Making FAIR trustworthy

*Anamika Chatterjee, Norwegian University of Science and Technology, Trondheim, Norway**

Prof (Dr.) Tsjalling Swierstra, Maastricht University, Maastricht, Netherlands

ABSTRACT:	1
<u>INTRODUCTION</u>	<u>2</u>
<u>METHODS AND OUTLINE</u>	<u>4</u>
<u>SECTION 1: THE FAIR PRINCIPLES</u>	<u>6</u>
<u>SECTION 2: SEMANTIC FREEDOM (COSTS TO PRODUCERS)</u>	<u>8</u>
<u>SECTION 3: EXPERT INTUITION (COSTS TO BIOCURATORS)</u>	<u>14</u>
<u>TRUST</u>	<u>17</u>
<u>CONCLUSION</u>	<u>18</u>
<u>REFERENCES</u>	<u>19</u>

Abstract:

The life sciences community is now increasingly leaning on the processing powers of machines to carry out advanced scientific research. So in order to adequately exploit the capabilities of machines in science, the FAIR (Findable, Accessible, Interoperable, Reusable) principles for scientific data management and stewardship have been proposed. These principles are to assist scientists in tweaking their established research routines so as to unlock the true potential of machines and contribute to better science.

However, through interviews with key scientist groups implicated by FAIR we have learned that doing what FAIR demands also presents certain epistemic concerns that discourage scientists to trust FAIR. One such concern is the loss of valuable knowledge in the translation of versatile human readable research output to a restricted, machine friendly language to enable machine action (semantic freedom). A second concern is evident in the displacement of human expertise by this increasing reliance on machines and the ensuing loss of knowledge contributed by these displaced experts (expert intuition).

Thus, through this article, we highlight how incorporating FAIR also presents an epistemic loss to the scientific community. But the lack of attention given to these concerns by the proponents of FAIR offers scientists who have to abide by FAIR sufficient reason to resist it. We thus propose that while the implementation of FAIR has so far been paternalistic and top-down, such concerns have also made the scientist sceptical. So by initiating a more balanced, open and honest discussion of not just the benefits and promises of FAIR but also such epistemic costs, FAIR could lay to rest reasons for such scepticism and foster trust within the stakeholders of FAIR.

36

37 Introduction

38

39 Science is a collective endeavour. Scientists depend on the work and results of others. As with any
40 collaboration, such dependence is inherently risky. Can the information and data collected by other scientists be
41 sufficiently trusted to be used? Of course, in practice scientists don't spend their days worrying about the
42 trustworthiness of their colleagues and their work. In general they are confident about the quality of published
43 work and data. This confidence is warranted because the scientific community has over the years developed
44 manifold ways to reduce the risk of depending on the work of others. A major role here is played by professional
45 ethics, embedded as tacit know how in practical routines, methods, institutions, regulations, and in a scientific
46 ethos that novice scientists gradually acquire by mimicking their colleagues. Thus professional ethics ensures
47 that under normal circumstances explicit and deliberate *trust* is not required, as scientists can safely rely on an
48 implied and routine *confidence* in the other members of their community.

49

50 When business goes as usual, little time and energy needs to be devoted to ethical reflection and
51 deliberation. The problem of trust only arises when practical routines become challenged and destabilized. Such
52 a destabilisation can know many causes. One such cause can be the discovery of a major scientific fraud. A more
53 structural challenge to the traditional scientific ethics comes from a changing socio-economic environment, like
54 increased competition, job insecurity, and financial pressure to produce results. Such factors regularly put
55 pressure on the traditional ethos of a disinterested quest for truth, and raise doubts regarding established norms
56 regulating data-sharing and careful testing. A similar challenge to the traditional scientific ethos comes from the
57 increasing dependence on co-financing by private partners, who work with a quite different, profit-oriented,
58 ethos.

59 In response to disruptions like these, more explicit attention is being given to research ethics and
60 integrity in research organisations. Research ethics no longer only applies to safeguarding the rights and interests
61 of research subjects, like in the classical case of medical research ethics, but it now also increasingly concerns
62 itself with ensuring the independence of the researchers, the way results are shared with the other members of
63 the community, and with the quality of the data. This increased attention for and broadening of research ethics
64 is now manifest in all scientific disciplines. In sum, as a result of various disruptions of the established ways of
65 doing science, the – partly ethical - question 'what is a good scientist' has come to the fore again. And in the
66 midst of this re-positioning of science it is no longer always possible to simply have confidence in the work of
67 one's colleagues. It now sometimes takes a conscious decision whether or not the reported findings of one's
68 colleagues are reliable or not. In other words, routine confidence gets replaced by conscious giving or
69 withholding trust.

70

71 In this article we want to focus on another source of disruption, and the issues of trust connected with
72 it. The professional ethos and ethics of science has come under pressure by the large scale introduction of
73 computers in science, and their rapidly increasing role in collecting and analysing data. This is particularly the

case for the life sciences, that over the past three decades have been profoundly transformed by the omics-revolutions, which were only possible thanks to the rapidly increasing computing power that allowed for the analysis of vast data sets.

This digitalisation of science is profoundly affecting the way science is done, and the ideas about how it should be done. Some within the community loudly and eloquently advocate new configurations of computers and human scientists, asking their human colleagues to adapt their practices to what computers can and cannot do. With the increased importance of collecting and analysing vast datasets, new requirements are formulated with regard to how to do science, and especially, how to handle data. All over the world we now see universities scrambling to come up with new rules, norms, practices, professions, etc. to store, find, handle and reuse data. As the importance of computers in analysing datasets for patterns and correlations increases, new work has to be done to enable them to play this beneficial role. In practice this means that humans are asked to adapt their scientific practices in order that computers can be used with the upmost efficiency and productivity.

