A Flexible Framework for Defining, Representing and Detecting Changes on the Data Web

Yannis Roussakis
FORTH-ICS
rousakis@ics.forth.gr

Ioannis Chrysakis
FORTH-ICS
hrysakis@ics.forth.gr

Kostas Stefanidis
FORTH-ICS
kstef@ics.forth.gr

Giorgos Flouris
FORTH-ICS
fgeo@ics.forth.gr

Yannis Stavrakas
ATHENA-IMIS
yannis@imis.athena-innovation.gr

ABSTRACT

The dynamic nature of Web data gives rise to a multitude of problems related to the identification, computation and management of the evolving versions and the related changes. In this paper, we consider the problem of change recognition in RDF datasets, i.e., the problem of identifying, and when possible give semantics to, the changes that led from one version of an RDF dataset to another. Despite our RDF focus, our approach is sufficiently general to engulf different data models that can be encoded in RDF, such as relational or multi-dimensional. In fact, we propose a flexible, extendible and data-model-independent methodology of defining changes that can capture the peculiarities and needs of different data models and applications, while being formally robust due to the satisfaction of the properties of completeness and unambiguity. Further, we propose an ontology of changes for storing the detected changes that allows automated processing and analysis of changes, cross-snapshot queries (spanning across different versions), as well as queries involving both changes and data. To detect changes and populate said ontology, we propose a customizable detection algorithm, which is applicable to different data models and applications requiring the detection of custom, user-defined changes. Finally, we provide a proof-of-concept application and evaluation of our framework for different data models.

1. INTRODUCTION

With the growing complexity of the WWW, we face a completely different way for creating, disseminating and consuming big volumes of information. Large-scale corporate, government, or even user-generated data are published and become available to a wide spectrum of users. DBpedia, Freebase, YAGO and Atlas are, among many others, examples of large data repositories, which store information about various entities, including their relationships. Typically, data in such datasets are represented using the RDF model [11], in which information is stored in triples of the form (*subject*, *predicate*, *object*), meaning that *subject* is related to *object* via *predicate*.

Dynamicity is an indispensable part of the current web; datasets are constantly evolving for several reasons, such as the inclusion of new experimental evidence or observations, or the correction of erroneous conceptualizations [20]. As an example, consider the detected number of changes between the versions 12.07 and 13.05, and 13.05 and 13.07 of Atlas, which are 879.5M and 801.2M triples, respectively, while between the versions 3.7 and 3.8, and 3.8 and 3.9 of the English DBpedia are 20.7M and 9.3M triples (Figure 7).

This constant evolution poses several research problems, which are related to the identification, computation, storage and management of the evolving versions. An important question is how to support complex changes, whose constituent changes are seemingly unrelated and may occur on disparate pieces of data, but together as a whole they have a semantically coherent meaning for an application domain. Clearly, understanding and detecting the changes of evolving datasets is a prerequisite to address these problems. In particular, finding the differences (deltas) between datasets has been proved to play a crucial role in various curation tasks, such as the synchronization of autonomously developed datasets versions [3], or the visualization of the evolution history of a dataset [12]. Deltas are also necessary in certain applications that require access to previous versions of a dataset to support historical or cross-snapshot queries (e.g., [18]), in order, for example, to identify past states of the dataset, understand the evolution process, or detect the source of errors in the current modelling.

In general, there are two ways to record the occurring changes. The first is by constantly *monitoring* the dataset and logging any change, which requires a closed and controlled system, where all changes pass through a dedicated application that records every change as it happens. A more flexible method includes *detecting* a posteriori the changes that happened. This avoids the need to have a controlled system, and allows remote users of a dataset to identify changes, even if they have no access to the actual change process or knowledge that it happened. Here, we focus on the latter case, which is more challenging and interesting, even though most of the proposed solutions (namely, the definition of a language of changes and the representation of the changes) are applicable in both scenarios.

Recognition of changes is based on 3 main pillars, namely *defining*, *representing* and *detecting* changes. The first pillar (defining changes) is related to the definition of a *language of changes*, which is based on the set of changes that are understandable by the system. Defining the language of changes includes identifying the types of changes that will be detected, their formal definition (e.g., their se-

mantics and parameters), and takes into account the need for complex, domain-specific changes. It has been argued that the language of changes as a whole, should be well-behaved in the sense of satisfying certain properties, namely being *complete and unambiguous*, so as to allow the generation of unique deltas in a deterministic manner [14]. The definition of changes is given in Section 3.

The second pillar (representing changes) is related to the representation scheme used for storing the detected changes. This is necessary to allow a persistent representation and storage of the detected changes, as well as to support navigation among versions, analysis of the deltas, cross-snapshot or historic queries, and the raising of changes as first class citizens in a multi-version repository. Our approach is based on the definition of an *ontology of changes* that satisfies these goals. The representation scheme of the proposed changes is given in Section 4.

The third pillar (detecting changes) is related to the algorithmic definition of the detection process. This process identifies the changes applied between any two versions, based on the actual change definitions in the language of changes, and stores the results in the ontology of changes. Our approach for change detection is based on the execution of appropriately defined SPARQL queries, whose answers determine the detected changes. Details on the detection process are given in Section 5.

A major challenge related to change detection and deltas is that no language of changes is suitable for all different applications and data models. There are two reasons for that. First, the different types of data available on the Web are often represented using different models, which call for different changes. Second, different uses (or users) of the data may require a different set of changes being reported, or the definition of special types of changes that happen often or are important for the operation of the data-driven application. Therefore, there is no one-size-fits-all solution for the problem of change recognition.

In this work, we propose a flexible framework for change recognition that can be applied to different data models and applications. Our intention is not to provide a single solution, but a generic methodology for defining the three components of a multi-version repository, through which different solutions, suitable for different needs, can be developed. Even though our framework is designed for RDF, our approach is applicable to any data model representable in RDF, i.e., RDF is used as a unifying underlying model to achieve this generality.

In overall, our contributions are as follows:

- Regarding the definition of changes, the objective of generality is met by providing two different types of changes, namely *simple* and *complex* (see Section 3). The former type is meant to capture fine-grained changes for the data model at hand, and should meet the requirements of completeness and unambiguity. The latter type is meant to capture more coarse-grained, or specialized, changes that are useful for the application at hand; this allows a customized behaviour of the change detection process, depending on the actual needs of the application. Complex changes are totally dynamic, and can be defined even at run-time, greatly enhancing the flexibility of our approach.

- For the purposes of representation, the proposed ontology of changes is designed to be generic enough, so as to be customizable for different languages of changes. This allows us to meet the aforementioned generality goal, as detailed in Section 4.

- The detection process is defined in a customizable manner, as the core detection algorithm is agnostic to the set of simple

or complex changes used. The customization for the actual language of changes employed is based on the definition of appropriate SPARQL queries (one per change) that come as parameters to the algorithm, thereby allowing new changes to be easily defined. Details are given in Section 5.

- An additional contribution is the application of this work for data models other than the pure RDF model, namely the multi-dimensional model that has been transformed to the RDF Data Cube vocabulary, while it could be also applied similarly to the relational model. This application is meant as a proof-of-concept that the proposed approach is indeed capable of supporting diverse needs, and is not exhaustive in terms of the types of applications or models that our framework supports.

- Finally, we provide experiments showing that the most crucial parameter in the change recognition process is the total number of detected changes rather than the size of the examined datasets. We prove also that the number of simple changes is proportional to the number of triples which will be inserted into the ontology. Except from the evaluation aspect of our framework, our experiments offer a detailed look on the evolution of real-world, big datasets, as we record changes and provide an analysis of their form.

The approach proposed in this paper extends our previous work on change detection [14], by providing a more generic change definition framework; here, we experience significantly improved performance (about 1 order of magnitude), scalability, as well as increased generality and applicability.

2. PRELIMINARIES

We consider two disjoint sets $\mathbb{U}$, $\mathbb{L}$, denoting the *URIs* and *literals* (we ignore here blank nodes that can be avoided when data are published according to the Linked Data paradigm). An *RDF triple* [11] is a tuple of the form *(subject, predicate, object)* and asserts the fact that *subject* is associated with *object* through *predicate*. The set $\mathbb{T} = \mathbb{U} \times \mathbb{U} \times (\mathbb{U} \cup \mathbb{L})$ is the set of all RDF triples. An *RDF dataset* $\mathcal{D}$ consists of a set of RDF triples. In the following, we denote by D_{old}, D_{new} the old and new versions, respectively, of a dataset $\mathcal{D}$.

SPARQL 1.1 [17] is the official W3C recommendation language for querying RDF graphs. The building block of a SPARQL statement is a *triple pattern tp* that is like an RDF triple, but may have a *variable* (prefixed with character ?) in any of its *subject*, *predicate*, or *object* positions; variables are taken from an infinite set of variables $\mathbb{V}$, disjoint from the sets $\mathbb{U}$, $\mathbb{L}$, so the set of triple patterns is: $\mathbb{TP} = (\mathbb{U} \cup \mathbb{V}) \times (\mathbb{U} \cup \mathbb{V}) \times (\mathbb{U} \cup \mathbb{L} \cup \mathbb{V})$. SPARQL triple patterns can be combined into *graph patterns gp*, using operators like *join* ("."), *optional* (OPTIONAL) and *union* (UNION) [1] and may also include *conditions* (using FILTER). In this work, we are only interested in SELECT SPARQL queries, which are of the form: "$SELECT\ v_1, \ldots, v_n\ WHERE\ gp$", where $n > 0$, $v_i \in \mathbb{V}$ and gp is a graph pattern.

For the evaluation of SPARQL queries, we follow the semantics discussed in [15, 1]. Evaluation is based on *mappings*, which are partial functions $\mu : \mathbb{V} \mapsto \mathbb{U} \cup \mathbb{L}$ that associate variables with URIs or literals (abusing notation, $\mu(tp)$ is used to denote the result of replacing the variables in tp with their assigned values according to μ). Then, the evaluation of a SPARQL triple pattern tp on a dataset $\mathcal{D}$ returns a set of mappings (denoted by $[[tp]]^{\mathcal{D}}$) such that $\mu(tp) \in \mathcal{D}$ for $\mu \in [[tp]]^{\mathcal{D}}$. This idea is extended to graph patterns by considering the semantics of the various operators (e.g.,

$[[tp_1 \; UNION \; tp_2]]^{\mathcal{D}} = [[tp_1]]^{\mathcal{D}} \cup [[tp_2]]^{\mathcal{D}})$. Given a SPARQL query "$SELECT \; v_1, \ldots, v_n \; WHERE \; gp$", its result when applied on $\mathcal{D}$ is $(\mu(v_1), \ldots, \mu(v_n))$ for $\mu \in [[gp]]^{\mathcal{D}}$. For the precise semantics and further details on the evaluation of SPARQL queries, the reader is referred to [15, 1].

3. DEFINING CHANGES

Our purpose in this work is to provide a change recognition method, which, given two (subsequent) dataset versions D_{old}, D_{new}, would produce their *delta* (Δ), i.e., a formal description of the changes that were made to get from D_{old} to D_{new}. A delta is based on a *language of changes* ($\mathcal{L}$), i.e., a set of formal definitions of the changes that the delta could contain and, subsequently, the change detection method understands and detects; these changes should correspond to the *evolution primitives* of the data model under consideration.

A language of changes, in its simplest form, consists of additions and deletions of elements (i.e., triples for the RDF case) from a dataset, i.e., changes of the form $Add(t)/Del(t)$. Such a delta is called a *low-level delta* [21] and can be easily computed as follows: $\Delta(D_{old}, D_{new}) = \{Add(t) \mid t \in D_{new} \setminus D_{old}\} \cup \{Del(t) \mid t \in D_{old} \setminus D_{new}\}$.

Low-level languages (and deltas) are easy to define and detect, and have several nice properties [21]; however, the representation of changes at the level of (added/deleted) triples, leads to a syntactic delta, which does not capture the semantics of a change and generates results that are not intuitive enough for the human user. For example, in the RDF context, the plain deletion of an individual (class instance) would correspond to a multitude of triple deletions (namely, all the triples that contain this URI, such as property instances); listing all these changes does not immediately convey the message that an individual was deleted, and the human observer may find it hard to decipher the intent of the actual changes that took place [14]. This problem is even more serious for other data models (e.g., multi-dimensional) that are translated into RDF; in this case, the low-level deltas, which are described in RDF jargon, need to be "translated back" into the original data model, making deciphering even more difficult.

To address this problem, *high-level deltas* have been proposed [14], which aim to describe changes at a more intuitive level, in order to make them more human-understandable. In the above example, the output would be a "delete individual" change, that immediately conveys the message that all triples that include a given URI (individual) were deleted. The main idea behind achieving this is to group low-level changes into high-level ones, under some conditions. For the purposes of this paper, we organize high-level changes in two major types, namely *simple* and *complex*, each of which represents changes with a certain granularity and role in the model. Details on these two types are given below.

3.1 Simple Changes

In general, the role of *simple changes* is to describe changes (evolution primitives) specific to the data model at hand; e.g., for the multi-dimensional model, we would expect changes like "Add-Dimension" or "Attach_Type_To_Measure". Using simple changes, the user can abstract from the syntactical and representational peculiarities of the underlying data model (including its possible translation to RDF format), thereby making deltas more intuitive.

A simple change (e.g., Attach_Type_To_Measure(m,t)), is composed of the *change name* (i.e., Attach_Type_To_Measure) and the *change parameters* (i.e., (m,t)). Its main characteristic is the triples that would be added/deleted from the RDF representation of the dataset when such a change occurs. These correspond to the triples

that are directly associated with said change and are assumed to be captured by the simple change. For example, the above change would be associated with the triple $(m, \mathsf{rdfs : range}, t)$ (which indicates that t is the type of measure m); if said triple is found in D_{new}, but not in D_{old}, then this indicates that a new type (t) has been attached to measure m (thus, Attach_Type_To_Measure(m,t) should be detected).

Triples of this type are required to be in one version but not in the other (i.e., in low-level delta) and are directly associated (*consumed*) by the corresponding simple change. Consumption in this respect means that the corresponding low-level change is captured (described) by said simple change.

A simple change can also have a number of logical conditions required for detection. For example, the triple $(m, \mathsf{rdfs : range}, t)$ may also indicate that a datatype (t) is attached to a dimension m. To determine whether the inclusion of such a triple in D_{new} corresponds to an Attach_Type_To_Measure(m,t) change (as opposed to an Attach_Datatype_To_Dimension(m,t) change), we need a condition that will determine whether m is a measure or a dimension. In this case, the Attach_Type_To_Measure(m,t) would require the presence of the triple $(m, \mathsf{rdf : type}, qb : measureProperty)$ in D_{new}. Note that, in general, conditions may refer to both the old and the new version. Unlike consumed triples, the triples appearing in conditions are not necessary added or deleted triples; their presence is necessary in either (or both) of the versions and their role is to disambiguate between similar changes.

