RDA Interest Group: FAIR for Machine Learning

Case Statement

1. Introduction

The idea of FAIR (findable, accessible, interoperable, and reusable) in the context of scientific data management and stewardship was developed in 2014 and turned into specific principles in 2016[1]. Along the way, the idea was generalized in concept to apply to both data and other digital scholarly objects, but it has become clear in practice that what works for data may not directly work for all other digital objects. For example, both previous and ongoing work show that many of the guiding FAIR principles need to either be re-written or reinterpreted for software, resulting in the FAIR principles for Research Software[2], already with an adoption commitment from different communities and institutes[3]. The FAIR principles also can apply to machine learning tools and models, though a direct application is not always possible as machine learning combines aspects of data, software and computational workflows.

There is a large amount of work around FAIR, both in RDA and elsewhere, initially focused on data and now software and other products but generally not ML models. There has already been some community work in this direction, both within RDA as well as beyond. This subject was first discussed at RDA during the RDA Virtual Plenary 16 as poster 31b (FAIR principles for ML models - https://doi.org/10.5281/zenodo.4271995). Additional activities include two dedicated BoF sessions at RDA VP17 (https://www.rd-alliance.org/defining-fair-machine-learning-ml) and RDA VP18 that aimed to capture the overall perspective on the topic. The discussion around FAIR for Machine Learning continued in further events under different domains; during the FAIR Festival, the efforts of FAIR4ML were presented together with similar initiatives for Software and Workflows. A position article was presented at the 2nd Workshop on Data and research objects management for Linked Open Science (DAMALOS 2021)[4]. Additionally, a dedicated BoF session was held in the context of SC21 ("Towards FAIR for Machine Learning Models", November 18th 2021), with a focus on FAIR for ML models, bringing together HPC ML researchers and the larger ML and information systems

---

[1] *Wilkinson, M., Dumontier, M., Aalbersberg, I. et al. The FAIR Guiding Principles for scientific data management and stewardship. Sci Data 3, 160018 (2016). https://doi.org/10.1038/sdata.2016.18*

[2] *Chue Hong, N. P., Katz, D. S., Barker, M., Lamprecht, A.-L., Martinez, C., Psomopoulos, F. E., Harrow, J., Castro, L. J., Gruenpeter, M., Martinez, P. A., Honeyman, T., et al. (2021). FAIR Principles for Research Software (FAIR4RS Principles). Research Data Alliance. https://doi.org/10.15497/RDA00065*

[3] *Martinez-Ortiz, Carlos, Katz, Daniel S., Lamprecht, Anna-Lena, Barker, Michelle, Loewe, Axel, Fouilloux, Anne, Wyngaard, Jane, Garijo, Daniel, Moldon, Javier, Castro, Leyla Jael, Wheeler, Daniel, Albers, Joost Rutger Demian, & Lee, Allen. (2022). FAIR4RS: Adoption support. Zenodo. https://doi.org/10.5281/zenodo.6258366*

[4] *Katz, Daniel S., Psomopoulos, Fotis and Castro, Leyla Jael, "Working Towards Understanding the Role of FAIR for Machine Learning", 2nd Workshop on Data and Research Objects Management for Linked Open Science, 2021, https://doi.org/10.4126/FRL01-006429415*

communities. Moreover the initiative was presented in the context of "Best Practices for Reusability of Machine Learning Models: Guideline and Specification" during ESIP 2021 in July 2021, as well as in the "Improving 'FAIRness' and 'Fairness' of AI/ML in Geoscience" session in January 2022. Finally, two informal Community Calls were organized, a kickoff in July 2021 and a follow-up in January 2022, aiming to organize and structure the efforts towards an RDA Interest Group.

2. User scenario(s) or use case(s) the IG wishes to address

This Interest Group will enable community members to discuss the various aspects of FAIR as applied to Machine Learning, looking both at domain specific and domain-agnostic use cases, and creating task forces and working groups as needed for specific guidance documents, recommendations, definitions and technical specification to that effect. Note that we intentionally use "task forces" here to distinguish from RDA working groups. In the context of the Interest Group, "Task Forces" are defined as groups that have a shorter duration and are more focused on a particular element of work, as compared to a more traditional RDA Working Group.

In order to ensure that these scenarios are valid across domains (e.g., health, earth science, physics, agriculture, materials science, energy, biology), individual Task Forces (TFs) may be initiated from within the IG that may be focused on particular domains could be initiated, each working in parallel on distinct topics.

