

# Integration of Semantic Web Technology in an Annotation-based Hypervideo System

Olivier Aubert, Pierre-Antoine Champin, and Yannick Prié

LIRIS, University Claude Bernard Lyon 1  
F-69622 Villeurbanne Cedex, France  
firstname.lastname@liris.cnrs.fr  
<http://liris.cnrs.fr/firstname.lastname>

**Abstract.** This article discusses the integration of semantic web technologies (ontology and inference) into audiovisual annotation based models and systems. The Advene project, aimed at all purpose hypervideo generation from annotated audiovisual documents, is used as a testbed. Advene principles and the Advene prototype are first presented, before a discussion on how ontology and reasoning have been easily integrated into the Advene framework. Some motivating examples are proposed, and our proposals and related works are discussed.

## 1 Introduction

This article has two primary goals and one angle of vision. The first goal is to present the Advene (Annotate Digital Video, Exchange on the NEt) project, rationale and prototype, as a powerful set of ideas and tools for designing and testing innovative uses of video through rich video annotation and hypervideo generation. The second goal is to present how semantic web technologies can be used in the Advene framework, in order to provide extra features mainly related to ontological reasoning. The angle of vision we adopt throughout the article is more related to the audiovisual annotation and hypermedia document engineering than to the semantic web. Our main concern is to study what semantic web technologies can bring to audiovisual annotation, hypervideo construction and audiovisual information systems.

We generally consider [1] that full audiovisual information systems (AVIS) should provide users with the possibility to search into video bases using indexes, to select from the retrieved documents the most relevant for the current task, and to reuse parts of documents, mainly in new hypermedia documents (simply watching them being a very simple reuse for instance). We also claim that annotations are the pivot to such audiovisual information systems, as they provide all the necessary information to search, select and manipulate audiovisual documents and fragments.

As a means to understand present annotation-based audiovisual information systems and to design future ones, we choose to use the hypervideo concept. *Hypervideos* [2] refer to hypermedia documents that are constructed from *annotated audiovisual documents* information: digital audiovisual documents (moving images and related sounds) and annotations for these audiovisual documents, which are digital pieces of information in relation with spatio-temporal fragments of the documents. Examples of hypervideos include web pages that refer to a video using some excerpts as key images,

video streams enriched with textual information and hyperlinks, reconstructed audiovisual streams, etc.

The next section of this article presents the Advene model and prototype for hypervideo engineering. The following section describes how we easily added semantic web reasoning capabilities to the Advene prototype thanks to its flexible model. The last section deals with some examples of semantic web-enriched uses of annotations in the context of Advene and more generally in the AVIS/hypervideo context.

## 2 Advene model and prototype

The Advene project aims at providing tools to exchange various analyses about movies stored in digital form (digital video files, DVDs...), and more importantly, offer the possibility to enhance and customize these analyses. Analyses are built upon annotations, which represent pieces of data of any type that are (spatio)temporally linked to the movie. The Advene prototype thus provides means to create and modify annotations, as well as to specify how they should be rendered in meaningful ways. Instead of exchanging the sole final form of an analysis, the Advene project makes it possible to rather exchange annotations and the specification of their visualisation, thus allowing end-users to customize data and visualisations in order to fit their needs.

We will see in this section how data is organized by the Advene model, and how the model is implemented in the current prototype<sup>1</sup>.

### 2.1 The Advene model

It is commonly agreed that the handling of audiovisual contents has to use metadata, the audiovisual data itself being not fitted to indexing or querying without any pre-processing. Of the various existing approaches, let us give an overview of two important standards – MPEG-7[3] and Annodex[4] – and see how our proposal relates to them.

MPEG-7 aims at being the standard representation format to exchange metadata associated to audiovisual streams. It provides means to link metadata to portions of audiovisual documents. The MPEG-7 standard defines standard metadata, mostly focusing on low-level descriptors automatically extractable from the audiovisual document (colors, textures, shapes, audio characteristics...), as well as a way to specify additional metadata through XML Schema. It is used by some vendors, but a common complaint is the complexity of its model, which makes it difficult to use for simple things or for interoperation with other standards [5].

The Annodex[4] projects aims at creating a *continuous media web* where metadata is embedded in audiovisual documents, making them indexable and searchable. Aiming at simplicity, its *Continuous Media Markup Language* (CMML) is inspired by HTML, and allows to quickly edit metadata. After edition, CMML files are combined with the audiovisual documents. Annodex solves the simplicity issue (using an HTML inspired syntax), at the expense of a lack of structure. Moreover, it merges metadata with the audiovisual document, making it harder to use different metadata for the same audiovisual document.

