

What's on this evening?

Designing User Support for Event-based Annotation and Exploration of Media

André Fialho¹, Raphaël Troncy², Lynda Hardman¹,
Carsten Saathoff³ and Ansgar Scherp³

¹ CWI, Amsterdam, The Netherlands, <firstname.lastname@cwi.nl>

² EURECOM, Sophia Antipolis, France, <raphael.troncy@eurecom.fr>

³ WeST Institute, Koblenz, Germany, <lastname@uni-koblenz.de>

Abstract. We present an event-based approach for users to explore, annotate and share media. We are constructing a web-based environment that allows users to explore and select events, including discovering meaningful, surprising or entertaining connections among them. We build a knowledge base of events from event directories that will be linked to the Linked Open Data (LOD) cloud, in conjunction with event and media ontologies. The approach is user-driven and, having carried out initial user inquiries, we are designing interfaces that support user-identified tasks while exploring the connections between users, multimedia content and events.

1 Introduction

As with all developing technologies, it is difficult to identify novel user needs that can be satisfied with emerging semantic web technologies. At the same time, it is difficult to develop the technology in specific directions without knowing what users are likely to want to do with the technology. In previous work we have identified comparison search tasks that can be supported using a combination of thesaurus-based linked data search and a modular user interface design [1], and also historical print annotation tasks [5] that can be supported using a combination of existing RDF data sets, semantic search functionality and task-oriented user interface.

In the context of the Petamedia⁴ Network of Excellence, we are exploring a similar method for designing an application that takes into account the “triple synergy” of users and their social networks, user-created content and metadata attached to this content in an application for supporting users in interacting with events. Events are a natural way for referring to any observable occurrence grouping persons, places, times and activities that can be described [10, 9]. Events are also observable experiences that are often documented by people

⁴ <http://www.petamedia.eu>

through different media (e.g. videos and photos). We explore this intrinsic connection between media and experiences so that people can search and browse through content using a familiar event perspective.

While wishing to support such functionality, we are aware that websites already exist that provide interfaces to such functionality, e.g. eventful.com, upcoming.org, last.fm/events, and facebook.com/events to name a few. These services have sometimes overlap in terms of coverage of upcoming events and provide social networks features to support users in sharing and deciding upon attending events. However, the information about the events, the social connections and the representative media are all spread and locked in amongst these services providing limited event coverage and no interoperability of the description. Our goal is to aggregate these heterogeneous sources of information using linked data, so that we can explore the information with the flexibility and depth afforded by semantic web technologies. Furthermore, we will investigate the underlying connections between events to allow users to discover meaningful, entertaining or surprising relationships amongst them. We also use these connections as means of providing information and illustrations about future events, thus enhancing decision support.

The work reported here uses an explorative user-centered design approach, where users are asked about real-world tasks they would like to carry out, and then asked for their opinions on specific technologies that they are familiar with and how these might be used to support the tasks. This approach ensures making design decisions that contribute towards an efficient, effective and satisfying user experience. Section 2 describes the method and the results of this user study, and presents the requirements for an event-based system for discovering and sharing media. Section 3 describes the event and media ontologies developed to support the semantic description of events extracted from event directories. Section 4 explains the design rationale of the interfaces, and gives some interface mockups to illustrate the types of task support and expected functionality we will be providing in the coming months. Finally, we give our conclusions and outline future work in Section 5.

2 User Need Assessment

We follow a user-centered design process consisting of an assessment of user needs and insights, identified through interaction with potential users at different stages of development. Our research starts by identifying who the users are, their interests, their goals and which tasks need to be supported in order to achieve these goals. We collect this information to define a first set of requirements and identify prospective scenarios that illustrate the environment task scope and a first design concept. The steps that follow consist of iterative cycles of re-design and evaluation until a satisfactory design is reached.

2.1 Method

The first step of our research was done in order to collect potential end-user experiences, opinions and interests while discovering, attending and sharing events, and user insights about potential web-based technologies that support these activities. We collected this initial input through an exploratory study with 28 participants (11 females). Participants were mostly students and researchers with ages varying from 23 to 47 years old. The study was done through an on-line survey with 8 questions divided into 2 sections. The same topics were then presented in discussion sessions with two groups of master students totaling 35 additional participants: One discussion was done with students (n=10) from an Interactive Multimedia Systems course and the other with students (n=25) from a Human-computer interaction for the Web course. The results from these discussions were used to validate the survey responses and to extend it with other collected insights.

