

## Article

# RDF Stream Taxonomy: Systematizing RDF Stream Types in Research and Practice

Piotr Sowiński <sup>1,2,\*</sup> , Paweł Szmeja <sup>2</sup> , Maria Ganzha <sup>1,2</sup>  and Marcin Paprzycki <sup>2</sup> 

- <sup>1</sup> Faculty of Mathematics and Information Science, Warsaw University of Technology, ul. Koszykowa 75, 00-662 Warsaw, Poland
- <sup>2</sup> Systems Research Institute, Polish Academy of Sciences, ul. Newelska 6, 01-447 Warsaw, Poland; pawel.szmeja@ibspan.waw.pl (P.S.); maria.ganzha@ibspan.waw.pl (M.G.); paprzyck@ibspan.waw.pl (M.P.)
- \* Correspondence: piotr.sowinski@ibspan.waw.pl

**Abstract:** Over the years, RDF streaming has been explored in research and practice from many angles, resulting in a wide range of RDF stream definitions. This variety presents a major challenge in discussing and integrating streaming systems due to a lack of a common language. This work attempts to address this critical research gap by systematizing RDF stream types present in the literature in a novel taxonomy. The proposed RDF Stream Taxonomy (RDF-STaX) is embodied in an OWL 2 DL ontology that follows the FAIR principles, making it readily applicable in practice. Extensive documentation and additional resources are provided to foster the adoption of the ontology. Three use cases for the ontology are presented with accompanying competency questions, demonstrating the usefulness of the resource. Additionally, this work introduces a novel nanopublications dataset, which serves as a collaborative, living state-of-the-art review of RDF streaming. The results of a multifaceted evaluation of the resource are presented, testing its logical validity, use case coverage, and adherence to the community's best practices, while also comparing it to other works. RDF-STaX is expected to help drive innovation in RDF streaming by fostering scientific discussion, cooperation, and tool interoperability.

**Keywords:** RDF stream; taxonomy; ontology; interoperability; SKOS; nanopublications



**Citation:** Sowiński, P.; Szmeja, P.; Ganzha, M.; Paprzycki, M. RDF Stream Taxonomy: Systematizing RDF Stream Types in Research and Practice. *Electronics* **2024**, *13*, 2558. <https://doi.org/10.3390/electronics13132558>

Academic Editors: Stefano Ferilli and Domenico Redavid

Received: 11 May 2024  
Revised: 22 June 2024  
Accepted: 27 June 2024  
Published: 29 June 2024



**Copyright:** © 2024 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

The Resource Description Framework (RDF) is a standard-backed foundation for modeling and exchanging knowledge graphs, with a wide range of practical applications [1,2]. The RDF model is based on the building blocks of the Web (e.g., it uses Internationalized Resource Identifiers for node names), making it very flexible and extensible [3]. The list of use cases for RDF has expanded greatly since its inception, with time moving beyond applications where the modeled knowledge is treated as a single, static dataset. More recently, the concept of RDF streams has gained traction in research and industrial applications. Generally, an RDF stream refers to a potentially infinite stream of data modeled using the RDF standard to which stream processing principles apply [4]. This research interest resulted in numerous contributions to the field, such as streaming serializations, querying and reasoning engines, programming frameworks, datasets, and other resources [5–8]. Without a doubt, streaming RDF solutions are in demand, and more innovation is still needed. However, the landscape of RDF streaming terminology is fragmented, with many dissimilar streaming task formulations and incompatible solutions.

The term “RDF stream” is used in the literature to refer to many different concepts, ranging from theoretical models [9] to very technical ones, focusing on the representation and practical usage [6,10,11]. As noted in past research works [12–14], RDF streaming systems need more standardization in stream exchange protocols and stream task definitions. This would be beneficial not only in the academic context but also in the facilitation of the distribution and reuse of RDF streams in practice. Additionally, it could enable streaming

tool interoperability through the inclusion of appropriate metadata. At the same time, if the research is to solve real-world problems, the variety of existing RDF streaming task formulations must be embraced, not limiting oneself to one specific type of RDF stream. However, up until now, there was very little academic discussion about the existing RDF stream types—what stream types are used, how they relate to each other, and which use cases are they applicable to. Our work aims to address this critical research gap, which has a large impact on the uptake of RDF streaming in general.

The main contribution of this work is the RDF Stream Taxonomy (RDF-STaX), embodied in an OWL 2 DL ontology. This novel taxonomy systematizes the existing RDF stream types on the basis of a broad state-of-the-art review. RDF-STaX aims to be a practical and useful resource for researchers and practitioners, enabling them to semantically annotate RDF stream types in publications, software, and streams published on the Web. To achieve this, RDF-STaX follows the community's best practices for ontology publishing and is accompanied by extensive documentation, including usage examples. The three envisioned use cases for RDF-STaX are presented with accompanying competency questions implemented in SPARQL. Two of these use cases were already implemented in practice—a living state-of-the-art review of RDF streaming using nanopublications and annotating stream types in RiverBench datasets. Finally, RDF-STaX is compared with other ontologies, highlighting the gap that it covers.

This work is organized as follows. Section 2 presents a comprehensive literature review of the topic. Section 3 introduces the proposed systematization of RDF streams, along with an alignment of these definitions to the state of the art. Section 4 describes the proposed RDF-STaX ontology. Section 5 presents the use cases envisioned for RDF-STaX, while Section 6 includes a multifaceted evaluation of the proposed ontology, including a comparison with other ontologies. Finally, Section 7 discusses the results and outlines future work directions, with Section 8 presenting the concluding remarks.

## 2. Literature Review

The primary goal of this paper is to introduce a taxonomy of RDF stream types, which naturally must be grounded in the state of the art to be relevant. As no comprehensive and up-to-date review of this topic is available, here, we attempt to summarize the most significant works that use the concept of RDF streams.

The scope of this study only includes RDF streams that can be fully represented using the RDF 1.1 W3C Recommendation [3]. This notably excludes theoretical stream models that do not have a direct translation to RDF, such as RSP-QL [9]. However, due to the importance of such models, they are briefly discussed to provide additional context. The reason for not considering purely theoretical models is that they are not directly comparable to models that have a fixed representation in RDF 1.1, as they operate on a different level of abstraction. Therefore, including both purely theoretical and fully representable RDF streams in the review would make the resulting taxonomy methodologically flawed. This scope decision is additionally motivated by the fact that only streams fully representable in RDF can be easily exchanged and published on the Web using already existing protocols, while the theoretical models have more limited applicability, for example, only within one specific system (e.g., a stream reasoning engine). This choice is thus natural due to this work's focus on solving the challenge of streaming tool interoperability and harmonizing the ways in which RDF streams are exchanged.

To cover the various areas of research related to RDF streaming, several survey papers and books were included in this review [5,7,8,15–21] on topics such as RDF data management, linked streaming data, and stream reasoning. For many of the examined works, determining the exact nature of the stream was not trivial due to unclear descriptions and a lack of standardized terminology. Therefore, not only were the papers examined but the code and documentation were too where available.

### 2.1. RDF Streaming Protocols

One prominent area of research that uses RDF streams is work focusing on streaming protocols for transmitting RDF data. One of the earlier examples is Streaming HDT (S-HDT) [22], which facilitates streaming RDF data in IoT devices. More specifically, S-HDT producers (IoT devices) emit streams of compressed, unnamed RDF graphs. RDSZ [23] is a later method that also attempts to stream compressed RDF graphs over the network. ERI [10] is a similar approach; however, it has several views of the RDF stream. At the highest level, ERI streams a continuous sequence of triples that is then divided into discrete blocks (RDF graphs). Then, on the lowest level, these graphs are decomposed into *subject-molecules* (graphs with a single RDF subject) that are then streamed over the network. Therefore, ERI has three different views of the streaming problem, with each corresponding to a different level in the protocol.

The RDF EXI communication protocol for constrained devices [24] focuses on streaming RDF graphs in a very efficient manner. Finally, the recent Jelly protocol [6] treats RDF streams similarly to ERI, streaming a sequence of RDF statements divided into discrete elements (stream frames). The main difference is that Jelly can also stream quad statements, making it generalizable to RDF datasets. It should be noted here that none of these works assume time annotations to be a necessary part of RDF streams, in contrast to the works described in the following subsection.

### 2.2. Stream Processing and Reasoning

There is a large body of research on the topics of RDF stream reasoning and stream processing, where a special focus is often placed on the time annotations of stream elements. We could find only one work in this category that did not consider the temporal aspect, an early proposal for a streaming SPARQL engine from 2007 by Groppe et al. [25]. The proposed engine executes queries over streams of RDF triples, in real time.

Most publications in this category focus on the inner workings of querying/reasoning engines and adopt task formulations that are not fully grounded in the RDF specification. The most popular such formulation is the stream of timestamped triples [9,21,26–33], where a stream is a sequence of  $\langle triple, timestamp \rangle$  tuples. There are many variations of this task, with time intervals instead of timestamps [26], multiple timestamps per triple [33,34], and different time domains. However, this model (and other equivalent or similar models, e.g., time-varying graphs [12,33,35]) by itself is not enough to fully realize and exchange an RDF stream—this would require encoding the time information in RDF, which is often omitted in these works.

The lack of grounding in the RDF specification of the timestamped triple model is a known problem [29], and solutions to it have been proposed. The most prominent formulation is the stream of timestamped named graphs, where each stream element is a named RDF graph. The timestamp (or other metadata) of this graph is added by making statements in the default graph about the named graph's name node [36–38]. This representation was possibly first proposed by Tappolet and Bernstein in 2009 [39] and was then expanded by the W3C RSP Community Group, which created a draft specification for the RSP Data Model [40]. The RSP Data Model uses the aforementioned formulation to define RDF streams but also mentions that RSP engines may consume streams of graphs, datasets, or named graphs and then turn them into RDF streams proper by associating each element with a timestamp. The RSP Data Model is the most mature and best-grounded model for representing timestamped streams, having been drafted by the community, and even proposing the formal semantics of RDF streams. However, it was never made into a W3C Recommendation, and the RSP Community Group at the time of this work's submission appears to be inactive: <https://lists.w3.org/Archives/Public/public-rsp/2023Mar/0000.html> (accessed on 4 June 2024).

A similar representation for timestamped streams was proposed in 2010 by Barbieri and Della Valle [41]. In the proposal, stream elements are named RDF graphs (called instantaneous graphs). These graphs are linked together with additional metadata placed in an

additional named graph dedicated to this purpose (called a stream graph). Each instantaneous graph's name node is associated with a timestamp triple using the `sld:receivedAt` property. This formulation differs from the RSP Data Model mostly by placing the metadata in a named graph instead of the default graph.