Of particular relevance to the IG is a direct interaction to other RDA Interest Groups, Working Group and Community of Practices that have a vested interest in ML, investigating potential overlaps and synergies within the RDA network. Examples of relevant RDA groups are:

1. Sensitive Data Interest Group
2. FAIR Digital Object Fabric Interest Group
3. Software Source Code Interest Group
4. FAIR Data Maturity Model Working Group

Use cases that the IG will be investigating include, but are not limited to, the topics below, structured across the three main output types. These activities aim to enrich the existing RDA list of use cases with ML-specific scenarios.

A. Guidance documents:

- Recommendations and best practices
    - Provide supporting documentation on the differences and connections between FAIR for Machine Learning, and fair (i.e., unbiased) machine learning.
    - Provide guidance to users who build ML models and want to share them
    - Provide guidance to repositories that store ML models (e.g., DLHub, OpenML) and want them to be consistently FAIR across multiple repositories; including how to address challenges on archiving and sharing large-scale models (e.g., language models, earth system models)

  - Guidance on downloading and executing FAIR ML models within your own institute.
  - Guidance on validating and executing FAIR ML models within your own institute, and sharing the validation results (e.g., https://fairmodels.org/)
  - Provide best practices and recommendations about long-term storage of large datasets considering availability, read/write speed and costs, including maximizing data reusability for ML through FAIR data.
- Metadata and Ontology
  - Guidance on common minimum metadata (vocabulary-agnostic and possibly mapped to a broadly used vocabulary) that ML models and tools should provide
  - Determine metadata and ontologies for ML models that can support search across multiple repositories
  - Guidance on applying the FAIR Digital Object Concept for AI, in connection with the RDA FAIR Digital Object Fabric IG
  - Guidance on FAIR metadata for federated learning (visiting data approach for sensitive datasets like patient data in healthcare)
  - Guidance on minimal metadata for ML workflows: hyperparameters needed for reproducibility (e.g., learning rate, batch size, optimization schemes, etc.); data format (HDF5, root, etc.) for training, validation and inference, and links to FAIR datasets; metrics to quantify reliability of AI predictions.
- Metrics and Tools
  - Harness and further develop domain-agnostic success metrics, commodity tools and manual means to automate FAIRification and FAIR assessment.
  - Guidance on data/model/workflow versioning
- Awareness
  - Highlighting the concrete benefits a FAIR (DO) concept could have on current issues in AI and ML (data preparation, explainable AI, etc.)
  - Provide guidance on FAIR principles for common AI/ML benchmarking datasets.
  - Provide guidance for AI-readiness with focus on machine actionability to enable automation of AI and ML workflows.
  - Document the landscape on ethics for FAIR and ML
  - Gather common challenges and identify technical gaps

B. Definitions, crosswalks, standards, and terminologies

- What are the parts of ML related to FAIR digital objects (ML can include data, software, and workflows)
- Review existing terminologies/ontologies that support the FAIR principles (for AI-ready data, to ML workflows) and highlight gaps

C. Technical infrastructure specifications

- Provide a (or use an existing) platform to share standards, tools, best practices
- Interface of FAIR with other Data Management Principles in enabling AI/ ML
- Provide machine learning model management (MLMM) to trace whole model life cycle
- Application of the FAIR Digital Object Concept for AI / ML data objects

3. Objectives

The overall aim of this Interest Group is to foster collaborations among researchers and developers who are interested in making machine learning (data, models, workflows, etc.) FAIR, along with those who contribute to the infrastructure and policies that support this. It will work closely with other FAIR RDA Groups (such as the FAIR for Research Software Working Group), as machine learning combines aspects of data and software, but is distinctly different from both.

Specifically, objectives of this IG are to:

1. Discuss where FAIR should apply to ML, considering the work in other working groups and focusing on gaps
2. Define and prioritize cases for new Task Forces and Working Groups
3. Ultimately, build a community of practice for information sharing about ML and FAIR pertaining to ML

4. Participation

The FAIR4ML IG is open to all RDA members to participate, and is of particular interest to researchers and professionals active in Machine Learning. People who are interested in the way the FAIR principles can be applied to/customized for different objects are also possible participants, as a key component of the IG will be to investigate the FAIR principles for ML. This IG is also of interest to RDA members who have a vested interest in adopting the outputs and recommendations of this IG.

Finally, this IG will establish a set of liaisons to other key RDA groups, initially from within the IG Chairs group, with the explicit task of maintaining a bi-directional awareness of the efforts. An initial list of these key groups is:

1. FAIR for Research Software Working Group (or the RDA entity which will maintain the WG output)
2. Software Source Code Interest Group
3. FAIR Digital Object Fabric Interest Group
4. FAIR Data Maturity Model Working Group

5. Outcomes

All outcomes to be published through the RDA platforms, and communicated across all relevant channels.