The first half of the survey aimed at identifying participants' personal experiences and behaviors. It invited them to recall memorable previously attended events (e.g. festivals, conferences, concerts, art galleries, exhibitions, gatherings) and to share their opinions and experiences regarding: (1) how events are discovered; (2) characteristics that support deciding rather or not to attend to an event; (3) how the event experiences are registered and shared; and (4) meaningful, surprising or entertaining relationships amongst events.

On the second part of the survey, participants' were asked to share opinions regarding existing web technologies in the context of the aforementioned activities. To better address the triple synergy paradigm (Section 1), we explored the concept of merging event directories, media directories and social networks. With that in mind, we asked participants to share their opinions regarding: (1) the perceived benefits and drawbacks of event directories (e.g. Eventful); (2) enabled possibilities, benefits and drawbacks of combining media sharing websites (e.g. YouTube and Flickr) with event directories; (3) enabled possibilities, benefits and drawbacks of combining social networks (e.g. Facebook and Twitter) with event directories; and (4) suggestions regarding desired and useful features.

Answers obtained from the survey were analyzed through affinity diagramming. The process consists of iterative clustering cycles which allow organizing the collected ideas into common themes, thus allowing to identify the most common opinions for each raised subject. The results from this first exploratory study are described in the following section.

2.2 Results

In this section, we present a summary of the results of our user study. The summary contains main reported experiences, interests and opinions around event related activities.

Past experiences. Concerning participants' experiences when discovering events, the vast majority reported to find out about events through invitations and recommendations from friends and colleagues. Traditional media such as posters,

flyers, news articles and television ads seem to play a major role when discovering events. Social networks were also reported to be used, with specific reference to event posting and invitation features. More seldom participants use event directories (e.g. Livenation, local city event directories, Ticketmaster, last.fm) or participate in mailing lists, newsletters and forums to obtain updates. The use of search engines was reported, specifically when they knew what to look for. Moreover, participants also rely on previously attended events or venues as reference for finding new events. During the group discussions, participants seemed to rely more heavily on social networks in comparison to the survey responses.

When deciding whether or not to attend to an event, participants seem to prioritize background information. Location was often referred to for orientation and because of distance constraints. Price was commonly mentioned to allow identifying cost-benefit ratios and due to budget constraints. Time of the event was a main decision factor, followed by information about who else would be joining the event, and more specifically, which friends will attend. The content of the event itself (e.g. type, performer, topic) and subjective factors such as fun, relevance, interest, atmosphere, target audience and reputation were also mentioned. Students from the group discussions, preferred the event attendance (“who’s joining?”) and price constraints over all other characteristics.

Regarding how participants register the experience, they often take pictures for sharing after the event. Less commonly, participants record short videos. As for the how they share the information, they most commonly talk to others, describing their experience. Participants share the collected media directly (e.g. file transferring, showing on the mobile) or use media directories and social networks such as Facebook, Flickr or YouTube.

Concerning relationships between events, the most referred to characteristics that motivate participants to look into related events were the event categories (e.g. type of event, topic, genre). Another important factor was the event attendees, to find other events they would attend. This could refer to groups of people (i.e. target audience, users with similar interests), but most importantly, individuals in their social networks. Other main event characteristics also mentioned were: location, performers, organizers and time/duration. Lastly, future events from repeated events was also seen as a strong relationship.

Existing technologies. Existing event directories (e.g. eventful) perceived benefit was clearly to be a single access point providing an overview of event information. Another reported benefit is that it supports opportunistic event discovery and facilitates exploration based on different contexts (e.g. location, popularity, categories). Other positive features include: social features (e.g. commenting, sharing events), notification of upcoming events, and shortcuts (e.g. ticket purchase). As for the drawbacks, the main reference was about the unreliability (i.e. unknown source) and incompleteness of information. In particular, low coverage of events and insufficient information for decision support (e.g. lack of location map, videos) have been mentioned. In contrast, the information overload was also seen as a potential drawback making it difficult to find specific events.

When presented the possibility of combining media and event directories, participants recognize benefits due to information enrichment. They claim it would help illustrate events with videos and pictures of past related events (i.e. past performances), other people's experiences, promotional (marketing) material, and so on. The main recognized value would be to give a better idea about the event's environment/atmosphere and provide visual information to support decision making. Participants said it would also support remembering and sharing past experiences. Drawbacks from the merger concern information overload and privacy issues while sharing personal media.