More formally, a simple change is defined as follows:

DEFINITION 1. *A simple change* $c(p_1, \ldots, p_n)$ *is defined as a tuple of the form* $\langle \delta^+, \delta^-, \phi_{old}, \phi_{new} \rangle$ *where:*

- *c is the name and* $p_1, \ldots, p_n \in \mathbb{V}$, $n \geq 0$, *are the parameters of the change,*

- $\delta^+, \delta^- \subseteq \mathbb{TP}$ *are called the consumed added and consumed deleted triples, respectively, and are sets of triple patterns,*

- ϕ_{old}, ϕ_{new} *are graph patterns, called the conditions related to* D_{old}, D_{new}, *respectively.*

In our running example, Attach_Type_To_Measure(m,t) has two parameters (m, t) and $\delta^+ = \{(m, \mathsf{rdfs : range}, t)\}$, $\delta^- = \emptyset$, $\phi_{old} =$ " ", $\phi_{new} =$ "$(m, \mathsf{rdf : type}, qb : measureProperty)$".

The structure of Definition 1 is used for defining the changes that the language of changes accepts, but any actual detection will give specific values (URIs or literals) to the parameters m, t of the change (e.g., Attach_Type_To_Measure(dm-measure:meas7v8t,dm-type:int)). This is captured with the following notion:

DEFINITION 2. *Consider a change* $c(p_1, \ldots, p_n)$. *Then, an assignment* $x_1, \ldots, x_n$ *of URIs/literals to variables* $p_1, \ldots, p_n$ *is called an instantiation of c and denoted by* $c(x_1, \ldots, x_n)$.

Now we are in position to formally define the detectability of a change instantiation. Intuitively, a change instantiation corresponds to a certain assignment of (some of) the variables in δ^+, δ^-, ϕ_{old}, ϕ_{new}; the assignment should be such that the conditions (ϕ_{old}, ϕ_{new}) are true in the underlying datasets, and the triples that correspond to the triple patterns in δ^+, δ^- are found in the low-level delta. More precisely:

DEFINITION 3. *A change instantiation* $c(x_1, \ldots, x_n)$ *of a simple change* $c(p_1, \ldots, p_n)$ *is detectable for the pair* D_{old}, D_{new} *iff there is a* $\mu \in [[\phi_{old}]]^{D_{old}} \cap [[\phi_{new}]]^{D_{new}}$ *such that for all* $tp \in \delta^+$: $\mu(tp) \in D_{new} \setminus D_{old}$ *and for all* $tp \in \delta^-$: $\mu(tp) \in D_{old} \setminus D_{new}$ *and for all i:* $\mu(p_i) = x_i$.
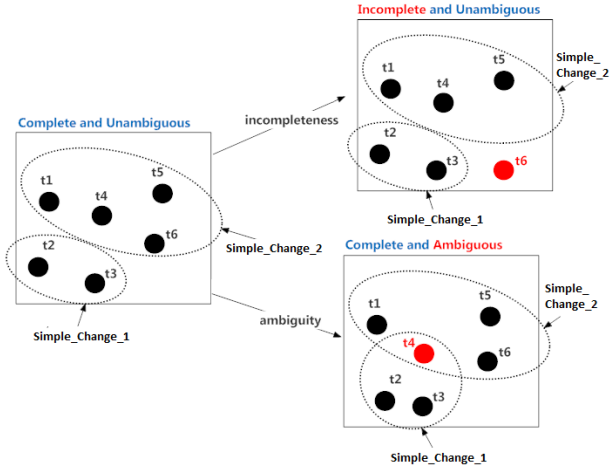
Figure 1: Visualization of Completeness and Unambiguity

Then, a detectable change *consumes* the triples that correspond to the triple patterns in δ^+, δ^-. Formally:

DEFINITION 4. *A detectable change instantiation $c(x_1, \ldots, x_n)$ of a simple change $c(p_1, \ldots, p_n)$ consumes $t \in D_{new} \setminus D_{old}$ (respectively, $t \in D_{old} \setminus D_{new}$) iff there is a $\mu \in [[\phi_{old}]]^{D_{old}} \cap [[\phi_{new}]]^{D_{new}}$ and a $tp \in \delta^+$ (respectively, $tp \in \delta^-$) such that $\mu(tp) = t$ and for all i: $\mu(p_i) = x_i$.*

The concept of consumption represents the fact that low-level changes are "assigned" to simple ones, essentially allowing a grouping (partitioning) of low-level changes into simple ones. To fulfil its purpose, this "partitioning" should be perfect, in the sense that all low-level changes should be associated to one, and only one, simple change. This is captured by the properties of *completeness* and *unambiguity*. Formally:

DEFINITION 5. *Consider a set of simple changes C. This set is called complete iff for any pair of versions D_{old}, D_{new} and for all $t \in (D_{new} \setminus D_{old}) \cup (D_{old} \setminus D_{new})$, there is an instantiation $c(x_1, \ldots, x_n)$ of some $c \in C$ such that $c(x_1, \ldots, x_n)$ is detectable and consumes t.*

DEFINITION 6. *Consider a set of simple changes C. This set is called unambiguous iff for any pair of versions D_{old}, D_{new} and for all $t \in (D_{new} \setminus D_{old}) \cup (D_{old} \setminus D_{new})$, if $c, c' \in C$ and $c(x_1, \ldots, x_n), c'(x_1', \ldots, x_m')$ are detectable and consume t, then $c(x_1, \ldots, x_n) = c'(x_1', \ldots, x_m')$.*

In a nutshell, completeness guarantees that all low level changes are associated with at least one simple change, thereby making the reported delta complete (i.e., not missing any change); unambiguity guarantees that no race conditions will emerge between simple changes attempting to consume the same low level change (see Figure 1 for a visualization of the notions of completeness and unambiguity). The combination of these two properties guarantees that the delta is produced in a deterministic manner and that it will properly reflect the changes that were actually performed.

3.2 Complex Changes

To guarantee completeness and unambiguity, simple changes must be defined at design time, and remain immutable between different applications; for changes corresponding to custom, application-specific evolution primitives, we use *complex changes*. Complex

changes can also be used to define changes that were not foreseen at design-time. This would give an extra flexibility to the user to describe more specialized, coarse-grained changes, which are relevant to a certain application, but not generic enough to be considered part of the "core" language of (simple) changes. This essentially allows extending the language of changes at run-time.

Complex changes are defined in a way similar to simple ones; there are certain differences though. First, complex changes are not directly associated with low level changes; instead of δ^+, δ^-, they include a set of simple changes, denoted by δ^s, which corresponds to the set of simple changes that should be detectable in order for said complex change to be detectable. This is necessary to make their definition more user-friendly, given that complex changes will be defined by the end-user who does not understand low level changes. Second, as complex changes can be freely defined by the user, it would be unrealistic to assume that they will have any quality guarantees, such as completeness or unambiguity. As a consequence, the detection process may lead to non-deterministic consumption of simple changes and conflicts; to avoid this, complex changes are associated with a *priority level*, which is used to resolve such conflicts.

Furthermore, complex changes support *associations*, i.e., correspondences between URIs/literals that are necessary to support changes like renames and merge/splits [14]. In general, an association α is a pair (X, Y) where X, Y are sets of URIs, literals or variables (called *URI/literal/variable association*, respectively). Associations can have one of the following forms (where v_i may be a URI, literal or variable):

- $\{v_1\} \rightsquigarrow \{v_2\}, v_1 \neq v_2$ (rename)

- $\{v_0\} \rightsquigarrow \{v_1, \ldots, v_n\}, v_i \neq v_j$ for $i \neq j$ (split)

- $\{v_1, \ldots, v_n\} \rightsquigarrow \{v_0\}, v_i \neq v_j$ for $i \neq j$ (merge)

A set of URI and literal associations should be provided as an input to the detection process. Such associations could be provided manually by the user, or automatically identified using, e.g., entity matching software [19]; detecting associations is out of the scope of this paper. Given a mapping μ and a variable association α, we will denote by $\mu(\alpha)$ the URI/literal association resulting from replacing all variables in α with their corresponding URI/literal according to μ.

Complex changes are useful in several applications. For example, various ontological applications refrain from deleting classes, but use a special subsumption relationship to a class "Obsolete" to indicate that a certain class (say, cl) should no longer be used. This action could be captured by the change Mark_as_Obsolete(cl), which would consume the simple change Add_SuperClass(cl, obs), where $obs = geneontology : ObsoleteClass$.

The formal concepts associated with complex changes are similar to the ones related to simple changes:

DEFINITION 7. *A complex change $c(p_1, \ldots, p_n)$ is defined as a tuple of the form $\langle \delta^s, \phi_{old}, \phi_{new}, \mathcal{A}, P \rangle$ where:*

- *c is the name and $p_1, \ldots, p_n \in \mathbb{V}, n \geq 0$, are the parameters of the change,*

- *δ^s is a set of simple changes called the related simple changes,*

- *ϕ_{old}, ϕ_{new} are graph patterns, called the conditions related to D_{old}, D_{new}, respectively,*

- *$\mathcal{A}$ is a set of variable associations,*

- *P is a number corresponding to the priority level of the complex change.*

In our running example, Mark_as_Obsolete(cl) has one parameter (cl) and $\delta^s = \{Add\_SuperClass(cl, obs)\}$, $\phi_{old} =$ "obs = geneontology:ObsoleteClass", $\phi_{new} =$ " ", $\mathcal{A} = \emptyset$, $P = 2$.

The definition of complex change instantiations is identical to Definition 2. For detectability, we have a series of definitions, so as to take into account priorities:

DEFINITION 8. *A change instantiation $c(x_1, \ldots, x_n)$ of a complex change $c(p_1, \ldots, p_n)$ is initially detectable for the pair D_{old}, D_{new} and the associations $\mathcal{A}^{D_{old}, D_{new}}$ iff there is a $\mu \in [[\phi_{old}]]^{D_{old}} \cap [[\phi_{new}]]^{D_{new}}$ such that for all $c'(\mu(p_1), \ldots, \mu(p_n)) \in \delta^s$, $c'(\mu(p_1), \ldots, \mu(p_n))$ is detectable, for all $\alpha \in \mathcal{A}$, $\mu(\alpha) \in \mathcal{A}^{D_{old}, D_{new}}$ and for all i, $\mu(p_i) = x_i$.*

DEFINITION 9. *An initially detectable change instantiation $c(x_1, \ldots, x_n)$ of a complex change $c(p_1, \ldots, p_n)$ consumes a simple change instantiation $c'(x'_1, \ldots, x'_m)$ of a simple change $c'(p'_1, \ldots, p'_m)$ iff $c'(p'_1, \ldots, p'_m) \in \delta^s$ and there is a $\mu \in [[\phi_{old}]]^{D_{old}} \cap [[\phi_{new}]]^{D_{new}}$ such that for all $\alpha \in \mathcal{A}$, $\mu(\alpha) \in \mathcal{A}^{D_{old}, D_{new}}$ and for all i, $\mu(p_i) = x_i, \mu(p'_i) = x'_i$.*

DEFINITION 10. *A change instantiation $c(x_1, \ldots, x_n)$ of a complex change $c(p_1, \ldots, p_n)$ is detectable for the pair D_{old}, D_{new} and the associations $\mathcal{A}^{D_{old}, D_{new}}$ iff it is initially detectable for the pair D_{old}, D_{new} and the associations $\mathcal{A}^{D_{old}, D_{new}}$ and there is no initially detectable change instantiation of another complex change with a higher priority that consumes the same simple change.*

4. REPRESENTING CHANGES

4.1 Motivation for the Ontology of Changes

We treat changes as first-class citizens in order to be able to perform queries analysing the evolution of datasets. Further, we are interested in performing combined queries, in which both the datasets and the changes should be considered to get an answer. To achieve this, the representation of the changes that are detected on the data cannot be separated from the data itself.

For example, consider the following query: "return all countries for which the unemployment rate of their capital city increased at a rate higher than the average increase of the country as a whole between versions D_{old} and D_{new}". This query requires access to the data (to identify countries and capitals) *and* to the changes (to describe the actual increase in the rates of unemployment for each city and country). Therefore, to answer it, the changes should be stored in a structured form and their representation should include connections with the actual entities (cities or capitals) that they refer to. Note that the definition of a cross-snapshot query language that treats changes as first-class citizens is outside the scope of this paper, however this work provides a formalism on which such a query language can be based. Therefore, we propose representing changes as special entities in an RDF dataset, with connections to the actual data, so that a detectable change can be associated with the corresponding data entities that it refers to.

More specifically, we propose an *ontology of changes* for storing the detected changes, thereby allowing a supervisory look of the detected changes and their association with the entities they refer to in the actual datasets, facilitating the formulation and the answering of queries that refer to both the data and their evolution.

In said ontology, the schema describes the definition of the changes ($c(p_1, \ldots, p_n)$ – see Definitions 1, 7), whereas information on detected changes (which are change instantiations – see Definition 2)
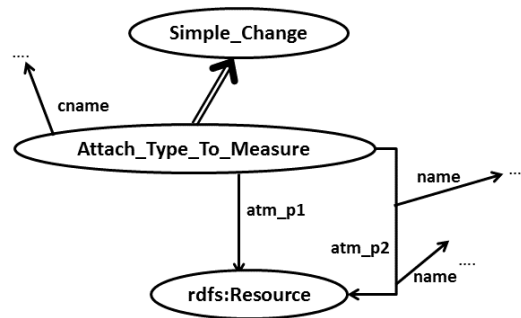

Figure 2: Representation of Simple Changes

appear at the instance level, instantiating the corresponding classes. Specifically, at schema level, we introduce one class for each simple and complex change $c(p_1, \ldots, p_n)$ that is understood and considered by the language of changes ($\mathcal{L}$), and at instance level, we introduce one individual for each detectable change $c(x_1, \ldots, x_n)$ in each pair of versions.

4.2 Ontology of Changes: Schema Level

For simple changes, the corresponding change definition is modelled as a subclass of the main class $Simple\_Change$, and is associated with adequate properties to represent its parameters (see Figure 2); each such property models the type of the parameter (e.g., whether it is a URI or a literal, via the classes rdfs:Resource, rdfs:Literal respectively) and its name (which is a descriptive name that allows a more intuitive interaction with the user, useful also during the construction of complex changes that consume said simple change). Also a descriptive name ($cname$) is captured for each type of change. Note that conditions do not need to be stored.

For complex changes, similar ideas are used; however, note that the information related to complex changes is generated on the fly at change creation time (in contrast to simple changes, which are inbuilt in the ontology at design time). In addition, more detailed information related to the change should be available in the ontology for the change detection process, because, unlike simple changes, this information is not known at design time and cannot be embedded in the code.

Each complex change is modelled as a subclass of the main class $Complex\_Change$, and is associated with adequate properties to represent its parameters (see Figure 3). Again, parameters have a type and a name, which are modelled in the same manner as in simple changes. In addition, each complex change is associated with the simple change(s) that it consumes, and includes also properties describing its (user-defined) descriptive name and its priority.

Finally, for each complex change, the SPARQL query used for its detection, is automatically generated at change definition time; this is done for efficiency, to avoid having to generate this query in every run of the detection process. The SPARQL query encapsulates all the relevant information about the detection of the change, so further information on the complex change (namely conditions and associations) need not to be stored.