1. Refined use cases relevant to FAIR 4 ML
2. Definitions on FAIR 4 ML
3. Recommendation and guidance documents on FAIR 4 ML
4. Technical specifications in support of FAIR 4 ML
5. Produce and offer training of the supporting outcomes on FAIR 4 ML (connection with RDA groups as well as external networks such as Pistoia Alliance and ELIXIR)

6. Mechanism

The Interest Group will be holding regular, monthly calls, aiming for the overall coordination of the IG, as well as facilitating the discussion within the community.

Any creation of a new Task Force or Working Group will be discussed and announced in the IG, both through the dedicated mailing list, but also within the regular TCs. The Task Forces themselves will maintain their own coordination, structure and schedules, but report and update the wider community through the IG calls.

Finally, general engagement with the wider community to be held through Town halls, webinars etc, organized ad hoc throughout the year, while at the same time, ensuring regular participation of the IG to the RDA Plenaries.

7. Timeline

Our main goals for the first 12 months of the IG are to:

1. Build the IG membership
2. Identify the main definitions relevant for FAIR for Machine Learning
3. Engage with relevant RDA IGs/WGs/CoP and external communities and projects
4. Create initial TFs/WGs

After approval of the IG, immediate focus will be given to the organization of a FAIR4ML session in the next RDA Plenary, aiming for critical mass from the RDA members.

8. Group Members

Potential co-chairs of the RDA FAIR4ML IG are:

1. Lars Eklund (Europe, SE)
2. Amy Nurnberger (Americas, USA)
3. Johan van Soest (Europe, NL)
4. Daniel S. Katz (Americas, USA)
5. Fotis Psomopoulos (Europe, GR)
6. David Elbert (Americas, USA)
7. Leyla Jael Castro (Europe, DE)
8. Gnana Bharathy (Australasia, AU)

First name/s	Surname	Preferred Pronouns	Institution/s	Country
Amy	Nurnberger	she/her	Massachusetts Institute of Technology	USA
Lars	Eklund	he/him	SND Swedish National Data Service and UPPMAX HPC center Uppsala University	Sweden

Johan	van Soest	he/him	Maastricht University	The Netherlands
Christine	Kirkpatrick	she/her	San Diego Supercomputer Center	USA
Gnana	Bharathy	he/him/his	Australian Research Data Commons and Univ. of Tech. Sydney	Australia
Matthijs	Sloep	he/him	Maastricht University	The Netherlands
Chris	Erdmann	he/him	AGU	USA
Huajin	Wang	she/her	Carnegie Mellon University	USA
Ge	Peng	she/her	University of Alabama in Huntsville	USA
Diego	Chialva	he/him/his	ERCEA	Belgium
Daniel S.	Katz	he/him	University of Illinois at Urbana-Champaign	USA
Chunhua	Liao	he/him	Lawerence Livermore National laboratory	USA
Núria	Queralt Rosinach	she/her	Leiden University Medical Center	The Netherlands
Vladimir	Makarov	he/him	Pistoia Alliance	USA
Eliu	Huerta	he/him/his	Argonne National Laboratory and The University of Chicago	USA
Fotis	Psomopoulos	he/him/his	Institute of Applied Biosciences, CERTH	Greece
Joaquin	Vanschoren	he/him	TU Eindhoven	Netherlands
Daniel	Wang	he/him	M.D. Anderson Cancer Center	USA
Junqi	Yin	he/him	Oak Ridge National Laboratory	USA
Michelle	Barker	she/her	Research Software Alliance	Australia
Yuhan “Douglas”	Rao	he/him	North Carolina State University	USA
Ian	Bruno	he/him	CCDC	UK
Francoise	Genova	she/her	CDS, Observatoire Astronomique de Strasbourg	France
David	Elbert	he/him	Johns Hopkins University	USA
Sanjay	Purushotham	he/him	University of Maryland Baltimore County	USA
Nicolas	Blumenröhr	he/him	Karlsruhe Institute of Technology	Germany
Andy	Gallo	he/him	GE Research	USA

			andrew.gallo@ge.com	
Geoffrey	Fox	he/him	Univ. of Virginia	USA
Murali	Emani	he/him	Argonne National Laboratory	USA
Ishwar	Chandramouliswaran	he/him	National Institutes of Health	USA
Leyla Jael	Castro	she/her	ZB MED Information Centre for Life Sciences	Germany
Rob	Quick	he/him	Indiana University	USA
Santosh Karthikeyan	Viswanathan	he/him	AstraZeneca	India
Tom	Pollard	he/him	MIT, PhysioNet	USA
Malin	Sandström	she/her	INCF, Karolinska Institute	Sweden