

Article

DPCat: Specification for an Interoperable and Machine-Readable Data Processing Catalogue Based on GDPR

Paul Ryan ^{1,2,3,*} , Rob Brennan ^{2,3}  and Harshvardhan J. Pandit ^{3,4,*} ¹ Uniphar PLC, D24 VO6K Dublin, Ireland² School of Computing, Dublin City University, D09 PX21 Dublin, Ireland; rob.brennan@dcu.ie³ ADAPT SFI Research Centre, D02 FX65 Dublin, Ireland⁴ School of Computer Science & Statistics, Trinity College Dublin, D02 PN40 Dublin, Ireland

* Correspondence: paul.ryan76@mail.dcu.ie (P.R.); pandith@tcd.ie (H.J.P.)

Abstract: The GDPR requires Data Controllers and Data Protection Officers (DPO) to maintain a Register of Processing Activities (ROPA) as part of overseeing the organisation's compliance processes. The ROPA must include information from heterogeneous sources such as (internal) departments with varying IT systems and (external) data processors. Current practices use spreadsheets or proprietary systems that lack machine-readability and interoperability, presenting barriers to automation. We propose the Data Processing Catalogue (DPCat) for the representation, collection and transfer of ROPA information, as catalogues in a machine-readable and interoperable manner. DPCat is based on the Data Catalog Vocabulary (DCAT) and its extension DCAT Application Profile for data portals in Europe (DCAT-AP), and the Data Privacy Vocabulary (DPV). It represents a comprehensive semantic model developed from GDPR's Article and an analysis of the 17 ROPA templates from EU Data Protection Authorities (DPA). To demonstrate the practicality and feasibility of DPCat, we present the European Data Protection Supervisor's (EDPS) ROPA documents using DPCat, verify them with SHACL to ensure the correctness of information based on legal and contextual requirements, and produce reports and ROPA documents based on DPA templates using SPARQL. DPCat supports a data governance process for data processing compliance to harmonise inputs from heterogeneous sources to produce dynamic documentation that can accommodate differences in regulatory approaches across DPAs and ease investigative burdens toward efficient enforcement.

Keywords: GDPR; data governance; semantic-web

Citation: Ryan, P.; Brennan, R.; Pandit, H.J. DPCat: Specification for an Interoperable and Machine-Readable Data Processing Catalogue Based on GDPR. *Information* **2022**, *13*, 244. <https://doi.org/10.3390/info13050244>

Academic Editors: Riccardo Albertoni and Peter Winstanley

Received: 11 April 2022

Accepted: 7 May 2022

Published: 10 May 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Many organisations are complex entities that perform heterogeneous processing on diverse personal data, often organised using multiple organisational units or outsourced processing partners and sometimes under the jurisdiction of multiple Data Protection Authorities (DPAs). Under the EU's General Data Protection Regulation (GDPR), organisations that act as a 'Data Controller' are obliged to create and maintain a "Register of Processing Activities (ROPA)" as a comprehensive record of personal data processing activities carried out under their responsibility (GDPR Art.30). The ROPA, as described in the GDPR, is a temporal snapshot of the organisation's practices and is the point of initiating communication or investigation regarding compliance such as with a DPA. It is thus an important part of the organisation's processes related to ensuring and documenting its compliance.

In practice, organisations struggle to keep accurate and up-to-date ROPAs [1]. They often fail to integrate the maintenance and management of the Register of Processing Activities into their day-to-day operations [1]. This can result in a breakdown in the GDPR accountability principle (GDPR Article 5.2) as there is a lack of clarity as to the who, how, and when the ROPA is updated. To assist organisations with their ROPA-related duties, DPAs have provided guidance and templates that intend to ease the task of understanding

requirements and harmonise the documentation through commonly used formats and environments such as spreadsheets [2,3]. In providing these templates, DPAs indicate what can be considered ‘good practice’ regarding what information should be documented within a ROPA. However, despite being based on a common legal obligation (GDPR Art.30), there is variance in the templates provided by DPAs where additional fields (not in the GDPR) are also encouraged to be documented [2]. An organisation operating in multiple jurisdictions is thus tasked with consolidating differing requirements from each DPA as either a distinct set of ROPA documents or a single combined one.

Furthermore, the exercise of gathering the information necessary to create a ROPA is not a one-off activity [4] as there may be several data sources both internally (e.g., departments) [5] and externally (e.g., data processors) [5,6]. Therefore, ROPA creation requires communication between these distinct units to collate information pooled from ‘heterogeneous sources’ into a singular location to produce a ROPA. This necessitates some form of information management process for the tasks associated with documents such as reading or viewing, writing all or parts of it, exchanging them between relevant stakeholders, and ensuring their correctness and availability (e.g., backups or version control).

To address such requirements, the market vendors offer dedicated solutions for ROPA management, often as part of a larger suite of the GDPR compliance tools [7]. This follows the increasing trend of organisations adopting regulatory technology (RegTech) [8,9] to assist with legal compliance and requirements. The utilisation of a ROPA is poised to be an important and key feature given its importance in the GDPR compliance processes.

However, these RegTech solutions are primarily centralised and proprietary, and they emphasise custom processes that cannot be utilised outside vendor-defined use cases. In particular, the information being exchanged between internal and external stakeholders has been poorly researched in academia and commercial offerings (see Section 2) despite the need for shared business and regulatory taxonomies for facilitating semantic interoperability [10] between stakeholders to identify feasible and compliant software solutions for data protection and privacy regulations [11,12].

There is a lack of ROPA-related explorations in academic research, with existing efforts limited to early-stage work involving enterprise architecture models [13] or data [14]. For larger projects that have focused on GDPR compliance with explicit requirements regarding non-proprietary technologies and focusing on interoperability (e.g., semantic web), there is a distinct absence of research addressing ROPA-related tasks despite overlapping with the same information requirements. In terms of ongoing work, the ONTOROPA project [12] proposes building a semantics-based ROPA with blockchain-based trust guarantees.

We propose an approach to solving these challenges, whereby we identify what data is required to complete the ROPA, who the ROPA stakeholders are, how they utilise the ROPA, and what information flows requiring interoperability and machine-readability of the ROPA are required. To address the identified challenges and their solutions, we present our work based on the following research objectives:

- RO1 Identify information and information flows relevant for a ROPA in terms of stakeholders based on the GDPR and EU DPAs guidelines and templates;
- RO2 Develop a machine-readable specification for representing and exchanging ROPA relevant information in an interoperable manner;
- RO3 Specify a mechanism for using developed machine-readable formats for aggregation, querying, validation, and exporting of information based on identified ROPA-related information flows.

Our previous work on this topic consisted of creating a semantic model of a ROPA [5]. In this, we evaluated the GDPR and six DPA templates and guidelines to identify a set of concepts required for the representation of ROPA-related information and proposed its formulation as a ‘*common semantic model*’ for representing commonality across the EU. We utilised the data privacy vocabulary (DPV) [15], developed by the W3C Data Privacy Vocabularies and Controls Community Group (DPVCG), as a vocabulary for representing identified concepts (Note: H. J. Pandit chairs DPVCG and is the editor of DPV

<https://www.w3.org/community/dpvcg/> as of 5 May 2022). We found and reported missing concepts to DPVCG, which subsequently extended the DPV with our contribution. We further developed our common semantic model into a proposal for establishing a 'Data Processing Catalogue (DPCat) [16] that utilises the Data Catalog Vocabulary (DCAT) [17] and its extension, the DCAT Application profile for data portals in Europe (DCAT-AP) [18], to represent the ROPA-related information in the form of 'datasets' and 'catalogues' that could be maintained, used, and shared consistently.

This article expands on our prior work to provide a more complete and feasible solution for establishing a common machine-readable and interoperable mechanism for a common representation of a ROPA. We extended the common semantic model to incorporate ROPA templates from all EU DPAs (17 of 31 DPAs have published templates) and updated the DPCat specification and the DPV to support representing this information. To demonstrate its practical application and usefulness, we applied the DPCat specification to ROPA documents published by the European Data Protection Supervisor (EDPS) for each identified use case (see Section 6). Finally, we go beyond state of the art by demonstrating the potential of our solution in realising the EU's 'Data Spaces' vision [19] by creating 'compliance-related specifications' that support representation (RDF), querying (SPARQL), validation (SHACL), and exchange (DCAT + DPV) of information.

The principal contributions of this paper are summarised as follows:

1. use cases exploring ROPA data governance and stakeholders (RO1);
2. A Common Semantic Model for ROPA (CSM-ROPA) representing information requirements from EU DPAs (RO2);
3. Data Processing Catalogue (DPCat) specifications for representing and exchanging ROPA-related information and provenance (RO2);
4. Demonstration of representation, querying, validation, and exchange of ROPA-related information using DPCat and semantic web technologies (RO3);
5. Discussion on the practicality and application of DPCat as a 'common mechanism' for exchanging compliance information.

All associated data in documents, analysis, code, and executable artefacts are available under an open and permissive licence at <https://w3id.org/dpcat/repo>.

The remainder of the paper is structured as follows: Section 2 discusses the state of the art and related work, and Section 3 describes the development of the Common Semantic Model for ROPA (CSM-ROPA) development. In Section 4, we discuss ROPA information flows and data governance requirements for ROPA. Section 5 describes the DPCat data processing catalogue to enable ROPA information sharing, aggregation, and querying for ROPA stakeholder interoperability. Section 6 provides an application use case to demonstrate the practicality and feasibility of DPCat. The remainder of the paper discusses the impact of our approach on real-world use cases based on enabling better automation and tooling for regulatory compliance and critically for authorities to ease investigative burdens towards effective enforcement, and we provide our conclusions and recommendations for future work.

2. State of the Art and Related Work

This section presents an overview of relevant work specifically regarding modelling, creation, and maintenance of ROPAs, tangentially regarding the GDPR-related machine-readable and interoperable information management and compliance processes.

2.1. Information Management Solutions for ROPA

The International Association of Privacy Professionals (IAPP), the largest global community for privacy and data protection professionals, reported that 65% of organisations relied on spreadsheets or completely manual solutions to maintain their ROPAs [20]. Another IAPP report found 169 vendors supplying ROPA-related information management services and software in 2020 [21]. This practice can be seen as reflective of the prevalence of maintaining compliance-related information in 'manual tools' such as spreadsheets

without using technological solutions that operate on them. Instead, the exception to these is proprietary solutions offered by vendors such as One Trust, Data Grail, and Transcend. These privacy vendors have seen the importance in offering tools that enable integration with their solutions; however, such integrations link into their locked ecosystem without the ability for organisations to control their data or to move it to an alternate technology provider. DPAs, in reaction to existing common practices, also provide spreadsheet templates that encourage the use of manual or vendor-specific solutions.

One of the main failings of organisations regarding ROPAs is devolving its maintenance to their data protection officer (DPO) and not having active involvement in its upkeep [1]. Best practice for a ROPA suggests complete involvement of stakeholders in the ongoing maintenance of the ROPAs [4] to provide the DPO with an accurate and up-to-date view of personal data processing carried out by that organisation [22,23]. This means that we need processes that assist DPOs and enable the engagement of stakeholders in the upkeep and review of ROPAs [14]. From this, we conclude there is demand for automation through technological solutions and that the market is responding to such needs with commercial offerings.

Furthermore, the Future of Privacy Forum (FPF) [24] reports that privacy and data protection technology providers also face a significant obstacle regarding lack of a common terminology [25]. Therefore, the demand for automation and technologies should be accompanied by requirements for common mechanisms and terminologies that can operate across processes and stakeholders and establish standardised mechanisms within the ecosystem. We can look towards the manufacturing and finance sectors, where the harmonisation achieved through commonality and standardisation has improved regulation and value chains [10].

In response, we identify information and information flows relevant for ROPA governance based on the GDPR and DPA guidelines and templates (*RO1*) and provide a common terminology for representing ROPA to overcome the lack of a common terminology (*RO2*).

2.2. GDPR Compliance Approaches Using Machine-Readable Metadata

In contrast to market reaction, the ROPA as a topic has received little attention within the academic and research communities despite evidence of a broad category of approaches addressing GDPR compliance. Rozenhal et al. [26] proposed ‘Enterprise Architecture’ as an ideal source for representing processing activities and technologies in an organisation. This is supported by Burmeister et al., who also investigated how enterprise architecture can provide a DPO with insights on organisational data processing activities concerning GDPR compliance [27]. Enterprise architecture has further explored sources of ROPA-related information in the next subsection.

The ONTOROPA project [12] proposes using semantic web ontologies and knowledge graphs for representing ROPA-related information, and using a blockchain to certify its integrity and authenticity. To address such challenges, research efforts at producing common terminology using semantic web vocabularies and ontologies have been developed [15,28]. Other approaches utilise such vocabularies to construct ‘legal knowledge bases’ and utilise them for compliance evaluation and monitoring which can help harmonise and facilitate a joint approach between legal departments and other stakeholders to identify feasible and compliant solutions around data protection and privacy regulations [11].

Several semantic-based projects provide ontologies, vocabularies, and policy languages that can be utilised to represent GDPR concepts. These mainly focus on terms referenced in the GDPR rights and obligations. Most projects focus on legal compliance evaluation rather than deployment and interoperability. They do not consider the critical aspect of how the information required is maintained or generated within/by organisations and the stakeholders and information flows involved in this process. Some notable outputs for this are: BPR4GDPR’s IMO [29], GDPRov [30], GConsent [31], DPV [15], GDPRtEXT [28], SPE-CIAL’s ontologies [32], and PrOnto [33]. A recent survey (2022) provides further overview

and details regarding ontologies and policy languages for representing information flows based on the GDPR [34].

BPR4GDPR (business process re-engineering and functional toolkit for GDPR compliance) [29,35] is a relevant ontology-based compliance methodology used to dictate and evaluate processes. It is based on advanced process mining from event logs of IT systems to discover, monitor, and improve processes without pre-modelling the processes before mining. BPR4GDPR thus creates a novel process-monitoring architecture with constraints for conformance checking and automated evolution of processes to satisfy the rules. It will take significant trials and development before the widespread deployment of advanced techniques like this in conventional organisations.

GDPRov [30], GConsent [31], and SPECIAL [32] provide ontologies for expressing GDPR-related concepts but do not incorporate ROPA requirements. GDPRtEXT [28] provides a vocabulary of GDPR concepts, of which some relate to ROPA (GDPR Art.30). PrOnto [33] provides concepts regarding data types, documents, agents and roles, purposes, legal bases (and more), but it is not available for reuse. DPV [15] also provides concepts regarding data categories, purposes, legal bases (and more), represents a community consensus, and is available for reuse.