4.3 Ontology of Changes: Data Level

The instance level is used to store the detectable simple and complex changes. Each detectable change between any two versions is represented using a different individual associated through adequate properties with all relevant information, namely, the versions between which the change was detected and the exact values of its
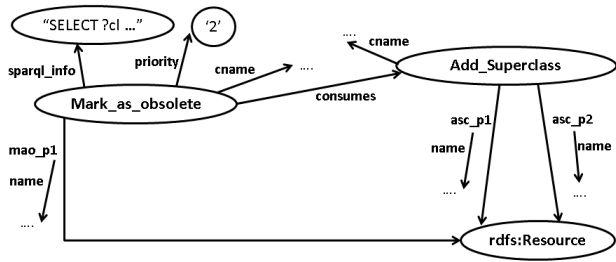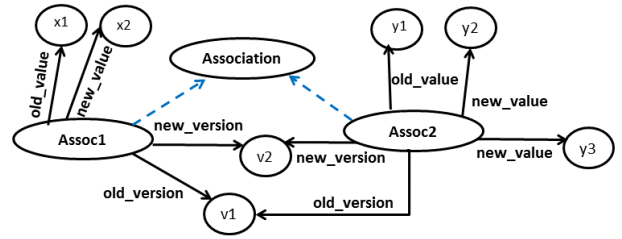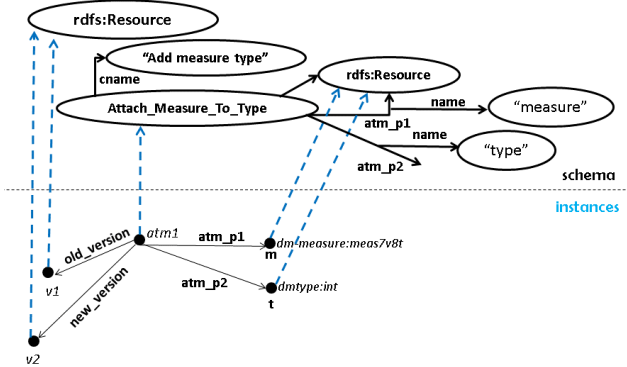
Figure 3: Representation of Complex Changes

Figure 4: Representation of Simple Change Detection

parameters (of the change instantiation). For complex changes, we additionally need to include consumption information. Examples of such instantiations for simple and complex changes are shown in Figures 4 and 5, for the detectable changes Attach_Type_To_Measure($dm - measure : meas7v8t, dm - type : int$) and Mark_as_Obsolete ($efo : EFO\_0004151$), respectively.

4.4 Ontology of Changes: Associations

Another useful feature of the ontology of changes is the storage of *associations*, which are necessary for the detection of complex changes. Associations are stored as instances classified under the class $Association$, and record the versions between which they are applicable, as well as the old and new values in the corresponding
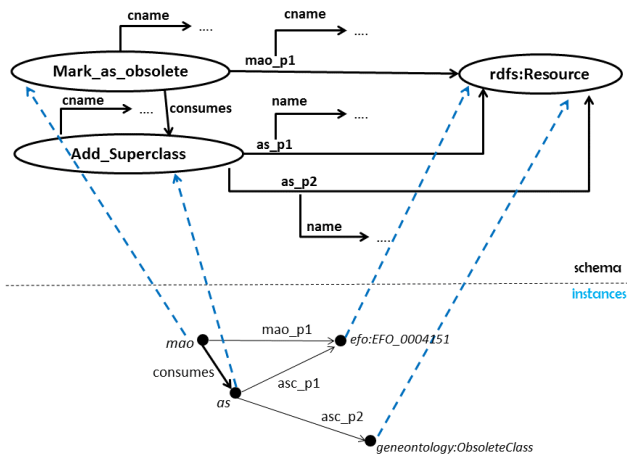

Figure 5: Representation of Complex Change Detection

Figure 6: Representation of Associations

association. For example, Figure 6 shows the representation of the associations $\{x1\} \rightsquigarrow \{x2\}$ and $\{y1\} \rightsquigarrow \{y2, y3\}$, which appear between versions $v1, v2$.

5. DETECTING AND STORING CHANGES

The change detection process is responsible for the detection of simple and complex changes between two dataset versions (and their corresponding associations, which are a necessary input for the detection of complex changes), as well as for the enrichment of the ontology of changes with information about the detectable changes. Thus, the process can be considered to comprise of two steps: the first is the identification of the detectable changes and the corresponding information (triples) to be inserted in the ontology of changes (*triple creation*), and the second is the actual ingestion of said information in the ontology of changes (*triple ingestion*).

To detect simple and complex changes, we rely on plain SPARQL queries, which are generated from the information drawn from the definition of the corresponding changes (Definitions 3, 10). For simple changes, this information is known at design time, so the query is loaded from a configuration file, whereas for complex changes, the corresponding query is generated once at change-creation time (run-time) and is loaded from the ontology of changes (see Figure 3). The results of the generated queries determine the change instantiations that are detectable; this information is further processed to determine the actual triples to be inserted in the ontology of changes. Recall that this depends on the actual detected change and its type (see Section 4).

More specifically, the SPARQL queries used for detecting a simple change are SELECT queries, whose returned values are the parameter values of the change; thus, for each parameter of the change definition, we put one variable in the SELECT clause. Then, the WHERE clause of the query includes the triple patterns that should (or should not) be found in each of the versions in order for a change instantiation to be detectable; more specifically, the triple patterns in δ^+ must be found in D_{new} but not in D_{old}, the triple patterns in δ^- must be found in D_{old} but not in D_{new}, and the graph patterns in ϕ_{old}, ϕ_{new} should be applied in D_{old}, D_{new}, respectively.

Let's use this methodology to construct the SPARQL query for the simple change Attach_Type_To_Measure(m,t) in which: $\delta^+ = \{(m, \mathsf{rdfs} : \mathsf{range}, t)\}, \delta^- = \emptyset, \phi_{old} = \text{"}, \phi_{new} = \text{"}(m, \mathsf{rdf} : \mathsf{type}, qb : measureProperty)\text{"}$. The corresponding SPARQL query is:

```
SELECT ?m ?t WHERE {
GRAPH D_new{ ?m rdf:type qb:MeasureProperty.
?m rdfs:range ?t.  }
FILTER NOT EXISTS
{ GRAPH D_old{ ?m rdfs:range ?t.  } }
}
```

Each result row of this query corresponds to the parameter values of one detectable change instantiation; for example, if the query returns the pair (dm-measure:meas7v8t,dm-type:int), then the change instantiation Attach_Type_To_Measure(dm-measure:meas7v8t,dm-type:int) is detectable.

The generation of the SPARQL queries for the complex changes follows a similar pattern. One difference is that the WHERE clause should also include the required associations, which are assumed to be already stored in the ontology of changes. A second important difference is that complex changes check the existence of simple changes (namely, those found in the δ^s part of the complex change definition) in the ontology of changes, rather than triples in the two versions (as is the case with simple changes detection); therefore, complex changes should be detected after the detection of simple changes and their storage in the ontology. Note also that the considered simple changes should not have been marked as "consumed" by other detectable changes of a higher priority; thus, it is important for queries associated with complex changes to be executed in a particular order, as implied by their priority.

Let's illustrate the above methodology to construct the SPARQL query for the complex change Mark_as_Obsolete(cl). This change has one parameter (cl) and $\delta^s = \{Add\_SuperClass(cl, obs)\}$, $\phi_{old} =$ "obs = geneontology:ObsoleteClass", $\phi_{new} =$ " ", $\mathcal{A} = \emptyset$, $P = 2$. The corresponding SPARQL query is:

```
SELECT ?cl WHERE {
GRAPH <changesOntology> {
?asc a co:Add_Superclass;
co:asc_p1 ?cl;
co:asc_p2 ?obs.
FILTER NOT EXISTS { ?cc co:consumes ?asc }
FILTER (?obs = "geneontology:ObsoleteClass").
}
}
```

As with simple changes, each query result corresponds to the parameter values of a detectable change instantiation; for example, if the query finds the simple change Add_Superclass(efo:EFO_0004151, geneontology:ObsoleteClass) which is not consumed by any complex change, then it will return efo:EFO_0004151, which implies that the change instantiation Mark _as_Obsolete(efo:EFO_0004151) is detectable.

Following detection, the information about the detectable (simple or complex) change instantiations must be stored in the ontology of changes. To do so, we have to process each result row to create the corresponding triple blocks, as specified in Section 4. This is done as a separate process that first stores the triple blocks in a file (on disk) and subsequently uploads them in Virtuoso using Virtuoso' bulk loading process (triple ingestion).

Note that the detection and storing of changes could be done in one step, if one used an adequately defined SPARQL update statement that identified the detectable change instantiations, created the corresponding triple blocks and inserted them in the ontology using a single statement. However, this approach turned out to be slower by 1-2 orders of magnitude, partly because it does not exploit bulk updates based on multiple threads, and also because bulk loading is much faster than inserting triples using SPARQL updates in Virtuoso.

The described process enjoys the characteristics put forward in the previous sections, namely:

Expressiveness: The algorithm supports the entire semantics of SPARQL, thereby allowing a very expressive method to define changes.

Extensibility: The addition of a new change (simple or complex) requires just the addition of the corresponding SPARQL query, with no further modifications on the source code.

Data Model Insensitivity: The process can be applied in any data model (e.g., multi-dimensional) which can be expressed in RDF (see Section 6) ignoring its internal representation; all peculiarities of the underlying data model are incorporated in the corresponding SPARQL queries, thereby allowing the code to be versatile.

Cross-Platform Character: The process requires very generic RDF data management operations (SPARQL support and RDF data import), so it can be implemented in any triple store.

6. APPLICATION OF THE FRAMEWORK

As mentioned above, our framework is flexible and generic enough to be applicable to any data model, which is transformable to the RDF format. Thus, to apply our framework to any given data model, one first has to determine how to transform said data model in RDF. The second step is the definition of the simple changes that would comprise the *language of changes* ($\mathcal{L}$); note that said changes should be complete and unambiguous for optimal results. Defining the simple changes amounts to identifying the SPARQL query that is necessary for detection, as well as the representation of the simple changes in the schema of the changes ontology; the latter can be easily automated.

Complex changes need not be defined at this point, as they are created at run-time; to facilitate the definition of complex changes one could create an adequate user-friendly interface, that could (potentially) sacrifice some of the expressive power of the complex change definition (Definition 7) in favour of user-friendliness; at any rate, the definition of such an interface is beyond the scope of this work.

Indicatively, we demonstrate, for visualization and experimentation purposes, the above process for the RDF and multi-dimensional data models; a similar methodology could be used for other data models, e.g., relational, but further applications of our framework are omitted due to space limitations.

6.1 RDF Model

The RDF data model is a popular model for exposing, sharing, and connecting pieces of data, information, and knowledge on the Semantic Web. Scientists from various areas of expertise use RDF to publish their scientific observations and measurements on the Web.

The case of RDF is straightforward in the sense that no data transformation is required. For the definition of simple changes, we used a set very similar to the so-called "basic changes" in [14]. The full list of defined simple changes, along with details on their definition and the corresponding SPARQL queries used for detection can be found in Appendix B.

6.2 Multi-dimensional Model

The multi-dimensional data model is useful for capturing statistics and information on data items along several dimensions. The bridge to RDF is achieved through the Data Cube vocabulary[1] which was proposed as a W3C recommendation to enable the publication of statistical data flows and other multi-dimensional data sets over the Web. The Data Cube vocabulary allows multi-dimensional data to enjoy all the advantages of the LOD and RDF technologies (knowledge sharing and interconnection) via its publication to

---

[1] http://www.w3.org/TR/vocab-data-cube

Table 1: Evaluated Datasets: Versions and Sizes

Dataset	Version	# Triples
RDF datasets		
ATLAS	v12.07	457.951.940
	v13.05	422.144.126
	v13.07	447.149.655
Dbpedia	v3.7	48.898.490
	v3.8	63.126.304
	v3.9	67.980.265
GO	24-03-2009	189.378
	22-09-2009	195.125
	20-04-2010	210.076
Multi-dimensional datasets		
4lqc	v1	229.338.925
	v2	243.336.594
	v3	243.449.874
1zph	v1	4.022.424
	v2	4.022.352
	v3	4.022.208
1bu4	v1	259.125
	v2	270.049
	v3	270.049

RDF, thereby allowing publishers or third parties to annotate and link to specified data, which can be flexibly combined across different datasets.

To apply our framework to the multi-dimensional model, we adopt the transformation proposed by the Data Cube vocabulary. The definition of simple changes is based on the identification of all entities of the Data Cube vocabulary (i.e., fact tables, dimensions, observations, codelists, hierarchies, measures, attributes) in order to express special or generic cases of changes that appear in each of those entities. To achieve completeness and unambiguity in our definition of simple changes, we take into account how data cube entities are connected and which sets of triple patterns denote the existence/characteristics of the corresponding entities. Again, the full list of the defined simple changes for the multi-dimensional model, along with details on their definition and the corresponding SPARQL queries used for detection can be found in Appendix C.

7. EXPERIMENTAL EVALUATION

Our framework was implemented and applied on the data models described in Section 6. The evaluation considered some representative real-world datasets of various sizes from these two data models, and its aims were the following:

- To identify the number and type of simple changes that usually occur in real world settings.

- To study the performance of our change detection process when dealing with real settings and quantify the effect of the size of the compared versions and the number of detected changes in the performance of the algorithm.

To our knowledge, this is the first time that change detection has been evaluated for datasets of this size and versatility (RDF and multi-dimensional).

7.1 Setting

For the management of linked data (e.g., storage of datasets and query execution), we worked with a scalable triple store, namely the open source version of Virtuoso Universal Server[2], v7.10.3209 (note that, our work is not bounded to any specific infrastructure or

triple-store). Virtuoso is hosted on a machine which uses an Intel Xeon E5-2630 at 2.30GHz, with 384GB of RAM running Debian Linux wheeze version, with Linux kernel 3.16.4. The system uses 7TB RAID-5 HDD configurations. From the total amount of memory, we dedicated 64GB for Virtuoso and 5GB for the implemented application. Moreover, taking into account that CPU provides 12 cores with 2 threads each, we decided to use a multi-threaded implementation; specifically, we noticed that the use of 8 threads during the creation of the RDF triples along with the ingestion process gave us optimal results for our setting. This was one more reason to select Virtuoso for our implementation, as it allows the concurrent use of multiple threads during ingestion. To eliminate the effects of hot/cold starts, cached OS information etc., each change detection process was executed 10 times and the average times were considered.

For our experimental evaluation, we used 3 RDF and 3 multi-dimensional datasets. The selected RDF datasets were ATLAS[3], a subset of the English Dbpedia[4] (consisting of article categories, instance types, labels and mapping-based properties) and GO[5]. The selected multi-dimensional datasets were provided by Datamarket and contain various statistics; in particular, 4lqc[6] contains weather measurements from the Global Historical Climatology Network, 1buh[7] contains numbers of employed persons between 15 and 64 years old taking time off over the last 12 months for family sickness or emergencies and 1zph[8] contains information about unemployment by sex, age, duration of unemployment and registration. All these measurements are taken from Eurostat. Table 1 summarizes the corresponding versions and their sizes.

7.2 Detected Changes

The first part of our evaluation studies the number and type of simple changes that appear in the evaluated datasets. The results are summarized in Figure 7, where we note the large number of changes which occurred during ATLAS evolution compared to the other datasets. This is explained by the fact that ATLAS contains experimental biological results and measurements that change over time, thus new versions are vastly different from previous ones. Moreover, note that the majority of changes (in all datasets except GO) are applied to the data level (e.g., Delete_Property_Instance), whereas in GO, we have also changes which are applied to the schema.