Regarding the possibilities afforded by merging social networks and event directories, some participants think that the main benefits are communication between users, and the sharing of more information (e.g. invitations, opinions, pictures). It was also said to facilitate viewing event attendance, identifying event popularity, and even provide an overview about friends' whereabouts. Live event information (e.g. real-time tweets and comments, live pictures/streams) updates were also seen as a positive afforded feature. Despite the benefits, some participants think the amount of information could clutter the service. Others pointed out that services such as Facebook already provide enough event sharing features. Suggestions included making use of existing social network profiles and/or extending these services.

Other features that users would appreciate having when dealing with events were broadly described with little overlap. Some of these features are: recommendations (based on past attendance, preferences, and from people with similar interests); better visualizations for exploring and searching events (e.g. map integration); the potential to combine categories and attributes while browsing; obtain more information about events and users (e.g. opinions, price and availability).

Conclusions. The opinions gathered seem to support the development of an environment that merges event directories, social networks and media sharing platforms. Moreover, this information enrichment is thought to provide better means of supporting the decision making process. This assumption is based on the possibility of allowing users to better experience an event by viewing associated media. On the other hand, social information obtained implicitly (behaviors) and explicitly (comments, reviews and ratings) provide better judgments of events in terms of attendance, shared interests and reputation. A common concern about information overload suggests that the interface should avoid cluttering and provide only necessary information. Furthermore, there is a need to support different visualizations and better browsing possibilities depending on user interests and constraints. Lack of event coverage and information completeness is another important identified issue that can be addressed using and combining multiple information sources. These issues, along with other identified user interests were translated into a set of requirements in order to guide the following steps of the environment design and development. The following section describes these requirements in more detail.

2.3 User Requirements

Based on this user study, we define a first set of requirements, translating user needs into functionalities that the system should support (Table 2.3). It is important to note that the requirements presented here are representative of users who participated in the previous studies. They should be complemented with other non-functional and functional requirements as described in existing design patterns and interface guidelines [4, 7].

Discovering Provide a comprehensive coverage of past and upcoming events Allow searching events based on tags (e.g. performer name, genre, title) Allow opportunistic discovery by filtering and combining properties (e.g. categories, location, time, price)
Inspecting Show complete background information about events (e.g. title, location, description, venue, performers, time, category, genre, availability, size) Allow identifying subjective aspects of events (e.g. popularity, fun, atmosphere, reputation) Show media associated to events for reliving experiences and for decision support Show who is joining or joined the event (attendance) Allow identifying related and repeated events
Visualizing Rely on traditional media information display (e.g. posters, flyers, ads) Show only the necessary information in a simple way Allow different visualizations and browsing contexts (e.g. time, location, people)
Enriching Allow creating events Allow associating pictures and videos with existing events Allow associating comments and opinions with existing events
Sharing Make use of existing social networks (e.g. Facebook, Twitter) Allow inviting and recommending events using existing services
Recommending & Preferences Allow receiving recommendation about events based on personal interests and behaviors Allow receiving recommendations based on other people's preferences and behaviors (collaborative filtering) Identify interests and preferences based on past event attendance

Table 1. Requirements

2.4 Scenarios

Scenarios are informal narrative descriptions that allow exploration and discussion of context, user needs and requirements [2]. For the purposes of our research, we created a number of scenarios, each covering a range of the aforementioned

requirements and illustrating prospective goals and tasks supported by the system. To better emphasize the context and allow better interpretation and inference of user needs we created four personas. The personas were inspired by the different participants in the previous exploratory study and describe attributes and background information about the actors involved in the scenario. Characteristics are representative of different age groups, professions, preferences, and commonly used event, media and social network sources.

We provide below four different scenarios, each from one different persona. Scenario 1: *Johnny was invited to a party by a friend and receives a link providing information about this event. He wants to know when and where this event will be and who else was invited. More importantly, he wants to know whether his closest friends confirmed to attend the event or not.*

Scenario 2: *Julie would like to go to a play on her favorite theater. She wants to see a comedy, hopefully playing the upcoming week. She has only been to a few comedies, but she remembers one she specifically enjoyed. Julie would like to see if there is something similar playing and read what other people say about it.*

Scenario 3: *Jack recorded a video with his mobile phone camera while he was attending the Haiti Relief concert from Radiohead given on 24 January 2010 in Los Angeles. He thinks it was a really nice experience and wants to share it on-line. He would also like to see what other pictures and videos were captured during the concert and see how other people experienced the show.*

Scenario 4: *Jessica is going to Paris on her honeymoon and she would like to see what will be happening there during her stay. She wants to do many different things, but cannot decide yet, so she wants to put these things on a “maybe” list in order to decide later. If possible, she would like to see videos of these events to make sure it has a cozy and romantic atmosphere.*

3 Event and Media Ontologies

In this section, we present the ontologies used for representing events, media and users metadata. We use the scenario 3 described above to illustrate these models.