In efforts that are focused on integration with software development methodologies and infrastructures, two notable efforts are TIRA—an OpenAPI extension for REST architectures [36]—and TILT—which provides concept integration within code for practical privacy engineering [37]. The German DPAs effort on identifying fields and requirements through a ‘Standard Data Protection Model (SDM)’ is also relevant here given its focus on information systems [38].

These existing efforts, specifically DCAT-AP) and DPV, provide the basis on which we develop a machine-readable specification to represent and facilitate the exchange of ROPA-relevant information in an interoperable manner (RO2) and to utilise it for aggregation, querying, validation, and exporting information based on identified ROPA-related information flows using semantic web technologies (RO3).

3. A Common Semantic Model for ROPAs (CSM-ROPA)

Despite a ROPA being based only on requirements established by the GDPR Art. 30, our prior work found variance amongst ROPAs templates provided by six DPAs in terms of what information needed to be documented. The additional fields were related to what the DPAs considered best practices to assist organisations in collecting and representing information from their various business processes. We harmonised the requirements from different templates to construct a ‘common semantic model’ for a ROPA (CSM-ROPA) to enable the representation of all DPA-specified ROPA information [2]. We then represented these information requirements through concepts from the data privacy vocabulary (DPV) [15] to provide an interoperable machine-readable vocabulary that can act as a mediation mechanism between stakeholders and tools operating on a ROPA and associated compliance processes. In this section, we present results from our extended work where we analysed and incorporated ROPA templates from all EU DPAs to create a single (and truly) ‘common semantic model’ for a ROPA and represented it using DPV to provide a consistent and interoperable specification for representing a ROPA and its relevant information.

3.1. Analysis of DPA ROPA Templates

The GDPR has 31 DPAs representing nations and member states from the EU and the EFTA EEA (Note: based on EDPB membership this consists of 31 DPAs from 27 EU member states, the EDPS, and 3 additional members comprising the EFTA EEA states; the German regional DPAs were considered part of the national German DPA). Each DPA provides guidance regarding the ROPA based on its basis in the GDPR Art.30, and some DPAs also provide templates to assist organisations with maintaining their ROPA documents. In our prior work analysing six DPA templates [2], we found that the DPA ROPA templates go beyond the GDPR Art.30 requirements, are not consistent with other DPA templates, and

represent a challenge in producing a ‘collective understanding’ of what information is required for maintaining a ROPA.

In this work, we expanded the analysis to all 31 DPAs, and found 17 DPAs provided ROPA templates varying in language and content (Note: Of 17 DPA templates, 5 used English. We used Google Translate to convert the rest to English and manually ensured consistency in translation between templates regarding terms used). On these 17 templates, we performed term extraction, semantic analysis, term frequency enumeration, de-duplication, and antonym/homonym identification. Templates with minimal information restricted their contents for conforming with the GDPR Art.30. Some templates, such as those provided by Belgian and Greek DPAs, were extensive in fields beyond what the GDPR or other DPAs suggested.

The exercise, carried out over 2020–2022, yielded 47 unique concepts representing information to be recorded in a ROPA. Of these, 18 concepts were related to the requirements defined in the GDPR Art.30, and the rest (29 concepts) were either supplementary to these or added by DPAs. An overview of the exercise is presented in Appendix A which shows the identified concepts and their relevance to each DPA template analysed. (Note: We could not discern a source or basis in law (EU or national) for concepts added by DPAs).

3.2. Developing a Semantic Model for ROPA Using DPV

In our previous work [2,3], we utilised DPV to represent terms identified from ROPA templates as machine-readable and interoperable concepts for use in information management and compliance-based approaches. Through this, we proposed a ‘Common Semantic Model for ROPA’ (CSM-ROPA). In this section, we describe our work in expanding the CSM-ROPA to cover additional requirements and concepts identified from the analysis of DPA ROPA templates and incorporate the updates made to DPV.

The DPV provides a semantic vocabulary consisting of hierarchical taxonomies of concepts relevant to the GDPR such as personal data, purposes, processing operations, technical and organisational measures, legal bases, and entities. We chose DPV as it provides the most comprehensive vocabulary for our purposes, is open and accessible, has ongoing development and mechanisms to submit contributions, and is familiar to the authors.

The process of representing identified concepts using DPV used the methodology [39], where for each term, we constructed a competency question to identify relevant DPV concepts. For example, the term *purposes* was phrased as the question: *What are the Purposes of processing?* We then identified whether the DPV contained the (semantically) exact concept—which we call an ‘*exact match*’, failing which we looked for the closest relevant term(s) which could be used as a substitute—called a ‘*partial match*’, and if any existing term could not represent the term—we considered it a ‘*new term*’ to be proposed to the DPVCG for inclusion in the DPV. Of the 47 unique concepts found through ROPA templates analysis, we found 44 exact matches, one partial match, and two new terms proposed and added to the DPV. Appendices A and B provide an overview of this outcome.

The output of this was the CSM-ROPA consisting of 47 concepts covering information requirements from the GDPR and DPA templates for representing a ROPA. CSM-ROPA, through the use of DPV concepts, provides the ability to express a ROPA as a machine-readable and interoperable ‘graph’ that can be utilised in technological solutions for automating processes associated with ROPA and the GDPR compliance. The CSM-ROPA data and analysis are available online at <https://w3id.org/dpcat/csm-ropa>.

4. Information and Data Governance for ROPA

The CSM-ROPA, described in the previous section, enables the representation of a ROPA in a machine-readable and interoperable manner and covers information requirements from the GDPR and DPA ROPA templates. However, a ROPA is not a single document in practice but is a related set of evolving information that must be periodically collected and maintained. The information required for maintaining a ROPA thus may

have one or more internal sources such as a department, unit, or assigned person where such ‘organisational units’ provide data about their respective processes and activities. A ROPA may also have one or more external sources—such as processors, contractors, vendors—where such ‘external entities’ provide the information required for establishing records of agreed activities and assurance of compliance obligations.

The ROPA provides the DPO with an important overview of the organisation’s practices [22] and is part of the DPO’s obligations regarding compliance (GDPR Art.39) [40]. This requires communication between internal stakeholders such as units or departments and external stakeholders such as DPAs, auditors, and certification bodies to collate necessary information for ROPA governance.

We present five use cases that explore the key stakeholders and their roles regarding the ‘heterogeneous sources’ in ROPA-related data governance. This follows the methodology from prior work [6] regarding identifying stakeholders and information flows related to GDPR compliance and establishing the utility of developing machine-readability and semantic interoperability mechanisms based on it. (Note: In this, we relied on P. Ryan’s experience as an active DPO for over 30 legal entities).

In our analysis, we considered the DPO as the nominated entity with responsibility within an organisation to oversee the ROPA-related processes as per the obligations from the GDPR (Art. 39). From this perspective, we explore possible combinations based on the existence or involvement of specific stakeholders and their effect on the DPO’s duties to collect and maintain ROPA-related information. We also considered a data controller as the primary type of organisation despite a data processor being required to maintain a ROPA and involve a DPO as a stakeholder. The data controller’s use cases are more complex than a data processor’s, and a solution satisfying a controller’s ROPA requirements can be trivially modified for use by a processor.

This exercise concludes with an argument for the expression of ROPA-related information in a machine-readable and interoperable format. Section 6 then presents DPCat as our solution to communicate or exchange ROPA-relevant information between stakeholders and assist in the automation of compliance processes.

4.1. Use Case U1: Data Controller

This use case, illustrated in Figure 1, represents a single data controller that maintains a ROPA (GDPR Art.30) for which it identifies and documents relevant processing activities conducted under its responsibility. In addition, as best practice, the controller must assess guidelines and templates provided by relevant DPA(s) and adapt its documentation processes accordingly to meet any additional suggestions or requirements. A ROPA produced by a controller is utilised by its DPO as part of the responsibility to oversee compliance. The ROPA may also be accessed by a DPA or an auditor (e.g., a certification body) as part of their correspondence with the controller or an investigation or auditing process.

The information flows between these stakeholders can involve: (i) a ROPA that conforms to the GDPR Art.30 requirements; (ii) a ROPA that conforms to a DPAs guidelines and templates; (iii) provenance, e.g., ROPA issuer, timestamps, contact details; and (iv) a selective part of ROPA, e.g., temporal period, specific processing activities.

4.2. Use Case U2 Data Controller with Internal Organisational Units

The second use case, presented in Figure 2, expands U1 with internal information flows through four ‘organisational units’ or departments: marketing, human resources, IT services, and web services, where relevant data for generating a ROPA must be collated into a common location [4]. U2 also involves potential follow-ups with each unit regarding maintenance of records per department, and establishing ‘points of contact’ and ‘responsible entity’. The external information flows, i.e., DPAs and auditors, stay the same as internal units are not separate legal entities subject to direct investigation.

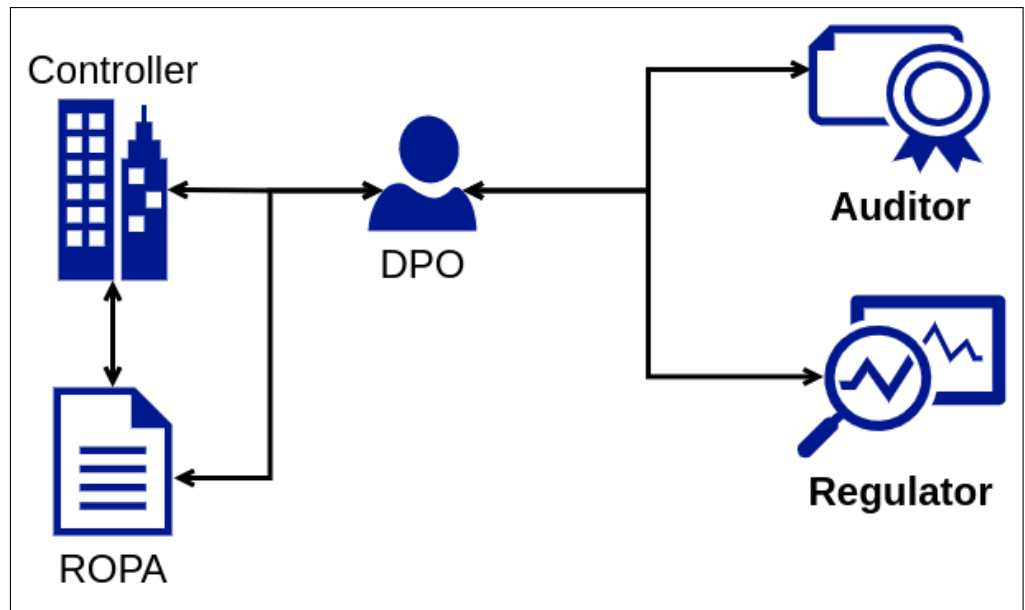


Figure 1. Basic generation of legal requirement ROPA.

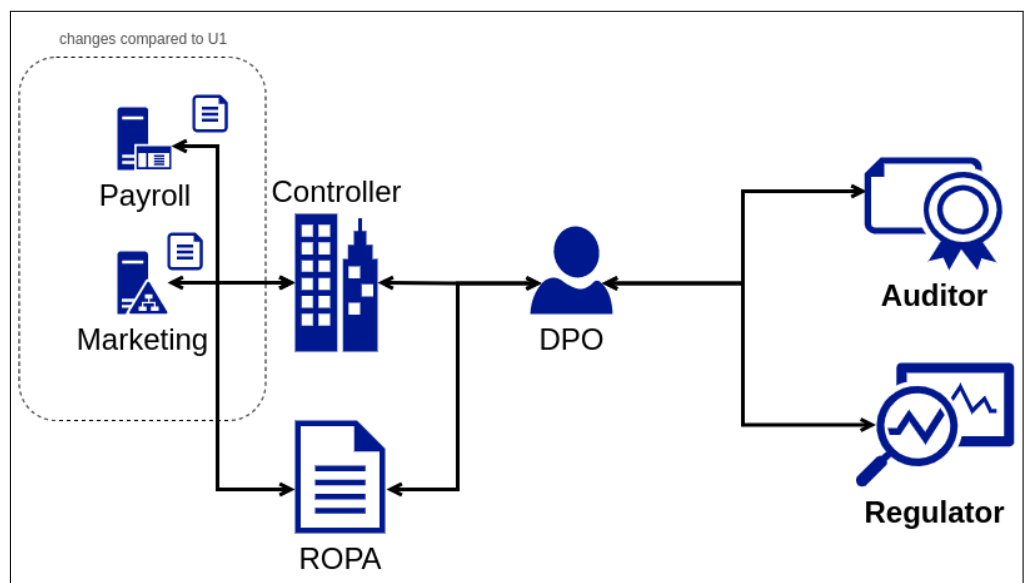


Figure 2. Organisational units updating and maintaining ROPA.

The key information flows between these stakeholders in addition to those in U1 may involve: (i) complete or partial ROPA information for each internal organisational unit; (ii) provenance, e.g., department as issuer, point of contact or responsible entity, timestamps, contact details; (iii) collation of department information into a common ROPA for external stakeholders; and (iv) a selective part of ROPA, e.g., specific department.

4.3. Use Case U3: Data Controller with Data Processors

The third use case, illustrated in Figure 3, has additional information flows where the controller and its DPO collect relevant information from appointed processors. In cases where a processor is common to all departments or is managed at the organisational level, U3 is an extension of U1. Where organisational units utilise specific external vendors (i.e., data processors), U3 is an extension of U2. In this, we consider the practical situations where data governance is often managed by internal units despite the GDPR associating data processors directly with a data controller.

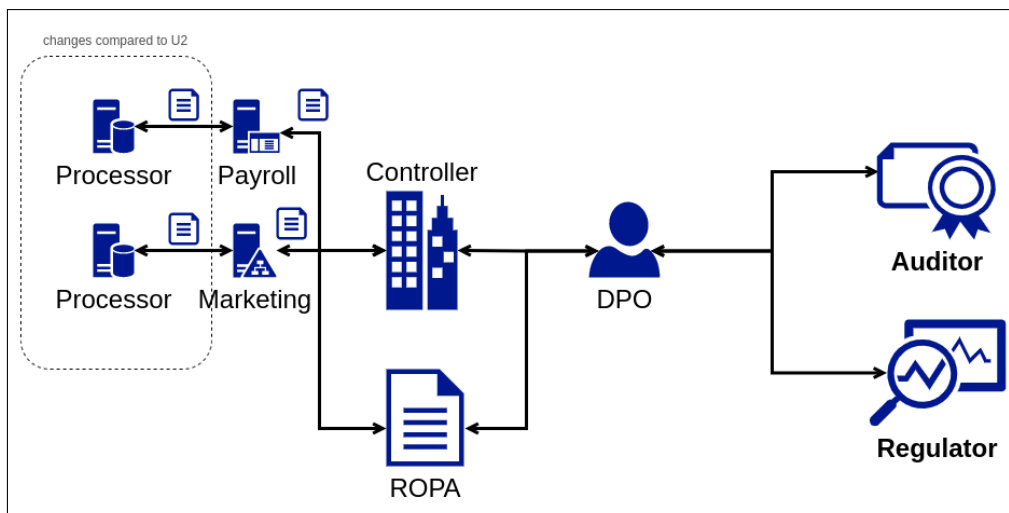


Figure 3. A data controller with organisational units and data processors.