Time-Annotated RDF (TA-RDF) [42] is an alternative RDF representation for the timestamped stream problem (more precisely, TA-RDF has a direct translation to RDF). It identifies each stream element with a URI node in an RDF graph. The data relevant to that stream element are then attached to the node. This boils down to each stream element being an RDF graph in which a single node serves as the “point of entry”. This model is notably less flexible than the RSP Data Model, as it requires the entire content of each stream element to be connected to this single node. A similar approach is used by Stream Containers [14], a framework for publishing RDF streams on the Web. There, a stream is a sequence of unnamed RDF graphs that are referenced by URIs. This is not the same as RDF named graphs, as each URI instead points to a node (the element's subject) within one of the unnamed graphs, and the temporal information is also contained within this graph.

One auxiliary RDF stream definition also appears in this category. The VoCaLS vocabulary for describing linked streams [43] defines an RDF stream as a potentially infinite sequence of RDF graphs and/or triples. This definition is used in practice in the RSP4J API [33], whose YASPER implementation can consume streams of RDF graphs or RDF triples.

### 2.3. Streaming Semantic Annotation and Translation

Several systems were developed for streaming semantic annotation (lifting) and translation. RMLStreamer [11,44] annotates data using the RML language [45], outputting a stream of RDF datasets. CARML [46] is a Java library that provides generic facilities for annotating data using RML. It supports outputting a stream of RDF quad statements (using RDF4J Rio [47]), which corresponds to a single RDF dataset. CARML does not natively support streams of datasets, and thus it operates on a different abstraction level than RMLStreamer. This limitation is addressed by the Semantic Annotation enabler [48] from the ASSIST-IoT project [49], which outputs a stream of RDF datasets, each corresponding to one input document. SPARQL-Generate [50] is an extension of the SPARQL language, which allows for generating RDF from various data sources. Its reference implementation <https://github.com/sparql-generate/sparql-generate> (accessed on 4 June 2024) supports outputting a stream of RDF triples (using Apache Jena RIOT [51]).

Last in this category is the Inter-Platform Semantic Mediator (IPSM) [52], which translates a stream of RDF graphs into another such stream, in real time. Although on a logical level IPSM processes RDF graphs, it physically outputs a stream of RDF datasets, where each dataset has two named graphs—one containing the translated RDF and the other with metadata about the translation itself.

### 2.4. Semantic Streaming Applications

Over the years, RDF streams have been employed in various real-life use cases. One such example is DBpedia-Live [53]—a now-defunct service that published changes to the DBpedia knowledge base in real time. Two streams were published—deleted and added triples. In each stream, the elements were RDF graphs in the form of N-Triples files.

RDF streams are also commonly used in semantic sensor networks [16]. In this context, the Graph of Things (GoT) [54] was proposed, which uses streams of RDF graphs. Another common use case is monitoring social media activity. Here, a survey by Keskiärrkkä and Blomqvist [15] noted that RDF streams should have RDF graphs as elements, as this approach better addresses the needs of their use case. However, they also acknowledge that RDF streams can sometimes have RDF triples as elements. Streams are also used for sending updates between decentralized social network servers, as in the ActivityPub protocol [55] standardized by the W3C. In ActivityPub, an item of activity is described with an RDF graph (in JSON-LD). These items form ordered streams of RDF graphs which are

then exchanged between the nodes. The Linked Data Event Streams (LDES) [56] proposal addresses publishing streams composed of elements, where each element is a fragment of an RDF graph, identified by a specific subject node. Finally, the recent proposal for a live open scientific knowledge graph [57] uses streams of RDF graphs to monitor scientific and social activity in real time.

### 2.5. Streaming I/O

The last broad category of RDF streaming solutions are those that process simple sequences of RDF statements (triples or quads) for input/output (I/O) operations. Two examples are the RDF4J Rio toolkit [47] and Apache Jena's RIOT [51], both of which can parse and serialize streams of statements, reading from or writing to a Java byte stream. The two most obvious serializations for this task are N-Triples and N-Quads; however, Jena also supports *streamed block formats* for Turtle and Trig, where statements are transmitted in small batches (blocks).

### 2.6. Summary

The review uncovered a wide variety of task formulations for RDF streams. It should be noted that most of the identified works use the exact same term (RDF stream) to describe very different concepts. The review also highlighted an important distinction between theoretical models of RDF streams (e.g., timestamped triples) and practical ones, making it clear that direct comparisons between the two are not possible. This distinction was also observed earlier in several works [5,14,29,34]. Another contrasting feature is the use of graphs or datasets as stream elements, versus singular triple or quad statements. This was also noticed in earlier works focusing on the practical aspects of implementing RDF streams [5,15].

Worth noting is that some works have several views of the same RDF stream. For example, ERI [10] views its RDF streams as a flow of triples, blocks (RDF graphs), or subject–molecules (RDF subgraphs), depending on the level of abstraction. This suggests that describing the type of an RDF stream depends on the context in which it is considered, or, metaphorically, one's perspective. In fact, an N-Triples file would be considered by RDF4J's Rio [47] as a stream of triples, while an application like DBpedia-Live [53] or IPSM [52] would see it as a single element in a stream of graphs.

The diversity of problem statements encountered in the literature clearly suggests that expecting there to be a single model of RDF streams may be unrealistic, especially given that each stream may be understood in several different ways, depending on the perspective or context. Each of the identified approaches is motivated by some use case and thus is valid. Therefore, we must seek solutions to better understand and harmonize this task diversity, which is the primary goal of RDF-STaX.

## 3. Proposed Systematization

In what follows, we propose a set of definitions for the types of RDF streams present in the literature and current practice. Then, Section 3.2 organizes these definitions in the RDF Stream Taxonomy, and Section 3.3 links the proposed stream types to the state-of-the-art review performed in the previous section. The basis for all following definitions is the RDF 1.1 W3C Recommendation [3], which formally defines terms such as RDF graphs and RDF triples, and must be implemented by any RDF 1.1-compliant software. We use it with the aim of making our definitions readily applicable in the context of existing, RDF 1.1-compliant software.

The model-theoretic semantics of RDF streams are not considered in this work. Instead, the proposed semiformal definitions can be used as a universal basis on top of which the formal semantics can be defined, as needed. This is motivated by, firstly, the semantics of RDF streams being entirely irrelevant for some of the discussed application areas (e.g., streaming protocols, streaming I/O). Secondly, only the semantics of RDF graphs are well-defined—there are no agreed-upon formal semantics for RDF datasets [58], and there are

similarly many ways to approach this topic for RDF streams. Therefore, we leave this topic for future discussion. The definitions are formulated in a semiformal manner similar to the RDF 1.1 Concepts and Abstract Syntax W3C Recommendation [3] or the RSP Data Model [40], with the goal of making them readily applicable for RDF practitioners.

### 3.1. RDF Stream Definitions

We start by defining a universal, abstract RDF stream that will serve as a common denominator for further discussion:

**Definition 1.** *An RDF stream is an ordered, potentially infinite sequence of RDF stream elements.*

The manner in which the sequence is ordered is not specified on purpose, as this depends on the use case. Often, the order will be defined by time, but a stream does not have to be time-dependent (in fact, many methods do not consider time at all—see Section 2). This definition also does not specify what exactly are the RDF stream elements, as this is left to the further definitions that are based on it. Thus, we specialize the definition further:

**Definition 2.** *A grouped RDF stream is an RDF stream whose elements are either RDF graphs or RDF datasets.*

**Definition 3.** *A flat RDF stream is an RDF stream whose elements are statements (either RDF triples or RDF quads).*

Grouped and flat RDF streams are still abstract, in the sense that their definitions are not concretized enough for practical implementation. However, they correspond to an important distinction between streams of groups of statements and streams of individual statements. With this, we can define the concrete types of RDF streams which can be used in practice:

**Definition 4.** *An RDF graph stream is a grouped RDF stream whose elements are unnamed (default) RDF graphs.*

In other words, an element of an RDF graph stream is a set of triples. This is a commonly used formulation that can be further specialized:

**Definition 5.** *An RDF subject graph stream is an RDF graph stream in which every element contains an IRI node (called the subject node) that uniquely identifies the graph in the stream. Every other node in the graph can be reached by traversing triples, starting from the subject node.*

RDF subject graph streams or similar formulations were proposed to handle time-stamped RDF streams (TA-RDF [42] and Stream Containers [14]). ERI [10] uses a similar formulation for subject–molecules, but its subject IRIs are unique only within one block. The proposed definition is slightly broader than the timestamped models, while still being useful—the advantage is that the element can be identified by an IRI, without using RDF datasets. The rest of the element can be discovered with what can be intuitively thought of as RDF graph connectivity.

We can then move on to streams of RDF datasets:

**Definition 6.** *An RDF dataset stream is a grouped RDF stream whose elements are RDF datasets.*

It should be noted here that although an RDF graph stream may seem like a special case of an RDF dataset stream, it is not. According to RDF 1.1, an RDF dataset is not a generalization of an RDF graph but rather a collection of RDF graphs. Thus, the definitions of RDF dataset streams and RDF graph streams are not directly related to the taxonomy. It

should be noted, however, that an RDF graph stream can be trivially transformed into an RDF dataset stream, by simply assuming that each graph in the former is the default graph in the dataset.

This definition can be further restricted to one named graph per element:

**Definition 7.** *An RDF named graph stream is an RDF dataset stream in which every element has exactly one named RDF graph pair  $\langle n, G \rangle$ , where  $G$  is an RDF graph, and  $n$  is the graph name. Apart from graph  $G$ , the dataset may contain any number of triples in the default graph.*

This is a simplified version of the RSP Data Model [40], without the temporal aspect. Although we could find only one reference to this stream type in the literature (RSP engines can consume such streams if an appropriate conversion is implemented [40]), it is an intuitively important specialization of the RDF dataset stream. Here, every element corresponds to one named RDF graph with some optional information (metadata) about it in the default graph. As with RDF subject graph streams, this has the advantage of identifying a stream element by IRI, albeit with more flexibility as to the element's contents.