We will explicate the new demands posed to researchers in the life sciences by discussing the so-called FAIR initiative. FAIR is the acronym for Findable, Accessible, Interoperable, and Reusable, and it captures the main four quality criteria for data storage (Wilkinson, Dumontier et al. 2016). To meet these new criteria, researchers who produce new data have to adapt their existing routines. We will argue that their willingness to do so has, amongst other things, to do with trust: are they willing to take the risk that they invest in the new data-driven paradigm which may not (sufficiently) deliver on its promises? Do they sufficiently trust the new data-driven science – with its hero role for the computer – to willingly embrace/comply with the new demands that the FAIR paradigm asks of them? This is an important question. Indeed, one can tell scientists that the *FAIR* way of doing science cannot be stopped, and that as scientists they are morally bound to change their ways to further scientific progress. But it is questionable whether the plea from authority or progress will be successful when these scientists are not willing to put their trust into the new paradigm. Scientists are professionals, and have a matching degree of freedom to drag their feet and to frustrate developments. If the new way of doing science is to become a success, it is therefore necessary to solicit scientists' compliance and trust.

We argue that to win the trust of researchers, it is necessary to engage in an open dialogue. This implies that those in favour of the shift to data-driven science cannot confine themselves to 'sending' alone. They will also have to listen – opening up to the concerns of those colleagues who have doubts whether the requested paradigm change will indeed lead to epistemically superior research. Rather than dismissing such doubts off hand as 'irrational resistance' and 'conservatism', they need to be explored in an open dialogue to establish how reasonable they are.

In this article, we want to give voice to two such 'reasonable' concerns with data driven science. Both relate to the overbearing role of computers in the organization of data and in the generation of hypotheses and even conclusions. We will show how in the FAIR discourse there is an unquestioned assumption that the

computer is superior to the human, and that scientific progress therefore requires that humans adapt to the requirements of the computer. From this perspective, resistance against these adaptations is dismissed as stemming from a fear to be made obsolete, or from an affront to the researcher's ego. Such concerns certainly may play a role, as may do a certain conservatism (scientists are only human, after all). But what is more interesting is that there are also *epistemic* concerns raised. The substitution of humans by computers, or their transformation into providers and carers for the computers, may exact a price in terms of scientific knowledge production.

Methods and Outline

In section 1 we present the FAIR approach, focusing on the demands it makes on the researchers to collaborate more closely together with computers. This analysis is based on the literature that has been published on the FAIR initiative since its conception In 2016. This includes three articles: 'FAIR principles for data management and stewardship' (Wilkinson, Dumontier et al. 2016); 'Cloudy, increasingly FAIR' (Mons and al. 2017); and 'A design framework and exemplar metrics for FAIRness' (Wilkinson, Sansone et al. 2018). We then delve into the costs of making data FAIR as have been identified in three reports authored by the Directorate General for Research and Innovation of the European Union. These are: *Cost/benefit analysis of FAIR: Cost of not having FAIR research data* (Publications Office March 2018); *Cost/benefit analysis of FAIR: Policy recommendations* (Publications Office May 2018) and *Turning FAIR data into a reality* (Publications Office November 2018).

In sections 2 and 3 we introduce two specifically epistemic costs that have received little attention in the literature and arise as a result of the ways FAIR envisions ideal relations between computers and human researchers. The first epistemic cost regards the restrictions placed on a researcher's semantic autonomy by the need to make data actionable by a machine. The second epistemic cost regards the loss of human expertise that goes into curating this data when such human roles/tasks are replaced/taken over with/by the machine. These costs become visible through our analysis of eleven semi-structured interviews with members from different cohorts of the new 'digital ecosystem' for biological research. These included data producers and users such as biologists and bioinformaticians, as well as, data stewards such as biocurators and software developers. We questioned them on their view of the strengths, but also the weaknesses, of the human-machine relations as propagated by the FAIR approach. Our analysis further builds on actively following the efforts of the Gene Regulation Consortium wherein members from these cohorts are working towards making data more machine actionable. ("GREEKC," 2016)

In section 4 we draw conclusions, returning to the issues of trust and compliance. We argue that the two epistemic costs we identified cannot be ignored or swiped off the table as mere manifestations of irrational resistance. The way these concerns are discussed and dealt with is therefore crucial for whether at least some groups of researchers are willing to take the jump, and to put their trust in the new way of doing science. There

150 should therefore be an open discussion about the best way human researchers and computers can collaborate
151 to produce sound scientific outcomes. It may well be that after a few years the research community will
152 unthinkingly, that is: confidently, work in accordance with the FAIR principles. But the only way to arrive there is
153 to acknowledge that in the current situation trust is all one can hope for, and the only way to gain trust is by
154 deserving it.

155

156

157

Section 1: The FAIR principles

In *FAIR Principles for Data Management and Stewardship* (Wilkinson, Dumontier et al. 2016) the authors propose the FAIR principles. They describe how the current ‘digital ecosystem’ for scholarly research data is populated with ‘digital objects’ that range from data, to tools, algorithms, and workflows. The reuse of these DOs proves to be a challenging and tedious task for humans, as well as for machines. While the variety of ways to produce, store and manage DOs hinder their reuse through automated means, the quantity of DOs surpasses manual means to do so. As a result, existing DOs too seldom fuel future research and progress. Recognising this challenge, Wilkinson et al. state “There is an urgent need to improve the infrastructure supporting the reuse of scholarly data.” (p.1) Producers of data need to learn to extend adequate care and attention towards their DOs *after* these have been published.