On the multi-dimensional model, the majority of changes in all datasets are of type Add_Dimension_Value_to_Observation and Delete_Dimension_Value_from_Observation. For both types of changes, we observe that almost the same number of changes is reported, which implies that the datasets' evolution mainly consists of changing dimension values upon observations; thus, creating a complex change to capture this pair of changes would immediately halve the number of reported changes.

7.3 Performance of Change Detection

The second part of our evaluation examines the performance of the detection process for the RDF and multi-dimensional testbeds; the results appear in Table 2. The reported performance is split into two parts, namely triple creation and triple ingestion; as explained in Section 5, the former includes the execution of the SPARQL

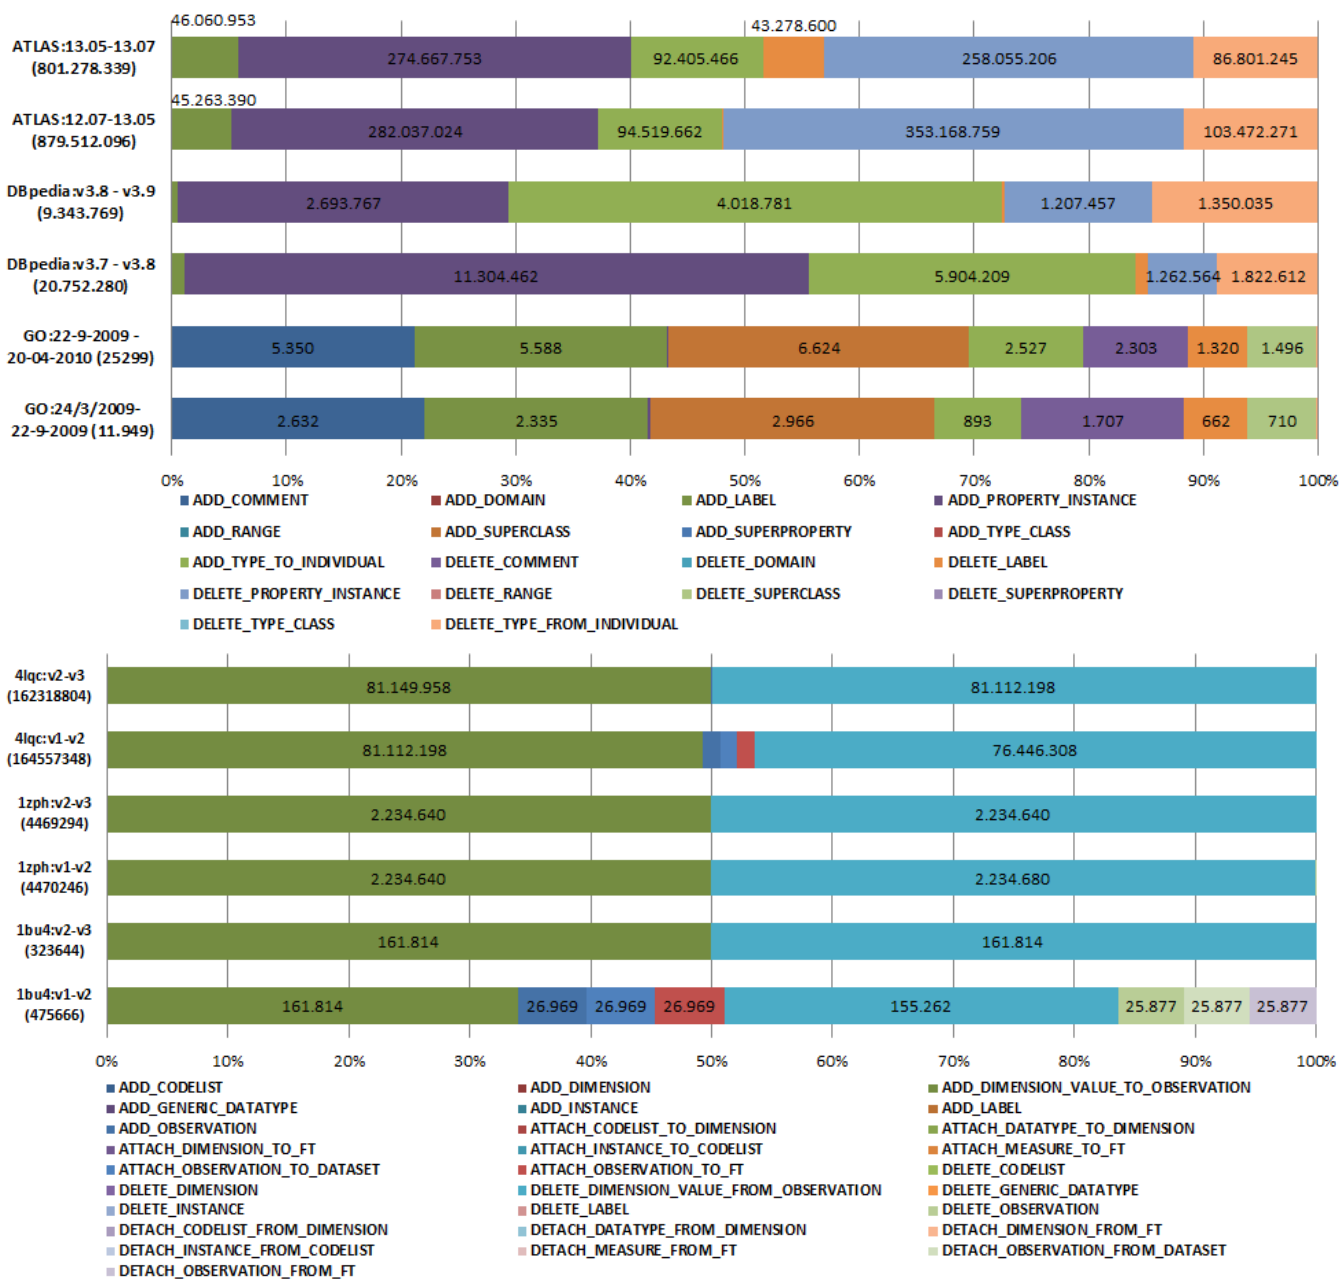**Figure 7: Detected Changes (RDF/Multi-dimensional)**

Table 2: Performance of Change Detection (RDF/Multi-dimensional)

Datasets and Versions	# Simple Changes	# Ingested Triples	Triple Creation (sec)	Triple Ingestion (sec)	Duration (sec)
RDF datasets					
ATLAS: v12.07-v13.05	879.512.096	4.980.846.857	7.294,73	9.654,46	16.949,20
ATLAS: v13.05-v13.07	801.278.339	4.539.113.055	5.953,07	9.243,61	15.196,68
DBpedia: v3.7-v3.8	20.752.280	116.327.560	465,26	143,49	608,74
DBpedia: v3.8-v3.9	9.343.769	50.620.418	296,20	72,94	369,14
GO: 24-03-2009-20-04-2010	11.949	59.179	0,80	0,50	1,31
GO: 22-09-2009-20-04-2010	25.299	124.262	0,96	0,57	1,52
Multi-dimensional datasets					
4lqc: v1-v2	164.557.348	978.012.302	1.651,03	994,85	2.645,87
4lqc: v2-v3	162.318.804	973.837.296	1.645,03	996,96	2.641,99
1zph: v1-v2	4.469.294	26.815.814	129,38	46,91	176,29
1zph: v2-v3	4.469.262	26.815.494	131,25	47,33	178,58
1bu4: v1-v2	475.666	2.642.549	7,42	6,60	14,02
1bu4: v2-v3	323.644	1.941.848	5,12	5,97	11,09

queries for detection and the identification of the triples to be inserted in the ontology of changes, whereas the latter is the actual enrichment of the ontology of changes.

The main conclusion from Table 2 is that the number of simple changes is a more crucial factor for performance than the sizes of the compared versions (see also Table 1). This observation is more clear in the DBpedia dataset, where the evolution between v3.7 and v3.8 produces about twice the number of changes than the evolution between v3.8 and v3.9; despite the fact that in the second case, we have to compare larger dataset versions, the execution time in the former case is almost twice as large. Note that this conclusion holds for both triple creation and ingestion.

More importantly, the performance of our approach is about 1 order of magnitude faster than the performance reported by [14], upon which this work was built, even if the reported simple changes are compatible in both works, as the used simple changes for the RDF model are similar to the so-called "basic changes" in [14]. For example, when comparing the performance of change detection over the GO versions v22-09-2009 and v20-04-2010, our approach needs 1,52 sec, while [14] requires 33,13 sec.

8. RELATED WORK

In general, approaches for change detection can be classified into low-level and high-level ones, based on the types of changes they support. Low-level change detection approaches report simple add/delete operations, which are not concise or intuitive enough to human users, while focusing on machine readability. [4] discusses a low-level detection approach for propositional Knowledge Bases (KBs), which can be easily extended to apply to KBs represented under any classical knowledge representation formalism. This work presents a number of desirable formal properties for change detection languages, like delta uniqueness, reversibility of changes and the ability to move backwards and forwards in the history of versions using the deltas. Similar properties appear in [21], where a low-level change detection formalism for RDFS datasets is presented, as well as in [14], upon which this work builds.

[8] describes a low-level change detection approach for the Description Logic $\mathcal{EL}$; there, the focus is on a concept-based description of changes, and the returned delta is a set of concepts whose position in the class hierarchy changed. [9] presents a formal low-level change detection approach for DL-Lite ontologies, which focuses on a semantical description of the changes. Recently, [6] introduced a scalable approach for reasoning-aware low-level change detection that uses a relational database management system, while [10] supports change detection between RDF datasets containing blank nodes. All these works result in non-concise, low-level deltas, which are difficult for a human to understand.

High-level change detection approaches provide more human-readable deltas. Although there is no agreed-upon list of changes that are necessary for any given context, various high-level operations, along with the intuition behind them, have been proposed [7, 13, 16]. However, these approaches do not present formal semantics of such operations, or of the corresponding detection process; thus, no useful formal properties can be guaranteed.

In [7, 13], a fixed-point algorithm for detecting changes, implemented in PromptDiff, is described. The algorithm incorporates heuristic-based matchers to detect changes between two versions, thus introducing uncertainty in the results, and obtaining a recall of 96% and a precision of 93%.

[16] proposes the Change Definition Language (CDL) as a means to define high-level changes. A change is defined and detected using temporal queries over a version log that contains recordings of the applied low-level changes. The version log must be updated whenever a change occurs; this overrules the use of this approach in non-curated or distributed environments. In our work, version logs are not necessary for the detection, as the delta can be produced a posteriori. [2] focuses on defining a formal way to represent high-level changes as sequences of triples, but does not describe a detection process or a specific language of changes. Finally, [5] proposes an interesting high-level change detection algorithm that takes into account the semantics of OWL.

9. CONCLUSIONS

In this paper, we proposed an approach to cope with the dynamicity of Web datasets via the management of changes between versions. We advocated in favour of a flexible, extendible and triplestore independent approach, which is suitable for any data model that is representable in RDF terms. Our approach prescribes the definition of data model-specific, as well as application-specific, changes, and their management (definition, storage, detection) in a manner that ensures the satisfaction of formal properties (like completeness and unambiguity), the flexibility and customization of the considered changes (via complex changes, which can be defined at run-time), as well as the easy configuration of a scalable detection mechanism (via a generic algorithm that builds on SPARQL queries that can be easily generated from the changes' definitions).

This work builds upon our previous work on the topic of change detection [14], by providing a more generic change definition framework, that is significantly more scalable and versatile than the one proposed in [14], as shown in Section 7.

10. ACKNOWLEDGMENTS

11. REFERENCES

[1] M. Arenas, C. Gutierrez, and J. Pérez. On the semantics of SPARQL. In *Semantic Web Information Management - A Model-Based Perspective*. Springer, 2009.

[2] S. Auer and H. Herre. A versioning and evolution framework for rdf knowledge bases. In *PSI*, 2007.

[3] R. Cloran and B. Irvin. Transmitting RDF graph deltas for a cheaper semantic Web. In *SATNAC*, 2005.

[4] E. Franconi, T. Meyer, and I. Varzinczak. Semantic diff as the basis for knowledge base versioning. In *NMR*, 2010.

[5] G. Groner, F. S. Parreiras, and S. Staab. Semantic recognition of ontology refactoring. In *ISWC*, 2010.

[6] D.-H. Im, S.-W. Lee, and H.-J. Kim. Backward inference and pruning for rdf change detection using rdbms. *J. Information Science*, 39(2):238–255, 2013.

[7] M. Klein, A. Proefschrift, M. Christiaan, A. Klein, and J. M. Akkermans. Change management for distributed ontologies. Technical report, VU University Amsterdam, 2004.

[8] B. Konev, D. Walther, and F. Wolter. The logical difference problem for description logic terminologies. In *IJCAR*, 2008.

[9] R. Kontchakov, F. Wolter, and M. Zakharyaschev. Can you tell the difference between DL-Lite ontologies? In *KR*, 2008.

[10] D.-H. Lee, D.-H. Im, and H.-J. Kim. A change detection technique for RDF documents containing nested blank nodes. In *PSI*, 2007.

[11] F. Manola, E. Miller, and B. McBride. RDF primer. www.w3.org/TR/rdf-primer, 2004.

[12] N. F. Noy, A. Chugh, W. Liu, and M. A. Musen. A framework for ontology evolution in collaborative environments. In *ISWC*, 2006.

[13] N. F. Noy and M. A. Musen. Promptdiff: A fixed-point algorithm for comparing ontology versions. In *AI*, 2002.

[14] V. Papavasileiou, G. Flouris, I. Fundulaki, D. Kotzinos, and V. Christophides. High-level change detection in RDF(S) KBs. *ACM Trans. Database Syst.*, 38(1), 2013.

[15] J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of SPARQL. In *ISWC*, 2006.

[16] P. Plessers, O. De Troyer, and S. Casteleyn. Understanding ontology evolution: A change detection approach. *Web Semant.*, 5(1):39–49, 2007.

[17] E. Prud'hommeaux, S. Harris, and A. Seaborne. SPARQL 1.1 Query Language. Technical report, W3C, 2013.

[18] K. Stefanidis, I. Chrysakis, and G. Flouris. On designing archiving policies for evolving RDF datasets on the Web. In *ER*, 2014.

[19] K. Stefanidis, V. Efthymiou, M. Herchel, and V. Christophides. Entity resolution in the Web of data. In *WWW*, 2014.

[20] J. Umbrich, M. Hausenblas, A. Hogan, A. Polleres, and S. Decker. Towards dataset dynamics: Change frequency of Linked Open Data sources. In *LDOW*, 2010.

[21] D. Zeginis, Y. Tzitzikas, and V. Christophides. On computing deltas of rdf/s knowledge bases. *ACM Trans. Web*, 5(3):14:1–14:36, 2011.

APPENDIX

A. CHANGE DETECTION ANALYSIS

A.1 Ontological Data Changes Analysis

The following tables provide the change detection analysis for all the considered datasets.

