

# Semantic-enabled Architecture for Auditable Privacy-Preserving Data Analysis

Fajar J. Ekaputra<sup>a,\*</sup>, Andreas Ekelhart<sup>b,c</sup>, Rudolf Mayer<sup>c,a</sup>, Tomasz Miksa<sup>c,a</sup>, Tanja Šarčević<sup>c,a</sup>, Sotirios Tsepelakis<sup>c</sup>, and Laura Waltersdorfer<sup>a</sup>

<sup>a</sup> *Information and Software Engineering Research Group, TU Wien, Vienna, Austria*  
*E-mails: fajar.ekaputra@tuwien.ac.at, laura.waltersdorfer@tuwien.ac.at*

<sup>b</sup> *Department of Information Systems and Operations Management, WU, Vienna, Austria*  
*E-mail: andreas.ekelhart@wu.ac.at*

<sup>c</sup> *SBA-Research, Vienna, Austria*  
*E-mails: rmayer@sba-research.org, tmiksa@sba-research.org, tsarcevic@sba-research.org, stsepelakis@sba-research.org*

**Abstract.** Small and medium-sized organisations face challenges in acquiring, storing and analysing personal data, particularly sensitive data (e.g., data of medical nature), due to data protection regulations, such as the GDPR in the EU, which stipulates high standards in data protection. Consequently, these organisations often refrain from collecting data centrally, which means losing the potential of data analytics and learning from aggregated user data.

To enable organisations to leverage the full-potential of the collected personal data, two main technical challenges need to be addressed: (i) organisations must preserve the privacy of individual users and honour their consent, while (ii) being able to provide data and algorithmic governance, e.g., in the form of audit trails, to increase trust in the result and support reproducibility of the data analysis tasks performed on the collected data.

Such an *auditable, privacy-preserving data analysis* is currently challenging to achieve, as existing methods and tools only offer partial solutions to this problem, e.g., data representation of audit trails and user consent, automatic checking of usage policies or data anonymisation. To the best of our knowledge, there exists no approach providing an integrated architecture for auditable, privacy-preserving data analysis.

To address these gaps, as the main contribution of this paper, we propose the WellFort approach, a semantic-enabled architecture for auditable, privacy-preserving data analysis which provides secure storage for users' sensitive data with explicit consent, and delivers a trusted, auditable analysis environment for executing data analytic processes in a privacy-preserving manner. Additional contributions include the adaptation of Semantic Web technologies as an integral part of the WellFort architecture, and the demonstration of the approach through a feasibility study with a prototype supporting use cases from the medical domain. Our evaluation shows that WellFort enables privacy preserving analysis of data, and collects sufficient information in an automated way to support its auditability at the same time.

**Keywords:** provenance, semantic web, privacy-preserving data analysis, auditability, dpv, consent management

## 1. Introduction

In Europe, the GDPR [1], in effect since May 2018, stipulates high standards in data protection and imposes substantial fines for non-compliance. Moreover, security breaches have the potential to go beyond fi-

nancial impact and ruin an organisation permanently, as lost reputation and trust cannot be regained easily. Thus, it is not enough for organisations to make sure that the data analysis techniques fulfil guarantees of privacy of the individuals representing the underlying records [2], but they also need to ensure that the use of data is compliant with the user consent. This is es-

---

\*Corresponding author. E-mail: fajar.ekaputra@tuwien.ac.at.

pecially important to keep track of explicit consent for sensitive user data as defined in GDPR art. 9.

Furthermore, to enable auditability, accountability, and transparency of systems used for data analysis, we must capture and understand the broad context in which the system operates: data sources and algorithms, as well as processes and services that the deployed systems depend on [3].

Therefore, without a strong background in IT security and big budgets, providing a secure platform for data storage and analysis is often beyond capabilities of small and medium-sized organisations.

As a consequence, organisations often refrain from collecting data centrally and, for example, offer applications that analyse data locally instead. While this reduces the attack surface, it means losing on the potential of data analytics and learning from the entire collected data, and thereby hinders innovative services and research. As an example, identifying trends over cohorts of users is not possible.

To this end, the main research problem discussed in this paper can be formulated as follows: *How to enable auditable, privacy-preserving data analysis systems?*

Further, we divided the problem into a number of research questions as follows:

- *What are the key characteristics of auditable, privacy-preserving data analysis systems?*
- *What are the key elements of an architecture for auditable, privacy-preserving data analysis systems?*
- *How can Semantic Web technologies be used to enable auditable, privacy-preserving data analysis systems?*

Recently, research on applying Semantic Web technologies to address challenges related to auditability and user privacy has been intensified. On the one hand, the emergence of the W3C PROV-O recommendation [4] has been well-received in the community as a de-facto standard for RDF representation of provenance trails, forming a basis for auditability. As a result, the development of methods and tools around PROV-O are growing rapidly, summarised in a recent survey paper [5]. On the other hand, the research on representation and compliance checking for user consent and personal data handling also gained traction. Semantic Web technologies, with ontologies and reasoners as key resources, have been proposed to facilitate representing [6–8] and automated compliance checking [9–11] on usage policy, which are key to enable consent-

check mechanisms for privacy-preserving data analysis.

In a wider research community, privacy-preserving data analysis is an active research area. Various approaches that have been proposed for privacy-preserving data analysis can be categorised according to the data lifecycle they are applied to, for example: *privacy-preserving data publishing* (PPDP) and *privacy-preserving data mining outputs* (PPDMO) [12]. In our work we rely on the methods from PPDP and PPDMO. PPDP relies on anonymising data records before publishing. The most prominent method is *k*-anonymity [13], including its extensions, *l*-diversity and *t*-closeness. Approaches towards PPDMO include *query inference control*, where either the original data or the output of the query are perturbed and *query auditing* where some queries are denied because of potential privacy breach. DataSHIELD [14] is an example of such a system, which applies both of the later techniques to a number of data mining functions.

Despite research on isolated topics, to the best of our knowledge, no integrated architecture for auditable privacy-preserving data analysis exists. The lack of such an approach can be attributed in part to the contradicting goals of collecting as much data as possible for auditability, while aiming for reduced data collection to minimise the impact of a data breach.

To fill in this gap, we propose a novel approach called WellFort, a semantic-enabled architecture for auditable, privacy-preserving data analysis which (i) provides secure storage for users' sensitive data with explicit consent, and (ii) delivers a trusted analysis environment for executing data analytic processes in a privacy-preserving manner. A novelty of our approach is that organisations do not have direct access to individual data, but only allow in aggregated or anonymised form. Organisations on the one hand can benefit from a large groups of individuals that are potentially willing to share their data for broader analytics tasks, such as those required by e.g., medical research. Users, on the other hand, benefit from a privacy-preserving and secure platform for their data, and can contribute to analytics/research projects in a secure manner. Finally, analysts (such as researchers) may obtain a detailed source of microdata, if data subjects give the corresponding consent.

The contribution of this paper is three-fold:

- We propose WellFort, a semantic-enabled architecture for auditable privacy preserving data analysis,

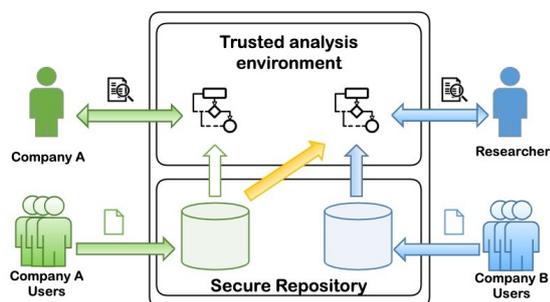


Fig. 1. WellFort overview scenario

- We adapt Semantic Web technologies as an integral part of the architecture, including: (a) ontology-based data representation of provenance, dataset metadata, and usage policies, (b) ontology development and extension methods, as well as (c) automatic usage policy compliance checking with OWL2 reasoning, and
- We demonstrate the feasibility of our approach on real-world use cases in the medical domain.

The rest of the paper is structured as follows. Section 2 presents the requirements for the architecture. Section 3 describes the WellFort conceptual architecture and the processes it supports. Section 4 reports on the Semantic Web methods adoption to support the architecture. Section 5 presents the prototype implementation. Section 6 investigates the feasibility of a WellFort-based prototype in four use cases, and Section 7 presents related work. Finally, conclusions and an outlook are given in Section 8.

## 2. Requirements

In this section, we introduce the main functional and non-functional requirements for our platform. These iteratively refined requirements are based on privacy standards as well as use cases and discussions with partners from the WellFort project<sup>1</sup>.

In the following, we describe a scenario (cf. Figure 1) to exemplify and motivate the requirements for auditable privacy-preserving data analysis described in the following sub-sections. Later, in Section 6, we will derive four concrete use cases from this scenario to conduct feasibility studies of our approach.

*Andrea is a user of two health apps, one from Company H, which monitors her heart rate, and another*

*one from Company M, which records her cholesterol level. She wants to collect her data from both apps, link and share her data, but only for research purposes. She also considers sharing her data for specific commercial purposes, like getting product recommendation, but wants to decide about it later.*

*However, since Company H and Company M are small companies, they do not want to manage the data themselves. They plan to store their users' data in a cloud platform without having to develop and run the platform themselves.*

*Barbara, a researcher from TU Wien, is planning to develop a machine-learning model on the relation between heart rate and cholesterol level for her research. She is looking for suitable data for her model, but at the same time, she wants to make sure that she is compliant with regulations.*

*Later on, Andrea asks Julia, an auditor of the cloud platform, regarding the use of her data. Since Andrea's data is used as part of Barbara's research, Julia provides Andrea with the information on the research setup.*

In the scenario above, the data captured by the application providers (i.e., *Company A and B*) can be categorised in the special category of personal data in GDPR art. 9. Thus, the cloud platform should facilitate users like *Andrea* to define explicit consent to their data. Furthermore, due to privacy reasons, application providers and researchers like *Barbara* should not get direct access to user data repositories, but still should be able to conduct data analysis in a privacy-preserving manner without compromising user consent. *Julia*, in the other hand, should be able to provide the audit information to *Andrea*, without having access to the content of her personal data.

### 2.1. Privacy-preserving Data Analysis

*Data analysis.* The platform must enable data analysts to conduct privacy-preserving analysis without disclosing individual user data. To reduce the attack surface, input and intermediate data beyond the analysis must not be persisted by the platform. It should be possible to analyse data originating from the same application (organization) and also to conduct cross-application (cross-data owner) analysis (e.g., correlating heart rate with cholesterol level, even though each was collected from a different application).

Finally, all data analysis computation on raw individual data must take place inside the platform and the analysis environments can be stored only for a limited

<sup>1</sup><https://www.sba-research.org/research/projects/wellfort/>

1 period of time, as defined by the data retention policy  
2 of the platform.

3 *Data publishing.* Analysts must not download any  
4 data in its original form, to fulfil regulatory require-  
5 ments, and to minimise the surface for privacy attacks,  
6 such as user re-identification. Instead, for offline anal-  
7 ysis, the platform should offer anonymized downloads,  
8 and the generation of synthetic data [15]. The platform  
9 must provide measures to prohibit downloading similar  
10 data multiple times by a single analyst, since this  
11 provides the channel for disclosure of sensitive infor-  
12 mation.

13 *Usage Policy Management.* Data analysis requires  
14 fine-grained consent from the involved users [1], and  
15 access control has to ensure that only authorised anal-  
16 ysts can work with the data. To this end, usage  
17 policies, which include both user consent and data  
18 handling setup need to be represented in a machine-  
19 readable format. Further, a mechanism to conduct an  
20 automatic usage policy checking between consent and  
21 data handling is required.

## 22 2.2. Auditability

23 *Provenance data capture and tracing.* Provenance  
24 of data must be traced in the whole data lifecycle  
25 within the platform. Automated collection of by who,  
26 when, and how data has been accessed, processed and  
27 analysed needs to be recorded, ranging from high level  
28 to precise trails. In addition, provenance must include  
29 information on the software components and their con-  
30 figuration used to modify and process data. The plat-  
31 form must also provide traces on how each result was  
32 computed. Finally, analysts must not be able to acquire  
33 or infer personal or otherwise sensitive data via the  
34 provenance trail.

35 *Provenance data inspection and analysis capability.*  
36 Collected provenance data streams must be accessible  
37 to users with auditing privileges through an interface,  
38 to be able to answer audit questions concerning data  
39 and its upload, usage, and analysis. Such a question  
40 could be to inspect which studies involved a user's per-  
41 sonal data, for what purposes, and how it was anal-  
42 ysed. Auditors should be able to access the data while  
43 preserving the privacy of users, therefore, no personal  
44 data must be acquired by the provenance feature.

45 *Metadata-retention after deletion.* Even after users  
46 revoke their consent due to their right to be forgotten,  
47 consent for past studies is still valid. In order to ful-  
48 fil the requirement of providing proof of correct data  
49 capturing and data analysis, minimum necessary meta-  
50 data must be kept in the provenance module  
51 for past studies that were conducted and based on now  
52 deleted information.

1 information must be kept in the provenance module  
2 for past studies that were conducted and based on now  
3 deleted information.  
4  
5

## 6 3. Conceptual architecture

7  
8 In this section, we describe the conceptual architec-  
9 ture of the platform for auditable privacy-preserving  
10 data analysis, which addresses the requirements de-  
11 fined in Section 2. More on the technical details will  
12 be described on Section 4 and Section 5.

13 The conceptual architecture of the platform is de-  
14 picted in Figure 2. There are three distinct actors:  
15

- 16 – *User* – store their data in the platform and give  
17 consent to analyse it. They use an application pro-  
18 vided by the organisation and do not interact with  
19 the platform directly, e.g., using a dedicated user  
20 interface.
- 21 – *Analyst* – can run experiments in the platform.  
22 They define what types of data will be used in the  
23 analysis and perform the actual analysis.
- 24 – *Auditor* – can analyse evidence collected to an-  
25 swer specific audit questions that depend on the  
26 purpose of the audit, e.g., a litigation case, or  
27 simple usage statistics for users wanting to know  
28 when and by whom their data was used.