3.1 The LODE ontology

The LODE ontology⁵ is a minimal model that encapsulates the most useful properties for describing events [9]. The goal of this ontology is to enable interoperable modeling of the “factual” aspects of events, where these can be characterized in terms of the *four Ws*: *What* happened, *Where* did it happen, *When* did it happen, and *Who* was involved. “Factual” relations within and among events are intended to represent intersubjective “consensus reality” and thus are not necessarily associated with a particular perspective or interpretation. This model thus allows us to express characteristics about which a stable consensus has been reached, whether these are considered to be empirically given or

⁵ <http://linkedevents.org/ontology/>

rhetorically produced will depend on one’s epistemological stance. We exclude, at this stage, properties for categorizing events or for relating them to other events through parthood or causal relations. We will see in the next section how these aspects, that belong to an interpretive dimension, can be handled through the Descriptions and Situations approach of the Event-Model-F [8].

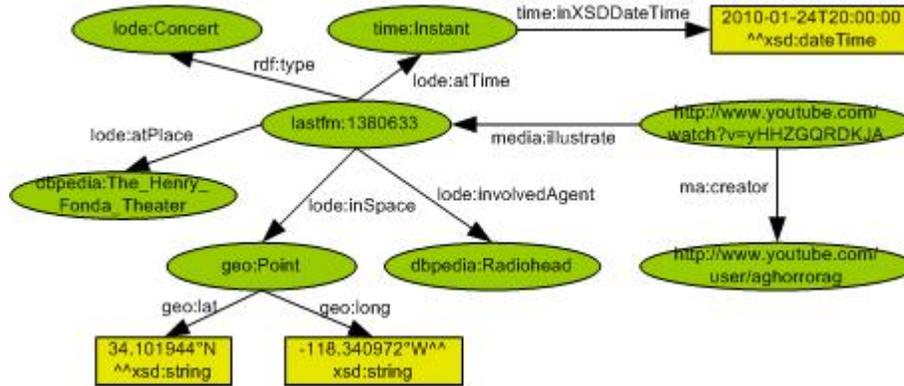


Fig. 1. The *Radiohead Haiti Relief Concert* described with LODE

The Figure 1 depicts the metadata attached to the event identified by 1380633 on last.fm according to the LODE ontology. More precisely, it indicates that an event of type *Concert* has been given on the 24th of January 2010 at 20:00 PM in the *Henry Fonda Theater* featuring the *Radiohead* rock band.

LODE is not yet another “event” ontology *per se*. It has been designed as an *interlingua* model that solves an interoperability problem by providing a set of axioms expressing mappings between existing event ontologies. Therefore, an OWL-aware agent would infer that the resource identified by *dbpedia:Radiohead* is a *dul:Agent* as described in the *Dolce Ultra Lite* ontology.

3.2 The Media Ontology

The Ontology for Media Resource currently developed by W3C is a core vocabulary which covers basic metadata properties to describe media resources⁶. It also contains a formal set of axioms defining mapping between different metadata formats for multimedia. In the Figure 1, we see that the video hosted on YouTube has for *ma:creator* the user *aghorrarag*.

The Ontology for Media Resource can then be used to attach different types of metadata to the media, such as the duration, the target audience, the copyright, the genre, the rating. Media Fragments can also be defined in order to have a smaller granularity and attach keywords or formal annotations to parts of

⁶ <http://www.w3.org/TR/mediaont-10/>

the video. The link between the media and the event is realized through the `illustrate` property, while more information about the user could be attached to his URI using for example the FOAF ontology.

3.3 The M3O and F Ontologies

The pattern-based ontologies M3O [6] and Event-Model-F [8] can also be used for associating a media with an event description, or for describing relationships between events. This results in a more complex description but brings more expressiveness for representing the context of an annotation such as stating its provenance. These ontologies are based on the foundational ontology DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering) [3].

