International FAIR Digital Objects Implementation Summit 2024

Full Papers

https://doi.org/10.52825/ocp.v5i.1419

Published: 18 Mar. 2025

The Five Safes RO-Crate

FAIR Digital Objects for Trusted Research Environments for Health Data Research

Stian Soiland-Reyes[1,*], Stuart Wheater[3], Thomas Giles[4], Jonathan Couldridge[4], Philip Quinlan[4], and Carole Goble[1]

[1]The University of Manchester, UK

[2]University of Amsterdam, The Netherlands

[3]Arjuna Technologies, UK

[4]The University of Nottingham, UK

*Correspondence: Stian Soiland-Reyes, soiland-reyes@manchester.ac.uk

Abstract. Trusted Research Environments (TREs) are secure locations in which health and other sensitive data are placed and made available for researchers to analyse under strict controls. TRE's in the UK operate under the Five Safes governance framework of safe data, safe people, safe projects, safe settings and safe outputs to protect data confidentiality. However, there is no standardised mechanism for streamlining the exchange of the metadata needed between analysis toolkits and TREs to follow Five Safes procedures. This lack of standardised interoperability is exacerbated when undertaking federated analysis across multiple TREs. The "Five Safes RO-Crate" digital object is a proposed approach for packaging the metadata needed for exchanging research requests and results between analysis tools and TRE providers, enabling them to operate Five Safe compliant processes. The approach has been piloted by the DARE UK TRE-FX project with commercial and open-source analysis toolkits and two health data TREs. The work will continue to be developed in Health Data Research UK's Federated Analytics work programme and incorporated into the TRE Blueprints currently being developed by EOSC-ENTRUST European Network of Trusted Research Environments and DARE-UK. Five Safes RO-Crate is an important component of the metadata middleware necessary for implementing scalable TRE federated analysis.

Keywords: RO-Crate, Trusted Research Environment, Five Safes Framework, Secure Data, Sensitive Data, HDR UK, Health Data, DARE UK

1. Background

Trusted Research Environments (TREs) [1] are secure locations managed by data custodians in which data are placed for researchers to analyse. TREs can be set up to host administrative data, hospital data or any other data that needs to remain securely isolated, and access controlled. Data is never moved. Instead, data remains in situ within

locally controlled TREs and is "visited" by analysis tools that have been pre-approved by data custodians and securely on-boarded into the TREs.

The UK TREs work within the "Five Safes" data governance framework [2]. Originated by the UK's Office of National Statistics [3], the framework is a set of principles that aim to protect data and enable data services to provide safe research access to data in full GDPR compliance (Table 1). In a nutshell, specific people are authorised for a specific project on specific data in a specific safe setting, and the disclosure of results is specifically screened, usually by people but in future semi-automatically given the appropriate metadata and policy expressions such as Open Digital Rights Language (ODRL) [4].

Table 1. *The Five Safes Framework; dimensions are scales within which a TRE can design and report "safe use".*

Dimension	Description	Action
Safe Data	Awareness of a disclosure risk in the data itself	Data is treated to protect any confidentiality concerns.
Safe projects	Projects make appropriate use of the data	Research projects are approved by data owners for the public good.
Safe people	Users can be trusted to use the data appropriately.	Researchers are trained and authorised to use data safely.
Safe settings	The access facility limits unauthorised use.	A secure lab environment prevents unauthorised use.
Safe outputs	The statistical, aggregated results do not disclose confidential information.	Outputs are screened and approved by non-disclosive validation.

Since 2020 the Five Safes has been the overriding framework for the design of new secure facilities and data sharing arrangements in the UK for public health and social sciences. It has become best practice in data protection whilst fulfilling the demands of open science and transparency, extending its reach beyond the UK with take-up in Canada, Australia and Europe [2].

1.1 Federated Analysis using multiple TREs

Federated analysis across multiple TREs, requesting and gathering the data needed from each one, widens the scope of research and makes more effective use of data. Researchers remotely visit data, run their same analysis at each site, and receive a local result, which can then be aggregated (Figure 1). However, geographical or governance boundaries, for example in devolved healthcare in the UK and across national borders in Europe, make this challenging. The current UK TRE environment has evolved assuming 1:1 interaction between researchers and data custodians rather than a more scalable approach where many researchers can interact with many data custodians.

