one or many GPUs are used for distributed training. It may also be the case that users prefer to use AI-accelerator machines and the AI model, hyperparameters, libraries and dependencies will have to be changed. These considerations have persuaded researchers to define FAIRness in the context of AI inference. These caveats were also discussed by Ravi *et al.*⁸, where a FAIR AI model was produced using distributed computing with GPUs, quantized with NVIDIA TensorRT, and trained from the ground up using the SambaNova DataScale system at the **ALCF AI Testbed** (<https://www.alcf.anl.gov/alcf-ai-testbed>). However, FAIRness of these different AI models was quantified at the inference stage.

Promise or role of privacy preserving and federated learning in the creation of FAIR AI datasets and models. Sample case: **PALISADE-X Project** (<https://www.palisadex.net>). The scope of applications in this project includes development of AI models using closed source/sensitive data and leveraging distributed secure enclaves. Current applications include biomedical data, but may be applicable to data from smart grids, national security, physics, astronomy, etc.

The development of FAIR AI tools for privacy preserving federated learning should be guided by several considerations. For instance, ethically sourced data (beyond human safety protection) should include attributes for the creation of AI models in a responsible manner. Furthermore, open, AI-driven discovery with protected data should be guided with clear principles and examples that demonstrate how to use data in a way that protects the privacy of individuals or organization. Ethical data sharing and automated AI-inference results should be regulated with input from interdisciplinary teams. Care should be taken to perform a thorough external validation of developed models to capture diversity and measure their applicability across different data distributions. In the case of personalized medicine, existing smart watches can identify markers that may identify suicidal behaviour. Should these results be readily shared with healthcare provider without input from individuals? These considerations demand thoughtful policy development and governance for datasets and AI models.

Ethical issues go well beyond biology, genomics and healthcare. For instance, in materials science and chemistry a recent article described a methodology to train an AI model to minimize drug toxicity, and then used to show potential misuse of maximizing toxicity for chemical weapons development⁵⁵.

Transparent/interpretable AI models are considered critical to facilitating the adoption of AI-driven discovery. Why is (or isn't) this possible/reasonable in view of the ever increasing complexity of AI models? AI models have surpassed human performance in image classification challenges^{56,57}. These algorithms process data, identify patterns and features in different ways to humans. When we try to understand what these AI models learn and how they make decisions, we should avoid using human-centric judgements on what is correct or acceptable. These algorithms need not work or “think” as humans to be promoted as reliable and trustworthy tools for scientific discovery and innovation. Rather, we should focus on defining clear, easy to follow, quantifiable principles to thoroughly examine AI predictions. At the same time, it is important to distinguish persuasive⁵⁸ from interpretable AI⁵⁹.

Scientific visualization is a powerful tool to explore and get new insights on how and what AI models learn; the interplay among data, a model's architecture, training and optimization schemes (when they incorporate domain knowledge) and hardware used; and what triggers a sharp response in an AI model that is related to new phenomena or unusual noise anomalies^{47,48}.

Explainability of AI models can be deemed crucial in scientific domains when the decision making process of deep learning models can be important to make them trustworthy and generalizable. Interpretability of deep neural networks is important to identify relative importance of features and identify information pathways within the network. With prohibitively large complexities of neural architectures, existing methods of explainable AI can be constrained by their lack of scalability and robustness. Domain specific approaches for developing novel methods in explainable AI need to be explored to ensure development of reliable and reusable AI models^{26,60}.

A number of strategies to create explainable AI models include the use and adoption of community-backed standards for effective reporting of AI applications. AI practitioners should also define use space of a model, and evaluate resource credibility using, e.g., these Ten Simple Rules⁶¹. It is also good practice to use well-known metrics to quantify the performance, reliability, reproducibility, and statistical soundness of AI predictions.

Current trends in explainable AI include the integration of domain knowledge in the design of AI architectures, training and optimization schemes, while also leaving room for serendipitous discovery^{62,63}. At the end of the day we expect AI to shed light on novel features and patterns hidden in experimental datasets that current theories or phenomenology have not been able to predict or elucidate⁶⁴. Exploring foundation AI models, such as GPT-4⁶⁵, provides new insights on what the model has learned, and helps understand concepts such as model memorization and deep generalization.