The key information flows between these stakeholders, in addition to U1 and U2, involve: (i) ROPA information from appointed processors; (ii) Provenance, e.g., sources, timestamps, contact details; (iii) Collation of information from heterogeneous sources into a common ROPA; and (iv) A selective part of ROPA, e.g., specific processor.

4.4. Use Case U4: Data Controller in a Joint Controllers Relationship

The fourth use case, illustrated in Figure 4, expands U3 with the Data Controller being in a joint controllers relationship with two or more controllers sharing the responsibility of processing as per GDPR Art.26. Similar to the possibility of associating processors with organisational units in U3, joint controllers can also similarly be associated with units for situations where the processing is limited to a unit’s activities. In U4, the controller and its DPO have additional information flows regarding collecting relevant information from other (joint) controllers and any potential follow-ups.

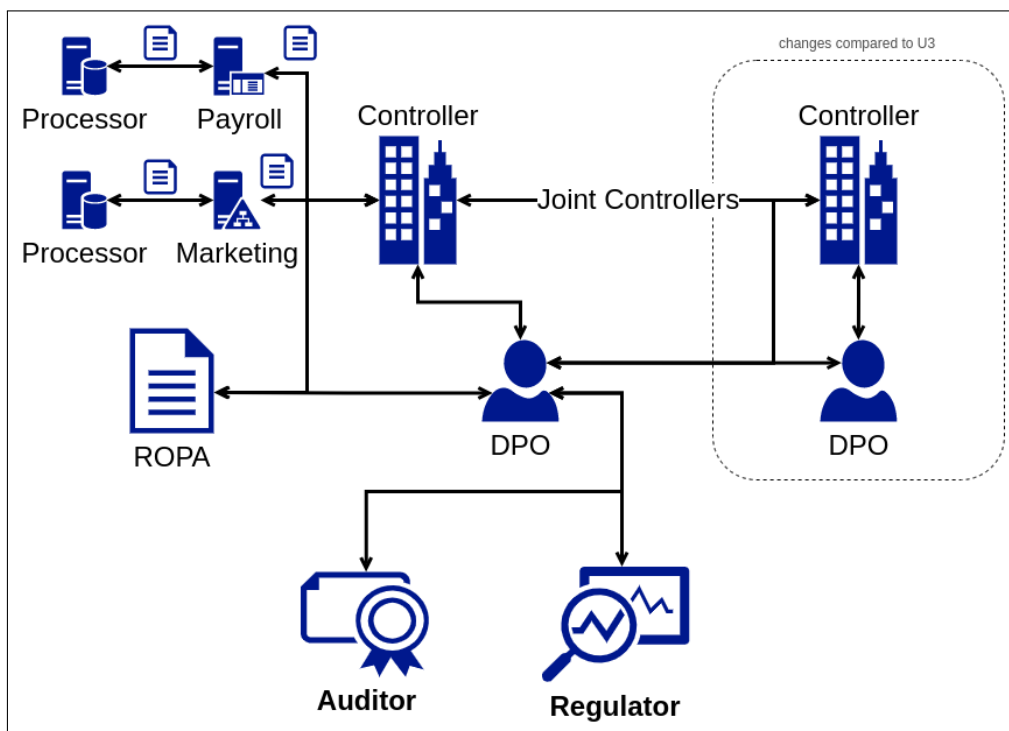


Figure 4. Data controller in a joint controller relationship.

The key information flows for these stakeholders, in addition to U3, involve: (i) ROPA information from joint controllers; (ii) provenance, e.g., sources, timestamps, contact details; (iii) collation of information from heterogeneous sources into a common ROPA; and (iv) a selective part of ROPA, e.g., specific (external) controller.

4.5. Use Case U5: DPO Overseeing Multiple Data Controllers

Use cases U1–U4 considered the perspective of a data controller that employs a DPO in terms of managing their ROPA information. In U5, illustrated in Figure 5, we consider the scenario of a DPO being an external organisation or individual providing ‘DPO-as-a-service’. We call this entity ‘External DPO’, and consider their duties as involving overseeing multiple organisations. The external DPO has to address U1–U4 for several organisations which translates to additional information flows. This is distinct from information flows associated with other external entities, i.e., DPAs or auditors, in that the external DPO requires information including internal organisational units and data governance processes for an accurate understanding and potential follow-up tasks.

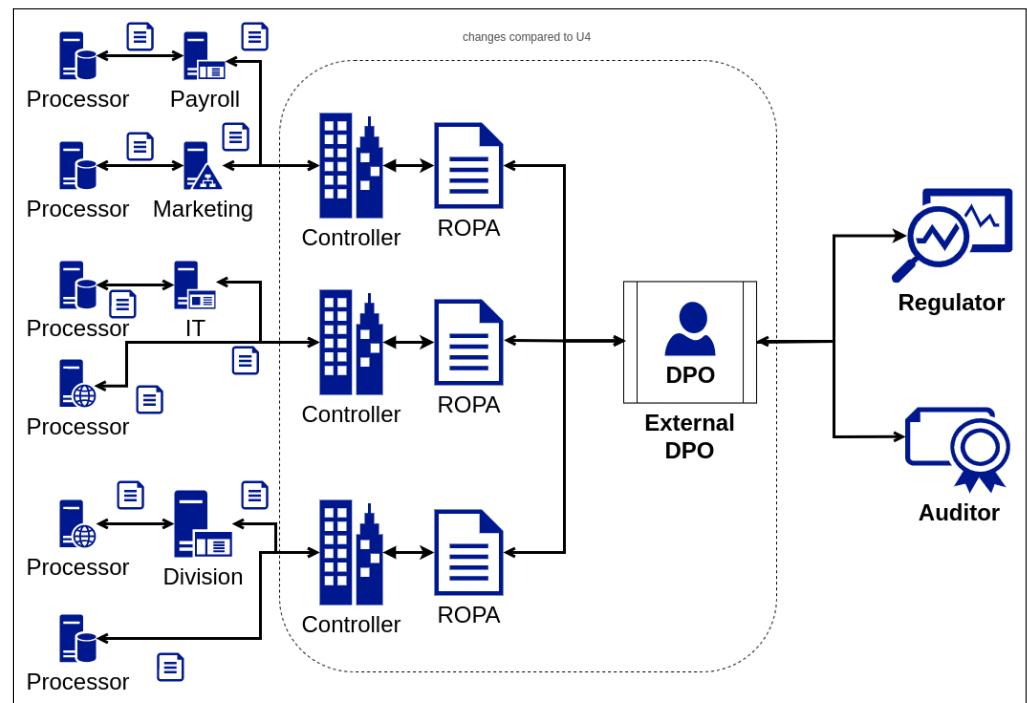


Figure 5. A DPO overseeing multiple data controllers.

The key information flows for these stakeholders, in addition to U1–U4, involve: (i) collect ROPA information from multiple organisations; (ii) produce ROPA for a specific controller; (iii) provenance, e.g., sources, timestamps, contact details; (iv) separation of ROPA-related information reflecting organisational units, e.g., departments; (v) a selective part of ROPA, e.g., specific department for a specific controller.

4.6. Requirements Analysis

Since the GDPR does not dictate or concern how a ROPA is generated or maintained as long as it meets the legal requirements, the organisation has the freedom to determine practices that suit its compliance approach and style. For example, an organisation may choose to maintain ROPAs centrally overseen by the DPO, where information from all sources is fetched externally and collated into a common document (e.g., a spreadsheet) and added to the information management system. Alternatively, an organisation may opt to maintain separate ROPA documents for each of its departments.

In either case, upon being asked by a DPA or an auditor, an organisation has to produce a ROPA for the specified criteria such as reflecting specific processing activities or for a certain time period. An organisation must first identify the relevant information from its ROPA documents and extract the required information. This task can involve manual efforts by the DPO or responsible entity unless the organisation utilises technological solutions that support such use cases and provide an easier workflow based on automation.

Identifying these distinct use cases enables us to collect and harmonise the requirements regarding the expression of information and interpret them for use of the CSM-ROPA for data governance. We indicate that a solution must:

1. Indicate the source of information, e.g., department, processor;
2. Present collation of ROPA from discrete, possibly partial information artefacts, e.g., purpose from the department, technical measures from processor;
3. Record provenance, e.g., timestamps;
4. Record organisational details, e.g., point of contact, responsible entity;
5. Maintain distinct records, e.g., department, processors, or temporal periods;

We present additional requirements for enabling systems to use this information:

6. *'Packaging'*: Sharing ROPA record(s) with internal or external stakeholders, e.g., department to DPO or processor to controller;
7. *'Querying'*: Retrieving partial information from ROPA, e.g., specific period or process;
8. *'Exporting'*: Generating ROPA documentation as per requirements, e.g., GDPR Art.30;
9. *'Customisation'*: Customising information storage, retrieval, and exporting based on a variance in requirements, e.g., additional information for specific DPA templates;
10. *'Assuring'*: Providing data integrity and other quality guarantees for records;

We present further additional requirements that motivate operational details:

11. *'Machine-readability'*: for using automation and tooling for information management;
12. *'Interoperability'*: for consistency in and interpretation across stakeholders;
13. *'Openness'*: for enabling adoption without lock-ins across technologies or providers;
14. *'Extendability'*: to enable customisation of a solution for a use case or contextual requirements, e.g., additional terms, new information requirements;
15. *'Verifiable'*: to support information management through validation of information in terms of correctness and completeness, e.g., all necessary fields are declared with valid information types, as well as to support compliance processes in ensuring validity and accountability, e.g., ensuring every processing has a purpose.

From this, we conclude that CSM-ROPA, while sufficient to represent information required to generate a ROPA, is insufficient to meet requirements for exchanging or using information amongst relevant stakeholders. In the next section, we present our use of this as a motivating factor for developing a 'catalogue' that can encapsulate the ROPA-related information and satisfy requirements for its maintenance and exchange with stakeholders.

5. DPCat: A Data Processing Catalogue

In the earlier sections, we presented a 'Common Semantic Model for ROPA' (CSM-ROPA) representing a consolidated set of information requirements based on an analysis of DPA ROPA templates that can be used as a machine-readable and interoperable vocabulary through the use of DPV (see Section 3). We then explored the sources and use of ROPA in terms of data flows between stakeholders and identified five use cases that provided further requirements regarding using CSM-ROPA in practical settings. In this section, we present the data processing catalogue (DPCat) specification, also published online (<https://w3id.org/dpcat>), that addresses identified requirements and facilitates governance of information from intra- and inter-organisational heterogeneous sources to enable representation of a ROPA in a machine-readable and interoperable manner.

DPCat extends the Data Catalog Vocabulary (DCAT)—a W3C standard for facilitating interoperability between data catalogues [17], with concepts identified in CSM-ROPA

using DPV to enable representation of ROPA and associated information as ‘catalogues’ and ‘datasets’, respectively, that can be recorded and exchanged between stakeholders. DPCat also aims to maintain compatibility with the ‘DCAT Application profile for data portals in Europe’ (DCAT-AP) [18] which represents the EU’s efforts at standardising catalogue metadata in data portals. The compatibility with DCAT-AP is based on our vision for DPCat to be usable in all DCAT-AP based catalogue information management tools and data portals to present a mechanism for sharing ROPA related information using an EU-advocated standard and to promote the possibility of reusing existing data portal infrastructures for compliance-related purposes - such as requirements for ROPA between controllers, processors, and DPAs.

Our prior work [2,3] regarding DPCat was based on the CSM-ROPA developed from six DPA ROPA templates that addressed two (U2: organisational units, U3: processors) of the five use cases. This work incorporates updated CSM-ROPA for 17 DPA ROPA templates, updates made to the DPV (from v0.2 to v0.5), and integration of DCAT and DCAT-AP requirements (e.g., cardinality) for compatibility.

5.1. DPCat Overview

DPCat, illustrated in Figure 6, distinguishes between a ROPA (as a document or an artefact) and ‘entries’ within a ROPA where each *entry* represents a specific context—such as a business process or data processing purpose. To represent these, we semantically extend the DCAT-AP concepts ‘*catalog*’ and ‘*dataset*’ as ‘*ROPA*’ and ‘*ROPAREcord*’, respectively. We also extend ‘*catalog*’ as ‘*ROPACatalog*’ to represent a collection of *ROPA* catalogues (i.e., a catalogue of catalogues) for when an organisation has multiple *ROPA* documents, e.g., representing different temporal periods or activities or organisational units (e.g., departments). Appendix C provides an overview of these concepts.

ROPAREcord is a *dcat:Dataset* that catalogues information to be documented in a *ROPA*, is akin to a ‘single row’ in a *ROPA* spreadsheet and represents a single record of processing. It is used as an instance of *dpv:PersonalDataHandling* to associate concepts such as purposes of processing or legal bases using the relevant DPV concepts identified from the CSM-ROPA analysis. To ensure compatibility with DCAT and DCAT-AP requirements and recommendations, such as a publisher being a *foaf:Agent*, DPCat declares the relevant DPV concepts as a subclass of DCAT-AP specified concepts. In a *ROPAREcord* instance, the concepts are coherent, i.e., all purposes apply to all personal data and are shared with all recipients and so on. To indicate separation, separate instances should be created.

A (*dpcat:*)*ROPA* represents a *dcat:Catalog* consisting of one or more *ROPAREcord* datasets and reflects the conventional perspective of ‘*ROPA* as a single document’ with each entry being a *ROPAREcord* within the catalogue. In both *ROPA* and *ROPAREcord*, the DCAT properties provide an association with relevant information such as the publisher indicating who had produced or provided that record, temporal annotations such as when the record was produced, or the time period represented, and annotations such as titles and descriptions. A *ROPACatalog* is the same as a *ROPA* in terms of being extended from *dcat:Catalog* and is used to bundle related *ROPA* catalogues together using *dcat:catalog* relation.

