

Annotating Relationships between Multiple Mixed-media digital objects by extending Annotea

Ronald Schroeter¹, Jane Hunter¹

¹School of ITEE, The University Of Queensland
{ronalds, jane}@itee.uq.edu.au

Abstract. Annotea provides an annotation protocol to support collaborative Semantic Web-based annotation of digital resources accessible through the Web. It provides a model whereby a user may attach supplementary information to a resource or part of a resource in the form of: either a simple textual comment; a hyperlink to another web page; a local file; or a semantic tag extracted from a formal ontology and controlled vocabulary. Hence, annotations can be used to attach subjective notes, comments, rankings, queries or tags to enable semantic reasoning across web resources. More recently tabbed Browsers and specific annotation tools, allow users to view several resources (e.g., images, video, audio, text, HTML, PDF) simultaneously in order to carry out side-by-side comparisons. In such scenarios, users frequently want to be able to create and annotate a link or relationship between two or more objects or between segments within those objects. For example, a user might want to create a link between a scene in an original film and the corresponding scene in a remake and attach an annotation to that link. Based on past experiences gained from implementing Annotea within different communities in order to enable knowledge capture, this paper describes and compares alternative ways in which the Annotea Schema may be extended for the purpose of annotating links between multiple resources (or segments of resources). It concludes by identifying and recommending an optimum approach which will enhance the power, flexibility and applicability of Annotea in many domains.

Keywords. Annotea, Annotation, Semantic Web, Relationships

1 Introduction

Simple Web annotation tools for annotating individual web objects have existed for over ten years [1, 2].

They began with annotation tools for attaching comments to web pages and textual documents, but then expanded to images and video, audio and 3D objects as more multimedia content was published on the Web. More recently, as many communities have formed online collaborative groups, annotation tools have transformed from asynchronous to synchronous - enabling real-time online discussions about resources. Figure 1 below illustrates the evolution of annotation tools over the past ten or so years.

In the past year, we have observed yet a new phase in the demands of users with respect to annotation tools. Our observation is related to the establishment of more online communities, who have established a consensus on exchangeable data standards, terminologies (defined through mark-up languages and machine-processable ontologies) and who want to be able to share and compare overlapping and related resources of many types. These resources may be of many media types (images, video, audio, multimedia, 3D), associated with specific disciplines (e.g., scientific models) or may comprise XML files used to represent shareable, exchangeable objects (e.g., scientific workflows).

To summarize, communities have been voicing a demand for annotation tools that enable a combination of the following:

1. The specification of links between whole objects or segments within objects and annotation of these links [3].
2. Support for annotating links between objects of the following types: images, video, audio, text, HTML, PDF, 3D objects and XML files;
3. Viewing of more than one object simultaneously to enable side-by-side comparison and association;
4. Annotations that are based on domain-specific terms from either controlled vocabularies or (OWL) ontologies. This enhances the ability for other application programs to process the annotations;
5. The ability to share these comparative interpretations and associations amongst communities of users through shared annotation servers, using a protocol such as Annotea.

Some specific examples that we have seen through our eResearch collaborations with different scientific communities include:

- In the humanities, film/media researchers want to link, compare and annotate segments between books, screenplays and different films and film versions;
- In molecular biology, researchers want to be able to relate and compare 3D protein structures – to discuss protein-protein docking interactions and protein function.
- In the geosciences, geologists want to be able to compare and annotate different types of computational models with still photos and videos of earth quakes.

The remainder of this paper is structured as follows. Section 2 provides an overview of previous related work and a description of the Vannotea tool developed by the authors, which has been the driver for the work described in this paper. Section 3 describes the existing Annotea protocol and the advantages of extending this to support new user demands. Section 4 describes extensions to Annotea to support machine-processable annotations (based on ontologies). Section 5 describes different possible approaches to extending Annotea to support the annotation of links between multiple objects. Section 6 concludes with a recommendation for the optimum approach and future work.

2 Previous Work

Significant previous work has focussed on the development of annotation tools. Figure 1 provides a 3D classification of existing annotation systems – classified according to:

- annotation level (x-axis) – from simple free-text annotations and tagging, to the attachment of local files and URLs, and more controlled annotations based on simple vocabularies and ontologies.
- Content type (y-axis) – text, HTML, images, video, audio and 3D objects
- Number of simultaneous resources (z-axis) - the ability to compare multiple files and annotate links between them.

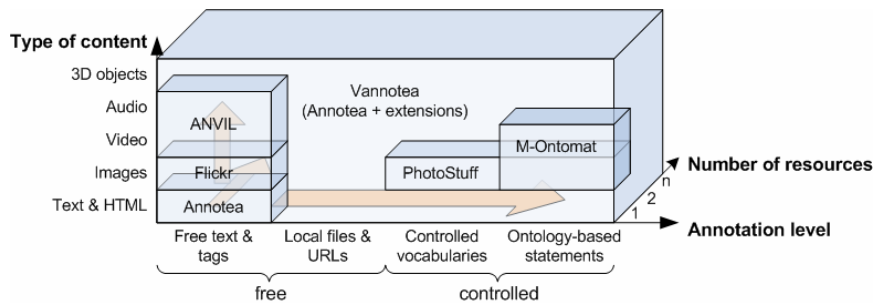


Figure 1. Annotation tools

2.1 Free Annotation Tools

Examples of simple free-text annotation tools are depicted in the left column of Figure 1. They include tools that are based on Annotea without extensions such as Annozilla¹ and Amaya². Flickr³ is an online management photo management and sharing application. It allows users to upload their photos and freely annotate them. ANVIL [4] is a stand-alone tool which allows free-text annotations of audio and video files. Many more tools could be mentioned here but the focus of this paper is on annotation systems that support ontology-based annotation of links between multiple web-accessible digital resources.