Table 3: DBpedia Dataset

	v3.7-v3.8	v3.8-v3.9
ADD_COMMENT	0	0
ADD_DOMAIN	0	0
ADD_LABEL	229805	49399
ADD_PROPERTY_INSTANCE	11304462	2693767
ADD_RANGE	0	0
ADD_SUPERCLASS	0	0
ADD_SUPERPROPERTY	0	0
ADD_TYPE_CLASS	0	0
ADD_TYPE_TO_INDIVIDUAL	5904209	4018781
DELETE_COMMENT	0	0
DELETE_DOMAIN	0	0
DELETE_LABEL	228628	24330
DELETE_PROPERTY_INSTANCE	1262564	1207457
DELETE_RANGE	0	0
DELETE_SUPERCLASS	0	0
DELETE_SUPERPROPERTY	0	0
DELETE_TYPE_CLASS	0	0
DELETE_TYPE_FROM_INDIVIDUAL	1822612	1350035

A.2 Multidimensional Data Changes Analysis

The following tables provide the change detection analysis for all the considered datasets.

Table 4: GO Dataset

	go:24/3/2009-22-9-2009	go:22-9-2009-20-04-2010
ADD_COMMENT	2632	5350
ADD_DOMAIN	0	0
ADD_LABEL	2335	5588
ADD_PROPERTY_INSTANCE	22	36
ADD_RANGE	0	0
ADD_SUPERCLASS	2966	6624
ADD_SUPERPROPERTY	0	0
ADD_TYPE_CLASS	0	0
ADD_TYPE_TO_INDIVIDUAL	893	2527
DELETE_COMMENT	1707	2303
DELETE_DOMAIN	0	0
DELETE_LABEL	662	1320
DELETE_PROPERTY_INSTANCE	0	0
DELETE_RANGE	0	0
DELETE_SUPERCLASS	710	1496
DELETE_SUPERPROPERTY	0	0
DELETE_TYPE_CLASS	0	0
DELETE_TYPE_FROM_INDIVIDUAL	22	55

Table 5: Atlas Dataset

	atlas:12.07-13.05	atlas:13.05-13.07
ADD_COMMENT	707	79
ADD_DOMAIN	42	18
ADD_LABEL	45263390	46060953
ADD_PROPERTY_INSTANCE	282037024	274667753
ADD_RANGE	48	20
ADD_SUPERCLASS	18455	5814
ADD_SUPERPROPERTY	239	18
ADD_TYPE_CLASS	12574	1813
ADD_TYPE_TO_INDIVIDUAL	94519662	92405466
DELETE_COMMENT	32	79
DELETE_DOMAIN	0	6
DELETE_LABEL	1017308	43278600
DELETE_PROPERTY_INSTANCE	353168759	258055206
DELETE_RANGE	0	6
DELETE_SUPERCLASS	1003	1112
DELETE_SUPERPROPERTY	14	16
DELETE_TYPE_CLASS	568	135
DELETE_TYPE_FROM_INDIVIDUAL	103472271	86801245

Table 6: 1bu4 Datasets

	v1-v2	v2-v3
ADD_CODELIST	1	0
ADD_DIMENSION	1	0
ADD_DIMENSION_VALUE_TO_OBSERVATION	161814	161814
ADD_GENERIC_DATATYPE	1	0
ADD_INSTANCE	6	0
ADD_LABEL	2	0
ADD_OBSERVATION	26969	0
ATTACH_CODELIST_TO_DIMENSION	1	0
ATTACH_DATATYPE_TO_DIMENSION	1	0
ATTACH_DIMENSION_TO_FT	7	7
ATTACH_INSTANCE_TO_CODELIST	6	0
ATTACH_MEASURE_TO_FT	1	1
ATTACH_OBSERVATION_TO_DATASET	26969	0
ATTACH_OBSERVATION_TO_FT	26969	0
DELETE_CODELIST	1	0
DELETE_DIMENSION	1	0
DELETE_DIMENSION_VALUE_FROM_OBSERVATION	155262	161814
DELETE_GENERIC_DATATYPE	1	0
DELETE_INSTANCE	5	0
DELETE_LABEL	2	0
DELETE_OBSERVATION	25877	0
DETACH_CODELIST_FROM_DIMENSION	1	0
DETACH_DATATYPE_FROM_DIMENSION	1	0
DETACH_DIMENSION_FROM_FT	7	7
DETACH_INSTANCE_FROM_CODELIST	5	0
DETACH_MEASURE_FROM_FT	1	1
DETACH_OBSERVATION_FROM_DATASET	25877	0
DETACH_OBSERVATION_FROM_FT	25877	0

Table 8: 4lqc Dataset

	v1 - v2	v2 - v3
ADD_DIMENSION_VALUE_TO_OBSERVATION	81112198	81149958
ADD_OBSERVATION	2332944	18880
ATTACH_DIMENSION_TO_FT	3	3
ATTACH_MEASURE_TO_FT	1	1
ATTACH_OBSERVATION_TO_DATASET	2332945	18880
ATTACH_OBSERVATION_TO_FT	2332945	18880
DELETE_DIMENSION_VALUE_FROM_OBSERVATION	76446308	81112198
DETACH_DIMENSION_FROM_FT	3	3
DETACH_MEASURE_FROM_FT	1	1

Table 7: 1zph Dataset

	v1 - v2	v2 - v3
ADD_DIMENSION_VALUE_TO_OBSERVATION	2234640	2234640
ADD_OBSERVATION	148	0
ATTACH_DIMENSION_TO_FT	6	6
ATTACH_MEASURE_TO_FT	1	1
ATTACH_OBSERVATION_TO_DATASET	148	0
ATTACH_OBSERVATION_TO_FT	148	0
DELETE_DIMENSION_VALUE_FROM_OBSERVATION	2234680	2234640
DELETE_OBSERVATION	156	0
DETACH_DIMENSION_FROM_FT	6	6
DETACH_MEASURE_FROM_FT	1	1
DETACH_OBSERVATION_FROM_DATASET	156	0
DETACH_OBSERVATION_FROM_FT	156	0

B.   SIMPLE CHANGES FOR RDF DATA MODEL

In this appendix, we list the simple changes referring to the RDF model. For each change and taking into account Definition 1 we present:

- Its name.

- The intuition it captures, described in terms of the RDF model.

- Its parameters, and the intuition behind each parameter.

- The SPARQL query which will be used for its detection.

- The consumed added and deleted triples $\delta^+, \delta^- \subseteq \mathbb{TP}$ which are essentially sets of triple patterns.

- The conditions related to D_{old}, D_{new} which form the respective graph patterns ϕ_{old}, ϕ_{new}.

For simplicity, we will write v1 to denote D_{old} and v2 to denote D_{new} in the following tables. Moreover, the conditions use clauses like "a does not appear in v1", rather than the more verbose (and formal) graph pattern: ϕ_{old}="FILTER NOT EXISTS { a ?p1 ?o1. ?s2 ?p2 a. ?s3 a ?o3}".

Change	**Add_Type_Class**(a)	**Delete_Type_Class**(a)
Intuition	Add object a of type rdfs : class	Delete object a of type rdfs : class
Parameters	a = The added object	a = The deleted object
SPARQL used for detection	```SELECT ?a WHERE {
GRAPH <v2> {
?a rdf:type rdf:Class. }
FILTER NOT EXISTS {
GRAPH <v1> {
?a rdf:type rdf:Class. }
} }```	```SELECT ?a WHERE {
GRAPH <v1> {
?a rdf:type rdf:Class.}
FILTER NOT EXISTS {
GRAPH <v2> {
?a rdf:type rdf:Class.}
} }```
δ^+	(a, rdf : type, rdfs : class)	$\emptyset$
δ^-	$\emptyset$	(a, rdf : type, rdfs : class)
ϕ_{old}	a does not appear in v1	−
ϕ_{new}	−	a does not appear in v2

Change	**Add_Type_Property**(a)	**Delete_Type_Property**(a)
Intuition	Add object a of type rdf : property	Delete object a of type rdf : property
Parameters	a = The added object	a = The deleted object
SPARQL used for detection	```SELECT ?a WHERE {
GRAPH <v2> {
?a rdf:type rdf:Property. }
FILTER NOT EXISTS {
GRAPH <v1> {
?a rdf:type rdf:Property. }
} }```	```SELECT ?a WHERE {
GRAPH <v1> {
?a rdf:type rdf:Property.}
FILTER NOT EXISTS {
GRAPH <v2> {
?a rdf:type rdf:Property.}
} }```
δ^+	(a, rdf : type, rdf : property)	$\emptyset$
δ^-	$\emptyset$	(a, rdf : type, rdf : property)
ϕ_{old}	a does not appear in v1	−
ϕ_{new}	−	a does not appear in v2

The changes *Add_Type_Individual* and *Delete_Type_Individual* are defined analogously with the exception that (a, rdf : type, rdfs : resource) should be in δ^+ (δ^-) instead of (a, rdf : type, rdf : property).

Change	**Add_Superclass**(a,b)	**Delete_Superclass**(a,b)
Intuition	Parent b of class a is added	Parent b of class a is deleted
Parameters	a = The class b = The new parent	a = The class b = The old parent
SPARQL used for detection	```SELECT ?a ?b WHERE {
GRAPH <v2> {
?a rdfs:subClassOf ?b }
FILTER NOT EXISTS {
GRAPH <v1> {
?a rdfs:subClassOf ?b }
} }```	```SELECT ?a ?b WHERE {
GRAPH <v1> {
?a rdfs:subClassOf ?b }
FILTER NOT EXISTS {
GRAPH <v2> {
?a rdfs:subClassOf ?b }
} }```
δ^+	(a, rdfs : subClassOf, b)	$\emptyset$
δ^-	$\emptyset$	(a, rdfs : subClassOf, b)
ϕ_{old}	−	−
ϕ_{new}	(a, rdf : type, rdfs : class). (b, rdf : type, rdfs : class)	(a, rdf : type, rdfs : class). (b, rdf : type, rdfs : class)

Change	**Add_Superproperty**(a,b)	**Delete_Superproperty**(a,b)
Intuition	Parent b of property a is added	Parent b of property a is deleted
Parameters	a = The property b = The new parent	a = The property b = The old parent
SPARQL used for detection	```SELECT ?a ?b WHERE {
GRAPH <v2> {
?a rdfs:subPropertyOf ?b }
FILTER NOT EXISTS {
GRAPH <v1> {
?a rdfs:subPropertyOf ?b }
} }```	```SELECT ?a ?b WHERE {
GRAPH <v1> {
?a rdfs:subPropertyOf ?b }
FILTER NOT EXISTS {
GRAPH <v2> {
?a rdfs:subPropertyOf ?b }
} }```
δ^+	(a, rdfs : subPropertyOf, b)	$\emptyset$
δ^-	$\emptyset$	(a, rdfs : subPropertyOf, b)
ϕ_{old}	−	−
ϕ_{new}	(a, rdf : type, rdf : property). (b, rdf : type, rdf : property)	(a, rdf : type, rdf : property). (b, rdf : type, rdf : property)

Change	Add_Type_To_Individual(a,b)	Delete_Type_From_Individual(a,b)
Intuition	Type b of individual a is added	Type b of individual a is deleted
Parameters	a = The individual b = The new type (class)	a = The individual b = The old type (class)
SPARQL used for detection	```SELECT ?a ?b WHERE {` `GRAPH <v2> { ?a rdf:type ?b.` `FILTER (?b != rdf:Class &&` `?b != rdfs:Property &&` `?b != rdfs:Resource). }` `FILTER NOT EXISTS {` `GRAPH <v1> { ?a rdf:type ?b.` `FILTER (?b != rdf:Class &&` `?b != rdfs:Property &&` `?b != rdfs:Resource). } }` `}```	```SELECT ?a ?b WHERE {` `GRAPH <v1> { ?a rdf:type ?b.` `FILTER (?b != rdf:Class &&` `?b != rdfs:Property &&` `?b != rdfs:Resource). }` `FILTER NOT EXISTS {` `GRAPH <v2> { ?a rdf:type ?b.` `FILTER (?b != rdf:Class &&` `?b != rdfs:Property &&` `?b != rdfs:Resource). } }` `}```
δ^+	$(a, \text{rdf : type}, b)$	$\emptyset$
δ^-	$\emptyset$	$(a, \text{rdf : type}, b)$
ϕ_{old}	−	−
ϕ_{new}	$(a, \text{rdf : type, rdfs : resource})$. FILTER (?b != rdf:Class && ?b != rdfs:Property && ?b != rdfs:Resource).	$(a, \text{rdf : type, rdfs : resource})$. FILTER (?b != rdf:Class && ?b != rdfs:Property && ?b != rdfs:Resource).

Change	Add_Property_Instance(a_1,a_2,b)	Delete_Property_Instance(a_1,a_2,b)
Intuition	Add property instance of property b	Delete property instance of property b
Parameters	a_1 = The subject a_2 = The object b = The property	a_1 = The subject a_2 = The object b = The property
SPARQL used for detection	```SELECT ?a1 ?b ?a2 WHERE {` `GRAPH <v2> { ?a1 ?b ?a2.` `FILTER(?b!=rdfs:subClassOf &&` `?b!=rdfs:subPropertyOf &&` `?b!=rdf:type &&` `?b!=rdfs:comment &&` `?b!=rdfs:label &&` `?b!=rdfs:domain &&` `?b!=rdfs:range) }` `FILTER NOT EXISTS {` `GRAPH <v1> { ?a1 ?b ?a2.` `FILTER(?b!=rdfs:subClassOf &&` `?b!=rdfs:subPropertyOf &&` `?b!=rdf:type &&` `?b!=rdfs:comment &&` `?b!=rdfs:label &&` `?b!=rdfs:domain &&` `?b!=rdfs:range) } }` `}```	```SELECT ?a1 ?b ?a2 WHERE {` `GRAPH <v1> { ?a1 ?b ?a2.` `FILTER(?b!=rdfs:subClassOf &&` `?b!=rdfs:subPropertyOf &&` `?b!=rdf:type &&` `?b!=rdfs:comment &&` `?b!=rdfs:label &&` `?b!=rdfs:domain &&` `?b!=rdfs:range) }` `FILTER NOT EXISTS {` `GRAPH <v2> { ?a1 ?b ?a2.` `FILTER(?b!=rdfs:subClassOf &&` `?b!=rdfs:subPropertyOf &&` `?b!=rdf:type &&` `?b!=rdfs:comment &&` `?b!=rdfs:label &&` `?b!=rdfs:domain &&` `?b!=rdfs:range) } }` `}```
δ^+	$(a1, b, a2)$	$\emptyset$
δ^-	$\emptyset$	$(a1, b, a2)$
ϕ_{old}	−	−
ϕ_{new}	FILTER (?b != rdfs:subClassOf && ?b != rdfs:subPropertyOf && ?b != rdf:type && ?b != rdfs:comment && ?b != rdfs:label && ?b != rdfs:domain && ?b != rdfs:range)	FILTER (?b != rdfs:subClassOf && ?b != rdfs:subPropertyOf && ?b != rdf:type && ?b != rdfs:comment && ?b != rdfs:label && ?b != rdfs:domain && ?b != rdfs:range)

Change	Add_Domain(a,b)	Delete_Domain(a,b)
Intuition	Domain b of property a is added	Domain b of property a is deleted
Parameters	a = The property b = The domain	a = The property b = The domain

SPARQL used for detection	`SELECT ?a ?b WHERE {`
	`GRAPH <v2> {`
	`?a rdfs:domain ?b }.`
	`FILTER NOT EXISTS {`
	`GRAPH <v1> {`
	`?a rdfs:domain ?b } }`
	`}`
		`SELECT ?a ?b WHERE {`
		`GRAPH <v1> {`
		`?a rdfs:domain ?b }.`
		`FILTER NOT EXISTS {`
		`GRAPH <v2> {`
		`?a rdfs:domain ?b } }`
		`}`

δ^+	(a, rdfs : domain, b)	$\emptyset$
δ^-	$\emptyset$	(a, rdfs : domain, b)
ϕ_{old}	−	−
ϕ_{new}	(a, rdf : type, rdf : property)	(a, rdf : type, rdf : property)

The changes *Add_Range* and *Delete_Range* are defined analogously with the exception that (a, rdfs : range, b) should be in δ^+ (δ^-) instead of (a, rdfs : domain, b).