29  
30 The architecture consists of three component groups,  
31 each serving a different purpose:

- 32 – *Secure Repository* – stores data uploaded by a  
33 user and allows the selection of data to be used  
34 in experiments by the analyst. There are no com-  
35 ponents in this group allowing to view or analyse  
36 data.
- 37 – *Trusted Analysis Environment* – selected data that  
38 fulfils experiment criteria, e.g., consent, fit for  
39 purpose, etc. is duplicated to this component for  
40 further analysis. This part of the platform pro-  
41 vides mechanisms allowing the analysts to con-  
42 duct their experiments in a privacy-preserving  
43 manner. Data selection is usually expressed via  
44 queries.
- 45 – *Audit Box* – collects and manages provenance  
46 data to support the auditability of processes  
47 within the *secure repository* and *trusted analy-  
48 sis environment*. The Knowledge Base can be ac-  
49 cessed via an interface to answer audit-related  
50 questions on personal data access and usage.  
51

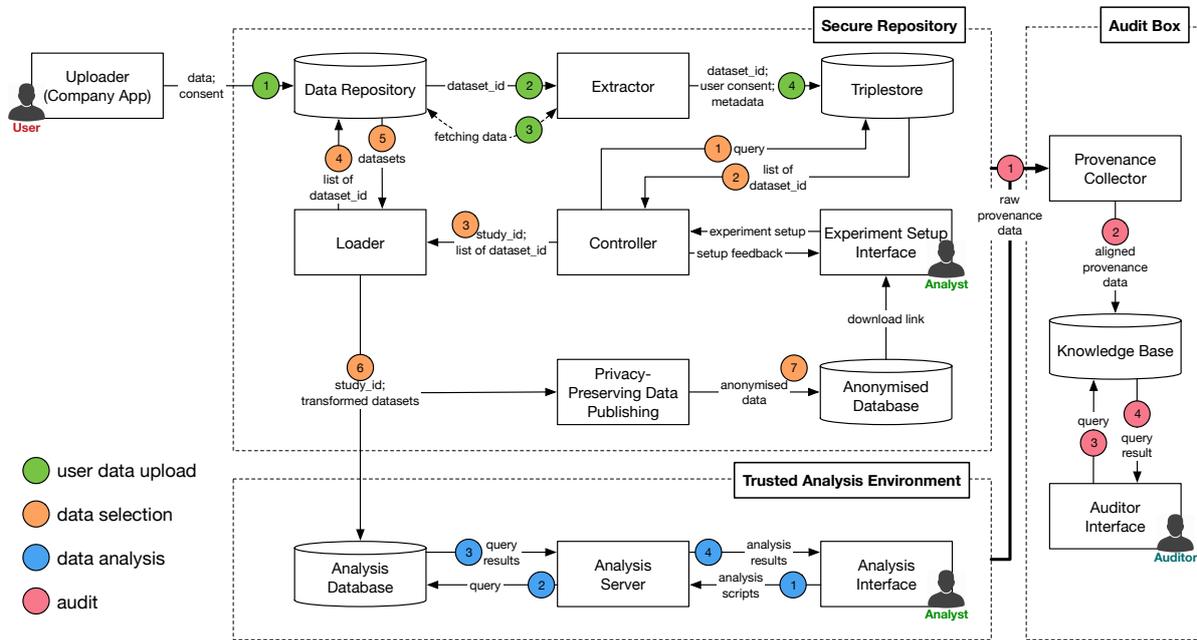


Fig. 2. WellFort conceptual architecture

In Figure 2, we use dashed lines to indicate the division into these three groups. Figure 2 further depicts four typical processes executed in the platform:

- *User data upload* – the process starts when a user’s application sends data to the platform. It extracts metadata from the data, and stores it together with the data and the consent indicated during the upload in the platform. Thus, every dataset uploaded to the platform is linked to a minimal set of information that allows for its retrieval.
- *Data selection* – analysts define the search criteria for data they want to use in their experiments. If the platform has enough data fulfilling their criteria (and also consent for usage), then the process loads the actual data into the *Trusted Analysis Environment* where the analyst obtains accurate aggregate results without having access to individual datasets. The search for data relies only on high-level information provided in the metadata. For example, *BloodGlucose*, which is one of the attributes in the actual user data, has only a boolean value in the corresponding metadata to indicate whether the dataset contains information on *BloodGlucose*. As an alternative, Analysts can download anonymised data and analyse it outside

(offline) of the platform. However, anonymisation can substantially reduce the quality of the data and the results of the offline data analysis thereupon.

- *Data analysis* – analysts process the data and produce results by submitting code to the platform. The platform will ensure that the analysts will not be able to identify or infer data subjects from the analysis.
- *Audit* – provenance information obtained from each component of the platform is sent to an independent component and organised in a knowledge base. Auditors query experiment details.

We numbered the steps of each process and marked them with colours: green (user data upload), orange (data selection), blue (data analysis), and pink (audit). We describe now each of the processes in detail to explain the role of each of the components.

### 3.1. User data upload

Following the terminology set by the Open Archival Information System (OAIS) reference model [16], this process is equivalent to the ingest process that transforms and enriches data uploaded by users, so that it can be later located and accessed according to the corresponding user consent.

1 The process starts by a user's application sending  
2 data together with the consent to the *Data Repository*.  
3 The *Data Repository* creates a record in which  
4 it stores the persistent identifier (PID), consent, and  
5 location of the received data. The PID is unique and  
6 never changes. Following the FAIR (Findable, Acces-  
7 sible, Interoperable, Reusable) principles [17], even if  
8 the data is deleted, the PID is kept, and the record is  
9 updated with information on why and when the data  
10 was removed.

11 The platform automatically creates metadata and  
12 stores it together with the consent in the *Triplestore*.  
13 The *Extractor* component receives the PID of the  
14 dataset and accesses the data. The generated metadata  
15 consists of a list of key-value pairs, where each key  
16 corresponds to an attribute located in data, e.g., heart  
17 rate, while the value is always boolean, i.e. true or  
18 false. Thus, the *Triplestore* contains PIDs of datasets,  
19 together with a high-level description of the presence  
20 of certain types of contents, and consent given. The  
21 *Triplestore* does not contain any user specific informa-  
22 tion. We provide more details on the metadata and con-  
23 sent representation in Section 4.1.1 and Section 4.1.2  
24 respectively.

25 It would be possible to use only a *Data Repository*  
26 and omit the *Triplestore*. However, the current de-  
27 sign has two major advantages. First, data is decou-  
28 pled from its metadata, and thus analysts can query  
29 the *Triplestore*, while the actual data can be kept of-  
30 fline, e.g., due to security measures. Second, by using  
31 the semantics provided by the underlying ontology, we  
32 can run advanced queries to identify suitable data for  
33 an experiment that go beyond simple string matching,  
34 e.g., by using subsumption reasoning for query expan-  
35 sions and usage policy compliance checking.

### 37 3.2. Data selection

40 Analysts use the *Experiment Setup Interface* to de-  
41 fine their experiments. They specify the purpose of the  
42 experiment, and which data attributes they want to use  
43 in their analysis. Based on this input, the *Controller* au-  
44 tomatically generates a query and sends it to the *Triple-*  
45 *store*. The *Triplestore* returns a list of dataset PIDs that  
46 fulfil the criteria (cf. Section 4.1 for details on how the  
47 result is produced). If the number of available datasets  
48 fulfils the requirements of the analyst (so that the study  
49 is meaningful from a statistical point of view), they can  
50 decide to start a new study and load the data into the  
51 *Trusted Analysis Environment*.

1 Each study has its own PID and links with the list  
2 of dataset PIDs, as well as metadata about the experi-  
3 ment. Similar to the PIDs used for datasets, the study  
4 PID allows us to refer to experiments in an unambigu-  
5 ous way.

6 The *Loader* receives the study PID and fetches the  
7 data from the *Data Repository*. Depending on the ac-  
8 tual implementation of the *Trusted Analysis Environ-*  
9 *ment*, the *Loader* transforms the data into a suitable  
10 format, e.g., database, pivot tables, etc. When the data  
11 is loaded, analysts receive information via the *Exper-*  
12 *iment Setup Interface* on how to access the *Trusted*  
13 *Analysis Environment*, e.g., an access URL, username  
14 and password.

15 Analysts can also decide to download either the k-  
16 anonymised data or synthetic data. To do so they use  
17 the *Experiment Setup Interface* to define their experi-  
18 ment. All the steps are the same as before, yet the data  
19 is not loaded in the the *Analysis Server* but transformed  
20 by the *Privacy Preserving Data Publishing* component  
21 and made available for download.

### 22 3.3. Data analysis

25 Analysts use the *Analysis Interface* to communicate  
26 with the *Analysis Server*, which in turn has access to  
27 the data stored in the *Analysis Database*. The *Analy-*  
28 *sis Interface* allows submitting code for data analysis  
29 and viewing the processing results. It also logs every  
30 activity of the analyst. The *Analysis Interface* can be  
31 realised as a web application running on a server side.

32 The *Analysis Server* inspects the code submitted by  
33 analysts and protects data from direct access. For ex-  
34 ample, the server allows only aggregation functions to  
35 be executed, and blocks requests to list properties of  
36 specific datasets (users). If the submitted scripts pass  
37 such requirements for privacy-preservation, then the  
38 results are calculated on the original data and returned  
39 to the *Analysis Interface*.

40 Please note that the *Analysis Interface* is part of the  
41 platform, and a direct access to the *Analysis Server* is  
42 not possible. Such a design is more secure, and allows  
43 to better track the analysis process.

44 A data analysis process can only be executed when  
45 the data has been loaded into the *Trusted Analysis En-*  
46 *vironment*, and when the analysts received their con-  
47 nection details. The *Analysis Database* receives only  
48 the data required for the specific study, keeps it for a  
49 limited time, and deletes it after the study is completed.  
50 During the analysis, there is no need to access the *Se-*  
51 *cure Repository*. Hence, thanks to loose coupling, each

group of components can be instantiated by deploying them on different nodes with varying levels of security, access, as well as computational power.

### 3.4. Audit

The *Secure Repository* and the *Trusted Analysis Environment* push provenance information to the *Provenance Collector* to separate this information from the components processing data, in which the provenance was captured. Such a design increases security of provenance information, because it prevents from non-authorised modification of provenance when one of the user-facing components is compromised.

Provenance capturing can be implemented in different ways and the captured information can be structured differently, hence, it is necessary to align representation of provenance information before loading it into the *Knowledge Base*.

The *Knowledge Base* stores provenance for each of the processes supported by the platform, specifically:

- *User Data Upload* – a persistent identifier of a dataset, consent, an automatically generated metadata summarising contents of the dataset, a timestamp.
- *Data Selection* – an identifier of the analyst, an identifier of the submitted query, query attributes (e.g. timestamp, study purpose, study description), query results, consent check results, and an identifier of the analysis (only if the consent check was successful and the analyst decides to analyse the data).
- *Data Analysis* – an identifier of the analysis (created in the previous step to link the query with the analysis), information on software environment and dependencies (e.g. operating system, software libraries, environment variables), script parameters, results of the analysis.

Auditors use the *Auditor Interface* to get access to provenance information organised in the *Knowledge Base*. They define queries to retrieve relevant provenance information when performing an audit, which we will describe in the next section.

## 4. Semantic-Web Methods for Auditable Privacy-preserving Data Analysis

In this section, we describe in detail how Semantic Web technologies contribute to the overall architec-

prefix	namespace
:	http://w3id.org/wellfort/ns/dpv#
meta:	http://w3id.org/wellfort/ns/meta#
wprv:	http://w3id.org/wellfort/ns/prov#
id:	http://w3id.org/wellfort/id/
dpv:	http://www.w3.org/ns/dpv#
prov:	http://www.w3.org/ns/prov#
dcat:	http://www.w3.org/ns/dcat#
dct:	http://purl.org/dc/terms/
xsd:	http://www.w3.org/2001/XMLSchema#
rdt:	https://github.com/End-to-end-provenance/ExtendedProvJson/blob/master/JSON-format.md

Table 1

The list of predefined prefixes and namespaces

ture presented in Section 3, in particular for consent and metadata management (Section 4.1) and auditability (Section 4.2). Note that from here on, for consistency reason we will use the prefixes and namespaces introduced in Table 1.

### 4.1. Consent and metadata management

In this section, we describe the mechanism to manage and utilise consent and metadata with Semantic Web technologies, to ensure that experiments can be executed with a sufficient amount of data without compromising the user consent of the datasets. The mechanism mainly resides within the *Secure Repository*, where the analysts interact with the system to setup their experiments.

When an analyst interacts with the *Experiment Setup Interface*, she needs to describe the *Data Handling* setup, which contains information about the required category of personal data, purpose, processing, recipient, and expiry time of the experiment. This information allows the data selection process to be conducted as the following:

*Step 1. Finding datasets with matching categories*, where we query the *Triplestore* for datasets. As a part of the step, the RDFS reasoning is used to expand the dataset selection based on the subsumption relations between data categories.

*Step 2. Usage policy compliance check*. In this step, we determine whether the Data Handling setup of the analyst is compliant with the user consent of datasets collected from Step 1. Therefore, it is necessary to automate the compliance check in order to speed up the process.

From the steps above, we identify three necessary components that need to be defined in order to facilitate the consent and metadata management, namely (i) Dataset metadata representation, (ii) Usage policy representation, and (iii) Automated compliance checking mechanism.

#### 4.1.1. Metadata Representation

There are several suitable ontologies that can be used to represent dataset metadata, and we therefore only need to adapt and extend those where required, e.g., Data Catalog Vocabulary (DCAT) [18] or VOID [19]. For our approach, we decided to base our metadata<sup>2</sup> on Data Catalog Vocabulary (DCAT) Version 2 [18] for our dataset metadata representation since it is the latest W3C recommendation.

We adopt two classes from DCAT, `dcat:Catalog` to describe a catalog of datasets stored within our *Data Repository*, and `dcat:Dataset` as the superclass of our `meta:Dataset` to represent each dataset. We argue that this subclass is necessary since we add additional properties specific to `meta:Dataset` as follows:

- `meta:hasConsent` that relates a dataset and active `dpv:Consent` given by a data subject for this dataset, and
- `meta:hasDataCategory`, where this relation indicates the categories of personal data contained within a dataset. Note that the value of this property may be an instance of several subclasses of `dpv:PersonalDataCategory`.

These two relations imply the use of the Data Privacy Vocabulary (DPV) [6] in our architecture that we will explain in the next subsection.

```
id:WellFortCatalog a dcat:Catalog ;
  dct:title "WellFort Catalog" ;
  dct:publisher id:TUWien ;
  dcat:dataset id:dataset-1, id:dataset-2 .
id:dataset-303 a meta:Dataset ;
  :hasConsent id:consent-1 ;
  :hasDataCategory id:category-1 ;
  dct:identifier "303" ;
  dct:issued "2011-12-05"^^xsd:date .
id:category-1 a :HeartRate, dpv:Age .
```

Listing 1: An excerpt of dataset metadata on WellFort

We provided an example metadata in Listing 1, which shows the single catalog in the platform and dataset-303, which contains personal information on

<sup>2</sup><http://w3id.org/wellfort/ns/metadata>

```
ObjectIntersectionOf (
  ObjectSomeValuesFrom (dpv:hasPersonalDataCategory
    SomePersonalDataCategory)
  ObjectSomeValuesFrom (dpv:hasProcessing SomeProcessing)
  ObjectSomeValuesFrom (dpv:hasPurpose SomePurpose)
  ObjectSomeValuesFrom (dpv:hasRecipient SomeRecipient)
  DataSomeValuesFrom (:expiryTime
    DatatypeRestriction (xsd:dateTime xsd:maxInclusive
      SomeDateTime)))
```

Listing 2: DPV adaptation of SPECIAL Policy Language for WellFort

heart rate and age. We deliberately removed the relation between users and datasets in our metadata.