2.2 Semantic Annotation Tools

Systems that support controlled vocabulary-based and ontology-based annotations of multimedia objects include the following:

¹ <http://annozilla.mozdev.org>

² <http://www.w3.org/Amaya>

³ <http://www.flickr.com>

PhotoStuff [5] is a tool that allows users to highlight regions within images, create instances from any ontology through sophisticated forms and link the instance to the region of the image. The users are able to perform the semantic annotation locally and then upload the RDF instance to a central database, where the RDF file - and therefore the whole graph including multiple instance statements - is then attributed to the user through his/her user account and time stamped for provenance data.

The **M-Ontomat-Annotizer** [6] provides ontology based image and video frame (and regions) annotation. This tool also supports initialization and linking of RDF(S) domain ontologies with low-level MPEG-7 visual descriptors.

Vannotea [3] is a collaborative tool that enables fine-grained annotation of objects of any media type, where the annotations themselves can be free-text, files or URLs or from a controlled vocabulary (e.g., WordNet) or ontology. As a result of user demand, Vannotea was recently extended to enable the viewing of multiple related objects simultaneously. Users in geographically distributed locations can share the Vannotea application and simultaneously view two or more videos or 3D objects through a user interface that allows side-by-side comparisons (Figure 2).

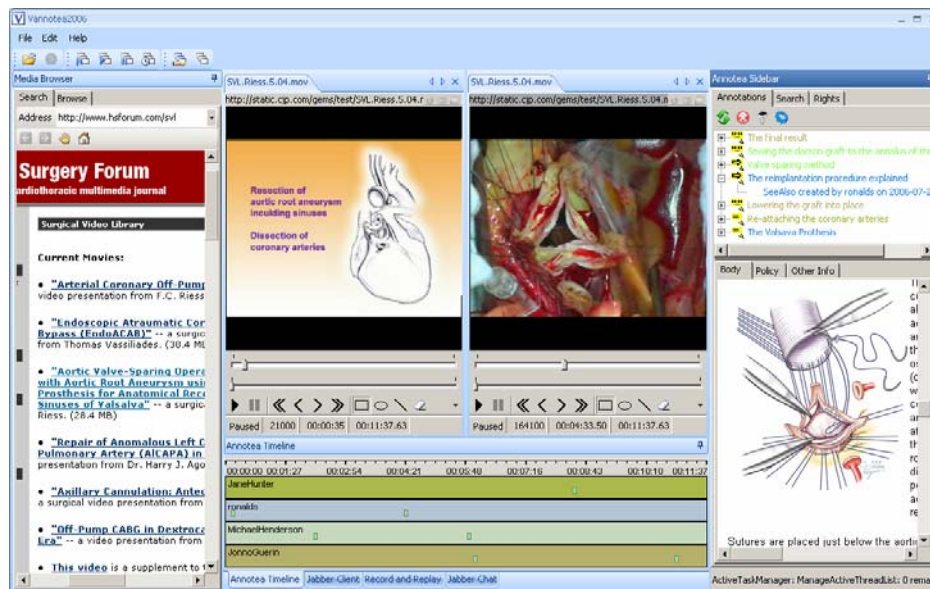


Figure 2. Screenshot of Vannotea

Hence the aim of this paper is to describe in detail the model we have chosen for storing the different types of annotations so that they can be attributed to individual users for provenance data. The model is based on Annotea [7] (described in the following Section 3) and extending it in the following directions:

- to allow controlled annotations (Section 4) - see arrow along the x-axis in Figure 1; and
- the annotations of links between parts of multiple digital object of any type (Section 5) – see arrows along the y- and z-axis in Figure 1.

3 Annotea

Through earlier work [8], we identified Annotea [7] as an ideal approach for implementing an annotation server. Annotea, in its original sense, is a Web-based annotation system that uses Resource Description Framework (RDF) to model free annotations as a set of statements or assertions made by the author about a particular webpage. These annotations are then stored in a HTTP enabled server, which enables clients to query, update, post, delete and reply to annotations.

A key strength of the Annotea protocol is that it uses open W3C standards such as RDF, XPointer, XLink and HTTP. The use of machine-processable RDF descriptions enables easy search, retrieval and linking of the annotations to related resources and services using semantic web technologies (e.g., OWL, SPARQL).