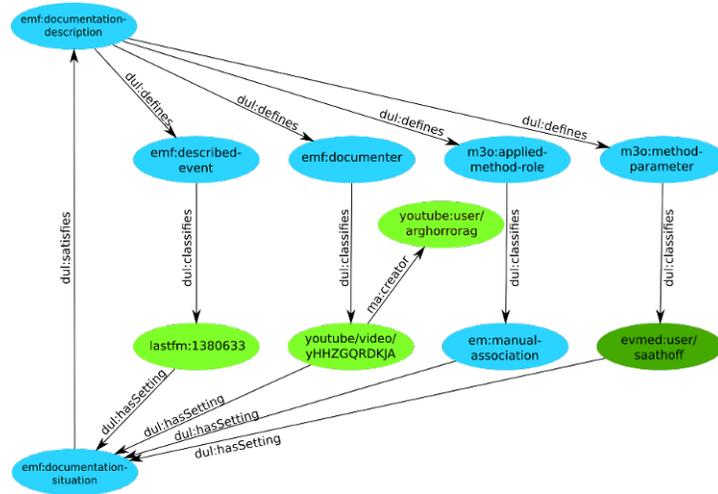


Fig. 2. Associating an event with a YouTube video and provenance information

The Figure 2 depicts the combination of the Event-Model-F Documentation Pattern and the M3O Provenance Pattern. While Media Ontology allows to specify the creator of the video, the patterns in Figure 2 express who actually created the association between the event and the video. We consider the Radiohead concert event to be a `described-event` and the YouTube video is classified as a `documenter`, expressing a reified `illustrates` relation. On the right hand side, we further detail this relation by adding a `em>manual-association` which is classified by an `applied-method-role` and the author of this manual association by adding a user `evmed:user/saathoff`, who is classified by the `method-parameter`. In summary, using the patterns of the Event-Model-F and the M3O, we can extend the LOD and Media ontologies with provenance information, making the distinction between the creator of some media or event and the creator of

the association between events and media, and even between the participants of this event.

3.4 Data Scraping and Semantization

We are populating these ontologies by scraping and semantifying data from event directories. As a first experiment, we explore the overlap in metadata between two popular web sites, namely Flickr as a hosting web site for photos and videos and Last.fm as a documentation of past and upcoming events. Explicit relationships between these two datasets exist using the `lastfm:event=XXX` machine tag. Hence, more than 1.5 millions photos are indexed with this tag yielding ten of thousands of events.

We use the Last.fm API to convert the event description into the LODE ontology (Section 3.1) and the Flickr API to convert the media description into the Ontology for Media Resources (Section 3.2). The result of this operation is a minimal description of the events where all values are strings (literals). Therefore, we perform an additional step in order to update this description into a truly linked data one. We invoke semantic web lookup services such as the dbpedia one in order to transform these strings into URI identifying unambiguously resources in the web of data. Hence, the "Radiohead" string is transformed into a dbpedia URI⁷ which provides additional information about the band such as its complete discography. This URI is declared to be `owl:sameAs` another identifier from the New York Times⁸ which provides information about the 38 associated articles from this newspaper to this band. The venue has also been converted into a dbpedia URI⁹ but has been augmented with geo-coordinates thus increasing the amount of information available in the LOD cloud for the benefit of all semantic web applications.

The linked data journey can be rich and long. One of the challenges we want to address is how to visualize these enriched interconnected datasets while still supporting the user tasks identified in the Section 2.3.

4 User Interface

In this section, we illustrate initial interface possibilities derived from the requirements and tasks presented in the Section 2.3. The interfaces are represented through low-fidelity prototypes. The prototypes allow exploring, refining and validating prospective concepts along with interface and interaction aspects through small studies with potential end-users and usability experts. Unsurprisingly, the sketches below correspond to the basic properties defined in the LODE ontology.

⁷ <http://dbpedia.org/page/Radiohead>

⁸ <http://data.nytimes.com/N12964944623934882292>

⁹ http://dbpedia.org/page/The_Henry_Fonda_Theater

4.1 Views and Perspectives

What - One prospective view is media centered and allows to quickly illustrate the event through associated media. In this view we display events through a representative images and convey different event characteristics (e.g. relevance, rating, popularity, etc) with one and/or more of the image properties, i.e., size and transparency. This approach has been used in other applications¹⁰ to represent clustered result sets or convey sorting by size on different contexts (Figure 3a).

When - Ordering can also be used to represent chronological event occurrence. In fact, the time centric view can be interpreted as the sorting of events chronologically (Figure 3b).