#### 4.1.2. Usage Policy Representation

One of the main requirement of our approach is to ensure that the personal data can be used for analysis without compromising user consent. To this end, there is a clear need to represent usage policies, both for user consent and analyst data handling, in a clear and concise manner. In recent years, a number of ontologies has been proposed to represent usage policies, e.g., DPV [6], PrOnto [8], SPECIAL [20] and SAVE [7], among others.

For our approach, we decided to use DPV as our basis for usage policy representation, due to (i) the community-based development of the vocabularies, and (ii) DPV does not recommend specific mechanism to use its concepts, providing adopters with a high degree of freedom to use and adapt it for their purpose. Furthermore, the DPV Community Group<sup>3</sup> involves researchers and practitioners with similar interests across the world and has become a strong candidate for usage policy representation.

We adapt the SPECIAL Policy Language [20] to structure DPV as usage policy. In this structure, a usage policy is composed of one or more *basic usage policies* each of which is an OWL2 expression of the form shown in Listing 2. We implement a small adjustment to the original structure by replacing the storage information with an explicit expiry time of policies, which is simpler but provide the necessary time limit on the use of data. We argue that the current structure is inline with the GDPR, which defines that policies shall specify clearly (i) which data are collected, (ii) what is the purpose of the collection, (iii) what processing will be performed, (iv) whether or not the data will be shared with others, and (v) for how long the data will be stored.

<sup>3</sup><https://www.w3.org/community/dpvcg/>

Within our approach, we use the `dpv:Consent` to represent user consent for their personal data and `dpv:PersonalDataHandling` to capture the analyst data handling. Other than these two classes, similar to SPECIAL, DPV provides taxonomies that represent general categories for each type of basic usage policy, which we briefly describe in the following:

- `dpv:PersonalDataCategory` to represent the category of personal data provided by a user and requested by an analyst.
- `dpv:Processing` to describe the categories of processes that are consented by a user and processes that are planned to be executed by an analyst.
- `dpv:Purpose` to describe the allowed purpose of the personal data processing (user) and the purpose of the data handling (analyst).
- `dpv:Recipient` to describe the designated recipient of personal data processing results (analyst) and the consented recipient of the results (user).

Furthermore, there are number of properties that we deemed necessary to be included in the consent and data handling setup:

- `dpv:hasPersonalDataCategory` to link either a consent or a data handling to the relevant personal data category.
- `dpv:hasPurpose` to link either a consent or a data handling to its purposes.
- `dpv:hasProcessing` to link either a consent or a data handling to the processing mechanisms.
- `dpv:hasRecipient` to link either a consent or a data handling to the target recipient.
- `:expiryTime` to store the expiry time for both user consent and data handling. We decided against using `dpv:expiry` to include `dpv:PersonalDataHandling` as a domain and to use `xsd:DateTime` as its range.
- `dpv:hasProvisionTime` to store the provision time of a user consent.
- `dpv:hasWithdrawalTime` to store the withdrawal time of a user consent.

While the generic concepts provided by DPV vocabularies cover a wide range of subjects, they might not cover the requirements for the fine-granular consent definition in domain-specific use cases. This is especially the case for the personal data categories and recipient, since the level of granularity can be different for various cases and would require an exten-



Fig. 3. An excerpt of the extended Data Privacy Vocabulary (DPV)

sion mechanism to be adopted. To address this issue, we adopt the vocabulary extension approach we previously developed in [21], where we adapted and extended the SPECIAL vocabularies in the smart-city domain. We reuse similar mechanism to extend DPV concepts with use case specific concepts to allow usage policy matching on fine-granular user consent for our evaluation in Section 6. An excerpt of the extended DPV<sup>4</sup> is shown in Figure 3.

#### 4.1.3. Automated Compliance Checking

In SPECIAL [20], policies and log events are described in semantically unambiguous terms from the same taxonomies defining usage policies, hence it facilitates automatic compliance checking. To allow automatic compliance checking, the usage policy enforced by a data controller contains the operations that are permitted within the data controller’s organisation.

In this work, we reuse the same mechanism as in SPECIAL approach [9], where the usage policy  $U_c$  attached to a `dpv:PersonalDataHandling` complies with the usage policy  $P_s$  in the data subject’s `dpv:Consent` if and only if all the authorisations in  $U_c$  are also authorised by  $P_s$ , that is,  $U_c$  complies with  $P_s$  if and only if

$$U_c \subseteq P_s \quad (1)$$

<sup>4</sup>Full version is available online: <http://w3id.org/wellfort/ns/dpv>

This implies checking whether the following axiom is entailed by the combined ontology containing a subset of DPV as described in Section 4.1.2 and the aforementioned auxiliary vocabularies:

$$\text{SubClassOf}(U_c P_s) \quad (2)$$

This entailment is supported by general inference engines for OWL 2 (e.g. HermiT and FaCT++) and further details of the compliance checking mechanism can be found in [9, 22]. Furthermore, Bonatti et al. [23] has developed a specialised reasoning engine that allows real-time OWL2 reasoning that enhanced the reasoning performance for the automatic consent check mechanism considerably.

The reasoning behind reusing the SPECIAL usage policy compliance checking mechanism for our approach is two-fold: (i) the approach is compatible with DPV, which is mainly due to the fact that DPV is developed based on the SPECIAL vocabularies, and (ii) our successful experience in adopting the SPECIAL approach in several of our prior works [10, 21, 24].

## 4.2. Auditability

In this subsection, we explain the extension made to accommodate auditability concerns and questions. First, we describe the methodology how to capture the provenance data in our architecture (Section 4.2.1). Next, we present a list of predefined provenance-focused competency questions from our scenario and requirement analysis (Section 4.2.2). In the following, we report on our extensions to the PROV-O data model, which are tailored to these concerns (Section 4.2.3).

### 4.2.1. Methodology

Miles et al. [25] offers a methodology for developing provenance-aware applications, arguing for the identification of actors and data flows of applications. We apply and adapt a condensed form to define our auditability needs:

1. *Definition of provenance competency questions:* These questions are to validate that the correct information pieces have been gathered to answer them and will be provided in Section 4.2.2.
2. *Identification of main actors, activities and entities:* During the definition of provenance questions, we identified the main actors, activities and entities.

3. *Extension of provenance data model:* We have chosen PROV-O as the foundation for our auditability. However, adaptations are needed to answer our defined competency questions. The process from Step 2 and 3 will be described in Section 4.2.3.

4. *Mapping of system data sources to provenance data:* In this step, audit data artefacts from the system have to be analysed and mapped accordingly to the extended PROV-O.

5. *Translation of competency questions into queries on the data:* This step is concerned with the translation of questions defined in Step 1 into queries. We will briefly describe the process from Step 4 and 5 in Section 4.2.4.

### 4.2.2. Auditability Competency Questions

This subsection concerns with the identification of audit-related questions from our scenario and requirements described in Section 2. In the process, we discussed the questions with both project members as well as our use case partners.

We group the questions according to the main actors: 1) *user-centric*, which are relevant for users to gain trust in the system and how data was used, which consent was provided, 2) *analyst-centric*, focusing on search parameters, system variables and the software environment of the analyst.

#### *User-centric questions:*

- If user data has been analysed, in which studies was it used?
- If user data has been analysed, which purposes did the study serve?
- What user consent was given for a dataset at a specific time?
- Which entities have analysed user data?

#### *Analysis-centric:*

- Which search attributes were specified for a certain study?
- Which software dependencies, deployment settings and software versions have been used when conducting a study?
- Are there differences in the results between two queries with the same search parameters but at a different time?
- Can these differences be related to user consent, additional data upload or a change in the platform/software environment?
- Which data transformation operations were conducted on data?

- Which software was used to analyse user data?
- What scripts were used to analyse user data?

#### 4.2.3. Extension of Provenance Data Model

We decided to adopt PROV-O<sup>5</sup> as the basis of our audit data model, due to its broad coverage of general provenance concepts, ease of use and extensibility [26]. However, while PROV-O is an established generic framework for provenance management, there is still a need to adapt and extend it to support the data tracing within the scope of our platform.

During the definition of provenance questions, we identified two specialised main actors, which influence the data flow: *users*, providing user data and associated user consent and *analysts* searching for and analysing user data. Another relevant agent is the *platform* itself, which provides the capabilities of data upload, management and analysis. Furthermore, we see in Figure 2 the main three generic activities: *user data upload*, *data selection* and *analysis*. Based on actors and activities, we identified the specialised entities such as the *query result* which captures the search parameter an analyst provides or the *environment* in which the analysis was conducted. Based on our analysis above, we have extended PROV-O with more specialised concepts in our platform as illustrated in Figure 4 and will be described as follows:

##### Subclass of *prov:Activity*

- Upload (*prov:Upload*) constitutes an upload activity of user data to the platform.
- TransformedStudyData (*wprv:TransformedStudyData*) describes the transformation of data from the secure repository to the data analysis.

##### Subclass of *prov:Agent*

- Platform (*wprv:Platform*) describes the platform characteristics, such as general version, data model version and other software information.
- Analysis Environment (*wprv:AnalysisEnvironment*) represents the analysis environment in which the analyst can conduct analysis activities of user data.

##### Subclass of *prov:Entity*

- Dataset (*dcat:Dataset*) represents a user dataset uploaded by an organisation to the platform.

- Query Result (*wprv:QueryResult*) describes characteristics of query results of queries issued by analysts.
- Study Purpose (*wprv:StudyPurpose*) provides the context of a study, such as recipient, purpose, processing type and duration.
- Study (*wprv:Study*) represents the analysis activities conducted by analysts.
- Script (*wprv:Script*) captures the scripts that analysts produce to analyse the provided data.
- Library (*wprv:Library*) collects dependencies such as software libraries that are used in scripts.
- Environment (*wprv:Environment*) states the environment in which scripts were processed such as operating system, architecture etc.
- Data (*wprv:Data*) describes the data that is produced by the scripts.

The specific extensions of the ontology support audit-related questions concerning the respective concepts.

#### 4.2.4. Data Source Mapping and Query Translation

The *Audit Box* component in WellFort architecture provides a basic building block for audit data management previously explained in Section 3.4.

In the *provenance collector* of the *Audit Box*, we define a set of mappings between the raw provenance data provided by the other two components of the architecture, namely *Secure Repository* and *Trusted Analysis Environment*, and our extended PROV-O data model to ensure that we can retrieve and collect the provenance data properly. Potential mapping technologies include RML for structured data [27] or ON-TOP [28].

Since we store the provenance data as RDF graphs, we naturally rely on the SPARQL query language to formulate auditor queries to the system. To this end, we support the auditors by implementing the competency questions described in Section 4.2.1 as predefined queries and provide it to the user in the *Auditor Interface*. Furthermore, it is also possible for auditors with advanced SPARQL capabilities to implement their own SPARQL queries on their audit process.

## 5. Prototype

Our overall guidelines when implementing the prototype of our platform were as follows:

<sup>5</sup><https://www.w3.org/TR/prov-o/>

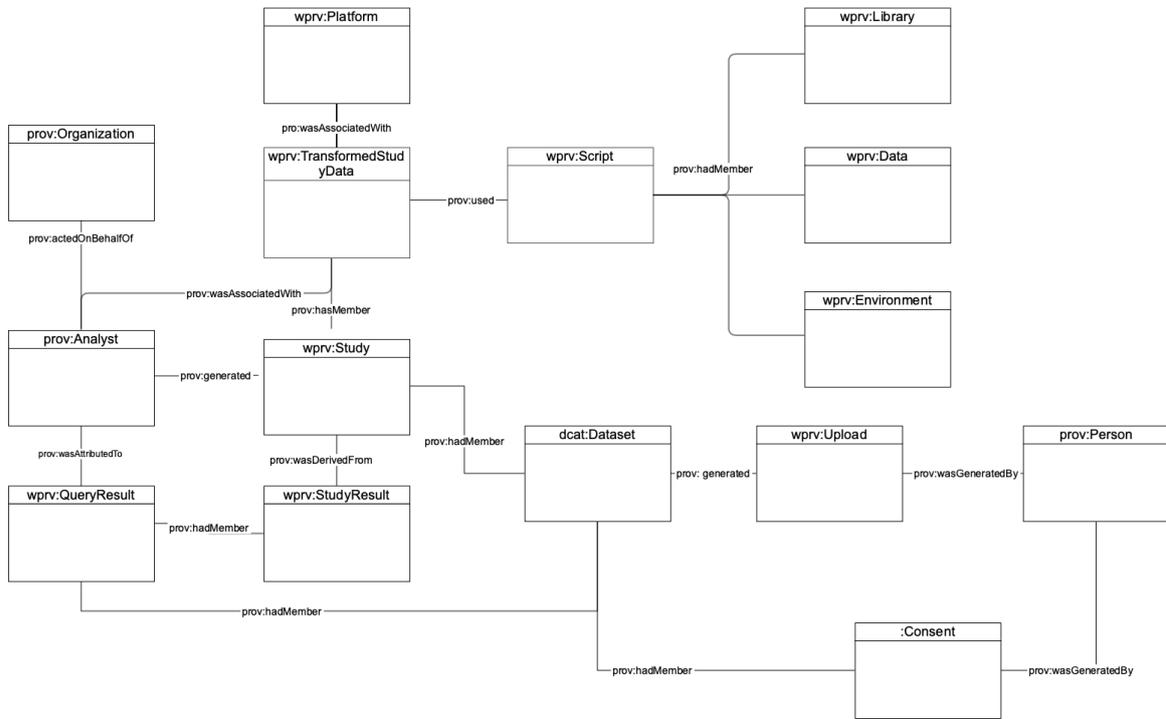


Fig. 4. WellFort PROV-O extension

- For data handling, we re-use proven data formats, standards, and ontologies as much as possible, to avoid common pitfalls when designing from scratch, and to foster re-use and interoperability with other systems.
- For software components, we re-use existing, well proven and tested components, e.g., for the basic repository system, database servers or analysis platform, and focus our effort on orchestration and integration of these components into a functional platform serving the intended purpose of providing auditable, privacy-preserving data analysis.