For common *ROPA*-related communication between stakeholders, such as associating a ‘point of contact’ (e.g., department or manager) for that information, DPCat uses DCAT relation *dcat:contactPoint*. Additionally, to adhere to the GDPR terminology, it uses the DPV properties to indicate controller (*dpv:hasDataController*), DPO (*dpv:hasDataProtectionOfficer*), and ‘responsible entity’ (*dpv:hasResponsibleEntity*). In this, the overlap between DCAT and DPV terms, such as the controller being the publisher or the DPO being the point of contact, may not always occur—such as when representing activities limited to a department where the point of contact is a member of that department who liaises with the DPO.

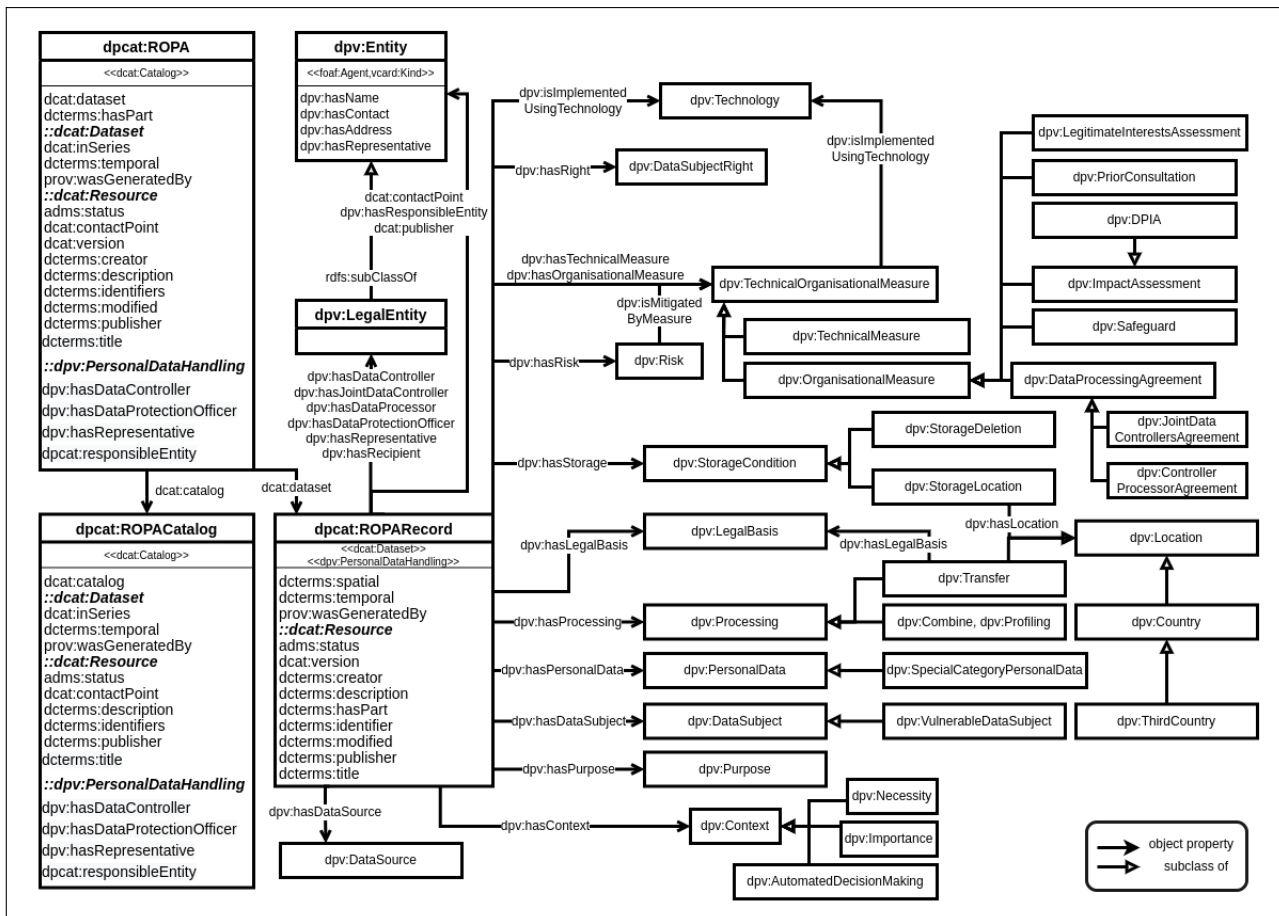


Figure 6. DPCat ROPA, ROPACatalog, and ROPACatalog overview.

5.2. Using DPCat for ROPA Information Management

As we elaborated on in Section 4, the information and data governance requirements within the use cases show a need for each entity to organise, maintain, and exchange relevant information to carry out ROPA-related processes. DPCat, as a specification, supports automation through integration into tools used for information storage and retrieval (e.g., databases) and information management practices (e.g., documents and data catalogues). It can represent all ROPA-related information or only catalogue metadata with links to the actual information stored externally (e.g., spreadsheets) as datasets. In either case, DPCat provides a consistent information structure that enables technological solutions such as querying, validation, and exporting (see next sections) to assist the relevant stakeholders in their tasks.

DPCat facilitates data governance for ROPA by incorporating the organisation’s structural and managerial requirements. For U1, where a ROPA has to be maintained at the organisational level, the ROPA and ROPACatalog data can be maintained centrally. For U2–U4, where there are heterogeneous sources of information, and it is desirable to record them in the same manner for provenance and follow-ups, the DCAT relations enable provenance of publishers and points of contact. In contrast, ROPACatalog enables collections of related information issued by, e.g., a department or a processor.

The semantics of DCAT provides flexibility in determining how ROPA information could be organised and stored without determining a single method or structure. For example, in addition to the structuring based on organisational units and external entities, it may be desirable to keep records based on contextual information—such as specific business processes related to a product or service. This can be achieved by creating additional ROPACatalog entries representing the other collection and linking them to

relevant *ROPA* entries. Through this, organisations can achieve multifaceted approaches in using the *ROPA* information without data duplication.

The use of technological solutions advocated by DPCat faces a hurdle in that the sources of information in *ROPA*-related workflows may not necessarily have the technical knowledge to produce consistent and valid metadata. For example, a DPO with the necessary legal knowledge does not necessarily have or is concerned with the underlying technicalities of information storage and retrieval beyond what is necessary to perform their duties. In such cases, existing information storage and management mechanisms such as databases and spreadsheets can continue to be used by DPCat being integrated into them rather than acting as a replacement. For example, using a SQL database, the information represented in its tables would utilise the DPCat as a schema with the input provided through existing means, e.g., input forms or importing spreadsheets using controlled structures. Alternatively, using an RDF-based solution such as a triple-store, the forms or spreadsheets could be converted to DPCat by utilising mappings.

We envision DPCat to be integrated into a typical workflow (i.e., U2–U4) for recording *ROPA* as follows. The source (e.g., department representative) generates a *ROPARecord* containing relevant information with provenance as the department. They use mechanisms available to them, e.g., a series of forms or a script that converts spreadsheets. This information is collated into a *ROPA* collection representing contextual grouping as determined by the organisational structure (e.g., maintained per department). For sources external to the organisation (e.g., a processor), the provided information is similarly stored in dedicated *ROPA* and *ROPARecord* entries and optionally integrated directly into relevant datasets (e.g., controller listing processor's technical measures in its *ROPA*). This can use technological solutions such as a database or a portal. To facilitate the structuring of *ROPA* records in an organised manner, *ROPACatalog* entries are used to collect and group *ROPA* entries according to some criteria, e.g., temporal period, legal counsel, or responsible managers.

5.3. Using DPCat for Querying and Validation

DPCat supports and enables a wide assortment of queries and validation approaches that utilise its metadata-based structure to perform information retrieval and verification tasks. DPCat can be a vital tool in technological solutions used for compliance-related processes through these. This section presents a few examples of queries and validation tasks that motivate the use of DPCat in an organisation's *ROPA*-related processes.

A common query associated with a *ROPA* is retrieving the GDPR Art.30 information for a specific context such as data transfers or covering some time period. DPCat supports such queries through DCAT and DPV metadata, e.g., indicating transfer locations as *dpv:DataTransfer* and *dpv:hasLocation* and DCAT *dcat:temporalPeriod* to perform time-based filtering. An example of this expressed as a SPARQL query is provided in Listing 1.

Similar to querying, DPCat also supports verification and validation of information, typically ensuring or assessing compliance with the GDPR. Validation refers to whether sufficient information is available, is in the correct form and format, and is sufficient according to some requirements. Verification refers to the evaluation of the information based on some norms such as specific obligations of the GDPR.

Constraints based on mandatory fields as prescribed by DCAT and DCAT-AP specifications also apply to DPCat. Therefore, data represented using DPCat can utilise existing validation and verification mechanisms for conformance to these standards. In addition, DPCat promotes the expression of the GDPR-specific constraints that are typically expressed as guidelines by DPAs and have been the subject of research by academic and commercial offerings. However, DPCat has an advantage over these existing solutions in that it also promotes interoperability between such verification mechanisms by virtue of being an interoperable specification for information to be verified.

Listing 1: SPARQL query retrieving ROPA records involving data transfers in a time period.

```

1 ?Entry a dpcat:ROPAREcord .
2 ?Entry dct:title ?title .
3 ?Entry dct:publisher/dpv:hasName ?publisher .
4 OPTIONAL { ?Entry dcat:contactPoint/dpv:hasName ?contact } .
5 ?Entry dct:created ?created .
6 ?Entry dpv:hasProcessing ?transfer .
7 ?transfer a dpv:Transfer .
8 # minimum date within which data transfer occurs
9 ?Entry dct:temporal/time:hasBeginning/time:inXSDDate ?start .
10 FILTER (?start < "2021-01-01"^^xsd:date) .
11 OPTIONAL {
12     # maximum date, if available
13     ?Entry dct:temporal/time:hasEnd/time:inXSDDate ?end .
14     FILTER (?end > "2022-01-01"^^xsd:date) .
15 }
16 OPTIONAL { ?transfer dpv:hasDataImporter/dpv:hasName ?importer . }
17 OPTIONAL { ?transfer dpv:hasDataExporter/dpv:hasName ?exporter . }

```

As an example of information validation typically involved for the GDPR, Listing 2 presents a SHACL constraint that ensures every *ROPAREcord* instance has an associated purpose. In addition, to ensure information is present and in correct form, SHACL constraints are also useful towards GDPR compliance such as for ensuring an appropriate legal basis as follows: (i) It must have a corresponding legal basis from the GDPR Art.6 (Note: Though the GDPR Art.30 does not require a legal basis in a ROPA, DPA guidelines strongly recommend it.); (ii) If processing involves special categories of personal data, it must additionally have a corresponding legal basis from the GDPR Art.9; (iii) If processing involves data transfers to non-EU locations, it must additionally have a corresponding legal basis from the GDPR Art.45, 46, or 49. We plan to provide such SHACL shapes for both information validation and the GDPR-based requirements verification in the future.

Listing 2: SHACL constraint to ensure every *ROPAREcord* has an associated purpose.

```

1 dpcat:Shape_EnsurePurpose
2   a sh:NodeShape ;
3   sh:name "Ensure every processing has a denoted Purpose "@en ;
4   sh:description "Ensure the dpv:hasPurpose property is defined and has
5     a value that is an instance of dpv:Purpose"@en ;
6   sh:targetClass dpcat:ROPAREcord ;
7   sh:property [
8     sh:path dpv:hasPurpose ;
9     sh:class dpv:Purpose ;
10    sh:minCount 1 ;
11  ] .

```

5.4. DPCat for Interoperable Information Exchange

DPCat provides a machine-readable and interoperable representation of information that an organisation can use to automate its ROPA management and associated tasks. In cases where information has heterogeneous sources, especially when involving external stakeholders such as processors and other controllers, DPCat can be utilised as a ‘standardised information representation’ for convenience in information flows. In this section, we explore the potential for such developments.

When they hire data processors, a data controller’s obligation includes maintaining information about the processing activities outsourced to the processor and some specifics regarding how they are carried out. For example, controllers may ask processors to provide the technical and organisational measures they implement to ensure sufficient safety and security in processing. Similarly, controllers may require information for data storage locations of data for cross-border data transfers. In cases where a processor contracts another (sub-)processor to carry out the processing, it has to maintain similar records of the sub-processor’s operations, but it also provides them to the controller as requested. In all these, information has to be periodical—maintained independently by the entity itself and

communicated to other entities as contextually necessary, and the other entities must also maintain this information independently. Such information flows and requirements are also necessary for a joint controller's relationships regarding involved controller(s).

If two entities communicating information for ROPA-related tasks use DPCat for their internal information representation, they can directly exchange ROPA information using DPCat-specified records. This is an ideal scenario. However, even if either or no entities do not use DPCat internally, DPCat can be utilised as a common specification for exchanging ROPA information between entities. In this case, the sender entity converts whatever internal representation it has into DPCat and sends it to the receiver entity to ensure that it can understand and interpret the information. The receiver converts DPCat-based information to whatever internal representation they utilise. Thus, DPCat offers advantages for ROPA information exchanges even if organisations do not wish to adopt it completely for internal processes. DPCat is also useful for DPAs and auditors in the same manner, where they can utilise it as an interoperable format for requesting information from organisations. The consistency and machine-readability of DPCat provide investigators with the potential for using automation and tools to reduce workload and repetitions.

6. Demonstration of DPCat in a Real-World Use Case

In this section, we demonstrate the application of DPCat in representing real-world ROPA documents published by the European Data Protection Supervisor (EDPS) [41] (EDPS), perform validations of them using SHACL, retrieve relevant information using SPARQL queries, and export it as RDF graphs as well as spreadsheets adhering to DPA templates. We provide evidence for the practicality and feasibility of DPCat and its benefits in ROPA information management processes. The data, code, and outputs are available online (<https://w3id.org/dpcat/demo/edps-ropa>).

6.1. Representing Information Using DPCat

EDPS is the DPA responsible for overseeing compliance by EU institutions, which consists of many employees across the various EU bodies and their associated personal data processing activities. The EDPS has published detailed ROPA documents based on the GDPR Art.30 requirements that provide transparency and accountability. As of March 2022, the EDPS has made available 58 ROPA document collections—with each consisting of one more PDF (format) document providing information in English regarding the processing operations. Collections are structured based on 'topics'—which can be a department (e.g., administrative and human resources or IT), processes (e.g., communication or public events), or specific measures (e.g., access to documents or physical security).