Holy grail of FAIR science. We identified the following objectives and end-goals of FAIR initiatives.

- As stated before, FAIR is not the end-goal. It is a journey of improving practices and adapting research resources along with technology innovations. FAIR contributes by enabling discovery and innovation. It will also help us identify best practices that lead to sustainability, lasting impact, and funding.
- Software, datasets, and AI models are all first class research objects. Investments and participation in FAIR activities should be considered for career advancement, tenure decisions, etc.
- Since digital assets cannot be open source forever (indefinite funding), FAIR initiatives should also inform what data, AI models and other digital assets should be preserved permanently.
- Leverage scientific data infrastructure to automate⁶⁶ the validation and assessment of the novelty and soundness of new AI results published in peer-reviewed publications.
- Create user friendly platforms that link articles with AI models, data, and scientific software to quantify the FAIRness of AI models, e.g., the [Physiome Project](https://journal.physiomeproject.org/) (<https://journal.physiomeproject.org/>), the [Center for Reproducible Biomedical Modeling](https://reproduciblebiomodels.org) (<https://reproduciblebiomodels.org>), and the [Garden project](#).
- Recent approaches have showcased how to combine data facilities, computing resources, FAIR AI models, and FAIR and AI-ready data to enable automated, AI-driven discovery⁸.

Creating FAIR discovery platforms for specific disciplines can possibly lead to silos, which would cut short the expected impact of FAIR initiatives. Therefore, synergies among ongoing efforts are critical to link AI model repositories, data facilities, and computing resources. This approach will empower researchers to explore and select available data and AI models. Following clear guidelines to publish and share these digital assets will facilitate the ranking of AI models according to their performance, ease of use and reproducibility; and for datasets according to their readiness for AI R&D and compatibility with modern computing environments. This approach is at the heart of the [Garden Project](#), which will deliver a platform in which FAIR AI models for materials science, physics, and chemistry are linked to FAIR data, and published in a format that streamlines their use on the cloud, supercomputing platforms or personal computers. AI Model Gardens will enable researchers to cross-pollinate novel methods and approaches utilized in seemingly disconnected disciplines to tackle similar challenges, such as classification, regression, denoising, forecasting, etc. As these approaches mature, and researchers adopt FAIR principles to produce AI-ready datasets, it will be possible to identify general purpose AI models, paving the way for the creation of *foundation AI models*, which are trained with broad datasets and may then be used for many downstream applications with relative ease^{67–69}. An exemplar of this approach in the context of materials science was presented by Hatakeyama-Sato and Oyaizu⁷⁰, in which an AI model was trained with diverse sources of information, including text, chemical structures, and more than 40 material properties. Through multitask and multimodal learning, this AI model was able to predict 40 parameters simultaneously, including numeric properties, chemical structures, and text.

Achieving the expected outcomes of FAIR initiatives requires coordinated scientific exploration and discovery across groups, institutions, funding agencies and industry. The Bridge2AI program is an example that such interdisciplinary, and multi-funding agency approach is indeed possible. Well defined, targeted efforts of this nature will have a profound impact in the practice of AI in science, engineering and industry, facilitating the cross-pollination of expertise, knowledge and tools. We expect that this document sparks conversations among scientists, engineers and industry stakeholders engaged in FAIR research, and helps define, implement and adopt an agreed-upon, practical, domain-agnostic FAIR framework for AI models and datasets that guides the development of scientific data infrastructure and computing approaches that are needed to enable and sustain discovery and innovation.

