FAIR Assessment Tools: Towards an "Apples to Apples" Comparisons

Authorship Community:

Mark D Wilkinson[1,3,*], Susanna-Assunta Sansone[2,4,*], Marjan Grootveld[2,5], Josefine Nordling[2,6,] Richard Dennis[2,7], David Hecker[2,8] on behalf of the EOSC FAIR Metrics subgroup

[1] Co-Chair, EOSC Task Force on FAIR Metrics and Data Quality
[2] Member, EOSC Task Force on FAIR Metrics and Data Quality
[3] Departamento de Biotecnología-Biología Vegetal, Escuela Técnica Superior de Ingeniería Agronómica, Alimentaria y de Biosistemas, Centro de Biotecnología y Genómica de Plantas. Universidad Politécnica de Madrid (UPM) - Instituto Nacional de Investigación y Tecnología Agraria y Alimentaria-CSIC (INIA-CSIC). Campus Montegancedo 28223 Pozuelo de Alarcón (Madrid), Spain. 0000-0001-6960-357X
[4] Oxford e-Research Centre, Department of Engineering Science, University of Oxford, UK. 0000-0001-5306-5690
[5] Data Archiving and Networked Services (KNAW-DANS), The Netherlands. 0000-0002-2789-322X
[6] CSC - IT Center for Science, Finland, 0000-0002-6974-2825
[7] Novo Nordisk Foundation Center for Stem Cell Medicine - reNEW, University of Copenhagen, Denmark, 0000-0002-4472-7194
[8] German Aerospace Center, Research Data Management, Germany, 0000-0003-3836-2800

* Corresponding authors: mark.wilkinson@upm.es, susanna-assunta.sansone@oerc.ox.ac.uk
** All TF members have had the opportunity to review and edit this document; invitations to be on the authorship list were open to all

EOSC Association AISBL

Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

Table of contents

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

Executive summary

As part of the EOSC Task Force on FAIR Metrics and Data Quality, the FAIR Metrics subgroup works to examine the uptake and application of metrics of FAIRness and the use and utility of FAIRness evaluations. A range of FAIR assessment tools is designed to measure compliance with established FAIR Metrics by measuring one or more Digital Objects (DO, including datasets and repositories). Unfortunately, the same DO assessment by different tools often exhibits widely different results because of independent interpretations of the Metrics, metadata publishing paradigms, and even the intent of FAIR itself.

In response to this *status quo*, the FAIR Metrics subgroup (represented by the authors of this report) brought together developers of several FAIR evaluation tools and enabling services (listed in the Acknowledgements) for a series of hands-on hackathon events to identify a common approach to metadata provision that could be implemented by all data publishers such as databases, repositories, and data catalogue managers. This led to identifying a process for (meta)data publishing based on well-established Web standards already in everyday use within data publishing communities, albeit not uniformly. A specification document describing the approach and a series of “Apples-to-Apples” (A2A[1]) benchmarks to evaluate compliance with this (meta)data publishing approach were created during a series of hackathon events. The authors of FAIR evaluation tools also began writing the code to ensure their independent tools would behave identically when encountering these A2A benchmark environments, thus helping to ensure that data publishers following this paradigm will be evaluated in a harmonized manner by all assessment tools; additional considerations for assessment tool harmonization are discussed later. This report explains the rationale for these workshop and hackathon events, outlines the outcomes and work done, describes the current status, and then discusses desirable next steps. The authors propose that the EOSC Association, the EOSC Task Force Long-term Data Preservation, and other groups and projects in Europe and worldwide consider this approach to deliver accurate means to all stakeholders that assist in achieving FAIRness of research data.

[1] Comparing apples to apples means comparing things that can reasonably be compared

EOSC Association AISBL

Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

1. Background

The information that contextualizes the rationale for the FAIR Evaluation Stakeholders' events is fully described in the "Community-driven Governance of FAIRness assessment: an Open Issue, an Open Discussion" whitepaper [1], and it is summarized in the following sections.

The status quo: A journey to a problem

FAIR compliance is currently "stuck" between being an increasingly common research and publishing requirement while remaining an unmeasurable set of ideals.

In the last few years, several activities have focused on developing FAIR metrics and maturity indicators, which describe features of a digital object that would correspond to its compliance with specific FAIR Principles [2], and tests that fully or semi-automatically apply these indicators. A range of FAIR assessment tools have been designed that measure FAIRness using combinations of tests to assess one or more DOs; as of Nov 2022, there are 22 tools described at FAIRassist [https://fairassist.org], and a comparison of a selection of those tagged as "automatic" is available at [3].

Unfortunately, the narrative FAIR principles are insufficient to circumscribe the valid mechanisms to achieve the behaviors they describe. This flexibility in the interpretation of FAIR led to a wide range of "FAIR" implementations. In parallel with this, FAIR assessment tool creators had no guidance on how to apply metrics to a digital object. They pursued independent interpretations of the original intent of each Principle, what needs to be tested, and how. As a result, the same digital object assessed by different tools often exhibits widely different outcomes, which results in frustration and confusion within the scholarly data publishing community.

Inconsistent FAIR Evaluation: The risks for all stakeholders

Without consistency in testing, there will be a tendency towards "gamification" of FAIR assessments.

Adherence to FAIRness is becoming a requirement of agencies and publishers who wish to maximize the utility and reuse of their data. However, given the status quo, mechanisms to evaluate FAIRness mustn't be misused nor misinterpreted, especially when these become a decision-making instrument in funding scenarios.

Nevertheless, research-performing organizations, researchers, and managers need objective measures and guidance to reach a reasonable level of FAIRness and to assist in interpreting the facets of FAIRness achieved by a given digital object. This is where FAIR validation/assessment/assistance tools would be helpful, but only if they are consistent with one another and, more importantly, trusted by the entire stakeholder community.