In the following, we provide details on our prototype and the technologies utilised.

### 5.1. Secure Data Repository

We implemented the *Data Repository* using the *Invenio framework*<sup>6</sup>, which is used by institutions to

build data repositories that enable publication of data according to FAIR principles. The *invenio* framework allows us to store and organise datasets uploaded from users' applications, assign PIDs to datasets, and manage access control. *Invenio* framework is a modular web application that is implemented in python. We added custom-built modules to achieve the needed functionality.

We implemented the *Extractor* and *Controller* modules as *Flask*<sup>7</sup> services inside the web application. They are responsible for automated extraction of metadata and consent, as well as managing the experiment. The *Experiment Setup Interface*, also built with *Flask*, is the only component the Analyst can interact with. It performs validation of the user input, provides pages for search queries, and functionality to setup the Analysis Environment. The *Loader* component was implemented with *watchdog*<sup>8</sup>, and is capable to supervise file system events, such as file creation. Its purpose is to *listen* for specific file patterns created (i.e., of per-

<sup>6</sup><https://inveniosoftware.org/>

<sup>7</sup><https://flask.palletsprojects.com>

<sup>8</sup><https://github.com/gorakhargosh/watchdog>

sonal data), that in turn invoke the import of the data to the analysis database.

The core part of the *Triplestore* component is built using SparkJava<sup>9</sup> and RDF4J<sup>10</sup>, with GraphDB<sup>11</sup> as its storage component. The *Triplestore* is responsible for retrieving a set of dataset ids based on the evaluation setup provided by an Analyst, taking into account the availability of the requested data, as well as the compliance between user consents and the setup.

To this end, we developed a *Usage Policy Compliance Check* engine to complement the *Triplestore* with a compliance check mechanism, which is built using OWL-API<sup>12</sup> and Hermit reasoner [29]. We explained the compliance check mechanism in Section 4.1.3, where user consent and experiment setup data is structured into OWL class restrictions and checked for their compliance using OWL2 reasoning, ensuring that only dataset ids with compliant consent are returned to the Analyst. We describe the data structure, which is based on SPECIAL Policy Language and DPV, on Section 4.1.2

The *Privacy-Preserving Data Publishing* module is implemented using open-source tools for data anonymisation. There are two methods enabling privacy-preserving download: (i) *k*-anonymisation [13] and (ii) synthetic data generation. For *k*-anonymisation we utilise ARX Data Anonymisation Tool<sup>13</sup>. The tool allows for extending the privacy guarantee of the downloaded data by supporting multiple additional data privacy models, such as *l*-diversity, *t*-closeness, etc. The choice of identifiers and quasi-identifiers is done inside the platform and is based on DPV vocabularies described in Section 4.2.3. Direct identifiers are described by `dpv:IdentifyingCategory` and are suppressed for the data download. Quasi-identifying attributes are those falling under a subset of class `dpv:PersonalDataCategory`, such as `dpv:Ethnicity`, `dpv:Geographic`, e.g. home address and `dpv:PhysicalCharacteristics`. Synthetic Data Vault (SDV)<sup>14</sup> is used for generating synthetic data in the *Privacy-Preserving Data Publishing* module. It is a Python library based on probabilistic graphical modelling and deep learning for generat-

<sup>9</sup><https://sparkjava.com/>

<sup>10</sup><https://rdf4j.org/>

<sup>11</sup><https://graphdb.ontotext.com/>

<sup>12</sup><https://owlcs.github.io/owlapi/>

<sup>13</sup><https://arx.deidentifier.org/>

<sup>14</sup><https://sdv.dev/>

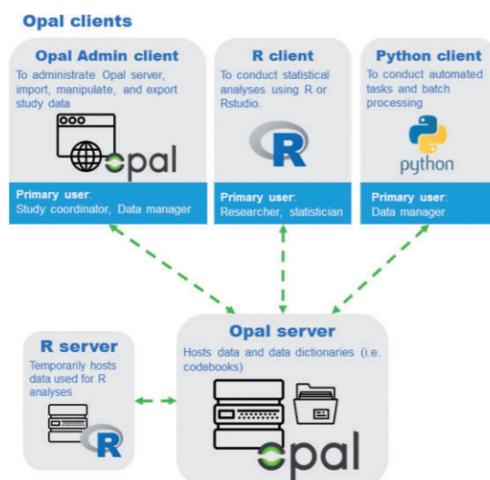


Fig. 5. Connection between Opal and clients [30]

ing synthetic data that has the same format and statistical properties as the original data.

## 5.2. Trusted Analysis Environment

The *Trusted Analysis Environment* contains the mechanisms for secure and privacy-preserving data analysis. The mechanism is provided by the analysis server, the analysis database, where the study data are temporarily stored, and the analysis interface, which allows the Analyst to interact with the data.

**Analysis Server** For this component, we rely on Opal<sup>15</sup>, which is a data management system that provides tools for importing, transforming, describing and exporting data. Data managers can securely import a variety of data types, e.g., text, numerical data, images, geo-localisation, etc. and formats, e.g., from SPSS or CSV.

Opal is built on REST web services, which allows connections from different software web clients using the HTTPS protocol, as shown in Figure 5. These clients are designed to ensure a broad range of specialised tasks to be achieved by different users. For example, a Python client allows data managers (in this case our platform administrator) to automate repetitive tasks such as data imports and access permission management. An R client, e.g., RStudio, allows conducting statistical analysis of the data stored on Opal. Opal

<sup>15</sup>Opal is open-source and freely available at [www.obiba.org](http://www.obiba.org) under a General Public License (GPL) version 3, and the metadata models and taxonomies that accompany them are available under a Creative Commons licence.

provides the access to the *Analysis Database* to the authenticated users. The *Loader* sets the user permissions for the data tables loaded into *Analysis Database* so that the Analyst can access them through the *Analysis Interface*. This allows the Analysts to operate only on the data that has been retrieved for studies they initiated. Client authentication is performed using either certificates, username/password, or token mechanisms.

**Analysis Database** In the *Trusted Analysis Environment*, Opal is connected to the *Analysis Database*, a PostgreSQL<sup>16</sup> database that temporarily stores the data relevant for the analyses. The tables from all studies conducted on our platform are imported as CSV and JSON files from the *Loader* and queried only through the *Analysis Interface*.

The details about the data formats within the *Analysis Database* are provided in Section 5.4.2.

**Analysis Interface** In our prototype platform, we utilise the R client of Opal as an interface where the Analyst conducts her studies on selected data using R DataSHIELD packages. DataSHIELD [14] is an open-source software that provides a modified R statistical environment linked to an Opal database. It was originally developed for remote, non-disclosive analysis of biomedical, healthcare and social-science data, although it can be used in any setting where micro-data (individual-level data) must be analysed but cannot or should not physically be shared with the research users for various ethico-legal reasons, or when the underlying data objects are too large to be shared. DataSHIELD implements >100 statistical methods<sup>17</sup> that embed privacy-protection traps, which prevent the Analyst from identifying individual data subjects or inferring the value of particular variables in given subjects based on analytic results. For instance, subsetting is functionally restricted by not generating any subset data set containing, by default, only 1-4 observations since results based on these subsets might be disclosive.

RStudio Server<sup>18</sup> is employed as a user interface of our platform, which an Analyst uses to connect to Opal and perform privacy-preserving analysis via DataSHIELD functions.

<sup>16</sup><https://www.postgresql.org/>

<sup>17</sup><https://rdrr.io/github/datashield/dsBaseClient/man/>

<sup>18</sup><https://rstudio.com/>

### 5.3. Audit-Box

The *Audit Box* collects and integrates provenance data from the entire platform, and provides an interface for Auditors to audit processes within the platform. *Audit Box* consists of three components of which the implementation will be described in the following.

```
{
  "description": "...",
  "consent": [ {
    "data_type": "Demographic,HealthRecord",
    "processing": "PseudoAnonymise",
    "purpose": "Security",
    "recipient": "CompanyH",
    "until": "2022-02-10"
  } ],
  "execution": {
    "invenio_ts": "2021-01-20T21:23:53",
    "end": "2021-01-20T21:23:53",
    "start": "2021-01-20T21:23:52"
  },
  "triggered_by": "api.py:40",
  "invenio_id": "378",
  "patient_id": "243",
  "upload_id": "0a5c1d64"
}
```

Listing 3: An example of raw provenance data

**Provenance Collector.** The task of this component is retrieving and transforming raw provenance data into an integrated RDF graph. In our prototype, we implement this component as a set of REST services. The services accept raw provenance data in JSON format from various components in the *Secure Repository* and *Trusted Analysis Environment* and transform it into PROV-O compliant RDF graphs.

Specifically for *Trusted Analysis Environment*, we rely on an existing tool to collect provenance from R-Scripts<sup>19</sup>. RDT format its provenance data using an extended PROV-DM<sup>20</sup> model and allows to collect all relevant information from the analysis environment such as operating system, language version, and script directory. We use RML mappings [27] and the caRML engine<sup>21</sup> to execute the transformation from the raw data into RDF graphs.

We provide an example of raw provenance data produced by *Secure Repository* (Listing 3) and the respective RDF graph result in Listing 4. These examples capture the provenance of the *upload* process, where a user finished uploading her personal data with a consent to the *Secure Repository*. Both the upload

<sup>19</sup><https://github.com/End-to-end-provenance/rdtLite>

<sup>20</sup><https://www.w3.org/TR/prov-dm/>

<sup>21</sup><https://github.com/carmel/carmel>

activity and the uploaded dataset is captured as instances of `prov:Activity` and `prov:Entity` respectively. We show in the resulted provenance graph that the upload process is associated with a python script `api.py` and executed in one second.

```

7 # upload activity
8 id:upload-0a5c1d64 a prov:Activity, wprv:Upload;
9   prov:startedAtTime "2021-01-20T21:23:52";
10  prov:endedAtTime "2021-01-20T21:23:53";
11  prov:wasAssociatedWith id:script-api-py-40 .
12
13 # uploaded dataset
14 id:dataset-303 a prov:Entity, dcat:Dataset;
15   prov:wasGeneratedBy id:upload-0a5c1d64;
16   prov:generatedAtTime "2021-01-20T21:23:53";
17   prov:wasAttributedTo id:patient-367 ;
18   :hasConsent id:consent-0a5c1d64.
19
20 # consent
21 id:consent-0a5c1d64 a dpv:Consent, owl:Class ;
22   owl:equivalentClass [ owl:intersectionOf (
23     [ a owl:Restriction ;
24       owl:onProperty dpv:hasPersonalDataCategory ;
25       owl:someValuesFrom [ owl:unionOf
26         ( dpv:Demographic dpv:HealthRecord ) ] ]
27     [ a owl:Restriction ;
28       owl:onProperty dpv:hasProcessing ;
29       owl:someValuesFrom dpv:PseudoAnonymise ]
30     [ a owl:Restriction ;
31       owl:onProperty dpv:hasPurpose ;
32       owl:someValuesFrom dpv:Security ]
33     [ a owl:Restriction ;
34       owl:onProperty dpv:hasRecipient ;
35       owl:someValuesFrom :CompanyH ]
36     [ a owl:Restriction ;
37       owl:onProperty dpv:expiryTime ;
38       owl:someValuesFrom [
39         a rdfs:Datatype ;
40         owl:onDatatype xsd:dateTime ;
41         owl:withRestrictions ( [ xsd:maxInclusive
42           "2039-04-15T23:59:59"^^xsd:dateTime ] )
43       ] ] ] ] ] .

```

Listing 4: The resulted RDF graph from Listing 3

*Knowledge Base.* We collected the provenance graph from the Provenance Collector into the Knowledge Base component. This component stores and manages the provenance data, which is necessary to support Auditors in conducting their tasks. Similar to the metadata store in the *Secure Repository*, we utilise GraphDB as our Knowledge Base for *Audit Box*. The communication between GraphDB and the Provenance Collector is handled using Eclipse RDF4J<sup>22</sup> that allows seamless I/O communication between them.

*Auditor Interface.* In our current prototype, the Auditor Interface relies on the SPARQL user interface of GraphDB, both for storing the pre-defined queries for the audit-related competency questions described

in Section 4.2.1 and to allows auditors to formulate their own queries related to their tasks.

#### 5.4. Data Model

This section explains the data model used in our prototype implementation: (i) a domain-specific vocabulary on health record and (ii) the analysis data model for our *Analysis Database*.

##### 5.4.1. Domain Specific Vocabularies: Health Record

The proposed conceptual architecture does not impose limitations on the formats of the data used in the Data Repository and therefore, the data model depends on the specific use case. For the prototype presented in this paper, we consider the following characteristics for our health record data model:

1. **Based on open-source standards** that guarantees stable specifications for platform users.
2. **Designed for interoperability and information exchange**, which is critical in the healthcare domain due to the heterogeneity of data and stakeholders [31]; and
3. **Widely used in the healthcare communities** to allow easier adoption for stakeholders that are already familiar with the data model.

Based on these criteria, we decided to use the Fast Healthcare Interoperability Resources (FHIR) developed by Health Level 7 International (HL7) as the data format in our Data Repository for our prototype. HL7-FHIR is the latest standard by HL7 for electronically exchanging healthcare information based on emerging industry approaches [32]. FHIR offers a number of advantages for prototype development, such as its solid foundation in Web Standards and supports for RESTful architectures.

In our use case, we used FHIR to represent Patient, Diagnostic Report, and Observation. While FHIR recommends the use of LOINC<sup>23</sup> for observation data, we decided to use the custom classification codes provided by our use case partners for our prototype for practical reason.

##### 5.4.2. Analysis Database

The privacy-preserving nature of the *Trusted Analysis Environment* comes with the price of reducing the data utility and flexibility of an Analyst compared to full data access. Another challenge is that the data orig-

<sup>22</sup><https://rdf4j.org/>

<sup>23</sup><https://www.hl7.org/fhir/valueset-observation-codes.html>

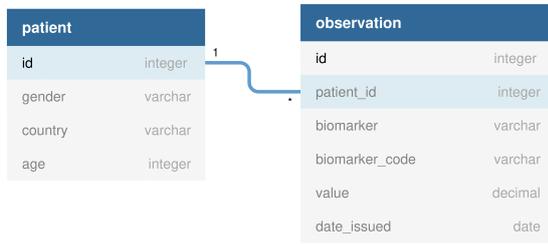


Fig. 6. Analysis Database schema

```

id,patientId,biomarker,value,biomarkerCode,issued
0,0,Diet,normal-diet,diet,2019-06-11T12:32:24Z
1,1,Diet,normal-diet,diet,2019-02-07T09:13:07Z
2,2,Diet,low-carb-diet,diet,2019-02-01T14:56:31Z
3,1,Vitamin B12,157.13,VIT-B12,2019-02-07T12:32:24Z

```

Fig. 7. An excerpt of data table CSV file for table *Observation*

inating from different sources needs to be transformed to the same common format, while keeping the data intuitive for Analysts to work with. Therefore, in the *Analysis Database* we aim to keep the most relevant data for the analysis in a rather concise but comprehensive format.