With this, we can move on to the timestamped variant:

**Definition 8.** *A timestamped named graph is an RDF dataset in which:*

- (1) *there is exactly one named RDF graph pair  $\langle n, G \rangle$ , where  $G$  is an RDF graph, and  $n$  is the graph name;*
- (2) *the default graph includes a timestamp triple  $\langle n, p, t \rangle$ , where  $p$  is a timestamp predicate that relates  $t$ , called the timestamp, and the graph  $G$ .*

**Definition 9.** *A timestamped RDF named graph stream is an RDF named graph stream in which every element is a timestamped named graph. The elements that share the same timestamp predicate  $p$  are ordered by the partial order associated with  $p$ .*

Definitions 8 and 9 were derived directly from the draft RSP Data Model [40]. Finally, we define concrete flat RDF stream types:

**Definition 10.** *A flat RDF triple stream is a flat RDF stream whose elements are triples.*

**Definition 11.** *A flat RDF quad stream is a flat RDF stream whose elements are quads.*

An “RDF quad” is a term that is often used by practitioners but not defined in the RDF 1.1 Recommendation. It is, however, defined in the RDF 1.2 Working Draft (16 April 2024) [59], and thus, we use the RDF 1.2 draft definition of quads here.

Flat RDF streams can be represented simply as N-Triples or N-Quads files; however, contrary to RDF graphs and datasets, they have a defined order of statements. Similarly to RDF graph and dataset streams, a triple stream is not a subclass of a quad stream, because a triple is not a special case of a quad. However, a trivial transformation from one to the other can be performed, by stating explicitly that the triples belong to the default graph.

Grouped and flat RDF streams are notably interrelated, as any RDF graph stream can be flattened into an RDF triple stream by simply listing and concatenating the statements in each of its elements—the same is true for RDF dataset streams and RDF quad streams. The reverse is also possible, the elements of a flat stream can be grouped into a grouped stream. Such operations are in fact common in streaming applications [5,6,10,15].

Finally, for all of the above definitions, their generalized variants can be trivially derived (e.g., generalized RDF graph stream) by using generalized triples and generalized datasets as their basis. We skip these definitions for the sake of brevity.

### 3.2. RDF Stream Taxonomy

The proposed definitions form a taxonomical structure of task formulations (RDF Stream Taxonomy—RDF-STaX), summarized in Figure 1. This overview allows us to

observe that each concrete RDF stream type uses one of the four element types, graph, dataset, triple, and quad, possibly with additional restrictions on them. In fact, these four types of elements seem to be the only conceivable ones that can be derived from the RDF 1.1 specification. For each of them, representing the stream element is a known and solved problem (RDF serializations), and thus, they can be easily used in practical applications. The abstract stream types are not directly usable in practice by themselves but rather serve as the basis of the taxonomy and a shared starting point for other definitions.

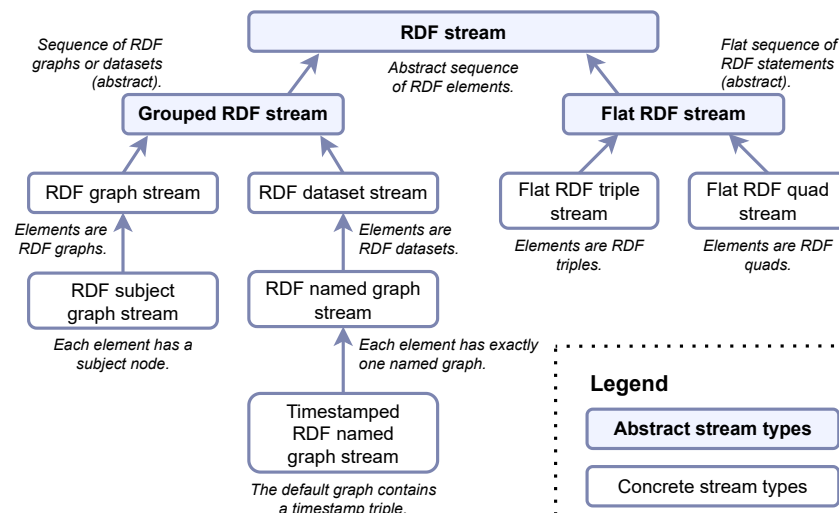


Figure 1. Overview of the RDF Stream Taxonomy (RDF-STaX).

### 3.3. Taxonomy Correspondence to the State of the Art

Finally, we link the proposed RDF stream types to the definitions used in state-of-the-art research and software—Table 1 presents such a summary. As noted in Section 2.6, a single RDF stream can be viewed from multiple perspectives, which is reflected in the table, with some items appearing more than once when such a situation occurs. In some cases, the correspondence is not ideal and there may be insignificant differences between the definitions proposed here and those in the literature. The short descriptions presented in the table are expanded upon in the nanopublications dataset, as described in Section 5.3.

Table 1. RDF stream type usage in research and software, according to RDF-STaX.

Stream type	Uses
RDF graph stream	ActivityPub [55], DBpedia-Live [53], ERI (blocks) [10], GoT [54], IPSM (logical level) [52], Jelly (stream frames) [6], Keskiärrkkä and Blomqvist [15], Live open scientific knowledge graphs [57], RDF EXI [24], RDSZ [23], RSP Data Model (can be consumed) † [40], RSP4J YASPER (can be consumed) [33], S-HDT [22], VoCaLS [43]
RDF subject graph stream	ERI (subject–molecule stream, only within a block) [10], LDES [56], TA-RDF [42], Stream Containers [14]
RDF dataset stream	IPSM (output) [52], Jelly (stream frames) [6], RMLStreamer [11,44], RSP Data Model (can be consumed) † [40], Semantic Annotation enabler [48]
RDF named graph stream	RSP Data Model (can be consumed) † [40]
Timestamped RDF named graph stream	Barbieri and Della Valle ‡ [41], Linked Data Notifications for RDF streams [36], RSP Data Model [40], Tappolet and Bernstein [39], TripleWave [37], Wu et al. [38]
Flat RDF triple stream	Apache Jena RIOT [51], ERI (high level) [10], Groppe et al. [25], Jelly (high level) [6], Keskiärrkkä and Blomqvist (mentioned) [15], RDF4J Rio [47], RSP4J YASPER (can consume) [33], SPARQL-Generate [50], VoCaLS [43]
Flat RDF quad stream	Apache Jena RIOT [51], CARMML [46], Jelly (high level) [6], RDF4J Rio [47]

† The RSP Data Model states that these stream types can be consumed by RSP engines by adding timestamps to each stream element. This is an optional feature of the RSP Data Model. ‡ The used representation differs from the definition in RDF-STaX by placing the metadata of the elements not in the default graph of the RDF dataset but in a dedicated named graph.



Different types of RDF streams appear to be applicable in different settings. In the context of network streaming applications, the most commonly used are grouped RDF streams. Flat RDF streams, on the other hand, are used for streaming I/O and internally by various applications.

#### 4. RDF-STaX Ontology

To make the proposed taxonomy readily applicable, it was modeled as an OWL 2 DL ontology (Figure 2), available at <https://w3id.org/stax/1.1.1/ontology> (accessed on 4 June 2024) under the Creative Commons Attribution 4.0 license. In the ontology, RDF stream types are instances of either the `stax:AbstractStreamType` or the `stax:ConcreteStreamType` class, with taxonomical relations being realized using SKOS in the OWL DL version [60,61]. The lightest ontology profile that could be used for RDF-STaX is the unofficial OWL 2 DL, due to RDF-STaX being based on SKOS. This balances usefulness in the contexts where ontological reasoning is used (by being compatible with a number of reasoners [62]), with uses where ease of annotation takes precedence over semantics (e.g., in stream metadata or knowledge graphs with very little enforced semantics).

Each stream type instance in the RDF-STaX ontology is accompanied by a label, description, formal definition, and usage example. The instances are not shown in Figure 2, but the structure they form is identical to the one presented in the taxonomy overview diagram (Figure 1), realized using the `skos:broader` property.

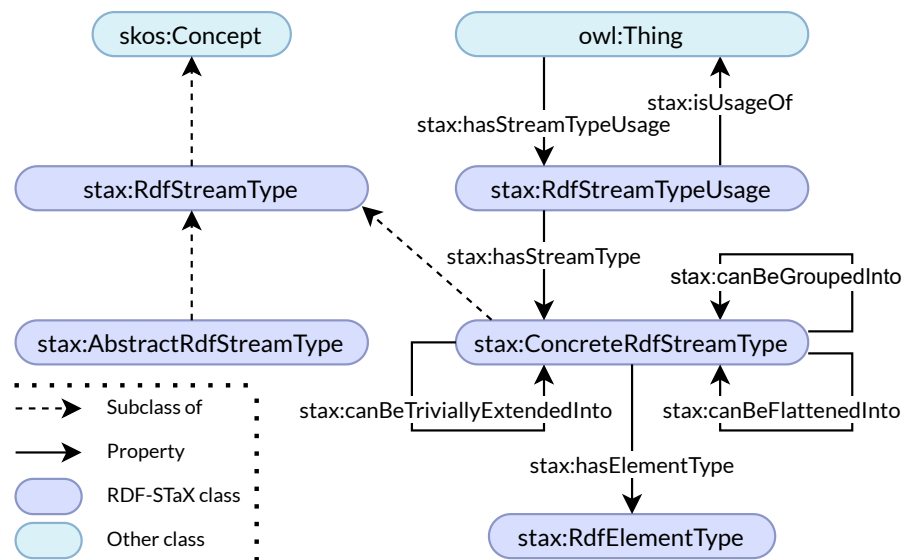


Figure 2. Classes and object properties in the RDF-STaX ontology.

It should be stressed that RDF-STaX does not define a class for RDF streams but rather classes for RDF stream *types*. For example, `stax:namedGraphStream` is an instance of the `stax:ConcreteRdfStreamType` class. The ontology also defines a class for RDF stream type usages (`stax:RdfStreamTypeUsage`), which can be thought of as subjective views of how an RDF stream is used in a specific context. This design philosophy is closely reflected in the expected usage patterns, which are described in the next subsection.