---

<sup>1</sup> available from <http://liris.cnrs.fr/advene/>

The Advene model somehow aims at bridging the gap between both approaches: it provides a way to link metadata to audiovisual documents. It does not impose any constraints on the nature of metadata, and keeps metadata separate from the audiovisual document, so that they can evolve and be exchanged independently from each other.

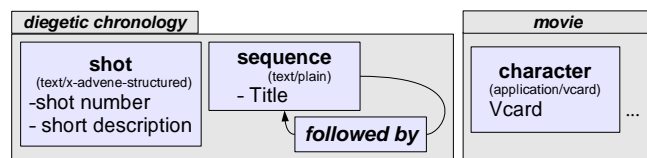
**Annotation structure** We developed the Advene model based on our reflexion about hypervideos [1]. An *Annotated Audiovisual Document* (AAD) is an audiovisual document augmented with metadata. Processing both the audiovisual document and its accompanying metadata gives *views* on the AAD, some of them qualifying as *hypervideos*: views of the AAD that on the one hand use information from both the audiovisual document *and* the annotation structure, and on the other hand give access to the temporality of the audiovisual document.

In the Advene model, described more precisely in [2], the annotation structure consists mainly of *annotations*, that contain data and are linked through a temporal (possibly spatio-temporal) fragment to a specific portion of the audiovisual document. The structure of data contained in the annotations is not specified by the model: it can be any type of data (simple text, structured information, audio documents, office documents...).

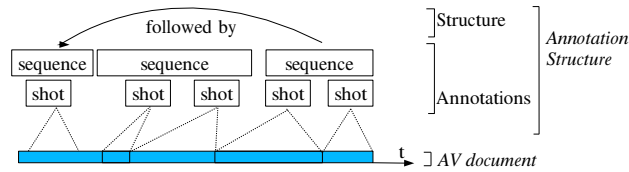
In order to be usable, while retaining their genericity, annotations are flexibly structured: *annotation types* define the kind of content (through a MIME-type specification) held by annotations. Multiple annotation types can be used to describe a number of analysis facets. Moreover, *relations* allow to link annotations with each other, and are specified by *relation types*. Relation types define the types of annotations that can be linked, as well as an optional content MIME type for relations.

As annotation types and relation types define a certain point of view in the document analysis, they are grouped as meaningful sets called *schemas*. An Advene schema thus defines annotation types and relation types that form together an analysis framework.

Let us illustrate this structure through a simple example. Consider a movie containing a lot of flashbacks. The analysis of the temporal relations of the various narrative sequences (also called *diegetic chronology*) can be used to discourse about the narrative structure. We define an Advene schema called *diegetic chronology*, that contains two annotation types: *shot*, that represents a shot as the basic unit in movies and *diegetic sequence* that represents a chronologically consistent unit. A relation type, *followed by*, will allow to link a sequence to the following one in the diegetic chronology. Another schema, called *movie*, contains among other types a *character* annotation type, that represents a character. Figure 1 sums up these schema, figure 2 present how it is possible to annotate a movie using the diegetic chronology schema.

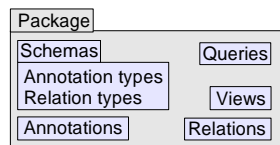


**Fig. 1.** The *diegetic chronology* and *movie* schema



**Fig. 2.** An annotated audiovisual document

One of the design goals of Advene is to allow users to specify themselves how they want the annotations to be rendered. The Advene model defines the notion of *view*, that represents a way to display annotations. Moreover, visualising data also means selecting the data to be visualised: a *query* represents a way to select elements from the annotation structure.



**Fig. 3.** Overview of the Advene model

Figure 3 gives an overview of the different elements of the Advene model. They are stored in documentary units called *packages*. A package is a document that holds all relevant information (schemas, annotations, queries and views) allowing to exchange, modify and visualise the metadata associated to an audiovisual document. Being separate from the audiovisual document, it can be modified and exchanged independently.

## 2.2 The Advene prototype

The Advene model is fairly generic. Some decisions regarding the implementation of query or view languages had to be made in the prototype.

**Visualising annotations in the Advene prototype** The Advene model defines a notion of view, without specifying what is in a view, which is a decision left to the implementation. The Advene prototype proposes three types of views: ad-hoc views (GUI views), static views (HTML templates) and dynamic views (set of rules allowing to dynamically modify the movie rendering).