A federated infrastructure would make it much easier for analysis tools to access multiple TREs. DARE UK (https://dareuk.org.uk/) is developing a blueprint for TRE federation [5] as is the European EOSC-ENTRUST project (https://eosc-entrust.eu/). The reality is that the ecosystem is highly autonomous and heterogeneous. Already well-established TREs use different software stacks and many popular analysis tools that need to access these multiple TREs are in widespread use. Currently each toolkit implements its own unique exchange standards, and each TRE develops a bespoke solution for the toolkits to access, all of which must be independently tested against

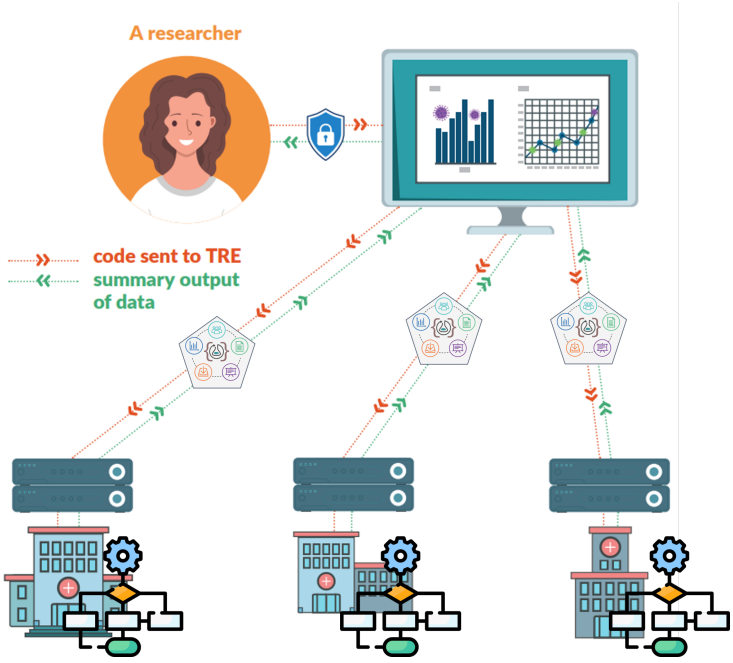

Figure 1. *Federated Analysis poses a query request to multiple TREs. Results are aggregated by the researcher. RO-Crates travel between analysis toolkits and TREs.*

the Five Safes framework. This is complicated, a blocker to interoperability, and scales poorly as further TREs are added to the requests.

Acknowledging the heterogeneity of established analysis platforms and TREs, the TRE-FX project (https://trefx.uk/) set out to streamline the flow of metadata needed for federated analysis using existing tooling and open standards. Scalable interoperability depends upon **standardised objects** flowing between services, as well as APIs and modularised component services that can be interchanged and mapped to multiple (existing) implementations. Interoperability also requires metadata harmonisation against community data standards such as the OMOP Common Data Model [6], but here we are interested in streamlining the flow of metadata that enables Five Safes processing rather than the harmonisation of results.

2. The Five Safes RO-Crate

To ensure that data access is safe, and the process is transparent, the **Five Safes RO-Crate** [7] packages the digital objects for research query requests and results with the metadata needed so that analysis clients and TREs can operate Five Safe compliant processes.

RO-Crate [8] is a community effort to establish a lightweight, web-native approach to packaging research data with their metadata. It has become a widely adopted framework for inter-service exchange, resource archiving, and reproducible reporting, used by many digital research infrastructures and their services. It is considered a "webby" implementation of the FDO Forum's FAIR Digital Objects, using Schema.org and Linked Data JSON-LD [9].

The Five Safes RO-Crate is a specialised profile of RO-Crate whereby encapsulated elements and metadata provide the necessary context for evaluating the safety and appropriateness of both data access and analysis (Table 2).

In practice, a Five Safes RO-Crate represents a unit of computational workflow-based access to sensitive information to enable trusted workflow execution in a TRE, from an authenticated workflow run request, through approval and review processes to a completed workflow execution. The request against the data in the TRE is thus expressed as a parameterisation of a pre-approved and pre-installed workflow. The idea of "code only" access to data has been proposed by the Goldacre Review [10] and pioneered by OpenSafely (https://www.opensafely.org/), the secure analytics platform for NHS electronic health records [11]. The Five Safes RO-Crate Profile builds upon the Workflow-Run-RO-Crate Profile [12], effectively making them a representation of trusted workflow provenance.