The ontology includes rich semantic relations between stream types. Properties `stax:canBeFlattenedInto` and `stax:canBeGroupedInto` relate grouped and flat stream types that can be easily converted into one another. For example, an RDF graph stream can be flattened into a flat RDF triple stream by concatenating its elements, while a flat RDF quad stream can be grouped into an RDF dataset stream by punctuating it. The property `stax:canBeTriviallyExtendedInto` is used for graph- and triple-based streams to indicate that they can be trivially turned into dataset- and quad-based streams, respectively. This transformation consists of stating explicitly that the triples are in the default graph. Object property chains are defined in the ontology to allow reasoning on these relations, exploiting

the taxonomical relations from SKOS [61]. These reasoning rules are presented below, in Manchester OWL Syntax [63]:

```
skos:broaderTransitive o stax:hasElementType
  SubPropertyOf: stax:hasElementType

skos:broaderTransitive o stax:canBeFlattenedInto
  SubPropertyOf: stax:canBeFlattenedInto

skos:broaderTransitive o stax:canBeGroupedInto
  SubPropertyOf: stax:canBeGroupedInto

skos:broaderTransitive o stax:canBeTriviallyExtendedInto
  SubPropertyOf: stax:canBeTriviallyExtendedInto
```

The ontology in version 1.1.1 includes 5 classes, 8 object properties, 15 individuals, and a total of 192 axioms. The axiom count was obtained with the `measure` command in ROBOT 1.9.6 [64] and does not include the imported SKOS vocabulary. The ontology's documentation <https://w3id.org/stax/1.1.1/ontology> (accessed on 4 June 2024) contains a detailed description of every class, property, and individual. Additional information facilitating ontology reuse is included, such as an explanation of the available ontology formats, download links, and versioning information.

#### 4.1. Ontology Usage Patterns

The proposed usage patterns for the RDF-STaX ontology reflect the philosophy of having a subjective view of how an RDF stream is used. These patterns are described in detail in the documentation <https://w3id.org/stax/1.1.1/use-it> (accessed on 4 June 2024), along with several illustrative examples and links to further resources.

The primary usage pattern in RDF-STaX is to create instances of the `stax:RdfStreamTypeUsage` class to annotate research works, software, or datasets, with each instance corresponding to a different view or use case of the stream. This pattern was selected due to the inherent subjectivity of determining the types of RDF streams, as discussed in Section 2.6. Below is an example of how this pattern may be used to annotate a DCAT [65] Dataset:

```
@prefix dcat: <http://www.w3.org/ns/dcat#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix stax: <https://w3id.org/stax/ontology#> .

_:dataset a dcat:Dataset ;
  # ... other properties of the dataset ...
  stax:hasStreamTypeUsage [
    a stax:RdfStreamTypeUsage ;
    stax:hasStreamType stax:datasetStream ;
    rdfs:comment "The data is a sequence of RDF datasets."@en
  ] , [
    a stax:RdfStreamTypeUsage ;
    stax:hasStreamType stax:flatQuadStream ;
    rdfs:comment "The data can be viewed as a flat sequence of RDF
  → quads."@en
  ] .
```

In some situations, it may be more convenient to annotate existing resources from an external perspective. RDF-STaX also supports this through the use of the `stax:isUsageOf` property. This is useful, for instance, when annotating published research papers or software with the RDF stream types that they use. This is illustrated in the example below, where the publication about IPSM [52] is annotated:

```

@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix stax: <https://w3id.org/stax/ontology#> .

_:usage a stax:RdfStreamTypeUsage ;
  rdfs:comment "The authors state that for their proposed system, 'the
  ↪ streaming element (i.e., a single message) (...) is a set of
  ↪ triples'. Therefore, internally it uses an RDF graph stream."@en ;
  stax:hasStreamType stax:graphStream ;
  stax:isUsageOf <https://doi.org/10.1109/ICSTCC.2017.8107003> .

```

#### 4.2. Alignments to Other Ontologies

The RDF-STaX ontology is aligned with other existing vocabularies with the use of semantic relations to foster the reuse of multiple ontologies linked together. Additionally, alignments can make it easier to understand the concepts introduced in RDF-STaX by comparing them to concepts already present in other vocabularies. However, as detailed in the comparison with the state of the art (Section 6.4), we could identify no other ontologies that would define the same concepts as RDF-STaX in a manner that would enable creating semantically strict, direct alignments (e.g., a subclassing relation). Instead, in cases where other ontologies define terms that are similar to those in RDF-STaX, a `skos:relatedMatch` alignment is provided. This essentially states that the meaning of the two concepts in different vocabularies is similar.

The following list enumerates the alignments that are provided in RDF-STaX release 1.1.1. Section 6.4 contains further discussion on this topic, comparing the meaning of RDF-STaX's terms with those of other ontologies.

- **SPARQL 1.1 Service Description vocabulary** [66]—the classes for RDF graphs (`sd:Graph`) and datasets (`sd:Dataset`) were aligned with the corresponding terms in RDF-STaX (`stax:graph` and `stax:dataset`, respectively), which are instances of class `stax:RdfElementType`.
- **VoCaLS** [43]—the `vocals:RdfStream` class was aligned with both the `stax:flatTripleStream` and `stax:graphStream` instances in RDF-STaX.
- **LDES** [56]—the `ldes:EventStream` class was aligned with the `stax:subjectGraphStream` instance in RDF-STaX.
- **VOID** [67]—the `void:Dataset` class was aligned with the `stax:flatTripleStream` instance in RDF-STaX.

#### 4.3. Availability and Sustainability

The sources of the RDF-STaX ontology and its documentation are hosted on GitHub <https://github.com/RDF-STaX/rdf-stax.github.io> (accessed on 4 June 2024). The repository includes continuous integration and deployment (CI/CD) scripts that automatically validate the ontology's profile, run competency question tests (see Section 5), generate documentation, run an OWL 2 reasoner to add more assertions, package the ontology into a number of formats, and publish it on the website: <https://w3id.org/stax> (accessed on 4 June 2024). Two versions of the ontology are published: OWL 2 DL, which can be used with reasoners (containing 249 triples), and OWL 2 Full (294 triples), which contains more metadata and alignments to other vocabularies. The latter is intended for use in contexts where ontological reasoning is not needed, e.g., for SPARQL querying. The additional metadata present there violates the OWL 2 DL profile, hence producing an OWL 2 Full ontology.

The ontology uses best practices for findability, accessibility, interoperability, and reusability (FAIR), developed by the community [68], compliance with which is evaluated in Section 6.3. All releases of the ontology, including the OWL 2 DL and OWL 2 Full variants, are distributed through the RDF-STaX website, under a permanent URL hosted by the W3C Permanent Identifier Community Group. The ontology is also archived in Zenodo [69] and registered in DBpedia Archivio <https://archivo.dbpedia.org/info?o=https://w3id.org/stax>

[//w3id.org/stax/ontology](https://w3id.org/stax/ontology) [70] (accessed on 4 June 2024) and Linked Open Vocabularies <https://lov.linkeddata.es/dataset/lov/vocabs/stax> [71] (accessed on 4 June 2024).

RDF-STaX is an open and freely licensed project, welcoming contributions to the ontology, its documentation, and tooling. A usage guide <https://w3id.org/stax/1.1.1/use-it> (accessed on 4 June 2024) with several use case scenarios and examples was prepared to facilitate adoption. To ensure long-term sustainability, a comprehensive contribution guide <https://w3id.org/stax/1.1.1/contributing> (accessed on 4 June 2024) was provided, along with a public issue tracker <https://github.com/RDF-STaX/rdf-stax.github.io/issues> (accessed on 4 June 2024). The project is entirely hosted on the permanently free infrastructure of GitHub and w3id.org, which should ensure its long-term stability.

## 5. Use Cases

In what follows, the three most representative use cases envisioned for RDF-STaX are described. The first use case is annotating published datasets and streams (for example, on the Web) to improve their accessibility and interoperability. The second use case focuses on achieving RDF streaming tool interoperability by embedding stream-type metadata within the stream itself. The final envisioned application is analyzing the state of the art (SOTA) of RDF streaming in research and software to easily compare different solutions and foster compatibility.

All of these use cases are described in detail below, along with accompanying competency questions (CQs) realized as SPARQL queries that should be run against the RDF-STaX ontology in the OWL 2 Full version. The CQ queries are available in RDF-STaX's documentation <https://w3id.org/stax/1.1.1/uses/cq/> (accessed on 4 June 2024) and are used to automatically and continuously test the ontology's validity with regard to the use cases, as described in Section 6.1. Two of the presented use cases (published dataset annotation and analyzing the SOTA of RDF streaming) are already realized in practice, as detailed below. The three presented use cases should only be considered a starting point for RDF-STaX, as the community may over time find new ways in which to use it. The resource is expected to evolve over time (through a process called ontology evolution in the literature [72]) to remain up to date with the community's needs. To enable this, RDF-STaX adopts a thoroughly open paradigm, inviting outside contributions.

### 5.1. Annotating Published Datasets and Streams

In this use case, RDF-STaX's stream types are used as metadata of datasets or streams published on the Web (for example, using the Linked Data mechanisms). This use case can be viewed from two sides—the publisher's and the end user's, which is reflected in the defined competency questions (see Table 2). The questions pertain to obtaining basic information about the defined stream types (CQ1.1 and CQ1.2). This can be used by the publisher to decide which stream type to use, and by the end user to understand what that stream type really is. Questions CQ1.3 and CQ1.5 play an analogous role but also aim to provide additional context by relating the terms in RDF-STaX to external sources (other ontologies). Finally, CQ1.4 serves the end user by providing links to examples of how each stream type can be used in practice.

This use case was implemented in RiverBench [73], an open RDF streaming benchmark suite, containing heterogeneous streaming datasets that are published using automated pipelines, following FAIR principles. RiverBench integrates with the RDF-STaX ontology, requiring all datasets to be annotated with their stream types. Each dataset has two stream types assigned, one for the grouped stream formulation and another for the flat. The consistency of the annotations is checked in CI/CD using SHACL rules [74] and the semantic relations defined in RDF-STaX (SKOS taxonomy properties and `stax:canBeFlattenedInto`). The datasets are then organized into benchmark profiles by their stream type. The published dataset documentation and RDF metadata based on DCAT [65] also includes the RDF stream type information.

**Table 2.** Competency questions for use case 1—annotating published datasets and streams. Full SPARQL queries and expectations for each test can be found on the ontology’s website: <https://w3id.org/stax/1.1.1/uses/cq/#use-case-1> (accessed on 4 June 2024).

#	Competency Question
CQ1.1	What are the names and descriptions of all RDF stream types?
CQ1.2	What is the definition for each stream type?
CQ1.3	What is the type of element of each concrete stream type? Provide additional references to external sources for each element type, if available.
CQ1.4	How can each of the concrete stream types be used? Provide a link to one example for each.
CQ1.5	What are the corresponding terms from other ontologies to those defined in the RDF-STaX ontology?