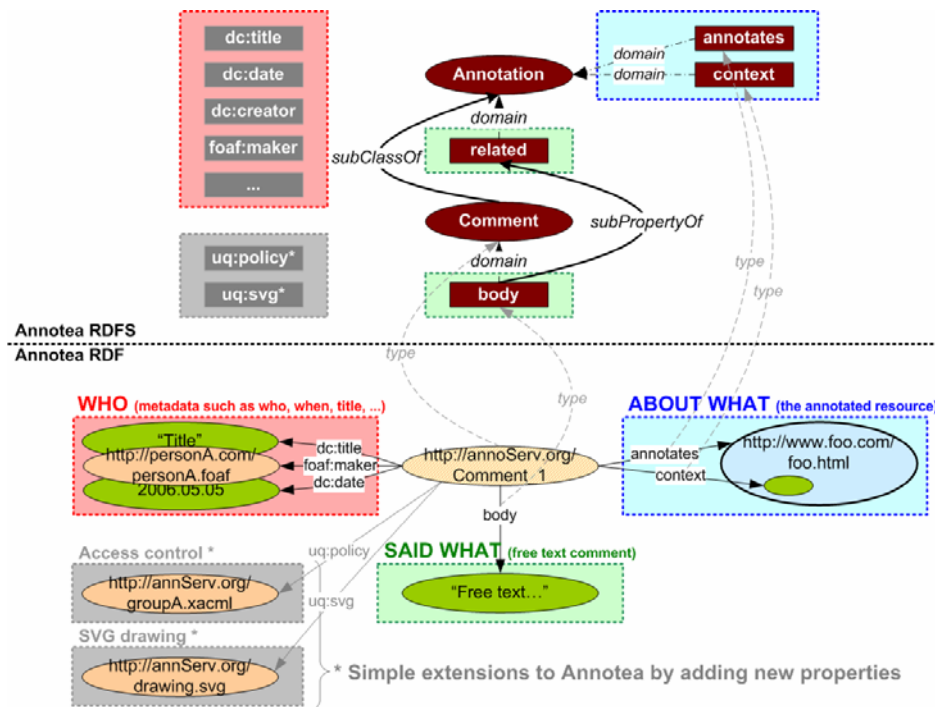


Figure 3. Annotea Schema and instance with access control and SVG extension

Figure 3 illustrates the RDF Schema of Annotea and an RDF instance – an Annotea object – separated by the dividing dotted line. The Annotea Schema introduces properties that point to the annotated Web document (`annotates`) and to a specific location within a structured Web document, thus describing the `context` of an annotation, for which Annotea uses the XPointer technology. Furthermore the specification provides a `related` property, which relates the resource representing the 'content' of an Annotation to the annotated resource.

Developers are encouraged to create new types of annotations by sub-classing from the `Annotation` class and creating sub-properties of the `related` property. Figure 3 shows such a new type, the `Comment` class and the `body` property. In essence, the RDF instance of a simple Annotea object will carry information about “*Who said what about which resource?*”

RDF/S allows easy addition of metadata properties from other schemas such as the Dublin Core (`title`, `date`, `creator`) and FOAF (Friend-of-a-Friend) namespaces, which are used to describe the provenance of an annotation.

We have also made our own, application specific extensions. For example, within Vannotea [3] we have extended it to allow annotations in the form of drawings on top of media types such as images or videos through the use of SVG.

Furthermore, we have added the functionality to enable users to apply fine-grained access control to their Annotations through XACML policies and implementing the Annotation Server as a Shibboleth Service Provider [9]. A survey of current Web-based annotation systems [10] reveals that they vary in the way in which annotations may be attached, the way in which they are presented and in the access control mechanisms. Some systems are designed for private use only whilst others permit sharing amongst groups and/or public access. None of the surveyed systems provide the kinds of fine-grained access control mechanisms that are achieved by our implementation and are required by collaborative teams of scientists engaging in eResearch.

Koivunen [11] introduced new `Bookmark` and `Topic` objects to Annotea. These social bookmarks and topics can be used for semantic authoring by letting ordinary users tag interesting web documents with their *own personal* concepts or folksonomies.

Currently, the `Bookmark` class is a separate, new class within Annotea. In our view, however, bookmarks are just a special type of annotations. Rather than a user attaching a free-text comment to a specific resource, the user can build their own folksonomy using the topic hierarchies and attach those topics to the resource. Semantically we are basically describing “*who used which topic about which resource?*”, rather than “*Who said what about which resource?*” as mentioned in the previous section.

Therefore we suggest to subclass the `Bookmark` class from the `Annotation` class as shown in Figure 4, which includes making the `recall` property a sub-property of the `annotates` property, and `hasTopic` a sub-property of `related`.

As a result, we will be able to query the Annotation Server to return any Annotation that is attached to a specific resource, whether it is of type `Comment`, `Bookmark`, or any other types that will follow in this paper.