We analysed EDPS ROPA documents and selected four (ids: 01, 05, 13, 55) that covered the U1–U4 use cases for departments, processors, joint controllers, and data transfers. We did not include the other documents despite their relevance due to the large labour and analysis efforts required and because the selected documents sufficed in demonstrating DPCat's application. The documents were PDFs intended for human comprehension and lacked consistent semantics, e.g., *purpose* field also contained *legal basis*.

We interpreted these documents and their structure as follows: each document (i.e., PDF) represented a single ROPA instance, and the information contained within it was structured using *ROPAREcord* instances. We utilised the criteria that each *ROPAREcord* would adhere to a single 'contextual entry' based on qualitative criteria regarding the complexity of information and separation of concerns. For example, document X specified two processors, which we interpreted as separate *ROPAREcord* instances for each processor to indicate the separation of concern in the controller's communication and data governance. The entire collection of documents and RDF graphs were then expressed as part of a single *ROPACatalog* instance reflecting the published set of records on EDPS' website.

The manually created RDF graphs were enhanced using the Apache Jena RDFS reasoner [42] to create a 'complete graph' for simplifying querying and validation. The limited RDFS reasoning was sufficient here to obtain the expansion of subclasses and subproperties

within the graph rather than generating inferences using an OWL reasoner. For storing the information and offering a querying interface, we utilised GraphDB Free Edition triple store [43], as it is a freely available triple-store compliant with relevant standards (e.g., SPARQL) and has several features for convenience, e.g., friendly interface, integrated reasoners, SHACL validation. A concise example of the *ROPARecord*, *ROPA*, and *ROPARecord* instances used to represent the EDPS ROPA documents and information is presented in Listing 3. An overview of data workflow is provided in Figure 7:

Listing 3: Summarised overview of *ROPA* and *ROPARecord* instances based on EDPS ROPA documents.

```

1 @prefix edps: <https://w3id.org/dpcat/examples/EDPS/vocab#> .
2 @prefix : <https://w3id.org/dpcat/examples/EDPS/05#> .
3 : a dpcat:ROPA ;
4   dct:title "Selection of staff"@en ;
5   dct:description "The purpose of the processing ..."@en ;
6   dct:created "2021-11-03"^^xsd:date ;
7   dct:identifier "05"^^xsd:string ;
8   dct:publisher edps:HRBA ;
9   dcat:contactPoint edps:HRBA ;
10  dcat:dataset :05-1, :05-2, :05-3, :05-4 .
11 :05-1 a dpcat:ROPARecord ;
12   dct:title "Selection and management of interim staff"@en ;
13   dct:description "Selection and management of interim staff"@en ;
14   dct:created "2022-02-16"^^xsd:date ;
15   dct:publisher edps:HRBA ;
16   dcat:contactPoint edps:HRBA ;
17   dpv:hasDataController edps:EDPS ;
18   dpv:hasResponsibleEntity edps:HRBA ;
19   dpv:hasPersonalData edps:InterimAgentLastName,
20     edps:InterimAgentFirstName, edps:InterimAgentCV ;
21   dpv:hasPurpose edps:MonitoringOf7YearRule, edps:
22     PreparationOfEmploymentContracts, edps:ExecutionOfContract ;
23   dpv:isImplementedUsingTechnology edps:SYSPER ;
24   dpv:hasRecipient edps:Managers, edps:HRBA ;
25   dpv:hasLegalBasis dpv:StaffContract ;
26   dpv:hasTechnicalOrganisationalMeasure edps:
27     AuthorisedPersonnelNeedToKnowBasis ;
28   dpv:hasDataSource edps:RandstadBelgium, edps:Daoust ;
29   dpv:hasStorage [
30     a dpv:StorageCondition ;
31     skos:editorialNote ""The data ... managed by Sysper""@en ;
32     dpv:hasDuration [
33       a dpv:StorageDuration, time:Duration ;
34       dpv:hasPurpose edps:EnsureReconstructionOfHistoryBySYSPER ;
35       dpv:hasPersonalData edps:DataRelatingToSYSPERService ;
36       time:numericDuration "5"^^xsd:decimal ;
37       time:unitType :unitYear ; ] ] ...

```

6.2. SHACL Shapes for DPV

For verification and validation of the generated RDF graphs, we first utilised the SHACL constraints provided with DCAT-AP specifications to ensure data correctness according to DCAT and DCAT-AP defined requirements, e.g., publishers being of type *foaf:Agent*. We then developed and utilised SHACL shapes representing the cardinality and type constraints to ensure correctness for DPCat's requirements. For executing the constraints, we utilised the open-source and freely available TopBraid SHACL tool [44].

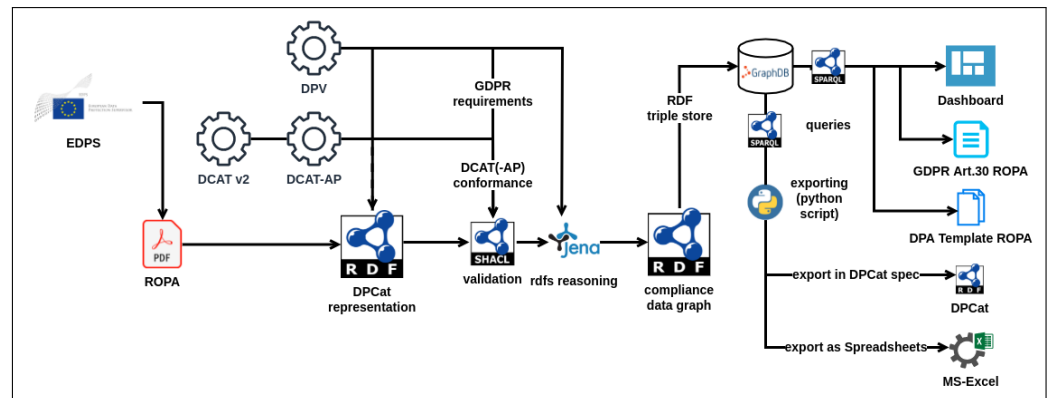


Figure 7. Data workflows in DPCat demo application.

In performing validation of the information, the shape constraints are based on DPCat, which utilises DPV and DCAT concepts to represent relevant information. However, neither DPV nor DPCat indicates what ‘shape’ some information must be represented. Consequently, there may be more than one ‘shape’ for a given scenario, often at arbitrary levels of complexity, which prevents a single set of common SHACL shapes from being developed and provided alongside the DPCat specification. For example, a SHACL constraint for ensuring data transfers is specified along with its appropriate location can be modelled in terms of *dpv:hasLocation* of *dpv:DataTransfer*. However, the *DataTransfer* instances could be used at any arbitrary node within the graph, making it difficult to define follow-up constraints such as the recipient of that transfer and its location.

A simple solution would be to associate all the relevant fields with a *ROPA* or *ROPARRecord* instance. A challenge in this is that all the DCAT-based structure may not be capable of incorporating all fields or that it would make DPCat too complex. An alternate approach would be identifying use cases for each concept’s use and defining specific SHACL shapes for how that information should be expressed using DPV. Given that this requires significant analyses and effort, for the purposes of this article, we limited our defined SHACL shapes for representing information from the EDPS documents. However, we argue for further research and development of such shapes so that they can be used to ensure data is consistently represented across use cases and implementations.

6.3. Querying ROPA Information

To simulate typical tasks performed by a DPO or a DPA, we utilised SPARQL queries for two use cases: (i) retrieval of information required by the GDPR Art.30; and (ii) overview of practices within an organisation in terms of various organisational units, purposes, legal bases, recipients, data transfers, etc. Here, query (i) relates to common compliance documentation procedures, and query (ii) shows the potential for DPCat to help create internal reports or dashboards based on ROPA information, e.g., for a DPO.

The first query, shown in Listing 4 with an output snippet in Table 1, retrieves ROPA information as per the GDPR’s Art.30.

The second query, shown in Listing 5 with an output snippet in Table 2, provides an overview of the organisation’s processing activities and relationships with external entities by retrieving relevant information from *ROPARRecord* instances.

Listing 4: Obtaining a GDPR Art.30 based overview using SPARQL.

```

1 SELECT DISTINCT ?Entry ?title ?purpose ?datasubject ?personaldata
2   ?recipient ?legalbasis ?transfer_location
3 WHERE {
4   ?Entry a dpcat:ROPARRecord .
5   ?Entry dct:title ?title .
6   ?Entry dpv:hasPurpose/skos:prefLabel ?purpose .
7   ?Entry dpv:hasDataSubject/skos:prefLabel ?datasubject .
8   ?Entry dpv:hasPersonalData/skos:prefLabel ?personaldata .
9   OPTIONAL { ?Entry dpv:hasRecipient/dpv:hasName ?recipient } .
10  OPTIONAL { ?Entry dpv:hasLegalBasis/skos:prefLabel ?legalbasis . }
11  OPTIONAL { ?Entry dpv:hasProcessing ?processing .
12             ?processing a dpv:Transfer .
13             ?processing dpv:hasLocation/skos:prefLabel ?transfer_location } }

```

Table 1. Query results (summarised) for the GDPR Art.30 information using DPCat.

Title	Purpose	Data Subject	Personal Data	Recipient	Legal Basis	Transfers
Selection of staff	Staff Selection	Job Applicants	Applicant CV	Selection Panel	Staff Reg. 2020	
Financial Transactions	Payment	Staff members	Physical Address			
Financial Transactions	Payment	Staff members	Credit Worthiness	AirPlus		Third Country
Financial Transactions	Budgetary commitments	Staff members	Job Applicant CV	ERCEA's Speedwell operators		
Financial Transactions	Budgetary commitments	Staff members	Bank Account	Local Profile Manager		
Financial Transactions	Payments	Staff members	Bank Account	The EDPS Financial team		

Listing 5: SPARQL query for overview based on GDPR Art.30 using DPCat.

```

1 SELECT DISTINCT ?org ?title ?purpose ?processor ?jointcontroller
2 WHERE {
3   ?record a dpcat:ROPARRecord ; dct:title ?title .
4   ?record dct:publisher/dpv:hasName ?org .
5   ?record dpv:hasPurpose/skos:prefLabel ?purpose .
6   OPTIONAL { ?record dpv:hasDataProcessor/dpv:hasName ?processor }
7   OPTIONAL {
8     ?record dpv:hasJointDataControllers/dpv:hasName ?jointcontroller
9   } }

```

6.4. Exporting a ROPA

To demonstrate how DPCat can facilitate information exchange and data governance within and between stakeholders, we provide two examples of information being exported. The first example exports information as DPCat-defined catalogues by using SPARQL CONSTRUCT queries to retrieve related information as an RDF graph. The SPARQL query and the resulting graph can be viewed online. Such exports help store information in the form of backups, copies, or graphs. It is also helpful in exchanging ROPA information between stakeholders such as those accompanying data governance between data controllers and data processors that all support use of DPCat.

The second example simulates automation of a DPO manually managing information in a spreadsheet. For this, we utilised a Python script that executed SPARQL queries and exported results into an MS-Excel (.xlsx) document based on DPA ROPA templates. While the output of a SPARQL query itself could also be exported as a CSV document, the use of Python in this case was to replicate the structure and contents of the DPA template and to operate over the more complex XLSX format that supports tabs within spreadsheets.

Table 2. Query results (summarised) for the GDPR Art.30 information using DPCat.

Department	Process	Purpose	Data Processor	Joint Controller
Human Resources, Budget, Administration (HRBA) Unit	Staff Selection	Select staff for the EDPS and EDPB Secretariat		
Human Resources, Budget, Administration (HRBA) Unit	Selection and management of interim staff	Monitoring of 7-year rule (EDPS Decision 13.12.2018)		
Human Resources, Budget, Administration (HRBA) Unit	Communicate staff selection	Select staff for the EDPS and EDPB Secretariat	Randstad Belgium SA/NV	
Human Resources, Budget, Administration (HRBA) Unit	Communicate staff selection	Select staff for the EDPS and EDPB Secretariat	Daoust SA/NV	
Human Resources, Budget, Administration (HRBA) Unit	Payment of Invoices for services	Payment of invoices for services		
Human Resources, Budget, Administration (HRBA) Unit	Communicate staff selection	Communicate staff selection	Randstad Belgium SA/NV	
Human Resources, Budget, Administration (HRBA) Unit	Communicate staff selection	Communicate staff selection	Daoust SA/NV	
Human Resources, Budget, Administration (HRBA) Unit	Administration of Access Requests	Administration of Access Requests		
Human Resources, Budget, Administration (HRBA) Unit	Financial Transactions	Financial Transactions	EC—DG-BUDG	EC—DG-BUDG

6.5. Analysis of Implementation and Lessons Learned

The application of DPCat to real-world ROPAs exposed inherent difficulties in constructing semantic representations due to inputs lacking or being loosely structured as opposed to strict structure that machine-based tools require. We discussed exploring this issue further with a proposed solution where a separate *registry of controlled vocabularies* is created by the organisation for the use case to first register their concepts such as the specific purpose used or data category processed and to then ensure the ROPA documents only used these concepts. However, we found this solution to significantly deviate from the organisational processes that lack such structured data collection methods. We consider this an open problem with the hope of better tooling being able to resolve it.

In representing the ROPA information using DPV, we faced hurdles in that the DPV as a vocabulary can support a wide range of data modelling styles. This presents barriers to the use of DPCat as a common information representation mechanism as two different organisations can model their data differently. While the common conceptual structure of DPV can assist in aligning the two models, it is better for the development of tools to have a consistent information structure. For this, we propose the creation of ‘DPV Shapes’ that provide suggested data modelling practices for modular use cases. Such shapes, expressed using SHACL, will foster commonality in how the DPV is used and will act as a common model for other modelling approaches that can be reduced or aligned to. In this, it is important to state one of the strengths of the DPV is its lack of rigid adoption requirements which provides an adopter the flexibility to use it within their use cases. The provision of shapes enables continued flexibility of the DPV as a vocabulary while providing guidelines for how it can be consistently used or made interoperable across different applications.

Finally, we faced challenges in determining a suitable mechanism for validation of DPCat-specified information. While we utilised SHACL shapes to demonstrate the potential for such validations based on information and GDPR compliance requirements, this area merits further exploration. In particular, SHACL constraints can be used for two categories of evaluation: first to check whether the necessary information is present and has expected values—similar to DCAT-AP SHACL shapes. The second is based on requirements drawn from the GDPR such as ensuring the correct legal basis is used. In these, the first is an inherent evaluation of *conformance* to a specification, as seen from the cardinality constraints in DCAT and DCAT-AP, while the second directly addresses GDPR compliance verification. This follows earlier research explorations demonstrating the use of SHACL constraints in ensuring information correctness and conformance for GDPR compliance [45].