At the time of submission, RiverBench’s 2.0.0 release includes 12 datasets that use each stream type in RDF-STaX, except the RDF named graph stream (however, there are examples of the timestamped variant). The full metadata dump of RiverBench 2.0.0 consists of 9280 triples, including 201 `stax:hasStreamTypeUsage` annotations. Detailed technical information about this use case along with examples are available on the RDF-STaX website: <https://w3id.org/stax/1.1.1/uses> (accessed on 4 June 2024).

### 5.2. Embedding Metadata in Streams

This use case focuses on embedding stream-type metadata within the stream itself, with the goal of enabling streaming system interoperability. For example, a stream producer would state that the stream is of a given type, and then the consumer would check if it can consume this type of stream or convert it into a consumable stream. As this use case focuses on automated handling of streams, the competency questions (see Table 3) exploit the reasoning capabilities of RDF-STaX. Thus, CQ2.1, CQ2.2, CQ2.3, and CQ2.4 all address different possible transformations a stream processor could perform on an RDF stream.

**Table 3.** Competency questions for use case 2—embedding metadata in streams. Full SPARQL queries and expectations for each test can be found on the ontology’s website: <https://w3id.org/stax/1.1.1/uses/cq/#use-case-2> (accessed on 4 June 2024).

#	Competency Question
CQ2.1	Which concrete stream types can be viewed as generalizations of other stream types?
CQ2.2	Which concrete stream types can be trivially extended (by assuming the default graph) into other stream types?
CQ2.3	Which concrete stream types can be flattened into other stream types?
CQ2.4	Which concrete stream types can be grouped to obtain other stream types?

### 5.3. Analyzing the State of the Art of RDF Streaming

The final use case pertains to analyzing the SOTA of RDF streaming by annotating published research works and software with the RDF stream types used by them. This task is thus very closely related to the SOTA analysis performed in this paper in Section 3.3. As with any other scientific field, however, the landscape of RDF streaming changes rapidly, and therefore, using the RDF-STaX ontology in combination with a knowledge base of research works can help automate some literature review processes. Therefore, the defined competency questions (see Table 4) tackle more complex issues that may be asked when analyzing the correspondences between the annotated research works. CQ3.1 and CQ3.2 pertain to the taxonomical structure of RDF-STaX, helping the user understand how the different stream types relate to each other. Meanwhile, CQ3.3 attempts to find pairs of essentially incompatible RDF stream formulations, i.e., those that cannot be trivially converted into each other. This would serve to identify methods that cannot easily interoperate.

**Table 4.** Competency questions for use case 3—analyzing the state of the art of RDF streaming. Full SPARQL queries and expectations for each test can be found on the ontology’s website: <https://w3id.org/stax/1.1.1/uses/cq/#use-case-3> (accessed on 4 June 2024).

#	Competency Question
CQ3.1	What are the taxonomical parents of each stream type, listed in order?
CQ3.2	Is there any stream type that is a taxonomical parent or child of itself?
CQ3.3	Which conversions between concrete stream types cannot be performed in any trivial way (grouping, flattening, extending)? Allow for multiple trivial transformations in series and take into account the taxonomical structure.

This use case was implemented within the scope of this work, with a nanopublication dataset of research works on RDF streaming. A nanopublication is a small unit of scientific knowledge published in RDF [75]—it can be an opinion, a measurement, or any other assertion. Nanopublications are a natural way to express statements about scientific works or software due to their strong provenance mechanisms and citeability. This fits well with the pattern of subjective statements about stream types, which is used in RDF-STaX. Therefore, to make it easier to create nanopublications with the RDF-STaX ontology, a template was prepared for Nanodash, an open nanopublication editor [76]. The template along with the accompanying manual <https://w3id.org/stax/1.1.1/nanopubs> (accessed on 4 June 2024) makes it easy to create and publish semantic assertions such as “X uses stream type Y because Z”, requiring the user to only fill out a simple online form.

The results from the survey conducted in this paper (as summarized in Section 3.3) were published as 35 nanopublications using the aforementioned mechanism, totaling 876 triples. The nanopublications include a more elaborated discussion of each stream type usage, providing more details than Table 1. The RDF-STaX nanopublications are packaged as a single knowledge graph in a CI/CD pipeline and published automatically on the website <https://w3id.org/stax/1.1.1/nanopubs> (accessed on 4 June 2024), where they can be downloaded and reused under the free CC BY 4.0 license. Anyone can contribute to the dataset by simply publishing nanopublications with the provided Nanodash template. Effectively, the dataset along with the publishing mechanism can serve as a collaborative, living state-of-the-art review, allowing one to easily compare and discuss works on RDF streaming, in a structured manner. A similar mechanism for living literature reviews was previously proposed by Wijkstra et al. [77] to whose work we refer the reader for more context.

## 6. Ontology Evaluation

This section presents the results of a multifaceted evaluation of the RDF-STaX ontology. Firstly, the evaluation includes checking use case coverage with the use of competency questions defined in Section 5. Secondly, the validity of the ontology is evaluated with regard to the OWL specification and its logical consistency. Thirdly, the findability, accessibility, interoperability, and reusability aspects of the ontology are evaluated, especially with regard to the current, community-developed best practices. Finally, RDF-STaX is compared with other similar resources, highlighting the advances offered by this work.

### 6.1. Use Case Coverage

This evaluation is based on the twelve competency question tests defined in Section 5. Each of these tests was written down as a SPARQL query paired with the expectation for the test, which specifies how many results should the query return when run against the RDF-STaX ontology. The tests are run in sequence by a CI/CD job implemented in the Python language, with the `rdflib` library [78]—source code is available on GitHub <https://github.com/RDF-STaX/ci-worker> (accessed on 4 June 2024) under the Apache 2.0 license. The job is triggered on every ontology release (including the development release) and pull request. If any test fails, the ontology will not be released, thus ensuring that it always covers the competency questions from the use cases. Further technical

implementation details on the tests are available in the documentation: <https://w3id.org/stax/1.1.1/contributing/#competency-question-tests> (accessed on 4 June 2024).

At the time of submission (RDF-STaX release 1.1.1), the CI/CD job reports that the ontology successfully passes all competency question tests. Additionally, the fact that the ontology was employed in two implemented applications (see Section 5) serves as additional validation of the ontology’s usefulness and practicality [79].

### 6.2. Logical and OWL Profile Validity

The purpose of this evaluation is to check if the ontology meets the requirements of specific OWL 2 profiles [80], which strike different balances between the expressive power of the ontology and reasoning efficiency. Similarly to the use case coverage evaluation, the compliance checks with OWL 2 profiles were implemented as CI/CD jobs, which run on any ontology release and pull request. Both distributions of the ontology (OWL 2 DL and OWL 2 Full) are checked to see if they match their respective profile. The check is performed using ROBOT’s [64] `validate-profile` command. Additionally, the logical consistency of the ontology is validated in CI/CD with ROBOT’s `reason` command, using the HermiT reasoner [81].


At the time of submission (RDF-STaX release 1.1.1), the CI/CD job reports that both distributions of the ontology (OWL 2 DL and OWL 2 Full) are compliant with their respective profiles, with no inconsistencies. For additional validation, RDF-STaX 1.1.1 was checked with the publicly available Ontology Pitfall Scanner (OOPS!) [82], which reported no critical issues.

### 6.3. Findability, Accessibility, Interoperability, and Reusability

The question of what constitutes a FAIR ontology in practice is a complicated one [83], with many possible interpretations for each of the FAIR Guiding Principles [84]. To best reflect the needs of the Semantic Web community and ensure impartiality, in this evaluation, we chose to rely on two well-known, publicly available, automatic FAIR evaluators: FOOPS! [85] and DBpedia Archivio [70]. The following results can be easily reproduced by inputting into either of them the RDF-STaX ontology IRI.

In FOOPS!, RDF-STaX obtained a score of 95%, with 8.83/9 points in findability, 3/3 in accessibility, 3/3 in interoperability, and 7.92/9 in reusability (Figure 3a). All of the issues reported by the tool are minor. In DBpedia Archivio, RDF-STaX obtained 4/4 stars (Figure 3b). It should be stressed here that Archivio’s manual states that such a result means the ontology has reached “minimum viability” in the context of FAIR but is not necessarily “good”. However, at the time of submission, this is the highest possible rating offered by DBpedia Archivio, with only 20.03% of all indexed ontologies having four out of four stars.

Title:	<input type="text" value="RDF Stream Taxonomy (RDF-STaX)"/>
URI:	<input type="text" value="https://w3id.org/stax/ontology"/>
License:	<input type="text" value="https://spdx.org/licenses/CC-BY-4.0"/>



(a)

Figure 3. Cont.

View Archived Ontology	Download Latest	Triples	Stars	Crawling Status
<a href="#">RDF Stream Taxonomy (RDF-STaX)</a>	<a href="#">owl, ttl, nt</a>	294	★★★★	

(b)

**Figure 3.** FAIR evaluation results for RDF-STaX: Both screenshots were taken on 4 June 2024. (a) FAIR evaluation results obtained from FOOPS! [85]. (b) FAIR evaluation results obtained from DBpedia Archivio [70].

#### 6.4. Comparison with the State of the Art

To the best of our knowledge, there is no other ontology or vocabulary that would model RDF stream types. As previously noted, the RDF-STaX ontology purposefully does not define classes for RDF streams but rather focuses on subjective RDF stream usages. This difference makes other ontologies nonoverlapping—instead, RDF-STaX can naturally complement them. As shown below, each existing ontology also covers only one or two selected formulations of the RDF streaming problem. Thus, RDF-STaX can provide a common semantic bridge to describe RDF stream types across different underlying vocabularies. To this end, the documentation includes examples <https://w3id.org/stax/1.1.1/use-it/> (accessed on 4 June 2024) of using RDF-STaX with these complementary ontologies (VoCaLS, VoID, LDES, and DCAT). The expected usage pattern is to define an RDF stream or a dataset with these ontologies as usual and add the `stax:hasStreamTypeUsage` property to indicate the stream type according to RDF-STaX. The same pattern can be used with any future ontology for describing streams or datasets.

In what follows, the relevant ontologies are discussed in detail, focusing on their differences and complementarities with RDF-STaX. In the discussion, only publicly available ontologies are included, as otherwise it would not be possible to reliably assess the ontology's scope and semantics.