Table 2. *The Five Safes RO-Crate metadata*

Dimension	**RO-Crate metadata**
Safe Data	Requested input parameters: to the pre-approved workflow that will execute against the TRE pre-prepared, confidential-safe data.
Safe projects	Responsible project: The project that the request is sent on behalf of, related to permission to use a TRE.
Safe people	Requesting Agent: The individual person who is requesting the workflow run and the organisation they are representing for access control purposes.
Safe settings	Requested Workflow Run: Pre-approved workflow pre-installed on the TRE.
Safe outputs	Output entities: described as a Workflow Run Crate profile with workflow provenance and aggregated results for disclosure control inspection.

Each Five Safes crate collects and carries, end to end, the metadata necessary for the workflow analysis execution. The crates are minted by the analysis toolkits with the selected pre-approved workflow analysis, and travel in and out of a TRE, collecting the disclosure-controlled results, along with the details of the process, to then be aggregated by the user (Figure 2, left). The final crates may be separately registered in a Data Use Register for audit purposes and combined into one "meta-crate" to gather all the results.

Included within the Five Safes RO-Crate specification are eight steps that ensure that the RO-Crate's metadata outlined in Table 2 are appropriately recorded and reviewed (Figure 2, right):

1. **Check Phase**: The Crate is checked for integrity and completeness before passing to the TRE.
2. **Validation Phase**: The Crate metadata is checked for validity according to the profile.
3. **Workflow Retrieval Phase**: The TRE retrieves the requested workflow, typically from a local proxy of pre-approved workflows.
4. **Sign-off Phase**: The TRE verifies the requesting user/project is permitted to execute a given workflow on the TRE's data. This may include manual inspection e.g. selection criteria in input parameters.
5. **Workflow Execution Phase**: The TRE records execution of the workflow using a workflow engine, recording results and any errors.
6. **Disclosure Phase**: The Crate goes through disclosure control before leaving the TRE. This may be manual or semi-automatic depending on the workflow.
7. **Publishing Phase**: The Crate is made ready for publishing outside of the TRE and delivery to the user and data usage registers.
8. **Receiving Phase**: The client or Data Usage Register may check the returned Crate for completeness of the above phases and record this.