*Ad-hoc views* are programmed views built in the GUI, that the user can configure. They feature standard views found in audiovisual software (time-line views, hierarchical data view, transcription view...).

*Static views* are XHTML templates that can be applied on the data. We are reusing the ZPT (Zope Page Templates) template system from the Zope platform [6]. This template system is oriented towards XML templates edition, using attributes in a dedicated namespace as processing instructions. Thanks to this attribute-based approach, both templates and result documents are valid XML documents, which allows us to process them with standard XML processing tools, like the *epoz* WYSISWYG browser-based editor [7] that has been integrated in the prototype.

Another component brought by the ZPT framework is the TALEX syntax, that proposes a simple, path-like addressing scheme to address elements from a data model. This approach does not try to be a full query language, such as XPath wrt. XML, but instead to provide a simple, user-accessible way of addressing elements. For instance, the expression `/package/annotationTypes/sequence/annotations/first/content/data` addresses the content of the first annotation of type *diegetic sequence*.

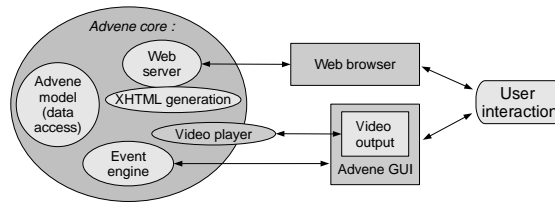
*Dynamic views* are able to dynamically change the way the movie is played, based on the annotations' content. Using a rule-based model similar to the filtering capabilities of e-mail software (Event-Condition-Action [8]), dynamic views allow the user to specify various actions to be executed when some events occur. The actions range from simple VCR-like functionality (pause, go to a position, stop...) to more elaborate video control (display captions – text or graphic – on the video, get a snapshot...), and also provide user-interaction facilities (information popups, navigation popups offering to go to another position...). The events are triggered by the annotation structure (annotation begin, annotation end...) or by user actions (player pause, player start...).

With this simple rule-based specification, it is possible to enrich the movie with information issued from the annotation structure, or even change the way the movie is played. The rule *When the event **annotation begin** occurs, display the annotation content as a caption if the annotation type is **sequence*** displays the sequence title over the video. The rule *When the event **annotation end** occurs, go to the beginning of the related annotation if there exists a **followed by** relation* will make a dynamic montage of the movie, restoring the diegetic chronology of the sequences.

*Queries* offer a way to select elements from the annotation structure. A simple query implementation has been integrated in the prototype, using the same framework as the dynamic views: elements matching a given condition can be extracted from a given set of elements. This approach has proved flexible enough to accommodate various needs in our experimentations: selecting elements based on their contents, their temporal relationships (through Allen relations) or their relations.

**Architecture of the Advene prototype** The open-source Advene prototype reuses standard software components: it embeds the versatile, open-source and cross-platform VLC video player [9], uses the ZPT template model from Zope, and uses a standard web browser to visualise the rendition of the ZPT templates. Figure 4 provides an overview of the prototype architecture.

The Advene prototype has been written in python, which proved an excellent choice for rapid development and experimentation. It provides a testbed for the development



**Fig. 4.** The Advene prototype architecture

of new ideas in the field of multimedia annotation handling and visualisation. It is being used in ongoing collaborations with researchers in human interactions (who study video recordings) or movies study, as well as by individual researchers that use audiovisual material. The following section describes how we have integrated OWL in Advene.

### 3 Integrating OWL in Advene

In this section, we demonstrate how OWL descriptions and inferences integrate smoothly inside the Advene model and prototype presented before. This will be illustrated with the example Advene schemas from the previous section.

We propose here a two-steps integration of OWL in Advene: exposing Advene elements as an OWL description, then performing inference over it and getting the results back into the Advene model. The first step is achieved by using Advene *views*, while the second one is performed by dedicated Advene *queries*.

#### 3.1 OWL views

Exposing Advene elements as OWL can be done in two non-exclusive ways: using views to “expose” Advene structures as OWL structures, or putting OWL statements inside annotations/relations (as their content).

**Viewing Advene structures as OWL** The first way is a direct application of the general notion of view in Advene, using OWL as a target format. It is straightforward in the current implementation since everything in an Advene package already has a URI, and since ZPT (the template language used to define static views) is able to produce any kind of XML document.

First, an XML view can do a straightforward “translation” of the Advene structure, according to an ontology of the Advene model (with classes such as Package, AnnotationType, Annotation, etc.). Such a view can be imported in any other Advene package and enhance it with OWL export capabilities.