The VoCaLS vocabulary [43] is intended to help publish streaming data using the Linked Data principles and to describe streaming services. It defines the `vocal:RDFStream` class for potentially infinite sequences of RDF graphs and/or triples. Instances of this class can be thus annotated as RDF graph streams or flat RDF triple streams in RDF-STaX. As RDF-STaX does not define a class for RDF streams (but rather has instances of RDF stream types), the two ontologies are entirely complementary. Due to the different semantics of the terms in VoCaLS and RDF-STaX, a direct semantic connection (e.g., subclassing) between the two ontologies cannot be made. Instead, a `skos:relatedMatch` alignment is provided, along with a suggested usage pattern in the documentation of RDF-STaX.

The Linked Data Event Streams (LDES) vocabulary [56] defines the `ldes:EventStream` class, which is a collection of stream elements—sets of RDF triples. Each element is identified with its subject node. Therefore, instances of this class can be annotated in RDF-STaX as RDF subject graph streams. Such an alignment was added to the ontology.

The VoID vocabulary [67] describes datasets with RDF data (not *RDF datasets* as per RDF 1.1), which are in fact sets of triples with additional context information. Thus, instances of class `void:Dataset` can be annotated as RDF-STaX's flat RDF triple streams. This alignment was added to the ontology.

The more general DCAT [65] vocabulary describes datasets and their distributions. These datasets may be RDF streams, but DCAT does not assume any specific format of the datasets annotated with it. Therefore, RDF-STaX can be used as a complementary tool to annotate the type of the stream contained within such a dataset—this is performed, for example, in RiverBench; see Section 5.1.

The RDF-STaX ontology also defines four element types (RDF graphs, triples, quads, and datasets). In this regard, we could identify only the SPARQL 1.1 Service Description vocabulary [66], which has the `sd:Graph` and `sd:Dataset` classes for RDF graphs and datasets. These classes were aligned with the corresponding terms in RDF-STaX, using the `skos:relatedMatch` property.



In summary, to the best of our knowledge, RDF-STaX is the first taxonomy and ontology that allows describing more than one type of RDF stream, making it a crucial step in enabling RDF streaming interoperability and understanding the field's state of the art.

## 7. Discussion and Future Work

Although the “RDF stream” is a term that is now ubiquitous in research, it is used to describe many different concepts. There is no universally accepted language for describing RDF streams, which leads to a situation in which RDF streaming solutions become hard to assess, compare, and connect together. For example, forming a pipeline composed of a stream annotator, translator, streaming protocol, and reasoner currently appears to be a tough challenge, as it is hard to determine what exactly is an RDF stream for each of these tools and whether these formulations are compatible with each other or not.

It appears that the variety of RDF stream formulations is here to stay. Each of the different definitions has its merits and use cases, and none of them should be disregarded. Instead, we call to embrace the variety of RDF streams by systematizing them in one shared taxonomy. The taxonomy, with its semantic relations (taxonomical, flattening, grouping, extending—see Section 4), allows one to reason about stream types not only in terms of subclassing but also in terms of possible conversions between them. This, we hope, will be a much-needed step in the direction of making RDF stream applications more compatible with each other.

RDF-STaX is embodied in an ontology that makes the taxonomy readily usable in practice. The ontology does not attempt to replace any existing vocabulary for describing streams or datasets but rather aims to complement them and serve as a shared semantic bridge between them. We hope that by being nonoverlapping by design and focused on a single purpose, RDF-STaX will remain useful and relevant in the long term. The three presented use cases, two of which were already realized in practice, are common and well known to the community, making the ontology all the more relevant.

Another result of this work is the nanopublications dataset, which is a collaborative, living state-of-the-art review of RDF streaming. Nanopublications, in our view, are a natural way to structure the intersubjective scientific discourse. It is possible, for example, to comment on other nanopublications and confirm or disagree with them, all in a structured manner. Here, we propose to use this powerful tool for a very narrow field of science (RDF streaming), but its possible future applications are certainly much broader [77]. We encourage other researchers to contribute their perspectives to the dataset or to comment on the existing assertions.

As for future work, the theory presented in this contribution notably does not tackle the formal semantics of RDF streams. As with RDF datasets, it is unlikely that there can be single, unifying semantics for streams, but, nonetheless, we see this area as worth exploring in the future. The presented survey of RDF streaming is also not a systematic review—it only serves as a basis for the discussion in this work. Producing a systematic review of RDF streaming would be very valuable to the community, and we hope that the proposed taxonomy and the provided nanopublication tooling will make creating it easier.

## 8. Conclusions

In this work, we propose the RDF Stream Taxonomy, which systematizes RDF stream definitions found in the literature. The taxonomy uncovered the variety of RDF streaming in research and practice, highlighting the open challenge of making the different applications compatible with each other. The proposed RDF-STaX ontology embodies the semantic relations between stream types in the taxonomy and makes the definitions readily applicable. The ontology positions itself as complementary to existing vocabularies and employs a flexible approach which consists of making subjective statements about the types of RDF streams. The ontology and its documentation are fully open and follow community-developed best practices for FAIR.

The ontology was designed with three use cases in mind, with supporting competency questions for each. Two of these use cases were already implemented—annotating research works in nanopublications and describing RDF stream types of streaming datasets. The nanopublications dataset is an added contribution of this work, constituting a living state-of-the-art review that anyone can contribute to by using the provided documentation and tooling.

**Author Contributions:** Conceptualization, P.S. (Piotr Sowiński), P.S. (Paweł Szmeja), and M.G.; data curation, P.S. (Piotr Sowiński); formal analysis, P.S. (Piotr Sowiński); funding acquisition, M.G. and M.P.; investigation, P.S. (Piotr Sowiński); methodology, P.S. (Piotr Sowiński) and P.S. (Paweł Szmeja); project administration, M.G. and M.P.; resources, M.G. and M.P.; software, P.S. (Piotr Sowiński); supervision, M.G.; validation, P.S. (Piotr Sowiński), P.S. (Paweł Szmeja) and M.G.; visualization, P.S. (Piotr Sowiński); writing—original draft, P.S. (Piotr Sowiński); Writing—review and editing, P.S. (Paweł Szmeja), M.G. and M.P. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Data Availability Statement:** The RDF-STaX ontology introduced in this study (along with the documentation and the nanopublications dataset) is openly available at <https://w3id.org/stax/1.1.1> (accessed on 4 June 2024) and in Zenodo at <https://zenodo.org/doi/10.5281/zenodo.10072907> (accessed on 4 June 2024), reference number 10072907.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

CI/CD	Continuous Integration and Continuous Delivery
CQ	Competency Question
DCAT	Data Catalog Vocabulary
FAIR	Findable, Accessible, Interoperable, Reusable
FOOPS!	Ontology Pitfall Scanner for FAIR
GoT	Graph of Things
I/O	Input/Output
IPSM	Inter Platform Semantic Mediator
IRI	Internationalized Resource Identifier
LDES	Linked Data Event Streams
OOPS!	Ontology Pitfall Scanner
OWL	Web Ontology Language
RDF	Resource Description Framework
RDF-STaX	RDF Stream Taxonomy
ROBOT	ROBOT is an OBO Tool
RML	RDF Mapping Language
RSP	RDF Stream Processing
SHACL	Shapes Constraint Language
SKOS	Simple Knowledge Organization System
SPARQL	SPARQL Protocol And RDF Query Language
TA-RDF	Time-Annotated RDF
URI	Uniform Resource Identifier
W3C	World Wide Web Consortium

## References

1. Pan, J.Z. Resource Description Framework. In *Handbook on Ontologies*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 71–90.
2. Hitzler, P. A review of the Semantic Web field. *Commun. ACM* **2021**, *64*, 76–83. [[CrossRef](#)]
3. Cyganiak, R.; Wood, D.; Lanthaler, M. RDF 1.1 Concepts and Abstract Syntax. W3C Recommendation, W3C. 2014. Available online: <https://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/> (accessed on 17 April 2024).
4. Kleppmann, M. *Designing Data-Intensive Applications: The Big Ideas behind Reliable, Scalable, and Maintainable Systems*; O'Reilly Media, Inc.: Sebastopol, CA, USA, 2017.