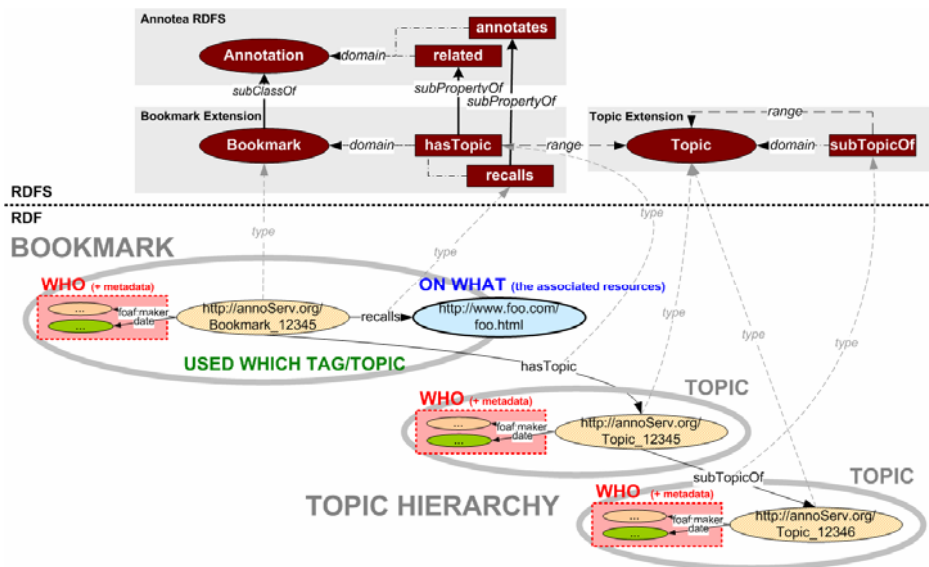


Figure 4. Bookmark and Topic Hierarchy

4 Ontology-based Annotations using Annotea

As mentioned earlier, various communities, especially within the field of eScience, are creating their domain-specific ontologies through group consensus. These ontologies can be hierarchical controlled vocabularies modelled in RDFS or more complex knowledge representations in OWL. This section illustrates how Annotea can be extended to allow users to take advantage of these formal concepts in order to create subjective semantic annotations. These formal annotations can aid in bridging the semantic gap between automatic recognition techniques that extract different low-level visual or audio features and highly subjective free-text annotations by humans. The line between objectivity and subjectivity is not always clear. When does a subjective annotation become objective and in whose eyes? As the provenance data for these controlled annotations is recorded, machines will be able to evaluate the objectivity of the individuals' semantic statements, based on trust relationships between the users and statistical calculations.

4.1 Controlled Vocabularies

One extension to Annotea using controlled vocabularies is to allow users to attach pre-defined ranking information to a specific web resource (e.g., the controlled terms "strong accept", "accept" and "reject" for a collaborative review process of scientific

papers). Another example is using terms from an ontology such as the WordNet Ontology⁴.

In any case, the controlled vocabulary or ontology is modelled in RDFS or OWL and publicly available over the web, so that an Annotea client can access and present it to the user when he/she wants to attach a controlled term to the resource. In Figure 5, the user searched for the term animal and then browses through the WordNet Ontology to navigate to a controlled vocabulary of a specific animal, which can then be attached to the resource (or part of the resource) that the user is currently viewing in his browser or Vannotea client.

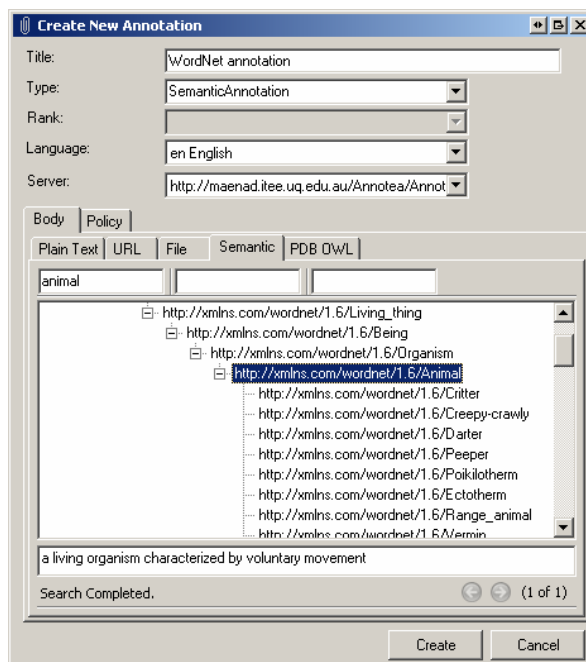


Figure 5. Screenshot of creating a link to a controlled term

This is very similar to creating bookmarks (see Figure 4), except that the topic is being replaced by a predefined controlled vocabulary. The benefit of using these controlled vocabularies is that we can perform searches using these terms, taking advantage of the ontology to infer that a “fish”, for example, is a subclass of an “animal”, and therefore returning all resources about a “fish” when querying for resources about “animals”. Since we store provenance data about who created the annotation, we envisage to take definitions of trust relationships inside a user’s FOAF profile into consideration when querying the Annotation Server, e.g., “retrieve the ranking information about a particular resource from all users that I trust and calculate an average rating”.

⁴ <http://xmlns.com/2001/08/wordnet>

4.2 Simple Formal Statements

Using the same interface depicted in Figure 5 users can also attach formal triple statements based on ontologies to a resource, i.e., relate a formal statement to the annotated resource and context. A statement is a more complex instance of an ontology compared to the controlled terms in the previous section. Figure 6 illustrates the schema of the `FormalStatement`. We introduce a `states` property which is a sub-property of `related` and has a range of `rdf:Statement`. The statement itself consists of a subject, predicate and object from an Ontology.

The example in Figure 6 shows a statement that says “*lion eats gazelle*” from a simple Wildlife Ontology which defines a lion being a subclass of a carnivore and a gazelle a subclass of a herbivore. As above, we can now perform ontology based searches to retrieve all video segments or images where a “*carnivore eats a herbivore*”, which would include scenes that were formally labelled as a “*lion eating a gazelle*”.