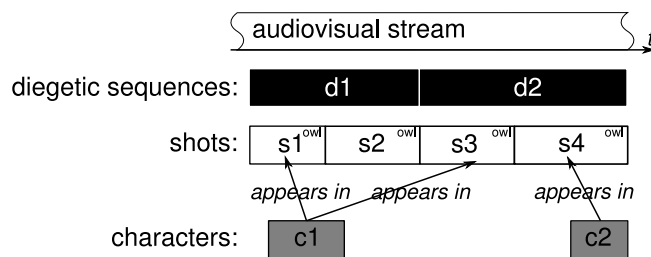
Beyond that simple translation, schema authors may want to devise more specific OWL views in order to embed the underlying semantics of their schema, so that the produced OWL statements are more richly describing the annotation structure. For example, we can imagine that the designer of *diegetic chronology* would represent the

binary relation *followed by* by an OWL property rather than by OWL instances<sup>2</sup>, and impose that *diegetic sequence* annotations be followed by at most one other sequence.

Another approach is to design an Advene schema on top of a pre-existing ontology in order to use that ontology in the context of video annotation. For example, one could want to describe people or characters appearing in a video with the FOAF<sup>3</sup> ontology, which allows to describe persons, relations between them, and the groups to which they belong. This can easily be done by defining a schema with an annotation type for each class of the ontology and a relation type for each property, then provide a view converting annotations complying to this schema into an OWL description complying with that ontology. This scenario shows that, instead of considering the OWL description as a by-product of the Advene package, one can consider Advene as a front-end tool for annotating videos with OWL.

**OWL statements in annotations/relations content** Putting OWL statements inside annotations/relations is also a straightforward application of Advene principles, which does not impose any data type on their content. One could for example add such a content in annotations of type *shot* in order to formally describe the depicted scene (e.g. with the ontomedia ontology [10], intended to describe fictional films). One could then query each annotation individually to perform inference over its content. But inference would not here take advantage of the fact that annotations are attached to a fragment of the video stream.

Yet the anchoring of OWL statements in the stream can be taken into account by *grouping* several contents, according to various criteria which we call *reasoning contexts*, and which can in turn be materialized by other annotations. For example, one might want to reason on the content of all shots temporally contained in a given diegetic sequence. Or, assuming that a relation *appears in* exist between *character* and *shot* annotations, one might want to reason on the content of all shots where a given character appears (see figure 5).



**Fig. 5.** Annotations of type *character* or *diegetic sequence* can be used as reasoning contexts for the OWL contained in *shot* annotations.

<sup>2</sup> Advene relations may involve more than two annotations, and they may have a content. The default is hence to represent them by OWL instances.

<sup>3</sup> <http://xmlns.com/foaf/0.1/>

Using annotations as reasoning contexts over multiple OWL-annotations can be achieved by defining specific views over the context annotations themselves, where an OWL ontology is generated, importing all the contents of the relevant annotations. TALES expressions and Advene queries can indeed be used in views to retrieve annotations based on temporal relationship or Advene relations. In the example of figure 5, applying the “temporal inclusion” view to *d1* would generate an ontology importing the contents of *s1* and *s2*, while applying it to *d2* would import *s3* and *s4*. On the other hand, a view using the *appears in* relation would generate an ontology importing *s1* and *s3* when applied to *c1*, and only *s4* when applied to *c2*. We see that OWL statements can be used in different context, depending on the point of view used to group them.

**Mixing the two** Of course, those two approaches can be mixed: specific OWL views can take advantage of both the annotation structure and contents to provide rich descriptions of the annotated video, as described in [11].

### 3.2 OWL Queries

We just saw how Advene elements can be viewed through an OWL description. In principle, any DL inference service [12] can be used to query that OWL description. However, we mainly focused on the use of A-box querying (see section 4), for it integrates smoothly with the notion of query in Advene, as we will see.

For this purpose, we use the SPARQL language<sup>4</sup>, a general purpose RDF query language, to query the *deductive closure* of OWL views (i.e. not only the expressed triples, but also all the triples which can be inferred from them). More precisely, we restrict ourselves to SELECT queries<sup>5</sup>. The sample query in figure 6 illustrates the main features of SPARQL. The PREFIX clauses define namespace prefixes used in the other clauses. The FROM clause locates the source of information to be queried. The WHERE clause describes a subgraph to be searched, where some nodes (whose name starts with a question mark ‘?’) are variables. The OPTIONAL sub-clause describes a part of the subgraph which is not required to match the query. Finally, the SELECT clause indicates which variables are to be returned (variables from the OPTIONAL clause may have a null value). The result of such a query is a list of tuples, each one being a binding of the selected variables, satisfying the query. We will now show how this is compatible with the notion of query in Advene.