5. Bonte, P.; Tommasini, R. Streaming linked data: A survey on life cycle compliance. *J. Web Semant.* **2023**, *77*, 100785. [[CrossRef](#)]
6. Sowiński, P.; Wasielewska-Michniewska, K.; Ganzha, M.; Paprzycki, M. (2022, October). Efficient RDF streaming for the edge-cloud continuum. In Proceedings of the 2022 IEEE 8th World Forum on Internet of Things (WF-IoT), Yokohama, Japan, 26 October–11 November 2022; pp. 1–8. [[CrossRef](#)]
7. Tommasini, R.; Bonte, P.; Spiga, F.; Della Valle, E. *Streaming Linked Data: From Vision to Practice*; Springer Nature: Berlin/Heidelberg, Germany, 2023.
8. Van Assche, D.; Delva, T.; Haesendonck, G.; Heyvaert, P.; De Meester, B.; Dimou, A. Declarative RDF graph generation from heterogeneous (semi-) structured data: A systematic literature review. *J. Web Semant.* **2023**, *75*, 100753. [[CrossRef](#)]
9. Dell’Aglio, D.; Della Valle, E.; Calbimonte, J.P.; Corcho, O. RSP-QL semantics: A unifying query model to explain heterogeneity of RDF stream processing systems. *Int. J. Semant. Web Inf. Syst. (IJSWIS)* **2014**, *10*, 17–44. [[CrossRef](#)]
10. Fernández, J.D.; Llaves, A.; Corcho, O. Efficient RDF interchange (ERI) format for RDF data streams. In Proceedings of the International Semantic Web Conference, Riva del Garda, Italy, 19–23 October 2014; Springer: Berlin/Heidelberg, Germany, 2014; pp. 244–259.
11. Oo, S.M.; Haesendonck, G.; De Meester, B.; Dimou, A. RMLStreamer-SISO: An RDF stream generator from streaming heterogeneous data. In Proceedings of the International Semantic Web Conference, Hangzhou, China, 24 October 2022; Springer: Berlin/Heidelberg, Germany, 2022; pp. 697–713.
12. Dell’Aglio, D.; Dao-Tran, M.; Calbimonte, J.P.; Le Phuoc, D.; Della Valle, E. A query model to capture event pattern matching in RDF stream processing query languages. In Proceedings of the European Knowledge Acquisition Workshop, Bologna, Italy, 19 November 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 145–162.
13. Le-Phuoc, D.; Polleres, A.; Hauswirth, M.; Tummarello, G.; Morbidoni, C. Rapid prototyping of semantic mash-ups through semantic web pipes. In Proceedings of the 18th International Conference on World Wide Web, Madrid, Spain, 20–24 April 2009; pp. 581–590.
14. Schraudner, D.; Harth, A. Stream Containers for Resource-oriented RDF Stream Processing. *arXiv* **2022**, arXiv:2202.13630.
15. Keskiärrkkä, R.; Blomqvist, E. Semantic complex event processing for social media monitoring—a survey. In Proceedings of the Social Media and Linked Data for Emergency Response (SMILE) Co-located with the 10th Extended Semantic Web Conference, Montpellier, France, 26–27 May 2013; CEUR Workshop Proceedings (May 2013).
16. Llanes, K.R.; Casanova, M.A.; Lemus, N.M. From sensor data streams to linked streaming data: A survey of main approaches. *J. Inf. Data Manag.* **2016**, *7*, 130.
17. Ma, Z.; Capretz, M.A.; Yan, L. Storing massive Resource Description Framework (RDF) data: A survey. *Knowl. Eng. Rev.* **2016**, *31*, 391–413. [[CrossRef](#)]
18. Modoni, G.E.; Sacco, M.; Terkaj, W. A survey of RDF store solutions. In Proceedings of the 2014 International Conference on Engineering, Technology and Innovation (ICE), Bergamo, Italy, 23–25 June 2014; IEEE: New York, NY, USA, 2014; pp. 1–7.
19. Özsü, M.T. A survey of RDF data management systems. *Front. Comput. Sci.* **2016**, *10*, 418–432. [[CrossRef](#)]
20. Su, X.; Gilman, E.; Wetz, P.; Riekkki, J.; Zuo, Y.; Leppänen, T. Stream reasoning for the Internet of Things: Challenges and gap analysis. In Proceedings of the 6th International Conference on Web Intelligence, Mining and Semantics, Nîmes, France, 13–15 June 2016; pp. 1–10.
21. Zhang, F.; Li, Z.; Peng, D.; Cheng, J. RDF for temporal data management—A survey. *Earth Sci. Inform.* **2021**, *14*, 563–599. [[CrossRef](#)]
22. Hasemann, H.; Krölller, A.; Pagel, M. RDF Provisioning for the Internet of Things. In Proceedings of the 2012 3rd IEEE International Conference on the Internet of Things, Wuxi, China, 24–26 October 2012; IEEE: New York, NY, USA, 2012; pp. 143–150.
23. Fernández, N.; Arias, J.; Sánchez, L.; Fuentes-Lorenzo, D.; Corcho, Ó. RDSZ: An approach for lossless RDF stream compression. In Proceedings of the European Semantic Web Conference, Riva del Garda, Italy, 19–23 October 2014; Springer: Berlin/Heidelberg, Germany, 2014; pp. 52–67.
24. Käbisch, S.; Peintner, D.; Anicic, D. Standardized and efficient RDF encoding for constrained embedded networks. In Proceedings of the European Semantic Web Conference, Bethlehem, PA, USA, 11–15 October 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 437–452.
25. Groppe, S.; Groppe, J.; Kukulenz, D.; Linnemann, V. A SPARQL engine for streaming RDF data. In Proceedings of the 2007 Third International IEEE Conference on Signal-Image Technologies and Internet-Based System, Shanghai, China, 16–18 December 2007; IEEE: New York, NY, USA, 2007; pp. 167–174.
26. Anicic, D.; Fodor, P.; Rudolph, S.; Stojanovic, N. EP-SPARQL: A unified language for event processing and stream reasoning. In Proceedings of the 20th International Conference on World Wide Web, Hyderabad, India, 28 March–1 April 2011; pp. 635–644.
27. Barbieri, D.F.; Braga, D.; Ceri, S.; Valle, E.D.; Grossniklaus, M. Querying RDF streams with C-SPARQL. *ACM SIGMOD Rec.* **2010**, *39*, 20–26. [[CrossRef](#)]
28. Bolles, A.; Grawunder, M.; Jacobi, J. Streaming SPARQL—extending SPARQL to process data streams. In Proceedings of the The Semantic Web: Research and Applications: 5th European Semantic Web Conference, ESWC 2008, Tenerife, Canary Islands, Spain, 1–5 June 2008; Proceedings 5; Springer: Berlin/Heidelberg, Germany, 2008; pp. 448–462.
29. Calbimonte, J.P.; Corcho, Ó. Evaluating SPARQL Queries over Linked Data Streams. In *Linked Data Management*; Chapman and Hall/CRC: Boca Raton, FL, USA, 2016; pp. 165–190.

30. Dell’Aglío, D.; Calbimonte, J.P.; Della Valle, E.; Corcho, O. Towards a unified language for RDF stream query processing. In Proceedings of the European Semantic Web Conference, Bethlehem, PA, USA, 11–15 October 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 353–363.
31. Komazec, S.; Cerri, D.; Fensel, D. Sparkwave: Continuous schema-enhanced pattern matching over RDF data streams. In Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems, Berlin, Germany, 16–20 July 2012; pp. 58–68.
32. Le-Phuoc, D.; Dao-Tran, M.; Xavier Parreira, J.; Hauswirth, M. A native and adaptive approach for unified processing of linked streams and linked data. In Proceedings of the International Semantic Web Conference, Bonn, Germany, 23–27 October 2011; Springer: Berlin/Heidelberg, Germany, 2011; pp. 370–388.
33. Tommasini, R.; Bonte, P.; Ongenaes, F.; Della Valle, E. RSP4J: An API for RDF stream processing. In Proceedings of the The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, 6–10 June 2021; Proceedings 18; Springer: Berlin/Heidelberg, Germany, 2021; pp. 565–581.
34. Dell’Aglío, D.; Calbimonte, J.P.; Balduini, M.; Corcho, O.; Della Valle, E. On correctness in RDF stream processor benchmarking. In Proceedings of the The Semantic Web–ISWC 2013: 12th International Semantic Web Conference, Sydney, NSW, Australia, 21–25 October 2013; Proceedings, Part II 12; Springer: Berlin/Heidelberg, Germany, 2013; pp. 326–342.
35. Le Phuoc, D.; Dao-Tran, M.; Le Tuan, A.; Duc, M.N.; Hauswirth, M. RDF stream processing with CQELS framework for real-time analysis. In Proceedings of the 9th ACM International Conference on Distributed Event-Based Systems, Oslo, Norway, 29 June–3 July 2015; pp. 285–292.
36. Calbimonte, J.P. Linked data notifications for RDF streams. In Proceedings of the Web Stream Processing Workshop (WSP 2017) and the 2nd International Workshop on Ontology Modularity, Contextuality, and Evolution (WOMoCoE 2017) Co-Located with 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria, 22 October 2017.
37. Mauri, A.; Calbimonte, J.P.; Dell’Aglío, D.; Balduini, M.; Brambilla, M.; Della Valle, E.; Aberer, K. TripleWave: Spreading RDF streams on the web. In Proceedings of the The Semantic Web–ISWC 2016: 15th International Semantic Web Conference, Kobe, Japan, 17–21 October 2016; Proceedings, Part II 15; Springer: Berlin/Heidelberg, Germany, 2016; pp. 140–149.
38. Wu, J.; Orlandi, F.; O’Sullivan, D.; Dev, S. A workflow to convert live atmospheric sensor data into linked data. In Proceedings of the IGARSS 2022–2022 IEEE International Geoscience and Remote Sensing Symposium, Kuala Lumpur, Malaysia, 17–22 July 2022; IEEE: New York, NY, USA, 2022; pp. 4086–4089.
39. Tappolet, J.; Bernstein, A. Applied temporal RDF: Efficient temporal querying of RDF data with SPARQL. In Proceedings of the European Semantic Web Conference, Crete, Greece, 31 May–4 June 2009; Springer: Berlin/Heidelberg, Germany, 2009; pp. 308–322.
40. RDF Stream Processing Community Group. RSP Data Model. Draft community group report, W3C RSP Community Group. Available online: <https://streamreasoning.org/RSP-QL/Abstract%20Syntax%20and%20Semantics%20Document/> (accessed on 17 April 2024).
41. Barbieri, D.F.; Valle, E. A proposal for publishing data streams as Linked Data. In Proceedings of the Linked Data on the Web Workshop, Raleigh, NC, USA, 27 April 2010.
42. Rodriguez, A.; McGrath, R.; Liu, Y.; Myers, J.; Urbana-Champaign, I. Semantic management of streaming data. *Proc. Semant. Sens. Netw.* **2009**, *80*, 80–95.
43. Tommasini, R.; Sedira, Y.A.; Dell’Aglío, D.; Balduini, M.; Ali, M.I.; Le Phuoc, D.; Della Valle, E.; Calbimonte, J.P. VoCaLS: Vocabulary and catalog of linked streams. In Proceedings of the The Semantic Web–ISWC 2018: 17th International Semantic Web Conference, Monterey, CA, USA, 8–12 October 2018; Proceedings, Part II 17; Springer: Berlin/Heidelberg, Germany, 2018; pp. 256–272.
44. Haesendonck, G.; Maroy, W.; Heyvaert, P.; Verborgh, R.; Dimou, A. Parallel RDF generation from heterogeneous big data. In Proceedings of the International Workshop on Semantic Big Data, Amsterdam, The Netherlands, 5 July 2019; pp. 1–6.
45. Dimou, A.; Vander Sande, M.; Colpaert, P.; Verborgh, R.; Mannens, E.; Van de Walle, R. RML: A generic language for integrated RDF mappings of heterogeneous data. *Ldow* **2014**, *8*, 1184.
46. CARML Contributors. CARML: A Pretty Sweet RML Engine, for RDF. 2023. Available online: <https://github.com/carml/carml> (accessed on 17 April 2024).
47. Eclipse Foundation, Inc. Parsing and Writing RDF with Rio. 2023. Available online: <https://rdf4j.org/documentation/programming/rio/> (accessed on 17 April 2024).
48. Szejma, P. ASSIST-IoT Semantic Annotation Enabler. 2023. Available online: [https://github.com/assist-iot/semantic\\_annotation](https://github.com/assist-iot/semantic_annotation) (accessed on 17 April 2024).
49. Szejma, P.; Fornés-Leal, A.; Lacalle, I.; Palau, C.E.; Ganzha, M.; Pawłowski, W.; Paprzycki, M.; Schabbink, J. ASSIST-IoT: A modular implementation of a reference architecture for the next generation Internet of Things. *Electronics* **2023**, *12*, 854. [[CrossRef](#)]
50. Lefrançois, M.; Zimmermann, A.; Bakerally, N. Flexible RDF generation from RDF and heterogeneous data sources with SPARQL-Generate. In Proceedings of the European Knowledge Acquisition Workshop, Bologna, Italy, 19–23 November 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 131–135.
51. Apache Software Foundation. Working with RDF Streams in Apache Jena. 2023. Available online: <https://jena.apache.org/documentation/io/streaming-io.html> (accessed on 17 April 2024).