However, Figure 6 also demonstrates the problem when using reification with `rdf:Statement`: the amount of triples explode [12] as every statement carries the same metadata (as indicated by the greyed out statements). Therefore we are currently investigating the use of named graphs as a mechanism for reasoning about provenance [13].

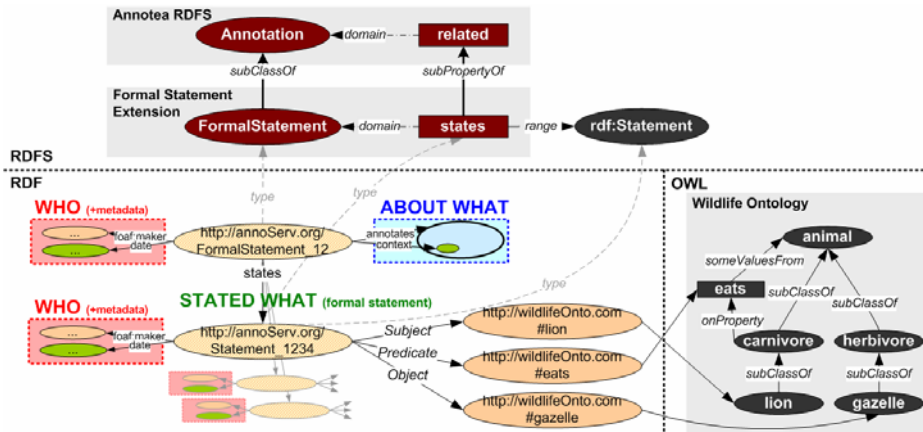


Figure 6. Formal Statement

5 Comparisons and Associations using Annotea

As mentioned earlier, tabbed Browsers and tools like Vannotea, allow users to view several objects (images, video, audio, text, HTML, PDF) simultaneously and carry out side-by-side comparisons. In such scenarios, users want to be able to annotate the link between two or more objects or between segments of multiple objects. For example, a user might want to annotate the link

- between a scene in an original film and the corresponding scene in a remake;
- between an image and a location (through Google maps URLs);
- between regions within several images;
- between structural components of two different 3D protein structures.

This section will investigate several approaches to model the annotation of such comparisons and associations within Annotea in a way that follows the best practices described earlier. A *comparison* annotates multiple resources and describes their similarities or dissimilarities whereas an *association* describes a user's mental connection between the resources. Although comparisons and associations are semantically different, they are both conceptually similar in the fact that they are about multiple resources, where the order of the resources is irrelevant and the description applies to the collection of resources.

As Annotea is based on RDF, it is very flexible with regards to adding/extending it to other properties as we have demonstrated earlier. Furthermore, since there is no cardinality defined for properties of an Annotea object, we can add multiple properties of the same type to the object. Therefore the most convenient way to relate an Association (or Comparison) to multiple resources would be to use multiple annotates properties.

However, if we want to create an Association object between two parts (contexts) of two resources, we run into the following problem: According to the Annotea Protocol [7], the context property is supposed to include an XPointer, e.g.:

```
<context>
  http://mydomain.com/foo.html#xpointer(id("Main")/p[2])
</context>
```

As we have identified in our earlier work with Vannotea, XPointer does not suit time continuous media. Instead, temporal fragment identifiers, which have been discussed for URIs⁵, could be used to refer to a time segment as follows:

```
<context>http://mydomain.com/foo.mpg#?t=15.2-18.7</context>
```

The two examples above show that the context includes information about the resource it refers to. Unfortunately, we cannot always assume for this to be the case. The context for images might be a definition of regions in SVG or some other format:

```
<context><svg id="SvgGdi_output"><g id="root_group">
  <rect height="102" id="77" width="79" x="95" y="125" />
</g></svg></context>
```

The context within 3D models might be an application specific string describing the zoom factor, position, angle, selected polygons or even selected atoms and molecules within JMOL⁶ models, e.g. using a JMOL script string:

```
<context>
  SAH 2001.OXT number:486,moveto 1 58 -106 993 83.7 158
</context>
```

⁵ <http://www.annotex.net/TR/draft-pfeiffer-temporal-fragments-02.html>

⁶ <http://jmol.sourceforge.net/>

This leads to the problem illustrated in Figure 7, where the literals denoting the context information have no formal connection to the resource they refer to.

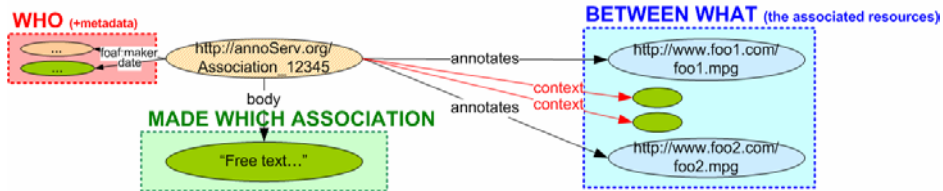


Figure 7. The context property

The following sections will illustrate several attempts to bypass this problem, each attempt with its own advantages and disadvantages, before providing a recommended solution.

5.1 Attempt 1: The isLinkedTo property

Attempt 1 is illustrated in Figure 8, in which every resource that is part of an association is annotated as a separate Association object, and all Association objects are then linked together by a new isLinkedTo property, which has a domain Association and a range Association.