Received: 17 October 2022; Accepted: 9 June 2023;
Published online: 26 July 2023

References

1. Wilkinson, M. D. *et al.* The FAIR guiding principles for scientific data management and stewardship. *Sci. Data* **3**, 160018, <https://doi.org/10.1038/sdata.2016.18> (2016).
2. Wilkinson, M. D. *et al.* A design framework and exemplar metrics for FAIRness. *Scientific Data* **5**, 180118, <https://doi.org/10.1038/sdata.2018.118> (2018).
3. Chue Hong, N. P. *et al.* FAIR principles for research software (FAIR4RS principles). *Research Data Alliance* <https://doi.org/10.15497/RDA00068> (2022).
4. Goble, C. *et al.* FAIR computational workflows. *Data Intelligence* **2**, 108–121, https://doi.org/10.1162/dint_a_00033 (2020).
5. Neubauer, M. S., Roy, A. & Wang, Z. Making Digital Objects FAIR in High Energy Physics: An Implementation for Universal FeynRules Output (UFO) Models. *SciPost Phys. Codebases* **13**, <https://doi.org/10.21468/SciPostPhysCodeb.13Y> (2023).
6. Bourne, P. E. *et al.* Playing catch-up in building an open research commons. *Science* **377**, 256–258, <https://doi.org/10.1126/science.abo5947> (2022).
7. Campo, E. M., Shankar, S., Szalay, A. S. & Hanisch, R. J. Now is the time to build a national data ecosystem for materials science and chemistry research data. *ACS Omega* **7**, 16, 13398–13402, <https://doi.org/10.1021/acsomega.2c00905> (2022).
8. Ravi, N. *et al.* FAIR principles for AI models with a practical application for accelerated high energy diffraction microscopy. *Scientific Data* **9**, 657, <https://doi.org/10.1038/s41597-022-01712-9> (2022).
9. Duarte, J. *et al.* FAIR AI Models in High Energy Physics. Preprint at <https://doi.org/10.48550/arXiv.2212.05081> (2022).
10. Chard, R. *et al.* Dhub: Model and data serving for science. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 283–292, <https://doi.org/10.1109/IPDPS.2019.00038> (2019).
11. Chard, R. *et al.* Funcx: A federated function serving fabric for science. In *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '20, 65–76, <https://doi.org/10.1145/3369583.3392683> (Association for Computing Machinery, New York, NY, USA, 2020).
12. Chard, K. *et al.* Globus nexus: A platform-as-a-service provider of research identity, profile, and group management. *Future Generation Computer Systems* **56**, 571–583, <https://doi.org/10.1016/j.future.2015.09.006> (2016).
13. Verma, G. *et al.* HPCFAIR: Enabling FAIR AI for HPC applications. In *IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC)*, 58–68, <https://doi.org/10.1109/MLHPC54614.2021.00011> (2021).
14. Liao, C. *et al.* HPC ontology: Towards a unified ontology for managing training datasets and AI models for high-performance computing. In *IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments*, 69–80, <https://doi.org/10.1109/MLHPC54614.2021.00012> (2021).
15. Brown, C. & Luszczek, P. SABATH GitHub: A software ecosystem for downloading and running ML/AI benchmarks. <https://github.com/icl-utk-edu/slip/tree/sabath>. Accessed: 2022-6-1.
16. Thiyyagalingam, J. *et al.* AI benchmarking for science: Efforts from the MLCommons science working group. In *HPC on Heterogeneous Hardware (H3) Workshop at ISC Conference*, 47–64, https://doi.org/10.1007/978-3-031-23220-6_4 (2023).
17. Thiyyagalingam, J., Shankar, M., Fox, G. & Hey, T. Scientific machine learning benchmarks. *Nature Reviews Physics* **4**, 413–420, <https://doi.org/10.1038/s42254-022-00441-7> (2022).
18. Fox, G., Hey, T. & Thiyyagalingam, J. Science data working group of MLCommons research. <https://mlcommons.org/en/groups/research-science/>. Accessed: 2020-12-3.
19. Blaiszik, B. *et al.* The Materials Data Facility: Data services to advance materials science research. *JOM* **68**, 2045–2052, <https://doi.org/10.1007/s11837-016-2001-3> (2016).
20. Blaiszik, B. *et al.* A data ecosystem to support machine learning in materials science. *MRS Communications* **9**, 1125–1133, <https://doi.org/10.1557/mrc.2019.118> (2019).