Unfortunately, the situation at this time is dire. Selecting just two of the automated FAIR assessment tools - The FAIR Evaluator [4] and F-UJI [5] - reveals issues related to their comparability. Figure 1 shows the output from The FAIR Evaluator (commercial version) (Figure 1A) and F-UJI (Figure 1B) when tested on the same URI representing the Catalog record from the Duchenne Parent Project registry's FAIR Data Point (Evaluator Metrics version 1.0.26; F-UJI Metrics version 0.4; September

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

2021). There are at least three significant issues revealed in this side-by-side comparison. First, the two assessment tools use different numbers of tests, and neither uses the complete set of 42 Maturity Indicators identified by the RDA Fair Data Maturity Model Working Group[2]. Second, the distribution of tests differs, with F-UJI having most of its tests in the "Reusability" domain.

In contrast, The FAIR Evaluator has most of its tests in the Findability and Interoperability domains.

Perhaps most striking, however, is the difference in score, where The FAIR Evaluator gives this FAIR Data Point record a near-perfect score. At the same time, F-UJI generates a deficient grade for the same digital resource.

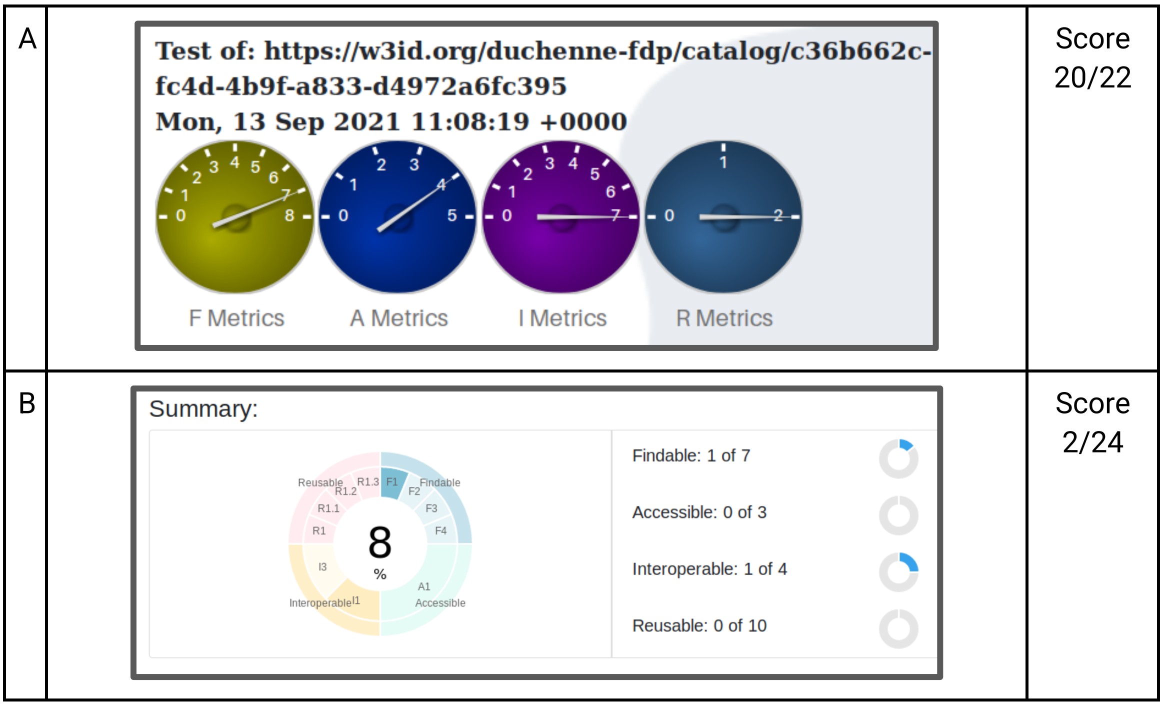

Figure 1: *The output display panels for The Evaluator (A) and F-UJI (B) when tested on the same URI, representing the Catalog record of the FAIR Data Point for the Duchenne Parent Project patient registry.*

This raises the question: What causes this inconsistency? The authors of F-UJI and The FAIR Evaluator were aware, from various personal communications, that their workflows for metadata "harvesting" - i.e., gathering the material that would be tested for its FAIRness - were non-identical. The degree to which this was caused by these differences, or the underlying design of the tests themselves, was unclear; nevertheless, there was ample reason to consider metadata harvesting as a starting point for harmonization since, without a common substrate for testing, differences in the tests themselves or their differential weighting between FAIR assessment tools cannot be delineated.

However, it might not be unexpected that different evaluation systems follow different metadata

2 https://www.rd-alliance.org/group/fair-data-maturity-model-wg/outcomes/fair-data-maturity-model-specification-and-guidelines-0

EOSC Association AISBL

Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

harvesting approaches due to the wide range of existing metadata publishing approaches. To appreciate the severity of the metadata provision problem, it is useful to examine the workflow for metadata retrieval for one of the most common use cases: a deposit identified by a Digital Object Identifier (DOI). A typical workflow to gather the metadata for a DOI-identified data deposit in Zenodo is shown in Figure 2.