Queries in Advene are used as *filters*: from a set of Advene elements (possibly the whole package), they select the subset of elements matching the query. SPARQL queries in Advene only require that the initial set of items is described in RDF (which has been made possible by the OWL views described previously), and that the URI bound to the variables in the result are converted back to the Advene element they identify<sup>6</sup>.

---

<sup>4</sup> <http://www.w3.org/TR/rdf-sparql-query/>

<sup>5</sup> SPARQL has other kinds of queries (CONSTRUCT, DESCRIBE) but they have different kinds of results, which do not fit in Advene queries

<sup>6</sup> Actually, there are two more slight differences: SPARQL queries return a set of *tuples* rather than a set of single elements, and those tuples may not only contain Advene elements but



```

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
SELECT ?x ?nx ?g
FROM <http://localhost:1234/packages/nelson/view/foaf_report>
WHERE {
  ?x rdf:type foaf:Person ;
     foaf:name ?nx ;
     foaf:knows ?y .
  ?y foaf:name "Ted Nelson" .
  OPTIONAL {
    ?g rdf:type foaf:Group ;
       foaf:member ?x .
  }
}

```

**Fig. 6.** A SPARQL query, retrieving the URI and name of every person knowing Ted Nelson, the URI of the group they belong to if any.

**Implementation** We have implemented a basic SPARQL support in Advene using Pellet<sup>7</sup> as an external inference and query engine: Pellet is indeed able to perform A-box querying with SPARQL over OWL models. Pellet accesses the OWL views and the query through the HTTP server embedded in the Advene core (see figure 3).

### 3.3 OWL in Advene: a working illustration

So as to demonstrate the feasibility of our approach for integrating OWL in Advene, we designed several schemas that correspond to the different possibilities we presented above. All these examples concern an example video featuring a ten minutes Ted Nelson speech at the ACM Hypertext 2003 conference. A first package had been designed as a proof of concept of various Advene features. We have extended this package with semantic web capabilities as a demonstration of the ideas presented in this article<sup>8</sup>. We will further refer to it as the *nelson-sw* package.

As a first proof of concept, we designed an ontology for the simple straightforward OWL translation of an Advene package, and a package containing the associated ZPT views for effective transformation. This package is imported by the *nelson-sw* package and provides it with generic OWL export according to the aforementioned ontology.

To go beyond that simple translation, we experimented the idea of enriching the resulting ontology with some knowledge that could enable reasoning over the package. For this, we created a “Virtual Montage” schema, defining a dynamic view driven by two binary relations: *boundaries* is used once to define the start and end of the virtual montage, and *jumpTo* is used between annotations in order to force the player to jump

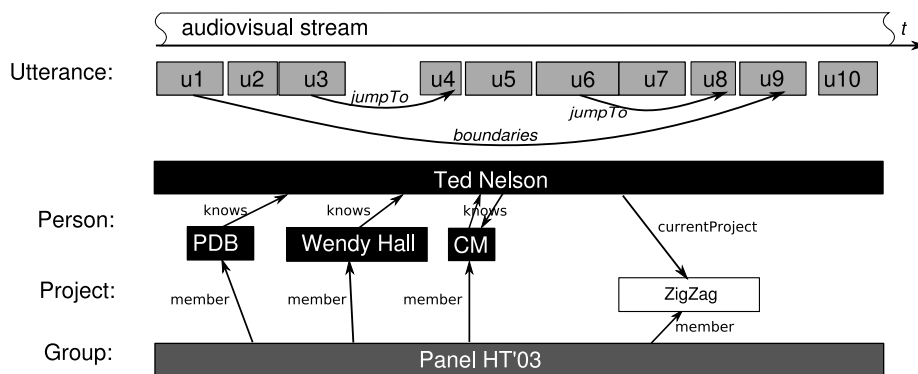
---

also *strings* (literals and unresolved URIs). However, the current implementation of static and dynamic views has no problems dealing with tuples and strings.

<sup>7</sup> <http://www.mindswap.org/2003/pellet/>

<sup>8</sup> <http://liris.cnrs.fr/advene/packages/nelson-sw/>

from the end of the first annotation to the beginning of the second one (see figure 7). OWL is used here to check that those relations are used in a consistent manner: there should be exactly one start and one end, and the jumps should form a single path (no loops, no branches) leading from the start to the end. A specific OWL view exports an ontology expressing those constraints in addition to the description of the annotation structure. The resulting OWL will hence be semantically consistent only if the virtual montage is valid. By importing this schema and view in the *nelson-sw* package, and by re-using *Utterance* annotations (created as a transcription of the discourse), we were able to build a two minutes long abstract of Ted Nelson’s speech, after checking that all the constraints for the montage were fulfilled.