21. Chen, Y. *et al.* A FAIR and AI-ready Higgs boson decay dataset. *Scientific Data* **9**, 31, <https://doi.org/10.1038/s41597-021-01109-0> (2022).
22. CERN. CERN Open Data Policy for the LHC Experiments. <http://opendata.cern.ch/docs/cern-open-data-policy-for-lhc-experiments> (2020).
23. Samuel, S., Löffler, F. & König-Ries, B. Machine learning pipelines: Provenance, reproducibility and FAIR data principles. In *Provenance and Annotation of Data and Processes*, 226–230, https://doi.org/10.1007/978-3-030-80960-7_17 (Springer, 2021).
24. Bailey, S. *et al.* Data and Analysis Preservation, Recasting, and Reinterpretation. Preprint at <https://doi.org/10.48550/arXiv.2203.10057> (2022).
25. Katz, D. S., Psomopoulos, F. E. & Castro, L. J. Working Towards Understanding the Role of FAIR for Machine Learning <https://doi.org/10.5281/zenodo.5594990> (2021).
26. Neubauer, M. S. & Roy, A. Explainable AI for High Energy Physics. Preprint at <https://doi.org/10.48550/arXiv.2206.06632> (2022).
27. Benelli, G. *et al.* Data Science and Machine Learning in Education. Technical Report. United States. <https://doi.org/10.2172/1882567> (2022).
28. Javier D. *Particle Physics and Machine Learning*. <https://jduarte.physics.ucsd.edu/capstone-particle-physics-domain>, <https://doi.org/10.5281/zenodo.4768815>.
29. U.S. White House Office of Science and Technology Policy. Materials Genome Initiative for Global Competitiveness. https://www.mgi.gov/sites/default/files/documents/materials_genome_initiative-final.pdf (2011).
30. U.S. White House Office of Science and Technology Policy. Materials Genome Initiative Strategic Plan. <https://www.mgi.gov/sites/default/files/documents/MGI-2021-Strategic-Plan.pdf> (2021).
31. Deagen, M. E., Brinson, L. C., Vaia, R. A. & Schadler, L. S. The materials tetrahedron has a “digital twin”. *MRS Bulletin* **47**, 379–388, <https://doi.org/10.1557/s43577-021-00214-0> (2022).
32. Blaiszik, B. 2021 AI/ML Publication Statistics and Charts. <https://doi.org/10.5281/zenodo.7057437> (2022).
33. Andersen, C. W. *et al.* OPTIMADE, an API for exchanging materials data. *Scientific Data* **8**, 217, <https://doi.org/10.1038/s41597-021-00974-z> (2021).
34. Brinson, L. *et al.* Polymer nanocomposite data: Curation, frameworks, access, and potential for discovery and design. *ACS Macro Letters* **9**, 1086–1094, <https://doi.org/10.1021/acsmacrolett.0c00264> (2020).
35. Bohr, A. & Memarzadeh, K. The rise of artificial intelligence in healthcare applications. *Artificial Intelligence in Healthcare* <https://doi.org/10.1016/B978-0-12-818438-7.00002-2> (2020).
36. Walsh, I. *et al.* DOME: recommendations for supervised machine learning validation in biology. *Nature Methods* **18**, 1122–1127, <https://doi.org/10.1038/s41592-021-01205-4> (2021).
37. Huerta, E. A. *et al.* Convergence of Artificial Intelligence and High Performance Computing on NSF-supported Cyberinfrastructure. *Journal of Big Data* **7**, 88, <https://doi.org/10.1186/s40537-020-00361-2> (2020).
38. Khan, A., Huerta, E. A. & Das, A. Physics-inspired deep learning to characterize the signal manifold of quasi-circular, spinning, non-precessing binary black hole mergers. *Physics Letters B* **808**, 135628, <https://doi.org/10.1016/j.physletb.2020.135628> (2020).
39. Huerta, E. A. *et al.* Accelerated, scalable and reproducible AI-driven gravitational wave detection. *Nature Astronomy* **5**, 1062–1068, <https://doi.org/10.1038/s41550-021-01405-0> (2021).
40. Chaturvedi, P., Khan, A., Tian, M., Huerta, E. A. & Zheng, H. Inference-Optimized AI and High Performance Computing for Gravitational Wave Detection at Scale. *Front. Artif. Intell.* **5**, 828672, <https://doi.org/10.3389/frai.2022.828672> (2022).
41. Dempsey, W., Foster, I., Fraser, S. & Kesselman, C. Sharing begins at home: How continuous and ubiquitous FAIRness can enhance research productivity and data reuse. *Harvard Data Science Review* **4**, <https://doi.org/10.1162/99608f92.44d21b86> (2022).