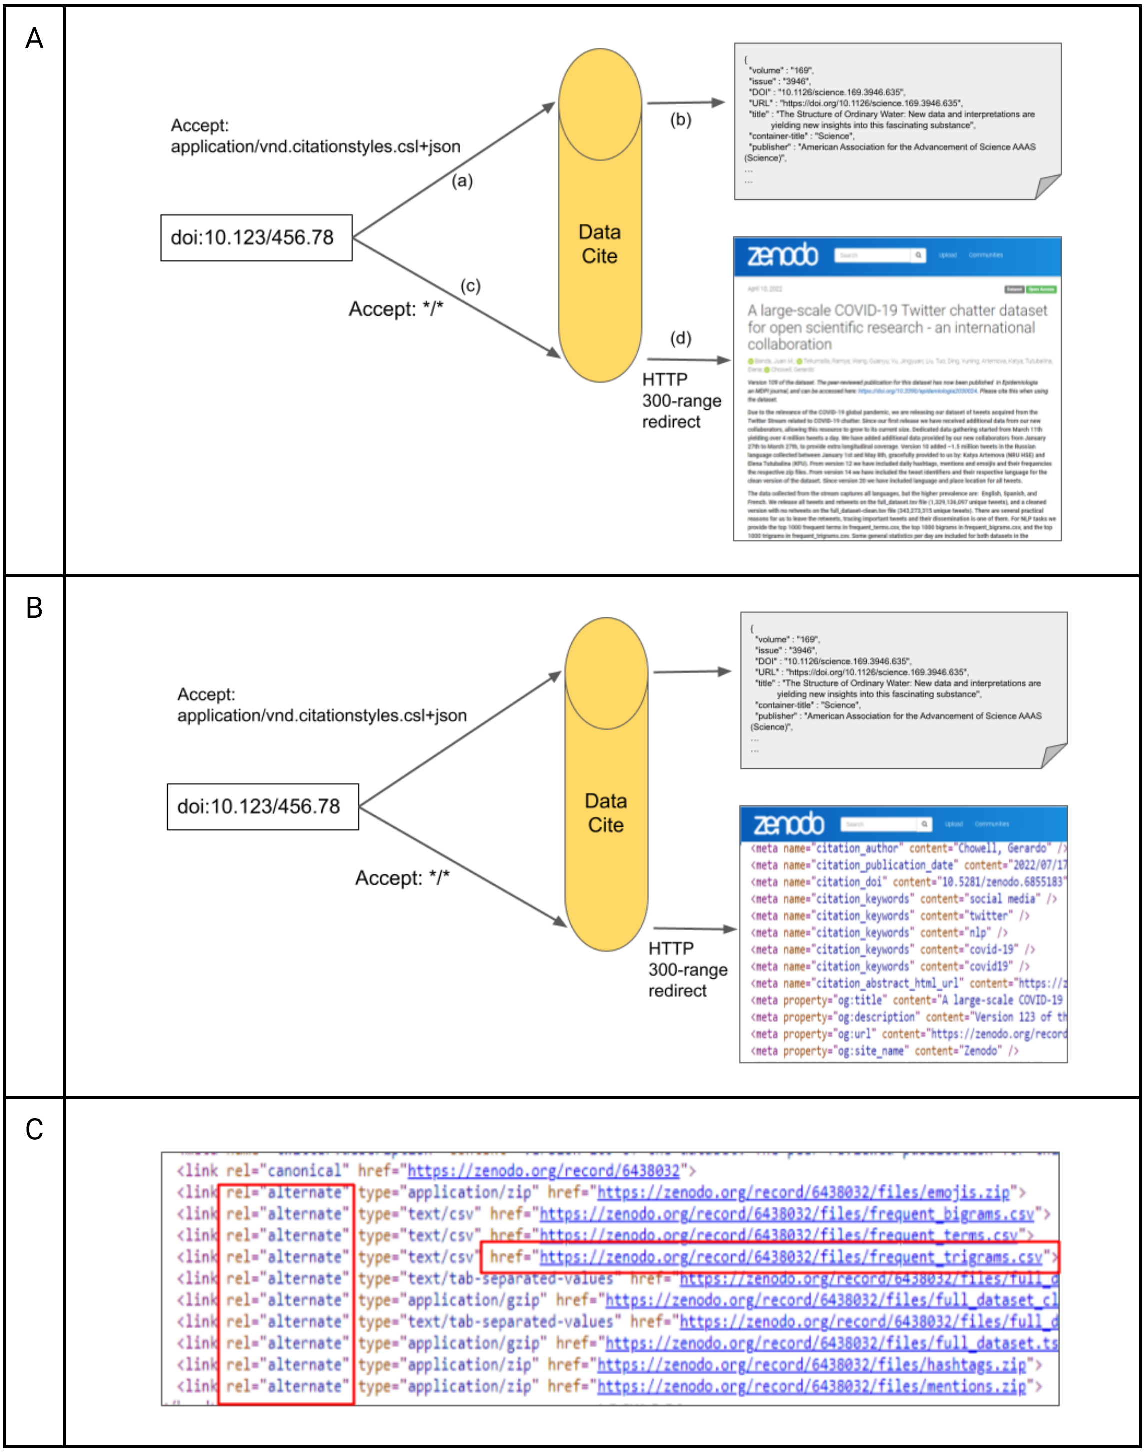

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

***Figure 2:** The workflow to gather deposit metadata starting from a DOI. **A**. The four-step process of content negotiation on the DOI provider (e.g., DataCite) begins with (a) Content Negotiation for vnd.citationstyles.csl+json, which returns (b) metadata about the deposit in DataCite's preferred JSON format. The process then continues with (c) a request for "the default representation" (*/*), which causes the DataCite servers to issue a 300-range HTTP redirect (d) to the Landing Page of the data repository - in this case, Zenodo. **B.** In this panel, the browser view of the Zenodo landing page in panel A has been changed to show a snippet of the HTML code itself, revealing embedded metadata in various formats. Visible are what seem to be Highwire Press standard tags in the upper portion of the displayed metadata and Open Graph ('og:') tags in the lower portion. **C.** Other portions of the embedded metadata record in the Landing Page take the form of Link Headers, with the screenshot highlighting headers of type "alternate." One example of an "alternate" link, highlighted in the red box, is a URL referring to one of the csv files in the Zenodo deposit.*

The Figure 2 legend details the harvesting steps, but several issues should be highlighted. First, in Figure 2A, the doi.org server is called with the DOI in the URL, and HTTP Content Negotiation is used to request Data Cite metadata from the server. While there are methods of advertising what the correct content type is for this Web call (i.e., vnd.citationstyles.csl+json), in the case of DataCite, the HEAD method call on the server returns a response that lacks this piece of information. As a result, the metadata harvester must already know about this option - or what other formats are allowed by DataCite - to retrieve this metadata (or it must guess). Second, in Figure 2B, various formats are used to embed metadata into the Zenodo Web page. This metadata is partially, but not fully, redundant with the metadata provided by DataCite. For example, in Figure 2C, we see metadata that uses the HTML Link Header standard, where the IANA defines the list of approved Link Relations (e.g., "alternate"). According to the IANA, the meaning of "alternate" is as follows:

> *"If the `alternate` keyword is used with the `type` attribute, it indicates that the referenced document is a reformulation of the current document in the specified format."*