The database consists of two default tables defined by the types of FHIR resources *Patient* and *Observation*:

- *Patient*: provides information about the users included in the study, such as sex, age and country.

- *Observation*: contains information about the specific measurements for a user. A biomarker can be any attribute contained in the source databases, for instance, blood sugar, activity level, etc. Each biomarker is represented with its code *biomarkerCode* and full name *biomarker*. A patient can have multiple measurements of the same biomarker measured at a different moment in time, so the attribute *issued* provides the information of when the specific measurement has been taken.

The two tables are connected by a foreign key *Observation:patientId - Patient:id*. The database schema can be seen in Figure 6.

As specified by Opal, each table is represented by (i) one CSV file called *data table* and (ii) one JSON file called *data dictionary*.

- *Data table*: the columns of this table correspond to the attributes in the actual table *Patient/Observation*. Each row is one data entry. An example data table file is shown in Figure 7.

- *Data dictionary*: this file is used to describe one data table. These tables must follow the scheme defined by Opal which imposes a set of fields such as *entityType*, *index*, *name*, *unit*, *valueType*, etc. for describing each of the attributes form *data table*.

## 6. Evaluation

In this section, we demonstrate the platform through representative use cases. In the summary we assess how well the demonstrated scenarios answer to the requirements posed in Section 2.

### 6.1. Data

For the evaluation we use data originating from two sources, both from SMEs that provide a health-care or lifestyle related application to end-users. For confidentiality reasons, we denote these sources as Company (or Datasource) M and Company H. A subset of the at-

Table 2  
Attributes of evaluation data

Datasource M		Datasource H	
attribute	type	attribute	type
Age	int	Birthday	date
Sex	text	Sex	text
Country	text	Sympathicus activity	dec
Blood Glucose	dec	Parasympathicus activity	dec
Lipids	dec	Heart rate	dec
Liver	dec	Pulse/Respiration Quotient	dec
Kidney	dec	Vitality index	dec

tributes provided by these two data sources are shown in Table 2. We generated 1200 synthetic datasets for 600 users based on the original data from both companies with varying consent configurations. Each dataset contains a single record for patient data (e.g., ID, age, and country) and health observation attributes (cf. Table 2) that we observe in the use cases for a single day. The synthetic datasets are generated using Synthetic Data Vault (SDV)<sup>24</sup> model trained with original datasets provided by companies M and H. Further details on the evaluation users, datasets and the associated user consent is given in Table 3, where **processing** and **expiry** refer to the dataset consent for processing and consent expiry time, respectively. We set

<sup>24</sup><https://sdv.dev/>

the consent of all user data to the default value of `dpv:hasPurpose` to `dpv:AcademicResearch` and `dpv:hasRecipient` to `TUWien`.

Table 3  
Synthetic users and datasets for the evaluation

users	dataset	processing	expiry
u0-49	ds0-49 (H) ds550-599(M)	dpv:Processing	2021-01-05
u50-199	ds50-199 (H) ds600-749 (M)	dpv:Processing	2021-01-10
u200-399	ds200-399 (H) ds750-949 (M)	dpv:Adapt, dpv:Combine, dpv:Analyse	2021-01-10
u400-499	ds400-499 (H) ds8=950-1049 (M)	dpv:Analyse	2021-01-10
u500-549	ds500-549 (H)	dpv:Processing	2021-01-10
u550-599	ds1050-1099 (M)	dpv:Processing	2021-01-10

## 6.2. Use cases

We define four use cases which will be processed in our platform:

- UC1: *Descriptive analysis of data from a single data source*
- UC2: *Descriptive analysis of data from multiple data sources*
- UC3: *Predictive modelling*
- UC4: *Data publishing*

UC1, UC2 and UC3 describe scenarios of using *Trusted Analysis Environment*, while UC4 describes the purpose of *Privacy-Preserving Data Publishing* component. UC1 and UC3 deal with data originating from a single data source, while presenting the spectrum of functionalities provided to the Analyst. UC2 focuses on a study utilising data from multiple sources. Within each use case, we provide a discussion on privacy concerns of the respective analyses.

We simulate the following scenario to introduce dynamic aspects:

- Initial state: all user datasets are uploaded and collected in the *Secure Repository*.
- Day 1 (05-01-2021): UC1 and UC2 are executed. UC1 analysis setup (expiry: 05-01-2021; processing: `dpv:Adapt` and `dpv:Analyse`) yields 450 datasets in total from Datasource H. UC2 analysis setup (expiry: 05-01-2021; processing: `dpv:Adapt`, `dpv:Combine`, and `dpv:Analyse`)

yields 800 datasets in total, 400 each from Datasource H and Datasource M.

- Day 2 (06-01-2021): User consent from 50 users expired (u0-u49). Furthermore, 50 users (u200-u249) removed `dpv:Adapt` from their consent
- Day 3 (07-01-2021): UC3 and UC4 are executed. Due to the changes and expiration of some user consents, UC3 analysis setup (expiry: 07-01-2021; processing: `dpv:Adapt` and `dpv:Analyse`) only yields 350 datasets from Datasource H, which is less than UC1 with an identical evaluation setup. UC4 analysis setup (expiry: 07-01-2021; processing: `dpv:Share`, `dpv:Anonymise`, and `dpv:Derive`) yields 200 datasets from Datasource M.

The flow for each of our use cases shares some common steps:

1. The **Analyst** initiates their study by selecting relevant attributes, specifying the study purpose and study expiry date.
  2. the **platform** yields the available datasets according to:
    - the user consent that needs to match with the Analyst's study purpose
    - the expiry date of data consent that needs to be later than the study expiry date.
- The number of the available datasets needs to exceed the predefined *threshold*, otherwise the analysis is considered to be potentially disclosive and the platform terminates the study setup. For our evaluation, we set the *threshold* to 10.
3. For each study, after its initialisation, the **platform** transforms the available datasets into the files describing two relational tables, *Patient* and *Observation*, as discussed in Section 5.4.2.
  4. For each study the **platform** defines the user permission in Opal for the Analyst that allows using data via DataSHIELD - this way each Analyst only sees the studies they initiated.
  5. When the study is initiated, the **Analyst** uses DataSHIELD and their credentials to connect to Opal and perform the analysis in case the use the *Trusted Analysis Environment*, or download their data if they use the *Privacy-Preserving Data Publishing* component.

### 6.2.1. UC1: Descriptive analysis of data from a single source

In our first use case, we assume the Analyst wants to conduct a set of descriptive studies to visualise se-

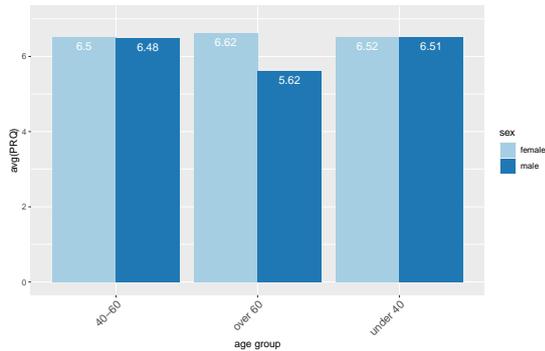


Fig. 8. UC1: Average PRQ of patients

lected attribute values across the population of platform users. Let this attribute be *Pulse/Respiration-Quotient (PRQ)* originating from the database of *Data-source H*. The Analyst then wants to study how the *PRQ* values differ between age groups and sex.

In our evaluation, the platform finds 450 participants with a set of records of the selected attributes, therefore the table *Patient* contains 450 records of users' age and sex, and table *Observation* 450 records of *PRQ*.

Figure 8 shows the results of the analysis, i.e. the average *PRQ* for different age groups (under 40, 40-60 and over 60) for male and female participants separately. The averages of *PRQ* are obtained by DataSHIELD functions and then plotted using the native R plotting package *ggplot*. The Analyst uses a set of data manipulation and statistical functions from DataSHIELD, for instance:

- Joining tables via `ds.cbind`: the Analyst needs to join the tables *Patient* and *Observation* by *patient\_id*.
- Subsetting via `ds.dataFrameSubset`: this function is used to subset the table joint in the previous step by the values of the attribute age and sex.
- Mean function via `ds.mean`: the Analyst uses this function to calculate the mean of each subset.

The listed DataSHIELD functions for subsetting and calculating the mean value are potentially disclosive functions. For example, applying subsetting repeatedly might lead to revealing the actual values from the table. DataSHIELD achieves privacy-preserving versions of these functions by prohibiting subsetting and calculating simple statistics, such as mean, on less than 4 data samples. Functionality for joining tables is implemented such that the method does not return any value to the Analyst, except for potential error mes-

sages. The joined table is persisted on the server in a variable, which the Analyst can use in other functions by name, therefore no values can be disclosed in the process.

### 6.2.2. UC2: Descriptive analysis of data from multiple sources

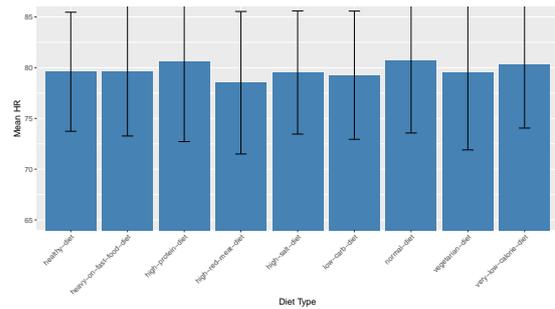


Fig. 9. Multiple source use case example: Mean and standard deviation of heart rate of users with different dieting profiles.

In this use case, the Analyst wants to conduct a study on attributes that are originating from separate data sources. For example, the Analyst wants to compare the resting heart rates of people with different dieting profiles, such as vegetarian diet, low-carbohydrate diet, etc. For that, the platform needs to collect the data from two applications. *Datasource M* provides the attribute *diet* and *Datasource H* provides the attribute *heart rate*. This analysis is only possible on users which use both applications. Users control the linking of their data and have to enable linking between those applications upon data submission (in their applications).

In our evaluation, the *Patient* table contains 400 user records and *Observation* table contains 800 records for these 400 users; 400 records of heart rate and 400 of diet type. The Analyst groups the patients by their diet type, and then for each dieting group calculates mean and standard deviation of heart rate values. To highlight a few steps, let us list the specific DataSHIELD functions that allow these operations:

- Subsetting via `ds.dataFrameSubset`: splitting the *Observation* table depending on the attribute 'biomarker' such that each subtable contains one type of biomarker, i.e. one table for diet, another table for HR.
- Joining tables via: `ds.merge`: joining the subtables from the previous step on attribute *patient\_id* and in that way pivot the *Observation* table into a table with the new set of columns: [*patient\_id*, *diet\_value*, *HR\_value*, ...]

- Aggregation via `ds.meanSdGp`: given that the Analyst subsets the pivoted table from the previous step by `diet_value`, this function obtains the mean and standard deviation of heart rate values for each dieting group.

The resulting analysis is shown in Figure 9. It shows the mean and standard deviation of the heart rate value for individuals in different dieting groups.

From the privacy-preserving perspective, the function `ds.merge` is achieved in the same manner as previously mentioned `ds.cbind`, by not returning any outputs to the Analyst. The implementations of aggregation functions are achieved similarly as subsetting, by prohibiting operations on less than 4 data samples.

### 6.2.3. UC3: Predictive modelling

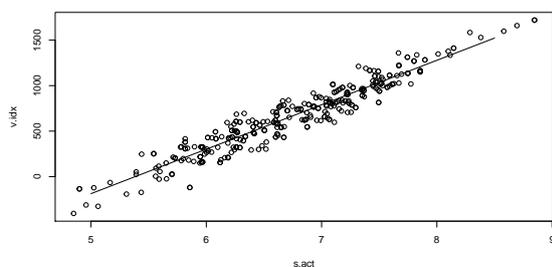


Fig. 10. UC3: Univariate GLM for target attribute Vitality Index and predictive attribute Sympathicus Activity

The Analyst wants to conduct a more sophisticated study to predict the value of Vitality Index (a measure of the resources that are available to the organism for activity and health maintenance). To that end, the Analyst wants to learn and apply a generalised linear model (GLM) that describes the linear combination of the attributes *sex*, *parasympathicus activity* and *sympathicus activity* and enables predicting the expected value of the attribute *Vitality Index* of a person. Such a predictive model functionality is provided as a DataSHIELD function, and thus readily available.

*Datasource H* provides the data about vitality index, sympathicus and parasympathicus activity and sex of its patients, thus this analysis task is a single source setting. The table Patient contains 350 records and table Observation 1050 records of vitality index, parasympathicus activity and sympathicus activity for the same set of patients.

The Analyst is now free to use any type and number of DataSHIELD and native R functionalities to obtain privacy-preserving results. We single out some specific DataSHIELD methods used in this study:

- Generalised Linear Model via `ds.glm`: this function is used for fitting GLM. The implementation supports binomial, gaussian and poisson error distribution functions.
- Visualisation via `ds.scatterPlot`: one of the non-disclosive plotting functions in DataSHIELD. This function may be used in this study to plot data points, taking into account two attributes, e.g., Vitality Index and Sympathicus Activity.

One option for the Analyst is to examine the relationship between Vitality Index and one of the aforementioned attributes by fitting a univariate GLM. The Analyst chooses sympathicus activity as a predictor attribute for Vitality index and Gaussian error distribution. Figure 10 shows the distribution of data points and the estimated linear model for target attribute Vitality Index, using Sympathicus Activity as a predictive attribute. Furthermore, to fit a more complex

```

$family
Family: gaussian
Link function: identity

$formula
[1] "v.idx ~ s.act + ps.act + sex"

$coefficients
      Estimate Std. Error  z-value  p-value  low0.95ci  high0.95ci
(Intercept) -2332.955688   61.37485  -38.011594  0.000000e+00 -2453.24818 -2212.66320
s.act        359.214501    15.10462   23.781766  5.157566e-125  329.60999   388.81901
ps.act       114.199106    11.34462   10.066368  7.779775e-24   91.96406   136.43415
sexmale     -8.615345      19.75788   -0.436046  6.628034e-01  -47.34008   30.10939
    
```

Fig. 11. UC3: Output of multivariate GLM for target attribute Vitality Index and predictive attributes Sympathicus Activity, Parasympathicus Activity and Sex

model, i.e. multivariate GLM, the Analyst in addition includes more predictive attributes: sex and parasympathicus activity. Figure 11 shows the output of the DataSHIELD function `ds.glm`. The Analyst sees the coefficient estimates, standard error, p-value indicating significance of a predictor for the model, etc.