42. FAIR4HEP. Cookiecutter4fair: v1.0.0, <https://doi.org/10.5281/zenodo.7306229> (2022).
43. Mons, B. *et al.* Invest 5% of research funds in ensuring data are reusable. *Nature* **578**, 491–491, <https://doi.org/10.1038/d41586-020-00505-7> (2020).
44. Wilkinson, M. D. *et al.* The fair guiding principles for scientific data management and stewardship. *Scientific data* **3**, 1–9, <https://doi.org/10.1038/sdata.2016.18> (2016).
45. Kasieczka, G. *et al.* The Machine Learning landscape of top taggers. *SciPost Physics* **7**, 014, <https://doi.org/10.21468/SciPostPhys.7.1.014> (2019).
46. Gupta, A., Huerta, E., Zhao, Z. & Moussa, I. Deep learning for cardiologist-level myocardial infarction detection in electrocardiograms. In Jarm, T., Cvetkoska, A., Mahnič-Kalamiza, S. & Miklavcic, D. (eds.) *8th European Medical and Biological Engineering Conference*, 341–355, https://doi.org/10.1007/978-3-030-64610-3_40 (Springer International Publishing, Cham, 2021).
47. Khan, A. *et al.* Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey. *Physics Letters B* **795**, 248–258, <https://doi.org/10.1016/j.physletb.2019.06.009> (2019).
48. Khan, A. *et al.* Deep transfer learning at scale for cosmology. <https://www.youtube.com/watch?v=8-jcflTZNdA&t=0s> (2018).
49. Roy, A. & Neubauer, M. S. Interpretability of an Interaction Network for identifying $H \rightarrow b\bar{b}$ jets. PoS, ICHEP2022 223, 11, <https://doi.org/10.22323/1.414.0223> (2022).
50. Wei, W. *et al.* Deep transfer learning for star cluster classification: I. application to the PHANGS-HST survey. *Monthly Notices of the Royal Astronomical Society* **493**, 3178–3193, <https://doi.org/10.1093/mnras/staa325> (2020).
51. Whitmore, B. C. *et al.* Star cluster classification in the PHANGS-HST survey: Comparison between human and machine learning approaches. *Monthly Notices of the Royal Astronomical Society* **506**, 5294–5317, <https://doi.org/10.1093/mnras/stab2087> (2021).
52. Rosofsky, S. G., Majed, H. A. & Huerta, E. A. Applications of physics informed neural operators. *Mach. Learn. Sci. Tech.* **4**, 025022, <https://doi.org/10.1088/2632-2153/acd168> (2023).
53. Kansky, K. *et al.* Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, 1809–1818, <https://doi.org/10.5555/3305381.3305568> (JMLR.org, 2017).
54. Rosofsky, S. G. & Huerta, E. A. Magnetohydrodynamics with Physics Informed Neural Operators. Preprint at <https://doi.org/10.48550/arXiv.2302.08332> (2023).
55. Urbina, F., Lentzos, E., Invernizzi, C. & Ekins, S. Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence* **4**, 189–191, <https://doi.org/10.1038/s42256-022-00465-9> (2022).
56. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Commun. ACM* **60**, 84–90, <https://doi.org/10.1145/3065386> (2017).
57. Chollet, F. Xception: Deep learning with depthwise separable convolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1800–1807, <https://doi.org/10.1109/CVPR.2017.195> (2017).
58. Gilpin, L. H. *et al.* Explaining explanations: An overview of interpretability of machine learning. *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)* 80–89, <https://doi.org/10.1109/DSAA.2018.00018> (2018).
59. The Royal Society. Explainable AI: the basics. Policy Briefing. <https://royalsociety.org/-/media/policy/projects/explainable-ai/AI-and-interpretability-policy-briefing.pdf> (2019).
60. Khot, A., Neubauer, M. S. & Roy, A. A Detailed Study of Interpretability of Deep Neural Network based Top Taggers. Preprint at <https://doi.org/10.48550/arXiv.2210.04371> (2022).
61. Erdemir, A. *et al.* Credible practice of modeling and simulation in healthcare: ten rules from a multidisciplinary perspective. *Journal of Translational Medicine* **18**, 1–18, <https://doi.org/10.1186/s12967-020-02540-4> (2020).
62. Stanev, V. G., Choudhary, K., Kusne, A. G., Paglione, J. & Takeuchi, I. Artificial intelligence for search and discovery of quantum materials. *Communications Materials* **2**, <https://doi.org/10.1038/s43246-021-00209-z> (2021).
63. Chen, B. *et al.* Automated discovery of fundamental variables hidden in experimental data. *Nature Computational Science* **2**, 433–442, <https://doi.org/10.1038/s43588-022-00281-6> (2022).
64. Davies, A. *et al.* Advancing mathematics by guiding human intuition with AI. *Nature* **600**, 70–74, <https://doi.org/10.1038/s41586-021-04086-x> (2021).
65. Brown, T. B. *et al.* Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, <https://doi.org/10.5555/3495724.3495883> (Curran Associates Inc., Red Hook, NY, USA, 2020).
66. Madduri, R. *et al.* Reproducible big data science: A case study in continuous fairness. *PLoS ONE* **14**, <https://doi.org/10.1371/journal.pone.0213013> (2019).
67. Bommasani, R. *et al.* On the Opportunities and Risks of Foundation Models. Preprint at <https://doi.org/10.48550/arXiv.2108.07258> (2021).
68. Chowdhery, A. *et al.* PaLM: Scaling Language Modeling with Pathways. Preprint at <https://doi.org/10.48550/arXiv.2204.02311> (2022).
69. OpenAI. GPT-4 Technical Report, <https://cdn.openai.com/papers/gpt-4.pdf>. Preprint at <https://doi.org/10.48550/arXiv.2303.08774> (2023).
70. Hatakeyama-Sato, K. & Oyaizu, K. Integrating multiple materials science projects in a single neural network. *Communications Materials* **1**, <https://doi.org/10.1038/s43246-020-00052-8> (2020).