This would therefore imply, to a metadata harvesting agent, that the various individual ZIP and CSV files that appear as the links for each 'alternate' keyword are each a reformulation of the entire Zenodo record (because the 'current document' in this context is the Landing Page of the deposit). This would be a troubling interpretation of that metadata if an agent were to make it. It would be straightforward to code an agent to interpret those links as pointers to data as a machine-readable way to identify which links on the Landing Page referred to data versus other kinds of Web hyperlinks; however, this interpretation would be constrained only to Zenodo records. Finally, two additional significant limitations exist in the example described above (and in many other similar examples). First, there is no reliable way for the data author to submit domain-specific metadata that a machine can discover because there is no pointer from the Landing Page to a specific metadata file within the data deposit. Second, the Landing Page itself is a source of ambiguity since the function or purpose of a Landing Page has not been universally defined. The Landing Page is where you arrive after various servers redirect you from the globally unique identifier (DOI in this case) to the end of the redirection chain. Does this imply that the Landing Page is what is identified by the DOI? Or is it the data record that the DOI identifies? Along the same line, does an "author" annotation on the landing page refer to the author of the data deposit or the landing page? This ambiguity cannot be easily managed because there are also DOIs that resolve directly to the data

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

record with no associated Landing Page; thus, the outcome of the resolution process is inconsistent in this regard, making it difficult to generalize a harvesting agent's behavior.

While this fairly typical scenario provides numerous examples of ambiguity and guesswork that a metadata harvesting agent experiences when attempting to comprehensively and accurately gather metadata about this deposit, this is just one scenario! Other ways of providing metadata include:

- **Embedded JSON or JSON-LD**
    - This introduces ambiguity since the "subject" of the JSON[-LD] statements is ambiguous - like the Landing Page problem, there can (and will) be metadata about the data deposit, in addition to there being metadata about the Landing Page itself (e.g., last update, publisher, etc.). These ambiguities are difficult for an agent to overcome.
- **Content Negotiation on the Landing Page**
    - Many providers of Landing Pages will allow Content Negotiation on the URI of the Landing Page itself and use this as an alternative (or in addition to) providing embedded metadata, thus encouraging the use of content negotiation as a means of obtaining metadata about the record. Strictly speaking, Content Negotiation only intends to retrieve a different form of the same resource(for example, a GIF, a JPEG, or a PDF of the same image). It is unclear if the use of Content Negotiation to provide metadata based on a Landing Page URI is truly in compliance with the intended use of this technology.
- **Link Headers and Typed Links**
    - Many providers will use Links - which have two distinct standard syntaxes and may appear in the HTTP Headers ("Link Headers") of the Web message or the HTML Headers of a Web page ("Typed Links") - as a method of publishing metadata, but also (and/or), of pointing to additional metadata records. This may include pointers to individual metadata files, but there is no guarantee that this is a comprehensive list of metadata files known by the data host. In that case, they use a special Link that points to a document called a "LinkSet," which should be the comprehensive list of all metadata linkages known by the publisher.
    - Linksets have two formats - an original text-based format and a more contemporary JSON-based format, and the publisher may use either.

There are other areas of ambiguity or error not detailed above, including that the number of redirection steps between a globally unique identifier and a Landing Page is not defined and cannot be determined by any "signal" provided by the server; thus, an agent might choose to attempts all paths of metadata gathering at each step of the redirection chain. And finally, while major contemporary repositories can generally be relied upon to behave according to Web standards, this is not true of smaller, boutique repositories, which may have been created for very specialist data and are published by domain experts with little experience with Web standards. Thus, common approaches such as Content Negotiation will often be troublesome because the host does not publish that they are capable of providing structured metadata via Content Negotiation, nor which

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

standard they are following, leading to guesswork and trial-and-error on the part of the agent.

Signposting

A straightforward approach to machine-readable metadata provision

Signposting [6, 7] is a set of design patterns that utilize the Web standard of Link Headers and their associated Link Relations to create an unambiguous environment within which an agent can “understand” where in the data/metadata ecosystem they are and how to navigate around it efficiently. It had an initial focus on the “long tail” of data publishers - those who may have the least experience building rich metadata infrastructures - but any participant can utilize the same patterns in the data publishing ecosystem at levels ranging from data infrastructure providers such as DataCite to generalist repositories such as Zenodo, Dataverse, to highly specialized data publishers.

EOSC Association AISBL

Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

2. The Apples-to-Apples Benchmark Environments

The workshops

As part of the EOSC FAIR Metrics and Data Quality Task Force (TF), the FAIR Metrics subgroup works to examine the utility and application of FAIR metrics and the problems with the existing set of FAIR assessment tools. In response to this challenge, the subgroup partnered with the GO FAIR Foundation to co-host a series of workshops[3] where these problems were discussed. The goal of the meetings was to identify and eventually code a shared solution to the metadata harvesting problem that could be implemented for all FAIR metrics and all FAIR assessment tools; thus, the meetings became known as the 'Apples-to-Apples' (A2A) meetings, referring to the English idiom regarding the futility of comparing entirely different things. Attendees represented both TF members as well as other community stakeholders, including representatives from several of the publicly available FAIR assessment tools (The FAIR Evaluator, F-UJI, FAIR Enough/SATIFYD[4] (manual), FAIR Enough[5] (automated), FAIR Checker[6], ENVRI-FAIR[7], and FAIRshake[8] [8]) and coders and designers from a variety of FAIR-oriented data infrastructure and standards projects such as RO-Crate [9], and FAIRsharing[9], matured under and endorsed by the RDA At the end of two half-day workshops, with participant presentations followed by open discussion, the driving question was decided: What is the minimum a data publisher needs to do to lead an agent, unambiguously, from identifier to metadata to data, or *vice versa*. Following a presentation from the Signposting project representatives, the workshop attendees agreed to pursue a limited subset of the Signposting design patterns that appear sufficient to fulfill this core requirement of FAIR. This subset of the full Signposting specification will be referred to as "FAIR Signposting"[10].

FAIR Signposting

FAIR Signposting proposes that a small number of Link Relations are sufficient to unambiguously lead an automated agent to discover and correctly retrieve a digital object's data record(s), unique identifier, and metadata. FAIR signposting consists of a subset of five recommended or required Link Relations from the full Signposting specification. Three of those were selected as the primary focus of the A2A Hackathons: *cite-as*, *describedby*, and *item*. Those Link Relations, and their defined usage, are provided in Table 1.