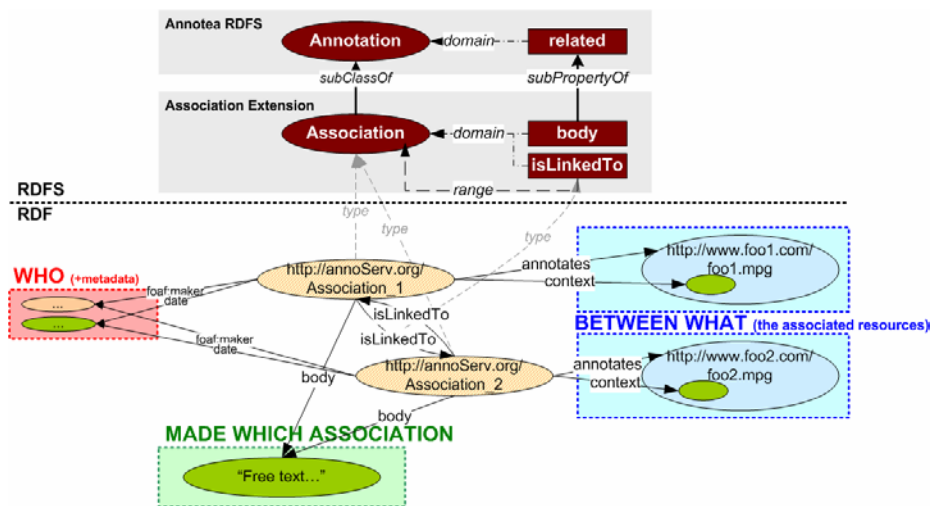


Figure 8. The isLinkedTo property

The advantage of this approach is that it is easy to implement as the additions don't require any changes to the current Annotation Server implementation. A general query such as "Give me all annotations that annotate this resource" (?Annotation

annotates "http://www.foo1.com/foo1.mpg") will retrieve AssociationItem_1 and its isLinkedTo property is automatically returned as part of it. However, to retrieve the resource AssociationItem_2 annotates, we would have to perform a nested query.

Another disadvantage is that the association as a whole cannot be addressed. In the example above it would have several addresses. This has several implications, not being able to delete or reply to the association are just a few. Furthermore, managing an update (modification) of the association becomes very cumbersome to implement as there would be many linked statements to fix. This might not be apparent in the above example, but an association might involve more than two resources, in which case deleting one AssociationItem would involve cleaning up n-1 links. Finally, the retrieval of all the links requires a recursive query, which is not supported by SPARQL.

5.2 Attempt 2: The AnnotationGroup object

Figure 9 illustrates Attempt 2, in which a newly introduced Group class (an rdf:Bag) links to all the AssociationItems. This means that the association as a whole can now be addressed through the AssociationGroup object. However, since the Group class is not a subclass of the Annotation class, it is not very useful. The query "Give me all annotations that annotate this resource" still returns an AssociationItem and a fairly complex and expensive nested query will need to retrieve the other AssociationItems through the AssociationGroup object.

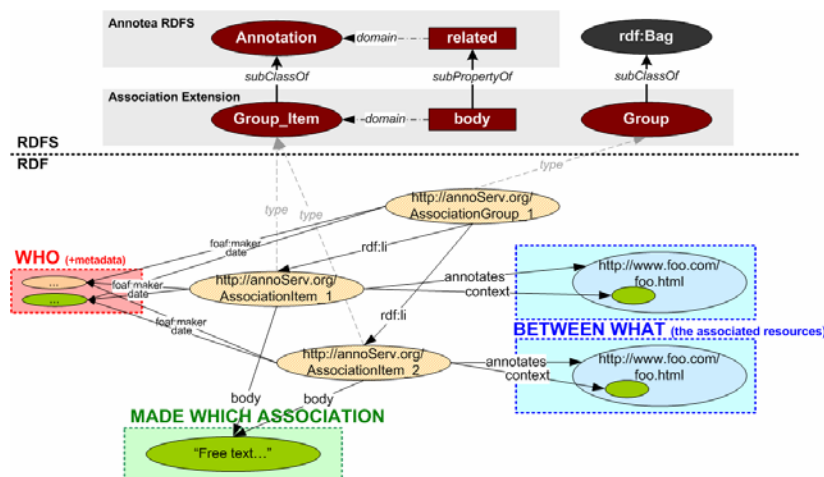


Figure 9. Annotation Group

Although this is likely to be supported by SPARQL using inferencing⁷, it should be avoided if possible.

⁷ Question 3.3. of the SPARQL FAQ (<http://thefigtrees.net/lee/sw/sparql-faq>)

5.3 Attempt 3: The Target object

Figure 10 shows the third attempt, which unlike the previous two, views an Association as a subclass of an Annotation, and tries to combine the context with the resource it refers to by introducing a new Target object. The advantage is that the association as a whole is addressable and there are also far less triples to manage.