They propose the FAIR (Findable, Accessible, Interoperable, Reusable) principles (from here on, FAIR) to improve management and stewardship of DOs. FAIR aims to enhance so called ‘machine actionability’

“Humans, however, are not the only critical stakeholders in the milieu of scientific data. Similar problems are encountered by the applications and computational agents that we task to undertake data retrieval and analysis on our behalf. These ‘computational stakeholders’ are increasingly relevant, and demand as much, or more, attention as their importance grows. One of the grand challenges of data-intensive science, therefore, is to improve knowledge discovery through assisting both humans, and their computational agents, in the discovery of, access to, and integration and analysis of, task-appropriate scientific data and other scholarly digital objects.” (p.2)

Making DOs FAIR would enable computers to autonomously decide upon DOs and allow them, and humans, to find, access, merge and reuse DOs. This would enhance the value of DOs, further knowledge discovery, and avoid the duplication of research efforts. Ultimately, the authors envision a FAIR ecosystem which would restrict the ambiguity and heterogeneity in how knowledge is represented across the ecosystem, reducing the barriers faced by machines in reusing DOs. They thus urge the entire scientific community to align established research routines and data management methods with FAIR. This diagnosis is echoed in *Cloudy, increasingly FAIR* (Mons and al. 2017): “The recognition that computers must be capable of accessing a data publication autonomously, unaided by their human operators, is core to the FAIR Principles.” (p.3) Again the authors promise that the FAIR ecosystem will enable evidenced research, weed out malpractice, and a wider inclusion of citizens and citizen scientists.

Recognising these attractive promises, many scientific institutions have by now embraced the FAIR principles. According to *A design framework and exemplar metrics for FAIRness* (Wilkinson, Sansone et al. 2018), FAIR has been endorsed by the G20; National Institutes of Health (NIH); Big Data 2 Knowledge (BD2K); Australian National Data Service (ANDS); The European Commission (EC), Science Europe and European Open Science Cloud (EOSC), Research Data Alliance (RDA), as well as, FORCE11. Aside from these, FAIR has also been adopted by

research universities affiliated with Elixir from various countries such as The Netherlands, Norway, Sweden and United Kingdom. Furthermore, the FAIR network now also extends to communities that develop open standards, policy leaders, editors, as well as, publishers like “the Data Seal of Approval, Springer Nature’s Scientific Data, Nature Genetics and BioMedCentral, PloS Biology, The BMJ, Oxford University Press’s Giga- Science, F1000Research, Wellcome Trust Open Research, Elsevier, EMBO Press and Ubiquity Press.” (p.6)

Despite these endorsements not all members of the scientific community have signed up to make their DOs FAIR. This is acknowledged in two reports published by the Directorate General of Research in EU, *Cost of not having FAIR research data* (Publications Office March 2018) and *Cost/Benefit analysis of FAIR* (Publications Office May 2018): “..many research performing organisations and infrastructures are still reluctant to apply the FAIR principles and share their datasets due to real or perceived costs, including time investment and money.” (p.9 & p.5) In a third report, *Turning FAIR data into reality* (Publications Office November 2018), an expert group has compiled FAIR’s progress, status and action plan.

The report *Cost of not having FAIR research data* tries to convince sceptics by demonstrating that currently the EU spends 10.2 billion euros because of *not* working by the FAIR guidelines. The costs are split amongst direct, indirect and opportunity costs. The direct costs are from the additional time spent by a researcher in finding, accessing, preparing and reusing existing DOs. These also include the storage costs i.e. the amount spent by an organisation in safe and secure storage and archival of their DOs. They also appear in license costs i.e. what the organisation currently spends in order to navigate around the various licenses encountered in the reuse of DOs. This also includes the cost of duplicating and repeating research efforts instead of reusing existing DOs. The report also identifies the indirect cost of retracting DOs that contain faulty, unreproducible or poor quality research that could have been avoided if DOs were FAIR. Additionally, the report highlights two opportunity costs of not making data FAIR. Without FAIR interdisciplinary collaboration remains a problem, and without FAIR unlocking new and different avenues for research and innovation economic growth opportunities are missed. To avoid all these costs, the scientific community should comply with FAIR.

The report *Cost/Benefit analysis of FAIR* draws policy recommendations from the first one. In order to benefit from long term advantages of FAIR, the short term costs in terms of time, money and investment must be borne first. The recommendations include ways and means in which funders, research infrastructures and research producing organisations can meet these short-term costs of implementing, complying with, and sustaining FAIR. The report urges stakeholders to think beyond a particular discipline or organisation and develop cooperative strategies to utilise knowledge across these boundaries. It also promotes cultural and behavioural changes needed to implement FAIR, and the use of artificial intelligence and automation for repetitive procedures of research.

Like the other two reports, *Turning FAIR data into reality* observes that compliance with FAIR relies on significant cultural and behavioural changes within the scientific community: “*Making FAIR data a reality requires*

a major change in the practice of many research communities, institutions and funders.” (p.11) This report discusses what is already in place for the FAIR ecosystem and what still more needs to be done. In two key areas considerable change is required: “1) the development, refinement and adoption of shared vocabularies, ontologies, metadata specifications and standards which are central to interoperability and reuse at scale; 2) the increased provision and professionalisation of data stewardship, data repositories and data services.” (p.12) We will return to these two requested changes below. We first want to draw attention to an important point.

All three reports acknowledge the existence of FAIR-sceptics who drag their feet. This scepticism is explained by pointing to the short term costs of introducing FAIR, both in financial and in practical terms. The proposed solutions to overcome this non-compliance are therefore purely financial and pragmatic. So, although the reports accept some resistance to FAIR as *understandable*, nowhere do they seriously explore the possibility that non-compliance may to some degree be *reasonable*. The unquestioned assumption is that FAIR is simply necessary for the progress of both science and society – goals that scientists can be requested to subscribe to. In the end, neither of the three reports seriously explores the possibility that some scientists resist FAIR for scientific, epistemic, reasons.

It is our contention that the purely pragmatic, incentive-based approach employed by FAIR will not be sufficient, as it bypasses these deeper, epistemic concerns in the scientific community regarding whether the transition to FAIR will indeed generate better science. We argue not only that until FAIR squarely faces these epistemic doubts, lack of compliance is likely to remain a problem, but also that the advocates of FAIR miss an important opportunity to learn from their sceptical fellow-researchers. In the next two sections we discuss two such epistemic costs, that arise from this effort to increase collaboration between humans and computers. These costs are related to the cultural and behavioural changes identified in the *Turning FAIR data into reality* report. We will show that according to researchers working in the field, FAIR’s requirement that scientists adopt a ‘shared vocabulary’ not only requires energy, but also comes with epistemic costs in the sense of making communication with colleagues more rather than less efficient and precise. Similarly, some researchers believe that the professionalisation of data also exacts an epistemic price in the form of restricting the space for valuable expert intuition.