Change	Add_Comment(a,b)	Delete_Comment(a,b)
Intuition	Comment b of object a is added	Comment b of object a is deleted
Parameters	a = The object b = The new comment	a = The object b = The old comment

SPARQL used for detection	`SELECT ?a ?b WHERE {`
	`GRAPH <v2> {`
	`?a rdfs:comment ?b }.`
	`FILTER NOT EXISTS {`
	`GRAPH <v1> {`
	`?a rdfs:comment ?b } }`
	`}`
		`SELECT ?a ?b WHERE {`
		`GRAPH <v1> {`
		`?a rdfs:comment ?b }.`
		`FILTER NOT EXISTS {`
		`GRAPH <v2> {`
		`?a rdfs:comment ?b } }`
		`}`

δ^+	(a, rdfs : comment, b)	$\emptyset$
δ^-	$\emptyset$	(a, rdfs : comment, b)
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Add_Label(a,b)	Delete_Label(a,b)
Intuition	Label b of object a is added	Label b of object a is deleted
Parameters	a = The object b = The new comment	a = The object b = The old comment

SPARQL used for detection	`SELECT ?a ?b WHERE {`
	`GRAPH <v2> {`
	`?a rdfs:label ?b }.`
	`FILTER NOT EXISTS {`
	`GRAPH <v1> {`
	`?a rdfs:label ?b } }`
	`}`
		`SELECT ?a ?b WHERE {`
		`GRAPH <v1> {`
		`?a rdfs:label ?b }.`
		`FILTER NOT EXISTS {`
		`GRAPH <v2> {`
		`?a rdfs:label ?b } }`
		`}`

δ^+	(a, rdfs : label, b)	$\emptyset$
δ^-	$\emptyset$	(a, rdfs : label, b)
ϕ_{old}	−	−
ϕ_{new}	−	−

C. SIMPLE CHANGES FOR MULTIDIMEN-SIONAL DATA MODEL

In this appendix, we list the simple changes refering to the multi-dimensional model. SPARQL queries have been expressed in RDF Data Cube vocabulary as soon as the data have been transformed respectively to this format.

For each change and taking into account Definition 1 we present:

- Its name.

- The intuition it captures, described in terms of the RDF model.

- Its parameters, and the intuition behind each parameter.

- The SPARQL query which will be used for its detection.

- The consumed added and deleted triples $\delta^+, \delta^- \subseteq \mathbb{TP}$ which are essentially sets of triple patterns.

Change	*Add_Codelist(c)*	*Delete_Codelist(c)*
Intuition	Add a new codelist	Delete a codelist
Parameters	c = The added codelist	c = The deleted codelist

SPARQL used for detection	`SELECT ?c WHERE {`
	`GRAPH <v2> {`
	`?c a skos:ConceptScheme.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?c a skos:ConceptScheme.`
	`}`
	`}`
	`}`
		`SELECT ?c WHERE {`
		`GRAPH <v1> {`
		`?c a skos:ConceptScheme.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?c a skos:ConceptScheme.`
		`}`
		`}`
		`}`

δ^+	$(c, \mathsf{rdf:type}, skos:ConceptScheme)$	$\emptyset$
δ^-	$\emptyset$	$(c, \mathsf{rdf:type}, skos:ConceptScheme)$
ϕ_{old}	–	–
ϕ_{new}	–	–

Change	*Attach_Codelist_to_Dimension(d, c)*	*Detach_Codelist_from_Dimension(d, c)*
Intuition	Assign a codelist to a dimension	Disassociate a codelist from a dimension
Parameters	d = The dimension in which the codelist is attached c = The codelist which is attached to dimension	d = The dimension in which the codelist is detached c = The codelist which is detached from dimension

SPARQL used for detection	`SELECT ?d ?c WHERE {`
	`GRAPH <v2> {`
	`?d qb:codeList ?c.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?d qb:codeList ?c.`
	`}`
	`}`
	`}`
		`SELECT ?d ?c WHERE {`
		`GRAPH <v1> {`
		`?d qb:codeList ?c.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?d qb:codeList ?c.`
		`}`
		`}`
		`}`

δ^+	$(d, qb:codeList, c)$	$\emptyset$
δ^-	$\emptyset$	$(d, qb:codeList, c)$
ϕ_{old}	–	–
ϕ_{new}	–	–

Change	Add_Dimension(d)	Delete_Dimension(d)
Intuition	Add a new dimension (not assigned to a fact table)	Delete a dimension
Parameters	d = The added dimension	d = The deleted dimension

SPARQL used for detection	`SELECT ?d WHERE {`
	`GRAPH <v2> {`
	`?d a qb:DimensionProperty.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?d a qb:DimensionProperty.`
	`}`
	`}`
	`}`
		`SELECT ?d WHERE {`
		`GRAPH <v1> {`
		`?d a qb:DimensionProperty.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?d a qb:DimensionProperty.`
		`}`
		`}`
		`}`

δ^+	$(D, \mathsf{rdf : type}, qb : DimensionProperty)$	$\emptyset$
δ^-	$\emptyset$	$(D, \mathsf{rdf : type}, qb : DimensionProperty)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Attach_Datatype_to_Dimension(d, t)	Detach_Datatype_from_Dimension(d, t)
Intuition	Associate a datatype with an existing dimension	Disassociate a datatype from an existing dimension
Parameters	d = The dimension in which the datatype is attached t = The datatype which is attached to dimension	d = The dimension in which the datatype is detached t = The datatype which is detached from dimension

SPARQL used for detection	`SELECT ?d ?t WHERE {`
	`GRAPH <v2> {`
	`?d a qb:DimensionProperty.`
	`?d rdfs:range ?t.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?d rdfs:range ?t.`
	`}`
	`}`
	`}`
		`SELECT ?d ?t WHERE {`
		`GRAPH <v1> {`
		`?d a qb:DimensionProperty.`
		`?d rdfs:range ?t.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?d rdfs:range ?t.`
		`}`
		`}`
		`}`

δ^+	$(d, \mathsf{rdfs : range}, t)$	$\emptyset$
δ^-	$\emptyset$	$(d, \mathsf{rdfs : range}, t)$
ϕ_{old}	−	$(d, \mathsf{rdf : type}, qb : DimensionProperty)$
ϕ_{new}	$(d, \mathsf{rdf : type}, qb : DimensionProperty)$	−

Change	Attach_Attr_to_Dimension(d, attr)	Detach_Attr_from_Dimension(d, attr)
Intuition	Associate an attribute property to a dimension	Disassociate an attribute property from a dimension
Parameters	d = The dimension in which the atribute is attached $attr$ = The attribute which is attached to dimension	d = The dimension in which the attribute is detached $attr$ = The attribute which is detached from dimension

SPARQL used for detection	`SELECT ?d ?attr WHERE {`
	`GRAPH <v2> {`
	`?d a qb:DimensionProperty.`
	`?d qb:attribute ?attr.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?d qb:attribute ?attr.`
	`}`
	`}`
	`}`
		`SELECT ?d ?attr WHERE {`
		`GRAPH <v1> {`
		`?d a qb:DimensionProperty.`
		`?d qb:attribute ?attr.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?d qb:attribute ?attr.`
		`}`
		`}`
		`}`

δ^+	$(d, qb : attribute, attr)$	$\emptyset$
δ^-	$\emptyset$	$(d, qb : attribute, attr)$
ϕ_{old}	−	$(d, \mathsf{rdf : type}, qb : DimensionProperty)$
ϕ_{new}	$(d, \mathsf{rdf : type}, qb : DimensionProperty)$	−

Change	**Add_Observation(o)**	**Delete_Observation(o)**
Intuition	Add an observation entity	Delete an observation entity
Parameters	o = The observation which is added	o = The observation which is deleted
| SPARQL used for detection | ```
SELECT ?o WHERE {
GRAPH <v2> {
?o a qb:Observation.
}
FILTER NOT EXISTS { GRAPH <v1> {
?o a qb:Observation.
}
}
}
``` | ```
SELECT ?o WHERE {
GRAPH <v1> {
?o a qb:Observation.
}
FILTER NOT EXISTS { GRAPH <v2> {
?o a qb:Observation.
}
}
}
``` |
δ^+	$(o, \mathsf{rdf:type}, qb:Observation)$	$\emptyset$
δ^-	$\emptyset$	$(o, \mathsf{rdf:type}, qb:Observation)$
ϕ_{old}	–	–
ϕ_{new}	–	–

Change	**Attach_Observation_to_FT(o, ft)**	**Detach_Observation_from_FT(o, ft)**
Intuition	Associate an observation to a fact table	Disassociate an observation from a fact table
Parameters	o = The observation which is attached to fact table ft = The fact table	o = The observation which is detached from fact table ft = The fact table
| SPARQL used for detection | ```
SELECT ?o ?ft WHERE {
GRAPH <v2> {
?o qb:dataSet ?ds.
?ds qb:structure ?ft.
}
FILTER NOT EXISTS { GRAPH <v1> {
?o qb:dataSet ?ds.
?ds qb:structure ?ft.
}
}
}
``` | ```
SELECT ?o ?ft WHERE {
GRAPH <v1> {
?o qb:dataSet ?ds.
?ds qb:structure ?ft.
}
FILTER NOT EXISTS { GRAPH <v2> {
?o qb:dataSet ?ds.
?ds qb:structure ?ft.
}
}
}
``` |
δ^+	$(o, qb:dataSet, ds), (ds, qb:structure, ft)$	$\emptyset$
δ^-	$\emptyset$	$(o, qb:dataSet, ds), (ds, qb:structure, ft)$
ϕ	–	–

Change	**Add_Measure_Value_to_Observation**(o, m, v)	**Delete_Measure_Value_from_Observation**(o, m, v)
Intuition	Add value to a measure in a specific observation	Delete a value from an observation
Parameters	o = The observation m = The measure on which the observation refers v = The value of the measure	o = The observation m = The measure on which the observation refers v = The value of the measure

SPARQL used for detection	`SELECT ?o ?m ?v WHERE {`
	`GRAPH <v2> {`
	`?o qb:dataSet ?ds.`
	`?ds qb:structure ?ft.`
	`?ft qb:component ?cs.`
	`?cs qb:measure ?m.`
	`?m a qb:MeasureProperty.`
	`?m rdfs:range ?v.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?o qb:dataSet ?ds.`
	`?ds qb:structure ?ft.`
	`?ft qb:component ?cs.`
	`?cs qb:measure ?m.`
	`?m a qb:MeasureProperty.`
	`?m rdfs:range ?v.`
	`}`
	`}`
	`}`
		`SELECT ?o ?m ?v WHERE {`
		`GRAPH <v1> {`
		`?o qb:dataSet ?ds.`
		`?ds qb:structure ?ft.`
		`?ft qb:component ?cs.`
		`?cs qb:measure ?m.`
		`?m a qb:MeasureProperty.`
		`?m rdfs:range ?v.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?o qb:dataSet ?ds.`
		`?ds qb:structure ?ft.`
		`?ft qb:component ?cs.`
		`?cs qb:measure ?m.`
		`?m a qb:MeasureProperty.`
		`?m rdfs:range ?v.`
		`}`
		`}`
		`}`

δ^+	$(o, qb : dataSet, ds)$, $(ds, qb : structure, ft)$, $(ft, qb : component, cs)$, $(cs, qb : measure, m)$, $(m, \mathsf{rdfs} : \mathsf{range}, v)$	$\emptyset$
δ^-	$\emptyset$	$(o, qb : dataSet, ds)$, $(ds, qb : structure, ft)$, $(ft, qb : component, cs)$, $(cs, qb : measure, m)$, $(m, \mathsf{rdfs} : \mathsf{range}, v)$
ϕ_{old}	–	$(m, \mathsf{rdf} : \mathsf{type}, qb : MeasureProperty)$
ϕ_{new}	$(m, \mathsf{rdf} : \mathsf{type}, qb : MeasureProperty)$	–

Change	**Add_Dimension_Value_to_Observation**(o, d, v)	**Delete_Dimension_Value_from_Observation**(o, d, v)
Intuition	Add value to a dimension in a specific observation	Delete value from a dimension in an observation
Parameters	o = The observation d = The dimension on which the observation refers v = The value of the dimension	o = The observation d = The dimension on which the observation refers v = The value of the dimension

SPARQL used for detection	`SELECT ?o ?d ?v WHERE {`
	`GRAPH <v2> {`
	`?o qb:dataSet ?ds.`
	`?ds qb:structure ?ft.`
	`?ft qb:component ?cs.`
	`?cs qb:dimension ?d.`
	`?d a qb:DimensionProperty.`
	`?d rdfs:range ?v.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?o qb:dataSet ?ds.`
	`?ds qb:structure ?ft.`
	`?ft qb:component ?cs.`
	`?cs qb:dimension ?d.`
	`?d a qb:DimensionProperty.`
	`?d rdfs:range ?v.`
	`}`
	`}`
	`}`
		`SELECT ?o ?d ?v WHERE {`
		`GRAPH <v1> {`
		`?o qb:dataSet ?ds.`
		`?ds qb:structure ?ft.`
		`?ft qb:component ?cs.`
		`?cs qb:dimension ?d.`
		`?d a qb:DimensionProperty.`
		`?d rdfs:range ?v.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?o qb:dataSet ?ds.`
		`?ds qb:structure ?ft.`
		`?ft qb:component ?cs.`
		`?cs qb:dimension ?d.`
		`?d a qb:DimensionProperty.`
		`?d rdfs:range ?v.`
		`}`
		`}`
		`}`

δ^+	$(o, qb : dataSet, ds)$, $(ds, qb : structure, ft)$, $(ft, qb : component, cs)$, $(cs, qb : dimension, d)$, $(d, \mathsf{rdfs} : \mathsf{range}, v)$	$\emptyset$
δ^-	$\emptyset$	$(o, qb : dataSet, ds)$, $(ds, qb : structure, ft)$, $(ft, qb : component, cs)$, $(cs, qb : dimension, d)$, $(d, \mathsf{rdfs} : \mathsf{range}, v)$
ϕ_{old}	–	$(d, \mathsf{rdf} : \mathsf{type}, qb : DimensionProperty)$
ϕ_{new}	$(d, \mathsf{rdf} : \mathsf{type}, qb : DimensionProperty)$	–