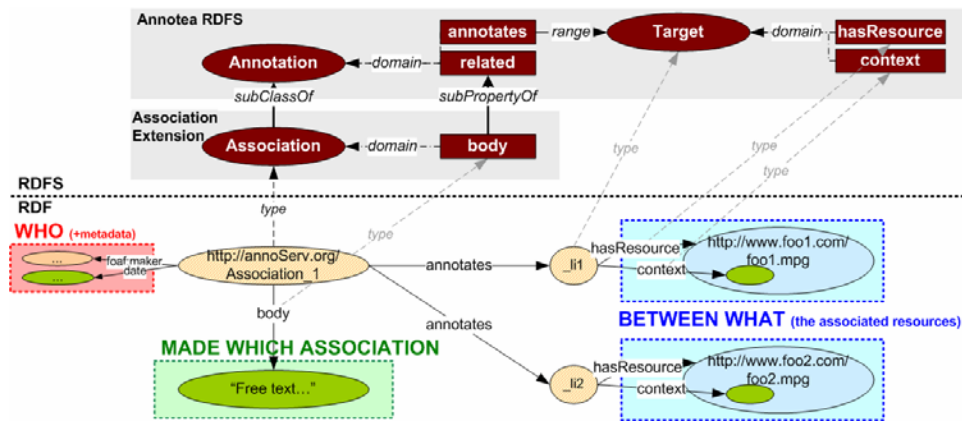


Figure 10. Target Class

On the other hand, it modifies the Annotea Schema, which renders this attempt backwards incompatible with annotations based on the original schema. The query "Give me all annotations that annotate this resource" no longer works, instead we have to extend it to follow the graph through the Target object (a blank node), e.g.:

```
?Annotation annotates _targetObject,
_targetObject hasResource "http://www.foo1.com/foo1.mpg"
```

5.4 Recommended solution: The Context object

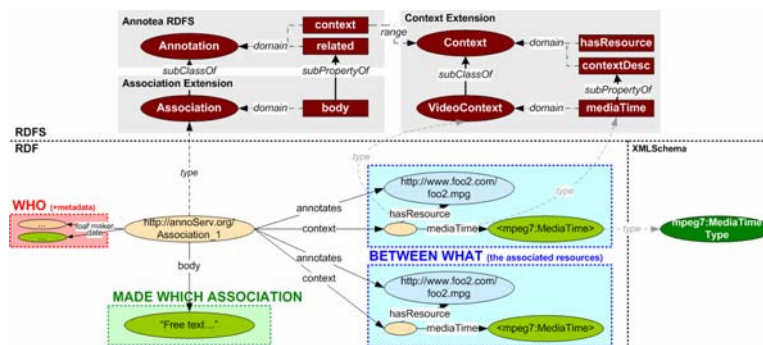


Figure 11. Context Class

Finally, the recommended solution is illustrated in Figure 11. The context property has a range Context class, which is the domain of two new properties, hasResource and contextDescription. Additionally, content type specific Context classes and contextDescription properties can be subclassed, e.g., VideoContext and mediaTime for resources that are videos.

The RDF instance in Figure 11 shows how the context property is pointing to a Context object (a blank node) which links to the same resource (hasResource) as the annotates property of the Association object. Additionally, it can contain any formalized or standardized description to represent the context, e.g., using the XML-Schema datatype MediaTimeType from the Multimedia Description Scheme⁸ (MDS) of the MPEG-7 standard [14]:

```
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:a="http://www.w3.org/2000/10/annotation-ns#"
  xmlns:dc="http://purl.org/dc/elements/1.0/"
  xmlns:foaf="http://xmlns.com/foaf/0.1/"
  <rdf:Description rdf:about="http://mydomain/Anno/10894">
    <rdf:type rdf:resource="http://www.w3.org/2000/10/annotation-
ns#Association" />
    <dc:creator>ronalds</dc:creator>
    <foaf:maker rdf:resource="http://www.my/~ronalds#ron" />
    <dc:date>2006-11-09T14:28:27Z</dc:date>
    <a:body rdf:resource="http://mydomain/Anno/body/10894" />
    <a:annotates>http://foo1.org/foo1.mpg</a:annotates>
    <a:annotates>http://foo2.org/foo2.mpg</a:annotates>
    <a:context>
      <a:hasResource>http://foo1.org/foo1.mpg</a:hasResource>
      <a:mediaTime rdf:parseType="XmlLiteral"
xmlns:mpeg7="urn:mpeg:mpeg7:schema:2001#">
        <mpeg7:MediaTime xsd:type="mpeg7:MediaTimeType">
          <mpeg7:MediaRelTimePoint mediaTimeUnit="PT1S">
10</Mpeg7:MediaRelTimePoint>
          <mpeg7:MediaIncrDuration mediaTimeUnit="PT1S">
5</Mpeg7:MediaIncrDuration>
          </mpeg7:MediaTime>
        </a:mediaTime>
      </a:context>
    <a:context>
      <a:hasResource>http://foo2.org/foo2.mpg</a:hasResource>
      <a:mediaTime rdf:parseType="XmlLiteral"
xmlns:mpeg7="urn:mpeg:mpeg7:schema:2001#">
        <mpeg7:MediaTime xsd:type="mpeg7:MediaTimeType">
          <mpeg7:MediaRelTimePoint mediaTimeUnit="PT1S">
22</Mpeg7:MediaRelTimePoint>
          <mpeg7:MediaIncrDuration mediaTimeUnit="PT1S">
5</Mpeg7:MediaIncrDuration>
          </mpeg7:MediaTime>
        </a:mediaTime>
      </a:context>
    </rdf:Description>
</rdf:RDF>
```

⁸ <http://m7itb.nist.gov/mds-2001.xsd>

This approach is backward compatible in the sense that the general query “*Give me all annotations that annotate this resource*” will return any `Association` object that has one link to the resource. The modified range of the `context` property doesn’t have to be defined within the Annotea Schema, but could be defined as a new sub-property of the `context` property within the Context extension.