Section 2: Semantic freedom (costs to producers)

To analyse the vast quantities of data typical of modern biology, machines are necessary. But to be helped, the biologist first has to enable the machine to help her. She has to prepare the input in a form that is digestible by the machine mind. FAIR therefore requires data producers to change their *semantic practices*. They have to transform the way in which they write, articulate, express and communicate their DOs. Not only does this require considerable extra effort, it also restricts how biologists can convey their findings.

An important role in this discussion is played by so-called contextual or meta-data. It would be difficult for a user to understand isolated or orphaned images, such as one taken under a microscope or captured through an electrophoretic assay. Cells from two different bacteria can look similar unless the user already has some information on the context. The same goes for orphaned or unlabelled graphs and tables. So any description should contain instructions on how the DO ought to be read and interpreted. Without this contextual description, a user could attach little significance or meaning to the DO, let alone reuse it. This context provides the user with a story within which data can be placed. It informs their judgements on the validity, quality and value of this data.

Let's first consider the situation before FAIR. Here a biologist can use quite some natural language to describe her data and context. This allows her considerable freedom and flexibility in how she represents knowledge within her DO. She can tweak, rearrange and adjust the language in a way that best communicates her findings (Nowak, Komarova et al. 2002). An ordinary knowledge of the alphabet, words, sentences, grammar, and common word processing tools such as Microsoft Word, suffices for a coherent and detailed account of one's research (Case, 1985). For instance, who was involved in the research and where was the data produced? What is their expertise or their research background? For whom and what purpose was the research conducted? What research gap(s) did the researchers aim to bridge? How did they do so? How do they justify their choice of methods, tools, reagents, specimen? How do they explain the figures, images and tables? How did they analyse this data? What inferences and conclusions do they draw from it? Piecing this information together gives a richly defined context around the DO and allows the biologist to target a specific intended audience for her DO. So, she can comfortably type out her text; create diagrams, tables and graphs; insert pictures of her results to her text; and cite and refer to the work of her colleagues, without having to explicate all nuances of her work.

This semantic freedom also enables the biologist to address specific users. She can bank on the data user's prior knowledge to use complex jargon, besides natural language. She can often build on what the intended readers already know as these are often colleagues with similar expertise. Some things can be left implicit, as the intended reader will fill in the blanks. Depending on her intended audience the biologist can explain the context in many different manners. Let's assume she needs to contextualise an interaction such as 'Human transcription factor AKT1 reduces the activity of target gene FOXO3 by attaching a phosphate group to its amino acid ser305'. For users interested in the causal direction of the interaction she could say 'Human transcription factor AKT1 *downregulates* target gene FOXO3 by attaching a phosphate group to its amino acid ser305'¹. These users would understand what 'downregulates' means. But when catering to users with knowledge of transcription she could say 'Human AKT1 downregulates FOXO3 by attaching a phosphate group to its amino acid ser305' because this type of user would already know that AKT1 is a transcription factor and FOXO3 is the target gene. Or if the users are well versed with biochemical processes she could modify the statement to 'Human AKT1 downregulates FOXO3 by *phosphorylating* its amino acid ser305'. Or if the users

¹ Statement taken from [Signor Database](#);

would already know much about amino acids she would then simply say ‘Human AKT1 downregulates FOXO3 by phosphorylating its ser305’.

Before FAIR, the biologist could use both natural language and jargon, and she could tailor her communication to specific audiences. This semantic freedom is incompatible with FAIR. While humans can rely on their semantic intuition to interpret and process semantically rich and diverse DOs, machines are not so capable (Wilkinson, Dumontier et al. 2016). Different styles of representing the same bit of knowledge baffle the machine. Natural language in the form of synonymous words, conjugated verbs, upper/lower cases, special characters, icons or bullet points – all characteristics of a naturally written DO - is far too complex for the machine:

“machine interpretability of all information about data will be a priority requirement. For example, FAIR Digital Objects will need to be ‘typed’, so that by reference to a given information source, a machine can determine what operations are possible against a given data type of which the data in question is confirmed as an instance. To facilitate machine processing at scale, metadata will need to be even more elaborate, and all relationships will need to be stable and semantically defined.” (*Turning FAIR data into reality*, p.45)

FAIR thus requires that producers make the contextual descriptions for their DOs more formal and that they cater to the *machine’s* capabilities instead of those of *human* users. Such formal knowledge representation styles are developed through grounding natural language within mathematical rules and logic which can then be shared with machines. Take for example formal ontologies. These offer fixed and specific definitions of the concepts common within a domain. Each of these concepts is then interlinked in a hierarchical order representing the logic underlying an ontology. Such ontologies help reduce the ambiguity and confusion around what is implied by a concept within different research projects (Stevens, Goble et al. 2000).