---

[3] The workshops, and hackathon events took place on the following dates: 2022-02-07, 2022-02-10, 2022-03-17, 2022-03-21, and 2022-04-26
[4] https://satifyd.dans.knaw.nl/
[5] https://w3id.org/fair-enough
[6] https://fair-checker.france-bioinformatique.fr/
[7] https://envri.eu/home-envri-fair/
[8] https://fairshake.cloud/
[9] https://fairsharing.org

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

Table 1: Link Relations used by FAIR Signposting	
Relation	**Usage**
cite-as	A one-to-one relationship between the entity and its globally unique identifier
describedby	A one-to-many relationship between the entity and all known metadata records about that entity
item	A one-to-many relationship between an entity representing a deposit and the data file(s) it contains.

FAIR Signposting further clarifies that a Landing Page should be interpreted by an agent as an information broker, pointing from the landing page to the canonical identifier, the records, and/or the metadata about the record. Thus the agent's interpretation of its environment is consistent and unambiguous, whether the agent arrived at that location from a Web search engine or through a process of identifier resolution. It allows Link Headers to appear in either (or both) the HTTP message headers or the <HEAD> portion of an HTML document, thus opening the possibility for Web software users to add their metadata references even if they do not have access to the Web server itself. Further details of the specification, for example, the recommended use of LinkSets, should be gained by reading the full FAIR Signposting specification [10].

This approach solves many of the outstanding problems faced by (meta)data harvesting agents. For example, if the identifier resolves to a data entity rather than to a Landing Page or if the agent has arrived directly at a data object after a search, it can simply examine the headers to find the identifier of the entity it is currently examining and the links to its metadata. Moreover, because the Links may be in the HTTP headers, it is possible to reliably employ this solution in the context of non-HTML digital objects, such as images or movies, which do not have generalized embedded metadata standards that support outward-facing links.

The Apples-to-Apples Hackathons

Key stakeholders had indicated to us, during the initial workshops, that their willingness to accommodate the outputs from these stakeholder meetings (i.e., to adopt the recommendations in their codebase) was contingent on:

1. A clear, well-crafted, and reasonable specification that demonstrably solved the problem
2. Evidence that the FAIR assessment tools all behaved identically (and positively) when encountering the proposed (meta)data publishing paradigm
3. Evidence that key stakeholders such as funding agencies and the EOSC supported this proposal as a means to maximize the benefit of their digital assets via consistent and reliable measures of their FAIRness.

With the first objective being achieved by the selection of FAIR Signposting in the initial set of workshops, we then focused on the second objective: harmonizing the various FAIR assessment tools' metadata harvesting workflows such that they all responded in the same way when

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

encountering a digital resource that follows the FAIR Signposting paradigm.

The hackathon split into two main groups: the FAIR Metric testing/evaluation group and a second group that focused on creating a comprehensive set of benchmark reference environments against which the evaluation tools could be tested and compared, which are collected in a benchmark repository[10]. The A2A benchmark repository is key as it allows us to continuously design new reference challenges as we encounter different interpretations of the FAIR signposting specification itself and/or as the scope of FAIR expands into new expert domains that may have specific domain standards that must be accommodated by FAIR Signposting. This ensures ongoing uniformity between FAIR assessment tools. Metadata harvesters that adhere to the A2A benchmarks are already available (for example, the FAIR Champion Harvester[11]).

[10] https://github.com/a2a-fair-metrics/a2a-fair-metrics/tree/main/benchmark
[11] https://w3id.org/FAIR-Champion-Harvester

EOSC Association AISBL

Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

3. Conclusions, so far

FAIR Signposting is a Web standards-compliant, easily implemented, explicit, unambiguous, and machine-readable approach to fulfilling the most foundational aspects of FAIR data publishing: Association of data with a globally unique identifier and an explicit association between data and its metadata. FAIR Evaluation systems are now being tooled to consistently and positively score resources published following these recommendations.

Figure 3 shows an interaction comparable to that in Figure 2 but using the FAIR Signposting approach. Because of FAIR Signposting, Content Negotiation is always guided by the content-type instructions available in the Link Header. Thus, it is not necessary to have prior knowledge of the data formats provided by, for example, DataCite, or even that DataCite is a metadata provider, because the agent is led to discover their metadata via the *describedby*.