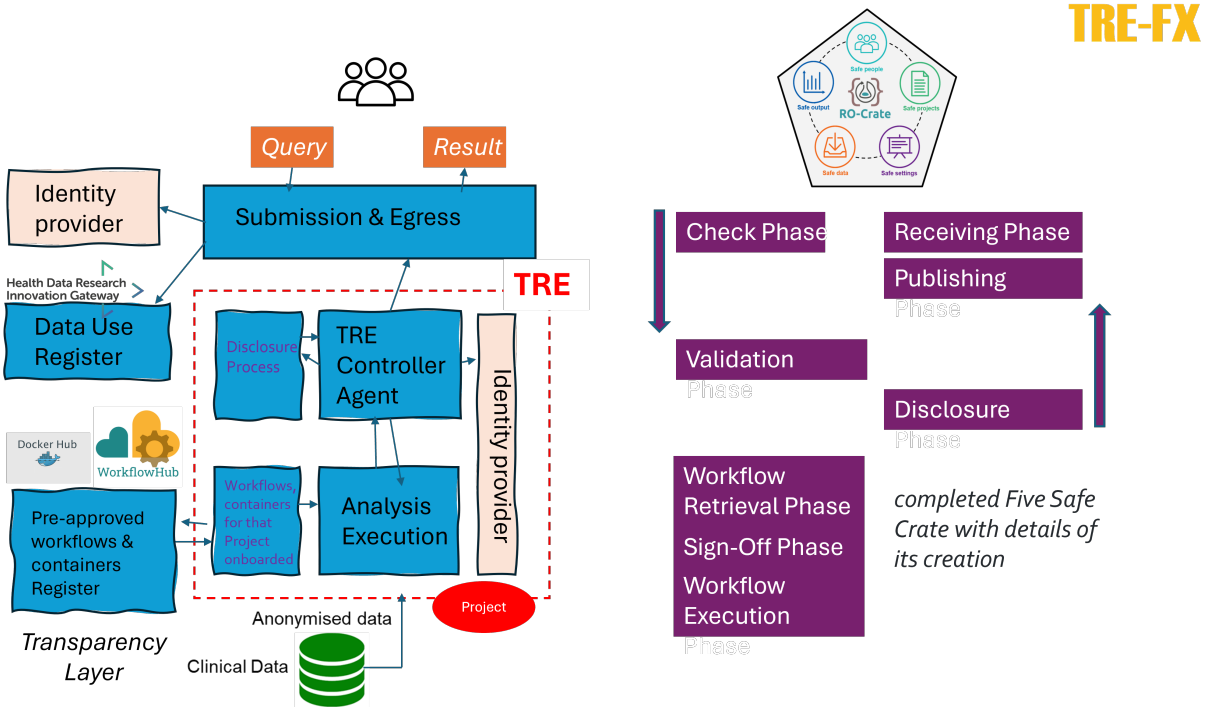

Figure 2. *The Five Safes RO-Crate travels in and out of the TRE (left) carrying the metadata and objects needed for Five Safes compliant analysis and reporting (right).*

2.1 FAIR in trusted research environments

It may seem counter-intuitive to follow the FAIR Guiding Practices [13] for sensitive data in a Trusted Research Environment, where the data cannot leave the premises. However, corresponding metadata is generally open, and authorization is permitted by FAIR principle *"A1.2 The protocol allows for an authentication and authorisation procedure, where necessary"*. For instance, UK's Clinical Practice Research Datalink [14] assigns DOIs to its CPRD datasets (https://www.cprd.com/digital-object-identifiers-dois-datasets), which landing pages [15] only provide metadata, and data access has to be justified and applied for.

Likewise, Five Safes RO-Crate is intended to create FAIR Digital Objects to document access and processing of such sensitive data, but not to contain them. Workflow results (e.g. summary statistics) are disclosure controlled and form part of the returned crate. The EOSC-ENTRUST and DARE UK initiatives are planning to work with TREs to establish citable dataset practices as shown by CPRD, adding a machine-readable RO-Crate containing only the metadata. The Five Safe Crate would then be able to link to this persistent identifier (PID) – and thus its RO-Crate – to globally identify the accessed sensitive data. However, currently we recognize that there is a steep learning curve for FAIR practices by TRE providers, as well as for data and software citation practices by TRE users. For this, mature FAIR technology and guidance that can be used without going into details of PIDs or Linked Data is essential.

An Five Safes RO-Crate is deployable in any FAIR repository like Zenodo or institutional research data repositories – for instance the *ro-crate-zenodo* tool [16] can upload an RO-Crate to Zenodo using its embedded metadata, translating it to the expected Datacite schema. However the crate can also be kept as a local record by individual researchers – the disclosed data is returned with its metadata – or by the TRE Provider, as evidence of data usage.

The HDR UK's Health Data Research Innovation Gateway [17] (https://healthdatagateway.org/) has a listing of health datasets, analysis scripts and records of data use. The Five Safe Crate can provide this information in a machine-readable way to the gateway, and work is planned for 2025 to attempt such integration. Further contextual metadata may be required (e.g. purpose of study), which is not easily inferrable by workflow engines within a TRE, but could be embedded in the crate by the TRE from the initial data usage request. Adding further FAIR aspects to the gateway (re-exposing the RO-Crate for instance) is then a poss8ible next step.

2.2 FDO approaches with Five Safes RO-Crate

Several aspects of FAIR Digital Objects [18] are tightly related to the Five Safes RO-Crate approach. Firstly, the Five Safes RO-Crate is by design machine-actionable (starting as a request to execute a workflow), and so serialises an *FDO Operation*. Secondly, the use of Workflow RO-Crates as FAIR Digital Objects [19] with Signposting is well-established and further utilised by Five Safe RO-Crate, as the workflow definition is listed by identifier from the submitted crate, and the referenced RO-Crate is retrieved from WorkflowHub in the *Workflow Retrieval Phase*.