Formal languages are complex and demand considerable work and effort to apply. As one of our interviewees explains: “computers are stupid! So you have to give them very precise information. Because if you just use an upper-case somewhere, the information could be lost to the computer.” [BC04] Another interviewee explicates how the same verb when conjugated differently appears as two distinct words to the machine “a computer won’t know that ‘phosphorylates’ and ‘phosphorylated’ indicate the same process of phosphorylation.” [BC11] She also explains how numerical or special characters such as hyphens (-) prove problematic for the machine. While these are commonly used in representing cell lines, species, measured readings or labels, they are often used differently in the same word. “Cell lines should also be specified in the same way. Because a cell line could be written with or without a hyphen.” [BC11]

These restrictions of formal language and formats stem from the machine’s sterile skills. The researcher now has to explicate everything in minute detail. She has to translate ‘Human transcription factor AKT1 reduces the activity of target gene FOXO3 by attaching a phosphate group to its amino acid ser305’ in such a manner that a machine can comprehend it. She needs to be mindful of cases, conjugations, special characters, free text spaces, fixed terminologies, definitions, and other do’s and don’ts of the format. Does the machine know what

she means by 'human'? By 'transcription factor'? 'Target gene'? 'Phosphate group'? 'Amino acid'? 'Ser305'? Does the machine understand that the 'species' is 'human' here? Could she write 'AKT1 *downregulates* FOXO3' or that 'FOXO3 is *downregulated* by AKT1'? What if she writes 'FOXO3' instead of 'FOXO3'? Or 'AKT-1' instead of 'AKT1'? Finally, does the machine's knowledge of these conceptual definitions align with hers?

As a result of these complications many data producers resist attempts to incorporate formal language in their routines. In 2008, *FEBS letters*, an open access journal, along with Elsevier undertook an experiment to assess authors' (producers') capabilities and inclination towards providing structured, formalised, machine actionable DOs (Arnaud, Andrew et al. 2008). For the *FEBS letters experiment* a cohort of 57 authors were invited to produce *Structured Digital Abstracts (SDAs)* for a published article. The results showed that authors took on average 68 mins to do so. While some authors raised concerns of this task delaying the publication process, others complained about the insufficient guidance they were given to complete the task. The authors proved insufficiently competent and knowledgeable about SDAs and how digital entries are made. Their submissions were assessed as being of low accuracy as compared to those of professional curators (Valencia, Andrew et al. 2008). One of the biocurators we interviewed recognized this lack of knowledge and competence:

"The idea was that the authors were supposed to write an abstract using terms from formal language. So information such as A interacts with B physically. And the authors just refused point blank to do it or messed it up because they didn't understand what we (biocurators) were doing. In a lot of cases it's just ignorance." [BC09]

Another interviewee added that the complexity of the task discourages authors.

"Well (formal language) is difficult.. it's a chore. That's what it is. It's something you have to be aware of. And it's already difficult for us (biocurators) to put people in the same line as to how to capture things in a very strict, very formal language that reflects a very specific scientific problem that is already super complex and its already a pain. Imagine if you want to do that for any scientific concept you can think about. So it's just a lot of work." [BC06]

The extra work demanded by FAIR is not trivial. Data producers would have to invest considerable time and effort into first learning the ropes of the formal language before they can effectively begin to use it for their DOs.

So, to an important extent, the resistance to formal language stems from the hesitance to take on this additional work. In the long run this type of resistance may indeed be overcome by providing education, financial incentives, and regulations. This is less evident of a second concern regarding the need to use formal language. According to some researchers, formal languages result in an epistemic loss. While they do enable machine action they also hinder the producer's efforts to *adequately represent* diverse, heterogeneous biological knowledge. As the proponents of formal language admit:

"Formality, both in the ontological and representation language sense, is a stern friend. A formal language has a well-defined interpretation of the world and a well-defined language with which to say things about that world ...As such, the formality (strictness) of the language is good—this precision

means that automatic reasoning with the symbols of the ontology can take place. This very strictness is, however, potentially restrictive” (Bodenreider and Stevens 2006, p.15)

This point was also stressed by our interviewees. One important issue is that much information has a *relative* character. Let us first illustrate this with a simple example. One standing on earth looking out at the moon believes the moon orbits the earth. But to one standing on the moon, the earth could be said to orbit around the moon! This perspectival relativity is quite common in biology. As one interviewee explains, “a protein modifies a certain function of another protein and this is an upregulation or a downregulation, because it modifies a specific molecular function. So, protein A upregulates protein B because it upregulates the cytoplasm localization of protein B. This is an *upregulation*. But from the opposite perspective, if I know that the protein in the cytoplasm is far from this function, then I know that the flux of the information has stopped, then I would say that this is a *downregulation*.” [BC04] This relativity is unworkable for formal systems, where terminologies and definitions need to be fixed. In the case of the earth and the moon there is a general consensus that the moon orbits the earth. But such a consensus is not always available. So unless there is a consensus between the producer and user on such concepts, structured DOs may be misinterpreted by the machine.

There is another related epistemic loss here. FAIR also restricts who the producer can intend her DO for. As the GO FAIR Initiative states “The rationale behind this [findable] principle is that someone should be able to find data based on the information provided by their metadata, even without the data’s identifier.” (FAIR 2019) But as FAIR is about machine actionability, this rationale becomes more a matter of *something* finding the DO rather than *someone*. By being FAIR DOs will be intended only for machines. This will subsequently limit the intended users to those who rely on machine actionability for their research.

The rigidity of formal formats involves a third epistemic loss. As an interviewee told us: “It’s quite frustrating that you cannot add another column to the tabular format and just include this additional bit of information that you now need.”[BC07] To manoeuvre around this frustration, data producers resort to the use of ‘free’ text fields to include all that cannot be included in the fixed columns and fields of the format. Free text fields allow producers more freedom in representing those bits of knowledge that are inadequately captured by the formal language and in the formal format. As one bioinformatician explains,

“When I ask the producers why they didn’t use other terms from the formal language that represents ‘phosphorylates’ they would tell me that there were two or three terms that were missing from the formal language but were needed by them, the producers, in order to explain their DOs. So, they preferred the free-text field instead. They also told me that they had talked to the developers of the formal language and format but that their suggestions weren’t implemented. [...] For a bioinformatician having such a text field is a nightmare! It requires sophisticated text mining procedures to extract meaningful data from this field. So for example a text field could say A phosphorylates B and inhibits it. But because ‘phosphorylates’ is not an ontology term this information has been put into a free text field.” [BC12]

419 To avoid natural language polluting the formal language, more and more formats now restrict the use of free
420 text spaces. From the perspective of machine-readability this may be a defensible strategy. But from an epistemic
421 perspective, such a restriction amounts to a loss of data, information, and knowledge. It is understandable that
422 some scientists protest the FAIR rules, as being detrimental to the aim of science: producing sound knowledge.