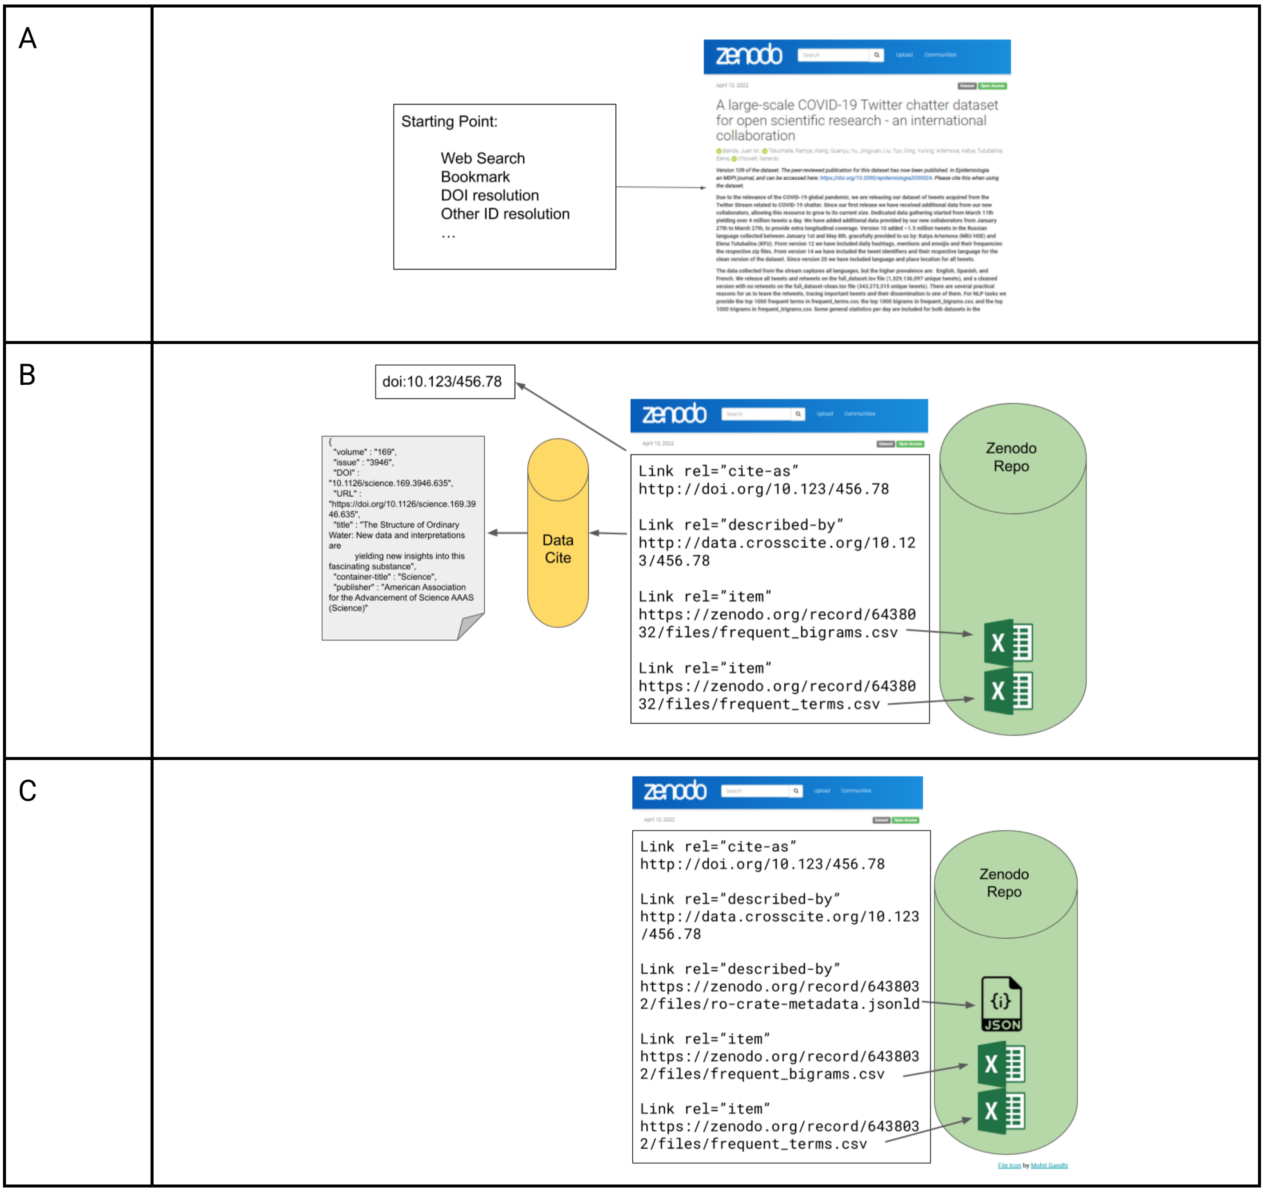

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

D	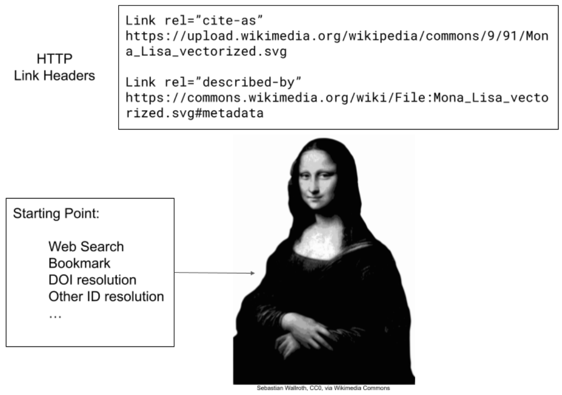

Figure 3. *Examples of metadata harvesting using FAIR Signposting; only a portion of each link header is shown to improve readability.* ***A:*** *A reimagined DOI example similar to that in Figure 2. An agent arrives at a Landing Page through some mechanism (search, bookmark, etc.).* ***B:*** *The Landing Page contains all necessary navigation links to obtain the data, the unique identifier for the data, and the metadata provided by DataCite.* ***C:*** *An example of how a data creator can provide domain-expert metadata (in this case, using an RO-Crate metadata document) that can be discovered and harvested by the agent.* ***D:*** *An example of how FAIR Signposting can be used to provide the same kinds of FAIR metadata for other kinds of digital objects, such as images, through HTTP Link Headers.*

header (Figure 3B). Figure 3B is notable also because it shows how it is now possible to unambiguously point "downwards" from a repository Landing Page to an individual data record within the deposit - the behaviour that was trying to be fulfilled by the *alternate* Links in Figure 2C is now unambiguously fulfilled by the *item* Link Header in Figure 3B. Moreover, Figure 3C reveals another significant advance over the *status quo* in that it is now possible for a data creator to publish their own domain-expert contextual metadata in a manner that can be automatically discovered by following a *describedby* Link Header pointing into their data deposit. Moving beyond citation metadata and making domain-specific metadata discoverable by computational agents will undoubtedly improve the discovery, interpretation, and reuse of problem-relevant data in the future - a key objective for all FAIR stakeholders. Finally, Figure 3D shows that this same approach can be used for a wide range of digital objects on the Web and that, no matter where a computational agent finds itself in the data ecosystem, it is able to navigate from there to all three critical FAIR facets - the unique identifier, the data, and its metadata.