For both univariate and multivariate analysis, the output of `ds.glm` are the learnt estimates of regression coefficients and their standard errors, both of which intrinsically do not disclose any personal information. Implementation of the privacy-preserving GLM model for the usage with the data schema of our platform does not differ from its original form in native R. However, a scatter plot generated from the original data is disclosive, therefore not permitted in DataSHIELD in its original form. There are rather two implementations of the privacy-preserving scatter plot available: (i) based on *k*-Nearest Neighbours (*k*NN) and (ii) based on adding random noise. In the first implementation, instead of the actual data point, *k*NN for

each data point is calculated and the average value is plotted.  $k$  is chosen by the Analyst, but cannot be set to a value smaller than 3. In the second implementation, a random normal noise is added to the  $x$  and  $y$  coordinate of each data point. In this analysis the  $k$ NN method is used with  $k = 3$ .

#### 6.2.4. UC4: Data publishing

In this use case, the Analyst wants to download anonymised data so that they can perform an offline analysis on the given data, using other data analysis tools rather than the *Trusted Analysis Environment*. Reasons for this might be of technical nature, such as programming language preference, familiarity and reliance on another analysis framework, required functionalities that are not covered by DataSHIELD, etc. Furthermore, the Analyst might be interested in the structure of the data and particular data values rather than statistical analysis, or wishes to avoid being time-constrained for conducting their studies.

The Analyst chooses attributes *sex*, *age*, *country* and *Total Cholesterol* provided by the Company M. In the  $k$ -anonymised dataset the direct identifiers such as user full name or any kind of IDs are not displayed at all. Other indirectly identifying attributes with personal information such as user country, age and sex, i.e. quasi-identifiers, are generalised according to the hierarchies predefined by the platform for quasi-identifying attributes as discussed in Section 5.1:

- *sex*
  - \* Level 1: \* (i.e. fully generalised)
- *age*
  - \* Level 1: [20,40[, [40,60[, [60, 81[
  - \* Level 2: [20, 81[
- *country*
  - \* Level 1: SE (South Europe), SEE (South-east Europe), CE (Central Europe), ...
  - \* Level 2: \*

Furthermore, the exact values of non-identifying attributes such as *Total Cholesterol* remain in their original form, since the possibility of re-identify persons from database based on this information is minute. Tables 4 and 5 show a subset of data before and after applying  $k$ -anonymisation, respectively, where only the later is the output of this analysis.

Table 4  
Original dataset

id	sex	age	country	Total Cholesterol
382	female	32	PT	3.38
24	female	32	BG	4.58
239	male	25	NL	4.86
253	female	38	AT	4.42
293	female	25	BG	5.08
77	male	63	PT	4.6
212	male	53	HR	5.45

Table 5  
Anonymised dataset

id	sex	age	country	Total Cholesterol
*	female	32	*	3.38
*	female	[20,81[	*	4.58
*	male	25	*	4.86
*	female	38	AT	4.42
*	female	[20,81[	BG	5.08
*	male	63	PT	4.6
*	male	[40,60[	SEE	5.45

### 6.3. Auditability

In this subsection, we focus on the evaluation of consent and provenance mechanisms that enable auditability of the platform within the context of the previously described use cases.

*Auditor Queries* We stated in Section 2.2 that provenance data must be automatically captured from each platform component to cover the whole data lifecycle within the platform. Here, we demonstrate the capability of the system in using the captured provenance data to answer Auditor queries.

We select four questions that arose from the context of the use cases to verify our designed provenance and consent capabilities. Several questions, i.e., Q1-Q3, are parts of the common auditability competency questions described in Section 4.2.1. However, in some cases, Auditor may require more complex queries that are not in the pre-defined list. On these cases, Auditor can formulate her own SPARQL queries, such as in the case of Q4.

- Q1: Which datasets were used in a specific study and which users are associated to the datasets?
- Q2: Which software has been used to analyse data in a specific study?
- Q3: Which query parameters were used for the selection of data in an analysis setup?

```

1 SELECT ?study ?data ?user WHERE {
2   ?study a wprv:Study;
3   prov:hadMember ?result.
4   ?result prov:wasGeneratedBy ?transform.
5   ?transform prov:used ?qr.
6   ?qr prov:hadMember ?data.
7   ?data prov:wasAttributedTo ?user.
8   filter(regex(str(?study), "id:study-b79d9da5.." ))
9 }

```

Listing 5: Q1- Linking datasets used in a study to users

Study	User	Dataset
id:studyd98458b8...	id:patient-367	id:dataset-303
id:studyd98458b8...	id:patient-318	id:dataset-154

Table 6

Results Q1- Datasets with associated patients which were used in a study

```

17 select ?library where {
18   ?study a wprov:Study;
19   prov:hasMember ?studydata.
20   ?studydata prov:wasInfluencedBy ?script.
21   ?script prov:hadMember ?library.
22   ?library rdt:name ?library.
23   filter(regex(str(?library), "library" ))
24   filter(regex(str(?study), "id:study-b79d9da5.." ))
25 }

```

Listing 6: Q2- Used libraries in a study

- Q4: *What is the ratio between analysed datasets in the studies and the total number of datasets in the triplestore?*

**Q1.** Query Q1 addresses concerns regarding consent and data usage for study purposes (cf. Listing 5). In case of suspicions regarding the correctness of given consent and analysed data, auditors can verify which data was used for which purposes stated in the study description.

Q1 consists of multiple graph traversals based on Prov-O relations, connecting datasets with the corresponding user. This query demonstrate the capability of an Auditor to establish such a connection between a study, user id, user dataset metadata, and consent attached to the dataset. Table 6 shows the answer to Q1, linking a dataset used in a specific study to a user. Based on this information, an Auditor can inform concerned users on which studies their data was used and under which purposes, similar to the situation described in the scenario provided in Section 2.1.

**Q2.** Query Q2 (cf. Listing 6) is concerned with the usage of software libraries applied in the analysis as described in *UC-2 descriptive analysis of data from multiple sources*.

Libraries
"library('ggplot2')"
"library('opalr')"
"library('DSOpal')"
"library('DSI')"
"library('dsBaseClient')"

Table 7

Results Q2- An excerpt of used libraries in a study

```

10 select distinct ?sp where {
11   ?q a wprv:Query.
12   ?q :hasDataHandling ?dh.
13   ?dh owl:equivalentClass/owl:intersectionOf ?cls.
14   ?cls rdf:rest*/rdf:first ?item.
15   ?item owl:onProperty dpv:hasPersonalDataCategory.
16   ?item owl:someValuesFrom ?nodes.
17   ?nodes owl:unionOf ?list.
18   ?list rdf:rest*/rdf:first ?searchparameters.
19 }

```

Listing 7: Q3- Retrieving search parameters used in an Analyst query

In our prototype, the provenance trails of the libraries, environment, and scripts executed on the R-server are collected using RDT-extension and stored in the Audit-box. Q2 provides an answer to the question of which software libraries are used in a specific study, with example results is shown in Table 7. The purpose of retrieving such information is to increase the study reproducibility and also to ensure the legitimate existence of the analysis results.

**Q3.** Query Q3 refers to the search parameters for personal data categories from UC3 predictive modelling (cf. Listing 7).

Before being able to conduct an analysis, an Analyst has to provide details on which data categories of personal data are required for her analysis (e.g. Gender, Age, Heartrate etc.; we also called them as *search parameters*) to check the availability of the data. This parameter, together with the rest of data handling setup, which includes the purpose of the analysis and expiry time, among others, are captured and stored in the *Audit Box*. Using this provenance information, we can retrieve specific search parameters used for a Query (see Table 8 for answers to Q3). This information can be used by an Auditor to check whether there are any suspicious activities conducted by a certain analysts, e.g., if there are too many queries with varying search parameters by one analysts during a short amount of time.

**Q4.** Query Q4 (cf. Listing 8) shows the ratio of utilised datasets in studies compared to the actual available datasets in the triple store. This ratio (see Ta-

Search Parameters
:VitalityIndex
:SymphaticusActivity
:ParaSymphaticusActivity
:Sex

Table 8

Results Q3- Search parameters of an Analyst query

```

select ?qr (?sumdata/?usdata as ?ratio) where { {
  Select ?qr (count(?data) as ?sumdata) where {
    ?data rdf:type dcat:Dataset ;
    prov:wasAttributedTo ?user.
    ?data ^prov:hadMember ?qr.
  } Group by ?qr } . {
  Select ?qr (count(?totdata) as ?usdata) where {
    ?totdata rdf:type dcat:Dataset .
  } Group by ?qr }
}

```

Listing 8: Results Q4- Retrieving the ratio between data sample and general user population

Study	Ratio
id:query-result-d98458b8..	0.01
id:query-result-907989fb..	0.018
id:query-result-38d7ae5b...	0.01

Table 9

Ratio of users analysed in a study compared to total user population

ble 9) can be regarded as an indicator how representative the study sample is in relation to the overall population and could serve as a robustness factor within the platform. To this end, there are two possible reasons why datasets were not available for analysis: (i) it could be that the consent was not applicable to the purpose given by the Analyst, or (ii) the dataset do not contain that the personal data categories required by the Analyst.

#### 6.4. Summary

In this section we described the behaviour of our platform throughout the concrete scenario that encompasses the dynamic aspects of the data and consent management, and 4 different data analysis use cases. The coherent user story allowed us to explore the application of our platform in a holistic manner, i.e. covering the behaviour of all components of our platform through one process and showing the feasibility of the proposed auditable privacy-preserving data analysis platform.

To summarise the evaluation, we now assess our platform according to the requirements defined in Sec-

tion 2.2, targeting particular components of the platform. In Table 10, we explicitly list the requirements and indicate to which extent they are satisfied by the described scenario: fully, partially or unsatisfied. The listed requirements follow the structure from Section 2.2, therefore **R1.x** denotes requirements regarding privacy-preserving data analysis, divided into 3 subcategories, (i) data analysis, (ii) data publishing and (iii) usage policy management, and **R2.x** denotes requirements regarding auditability: (i) provenance data capture and tracing, (ii) provenance data inspection and analysis capability and (iii) metadata-retention after deletion.

Requirements R1.1, R1.2 and R1.5 are fully satisfied in the platform, which is shown by each use case from our evaluation. The data for analysis is stored in the transformed format inside the analysis database only, and the outputs of the analysis are shown to the Analyst and not stored in the platform at all. Both of these intermediate products of data analysis are part of the analysis environment and are deleted together with the environment, which is defined by the data expiry. UC1 and UC3 directly support the requirement R1.2. R1.3, however, faces the limitations regarding user linkage across different platforms that we already discussed in the scope of UC2. Our platform offers the infrastructure to combine data with different schemes of data originating from different sources. However, record linkage methods impose some privacy and accuracy concerns [33], therefore we give the user full control over linking their data. As our solution depends on user engagement, we evaluate it as partially satisfied. Furthermore, requirements R1.6, R1.7 and R1.8 are described via UC4. R1.9 aims to address the potential surface for privacy attacks when malicious has the access to multiple similar anonymised or synthesised data. We will address this issue as part of the future work by employing a query-auditing method. R1.10 and R1.11 are discussed in Sections 4.1 and 5.2, respectively. The fulfilment of R1.12 was demonstrated in the data selection scenario in Section 6.2, where the number of retrieved dataset for each UC depends on the data handling setup and the dynamics of user consent.

The second part of requirements concerns the audibility and provenance of the system: R2.1, R2.2 and R2.5 are fully satisfied, shown through the automated capturing of provenance trails and the provenance management and analysis capabilities of the audit box. The Secure Repository and Trusted Analysis Environment both send log files in a predefined for-

Table 10  
Evaluation on the requirements

Requirement	full	partial	unsatisfied
R1.1: Input and intermediate data are not persisted by the platform	•		
R1.2: Analysis from a single data source	•		
R1.3: Analysis from multiple data sources		•	
R1.4: All computation on original data is inside the platform	•		
R1.5: Platform limits the period of time for storing analysis environment	•		
R1.6: No data download in original form	•		
R1.7: Anonymised data download	•		
R1.8: Synthetic data download	•		
R1.9: Measures for prohibiting the download of similar data			•
R1.10: Fine-grained consent	•		
R1.11: User access control	•		
R1.12: Automatic usage policy compliance checking	•		
R2.1: Data provenance traced in the whole data lifecycle	•		
R2.2: Automated trails of data access, processing and analysis.	•		
R2.3: Provenance of software components		•	
R2.4: Traces are available for each result		•	
R2.5: Trails do not contain private user data	•		
R2.6: Provenance data streams are accessible for users with auditing privileges	•		
R2.7: No personal data can be acquired by auditors	•		
R2.8: Users' consent for studies is still valid after revoking their consent	•		
R2.9: Minimum meta-information is kept for the past studies	•		

mat, which is then mapped on our provenance model. R2.3 is currently enabled through RDT-Lite, a library capturing provenance in R-Scripts. However, we further extended this functionality to other software environments. Regarding R2.6, the audit box is a separate component of our architecture, where only qualified users have access to. Furthermore, only metadata is captured, no personal data can be acquired by auditors through the audit box, which addresses R2.7. R2.8 and R2.9 are concerned with past studies: Only data with valid consent is analysed through the platform. However, it may occur that users whose data has been analysed revoke their consent after a study has been conducted. In this case, the applicable data will be removed, and only pointers will remain for past studies.

## 7. Related Work

This section discuss related work from three perspectives: privacy-preserving data analysis (Section 7.1, auditability (Section 7.2), and consent management (Section 7.3).

### 7.1. Privacy-preserving Data Analysis

For achieving privacy-preserving data analysis, often two complementary approaches are considered - privacy-preserving data publishing (PPDP) [34], and privacy-preserving computation, such as secure-multiparty computation or homomorphic encryption.

Regarding PPDP, early approaches include e.g., k-anonymity [13], where data is transformed so that individual records can not be distinguished from  $k - 1$  other records, and thus re-identification becomes difficult. Several implementations perform these transformations, e.g., ARX [35].

Differential privacy [36] aims at maximising the accuracy of responses to queries to databases, while minimising the likelihood of being able to identify the records used to answer them – whether a specific instance is included in the dataset should not increase the privacy of the individual represented. Systems implementing differential privacy include e.g., GUPT [37], PINQ [38], Airavat [39].