423

424 In sum, although, FAIR insists that it is crucial for contextual descriptions to be 'generous' and 'extensive', it does
425 require them to be formal and machine actionable. FAIR's restrictions on the semantic freedom of a data
426 producer involve epistemic costs. Complying with FAIR will influence how producers describe their DOs, and how
427 and by whom their DOs will be interpreted. The exclusion of information that cannot be accommodated by
428 formal language is for some researchers a reason for non-compliance.

Section 3: Expert intuition (costs to biocurators)

If the first cost of FAIR arises in matching a human's communicational expertise to that of a machine, a second cost arises with the attempt to transfer specialist expertise to the machine. This results, as we will show, in the displacement of *the human expert's intuition*. As noted earlier, the authors of *Turning FAIR data into reality* acknowledge that in modern data-centric biology a crucial role is played by so-called data stewards. These are defined as those experts who have "a set of skills to ensure data are properly managed, shared and preserved throughout the research lifecycle and in subsequent storage. Wilkinson et al 2016 similarly stress that these humans are key to achieving the ultimate goal of FAIR for "...machines to be capable of autonomously and appropriately acting when faced with the wide range of types, formats, and access-mechanisms/protocols that will be encountered during their *self-guided exploration of the global data ecosystem*." (p.4) (Wilkinson, Dumontier et al. 2016) During the active research process, this could involve data cleaning to remove inconsistencies in data sets, organising and structuring data, adding or checking metadata, and resolving data management issues." (p.46) So, FAIR raises the demand for human data stewards. But, and this is the crucial point: this is supposed to be a temporary phase. FAIR's premise is that ultimately the role of these data stewards will eventually be taken over by the machine.

We have interviewed a particular group of data stewards, called 'biocurators'. "Biocuration involves the analysis, interpretation and integration of biological information into data repositories, primarily to add value by annotating and interconnecting research data and results within a common biological framework." (Biocuration 2018) To convert human readable DOs to machine accessible versions, these experts go through numerous DOs pertaining to a particular research problem, extracting relevant knowledge from them and presenting this knowledge in a structured, formalised manner. They also provide a comprehensive overview of the existing knowledge within a particular research domain. In these overviews they identify the potential knowledge gaps and shed light on what has been done, what needs to be done, and how it can be done. Thus, their contribution serves the dual purpose of enabling machine action, as well as furthering the overall state of knowledge within a domain.

The promise of FAIR is that these roles will eventually be taken over by the machine. However, how likely is it that a machine can imitate the epistemic contribution of biocurators? In this section, we will explore the expertise employed by these biocurators, and argue that this probably cannot be mimicked by the machine. If the machine substitutes these humans, this would constitute a second epistemic loss.

Biocuration requires considerable expertise. For example, when a biocurator has to curate a DO that captures complex molecular interactions, she would first have to find the DO. For this she needs to know her way around a computer, a search engine, a PDF reader like Adobe, a text application like Microsoft Word, and so on. She also needs to be proficient in the language of the DO, say English. This proficiency not only helps her conduct the search but also determines how well she can read, comprehend and interpret the DO. She also relies on her more specialist expertise to assess and assimilate DOs, and to enrich the total state of knowledge. She needs to be able

to gather crucial details about the DO and to assess its overall quality, relevance, and utility towards solving the problem. She needs to be able to trace information such as the species, the construct, the identifiers, or the standard nomenclature of the entities involved in the research. Or to spot mistakes and gaps in DOs such as incorrectly labelled figures or missing supplementary material that could contribute to faulty results during reuse. She needs to be able to assess if a DO employs standardised terminology, languages and formats. Only based on all this varied expertise, she can eventually make an expert judgement regarding whether the DO meets the specific quality, technical and epistemic criteria relevant to the problem.

Moreover, she is in a unique position to provide a comprehensive overview of the knowledge available within a research domain. Based on her reading of numerous DOs, she is able to draw meaningful connections, correlations and patterns between several DOs. Thus she can provide an integrated, interconnected and layered review of the area enriched with her own understanding of the specific problem (Rhee, Maria et al. 2008). It is perhaps for this reason that biocurators have been described as “the museum cataloguers of the Internet age: they turn inert and unidentifiable objects (now virtual) into a powerful exhibit from which we can all marvel and learn.” (Bourne and McEntyre 2006) The International Society for Biocuration also propagates this view in *Biocuration: distilling data into knowledge* (Biocuration 2018): “Biocurators strive to distil the current ‘best view’ from conflicting sources..” The biocurators’ expertise facilitates them to manage and make an expert judgement regarding the worth of a DO and to make an epistemic contribution by providing a rich overview of a problem area.

FAIR claims that eventually all this expertise will be transposed to the machine. Expert systems based on algorithms will assess, judge and mimic human expertise to make decisions regarding the data they encounter (Ericsson, Hoffman et al. 2018). Such expert systems are expected to mimic the decision making process of a human expert. If that proves to be true, machines will then be responsible for the management, structuring and proofing of DOs, as well as ensuring it is FAIR. The expert system will check DOs for inconsistencies; ensure DOs have employed standard formats, languages and software; demand the required information such as species and identifiers; and essentially prepare DOs for the FAIR ecosystem. For this technological dream to materialize, the entire decision making process of a biocurator will have to be explicated and translated. Not only how she makes her assessment and judgement of a DO, but also how she provides an overview of the research area.