The incremental passing and augmentation of Five Safes RO-Crate within the architecture (Figure 2) is not using any FDO-based APIs, but build on established GA4GH APIs [20] like Task Execution service (https://www.ga4gh.org/product/task-execution-service-tes/) where the RO-Crate is passed as a bitstream payload. This could be re-architected to use Signposting FDOs [9], indeed the planned future architecture will simplify Five Safe Crate handling by introducing a microservice that will operate on the RO-Crate as a held digital object, with the rest of the architecture passing its reference. This will however not use PIDs or be public, as the crate at this stage has not been through the Disclosure phase.

3. Five Safes RO-Crate pilots

The Five Safes RO-Crates Profile is intended as an approach for standardised metadata exchange that is readily adoptable by pre-existing platforms and installations. To demonstrate this, we piloted across different federated tooling and existing federated discovery tools.

Health Data Research UK has an existing cohort discovery tool, based on technology from BC Platforms (https://www.bcplatforms.com/) and the query/response is in a proprietary format. The GA4GH Beacon [21] is an internationally accepted format for constructing a query to discover data, which is open and developed by the community. Both standards are different, and each requires software to be installed in order to receive and respond to queries. Our test was to demonstrate how an RO-Crate could be used such that an install of a single tool could enable visibility in both networks.

Computational workflows were used to simplify and harmonise analysis execution in the different TREs. Workflows are registered in the WorkflowHub registry (https://workflowhub.eu/) for transparency, one for the HDR UK tool [22] and one for the Beacon [23] Utilising expertise in the NIHR Nottingham Biomedical Research Centre, we installed ELIXIR's Workflow Execution Service (WfExS) (https://github.com/inab/WfExS-backend) [24] with software developed by the University of Nottingham, HUTCH TRE-FX (https://health-informatics-uon.github.io/hutch-trefx/), that can receive an RO-Crate and submit to WfExS for processing. We were able to successfully demonstrate that a data partner could be visible in the two different data

discovery networks, by simply installing one software tool (HUTCH with WfExS) and it could process a Five Safes RO-Crate and return the results.

In addition, we piloted with two widely used privacy-preserving analysis toolkits – the open-source software DataSHIELD (https://www.datashield.org/) [25] and the commercial platform BitFount (https://www.bitfount.com/ – with two TREs (the SAIL Databank in Wales (https://saildatabank.com/) and at the University of Nottingham in England) using synthetic data. Both analysis toolkits independently adapted their platforms to dispatch and receive Five Safes RO-Crates. We were able to demonstrate how BitFount could submit their analyses to HUTCH via an RO-Crate instead of their normal processing mechanism.

4. Discussion and Next Steps

The adoption of Five Safe RO-Crates would be a significant step towards supporting federated analysis through standardised information flow as carriers of metadata between TREs and analysis toolkits. RO-Crates can help to scale up and scale out federated analysis and provide a platform for semi-automated governance processing for authorisation and disclosure control [26].

The Crates provide an auditable history for data generated within a TRE, supporting reproducibility, transparency and trust and harmonising record-keeping practices. For researchers and toolkit providers, the provenance information gives end-users confidence in the data they have generated where it is not easy to go back and check the original source data and processing steps. For TRE operators, the metadata supports an eight-step recording and reviewing process. The metadata audit trail provides evidence of disclosure review as well as a record of how the data is used within the federated TRE system while still accommodating existing TREs governance processes. Public Involvement & Engagement activities have repeatedly emphasised the importance of clear and transparent communication, accountability, and robust data governance to win over public and data custodian trust.

RO-Crates are metadata middleware, intended for smoothing the exchange, archiving, citing and reporting of digital entities. Our pilots have shown that the Five Safes RO-Crates Profiles approach is developer sympathetic; adoptable by existing systems and their software stacks enabling the reuse of existing services and tools; and is flexible enough to support different implementations.

This pilot work is now moving into the Health Data Research UK's Federated Analytics programme and will be incorporated into the European EOSC-ENTRUST project which aims to build a European Network of Trusted Research Environments. There is much to do including: pre-approval processes and configurations for the analysis; agreement on how Crates should be digitally signed, and by who; and developing methods to ensure that the contents of the Five Safe RO-Crate are the objects described and have not been modified. A common request from EOSC-ENTRUST and HDR UK we are now taking into consideration is to be able to make Five Safe RO-Crate also where there is no workflow engine involved, but the used software applications or other detailed execution log traces are known. Analysis toolkits and TREs need support to adapt their platforms to be RO-Crate compliant and tools for creating and viewing Crates. There are different patterns of federated analysis to be supported. Future work by DARE UK Phase 2 and CLIMATE-ADAPT4EOSC project will investigate using Open Digital Rights Language (ODRL) [4] as a machine-actionable policy expression language, which will be informed by the metadata in the Five Safes RO-Crate.