7. Discussion

7.1. Impacts on ROPA-Related Research

Our literature review (Section 2) shows that although GDPR compliance is well studied, there has been a lack of academic research specifically addressing ROPAs. Labadie and Legner [11] identify “maintenance of records of processing activities” as a core GDPR data management (sub-)capability for organisations. DPCat transforms this into an IT system capability by extending Labadie and Legner’s model. In addition, DPCat adds sub-capabilities for aggregation of diverse accountability data, exchange of machine-readable ROPA information with stakeholders, generation of DPA-specific compliance records, and assurance mechanisms for ROPA data quality (Figure 8).

System Capabilities				
Define protected data scope	Identify data objects	Classify data attributes	Locate data records	
Manage Consent	Implement consent items	Record consent instances	Distribute consent	Enforce consent-based processing
Enable Data Processing Rights	Delete data	Pseudonymize data	Transmit data in standardized form	
Maintain ROPA	Aggregate accountability data	Exchange standardised data with stakeholders	Generate DPA-specific records	Assure data quality
Organisational Capabilities				
Orchestrate Data Protection Activities	Assume data protection responsibilities	Oversee data protection activities	Control compliance of external processors	
Demonstrate Compliant Data Processing	Maintain records of processing activities	Maintain documentation of system landscape	Supervise sensitive processing activities	
Disclose Information	Disclose information to individuals	Disclose information to authorities		

Figure 8. Labadie and Legner capability model for data management in GDPR [11] extended with “Maintain ROPA” system capabilities using DPCat—highlighted with bold text and red borders.

In addition, the analysis of DPA ROPA templates (Section 3) has demonstrated the previously undocumented extent of variance between DPA approaches which impacts DPOs and the GDPR-aware system designers or integrators. The creation of CSM-ROPA provides an ontological structure for representing ROPA-related information. By utilising DPV as a community-endorsed specification that draws on the skills, requirements, and expertise of DPOs, legal experts, and technologists, it provides a strong basis for establishing an agreement of semantics to address the gap identified by existing efforts [11,24,25]. In addition, CSM-ROPA supports existing approaches based on semantics by providing a target ontology, such as for enterprise architecture models proposed by Huth [13].

By moving ROPA processes to a data cataloguing approach, the DPCat specification facilitates adoption of modern metadata-driven data governance [40]. This in turn motivates adoption of data stewardship to support intra- and inter-organisational heterogeneous sources of information and compliance [11]. It also aligns with recent EU recommendations on placing data governance at the centre of personal and AI-based data processing [46]. These advances are significant when compared with low levels of governance and automation of ROPAs exhibited by most organisations to date and provide a step towards effective technologies and tools for data protection.

7.2. Impacts on Real-World GDPR Compliance

This section discusses impacts of DPCat on ROPA governance: enforcing consistency on current largely manual exercises for ROPA data maintenance, the impact on privacy and data protection software, and the impact on data exchange within large organisations.

The application of DPCat to EDPS ROPAs highlighted the cumbersome efforts required in disentangling machine-readable data from human-oriented documentation, even when data is in semi-structured tabular form. The human creators of ROPA information tried to enforce consistency and ease of presentation through such structures, but the variations and inconsistencies in readily using these as machine-readable information were apparent when converting them to RDF using DPCat. In addition, cross-references or inferences had to be made from other PDF documents in the dataset, leading to difficulty of validation and more opportunities for inconsistency. By providing a flexible, layered approach to machine-readable ROPA information collection, DPCat supports analysis and progressive integration or automation of processes from lightweight metadata-oriented approaches that emphasise DCAT-AP and the GDPR Article 30 requires more detailed knowledge models of accountability information that expand into full DPV models of personal data handling and beyond to other ontologies or datasets. No matter what level of modelling is selected, DPCat will provide significant advantages in terms of the ease and consistency of ROPA data maintenance in an organisation and thus, the cost of compliance.

The commercial software offerings for data protection [7] consist of tools supporting organisational units that are primarily aimed at DPOs or compliance units that maintain documentation for the organisation. However, they are information silos as they lack interoperability with other systems. In response, software companies may develop APIs or adapters to connect these to other systems, which will require updates and integration efforts for every system used by all stakeholders. This becomes particularly problematic when organisations appoint new processors, acquire companies, or add new systems. In contrast, DPCat provides a single integration point for ROPA information that can be used by any data protection, compliance, and operational business system.

DPCat also addresses the information needs of real-world intra-organisation use cases that are absent from existing literature (U1). For example, one co-author is a DPO for an organisation consisting of five divisions, with their own DPOs, containing 29 affiliated legal entities spread across the UK, the EU, and the USA. As would be the case in such organisations, there is a large amount of intra-group information flows. The organisation utilises shared services such as IT support, data analytics services, and human resources. There are also a large number of appointed processors providing services to affiliates. The organisation has a limited number of joint controller relationships, and intra-organisation processing is more common. Here, the challenge for the DPO [14] is ensuring the complexity of processing activities is collected and accurately reflected in the numerous ROPA documents. For this, the organisation uses several standalone proprietary solutions that still create a large dependence on manual effort and documents. Applying the DPCat approach to ROPA data management in such a landscape enables greater automation and spans the heterogeneous IT systems involved, both for compliance and business operations. Coping with an ever-changing diversity of internal data processing links is the key to empowering the DPO to monitor the personal data processing, communicate with stakeholders, and identify non-compliance. This then facilitates better management of external relations and compliance activities governed by contracts and law.

7.3. Practical Challenges for DPCat Deployment

7.3.1. Requirement for Enhanced Data Governance

Despite many organisations embracing the productivity and agility gains of digitalisation, they continue to struggle with the basic principles of data governance [10]. The agreed uses that data is put to must be clearly defined, and the organisation must ensure that the use of data positively relates to the regulatory environment. Organisations need to define the agreed behaviours and policies for data quality, who will access the data, how

data is interpreted, and how long the data will be retained. The challenge organisations face regarding personal data is locating, classifying, and cataloguing accountability data. DPCat provides an incentive to deploy a machine-readable data catalogue platform such as Comprehensive Knowledge Archive Network (CKAN) [47] that provides user-friendly interfaces and tools supporting all these activities. In addition, the presence of an actively maintained data catalogue is a spur to wider data governance activities in the organisation.

When examining the federated or distributed aspects of ROPA data governance deployment, the role of processors is significant. They will need to contribute updates to ROPA as per the DPCat specification. There may be resistance to this activity, or the processor may need technical assistance to meet the DPCat requirements. A wide adoption of DPCat by stakeholders would bring great benefit, but even if one company chooses to use it, it brings benefit, as DPCat can be utilised as an export format. DPCat provides some assistance here by being a single comprehensive integration point; if processors can comply with it, then unlike proprietary solutions, this integration cost should be a reusable effort that can be used for many customers playing the role of the data controller.

7.3.2. Agreed Semantics

The need for an agreed ontology to describe data processing activities is key to DPCat's success. This lack of common understanding of privacy terms is limiting the growth of the privacy tech industry [24]. DPCat provides a solution that vendors could adopt since it is based on two existing standards, DCAT and DPV. Nonetheless, vendors are typically driven by their customers or regulators to adopt open standards. Hence, the importance of the role of DPAs is discussed in the next section. Unless DPAs get involved, it will be up to DPOs and other privacy software customers to demand interoperable solutions for ROPA management.

7.3.3. Role of DPAs

DPAs could have a significant role in automated regulation, making compliance easier to achieve [48]. When we compare the success of RegTech, we see regulators that enable and facilitate digital compliance, actively promote and enable digital regulatory compliance standards, and act as enablers for the automation of regulation to actively create an environment for digital compliance [10]. Adopting DPCat for data controller to DPA communication would benefit the DPA when auditing ROPAs as audit, a breach investigation, or inspections. To achieve the adoption of DPCat, DPAs would need to move towards a symbiotic relationship with technology innovators and organisations that process personal data and develop open-source compliance tools, digital regulations, and sandboxes [49] as well as agreed common semantic vocabularies such as DPCat [10]. A proactive DPA could certainly speed up the use of DPCat, and there are certainly some moves toward technology, such as an online DPIA template [50]. However, we are very much at the early stages of automated GDPR regulatory compliance.

7.4. Limitations

While the approaches motivated by CSM-ROPA and DPCat provide promising solutions to the challenges in data governance associated with maintenance and use of ROPA towards GDPR compliance requirements, it also has certain limitations that need to be addressed to ensure it is effective in practice. In this section, we discuss identified limitations and propose future efforts toward addressing them.

7.4.1. Limitations of Scope

The DPCat specification reflects the information requirements derived from CSM-ROPA, which was constructed based on the GDPR requirements, and DPA guidelines and templates regarding ROPA. While this makes DPCat sufficient to carry out tasks associated with ROPA, it does not consider the relevance and overlap of information between a ROPA and other compliance documents such as DPIA (data protection impact assessment), TIA

(transfer impact assessment), data breach records, and controller–processor or controller–controller agreements.

In each of these, there is an obvious overlap with some of the information stored in a ROPA and the necessity to link these to the ROPA itself. For example, a DPIA may concern several processing activities that are spread across distinct *ROPAR* instances. While the ROPA can link to the DPIA document trivially through single information, it is advantageous for the DPIA information to be expressed similarly as the ROPA information to enable better information interoperability and governance and motivate the creation of tools that can work on all compliance-based activities using the same information. This can be achieved by further developing DCAT-based solutions for all of the information necessary in legal compliance tasks for the abovementioned requirements.

7.4.2. Limitations of Vocabulary

The DPV forms an important aspect of DPCat in that it provides the vocabulary for representing the GDPR-associated terms in a machine-readable and interoperable manner. Therefore, any limitations of DPV will also be reflected within the capabilities of DPCat. Given that the DPV is a community-managed resource (through the DPVCG), there is a forum for proposing additions and enrichments as needed for DPCat’s applications. However, better alignment between DPCat and DPV versions will have to be established to provide the reliability of DPCat’s usage and interpretation—for example, by pinning DPCat’s use of DPV to a specific version.

7.4.3. Limitations of Jurisdiction

DPCat as a solution is EU-centric in that it directly addresses (only) the GDPR requirements. However, there may be a wider need for organisations to document their processing activities in a different jurisdiction or a jurisdiction-agnostic manner. For addressing cases, DPCat may be supplemented with extensive modifications such as an adopter’s own jurisdiction-specific vocabulary which may bring about incompatibility between implementations. A solution would be to develop DPCat into a domain and jurisdiction-agnostic specification and then provide the GDPR specific concepts as an extension. This reflects current work regarding extending DCAT to DCAT-AP, and the provision of the GDPR-specific concepts such as through the GDPR Extension for DPV (DPV-GDPR) [51] that are separate from the ‘main’ DPV.

7.4.4. Limitations of ‘Data Shapes’

As mentioned earlier in Section 6, the querying and validation of information require consistency or foreknowledge regarding how the data is structured or ‘shaped’. Without this, the resulting SPARQL queries and SHACL shapes can be difficult to express or become complex without this. To ensure the consistency of DPCat implementation, especially for information exchange, it is vital to enforce the consistency of the underlying information. While DCAT (and DCAT-AP) provide this consistency to the expression of catalogues and datasets as resources, the lack of such consistency in the expression of DPV-specified information needs to be addressed. For this, we propose the development of use cases that define expectations and requirements for information, e.g., “a data transfer must specify location”, to create corresponding ‘SHACL shapes’ that harmonise how different implementers should utilise DPCat-specified information. This activity can be undertaken by the DPVCG for the larger benefit of all DPV adopters or, failing that, within DPCat to ensure its consistency in application.

8. Conclusions

The heterogeneity of data sources representing the organisation’s data processing activities presents significant challenges when completing a ROPA. Our research sought to establish the extent to which the DPCat specification for an interoperable and machine-

readable data processing catalogue based on DCAT-AP and DPV could overcome the heterogeneity of sources to facilitate the preparation of a ROPA.

We have shown that the DPCat specification enables more automation for realistic distributed ROPA maintenance use cases, leading to stronger regulatory compliance. DPAs, who already struggle with lack of funding and resources [52], can benefit from the use of DPCat that could ease the investigative burden required for effective enforcement. In pursuit of our first research objective (*RO1*) to identify the information necessary to represent ROPAs, we reviewed 17 ROPA templates across 31 DPAs. Our analysis identified 47 unique GDPR concepts, with templates requiring a minimum of 18 concepts up to a maximum of 32 concepts. Over the past 3 years, the DPV has been enhanced to express these concepts, and currently, 44 of the 47 concepts can be expressed exactly, and 1 can be partially expressed. The 2 remaining concepts are with the DPVCG for consideration.

For the second research objective (*RO2*), we presented the Data Processing Catalogue (DPCat) specification that facilitates governance and maintenance of data from intra- and inter-organisational heterogeneous sources to enable representation of information related to ROPA. Its application to EDPS ROPA demonstrated how DPCat could be utilised as a machine-readable solution to overcome conventional limitations for when data is maintained in documents, or proprietary systems lack machine-readability and interoperability.

The EDPS application also showed how DPCat enabled a data controller/processor to describe processing activities using a standardised model and vocabulary that facilitated aggregation, querying, validation, and exporting from heterogeneous sources (*RO3*). We used SHACL to ensure correctness, and SPARQL to query and export information for the GDPR articles and DPA templates. Through this, we established the data quality governance process for ROPA by harmonising inputs from heterogeneous sources and producing dynamic documentation that accommodates differences in regulatory approaches across DPAs.

In addition to formulating a research problem, we also explored the potential impact in real-world situations through the use case, application, discussions, and identification of concrete future directions to ensure practical benefits from implementing our work. In addition, as DPCat is an interoperable machine-readable record of the personal data processing activities of organisations, it offers avenues of future research, such as the generation of privacy notices, DPIAs, automatic supplier due diligence checking and international transfer compliance assessments from a common information model. The approach taken by DPCat, though being based on the GDPR, also has benefits for documentation requirements and compliance obligations from other laws, such as the California Consumer Protection Act (CCPA) [24].