Change	**Add_Hierarchy**(*h*)	**Delete_Hierarchy**(*h*)
Intuition	Add a new hierarchy	Delete an hierarchy
Parameters	h = The added hierarchy	h = The deleted hierarchy

SPARQL used for detection	<pre>SELECT ?h WHERE {
	GRAPH <v2> {
	?h a qb:HierarchicalCodeList.
	FILTER NOT EXISTS
	{ ?h a skos:ConceptScheme. }
	}
	FILTER NOT EXISTS { GRAPH <v1> {
	?h a qb:HierarchicalCodeList.
	FILTER NOT EXISTS
	{ ?h a skos:ConceptScheme. }
	}
	}
	}</pre>
		<pre>SELECT ?h WHERE {
		GRAPH <v1> {
		?h a qb:HierarchicalCodeList.
		FILTER NOT EXISTS
		{ ?h a skos:ConceptScheme. }
		}
		FILTER NOT EXISTS
		{ GRAPH <v2> {
		?h a qb:HierarchicalCodeList.
		FILTER NOT EXISTS
		{ ?h a skos:ConceptScheme. }
		}
		}
		}</pre>

δ^+	$(h, \mathsf{rdf:type}, qb:HierarchicalCodeList)$	$\emptyset$
δ^-	$\emptyset$	$(h, \mathsf{rdf:type}, qb:HierarchicalCodeList)$

ϕ_{old}	–	<pre>FILTER NOT EXISTS
		{ ?h a skos:ConceptScheme. }</pre>

ϕ_{new}	<pre>FILTER NOT EXISTS
	{ ?h a skos:ConceptScheme. }</pre>
		–

Change	**Attach_Hierarchy_to_Dimension**(*d*, *h*)	**Detach_Hierarchy_from_Dimension**(*d*, *h*)
Intuition	Associate an hierarchy to a dimension	Disassociate an hierarchy from a dimension
Parameters	d = The dimension in which the hierarchy is attached h = The hierarchy which is attached to dimension	d = The dimension in which the hierarchy is detached h = The hierarchy which is detached from dimension

SPARQL used for detection	<pre>SELECT ?d ?h WHERE {
	GRAPH <v2> {
	?d qb:codeList ?h.
	?h a qb:HierarchicalCodeList.
	FILTER NOT EXISTS
	{ ?h a skos:ConceptScheme. }
	}
	FILTER NOT EXISTS { GRAPH <v1> {
	?d qb:codeList ?h.
	}
	}
	}</pre>
		<pre>SELECT ?d ?h WHERE {
		GRAPH <v1> {
		?d qb:codeList ?h.
		?h a qb:HierarchicalCodeList.
		FILTER NOT EXISTS
		{ ?h a skos:ConceptScheme. }
		}
		FILTER NOT EXISTS { GRAPH <v2> {
		?d qb:codeList ?h.
		}
		}
		}</pre>

δ^+	$(d, qb:codeList, h)$	$\emptyset$
δ^-	$\emptyset$	$(d, qb:codeList, h)$
ϕ_{old}	–	$(h, \mathsf{rdf:type}, qb:HierarchicalCodeList)$
ϕ_{new}	$(h, \mathsf{rdf:type}, qb:HierarchicalCodeList)$	–

Change	**Add_Instance**(i)	**Delete_Instance**(i)
Intuition	Add a new instance	Delete an instance
Parameters	i = The added instance	i = The deleted instance
SPARQL used for detection	```SELECT ?i WHERE {` `GRAPH <v2> {` `?i a skos:Concept.` `}` `FILTER NOT EXISTS { GRAPH <v1> {` `?i a skos:Concept.` `}` `}` `}```	```SELECT ?i WHERE {` `GRAPH <v1> {` `?i a skos:Concept.` `}` `FILTER NOT EXISTS { GRAPH <v2> {` `?i a skos:Concept.` `}` `}` `}```
δ^+	$(i, \mathsf{rdf} : type, skos : Concept)$	$\emptyset$
δ^-	$\emptyset$	$(i, \mathsf{rdf} : type, skos : Concept)$
ϕ_{old}	–	–
ϕ_{new}	–	–

Change	**Attach_Instance_to_Codelist**(c, i)	**Detach_Instance_from_Codelist**(c, i)
Intuition	Associate a new instance with a codelist	Disassociate an instance from a codelist
Parameters	c = The codelist in which the instance is attached i = The instance which is attached to codelist	c = The codelist in which the instance is attached i = The instance which is attached to codelist
SPARQL used for detection	```SELECT ?c ?i WHERE {` `GRAPH <v2> {` `?i a skos:ConceptScheme.` `?c skos:inScheme ?i.` `}` `FILTER NOT EXISTS { GRAPH <v1> {` `?c skos:inScheme ?i.` `}` `}` `}```	```SELECT ?c ?i WHERE {` `GRAPH <v1> {` `?i a skos:ConceptScheme.` `?c skos:inScheme ?i.` `}` `FILTER NOT EXISTS { GRAPH <v2> {` `?c skos:inScheme ?i.` `}` `}` `}```
δ^+	$(i, skos : inScheme, c)$	$\emptyset$
δ^-	$\emptyset$	$(i, skos : inScheme, c)$
ϕ_{old}	–	$(c, \mathsf{rdf} : type, skos : ConceptScheme)$
ϕ_{new}	$(c, \mathsf{rdf} : type, skos : ConceptScheme)$	–

Change	**Attach_Instance_to_Hierarchy**(h, i)	**Detach_Instance_from_Hierarchy**(h, i)
Intuition	Associate a new instance with a hierarchy	Disassociate a new instance from a hierarchy
Parameters	h = The hierarchy in which the instance is attached i = The instance which is attached to hierarchy	h = The hierarchy in which the instance is attached i = The instance which is attached to hierarchy
SPARQL used for detection	```SELECT ?h ?i WHERE {` `GRAPH <v2> {` `?h a qb:HierarchicalCodeList.` `?i skos:inScheme ?h.` `}` `FILTER NOT EXISTS { GRAPH <v1> {` `?i skos:inScheme ?h.` `}` `}` `}```	```SELECT ?h ?i WHERE {` `GRAPH <v1> {` `?h a qb:HierarchicalCodeList.` `?i skos:inScheme ?h.` `}` `FILTER NOT EXISTS { GRAPH <v2> {` `?i skos:inScheme ?h.` `}` `}` `}```
δ^+	$(i, skos : inScheme, h)$	$\emptyset$
δ^-	$\emptyset$	$(i, skos : inScheme, h)$
ϕ_{old}	–	$(h, \mathsf{rdf} : type, qb : HierarchicalCodeList)$
ϕ_{new}	$(h, \mathsf{rdf} : type, qb : HierarchicalCodeList)$	–

Change	Attach_Instance_to_Parent(i, p)	Detach_Instance_to_Parent(i, p)
Intuition	Associate an instance with its parent	Disassociate an instance from its parent
Parameters	i = The instance p = The parent which is added to instance	i = The instance p = The parent which is deleted from instance
SPARQL used for detection	```SELECT ?i ?p WHERE {
GRAPH <v2> {
?i skos:broaderTransitive ?p.
}
FILTER NOT EXISTS { GRAPH <v1> {
?i skos:broaderTransitive ?p.
}
}
}```	```SELECT ?i ?p WHERE {
GRAPH <v1> {
?i skos:broaderTransitive ?p.
}
FILTER NOT EXISTS { GRAPH <v2> {
?i skos:broaderTransitive ?p.
}
}
}```
δ^+	$(i, skos:broaderTransitive, p)$	$\emptyset$
δ^-	$\emptyset$	$(i, skos:broaderTransitive, p)$
ϕ_{old}	–	–
ϕ_{new}	–	–

Change	Add_Measure(m)	Delete_Measure(m)
Intuition	Add a new measure	Delete a measure
Parameters	m = The measure which is added	m = The measure which is deleted
SPARQL used for detection	```SELECT ?m WHERE {
GRAPH <v2> {
?m a qb:MeasureProperty.
}
FILTER NOT EXISTS { GRAPH <v1> {
?m a qb:MeasureProperty.
}
}
}```	```SELECT ?m WHERE {
GRAPH <v1> {
?m a qb:MeasureProperty.
}
FILTER NOT EXISTS { GRAPH <v2> {
?m a qb:MeasureProperty.
}
}
}```
δ^+	$(m, \mathsf{rdf:type}, qb:MeasureProperty)$	$\emptyset$
δ^-	$\emptyset$	$(m, \mathsf{rdf:type}, qb:MeasureProperty)$
ϕ_{old}	–	–
ϕ_{new}	–	–

Change	Attach_Type_to_Measure(t, m)	Detach_Type_from_Measure(t, m)
Intuition	Associate a new datatype to a measure	Disassociate a datatype from a measure
Parameters	t = The added type to measure m = The measure in which the type is added	t = The deleted type from measure m = The measure in which the type is deleted
SPARQL used for detection	```SELECT ?m ?t WHERE {
GRAPH <v2> {
?m a qb:MeasureProperty.
?m rdfs:range ?t.
}
FILTER NOT EXISTS { GRAPH <v1> {
?m rdfs:range ?t.
}
}
}```	```SELECT ?m ?t WHERE {
GRAPH <v1> {
?m a qb:MeasureProperty.
?m rdfs:range ?t.
}
FILTER NOT EXISTS { GRAPH <v2> {
?m rdfs:range ?t.
}
}
}```
δ^+	$(m, \mathsf{rdfs:range}, t)$	$\emptyset$
δ^-	$\emptyset$	$(m, \mathsf{rdfs:range}, t)$
ϕ_{old}	–	$(m, \mathsf{rdf:type}, qb:MeasureProperty)$
ϕ_{new}	$(m, \mathsf{rdf:type}, qb:MeasureProperty)$	–

Change	**Add_Fact_Table(** ft **)**	**Delete_Fact_Table(** ft **)**
Intuition	Add a new fact table	Delete a fact table
Parameters	ft = The fact table is added	ft = The fact table which is deleted

SPARQL used for detection	<pre>SELECT ?ft WHERE {
	GRAPH <v2> {
	?ft a qb:DataStructureDefinition.
	}
	FILTER NOT EXISTS { GRAPH <v1> {
	?ft a qb:DataStructureDefinition.
	}
	}
	}</pre>
		<pre>SELECT ?ft WHERE {
		GRAPH <v1> {
		?ft a qb:DataStructureDefinition.
		}
		FILTER NOT EXISTS { GRAPH <v2> {
		?ft a qb:DataStructureDefinition.
		}
		}
		}</pre>

δ^+	$(ft, \mathsf{rdf:type}, qb:DataStructureDefinition)$	$\emptyset$
δ^-	$\emptyset$	$(ft, \mathsf{rdf:type}, qb:DataStructureDefinition)$
ϕ_{old}	$-$	$-$
ϕ_{new}	$-$	$-$

Change	**_Add_Label_**(s, o)	**_Delete_Label_**(s, o)
Intuition	Add a new label	Delete a label
Parameters	s = The subject in which the label is added o = The added label	s = The subject in which the label is deleted o = The deleted label
SPARQL used for detection	```SELECT ?s ?o WHERE {
GRAPH <v2> {
?s ?p ?o.
filter (?p = rdfs:label).
}
FILTER NOT EXISTS { GRAPH <v1> {
?s ?p ?o
}
}
}```	```SELECT ?s ?o WHERE {
GRAPH <v1> {
?s ?p ?o.
filter (?p = rdfs:label).
}
FILTER NOT EXISTS { GRAPH <v2> {
?s ?p ?o
}
}
}```
δ^+	$(a, \text{rdfs}: \text{label}, b)$	$\emptyset$
δ^-	$\emptyset$	$(a, \text{rdfs}: \text{label}, b)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Attach_Measure_to_Fact_Table(m, ft)	Detach_Measure_from_Fact_Table(m, ft)
Intuition	Associate a measure property to fact table	Disassociate a measure property from a fact table
Parameters	m = The added measure to fact table ft = The fact table in which the measure is added	m = The deleted measure to fact table ft = The fact table in which the measure is deleted
| SPARQL used for detection | ```
SELECT ?ft ?m WHERE {
GRAPH <v2> {
?ft qb:component ?cs.
?cs qb:measure ?m.
}
FILTER NOT EXISTS { GRAPH <v1> {
?ft qb:component ?cs.
?cs qb:measure ?m.
}
}
}
``` | ```
SELECT ?ft ?m WHERE {
GRAPH <v1> {
?ft qb:component ?cs.
?cs qb:measure ?m.
}
FILTER NOT EXISTS { GRAPH <v2> {
?ft qb:component ?cs.
?cs qb:measure ?m.
}
}
}
``` |
δ^+	$(ft, qb : component, cs), (cs, qb : measure, m)$	$\emptyset$
δ^-	$\emptyset$	$(ft, qb : component, cs), (cs, qb : measure, m)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Attach_Dimension_to_Fact_Table(d, fT)	Detach_Dimension_from_Fact_Table(d, ft)
Intuition	Associate a dimension property to fact table	Disassociate a dimension property from a fact table
Parameters	d = The added dimension to fact table ft = The fact table in which the dimension is added	d = The deleted dimension to fact table ft = The fact table in which the dimension is deleted
| SPARQL used for detection | ```
SELECT ?d ?ft WHERE {
GRAPH <v2> {
?ft qb:component ?cs.
?cs qb:dimension ?d.
}
FILTER NOT EXISTS { GRAPH <v1> {
?ft qb:component ?cs.
?cs qb:dimension ?d.
}
}
}
``` | ```
SELECT ?d ?ft WHERE {
GRAPH <v1> {
?ft qb:component ?cs.
?cs qb:dimension ?d.
}
FILTER NOT EXISTS { GRAPH <v2> {
?ft qb:component ?cs.
?cs qb:dimension ?d.
}
}
}
``` |
δ^+	$(ft, qb : component, cs), (cs, qb : dimension, d)$	$\emptyset$
δ^-	$\emptyset$	$(ft, qb : component, cs), (cs, qb : dimension, d)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Add_Attribute(attr)	Delete_Attribute(attr)
Intuition	Add a new attribute	Delete an attribute
Parameters	$attr$ = The attribute which is added	$attr$ = The attribute which is deleted
| SPARQL used for detection | ```
SELECT ?attr WHERE {
GRAPH <v2> {
?attr a qb:AttributeProperty.
}
FILTER NOT EXISTS { GRAPH <v1> {
?attr a qb:AttributeProperty.
}
}
}
``` | ```
SELECT ?attr WHERE {
GRAPH <v1> {
?attr a qb:AttributeProperty.
}
FILTER NOT EXISTS { GRAPH <v2> {
?attr a qb:AttributeProperty.
}
}
}
``` |
δ^+	$(attr, \mathsf{rdf} : type, qb : AttributeProperty)$	$\emptyset$
δ^-	$\emptyset$	$(attr, \mathsf{rdf} : type, qb : AttributeProperty)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Attach_Attr_to_Measure$(attr, m)$	Detach_Attr_from_Measure$(attr, m)$
Intuition	Associate an attribute with an existing measure	Disassociate an attribute from a measure
Parameters	$attr$ = The attribute which is attached to measure m = The measure which is attached	$attr$ = The attribute which is detached from measure m = The measure in which is the measure is detatched
| SPARQL used for detection | ```
SELECT ?attr ?m WHERE {
GRAPH <v2> {
?m a qb:MeasureProperty.
?m qb:attribute ?attr.
}
FILTER NOT EXISTS { GRAPH <v1> {
?m qb:attribute ?attr.
}
}
}
``` | ```
SELECT ?attr ?m WHERE {
GRAPH <v1> {
?m a qb:MeasureProperty.
?m qb:attribute ?attr.
}
FILTER NOT EXISTS { GRAPH <v2> {
?m qb:attribute ?attr.
}
}
}
``` |
δ^+	$(m, qb : attribute, attr)$	$\emptyset$
δ^-	$\emptyset$	$(m, qb : attribute, attr)$
ϕ_{old}	−	$(m, \mathsf{rdf : type}, qb : MeasureProperty)$
ϕ_{new}	$(m, \mathsf{rdf : type}, qb : MeasureProperty)$	−