4. Next Steps

Compliance with metadata standards

While FAIR Signposting provides a uniform way for FAIR Assessment tools to gather metadata, that metadata itself should adhere to some community-accepted standards. We suggest that resources like FAIRsharing, as a global registry of community-used data, metadata, and identifier standards for all disciplines, are well-positioned to fulfill the role of being a validator used by FAIR Assessment tools as a look-up service for the thousands[12] of types and names of models/formats, terminologies, minimal information checklists, and identifier schemas (key elements of the FAIR Principles) that FAIR Assessment tools encounter during the assessment process. A limitation on this objective, at this time, is that FAIR Assessment tools can only determine if a research output complies with relevant data and metadata standards by checking against the name and type of that standard, but not its structure or formulation unless these are "hard-coded" into the tool. While FAIR Signposting allows (and encourages) the metadata provider to declare which standard they are following through the use of the "profile" attribute, it is not required. The profile attribute will work best in the case of generic, domain-agnostic standards, such as DCAT or schema.org, which are in common use by metadata catalogues and pages published on the Web, and where robust technologies such as XML schemas or a variety of Linked Data "shape" tests can be automatically applied to validate compliance. This will not be the case for the more than a thousand dataset-level (data and metadata) community standards, such as MAGE-Tab[13] (a model and exchange format for microarray datasets), or CharDM[14] (a model for astronomical dataset characterization) that communities will want to assess as part of their evaluation of FAIR Principle R1.3 (adherence to community standards). Currently, these dataset-level standards could only be tested by profile-name-lookup in FAIRsharing (MAGE-Tab or CharDM) but could not be validated further. Only when standards are 'computable' can we have the necessary quantitative and verifiable measures of the degree to which a research output meets the standard[11]. This is an extremely challenging issue because of its scale and its frequent association with very specialist communities. Then, expert communities will each need to decide the granularity at which FAIR assessments should be made. If computable standards validators are required, they should be developed by those communities and similarly registered in FAIRsharing, such that FAIR Assessment tools can automatically detect, retrieve and apply them.

Bringing in key stakeholders

The workshops described above were carried out by a relatively small group of people who already had a stake in developing, testing, and implementing automated tools for assessing the level of FAIRness of digital objects. We could call these "builders of FAIR assessment tools" one type of stakeholder. A second essential stakeholder type is data publishers, such as data repositories. After all, the data to be assessed are in their holdings, and it will be the repositories that would have to

[12] https://fairsharing.org/standards
[13] https://doi.org/10.25504/FAIRsharing.ak8p5g
[14] https://doi.org/10.25504/FAIRsharing.yYjZWb

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

implement the Signposting link relations in their services. Third, FAIR and Open data are a crucial part of the EOSC's vision and ambition, and exploration of the FAIR assessment tool harmonization problem was a charter objective of the Task Force that co-organized the workshops and hackathons described here. Therefore we see the EOSC Association playing an important stakeholder role, such as driving the improvement and possibly the convergence of FAIR assessment tools through supporting these kinds of straightforward approaches, as well as ensuring that these common practices are embedded in current and future projects.

Furthermore, data producers and data consumers, as well as research funders, have an interest in learning about the level of FAIRness of digital objects. Research funders because they increasingly demand that the data they fund are FAIR; researchers and research teams that produce data because they want to meet such demands or practice Open Science; and – potential – data consumers because the FAIR score of a specific dataset can play a role in their decision to reuse, or not, a dataset. However, these stakeholders are more distant from the technical and policy recommendations made in this report. In the end, they stand to benefit from FAIR scores that are transparent, fair, and well-understood, but they do not have an immediate role in the uptake beyond encouraging the other stakeholder groups by expressing their needs.

We thus identify three *immediate* stakeholder groups:

1. Builders of FAIR assessment tools
2. Data publishers
3. EOSC Association

Concerning stakeholder 1: The FAIR Signposting specification, described earlier, already defines what technical recommendations FAIR assessment tool builders need to follow in their own codebases. An important next step will be the moment when the participants from the earlier workshops can demonstrate that the scores of their respective FAIR Assessment tools show less variation than before. Moreover, we would encourage emerging new FAIR assessment tools to be tested against the same benchmarking environments to reduce divergence in the interpretation and practice of FAIR going forward. This will be an ongoing process in the coming years, with no doubt new input coming from tools that aim to assess the FAIRness level of other research output than data, such as software or workflows, who should add their novel requirements to the A2A benchmarking environment. Finally, we strongly encourage the designers of tests and assessment tools to be fully transparent about their testing process, such that differences between assessment outputs can be more easily examined and explained.

Concerning stakeholder 2: Data publishers, such as data repositories, already play a crucial role in both making and keeping data FAIR: they usually provide the persistent identifier, offer support for – or even demand – standardized metadata and usage licences, and long-term repositories will also retain support for the older file format. With the growing adoption of FAIR ambitions and requirements, this role grows. Data publishers are clearly autonomous in the way they set up their repository and their business model. The decision on whether to implement Signposting, therefore, rests with them. Therefore reasoning along the following lines might be recommended: "The demand for FAIR is growing. Data producers, i.e., researchers, must demonstrate to their funders that the data are as FAIR as possible. Automated FAIR assessment tools have their limitations but will increasingly be used. As a data publisher, you can support your research community by aligning your repository structure so that such tools work optimally, which means implementing Signposting if you haven't done so already. This is in line with recommendations made by the Confederation of Open Access Repositories (COAR[12]) where Signposting is included among their recommended technological solutions." It is worth noting that for many (perhaps most) repositories, Link Headers are already in common usage, as shown for Zenodo in Figure 2C. It is primarily a task of harmonizing agreement around the required links and their expected usage, which we suspect will pose a very low barrier to adoption. In addition, registries, such as FAIRsharing, which participated in the events described here, will need to continue playing a key role in implementing search and retrieval functions for the information needed to assess repositories.

Concerning stakeholder 3: In its report "Recommendations on FAIR metrics for EOSC"[13], the EOSC Executive Board FAIR Working Group adopts an understandably cautious attitude towards FAIR metrics: "FAIR metrics should not be used to judge unfairly; the development of maturity over time and across communities has to be supported and binary judgments avoided." (p. 20) We agree with this and want to highlight that, by extension, this caution carries over from metrics to assessment tools implementing metrics. Typically, FAIR assessment scores need an associated narrative explanation, guidance, and/or consultancy. While this aspect was not in the scope of the Apples-to-Apples workshops, unambiguous benchmark environments will reduce the current inconsistencies in the scores that the various FAIR Assessment tools yield and thereby help to reduce confusion and possible misinterpretation of FAIR scores. Indeed, the consistency that is already emerging through the use of the A2A benchmarks is leading key authors of FAIR to begin harmonizing their error-reporting structures. This will, in turn, allow an ecosystem of visualization and user-support tools to emerge, built on harmonized FAIR assessment tool reports; data users and publishers will potentially gain the most from this since it will encourage the focus of FAIR evaluation to shift from quantitative assessment to qualitative advice, thus reducing the incentive to "game the system". The (cross-)community aspect will also be addressed in new projects, including the FAIR-IMPACT project[15], as well as other generic and discipline-specific projects, and collaboration with domain communities will be critical. FAIR-IMPACT will also work on identifying and mitigating biases between FAIR assessment tools. It is expected that this and other projects will endorse and implement the work by this TF on FAIR Metrics.