Author Contributions: Conceptualization: P.R., R.B. and H.J.P.; Methodology: P.R., R.B. and H.J.P.; software: P.R. and H.J.P.; Validation: P.R., R.B. and H.J.P.; Formal analysis: P.R., R.B. and H.J.P.; Investigation: P.R., R.B. and H.J.P.; Resources: H.J.P.; Data curation: P.R. and H.J.P.; Writing—original draft preparation: P.R., R.B. and H.J.P.; writing—review and editing: P.R., R.B. and H.J.P.; Visualization: P.R., R.B. and H.J.P.; Supervision: R.B. and H.J.P.; Project Administration: P.R., R.B. and H.J.P.; Funding Acquisition: P.R., R.B. and H.J.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research has received funding from Uniphar PLC, and the ADAPT Centre for Digital Content Technology which is funded under the SFI Research Centres Programme (Grant 13/RC/2106_P2) and co-funded by the European Regional Development Fund. Harshvardhan J. Pandit has received funding under the Irish Research Council’s Government of Ireland Postdoctoral Fellowship Grant#GOIPD/2020/790.

Institutional Review Board Statement: Not Applicable.

Informed Consent Statement: Not Applicable.

Data Availability Statement: All resources presented in this work are made available at <https://w3id.org/dpcat/repo> including the DPCat specification, CSM-ROPA outcome, code, and data.

Table A1. Cont.

GDPR	Field	A.30	BE	GR	GB	PL	CY	FR	PT	DE	DK	LU	FI	CZ	HR	IT	LT	SI	SK
30(1)(a)	Data Protection Officer Contact	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
44–47	Nature of Transfer	×	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.1(f)	Legitimate interests	×	✓	✓	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
6.1(f)	Legitimate interests assessment	×	✓	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
6, 14, 30.1(b)	Data Combination	×	✓	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	Nos. Fields	16	32	31	29	25	23	23	23	22	21	21	19	18	18	18	18	18	17

Appendix B. CSM-ROPA: Mapping with DPV Concepts

The following table summarises the mapping between CSM-ROPA fields and DPV concepts. The column ‘GDPR’ specifies the relevant clause in the GDPR, ‘DPV’ specifies relevant concepts within DPV for expressing field information, ‘Map.’ refers to mapping outcome: *E* indicating Exact mapping, i.e., the concept existed in DPV and could be used as is, *Pt* indicating Partial mapping, i.e., the concept did not exist exactly, but another concept was similar in context, and *S* for indicating the concept did not exist and has been proposed for inclusion. The columns ‘DC’ and ‘DP’ represent the necessity of field for data controllers and data processors, respectively, where *M* indicates Mandatory, i.e., a minimum requirement for ROPA as set out in Article 30 or as required for DPCAT functionality; *C* indicates Conditional, i.e., a minimum requirement for ROPA as set under Article 30 (if applicable); *R* indicates Recommended, i.e., a non-legal requirement for ROPA that assists the organisation in meeting the accountability principle, recommended by DPA guidelines; and *O* indicates Optional, i.e., a term found on a ROPA template that has no legal requirement for inclusion, nor any direct/supplementary role in demonstrating accountability.

Table A2. Mapping of CSM-ROPA fields with DPV Concepts.

GDPR	Field	DPV	Map.	DC	DP
5	Location of personal data	dpv:StorageLocation	E	R	R
5.1	Data Sources	dpv:DataSource	E	R	O
6.1	Legal basis	dpv:LegalBasis	E	M	O
6.1	Link to record of consent	dpv:Consent	E	R	R
9.1	Special Personal Data	dpv:SpecialCategoryPersonalData	E	R	O
9.1	Vulnerable Data Subjects	dpv:VulnerableDataSubject	E	R	O
22.1	Automated decision-making or profiling	dpv:AutomatedDecisionMaking	E	R	R
26.1	Joint Controller agreement	dpv:JointDataControllersAgreement	E	R	N/A
28	Data Processors	dpv:DataProcessor	E	R	M
28.3	Data Processing Contract	dpv:DataProcessingAgreement	E	R	R
28.3	Data processor contract	dpv:ControllerProcessorAgreement	E	R	R
30.1	Status of processing	dpv:Status	S	M	M
32	Tech/Org measures implementation	dpv:Technology	E	R	R
32	Security measures	dpv:TechnicalOrganisationalMeasure	E	R	R
32	Technologies used	dpv:Technology	E	R	R
33.5	Data Breach	dpcat:DataBreachRecord	S	R	R
35	Risk management	dpv:RiskMitigationMeasure	E	R	O
35	Relevant DPIA	dpv:DPIA	E	R	R
35	DPIA Results	dpv:DPIA	E	R	O
36.1	Impact Assessments	dpv:ImpactAssessment	E	R	R
36.1	Prior Consultations	dpv:Consultation	E	R	R

Table A2. Cont.

GDPR	Field	DPV	Map.	DC	DP
37.6	External DPO organisation	dpv:DataProtectionOfficer	E	R	R
–	Name of Business Process	dpv:PersonalDataHandling	Pt	O	O
–	Owner of Process	dct:contactPoint	E	M	M
–	Type of Processing	dpv:Processing	E	O	O
13, 14, 15	Data Subject Rights	dpv:DataSubjectRight	E	R	O
28, 30.1(c)	Data Categories Transfer to Third Parties	dpv:Transfer, dpv:PersonalData	E	R	R
30.1(a)	DPO contact	dpv:hasName, dpv:hasContact	E	MC	MC
30.1(a)	Representative	dpv:Representative	E	MC	N/A
30.1(a)	Representative contact	dpv:hasName, dpv:hasContact	E	MC	N/A
30.1(a)	Name of joint controller	dpv:JointDataController	E	MC	N/A
30.1(a)	Contact details of joint controller	dpv:hasName, dpv:hasContact	E	MC	N/A
30.1(b)	Purposes of processing	dpv: Purpose	E	M	O
30.1(b)	Main/Auxiliary Processing	dpv:Importance (Primary, Secondary)	E	O	O
30.1(c)	Personal Data Categories	dpv:PersonalDataCategory	E	M	M
30.1(c)	Categories of data subjects	dpv:DataSubject	E	M	M
30.1(d)	Categories of Recipients	dpv:Recipient	E	MC	MC
30.1(e)	Third Countries Data Transfer	dpv:ThirdCountry	E	MC	MC
30.1(e)	Appropriate Safeguards	dpv:Safeguard	E	MC	MC
30.1(f)	Retention/Deletion Periods	dpv:StorageDuration,	E	M	O
30.1(g)	Technical and organisational measures	dpv:TechnicalOrganisationalMeasure	E	M	M
30(1)(a)	Data Controller contact	dpv:hasName, dpv:hasContact	E	M	M
30(1)(a)	Data Protection Officer	dpv:DataProtectionOfficer	E	MC	MC
44–47	Nature of Transfer	dpv:DataTransferLegalBasis	E	MC	MC
6.1(f)	Legitimate interests	dpv:LegitimateInterest	E	R	R
6.1(f)	Legitimate interests assessment	dpv:LegitimateInterestAssessment	E	R	R
6, 14, 30.1(b)	Data Combination	dpv:Combine	E	R	O

Appendix C. DPCat Specification

The following tables summarise the *ROPA*, *ROPACatalog*, and *ROPAREcord* fields in the DPCat specification. The ‘Card.’ columns refer to the cardinality of the field, and ‘Nec.’ columns refer to necessity requirements for the fields, where *M* indicates Mandatory; *C* indicates Conditional, i.e., if applicable; *R* indicates Recommended; and *O* indicates Optional.

Table A3. DPCat *ROPA* and *ROPACatalog* fields.

Title	Relation	Domain	Range	Card.	Nec.
Datasets in Catalog	dcat:dataset	dpcat:ROPA(Catalog)	dpcat:ROPAREcord	0..n	M
Description	dct:description	dpcat:ROPA(Catalog)	rdfs:Literal	1..n	M
Issued	dct:issued	dpcat:ROPA(Catalog)	rdfs:Literal (XSD date/time)	0..1	R
Publisher	dct:publisher	dpcat:ROPA(Catalog)	foaf:Agent	1..1	M
Title	dct:title	dpcat:ROPA(Catalog)	rdfs:Literal	1..n	M
Contact Point	dcat:contactPoint	dpcat:ROPA(Catalog)	vcard:Kind	0..n	R
Temporal coverage	dct:temporal	dpcat:ROPA(Catalog)	dct:PeriodOfTime	0..n	O
Data Controller	dpv:hasDataController	dpcat:ROPA(Catalog)	dpv:DataController	1..1	M
DPO for Catalog	dpv:hasDataProtectionOfficer	dpcat:ROPA(Catalog)	dpv:Data ProtectionOfficer	0..1	MC
Representative	dpv:hasRepresentative	dpcat:ROPA(Catalog)	dpv:Representative	0..1	MC
Responsible Entity	dpcat:responsible Entity	dpcat:ROPA(Catalog)	dpv:Entity	0..n	O
Catalogs	dcat:catalog	dcat:ROPACatalog	dpv:ROPA	0..n	M

Table A4. DPCat ROPARecord fields.

Title	Relation	Domain	Range	Card.	Nec.
Contract Point	dcat:contactPoint	dpcat:ROPARecord	vcard:Kind	0...n	R
Description	dct:description	dpcat:ROPARecord	rdfs:Literal	1...n	M
Identifier	dct:identifier	dpcat:ROPARecord	rdfs:Literal	0...n	O
Date Issued	dct:issued	dpcat:ROPARecord	rdfs:Literal (datetime)	0...1	O
Publisher	dct:publisher	dpcat:ROPARecord	foaf:Agent	0...1	R
Temporal coverage	dct:temporal	dpcat:ROPARecord	dct:PeriodOfTime	0...n	R
Title	dct:title	dpcat:ROPARecord	rdfs:Literal	1...n	M
Joint Controller	dpv:hasJointDataControllers	dpcat:ROPARecord	dpv:LegalEntity	0...n	MC
Business Process	dpv:hasPersonalDataHandling	dpcat:ROPARecord	dpv:PersonalDataHandling	0...1	R
Process Owner	dcat:contactPoint	dpcat:ROPARecord	vcard:Kind	0...n	R
Purposes	dpv:hasPurpose	dpcat:ROPARecord	dpv:Purpose	1...n	M
Legal Basis	dpv:hasLegalBasis	dpcat:ROPARecord	dpv:LegalBasis	1...n	M
Type of Processing	dpv:hasProcessing	dpcat:ROPARecord	dpv:Processing	1...n	R
Personal Data	dpv:hasPersonalData	dpcat:ROPARecord	dpv:PersonalData	1...n	M
Special Personal Data Categories	rdfs:subClassOf	dpv:SpecialCategoryPersonaldata	dpv:PersonalData	1...n	R
Data Subjects	dpv:hasDataSubject	dpcat:ROPARecord	dpv:DataSubject	1...n	M
Vulnerable Data Subjects	rdfs:subClassOf	dpv:VulnerableDataSubject	dpv:DataSubject	0...n	R
Data Retention/Deletion Periods	dpv:hasStorage	dpcat:ROPARecord	dpv:StorageDuration	1...n	M
Data Combination	rdfs:subClassOf	dpv:Combine	dpv:Processing	0...n	R
Source of Data	dpv:hasDataSource	dpcat:ROPARecord	dpv:DataSource	1...n	R
Processor	dpv:hasDataProcessor	dpcat:ROPARecord	dpv:LegalEntity	0...n	M
Data Processing Contract	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:DataProcessingAgreement	0...n	R
Recipients	dpv:hasRecipient	dpcat:ROPARecord	dpv:LegalEntity	1...n	MC
Third countries for Transfers	dpv:hasThirdCountry	dpv:Transfer	dpv:ThirdCountry	0...n	MC
Nature of Transfer	dpv:hasLegalBasis	dpv:Transfer	dpv:LegalBasis	0...n	MC
Safeguards	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:Safeguard	0...n	MC
Risk management	dpv:hasRisk, dpv:isMitigatedByMeasure	dpcat:ROPARecord	dpv:Risk, dpv:RiskMitigationMeasure	0...n	R
Technical/Organisational measures	dpv:hasTechnicalOrganisationalMeasure	dpcat:ROPARecord	dpv:TechnicalOrganisationalMeasure	1...n	M
Data Subject Rights	dpv:hasRight	dpcat:ROPARecord	dpv:DataSubjectRight	1...n	R
Legitimate interests	dpv:hasLegalBasis	dpcat:ROPARecord	dpv:LegitimateInterest	0...n	R
Legitimate Interests Assessment	dpv:hasOrganisationalMeasure	dpv:LegitimateInterest	dpv:LegitimateInterestAssessment	0...n	R
Automated decision-making	dpv:hasContext	dpv:Processing	dpv:AutomatedDecisionMaking	0...n	R
Profiling	rdfs:subClassOf	dpv:Profiling	dpv:Processing	0...n	R
Record of Consent	dpv:hasLegalBasis	dpcat:ROPARecord	dpv:Consent	0...n	R
Location of Personal Data	dpv:hasStorage	dpcat:ROPARecord	dpv:StorageLocation	1...n	R
Status of Processing	dpv:hasContext	dpcat:ROPARecord	dpv:Status	1...n	R
Relevant Personal Data Breach	dpcat:associatedWithDataBreach	dpcat:ROPARecord	dpcat:DataBreach	0...n	R
Impact Assessment	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:ImpactAssessment	0...n	R
Prior Consultation	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:Consultation	0...n	R
Main/Auxiliary Processing	dpv:hasContext	dpcat:ROPARecord	dpv:Importance	1...n	R

Table A4. Cont.