Recently, synthetisation of data, which creates an artificial dataset that however still preserves global properties of the original data, is emerging as an alternative

1 method [15]. One approach is e.g., the Synthetic Data  
2 Vault [40].

3 While these solutions provide building blocks of  
4 privacy-preserving data analysis, they do not provide  
5 a complete system comprising data collection, consent  
6 management, data transformation and the actual  
7 privacy-preserving data analysis. We will review re-  
8 lated work that addresses these aspects to some de-  
9 gree. A recent report on Trusted Research Environ-  
10 ments [41] elaborates on requirements for such envi-  
11 ronments in the health domain.

12 A system utilising block-chain to trace the usage of  
13 federated data sources in an analysis process is pre-  
14 sented in [42], also providing a mechanism to check  
15 for consent.

16 The DEXHELPP project [43] focuses on safe-haven  
17 like settings, e.g., allowing the user a relatively unre-  
18 stricted access to data inside a specific remote comput-  
19 ing environment, but does not address issues of data  
20 integration, or consent management.

21 DataSHIELD [14] is a system allowing analysis  
22 of data from federated sources, enabling privacy by  
23 restricting the analyst to only executing aggregate  
24 queries. It does not address aspects of data and con-  
25 sent management. In our prototypical implementation,  
26 we utilise DataSHIELD to provide the functionality  
27 of the *Trusted Analysis Environment*. A similar sys-  
28 tem is ViPAR [44], a secure analysis platform for fed-  
29 erated data. The system offers a central server host  
30 where the data is virtually pooled via the Secure Shell  
31 protocol from remote research datasets hosted by re-  
32 search institutions. The data can then be used for an-  
33 alysis through a secure analysis portal, and it is deleted  
34 once the analysis is done. The central pooling com-  
35 ponent is the major difference between DataSHIELD  
36 and ViPAR. DataSHIELD applies the analysis to each  
37 of the remote sources separately and then combine the  
38 analysis results, while ViPAR combines the data itself.

39 BioSHaRE project [45] (Biobank Standardisation  
40 and Harmonisation for Research Excellence in the Eu-  
41 ropean Union) is a collaborative project that, among  
42 other includes DataSHIELD as a privacy-preserving  
43 analysis tool. Others systems used in the project are  
44 DataHSaPER for database federation and harmonisa-  
45 tion and OBiBa as information technology tools.

## 46 7.2. Auditability

47 Traditionally, system logs have been collected and  
48 analysed to conduct system audits and hence, mak-  
49 ing the underlying systems more transparent. How-  
50  
51

1 ever, the analysis of such logs is cumbersome and te-  
2 dious. Moreover, the right to explain as postulated by  
3 the European Union General Data Protection Regula-  
4 tion (GDPR) [1] and similar, other advances for algo-  
5 rithmic accountability demand better suited methods  
6 to provide context and reasoning. Therefore, semantic  
7 technologies seem to be a promising method to capture  
8 and manage data lineage in a structured way to provide  
9 such improved support in auditability questions.

10 Various frameworks and solutions to capture work-  
11 flows, such as Taverna [46] or data models to collect  
12 provenance data such as PROV-DM [47] enable im-  
13 proved analysis and management. Provenance data can  
14 also support the reproducibility of research results by  
15 capturing the experiment context, for example by us-  
16 ing ontologies [48] [49]. Nevertheless, due to rising  
17 complexity of workflows, experiments and technical  
18 setups, the need for automated and specific provenance  
19 collection and analysis suited to specific environments  
20 arises.

21 Language-specific solutions, such as noworkflow  
22 [50] for python scripts, or rdt for R scripts among oth-  
23 ers [51] have been developed and offer richer support  
24 such as visualisation. Prov-O serves as a fundamental  
25 building block for capturing provenance-related data.  
26 Its generic nature however, calls for extension and  
27 adaption to specific contexts to be able to benefit from  
28 contextual information: One example of such an ex-  
29 tension is presented in the domain of smart city appli-  
30 cations by [52]. However, the integration of heteroge-  
31 neous data stemming from different services and sys-  
32 tems is needed to capture the entire data lifecycle, that  
33 is common to industrial and academic settings.

34 ProvStore, developed to be the first public reposi-  
35 tory of provenance documents showcases the need for  
36 such improved data management [53]. Use cases re-  
37 lated to auditing and provenance include the possibil-  
38 ity to prove adherence to legislative regulation (e.g.  
39 GDPR) as suggested by [54] to visualisation of depen-  
40 dencies or debugging capabilities [55]. However, so  
41 far, most tools and approaches are very specific to cer-  
42 tain environments or setups without integrating hetero-  
43 geneous provenance trails, applicable to our use case.

## 44 7.3. Consent Management

45 The European General Data Protection Regulation  
46 (GDPR) [1] defines a set of obligations for controllers  
47 of personal data. Among other requirements, GDPR  
48 requires data controllers to obtain explicit consent for  
49 the processing of personal data from data subjects. Fur-  
50  
51

thermore, organisations must be transparent about their processing of personal data and demonstrate that their systems complies with usage constraints specified by data subjects.

The traditional representation of user consent in the form of human-readable description does not allow for automatic processing. Formal policy languages are designed to unambiguously represent usage policies, which makes it possible to automatically verify whether data processing is covered by consent given by data subjects. Here, we will briefly review the current alternative for (i) usage policy representation, and (ii) GDPR compliance tools.

*Usage Policy Representation.* There are several potential candidates for the formal representation of usage policies. KAoS [56] is a general policy language which adopts a pure ontological approach, whereas Rei [57] and Protune [58] use ontologies to represent concepts, the relationships between these concepts and the evidence needed to prove their truth, and rules to represent policies. PrOnto [8] is another ontology developed based on GDPR to represent a legal knowledge modelling of GDPR conceptual cores, i.e., privacy agents, data types, types of processing operations, rights and obligations. SPECIAL<sup>25</sup>, a European H2020 project, have developed the SPECIAL Policy Language [20] around the recent OWL2 standard with the goal to adequately trade off expressiveness and computational complexity for usage policy checking. Finally, there is the SAVE ontology [7], which integrates and aligns to DPV [6], ODRL [59], and ORCP [60] to provide a fine-grained representation of privacy policies. While in essence any of these vocabularies can be used for usage policy representation in our use case, none of them is community-based and geared towards open standards, unlike DPV [6]. Furthermore, it is important that the representation would allow for automatic compliance check, which will be discussed next.

*GDPR compliance tools.* The Information Commissioner's Office (ICO) in the UK [61], Microsoft [62], and TrustArc<sup>26</sup> have developed compliance checking tools that enable companies to assess the compliance of their applications and business processes by completing a predefined questionnaire. These approaches are manual with regards to the compliance checking, which are not suitable for our use cases, where compliance checking happens during the use of the platform and must be automated.

For automatic usage policy compliance checking, other than the SPECIAL compliance check mechanism [9] that we use in our approach, MIREL<sup>27</sup> and DAPRECO<sup>28</sup> projects extend PrOnto with Legal-RuleML [63] to allow legal reasoning through Defeasible Logic theory. The goal for this approach is to conduct complex legal reasoning and compliance, which is not the goal in our use cases. While do not explicitly designed for GDPR, PrivOnto [64] provides a framework that could facilitate analysis of privacy policies, including usage policy compliance. Another approach taken by RUNE is to utilise SHACL to allow compliance checks for both permissions and prohibitions [11]. However, while the approach is promising, the implementation of the approach is currently not publicly available.

## 8. Conclusion and Future Work

In this paper we discussed the problem of *enabling auditable privacy-preserving data analysis systems*, inspired by the challenges faced by many small and medium-sized organisations in acquiring, storing and analysing personal data due to data protection regulations. We detailed the problem into a number of research questions that we addressed in this paper. We discuss each of the research questions and our approach to address it in the following.

*What are the key characteristics of auditable privacy-preserving data analysis systems?* We described a scenario to exemplify the requirements for auditable privacy-preserving data analysis and motivate a number of functional and non-functional requirements for our approach. These are the results of our iterative process of identification and refining the requirements based on privacy standards as well as use cases and discussions with partners from the WellFort project.

*What are the key elements of an architecture for auditable privacy-preserving data analysis systems?* We proposed a novel approach called WellFort, a semantic-enabled architecture for auditable privacy-preserving data analysis which provides secure storage for users' sensitive data with explicit consent as well as a trusted analysis environment for executing data analytics processes in a privacy-preserving manner.

<sup>25</sup><https://www.specialprivacy.eu/>

<sup>26</sup><http://bit.ly/3tTlPq>

<sup>27</sup><https://www.mirelproject.eu/>

<sup>28</sup><https://www.fnr.lu/projects/data-protection-regulation-compliance/>

1 *How can Semantic Web technologies be used to en-*  
 2 *able auditable privacy-preserving data analysis sys-*  
 3 *tems?* We detailed how we select and utilise a set of  
 4 Semantic Web technologies to support our conceptual  
 5 architecture for auditable privacy-preserving data anal-  
 6 ysis systems. We found that for some tasks, e.g., the  
 7 automatic usage policy compliance checking, the tech-  
 8 nology already has a high-level of maturity. For others,  
 9 some adaptations are still required. Nevertheless, it is  
 10 encouraging to see that research results can be adapted  
 11 for real-world use cases.

12 In our evaluation, we demonstrated the feasibility  
 13 of our approach on four real-world use cases in the  
 14 medical domain. We show that the platform prototype  
 15 is able to track the entire data lifecycle, starting from  
 16 the deposit of data, to producing results in a scientific  
 17 study. Thus, the architecture provides means to inspect  
 18 which data was used, whether the consent was granted,  
 19 or which specific libraries and code were used to pro-  
 20 duce the results. This, in turn, increases reproducibility  
 21 and trust in the results, and can be used as evidence in  
 22 litigation cases to discharge claims.

23 The future work will focus on adoption of the pro-  
 24 posed platform to application domains in which au-  
 25 ditability and privacy are the key concerns. This will  
 26 require changes in the data model used in the reposi-  
 27 tory and development of methods to automate knowl-  
 28 edge graph generation from new types of data.

29 Another branch of future work is to investigate the  
 30 possibility of adopting Semantic Web technologies for  
 31 the *Trusted Analysis Environment*. Recent works in  
 32 this direction has emerged in the Semantic Web com-  
 33 munity, e.g., on privacy-aware query processing [65,  
 34 66]. Additionally, we will focus on improving the qual-  
 35 ity of *Trusted Analysis Environment*. This includes em-  
 36 bedding privacy-preserving record linkage to improve  
 37 the analysis of data originating from different com-  
 38 panies, and extending the repertoire of DataSHIELD  
 39 functionalities.

40 Furthermore, we plan to continue our work on the  
 41 Audit-Box to develop a multi-purpose RDF-based au-  
 42 dit toolbox to support auditability of data analysis sys-  
 43 tems. Such toolbox facilitates capturing and auditing  
 44 provenance trail data from heterogeneous sources, e.g.,  
 45 NoWorkflow [50] for python scripts. To this end, we  
 46 also consider the possibility of linking the toolbox with  
 47 our prior work on RDF data generation from security  
 48 log data [67].

49 **Acknowledgements.** This work was sponsored by  
 50 the Austrian Research Promotion Agency FFG un-  
 51

der grant 871267 (WellFort) and 877389 (OBARIS).  
 The COMET center SBA Research (SBA-K1) is  
 funded within the framework of COMET — Compe-  
 tence Centers for Excellent Technologies by BMVIT,  
 BMDW, and the federal state of Vienna, managed by  
 the FFG. Furthermore, this work was sponsored by the  
 Austrian Science Fund (FWF) and netidee SCIENCE  
 under grant P30437-N31. The authors thank the fun-  
 ders for their generous support.

## References

- [1] European Parliament and Council of European Union (2016), Regulation (EU) 2016/679, *Official Journal of the European Union* **59** (2016), 156. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>.
- [2] F.D. McSherry, Privacy integrated queries: an extensible platform for privacy-preserving data analysis, in: *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, 2009, pp. 19–30.
- [3] ACM, Statement on algorithmic transparency and accountability, *Association for Computing Machinery US Public Policy Council (USACM)* (2017), 1–2. ISBN 1202355129. [https://www.acm.org/binaries/content/assets/public-policy/2017\\_{\\_}usacm\\_{\\_}statement\\_{\\_}algorithms.pdf](https://www.acm.org/binaries/content/assets/public-policy/2017_{_}usacm_{_}statement_{_}algorithms.pdf).
- [4] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik and J. Zhao, *Prov-o: The prov ontology* (2013).
- [5] L.F. Sikos and D. Philp, Provenance-Aware Knowledge Representation: A Survey of Data Models and Contextualized Knowledge Graphs, *Data Science and Engineering* **5(3)** (2020), 293–316. doi:10.1007/s41019-020-00118-0.
- [6] H.J. Pandit and A. Polleres, Data Privacy Vocabulary (DPV), 2021. <https://dpvcg.github.io/dpv/>.
- [7] K. Krasnashchok, M. Mustapha, A. Al Bassit and S. Skhiri, Towards Privacy Policy Conceptual Modeling, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **12400 LNCS** (2020), 429–438. ISBN 9783030625214. doi:10.1007/978-3-030-62522-1\_32.
- [8] M. Palmirani, M. Martoni, A. Rossi, C. Bartolini and L. Robaldo, Pronto: Privacy ontology for legal compliance, in: *Proc. 18th Eur. Conf. Digital Government (ECDG)*, 2018, pp. 142–151.
- [9] P.A. Bonatti, Fast compliance checking in an OWI2 fragment, *IJCAI International Joint Conference on Artificial Intelligence 2018-July* (2018), 1746–1752. ISBN 9780999241127. doi:10.24963/ijcai.2018/241.
- [10] J.D. Fernández, F.J. Ekaputra, P.R. Aryan, A. Azzam and E. Kiesling, Privacy-aware Linked Widgets, in: *Companion Proceedings of The 2019 World Wide Web Conference on - WWW '19*, S. Amer-Yahia, M. Mahdian, A. Goel, G.-J. Houben, K. Lerman, J.J. McAuley, R. Baeza-Yates and L. Zia, eds, ACM Press, New York, New York, USA, 2019, pp. 508–514. ISBN 9781450366755. doi:10.1145/3308560.3317591.