It is however highly doubtful whether such algorithms will ever be developed (Dreyfus and Dreyfus 2009 [1986]), and they will certainly not be in the short run. Although the domain of AI has made impressive progress in the past decades, especially thanks to the Big Data revolution, it is hotly debated whether human know how can ever be adequately captured in the form of logical rules. And so far, in the case of biocuration, manually curated content has been associated with high quality, reliable and trustworthy knowledge as compared to that curated by a machine. As the *FEBS Letters* experiment (Valencia, Andrew et al. 2008) confirms, “..given the results of the systems on interaction pairs, it is also clear that complex curation needs cannot be transferred to purely automated approaches.” (p.3) The experiment demonstrated that the accuracy of human curation couldn’t be matched by that of automated curation. Most researchers in the field still regard manual curation as the gold

standard and employ a team of manual curators to filter content that has been automatically curated (Bodenreider and Stevens 2006, Consortium 2018).

Our interviewees supported this sceptical view regarding the chance that biocuration could be automated. As biocurator BC06 explains, automated curation presents a high risk of false positive knowledge entering the database: “All the automatic methods we have used to mimic what we do absolutely suck at it. They are really bad. They have the problem of identifying unreliable information and they introduce an amount of mistakes and are rife with false positives. For example, there is a gene called SDS and SDS is also a detergent very commonly used in the laboratory. But who else but a human can distinguish between SDS the detergent and SDS the protein? And not even every human can do that..” [BC06] Another biocurator gives another example. “Sometimes even a small or low throughput paper can take hours to curate. Because you cannot tell what they are working with. And sometimes they say ATR but ATR could be 5 different acronyms. And then you have to read the entire paper carefully hoping that there is some indication on which of these proteins the authors are talking about.” [BC10] Both confirm that such confusions can only be detected and solved by a human biocurator using her expert intuition, informed by years of accumulating experience.

Besides false positives there are also false negatives. Automated curation methods could also reject DOs that are actually valuable and relevant. As they rely on rigid, inflexible rules such DOs get discarded from their results owing to minor technicalities or a misinterpretation by the machine. BC01 indicates how the machine could reject contradictions and refutations in DOs, thus requiring a human expert to step in and solve the problem. “There are a lot of refutations. Science is based on that. You discover something, you put the theory on a paper and if somebody is repeating the experiment he may say “ok, it’s right” or “it’s not”. The computer cannot decide if this is right or wrong. You have to take a look based on general knowledge. So if we find any paradox or any contradiction we try to cope with it based in our previous experience and knowledge.”[BC01] Moreover, this decision rests on the gut feeling (or intuition) acquired through years of experience and discussing these issues with colleagues. This is what enriches their holistic knowledge of the field so as to *distil the best view*. But this intuition is hard to explicate, capture and formalize in a set of rules for the machine. Assessing the quality of a DO may be less a matter of applying the right algorithms than a matter of skill: impossible to formalize, but nonetheless non-arbitrary.

Instances such as these exemplify the value and importance of human expertise in carrying out biocuration. FAIR underestimates biocurators as a source of valuable knowledge. Their expertise cannot be adequately captured through formalised, structured, machine comprehensible rules and guidelines. The replacement of human biocurators with machine methods would therefore result in an epistemic loss, and it is therefore reasonable to resist their replacement by a machine.

Trust

In the previous two sections we discussed two *epistemic* costs of introducing FAIR –restrictions on semantic freedom and the displacement of the human expertise of the biocurators. Both contribute substantially to scientific knowledge. Both are incompatible with the ultimate goal of FAIR to construct an interoperable, machine actionable ecosystem for scholarly research data. Yet, as far as we were able to establish, neither of these two epistemic costs have been recognised or acknowledged, whereas the other, financial and pragmatic costs did receive attention. So what does this have to do with trust in FAIR? When legitimate concerns in the field are ignored by the proponents of FAIR, this is a cause for distrust and non-compliance. More precisely: not recognising these epistemic costs of FAIR easily results in a decline in institutional trust. Institutional trust is the trust people have in institutions that exercise power and authority over them. As in the case of citizens who have to be able to trust the public institutions of their country. Or researchers who have to be able to trust academic institutions such as universities or research establishments.

In any trust relationship the trustor accepts being vulnerable to certain risks. As Annette Baier writes in *Trust and Antitrust*

“One leaves others an opportunity to harm one when one trusts, and also shows one’s confidence that they will not take it. Reasonable trust will require good grounds for such confidence in another’s good will, or at least the absence of good grounds for expecting their ill will or indifference. Trust, then, on this first approximation, is accepted vulnerability to another’s possible but not expected ill will (or lack of good will) toward one.” (p.6)

In a successful trust relation, the trustor expects that these risks have been mitigated by the trustee. This also applies to the relation between researchers and biocurators on the one hand, and the proponents of FAIR on the other. And a minimum requirement for the trustee to earn the trust of the trustor, is that the former recognizes and acknowledges the risk perception of the latter. You cannot demand someone’s trust if you are not willing to take the other’s concerns seriously.

The lack of timely disclosure and preparation by FAIR’s advocates of the two, plausible, epistemic costs we identified, weaken the trust relationship with some of the researchers and biocurators. They don’t feel their concerns are taken seriously. As a result it would be unreasonable for them to trust that the FAIR system is adequately prepared to deal with these epistemic costs when they become manifest. Researchers and biocurators could rather believe they have good reasons for caution, scepticism and even doubt on whether the institution has been completely honest and transparent with them. Or whether the FAIR institution does have the best interests of science in mind. This lack of trust easily helps explain the lack of compliance with and the resistance to FAIR.

How to avoid this distrust? FAIR will have to become more sensitive to the epistemic costs of FAIR, and that non-compliance is therefore not merely understandable but – at least possibly – also reasonable in the sense that

non-compliance can be informed by the same values of good science that FAIR's proponents adhere to. Trust can only be built by accepting this symmetry between those in favour of FAIR and those being sceptical about (aspects of) it. To earn trust it doesn't suffice to top-down request compliance. Instead, it will be necessary to step down into the arena and initiate an open dialogue with the trustors about their epistemic or other concerns. Such a stakeholder dialogue could well help uncover even more such concerns. Part of the implementation of FAIR should therefore be the creation of a platform where affected researchers and other stakeholders can voice and collectively explore the potentials, but also the risks of transforming biology into a machine-dominated enterprise. In the end, only dialogue constitutes a viable basis for mutual trust. And only when stakeholders trust each other, it will be possible to develop collaborative strategies to constructively deal with concerns – by devising technological solutions directed at the machines and by devising social solutions directed at the humans. Or even only making transparent inevitable trade-offs between incompatible values..