**Fig. 7.** An extract of the *nelson-sw* package demonstrating the “Virtual Montage” and “FOAF” schemas

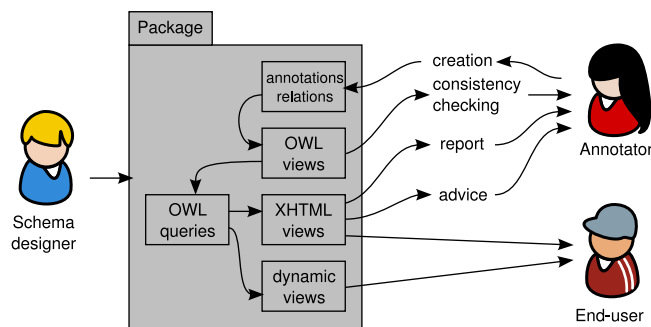
Another possibility we mentioned in section 3.1 was to design an Advene schema on top of a pre-existing ontology so as to use that ontology for audiovisual annotation with Advene. As an example, we used the FOAF ontology, considering three of its concepts (Person, Group and Document) and the relations between them for being represented respectively as annotation types and relation types from a dedicated schema, which has also been used in the *nelson-sw* package (see figure 7). A package using that schema can then be used as any other Advene package, enhanced with validation and inference services thanks to the associated ontology (see section 4); but Advene can also be seen in this case as a front-end for annotating a video with FOAF, considering the OWL view as the primary output target (with some added value, such as e.g. the automatic generation of a “depiction” property linking instances of Person to a snapshot of the annotated video fragment).

Note also that, beyond the translation of “external” structures (i.e. inter-linked annotations and relations), the “internal” structure (i.e. content) of annotations and possibly relations, can also be used in the OWL generation. We mentioned that annotations can contain OWL data. In our schema, they contain a simply structured text where each line has the form *key=value*. The OWL view takes advantage of this internal structure,

by recognizing a number of FOAF properties (name, homepage, currentProject...) and generating the appropriate statements about the corresponding resource. We can also imagine that annotations with type Person would simply contain the URL of the person's FOAF profile, which would then be imported in the OWL view.

## 4 Using OWL in Advene

In this section we explore the benefits, from the point of view of video annotation, of OWL-enhanced Advene. We do so by presenting a number of prospective scenarios made simple with OWL inference when they would have been complicated, when not unfeasible, with more “classical” queries and views. As illustrated in figure 8, some scenarios are focused on helping the annotator in her task, while others are focused on the end-user.



**Fig. 8.** The schema designer provides a schema (not represented) with OWL views, OWL queries, XHTML views and dynamic views adapted to that schema. The annotator creates annotations complying with the schema. Some of the views help her in her annotation task, while other views are aimed at the end-user.

### Checking consistency

We already mentioned that OWL allows schema designers to express semantic constraints on the use of their schemas. For example, the aforementioned “Virtual Montage” schema states in OWL that no annotation should be the source or the destination of more than one jump (i.e. that the montage should be linear). OWL consistency checking can then be used by the annotator to ensure that her annotation structures complies with the underlying semantics of the schema. Some engines, including Pellet, even provide human-readable explanations of why a given OWL description is inconsistent (see figure 9, section “Consistency”). Indeed, inconsistency will generally lead to unsatisfactory results from the other views provided with the schema; in the case of the virtual montage, only one path will be followed when playing the video.

## Reporting

Specific queries and static views can also be provided by the schema designer to the annotator, for her to check that everything complies with the intended semantics. A first advantage over plain consistency checking is that the “errors” can be more specifically explained in those views. Another advantage is that SPARQL queries can be more expressive than OWL classes<sup>9</sup>. A third advantage, and an important one, is that such reporting views provide a *finer grain* than boolean consistency: it can be only *suggested* annotations respect certain constraints, though they can in some circumstances violate them.

In the *nelson-sw* package, this kind of inference has been used to detect ill-used Advene relations. In FOAF, members of a group can be persons or other groups –more precisely, any *Agent*, a common superclass of Person and Group. Since Advene has no notion of superclass between annotation types, the Foaf Advene schema allows its “member” relation to link a group with any other annotation, so that annotations of type Person or Group can be used. It follows that even a Project can become a member of a group (see figure 7), though Projects are not Agents<sup>10</sup>.