Title	Relation	Domain	Range	Card.	Nec.
Joint Controller Agreement	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:JointDataControllers Agreement	0...n	R
Data Processor Contract	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:ControllerProcessor Agreement	0...n	R
Information System for Tech/Org measure	dpv:isImpementedUsingTechnology	dpv: TechnicalOrganisa-tionalMeasure	dpv:Technology	1...n	R
Security Measures	dpv:hasTechnicalOrganisational Measure	dpcat:ROPARecord	dpv:TechnicalOrganisational Measure	1...n	R
Relevant DPIA	dpv:hasOrganisationalMeasure	dpcat:ROPARecord	dpv:DPIA	0...n	R
DPIA Results	dpv:hasOrganisational Measure	dpcat:ROPARecord	dpv:DPIA	0...n	R
System or software name	dpv:isImpementedUsing Technology	dpcat:ROPARecord	dpv:Technology	1...n	R

References

- Castlebridge. Registers of Processing Activities. 2020. Available online: <https://castlebridge.ie/research/2020/ropa-report/> (accessed on 10 March 2022).
- Ryan, P.; Pandit, H.J.; Brennan, R. A Common Semantic Model of the GDPR Register of Processing Activities. In *Frontiers in Artificial Intelligence and Applications*; Villata, S.; Harašta, J.; Křemen, P., Eds.; IOS Press: Amsterdam, The Netherlands, 2020. [CrossRef]
- Ryan, P.; Brennan, R. Demonstrating GDPR Accountability with CSM-ROPA: Extensions to the Data Privacy Vocabulary. In Proceedings of the 24th International Conference Enterprise Information Systems (ICEIS '21), Online, 2021. Available online: <https://www.insticc.org/node/TechnicalProgram/iceis/2021/presentationDetails/103905> (accessed on 10 March 2022).
- Records of Processing and Lawful Basis—ICO. 2022. Available online: <https://ico.org.uk/for-organisations/accountability-framework/records-of-processing-and-lawful-basis/> (accessed on 10 March 2022).
- Ryan, P.; Pandit, H.; Brennan, R. Building a Data Processing Activities Catalog: Representing Heterogeneous Compliance-Related Information for GDPR Using DCAT-AP and DPV. In *Further with Knowledge Graphs*; IOS Press: Amsterdam, The Netherlands, 2021; pp. 169–182. [CrossRef]
- Pandit, H.J.; Debruyne, C.; O’Sullivan, D.; Lewis, D. An Exploration of Data Interoperability for GDPR. *Int. J. Stand. Res. (IJSR)* **2018**, *16*, 1–21. [CrossRef]
- Chiavetta, R. Privacy Tech Vendor Report. Technical Report; International Association of Privacy Professionals (IAPP). 2020. Available online: <https://iapp.org/resources/article/privacy-tech-vendor-report/> (accessed on 10 March 2022).
- OneTrust. IDC Releases First Worldwide Data Privacy Management Software Market Shares Report—OneTrust. 2020. Available online: <https://www.onetrust.com/blog/idc-releases-first-worldwide-data-privacy-management-software-market-shares-report/> (accessed on 10 March 2022).
- Buckley, R.P.; Arner, D.W.; Zetzsche, D.A.; Weber, R.H. The Road to RegTech: The (Astonishing) Example of the European Union. *J. Bank. Regul.* **2020**, *21*, 26–36. [CrossRef]
- Butler, T.; O’Brien, L. Understanding RegTech for Digital Regulatory Compliance. In *Disrupting Finance*; Lynn, T., Mooney, J.G., Rosati, P.; Cummins, M., Eds.; Springer International Publishing: Cham, Switzerland, 2019; pp. 85–102. [CrossRef]
- Labadie, C.; Legner, C. Understanding Data Protection Regulations from a Data Management Perspective: A Capability-Based Approach to EU-GDPR. In Proceedings of the Wirtschaftsinformatik, Siegen, Germany, 24–27 February; p. 15. Available online: <https://aisel.aisnet.org/wi2019/track11/papers/3/> (accessed on 10 March 2022).
- Martínez González, M.M.; Alvite Díez, M.L.; Casanovas, P.; Casellas, N.; Sanz, D.; Aparicio de la Fuente, A. State of the Art and Ambition. Technical Report D3.1; OntoROPA Project. 2021. Available online: <https://ontoropa.wordpress.com/> (accessed on 10 March 2022).
- Huth, D.; Tanakol, A.; Matthes, F. Using Enterprise Architecture Models for Creating the Record of Processing Activities (Art. 30 GDPR). In Proceedings of the 2019 IEEE 23rd International Enterprise Distributed Object Computing Conference (EDOC), Paris, France, 28–31 October 2019; pp. 98–104. [CrossRef]
- Korff, D.; Georges, M. *The Data Protection Officer Handbook*; SSRN Scholarly Paper ID 3428957; Social Science Research Network: Rochester, NY, USA, 2019. Available online: <https://papers.ssrn.com/abstract=3428957> (accessed on 10 March 2022).
- Pandit, H.J.; Polleres, A.; Bos, B.; Brennan, R.; Bruegger, B.; Ekaputra, F.J.; Fernández, J.D.; Hamed, R.G.; Lizar, M.; Schlehahn, E.; et al. Creating A Vocabulary for Data Privacy. In Proceedings of the the 18th International Conference on Ontologies, DataBases, and Applications of Semantics (ODBASE2019), Rhodes, Greece, 22–23 October 2019; p. 17. [CrossRef]
- Data Catalog Vocabulary (DCAT)—Version 2. Available online: <https://www.w3.org/TR/vocab-dcat-2/> (accessed on 10 March 2022).
- Cox, S.; Browning, D.; Albertoni, R.; Beltran, A.G.; Winstanley, P.; Perego, A. Data Catalog Vocabulary (DCAT)—Version 2. W3C Recommendation, W3C. 2020. Available online: <https://www.w3.org/TR/vocab-dcat-2/> (accessed on 10 March 2022).

18. Dekkers, M.; Peristeras, V.; Loutas, N.; Sofou, N.; Van Nuffelen, B. DCAT Application Profile for Data Portals in Europe Version 2.1.0. Semantic Interoperability Community (SEMIC), Directorate-General for Informatics (DIGIT), European Commission. 2019. Available online: <https://op.europa.eu/en/web/eu-vocabularies/dcat-ap> (accessed on 10 March 2022).
19. Strategy for Data | Shaping Europe's Digital Future. 2022. Available online: <https://digital-strategy.ec.europa.eu/en/policies/strategy-data> (accessed on 10 March 2022).
20. Measuring Privacy Operations 2019 Cookies, Local vs. Global Compliance, DSARs and More - IAPP and TrustArc. 2020. Available online: https://iapp.org/media/pdf/resource_center/measuring_privacy_operations_2019.pdf (accessed on 10 March 2022).
21. Bracy, J. Privacy Tech Vendor Report. Technical Report; International Association of Privacy Professionals (IAPP). 2021. Available online: <https://iapp.org/resources/article/privacy-tech-vendor-report/> (accessed on 10 March 2022).
22. The Value of Investing in Well-Constructed Records of Processing Activities—IAPP. Available online: <https://iapp.org/news/a/the-value-of-investing-in-well-constructed-recordings-of-processing-activities/> (accessed on 10 March 2022).
23. Drewer, D.; Miladinova, V. The Canary in the Data Mine. *Comput. Law Secur. Rev.* **2018**, *34*, 806–815. [CrossRef]
24. Sparapani, T.; Sherman, J. Privacy Tech's Third Generation A Review of the Emerging Privacy Tech Sector. 2021. Available online: https://fpf.org/wp-content/uploads/2021/06/FPF-PTA-Report_Digital.pdf (accessed on 10 March 2022).
25. Khatri, V.; Brown, C.V. Designing Data Governance. *Commun. Acm* **2010**, *53*, 148–152. [CrossRef]
26. Rozehnal, P.; Novák, V. The Core Of Enterprise Architecture As A Management Tool: Gdpr Implementation Case Study. In Proceedings of the 26th Interdisciplinary Information Management Talks, Kutná Hora, Czech Republic, 5–7 September 2020. Available online: http://idimt.org/wp-content/uploads/proceedings/IDIMT_proceedings_2018.pdf (accessed on 10 March 2022).
27. Burmeister, F.; Drews, P.; Schirmer, I. A Privacy-driven Enterprise Architecture Meta-Model for Supporting Compliance with the General Data Protection Regulation. In Proceedings of the Hawaii International Conference on System Sciences 2019 (HICSS-52), Maui, HI, USA, 8–11 January 2019. Available online: <http://hdl.handle.net/10125/60040> (accessed on 10 March 2022).
28. Pandit, H.J.; Fatema, K.; O'Sullivan, D.; Lewis, D. GDPRtEXT—GDPR as a Linked Data Resource. In *The Semantic Web*; Gangemi, A., Navigli, R., Vidal, M.E., Hitzler, P., Troncy, R., Hollink, L., Tordai, A., Alam, M., Eds.; Springer International Publishing: Cham, Switzerland, 2018; Volume 10843, pp. 481–495. [CrossRef]
29. Business Process Re-Engineering and Functional Toolkit for GDPR Compliance (BPR4GDPR H2020 Project). Available online: <https://www.bpr4gdpr.eu/> (accessed on 10 March 2022).
30. Pandit, H.J.; Lewis, D. Modelling Provenance for GDPR Compliance Using Linked Open Data Vocabularies. In Proceedings of the 5th Workshop on Society, Privacy and the Semantic Web—Policy and Technology (PrivOn2017) (PrivOn), Vienna, Austria, 22 October 2017. Available online: http://ceur-ws.org/Vol-1951/PrivOn2017_paper_6.pdf (accessed on 10 March 2022).
31. Pandit, H.J.; Debruyne, C.; O'Sullivan, D.; Lewis, D. GConsent—A Consent Ontology Based on the GDPR. In *The Semantic Web*; Hitzler, P.; Fernández, M., Janowicz, K., Zaveri, A., Gray, A.J., Lopez, V., Haller, A., Hammar, K., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2019; pp. 270–282. [CrossRef]
32. Bonatti, P.A.; Kिरrane, S.; Petrova, I.M.; Sauro, L. Machine Understandable Policies and GDPR Compliance Checking. *KI Künstliche Intell.* **2020**, *34*, 303–315. [CrossRef]
33. Palmirani, M.; Martoni, M.; Rossi, A.; Bartolini, C.; Robaldo, L. PrOnto: Privacy Ontology for Legal Compliance. In Proceedings of the 18th European Conference on Digital Government ECDG 2018, Santiago, Spain, 25–26 October 2018; p. 10. Available online: <https://ora.uniurb.it/handle/11576/2691050> (accessed on 10 March 2022).
34. Esteves, B.; Rodriguez-Doncel, V. Analysis of Ontologies and Policy Languages to Represent Information Flows in GDPR. *Semantic Web J.* **2022**, *Forthcoming*. Available online: <http://www.semantic-web-journal.net/content/analysis-ontologies-and-policy-languages-represent-information-flows-gdpr-1> (accessed on 10 March 2022).
35. Lioudakis, G.V.; Koukovini, M.N.; Papagiannakopoulou, E.I.; Dellas, N.; Kalaboukas, K.; de Carvalho, R.M.; Hassani, M.; Bracciale, L.; Bianchi, G.; Juan-Verdejo, A.; et al. Facilitating GDPR Compliance: The H2020 BPR4GDPR Approach. In *Digital Transformation for a Sustainable Society in the 21st Century*; Pappas, I.O., Mikalef, P., Dwivedi, Y.K., Jaccheri, L., Krogstie, J., Mäntymäki, M., Eds.; IFIP Advances in Information and Communication Technology; Springer International Publishing: Cham, Switzerland, 2020; pp. 72–78. [CrossRef]
36. Grunewald, E.; Wille, P.; Pallas, F.; Borges, M.C.; Ulbricht, M.R. TIRA: An OpenAPI Extension and Toolbox for GDPR Transparency in RESTful Architectures. In Proceedings of the International Workshop on Privacy Engineering (IWPE'21), Vienna, Austria, 6–10 September 2021. [CrossRef]
37. Grunewald, E.; Pallas, F. TILT: A GDPR-Aligned Transparency Information Language and Toolkit for Practical Privacy Engineering. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event, 3–10 March 2021; pp. 636–646. [CrossRef]
38. Rost, M.; Weichert, R. The Standard Data Protection Model. Technical Report Version 2.0b, Conference of the Independent Data Protection Supervisory Authorities of the Federation and the Länder, 2020. Available online: <https://www.datenschutz-mv.de/datenschutz/datenschutzmodell/> (accessed on 10 March 2022).
39. Scharffe, F.; Fensel, D. Correspondence Patterns for Ontology Alignment. In Proceedings of the International Conference on Knowledge Engineering and Knowledge Management, Sicily, Italy, 29 September–3 October 2008; pp. 83–92. [CrossRef]

40. Brous, P.; Janssen, M.; Krans, R. Data Governance as Success Factor for Data Science. In *Responsible Design, Implementation and Use of Information and Communication Technology*; Hattingh, M., Matthee, M., Smuts, H., Pappas, I., Dwivedi, Y.K., Mäntymäki, M., Eds.; Springer International Publishing: Cham, Switzerland, 2020; pp. 431–442. [CrossRef]
41. Records Register | European Data Protection Supervisor. 2022 Available online: https://edps.europa.eu/about/data-protection-within-edps/records-register_en (accessed on 10 March 2022).
42. Apache Jena Version 4.4.0. 2022. Available online: <https://jena.apache.org/> (accessed on 10 March 2022).
43. GraphDB Free Version 9.7.0-1. 2021. Available online: <https://graphdb.ontotext.com/> (accessed on 10 March 2022).
44. TopBraid SHACL API Version 1.3.2. TopQuadrant, Inc. 2022. Available online: <https://github.com/TopQuadrant/shacl> (accessed on 10 March 2022).
45. Pandit, H.J.; O’Sullivan, D.; Lewis, D. Test-Driven Approach Towards GDPR Compliance. In Proceedings of the 15th International Conference on Semantic Systems (SEMANTiCS2019), Karlsruhe, Germany, 9–12 September 2019. [CrossRef]
46. Coordinated Plan on Artificial Intelligence 2021 Review | Shaping Europe’s Digital Future. 2021. Available online: <https://digital-strategy.ec.europa.eu/en/library/coordinated-plan-artificial-intelligence-2021-review> (accessed on 10 March 2022).
47. CKAN—The Open Source Data Management System. Available online: <http://ckan.org/> (accessed on 10 March 2022).
48. Ryan, P.; Crane, M.; Brennan, R. Design Challenges for GDPR RegTech. In Proceedings of the 22nd International Conference on Enterprise Information Systems, Prague, Czech Republic, 5–7 May 2020; pp. 787–795. [CrossRef]
49. Arner, D.W.; Barberis, J.N.; Buckley, R.P. The Evolution of Fintech: A New Post-Crisis Paradigm? *SSRN Electron. J.* **2015**. [CrossRef]
50. Privacy Impact Assessment (PIA) | CNIL). 2021. Available online: <https://www.cnil.fr/en/privacy-impact-assessment-pia> (accessed on 10 March 2022).
51. DPV-GDPR: GDPR Extension for DPV. 2022. Available online: <https://w3c.github.io/dpv/dpv-gdpr/> (accessed on 10 March 2022).
52. Taylor, C. GDPR at Risk of Failing Due to Underfunding of Regulators, Study Finds. 2020. Available online: <https://www.irishtimes.com/business/technology/gdpr-at-risk-of-failing-due-to-underfunding-of-regulators-study-finds-1.4238927> (accessed on 10 March 2022).