The EOSC Association, with its Task Force on Long-Term Data Preservation[16], is the appropriate body for endorsing and promoting the processes recommended here, in which both the repository community and FAIR assessment tool builders put effort into better aligning the provisioning, harvesting, and interpreting of Digital Object information, to achieve optimal and transparent FAIR assessment scores.

[15] https://fair-impact.eu/
[16] https://www.eosc.eu/sites/default/files/tfcharters/eosca_tflongtermdatapreservation_draftcharter_20210614.pdf

EOSC Association AISBL

Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931

Sustainability

With multiple stakeholders having vital and diverse interests in the long-term reliability of FAIR assessment tools, it seems likely that this broad community can maintain the A2A benchmarking environment for a considerable time (it is currently a public, no-cost GitHub repository with a thin server layer around it). A more substantial concern is ongoing curation as new standards, and domain/community-specific evaluation metrics are designed. We suggest that two bodies - existing and emergent - will sufficiently address this burden. FAIRsharing, for example, covers all disciplines and focuses on carefully curating standards, linking them to repositories that implement them and policies that recommend them, and delivering informative knowledge graphs [e.g.,17]. This is done with the help of community curators, including representatives of the EOSC Science Clusters, via the RDA/EOSC-Future-funded FAIRsharing Community Curation Programme[18]. Thus FAIRsharing is an example of a resource that can be relied upon to ensure that newly minted "standards" are formally recognized, genuinely supported by their respective communities, and implemented by repositories. Furthermore, there is an emergent governance initiative around FAIR metrics and testing [1] whose role will be to rigorously curate FAIR metrics emerging from specialist communities and create the A2A benchmarks against which the code implementations of those metrics will be validated. Therefore, it seems that the outcomes of these workshops and hackathons will be sustainable in the short-to-medium term with the efforts of these two bodies.

Acknowledgments

We thank the GO FAIR Foundation for supporting and hosting these meetings, in particular, Erik Schultes and Jacintha Schultes. Meeting participants who are not members of the FAIR Metrics and Data Quality Task Force but without whom none of this would have been achieved: Herbert Van de Sompel, Leonidas Pispiringas, Tim Smith, Sonia Barbosa, Wilko Steinhoff, Allyson Lister, Milo Thurston, Philippe Rocca-Serra, Avi Ma'ayan, Carole Goble, Ceilyn Boyd, Kristian Garza, Peter Doorn, Alban Gaignard, Thomas Rosnet, Stian Soiland-Reyes, Antonis Lempesis, Andreas Czerniak, Luiz Bonino, Michel Dumontier, Vincent Emonet, Robert Huber, Barbara Magagna, Marie-Dominique Devignes.

[17] https://fairsharing.org/graph/1024
[18] https://fairsharing.org/community_curation

References

[1] Wilkinson MD, Sansone SA, Méndez E, David R, Dennis R, et al. Community-driven Governance of FAIRness Assessment: An Open Issue, an Open Discussion, https://doi.org/10.5281/zenodo.7390482
[2] Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, et al. The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data 2016;3:160018.
[3] {Aguilar F, Gomes J, Bernal I, Steinhoff W, Tykhonov V, EOSC-SYNERGY}. EOSC-SYNERGY EU DELIVERABLE D3.5: Final report on technical framework for EOSC FAIR data principles implementation 2022.
[4] Wilkinson MD, Dumontier M, Sansone S-A, Bonino da Silva Santos LO, Prieto M, Batista D, et al. Evaluating FAIR maturity through a scalable, automated, community-governed framework. Sci Data 2019;6:174.
[5] Huber R, Devaraju A. F-UJI : An Automated Tool for the Assessment and Improvement of the FAIRness of Research Data, ui.adsabs.harvard.edu; 2021, p. EGU21–15922.
[6] Klein M, Shankar H, Van de Sompel H. Signposting for Repositories. Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries, New York, NY, USA: Association for Computing Machinery; 2018, p. 355–6.
[7] Van de Sompel H, Rosenthal DSH, Nelson ML. Web Infrastructure to Support e-Journal Preservation (and More). arXiv [csDL] 2016.
[8] Clarke DJB, Wang L, Jones A, Wojciechowicz ML, Torre D, Jagodnik KM, et al. FAIRshake: Toolkit to Evaluate the FAIRness of Research Digital Resources. Cels 2019;9:417–21.
[9] Soiland-Reyes S, Sefton P, Crosas M, Castro LJ, Coppens F, Fernández JM, et al. Packaging research artefacts with RO-Crate. Data Science 2022;5:97–138. https://doi.org/10.3233/ds-210053.
[10] Van de Sompel H, Klein M, Jones S, Nelson ML, Warner S, Devaraju A, et al. FAIR Signposting Profile - Signposting the Scholarly Web 2022. https://signposting.org/FAIR/ (accessed July 5, 2022).
[11] Batista D, Gonzalez-Beltran A, Sansone S-A, Rocca-Serra P. Machine actionable metadata models. Sci Data 2022;9:592.
[12] Confederation of Open Access Repositories. COAR next generation repositories: Technologies. COAR next Generation Repositories: Technologies n.d. https://ngr.coar-repositories.org/technology/ (accessed July 8, 2022).
[13] Commission E, for Research D-G, Innovation, Aronsen J, Beyan O, Harrower N, et al. Recommendations on FAIR metrics for EOSC. Publications Office; 2021.

EOSC Association AISBL
Rue du Luxembourg 3, BE-1000 Brussels, Belgium
+32 2 537 73 18 | info@eosc.eu | www.eosc.eu
Reg. number: 0755 723 931 | VAT number: BE0755 723 931