52. Ganzha, M.; Paprzycki, M.; Pawłowski, W.; Szmaja, P.; Wasielewska, K. Streaming semantic translations. In Proceedings of the 2017 21st International Conference on System Theory, Control and Computing (ICSTCC), Sinaia, Romania, 19–21 October 2017; IEEE: New York, NY, USA, 2017; pp. 1–8.
53. Morsey, M.; Lehmann, J.; Auer, S.; Stadler, C.; Hellmann, S. DBpedia and the live extraction of structured data from Wikipedia. *Program* **2012**, *46*, 157–181. [[CrossRef](#)]
54. Le-Phuoc, D.; Quoc, H.N.M.; Quoc, H.N.; Nhat, T.T.; Hauswirth, M. The graph of things: A step towards the live knowledge graph of connected things. *J. Web Semant.* **2016**, *37*, 25–35. [[CrossRef](#)]
55. Tallon, J.; Webber, C. ActivityPub. W3C Recommendation, W3C. 2018. Available online: <https://www.w3.org/TR/2018/REC-activitypub-20180123/> (accessed on 26 June 2024).
56. Van Lancker, D.; Colpaert, P.; Delva, H.; Van de Vyvere, B.; Meléndez, J.R.; Dedecker, R.; Michiels, P.; Buyle, R.; De Craene, A.; Verborgh, R. Publishing base registries as linked data event streams. In Proceedings of the International Conference on Web Engineering, Biarritz, France, 18–21 May 2021; Springer: Berlin/Heidelberg, Germany, 2021; pp. 28–36.
57. Le-Tuan, A.; Franzreb, C.; Le Phuoc, D.; Schimmler, S.; Hauswirth, M. Towards Building Live Open Scientific Knowledge Graphs. In Proceedings of the Companion Proceedings of the Web Conference, Lyon, France, 25–29 April 2022; pp. 443–447.
58. Zimmermann, A. RDF 1.1: On Semantics of RDF Datasets. W3C Note, W3C. 2014. Available online: <https://www.w3.org/TR/2014/NOTE-rdf11-datasets-20140225/> (accessed on 28 June 2024).
59. Hartig, O.; Champin, P.A.; Kellogg, G. RDF 1.2 Concepts and Abstract Syntax. W3C Working Draft, W3C. 2024. Available online: <https://www.w3.org/TR/2024/WD-rdf12-concepts-20240416/> (accessed on 28 June 2024).
60. Jupp, S.; Bechhofer, S.; Stevens, R. SKOS with OWL: Don't be Full-ish! In Proceedings of the OWLED, Washington, DC, USA, 1–2 April 2008; Volume 432, p. 2009-2.
61. Miles, A.; Matthews, B.; Wilson, M.; Brickley, D. SKOS core: Simple knowledge organisation for the web. In Proceedings of the International Conference on Dublin Core and Metadata Applications, Madrid, Spain, 12–15 September 2005; pp. 3–10.
62. Singh, G.; Bhatia, S.; Mutharaju, R. OWL2Bench: A benchmark for OWL 2 reasoners. In Proceedings of the The Semantic Web—ISWC 2020: 19th International Semantic Web Conference, Athens, Greece, 2–6 November 2020; Proceedings, Part II 19; Springer: Berlin/Heidelberg, Germany, 2020; pp. 81–96.
63. Horridge, M.; Drummond, N.; Goodwin, J.; Rector, A.L.; Stevens, R.; Wang, H. The Manchester OWL syntax. In Proceedings of the OWLed, Athens, GA, USA, 10–11 November 2006; Volume 216.
64. Jackson, R.C.; Balhoff, J.P.; Douglass, E.; Harris, N.L.; Mungall, C.J.; Overton, J.A. ROBOT: A Tool for Automating Ontology Workflows. *BMC Bioinform.* **2019**, *20*, 407. [[CrossRef](#)] [[PubMed](#)]
65. Browning, D.; Beltran, A.G.; Perego, A.; Winstanley, P.; Cox, S.; Albertoni, R. Data Catalog Vocabulary (DCAT)—Version 3. W3C Working Draft, W3C. 2023. Available online: <https://www.w3.org/TR/2023/WD-vocab-dcat-3-20230307/> (accessed on 28 June 2024).
66. Williams, G. SPARQL 1.1 Service Description. W3C Recommendation, W3C. 2013. Available online: <https://www.w3.org/TR/2013/REC-sparql11-service-description-20130321/> (accessed on 28 June 2024).
67. Zhao, J.; Alexander, K.; Hausenblas, M.; Cyganiak, R. Describing Linked Datasets with the VoID Vocabulary. W3C Note, W3C. 2011. Available online: <https://www.w3.org/TR/2011/NOTE-void-20110303/> (accessed on 28 June 2024).
68. Garijo, D.; Poveda-Villalón, M. Best Practices for Implementing FAIR Vocabularies and Ontologies on the Web. In *Applications and Practices in Ontology Design, Extraction, and Reasoning*; IOS Press: Amsterdam, The Netherlands, 2020; pp. 39–54.
69. Sowiński, P. RDF-STaX/rdf-stax.github.io. 2024. Available online: <https://zenodo.org/records/11476591> (accessed on 28 June 2024).
70. Frey, J.; Streitmatter, D.; Götz, F.; Hellmann, S.; Arndt, N. DBpedia Archivo: A web-scale interface for ontology archiving under consumer-oriented aspects. In Proceedings of the Semantic Systems. In the Era of Knowledge Graphs: 16th International Conference on Semantic Systems, SEMANTiCS 2020, Amsterdam, The Netherlands, 7–10 September 2020; Proceedings 16; Springer International Publishing: Berlin/Heidelberg, Germany, 2020; pp. 19–35.
71. Vandenbussche, P.Y.; Atemez, G.A.; Poveda-Villalón, M.; Vatan, B. Linked Open Vocabularies (LOV): A gateway to reusable semantic vocabularies on the Web. *Semant. Web* **2017**, *8*, 437–452. [[CrossRef](#)]
72. Zablioth, F.; Antoniou, G.; d'Aquin, M.; Flouris, G.; Kondylakis, H.; Motta, E.; Plexousakis, D.; Sabou, M. Ontology evolution: A process-centric survey. *Knowl. Eng. Rev.* **2015**, *30*, 45–75. [[CrossRef](#)]
73. Sowiński, P.; Ganzha, M.; Paprzycki, M. RiverBench: An Open RDF Streaming Benchmark Suite. *arXiv* **2023**, arXiv:2305.06226.
74. Knublauch, H.; Kontokostas, D. Shapes Constraint Language (SHACL). W3C Recommendation, W3C. 2017. Available online: <https://www.w3.org/TR/2017/REC-shacl-20170720/> (accessed on 28 June 2024).
75. Kuhn, T.; Barbano, P.E.; Nagy, M.L.; Krauthammer, M. Broadening the scope of nanopublications. In Proceedings of the The Semantic Web: Semantics and Big Data: 10th International Conference, ESWC 2013, Montpellier, France, 26–30 May 2013; Proceedings 10; Springer: Berlin/Heidelberg, Germany, 2013; pp. 487–501.
76. Kuhn, T.; Taelman, R.; Emonet, V.; Antonatos, H.; Soiland-Reyes, S.; Dumontier, M. Semantic micro-contributions with decentralized nanopublication services. *PeerJ Comput. Sci.* **2021**, *7*, e387. [[CrossRef](#)]
77. Wijkstra, M.; Lek, T.; Kuhn, T.; Welbers, K.; Steijaert, M. Living literature reviews. In Proceedings of the 11th Knowledge Capture Conference, Virtual Event, 2–3 December 2021; pp. 241–248.
78. RDFLib Contributors. RDFLib 7.0.0. Available online: <https://rdflib.readthedocs.io/en/stable/> (accessed on 17 April 2024).

79. Sowiński, P.; Wasielewska-Michniewska, K.; Ganzha, M.; Paprzycki, M.; Bădică, C. Ontology Reuse: The Real Test of Ontological Design. In Proceedings of the New Trends in Intelligent Software Methodologies, Tools and Techniques, Kitakyushu, Japan, 20–22 September 2022; IOS Press: Amsterdam, The Netherlands, 2022. [CrossRef]
80. Fokoue, A.; Wu, Z.; Motik, B.; Horrocks, I.; Grau, B.C. OWL 2 Web Ontology Language Profiles (Second Edition). W3C Recommendation, W3C. 2012. Available online: <https://www.w3.org/TR/2012/REC-owl2-profiles-20121211/> (accessed on 17 April 2024).
81. Glimm, B.; Horrocks, I.; Motik, B.; Stoilos, G.; Wang, Z. HermiT: An OWL 2 Reasoner. *J. Autom. Reason.* **2014**, *53*, 245–269. [CrossRef]
82. Poveda-Villalón, M.; Gómez-Pérez, A.; Suárez-Figueroa, M.C. OOPS! (Ontology Pitfall Scanner!): An On-line Tool for Ontology Evaluation. *Int. J. Semant. Web Inf. Syst.* **2014**, *10*, 7–34. [CrossRef]
83. Poveda-Villalón, M.; Espinoza-Arias, P.; Garijo, D.; Corcho, O. Coming to terms with FAIR ontologies. In Proceedings of the International Conference on Knowledge Engineering and Knowledge Management, Bolzano, Italy, 16–20 September 2020; Springer: Berlin/Heidelberg, Germany, 2020; pp. 255–270.
84. Wilkinson, M.D.; Dumontier, M.; Aalbersberg, I.J.; Appleton, G.; Axton, M.; Baak, A.; Blomberg, N.; Boiten, J.W.; da Silva Santos, L.B.; Bourne, P.E.; et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci. Data* **2016**, *3*, 160018. [CrossRef] [PubMed]
85. Garijo, D.; Corcho, O.; Poveda-Villalón, M. FOOPS!: An Ontology Pitfall Scanner for the FAIR principles. In Proceedings of the ISWC (Posters/Demos/Industry), Virtual Conference, 24–28 October 2021.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.