An OWL view of the package is hence inconsistent when referring to the strict FOAF ontology. However, we provide with the package a more permissive version of the FOAF ontology, where all class disjunction axioms have been replaced by a new class named Absurd. This class is defined to contain the intersection of all pairs of (originally) disjoint classes. Hence, an instance of Project which should also be an Agent (because it is a member of a group) does not make that ontology inconsistent, but is bound to belong to the Absurd class. A list of absurd instances can then be given to the annotator for her to check (see figure 9, section “Report”).

## Advising

Moreover, some static views can even *advise* completions or modifications of the structure or content of Advene elements, based on *annotation patterns* expected by the schema designer.

For example, in the *nelson-sw* package, the “knows” Advene relation enables the annotator to state that a person knows another one. But the FOAF profile of a person (represented by the content of the Person annotations) may also contain this kind of information, under the form “I know the person whose homepage is ...”. In FOAF, the “homepage” property unambiguously identifies a person. Hence, the OWL reasoner can infer that two persons are the same, then reconstruct a social network by aggregating local information contained in each profile. In our example package, the profile of Ted Nelson states that he knows someone whose homepage happens to be the homepage of Wendy Hall (also explicitly represented by an annotation). The OWL reasoner will

---

<sup>9</sup> More precisely, class expressions can not contain any cycle in the graph representing their instances, while SPARQL query patterns can involve cycles. On the other hand, OWL classes can be reasoned about independently of their instances, while queries only cope with instances.

<sup>10</sup> This constraint is not actually stated in the original FOAF ontology, it seems reasonable enough, and serves the purpose of a simple example.

hence infer that Ted knows Wendy. A specific view will use this inferred implicit knowledge and suggest the annotator to make it explicit by putting a “knows” Advene relation between the annotations representing those two people (see figure 9, section “Report”).

## Consistency

**Consistent:** No


**Reason:** Individual <http://iris.cnrs.fr/advene/packages/nelson-sw/unstable/index.xml#a505> is forced to belong to class <http://xmlns.com/foaf/0.1/Agent> and its complement

## Report

The following are inconsistent with the ontology. Check the relations.


- [a505 \(type Project\)](#)
  - currentProject of [a502 \(type Person\)](#)
  - member of [a504 \(type Group\)](#)

## Advice



Ted Nelson knows

- [Cathy Marshall](#)
- [Wendy Hall](#)



Paul de Bra knows

- [Ted Nelson](#)

Fig. 9. A static view generated from the *nelson-sw* package

## Generating more hypermedia

We just saw how static views can be generated thanks to OWL inference for helping annotators in their task. The same mechanism can of course be used to generate hypermedia aimed at end users. Non-trivial queries can be used to extract a set of elements (e.g. “all persons known by a given person”), then use them to generate an appropriate static view (list of those person with a photo extracted from the video and a link to their homepage, graphical representation of their social network, ...) or dynamic view (the “Virtual Montage” package, link from a sequence depicting a person to the sequence depicting people he or she knows, ...).

## 5 Discussion and related work

Despite the long acknowledged need for semantically annotating multimedia documents, the unification of multimedia annotation standards with Semantic Web technologies is still a work in progress. An alleged difficulty for this unification is the lack of interoperability between standards [5, 13], especially between XML-based MPEG-7 and RDF-based OWL. Interoperability has however to be achieved since Semantic Web technologies “as is” are not quite adapted to multimedia annotation —though some approaches attempt to fill that gap [14]. Various approaches have hence been proposed, either to convert MPEG-7 structures into RDF based language in order to be able to reason about them [15], or to embed OWL ontologies into MPEG-7 structures in order to take advantage of standard-compliant tools, while retaining the semantics of the description [16]. While the former may be compared with the first approach presented in section 3.1, the latter can be compared to our proposition (in the same section) of designing Advene schemas according to existing ontologies.

Another hindrance to the large adoption of multimedia annotation in general is the complexity of the dominant standard MPEG-7 [5]. It is interesting to note that Semantic Web annotation is often the target of the same criticism, as working with (sometimes big) formal ontology requires some training for unskilled users. Advene eschews both by relying on a simple and extensible model for video annotation, and by not relying on formal ontologies from the bottom; we rather propose to use third-party or ad-hoc OWL ontologies on an opportunistic basis, i.e. when (and if) they can prove useful in a given context. By doing so, we argue that Advene meets the requirements for practical multimedia annotation expressed by [17].