- [11] A. Al Bassit, K. Krasnashchok, S. Skhiri and M. Mustapha, Automated Compliance Checking with SHACL, in: *Submitted to ESWC 2021*, 2021.
- [12] R. Mendes and J.P. Vilela, Privacy-preserving data mining: methods, metrics, and applications, *IEEE Access* **5** (2017), 10562–10582.
- [13] P. Samarati, Protecting respondents identities in microdata release, *IEEE transactions on Knowledge and Data Engineering* **13**(6) (2001), 1010–1027.
- [14] A. Gaye, Y. Marcon, J. Isaeva, P. LaFlamme, A. Turner, E.M. Jones, J. Minion, A.W. Boyd, C.J. Newby, M.-L. Nuotio et al., DataSHIELD: taking the analysis to the data, not the data to the analysis, *International journal of epidemiology* **43**(6) (2014), 1929–1944.
- [15] S.M. Bellovin, P.K. Dutta and N. Reitinger, Privacy and synthetic datasets, *Stan. Tech. L. Rev.* **22** (2019), 1, Publisher: HeinOnline.
- [16] C. Secretariat and N. Office, Reference Model for an Open Archival Information System (OAIS), 2012.
- [17] M. Wilkinson, M. Dumontier, I. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J. Boiten, L. da Silva Santos, P. Bourne, J. Bouwman, A. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C. Evelo, R. Finkers, A. Gonzalez-Beltran, A. Gray, P. Groth, C. Goble, J. Grethe, J. Heringa, P. 't Hoen, R. Hoofit, T. Kuhn, R. Kok, J. Kok, S. Lusher, M. Martone, A. Mons, A. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M. Swertz, M. Thompson, J. Van Der Lei, E. Van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao and B. Mons, The FAIR Guiding Principles for scientific data management and stewardship, *Scientific Data* **3** (2016), 160018. doi:10.1038/sdata.2016.18.
- [18] World Wide Web Consortium, Data Catalog Vocabulary (DCAT) - Version 2, 2020, pp. 1–16. <http://www.w3.org/TR/vocab-dcat/>.
- [19] R. Cyganiak, J. Zhao, A. Keith and M. Hausenblas, Vocabulary of Interlinked Datasets (void), 2011. <http://vocab.deri.ie/void>.
- [20] P.A. Bonatti, S. Kirrane, I. Petrova, L. Sauro and E. Schlehahn, D2.5 - Policy Language V2, Zenodo, 2018. doi:10.5281/zenodo.2545177.
- [21] J.D. Fernández, M. Sabou, S. Kirrane, E. Kiesling, F.J. Ekaputra, A. Azzam and R. Wenning, User consent modeling for ensuring transparency and compliance in smart cities, *Personal and Ubiquitous Computing* (2020), 1–22.
- [22] S.K. (WU), P.B. (CeRICT), J.D.F. (WU), C.G. (CeRICT), L.S. (CeRICT), D.D. (CeRICT), I.P. (CeRICT) and I.S. (CeRICT), D2.8 - Transparency and Compliance Algorithms V2, Zenodo, 2018. doi:10.5281/zenodo.2543622.
- [23] P.A. Bonatti, L. Ioffredo, I.M. Petrova, L. Sauro and I.R. Siahhaan, Real-time reasoning in OWL2 for GDPR compliance, *Artificial Intelligence* **289** (2020), 103389. doi:<https://doi.org/10.1016/j.artint.2020.103389>. <https://www.sciencedirect.com/science/article/pii/S0004370220301399>.
- [24] F.J. Ekaputra, P.R. Aryan, E. Kiesling, C. Fabianek and E. Gringinger, Semantic Containers for Data Mobility: A Seismic Activity Use Case., in: *SEMANTICS Posters&Demos*, 2019.
- [25] S. Miles, P. Groth, S. Munroe and L. Moreau, Prime: A methodology for developing provenance-aware applications, *ACM Transactions on Software Engineering and Methodology (TOSEM)* **20**(3) (2011), 1–42.
- [26] L. Moreau, P. Groth, J. Cheney, T. Lebo and S. Miles, The rationale of PROV, *Journal of Web Semantics* **35** (2015), 235–257.
- [27] A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens and R. Van de Walle, RML: a generic language for integrated RDF mappings of heterogeneous data, in: *Ldow*, 2014.
- [28] D. Calvanese, B. Cogrel, S. Komla-Ebri, R. Kontchakov, D. Lanti, M. Rezk, M. Rodriguez-Muro and G. Xiao, Ontop: Answering SPARQL queries over relational databases, *Semantic Web* **8**(3) (2017), 471–487.
- [29] B. Glimm, I. Horrocks, B. Motik, G. Stoilos and Z. Wang, Hermit: an OWL 2 reasoner, *Journal of Automated Reasoning* **53**(3) (2014), 245–269.
- [30] D. Doiron, Y. Marcon, I. Fortier, P. Burton and V. Ferretti, Software Application Profile: Opal and Mica: open-source software solutions for epidemiological data management, harmonization and dissemination, *International journal of epidemiology* (2017).
- [31] T. Benson and G. Grieve, *Principles of health interoperability: SNOMED CT, HL7 and FHIR*, Springer, 2016.
- [32] FHIR: Fast healthcare interoperability resources. <http://hl7.org/implementation/standards/fhir/>.
- [33] C. Clifton, M. Kantarcioglu, A. Doan, G. Schadow, J. Vaidya, A. Elmagarmid and D. Suciu, Privacy-preserving data integration and sharing, in: *Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, 2004, pp. 19–26.
- [34] B.-C. Chen, D. Kifer, K. LeFevre and A. Machanavajjhala, Privacy-Preserving Data Publishing, *Foundations and Trends® in Databases* **2**(1–2) (2009), 1–167. doi:10.1561/19000000008.
- [35] F. Prasser and F. Kohlmayer, Putting Statistical Disclosure Control into Practice: The ARX Data Anonymization Tool, in: *Medical Data Privacy Handbook*, Springer International Publishing, Cham, 2015, pp. 111–148. ISBN 978-3-319-23632-2 978-3-319-23633-9. doi:10.1007/978-3-319-23633-9\_6.
- [36] C. Dwork, Differential privacy: A survey of results, in: *International conference on theory and applications of models of computation*, Springer, 2008, pp. 1–19.
- [37] P. Mohan, A. Thakurta, E. Shi, D. Song and D. Culler, GUPT: privacy preserving data analysis made easy, in: *Proceedings of the 2012 international conference on Management of Data - SIGMOD '12*, ACM Press, Scottsdale, Arizona, USA, 2012, p. 349. ISBN 978-1-4503-1247-9. doi:10.1145/2213836.2213876.
- [38] F.D. McSherry, Privacy integrated queries: an extensible platform for privacy-preserving data analysis, in: *Proceedings of the 35th SIGMOD international conference on Management of data - SIGMOD '09*, ACM Press, Providence, Rhode Island, USA, 2009, p. 19. ISBN 978-1-60558-551-2. doi:10.1145/1559845.1559850.
- [39] I. Roy, S.T.V. Setty, A. Kilzer, V. Shmatikov and E. Witchel, Airavat: Security and Privacy for MapReduce, in: *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, NSDI'10, USENIX Association, USA, 2010, event-place: San Jose, California.
- [40] N. Patki, R. Wedge and K. Veeramachaneni, The Synthetic Data Vault, in: *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, IEEE, Montreal,

- QC, Canada, 2016, pp. 399–410. ISBN 978-1-5090-5206-6. doi:10.1109/DSAA.2016.49.
- [41] UK Health Data Alliance, Trusted Research Environments – A strategy to build public trust and meet changing health data science needs, 2020. [https://ukhealthdata.org/wp-content/uploads/2020/07/200723-Alliance-Board-Paper-E\\_TRE-Green-Paper.pdf](https://ukhealthdata.org/wp-content/uploads/2020/07/200723-Alliance-Board-Paper-E_TRE-Green-Paper.pdf).
- [42] M. Koscina, D. Manset, C. Negri and O. Perez, Enabling trust in healthcare data exchange with a federated blockchain-based architecture, in: *IEEE/WIC/ACM International Conference on Web Intelligence on - WI '19 Companion*, ACM Press, Thessaloniki, Greece, 2019, pp. 231–237. ISBN 978-1-4503-6988-6. doi:10.1145/3358695.3360897.
- [43] N. Popper, F. Endel, R. Mayer, M. Bicher and B. Glock, Planning Future Health: Developing Big Data and System Modelling Pipelines for Health System Research, *SNE Simulation Notes Europe* **27**(4) (2017), 203–208. doi:10.11128/sne.27.tn.10396.
- [44] K.W. Carter, K. Carter, R. Francis, M. Bresnahan, M. Gissler, T. Grønberg, R. Gross, N. Gunnes, G. Hammond, M. Hornig, C. Hultman, J. Huttunen, A. Langridge, H. Leonard, S. Newman, E. Parner, G. Petersson, A. Reichenberg, S. Sandin, D. Schendel, L. Schalkwyk, A. Sourander, C. Steadman, C. Stoltenberg, A. Suominen, P. Surén, E. Susser, A. Sylvester Vethanayagam and Z. Yusof, ViPAR: a software platform for the Virtual Pooling and Analysis of Research Data, *International Journal of Epidemiology* **45**(2) (2016), 408–416. doi:10.1093/ije/dyv193.
- [45] D. Doiron, P. Burton, Y. Marcon, A. Gaye, B.H.R. Wolfenbuttel, M. Perola, R.P. Stolk, L. Foco, C. Minelli, M. Waldenberger, R. Holle, K. Kvaløy, H.L. Hillege, A.-M. Tassé, V. Ferretti and I. Fortier, Data harmonization and federated analysis of population-based studies: the BioSHaRE project, *Emerging Themes in Epidemiology* **10**(1) (2013), 12. doi:10.1186/1742-7622-10-12.
- [46] T. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, M. Greenwood, T. Carver, K. Glover, M.R. Pocock, A. Wipat et al., Taverna: a tool for the composition and enactment of bioinformatics workflows, *Bioinformatics* **20**(17) (2004), 3045–3054.
- [47] K. Belhajjame, R. B'Far, J. Cheney, S. Coppens, S. Cresswell, Y. Gil, P. Groth, G. Klyne, T. Lebo, J. McCusker et al., ProvdM: The prov data model, *W3C Recommendation* (2013).
- [48] T. Miksa and A. Rauber, Using ontologies for verification and validation of workflow-based experiments, *Web Semantics: Science, Services and Agents on the World Wide Web* **43**(Supplement C) (2017), 25–45. doi:<https://doi.org/10.1016/j.websem.2017.01.002>.
- [49] B. Goesswein, T. Miksa, A. Rauber and W. Wagner, Data Identification and Process Monitoring for Reproducible Earth Observation Research, in: *IEEE eScience 2019*, IEEE, 2019, pp. 28–38. doi:10.1109/eScience.2019.00011.
- [50] J.F. Pimentel, L. Murta, V. Braganholo and J. Freire, noWorkflow: a tool for collecting, analyzing, and managing provenance from python scripts, *Proceedings of the VLDB Endowment* **10**(12) (2017).
- [51] A.M. Ellison, E.R. Boose, B.S. Lerner, E. Fong and M. Seltzer, The End-to-End Provenance Project, *Patterns* **1**(2) (2020), 100016.
- [52] H. Packer, D. Diochnos, M. Rovatsos, Y. Gal and L. Moreau, Semantics and Provenance for Accountable Smart City Applications, *Semantic Web–Interoperability, Usability* (2014).
- [53] T.D. Huynh and L. Moreau, ProvStore: a public provenance repository, in: *International Provenance and Annotation Workshop*, Springer, 2014, pp. 275–277.
- [54] B.E. Ujcich, A. Bates and W.H. Sanders, A provenance model for the European union general data protection regulation, in: *International Provenance and Annotation Workshop*, Springer, 2018, pp. 45–57.
- [55] J.F. Pimentel, J. Freire, L. Murta and V. Braganholo, A survey on collecting, managing, and analyzing provenance from scripts, *ACM Computing Surveys (CSUR)* **52**(3) (2019), 1–38.
- [56] A. Uszok, J. Bradshaw, R. Jeffers, N. Suri, P. Hayes, M. Breedy, L. Bunch, M. Johnson, S. Kulkarni and J. Lott, KAoS policy and domain services: Toward a description-logic approach to policy representation, deconfliction, and enforcement, *Proceedings - POLICY 2003: IEEE 4th International Workshop on Policies for Distributed Systems and Networks* (2003), 93–96. ISBN 0769519334. doi:10.1109/POLICY.2003.1206963.
- [57] L. Kagal, T. Finin and A. Joshi, A policy language for a pervasive computing environment, *Proceedings - POLICY 2003: IEEE 4th International Workshop on Policies for Distributed Systems and Networks* (2003), 63–74. ISBN 0769519334. doi:10.1109/POLICY.2003.1206958.
- [58] P. Bonatti, J.L. De Coi, D. Olmedilla and L. Sauro, A rule-based trust negotiation system, *IEEE Transactions on Knowledge and Data Engineering* **22**(11) (2010), 1507–1520. doi:10.1109/TKDE.2010.83.
- [59] R. Iannella and S. Villata, ODRL Information Model 2.2, *W3C recommendation* (2018). <https://www.w3.org/TR/odrl-model/>.
- [60] S. Kirrane, M. De Vos and J. Padget, ODRL Regulatory Compliance Profile 0.2, Technical Report, Vienna University of Economics and Business, 2020.
- [61] Information Commissioner's Office, Getting ready for the GDPR, 2020. <https://ico.org.uk/for-organisations/data-protection-self-assessment/>.
- [62] Microsoft Trust Center, Detailed GDPR assessment, 2017. <http://aka.ms/gdprdetailedassessment>.
- [63] T. Athan, G. Governatori, M. Palmirani, A. Paschke and A. Wyner, LegalRuleML: Design principles and foundations, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **9203** (2015), 151–188. ISBN 9783319217673. doi:10.1007/978-3-319-21768-0\_6.
- [64] A. Oltramari, D. Piraviperumal, F. Schaub, S. Wilson, S. Cherivirala, T.B. Norton, N.C. Russell, P. Story, J. Reidenberg and N. Sadeh, PrivOnto: A semantic framework for the analysis of privacy policies, *Semantic Web* **9**(2) (2018), 185–203.
- [65] K.M. Endris, Z. Almhithawi, I. Lytra, M.-E. Vidal and S. Auer, BOUNCER: privacy-aware query processing over federations of RDF datasets, in: *International Conference on Database and Expert Systems Applications*, Springer, 2018, pp. 69–84.
- [66] R. Taelman, S. Steyskal and S. Kirrane, Towards Querying in Decentralized Environments with Privacy-Preserving Aggregation, *arXiv preprint arXiv:2008.06265* (2020).
- [67] A. Ekelhart, F.J. Ekaputra and E. Kiesling, The SLOGERT Framework for automated Log Knowledge Graph Construction (2020).