Acknowledgements

E.A.H.: This work was supported by the FAIR Data program of the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under contract number DE-AC02-06CH11357. **B.B.:** This work was supported by the National Science Foundation under NSF Award Numbers: 1931306 and 2209892. **C.K.:** This work was supported by the National Science Foundation under NSF Award Number: 1916481 “BD Hubs: Collaborative Proposal: West: Accelerating the Big Data Innovation Ecosystem” and NSF Award Number: 2226453 “Disciplinary Improvements: AI Readiness, Reproducibility, and FAIR: Connecting Computing and Domain Communities Across the ML Lifecycle”. **D.S.K., V.K., M.S.N., A.R.:** This work was supported by the FAIR Data program of the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under contract number DE-SC0021258. **G.F., S.J.:** This work was supported by the FAIR Data program of the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under contract number DE-SC0021352. **C.D., L.H.:** The ESCAPE project has received funding from the Horizon 2020 research and innovation programme, Grant Agreement no. 824064. The EOOSC-Future project has received funding from the Horizon 2020 research and innovation programme, Grant Agreement no. 101017536. **C.D.** has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 101002463) and from the Swedish Research Council. **M.E.** The HPCFAIR project is supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Program under Award Number DE-SC0021293.

Author contributions

E.A.H. conceived the convergence of ideas and visions around FAIR initiatives, and their discussion among leads of these efforts, which is documented in this manuscript. All authors reviewed and contributed to the manuscript.

Competing interests

The authors declare the following competing interests: They are funded by the U.S. Department of Energy and/or the National Science Foundation (as described in detail in the Acknowledgements section) to lead the definition and application of FAIR principles for scientific data, AI models, research software, and workflows. They are the lead developers of scientific data infrastructure used to enable these advances, including Globus, funcX (now Globus Compute), the Data and Learning Hub for Science (DLHub), CookieCutter4FAIR, APPFL: Open-Source Software Framework for Privacy-Preserving Federated Learning, the Garden Project, the RDA FAIR for Machine Learning Interest Group, FARR: FAIR in ML, AI Readiness, & Reproducibility Research Coordination Network, and the ELIXIR Machine Learning focus group.

Additional information

Correspondence and requests for materials should be addressed to E.A.H.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© UChicago Argonne, LLC, Operator of Argonne National Laboratory 2023