Federated analysis requires the challenging interaction of digital technology, data governance policy and data custodian practice. We believe that Five Safes RO-Crates will play an important future role in facilitating this interaction.

Data availability statement

The Five Safes RO-Crate Specification [7] is available at https://w3id.org/5s-crate/ as Open Access according to the MIT Licence.

Underlying and related material

Software implementing Five Safes RO-Crate for Federated Analytics in HDR UK is available in GitHub https://github.com/federated-analytics with Open Source licenses.

Author contributions

Stian Soiland-Reyes: Conceptualization, Investigation, Methodology, Software, Supervision, Validation, Writing – review & editing; Stuart Wheater: Investigation, Software, Validation; Thomas Giles: Conceptualization, Methodology, Project administration; Jonathan Couldridge: Investigation, Software, Validation; Philip Quinlan: Conceptualization, Funding acquisition, Project administration, Supervision, Writing – original draft; Carole Goble: Conceptualization, Funding acquisition, Project administration, Supervision, Writing – original and final draft.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was funded by UK Research & Innovation [Grant Number MC_PC_23007] as part of Phase 1 of the DARE UK (Data and Analytics Research Environments UK) programme, delivered in partnership with Health Data Research UK (HDR UK) and Administrative Data Research UK (ADR UK); and by the European Commission's Horizon Europe programme for grant HORIZON-INFRA-2021-EMERGENCY-01 101046203 "Beyond COVID" (https://doi.org/10.3030/101046203).

Acknowledgements

We would like to acknowledge the rest of the TRE-FX project team, including: Sam Cox, Daniel Lea, Vasiliki Panagi, Tim Beck (University of Nottingham), Blaise Thomson, Naaman Tammuz (Bitfount), Jillian Beggs, Christian Cole (University of Dundee), Suzy Gallier, Rima Doal (University Hospitals Birmingham NHS Foundation Trust), Justin Biddle, Simon Thompson (University of Swansea), Becca Wilson (University of Liverpool), Elizabeth Sapey (The University of Birmingham), and Emily Jefferson (HDR UK).

References

[1] UK Health Data Research Alliance and NHSX, "Building trusted research environments – principles and best practices; towards TRE ecosystems," HDR UK, Tech. Rep., 2021. DOI: 10.5281/zenodo.5767586.

[2] E. Green and F. Ritchie, "The present and future of the Five Safes framework," *Journal of Plantation Crops*, vol. 13, no. 2, Nov. 2023. DOI: 10.29012/jpc.831.

[3] P. Stokes. "The 'Five Safes' – data privacy at ONS," Accessed: Oct. 28, 2024. [Online]. Available: https://blog.ons.gov.uk/2017/01/27/the-five-safes-data-privacy-at-ons/.

[4] R. Iannella, S. Villata, and Permissions & Obligations Expression Working Group, "ODRL information model 2.2," W3C, Feb. 2018. [Online]. Available: https://www.w3.org/TR/2018/REC-odrl-model-20180215/.

[5] DARE UK. "Federated architecture blueprint," Accessed: Oct. 28, 2024. [Online]. Available: https://dareuk.org.uk/our-work/federated-architecture-blueprint/.

[6] OHDSI. "Standardized data: The OMOP Common Data Model." OHDSI – Observational Health Data Sciences and Informatics, Accessed: Oct. 28, 2024. [Online]. Available: https://www.ohdsi.org/data-standardization/.

[7] S. Soiland-Reyes, S. Wheater, T. Giles, C. Goble, and P. Quinlan, "TRE-FX technical documentation - Five Safes RO-Crate," DARE UK, Tech. Rep., 2023. DOI: 10.5281/zenodo.10376350.

[8] S. Soiland-Reyes et al., "Packaging research artefacts with RO-Crate," *Data Science*, vol. 5, no. 2, pp. 97–138, Jul. 2022. DOI: 10.3233/ds-210053.