Finally, the Advene architecture can provide the functionalities targeted by other approaches: controlling and checking the structure of annotations [16] as seen in section 4, semantic information retrieval [14] thanks to OWL queries, virtual montage [18]. But advantage can also be taken from semantic annotation by other uses, such as enriched video viewing, which are not, to our knowledge, addressed by this community.

## 6 Conclusion

In this article we have presented some ideas for integrating semantic web technologies in an annotation-based hypervideo system, most of them already implemented in the Advene prototype, thanks to its versatility and to the simplicity of the Advene model. Current work on the “semantic web side” of the Advene project entails smoother integration of OWL-queries in the prototype and graphical editing of such queries, design of more OWL-views for consistency checking, reporting and advising, design of reasoning-enriched dynamic views, and theoretical study of the notion of “reasoning context”. The Advene prototype is freely downloadable and extensible, and we encourage anybody to use it for testing new ideas on multimedia and semantic web.

## References

1. Aubert, O., Prié, Y.: From video information retrieval to hypervideo management. In: Corimedia, the international workshop on multidisciplinary image, video, and audio retrieval and mining, Sherbrooke, Canada (2004) 10 pp.
2. Aubert, O., Prié, Y.: Advene: active reading through hypervideo. In: ACM Hypertext'05. (2005)
3. Sanchez, J.M.M., Koenen, R., Pereira, F.: MPEG-7: The Generic Multimedia Content Description Standard, Part 1. *IEEE Multimedia Journal* **9**(2) (2002) 78–87
4. Pfeiffer, S., Parker, C., Schremmer, C.: Annodex: a simple architecture to enable hyperlinking, search and retrieval of time-continuous data on the web. In: 5th ACM SIGMM International workshop on Multimedia information retrieval. (2003) 87–93
5. van Ossenbruggen, J., Nack, F., Hardman, L.: That obscure object of desire: Multimedia metadata on the web, part 1. *IEEE MultiMedia* **11**(4) (2004) 38–48
6. Zope Corporation: Zope Page Templates reference. (2004) <http://www.zope.org/Documentation/Books/ZopeBook/2.6Edition/AppendixC.stx>.
7. Jablonski, M.: Epöz, a cross-browser WYSIWYG editor for Zope. (2003) <http://epöz.sourceforge.org/>.
8. Paton, N.W., ed.: *Active Rules in Database Systems*. Springer Verlag, New York (1999)
9. Fallon, H., de Lattre, A., Bilien, J., Daoud, A., Gautier, M., Stenac, C.: *VLC User Guide*. VideoLAN Project. (2003)
10. Lawrence, F., Tuffield, M.M., Jewell, M.O., Prügel-Bennett, A., Millard, D.E., Nixon, M.S., Schraefel, M., Shadbolt, N.R.: *OntoMedia - Creating an Ontology for Marking Up the Contents of Heterogeneous Media*. In: *Proceedings of Ontology Patterns for the Semantic Web ISWC-05 Workshop*, Galway, Ireland (2005)
11. Troncy, R.: *Integrating Structure and Semantics into Audio-visual Documents*. In: *Second International Semantic Web Conference (ISWC2003)*, Sanibel Island, Florida, USA, Springer (2003) 566–581
12. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.F., eds.: *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press (2003)
13. van Ossenbruggen, J., Stamou, G., Pan, J.Z.: *Multimedia Annotations and the Semantic Web*. In: *Proc. of the International Workshop on Semantic Web Case Studies and Best Practices for eBusiness (SWCASE)*. (2005)
14. Isaac, A., Troncy, R.: *Using several ontologies for describing AV documents : a case study in the medical domai*. In: *2nd European Semantic Web Conference, Workshop on Multimedia and the Semantic Web*, Heraklion, Crete (2005)
15. Hunter, J.: *Adding Multimedia to the Semantic Web - Building an MPEG-7 Ontology*. In: *International Semantic Web Working Symposium (SWWS)*, Stanford (2001)
16. Troncy, R., Carrive, J.: *A reduced yet extensible audio-visual description language*. In: *Proceedings of ACM Document Engineering*. (2004) 87–89
17. Geurts, J., van Ossenbruggen, J., Hardman, L.: *Requirements for practical multimedia annotation*. In: *Workshop on Multimedia and the Semantic Web*, Heraklion, Crete (2005) 4–11 part of 2nd European Semantic Web Conference.
18. Bocconi, S., Nack, F., Hardman, L.: *Supporting the generation of argument structure within video sequences*. In: *ACM Hypertext'05*. (2005) 75–84