6 Conclusion

In this paper we have demonstrated various ways of extending the Annotea Schema to enable annotation of links between segments of multiple objects. We have shown that careful considerations need to be made as the flexibility of RDF and the ease to add and extend new classes and properties might have wide reaching implications in the big picture if not being thought through cautiously.

When extending Annotea, we recommend the following best practices:

- Reuse existing Annotea classes and properties as well as RDF and XML Schema (built-in) types where possible;
- A general query such as “Give me all Annotations that annotate this resource” should always return ALL objects that are sub-classes of `Annotation`, e.g. `Comment`, `Bookmark`, `Ranking`, `FormalStatement`, `Association`. It is then up to the client to display the different objects accordingly;
- A general query like above should return all the information needed to enable clients to display an appropriate overview/list, i.e., avoiding nested queries on the server-side to retrieve additional information where possible;
- Avoid unnecessary explosion of triples in the triple-store;
- Investigate SPARQL implications.

7 Future work

In the future, we will investigate the following:

- The use named graphs instead of triple hungry reification where possible to avoid triple explosion in the RDF store.
- Consider definitions of trust relationships inside a user’s FOAF profile when querying the Annotation Server, e.g. filter out annotations that are not trusted.
- Combine Comparisons/Associations with ontologies, which should be a straight forward task to do based on the work presented in this paper.
- Add generic HTML based ontology browsers and forms to create more complex ontology instances and embed it into the user interface of our Annotea Sidebar and Vannotea.

References

1. Roscheisen, M., Mogensen, C., and Winograd, T.: Shared Web Annotations as a Platform for Third-Party Value-Added, Information Providers: Architecture, Protocols, and Usage Example. Computer Science Dept., Stanford University. Technical Report CSDTR/DLTR STAN-CS-TR-97-1582 (1994).
2. Davis, J. and Huttenlocher, D.: The CoNote System for Shared Annotations. (1995) <http://www.cs.cornell.edu/home/dph/annotation/annotations.html>.
3. Schroeter, R., Hunter, J., Guerin, J., Khan, I., and Henderson, M.: A Synchronous Multimedia Annotation System for Secure Collaboratories. In 2nd IEEE International Conference on E-Science and Grid Computing (eScience 2006). Amsterdam, Netherlands (2006)
4. Kipp, M.: Anvil - A Generic Annotation Tool for Multimodal Dialogue. In 7th European Conference on Speech Communication and Technology (Eurospeech). Aalborg (2001) 1367-1370
5. Halaschek-Wiener, C., Golbeck, J., Schain, A., Grove, M., Parsia, B., and Hendler, J.: Annotation and provenance tracking in semantic web photo libraries In International provenance and annotation workshop. Chicago (2006)
6. Petridis, K., Kuehn, K., Handschuh, S., Bloehdorn, S., Saathoff, C., Avrithis, Y., Kompatsiaris, Y., and Staab, S.: Semantic Annotation of Images and Videos for Multimedia Analysis and Retrieval. In Lecture Notes in Computer Science, vol. 3532: Lecture Notes in Computer Science (2005) 592-607.
7. Koivunen, M.-R. and Kahan, J.: Annotea: an open RDF infrastructure for shared Web annotations. In Proceedings of the 10th international conference on World Wide Web. Hong Kong. ACM Press (2001)
8. Schroeter, R., Hunter, J., and Kosovic, D.: Vannotea - A Collaborative Video Indexing , Annotation and Discussion System For Broadband Networks. In Knowledge Markup and Semantic Annotation Workshop, K-CAP 2003. Sanibel, Florida (2003)
9. Khan, I., Schroeter, R., and Hunter, J.: Implementing a Secure Annotation Service. In International Provenance and Annotation Workshop (IPAW2006). Chicago (2006)
10. Uren, V., Cimiano, P., Iria, J., Handschuh, S., Vargas-Vera, M., Motta, E., and Ciravegna, F.: Semantic annotation for knowledge management: Requirements and a survey of the state of the art. Web Semantics: Science, Services and Agents on the World Wide Web, vol. 4 (2006) 14-28
11. Koivunen, M.-R.: Semantic Authoring by Tagging with Annotea Social Bookmarks and Topics In The 5th International Semantic Web Conference (ISWC2006) - 1st Semantic Authoring and Annotation Workshop (SAAW2006). Athens, GA, USA (2006)
12. Carroll, J. J. and Stickler, P.: RDF triples in XML. In 13th international World Wide Web conference. New York, NY, USA (2004) 412 - 413
13. Watkins, E. R. and Nicole, D. A.: Named Graphs as a Mechanism for Reasoning about Provenance. In APWeb 2006: 8th Asia-Pacific Web Conference. Harbin, China (2006) 943-948
14. Salembier, P. and Smith, J. R.: MPEG-7 multimedia description schemes. IEEE Transactions on Circuits and Systems for Video Technology, vol. 11 (2001) 748-759