Change	Attach_Observation_to_Dataset(o, ds)	Detach_Observation_from_Dataset(o, ds)
Intuition	Associate an observation with an existing dataset	Disassociate an observation from a dataset
Parameters	o = The observation which is attached to dataset ds = The dataset which in which the observation is attached	o = The observation which is detached from dataset ds = The dataset in which is the observation is detatched
| SPARQL used for detection | ```
SELECT ?o ?ds WHERE {
GRAPH <v2> {
?o qb:dataSet ?ds.
}
FILTER NOT EXISTS { GRAPH <v1> {
?o qb:dataSet ?ds.
}
}
}
``` | ```
SELECT ?o ?ds WHERE {
GRAPH <v1> {
?o qb:dataSet ?ds.
}
FILTER NOT EXISTS { GRAPH <v2> {
?o qb:dataSet ?ds.
}
}
}
``` |
δ^+	$(o, qb : dataSet, ds)$	$\emptyset$
δ^-	$\emptyset$	$(o, qb : dataSet, ds)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Add_Inscheme(x, s)	Delete_Inscheme(x, s)
Intuition	Add a scheme in a term	Delete a scheme from a term
Parameters	x = The term which is associated to the scheme s = The associated scheme	x = The term which is deleted from associated scheme s = The associated scheme

SPARQL used for detection	`SELECT ?x ?s WHERE {`
	`GRAPH <v2> {`
	`FILTER NOT EXISTS { {`
	`?s rdf:type skos:ConceptScheme.}`
	`UNION`
	`{?s rdf:type`
	`qb:HierarchicalCodeList.}`
	`}`
	`?x skos:inScheme ?s.`
	`}`
	`FILTER NOT EXISTS { GRAPH <v1> {`
	`?x skos:inScheme  ?s.`
	`}`
	`}`
	`}`
		`SELECT ?x ?s WHERE {`
		`GRAPH <v1> {`
		`FILTER NOT EXISTS {`
		`{?s rdf:type`
		`skos:ConceptScheme.}`
		`UNION`
		`{?s rdf:type`
		`qb:HierarchicalCodeList.}`
		`}`
		`?x skos:inScheme ?s.`
		`}`
		`FILTER NOT EXISTS { GRAPH <v2> {`
		`?x skos:inScheme ?s.`
		`}`
		`}`
		`}`

δ^+	$(x, skos : inScheme, s)$	$\emptyset$
δ^-	$\emptyset$	$(x, skos : inScheme, s)$

ϕ_{old}	−	`FILTER NOT EXISTS {`
		`{?s rdf:type skos:ConceptScheme.}`
		`UNION`
		`{?s rdf:type`
		`qb:HierarchicalCodeList.}`
		`}`

ϕ_{new}	`FILTER NOT EXISTS {`
	`{?s rdf:type`
	`skos:ConceptScheme.}`
	`UNION`
	`{?s rdf:type`
	` qb:HierarchicalCodeList.}`
	`}`
		−

Change	*Add_Unknown_Property(s, p, o)*	*Delete_Unknown_Property(s, p, o)*
Intuition	Add a new (unknown) property with specified subject and object related	Delete a property
Parameters	s = The subject in which the property is added p = The property o = The object which is related with the subject via the property	s = The subject in which the property is deleted p = The property o = The object which is related with the subject via the property

SPARQL used for detection	```
	SELECT ?s ?p ?o WHERE {
	GRAPH <v2> {
	{FILTER NOT EXISTS
	{?p rdfs:subPropertyOf rdfs:label}}
	UNION
	{FILTER (?p != rdfs:label).}
	?s ?p ?o.
	FILTER(?p != rdfs:label).
	FILTER(?p != rdfs:range).
	FILTER(?p != skos:inScheme).
	FILTER(?p != skos:broaderTransitive).
	FILTER(?p != qb:codeList).
	FILTER(?p != qb:component).
	FILTER(?p != qb:dimension).
	FILTER(?p != qb:measure).
	FILTER(?p != qb:attribute).
	FILTER(?p != qb:dataSet).
	FILTER(?p != qb:structure).
	FILTER(?p != rdf:type
	|| skos:Concept).
	FILTER (?p != rdf:type
	|| skos:ConceptScheme).
	FILTER (?p != rdf:type
	|| qb:AttributeProperty).
	FILTER (?p != rdf:type
	|| qb:CodedProperty).
	FILTER (?p != rdf:type
	|| qb:DimensionProperty).
	FILTER (?p != rdf:type
	|| qb:DataStructureDefinition).
	FILTER (?p != rdf:type
	|| qb:HierarchicalCodeList).
	FILTER (?p != rdf:type
	|| qb:Observation). }
	FILTER NOT EXISTS
	{ GRAPH <v1> {
	?s ?p ?o.
	}
	}
	}
	``
		``
		SELECT ?s ?p ?o WHERE {
		GRAPH <v1> {
		{FILTER NOT EXISTS
		{?p rdfs:subPropertyOf rdfs:label}}
		UNION
		{FILTER (?p != rdfs:label).}
		?s ?p ?o.
		FILTER(?p != rdfs:label).
		FILTER(?p != rdfs:range).
		FILTER(?p != skos:inScheme).
		FILTER(?p != skos:broaderTransitive).
		FILTER(?p != qb:codeList).
		FILTER(?p != qb:component).
		FILTER(?p != qb:dimension).
		FILTER(?p != qb:measure).
		FILTER(?p != qb:attribute).
		FILTER(?p != qb:dataSet).
		FILTER(?p != qb:structure).
		FILTER(?p != rdf:type
		|| skos:Concept).
		FILTER (?p != rdf:type
		|| skos:ConceptScheme).
		FILTER (?p != rdf:type
		|| qb:AttributeProperty).
		FILTER (?p != rdf:type
		|| qb:CodedProperty).
		FILTER (?p != rdf:type
		|| qb:DimensionProperty).
		FILTER (?p != rdf:type
		|| qb:DataStructureDefinition).
		FILTER (?p != rdf:type
		|| qb:HierarchicalCodeList).
		FILTER (?p != rdf:type
		|| qb:Observation). }
		FILTER NOT EXISTS
		{ GRAPH <v1> {
		?s ?p ?o.
		}
		}
		}
		``

δ^+	(s, p, o)	$\emptyset$
δ^-	$\emptyset$	(s, p, o)

ϕ_{old}	—	{FILTER NOT EXISTS {?p rdfs:subPropertyOf rdfs:label}} UNION {FILTER (?p != rdfs:label).} ?s ?p ?o. FILTER(?p != rdfs:label). FILTER(?p != rdfs:range). FILTER(?p != skos:inScheme). FILTER(?p != skos:broaderTransitive). FILTER(?p != qb:codeList). FILTER(?p != qb:component). FILTER(?p != qb:dimension). FILTER(?p != qb:measure). FILTER(?p != qb:attribute). FILTER(?p != qb:dataSet). FILTER(?p != qb:structure). FILTER(?p != rdf:type || skos:Concept). FILTER (?p != rdf:type || skos:ConceptScheme). FILTER (?p != rdf:type || qb:AttributeProperty). FILTER (?p != rdf:type || qb:CodedProperty). FILTER (?p != rdf:type || qb:DimensionProperty). FILTER (?p != rdf:type || qb:DataStructureDefinition). FILTER (?p != rdf:type || qb:HierarchicalCodeList). FILTER (?p != rdf:type || qb:Observation).
ϕ_{new}	{FILTER NOT EXISTS {?p rdfs:subPropertyOf rdfs:label}} UNION {FILTER (?p != rdfs:label).} ?s ?p ?o. FILTER(?p != rdfs:label). FILTER(?p != rdfs:range). FILTER(?p != skos:inScheme). FILTER(?p != skos:broaderTransitive). FILTER(?p != qb:codeList). FILTER(?p != qb:component). FILTER(?p != qb:dimension). FILTER(?p != qb:measure). FILTER(?p != qb:attribute). FILTER(?p != qb:dataSet). FILTER(?p != qb:structure). FILTER(?p != rdf:type || skos:Concept). FILTER (?p != rdf:type || skos:ConceptScheme). FILTER (?p != rdf:type || qb:AttributeProperty). FILTER (?p != rdf:type || qb:CodedProperty). FILTER (?p != rdf:type || qb:DimensionProperty). FILTER (?p != rdf:type || qb:DataStructureDefinition). FILTER (?p != rdf:type || qb:HierarchicalCodeList). FILTER (?p != rdf:type || qb:Observation).	—

Change	*Add_Generic_Datatype*(x, t)	*Delete_Generic_Datatype*(x, t)
Intuition	Add a data type to a given subject	Delete a data type from a subject
Parameters	x = The subject in which the datatype is added t = The added datatype	x = The subject in which the datatype is deleted t = The deleted datatype
SPARQL used for detection	```SELECT ?x ?t WHERE {
GRAPH <v2> {
?x rdfs:range ?t.
}
FILTER NOT EXISTS { GRAPH <v1> {
?x rdfs:range ?t.
}
}
}```	```SELECT ?x ?t WHERE {
GRAPH <v1> {
?x rdfs:range ?t.
}
FILTER NOT EXISTS { GRAPH <v2> {
?x rdfs:range ?t.
}
}
}```
δ^+	$(x, \mathsf{rdfs : range}, t)$	$\emptyset$
δ^-	$\emptyset$	$(x, \mathsf{rdfs : range}, t)$
ϕ_{old}	−	−
ϕ_{new}	−	−

Change	Add_Generic_Attribute(x, attr)	Delete_Generic_Attribute(x, attr)
Intuition	Add a generic attribute to a given subject	Delete a generic attribute from a subject
Parameters	x = The subject in which the attribute is added $attr$ = The added attribute	x = The subject in which the attribute is deleted $attr$ = The deleted attribute
| SPARQL used for detection | ```
SELECT ?x ?attr WHERE {
GRAPH <v2> {
FILTER NOT EXISTS {
{?attr rdf:type
qb:DimensionProperty.}
UNION
{?attr rdf:type
qb:MeasureProperty.}
UNION
{?attr rdf:type
qb:CodedProperty.}
}
?x qb:attribute ?attr.
}
FILTER NOT EXISTS { GRAPH <v1> {
?x qb:attribute ?attr.
}
}
}
``` | ```
SELECT ?x ?attr WHERE {
GRAPH <v1> {
FILTER NOT EXISTS {
{?attr rdf:type
qb:DimensionProperty.}
UNION
{?attr rdf:type
qb:MeasureProperty.}
UNION
{?attr rdf:type
qb:CodedProperty.}
}
?x qb:attribute ?attr.
}
FILTER NOT EXISTS { GRAPH <v2> {
?x qb:attribute ?attr.
}
}
}
``` |
| δ^+ | $(x, qb : attribute, attr)$ | $\emptyset$ |
| δ^- | $\emptyset$ | $(x, qb : attribute, attr)$ |
| $\phi_{old}$ | − | ```
FILTER NOT EXISTS {
{?attr rdf:type
qb:DimensionProperty.}
UNION
{?attr rdf:type
qb:MeasureProperty.}
UNION
{?attr rdf:type
qb:CodedProperty.}
}
?x qb:attribute ?attr. }
``` |
| $\phi_{new}$ | ```
FILTER NOT EXISTS {
{?attr rdf:type
qb:DimensionProperty.}
UNION
{?attr rdf:type
qb:MeasureProperty.}
UNION
{?attr rdf:type
qb:CodedProperty.}
}
?x qb:attribute ?attr. }
``` | − |

Change	*Add_Generic_Value_to_Observation*(o, p, v)	*Delete_Generic_Value_from_Observation*(o, p, v)
Intuition	Add a generic value to observation	Delete a generic value from an observation
Parameters	o = The observation p = The property related as predicate of the observation v = The added value	o = The observation p = The property related as predicate of the observation v = The deleted value
SPARQL used for detection	```SELECT ?o ?p ?v WHERE {\nGRAPH <v2> {\nFILTER NOT EXISTS {\n{?p rdf:type qb:DimensionProperty.}\nUNION\n{?p rdf:type qb:MeasureProperty.}\n}\n}\nGRAPH <v2> {\n?o qb:dataSet ?ds.\n?ds qb:structure ?ft.\n?ft qb:component ?cs.\n?cs qb:componentProperty ?p.\n?p rdfs:range ?v.\n}\nFILTER NOT EXISTS { GRAPH <v1> {\n?o qb:dataSet ?ds.\n?ds qb:structure ?ft.\n?ft qb:component ?cs.\n?cs qb:componentProperty ?p.\n?p rdfs:range ?v.\n}\n}\n}```	```SELECT ?o ?p ?v WHERE {\nGRAPH <v1> {\nFILTER NOT EXISTS {\n{?p rdf:type qb:DimensionProperty.}\nUNION\n{?p rdf:type qb:MeasureProperty.}\n}\n}\nGRAPH <v1> {\n?o qb:dataSet ?ds.\n?ds qb:structure ?ft.\n?ft qb:component ?cs.\n?cs qb:componentProperty ?p.\n?p rdfs:range ?v.\n}\nFILTER NOT EXISTS { GRAPH <v2> {\n?o qb:dataSet ?ds.\n?ds qb:structure ?ft.\n?ft qb:component ?cs.\n?cs qb:componentProperty ?p.\n?p rdfs:range ?v.\n}\n}\n}```
δ^+	$(o, qb : dataSet, ds)$, $(ds, qb : structure, ft)$, $(ft, qb : component, cs)$, $(cs, qb : componentProperty, p)$, $(p, \mathsf{rdfs} : \mathsf{range}, v)$	$\emptyset$
δ^-	$\emptyset$	$(o, qb : dataSet, ds)$, $(ds, qb : structure, ft)$, $(ft, qb : component, cs)$, $(cs, qb : componentProperty, p)$, $(p, \mathsf{rdfs} : \mathsf{range}, v)$
ϕ_{old}	−	```FILTER NOT EXISTS {\n{?p rdf:type qb:DimensionProperty.}\nUNION\n{?p rdf:type qb:MeasureProperty.}\n}```
ϕ_{new}	```FILTER NOT EXISTS {\n{?p rdf:type qb:DimensionProperty.}\nUNION\n{?p rdf:type qb:MeasureProperty.}\n}```	−