[9] S. Soiland-Reyes, P. Sefton, S. Leo, L. J. Castro, C. Weiland, and H. Van de Sompel, "Practical webby FDOs with RO-crate and FAIR signposting: Experiences and lessons learned," in *International FAIR Digital Objects Implementation Summit 2024*, vol. (this issue), TIB Open Publishing, 2024.

[10] B. Goldacre and J. Morley, "Better, broader, safer: Using health data for research and analysis," Department of Health and Social Care, A review commissioned by the Secretary of State for Health and Social Care, Apr. 2022. Accessed: Oct. 28, 2024. [Online]. Available: https://www.gov.uk/government/publications/better-broader-safer-using-health-data-for-research-and-analysis.

[11] L. Nab et al., "OpenSAFELY: A platform for analysing electronic health records designed for reproducible research.," *Pharmacoepidemiology and Drug Safety*, vol. 33, no. 6, e5815, Jun. 2024. DOI: 10.1002/pds.5815.

[12] S. Leo et al., "Recording provenance of workflow runs with RO-Crate," *PLOS One*, vol. 19, no. 9, e0309210, Sep. 2024. DOI: 10.1371/journal.pone.0309210.

[13] M. D. Wilkinson et al., "The FAIR Guiding Principles for scientific data management and stewardship," *Scientific Data*, vol. 3, no. 1, Mar. 2016. DOI: 10.1038/sdata.2016.18.

[14] E. Herrett et al., "Data Resource Profile: Clinical Practice Research Datalink (CPRD)," *International Journal of Epidemiology*, vol. 44, no. 3, pp. 827–836, Jun. 2015. DOI: 10.1093/ije/dyv098.

[15] Clinical Practice Research Datalink, *CPRD Aurum september 2024*, 2024. DOI: 10.48329/ZRVZ-6A47.

[16] E. Chadwick and S. Soiland-Reyes, *Rocrate-zenodo*, version v0.1.1, Aug. 2024. DOI: 10.5281/zenodo.13365999. [Online]. Available: https://github.com/ResearchObject/ro-crate-zenodo/.

[17] G. Reilly and S. Varma, "Health Data Research Innovation Gateway," *ITNOW*, vol. 63, no. 2, pp. 60–63, Jun. 2021. DOI: 10.1093/itnow/bwab061.

[18] I. Anders et al., "FAIR digital object technical overview," Version PEN 2.0, Proposed Recommendation Full FDO Overview PEN-2.0-v2, Jan. 18, 2023. DOI: 10.5281/zenodo.7824714.

[19] S. Soiland-Reyes, E. Chadwick, F. Bacall, J. M. Fernández, B. Grüning, and H. Bayındır, "EuroScienceGateway D2.1: Reproducible FAIR Digital Objects for Workflows," Zenodo, Deliverable, Aug. 2024. DOI: 10.5281/zenodo.13225792.

[20] H. L. Rehm et al., "GA4GH: International policies and standards for data sharing across genomic research and healthcare," *Cell Genomics*, vol. 1, no. 2, p. 100 029, Nov. 2021. DOI: 10.1016/j.xgen.2021.100029.

[21] M. Fiume et al., "Federated discovery and sharing of genomic data using beacons.," *Nature Biotechnology*, vol. 37, no. 3, pp. 220–224, Mar. 2019. DOI: 10.1038/s41587-019-0046-x.

[22] V. Panagi, "Rquest-omop-worker," *WorkflowHub*, 2023, Common Workflow Language. DOI: 10.48546/workflowhub.workflow.471.2.

[23] V. Panagi, *Beacon-workflow*, Common Workflow Language, WorkflowHub, 2024. Accessed: Oct. 28, 2024. [Online]. Available: https://workflowhub.eu/workflows/882?version=3.

[24] J. M. Fernández, L. Rodríguez-Navas, and S. Capella-Gutiérrez, "Secured and annotated execution of workflows with WfExS-backend," *F1000Research*, 2022. DOI: 10.7490/f1000research.1119198.1.

[25] I. Budin-Ljøsne et al., "DataSHIELD: An ethically robust solution to multiple-site individual-level data analysis.," *Public health genomics*, vol. 18, no. 2, pp. 87–96, 2015. DOI: 10.1159/000368959.

[26] J. Smith et al., "SACRO: Semi-automated checking of research outputs," DARE UK, Tech. Rep., 2023. DOI: 10.5281/zenodo.10055365.