Conclusion

The ecosystem for scholarly research DOs is in dire need of management and stewardship. FAIR proposes principles to enhance the Findability, Accessibility, Interoperability and Reusability of DOs. Humans can no longer do this without the help of machines, and this is why FAIR foregrounds machine-actionability. However, this transformation of the research system not only comes with pragmatic and financial costs, but may also hamper the goal of producing valid new scientific knowledge. We showed how machine-actionability not only puts strict demands on the researchers, but may also severely hamper their mutual communications by disallowing them to apply context-sensitive natural language geared towards specific audiences with a pre-understanding of the subject matter. In the end this restriction of their *semantic freedom* has to be acknowledged as an epistemic cost of the FAIR transition. Only such an acknowledgement can provide the basis for a constructive dialogue about the question whether this cost is or is not too high a price for the epistemic benefits that machine-actionability undoubtedly also brings. Similarly we showed that human expertise – we discussed biocurators but the point is more general - may well not be translatable in machine-readable algorithms. The biocurators apply both articulate knowledge and intuitive skills when assessing a DO. Their judgement is not technical in the sense of checking a list of predefined criteria. The attempt to replace human experts by machine intelligence is therefore likely to incur a second epistemic cost, in the form of a loss of valuable *human, partly intuitive, expertise*.

From our interviews with researchers, bioinformaticians, and biocurators it became clear that these epistemic concerns do inform their doubts and sometimes non-compliance with FAIR. They don't trust the proponents of FAIR to take these concerns sufficiently seriously. While they accept the basic premise of FAIR as reasonable, they think that it is important to consider these epistemic costs as well. Instead of implementing FAIR in a somewhat paternalistic and top-down fashion, the implementation of FAIR demands an open and non-prejudiced discussion about both the epistemic benefits and the epistemic costs. In the end such a discussion is

the only way to generate sufficient trust amongst members of the research eco-system to develop new collaborations between humans and machines.

References

- Arnaud, C., C.-A. Andrew, L. Luana and C. Gianni (2008). "Linking entries in protein interaction database to structured text: The FEBS Letters experiment - Ceol - 2008 - FEBS Letters - Wiley Online Library."
- Biocuration, I. S. f. (2018). "Biocuration: Distilling data into knowledge." PLoS Biol **16**(4): e2002846.
- Biocuration, I. S. f. (2018). "Biocuration: Distilling data into knowledge." PLoS biology **16**(4): e2002846.
- Bodenreider, O. and R. Stevens (2006). "Bio-ontologies: current trends and future directions." Brief Bioinform **7**(3): 256-274.
- Bourne, P. and J. McEntyre (2006). "Biocurators: Contributors to the World of Science." Plos Computational Biology.
- Consortium, U. (2018). "UniProt: a worldwide hub of protein knowledge." Nucleic Acids Research **47**(D1).
- Ericsson, K. A., R. R. Hoffman, A. Kozbelt and A. M. Williams (2018). The Cambridge handbook of expertise and expert performance, Cambridge University Press.
- FAIR, G. (2019). from <https://www.go-fair.org/fair-principles/f2-data-described-rich-metadata/>.
- Mons, B. and e. al. (2017). "Cloudy, increasingly FAIR; revisiting the FAIR Data guiding principles for the European Open Science Cloud." Information Services & Use **37**(1): 49-56.
- Nowak, M. A., N. L. Komarova and P. Niyogi (2002). "Computational and evolutionary aspects of language." Nature **417**(6889): 611-617.
- Publications Office, E. U. (March 2018, 2019-01-16). "Cost-benefit analysis for FAIR research data : cost of not having FAIR research data." from

<https://op.europa.eu:443/en/publication-detail/-/publication/d375368c-1a0a-11e9-8d04-01aa75ed71a1>.

- Publications Office, E. U. (May 2018, 2019-01-16). "Cost-benefit analysis for FAIR research data : policy recommendations." from <https://op.europa.eu:443/en/publication-detail/-/publication/d3766478-1a09-11e9-8d04-01aa75ed71a1/language-en>.
- Publications Office, E. U. (November 2018, 2018-11-26). "Turning FAIR into reality : final report and action plan from the European Commission expert group on FAIR data." from <https://op.europa.eu:443/en/publication-detail/-/publication/7769a148-f1f6-11e8-9982-01aa75ed71a1/language-en/format-PDF>.
- Rhee, D. H., C. Maria, F. Petra, G. Takashi, H. Linda, H. Winston, P. H. David, K. Renate, S. Mary, P. Susan St, T. Simon, W. Owen and Y. Seung (2008). "The future of biocuration." *Nature* **455**(7209): 47-50.
- Stevens, R., C. A. Goble and S. Bechhofer (2000). "Ontology-based knowledge representation for bioinformatics." *Brief Bioinform* **1**(4): 398-414.
- Valencia, F. L., C.-a. Andrew, A. M. Scott, C. Arnaud, K. Martin, L. Luana, H. Lynette, C. Gianni and Alfonso (2008). "The FEBS Letters /BioCreative II.5 experiment: making biological information accessible." *Nature Biotechnology* **28**(9): 897-899.
- Wilkinson, M. D., M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos and P. E. Bourne (2016). "The FAIR Guiding Principles for scientific data management and stewardship." *Scientific data* **3**.
- Wilkinson, M. D., S.-A. Sansone, E. Schultes, P. Doorn, L. O. B. da Silva Santos and M. J. S. d. Dumontier (2018). "A design framework and exemplar metrics for FAIRness." *Scientific Data* **5**.