Where - A location centric view can be used to represent where the events occur geographically to orient the user and convey distance. The use of maps is commonly used to visualize such information (Figure 3c).

Who - Events are intrinsically bound to a social component. Users want to know who will be attending to an event when deciding to attend to it. In this context, a people centric view would be relevant to explore the relationships between users and events. Alternatively we can combine attendance information to other views such as location, allowing users to browse for friends on a map and identify their attended events. It could also be used to provide means of visualizing event popularity ,e.g. identify the cities hot-spots on a map, indicate visual cues of popularity according to number of attendees.

In order to allow users to relive experiences from events attended in the past, follow future confirmed events, and keep track of authored events, it is necessary to display events in the context of the users own attendance and ownership. For this reason we will support a “my events” feature with overall browsing possibilities. If several views are to be supported one challenge that can arise concerns transitions between these different views. This is specifically important for facet browsing, due to sudden disappearance of items during navigation [7]. Animated transitions could be used in order to allow the users to maintain orientation during such navigational changes.

4.2 Search Interface

When discovering events we believe users will also rely on browsing, which allow them to analyze large sets of event sets, and narrow them according to their interests and constraints. Overall, we believe users will have different information or browsing/search needs as follows:

- Navigational - when the intention is to reach a particular known event;
- Contextual Browsing - discover one or more events given a specific context (e.g. by location, performer, type, time);
- Entertainment Browsing - serendipitous and opportunistic discovery of events;

¹⁰ See for example <http://www.jinni.com> or <http://www.ted.com/talks>.

Since it is often easier to recognize a word or name than it is to think up that term, it is useful to prompt users with information related to their needs. Based on that principle, we will explore the use of hierarchical faceted metadata which will allow users to browse through multiple categories, each corresponding to different dimensions of the collection [4]. As a general guideline and given users' request, we will avoid empty results during search. Faceted browsing can avoid empty results by restricting the available filtering options in the given focus to only those which lead to non-empty results (poka-yoke principle) [7]. Consequently, the user is visually guided through an interactive query refinement process, while visualizing the number of results in different categories. Additionally, we will explore the information afforded by linked data to display results which are closely related to the user interests. For example, if during a search, no Jazz concerts are available in Amsterdam, we will show other events from nearby cities, time period or even other type of events closely related to Jazz music.



Fig. 3. Interface views illustrating a set of events under: (A) media centric perspective; (B) a chronological perspective; and (C) location centric perspective

While trying to reach a specific event, traditional keyword search can be done through entry forms. Dynamic term suggestions or auto-completion can be used to provide rapid and effective user feedback by suggesting a list matching terms

as the user types the message. Semantic auto-completion extends this method by providing means of clustering the terms according to different categories or facets [5]. Keyword search can also be integrated to faceted browsing and extend the defined classification options. In this context, it is important to indicate if the search will act as a keyword filter or if it will match the classification terms [7]. In regards to event attributes at initial search constraint definitions, time, place and event type seem to be the core indispensable inputs. A potential solution is to always display these attributes during the whole searching/browsing process to enable zooming-in and out from a search result set at any point. Since time period is a range variable input, a common solution is to use a timeline slider control input [7].

4.3 Event Representation

When representing an event instance, we show all information needed to support the decision making process (e.g. Figure 4). Since experiences are centered around media content, we wish to explore different media that better illustrate the event to end-users. Some information that can support decision making are the following.

- background information (e.g. performers, topic, genre, price, attendance list, etc)
- subjective or computed attributes (e.g. reputation, fun, atmosphere, audience)
- user opinions, comments and ratings (strangers and friends)
- representative media (ads, media from past related events, media from the audience, etc.)

Apart from the inspection of the event instance, other conceptual classes (e.g. users, venues, performers, media) should also have accessible views, so that the user can obtain more information about these instances and explore events related to them. In future work we will also identify what are the relevant associated information and how to represent navigation from and to these nodes.

4.4 Enriching Information

Regarding event content enrichment, interfaces that allow users to add/upload information and assign such information to events will be investigated and explored in future studies.

One of the required enrichment features refers to assigning user attendance and keeping track of the users' previously attended events. This information can be used so that the user can easily access past experiences. Moreover, attended events may be used to identify user interests for recommendation and personalization of the facet-pears during search [7]. In order to keep track of events, we will give options to allow users to say if they were in a past event (e.g. *I was there*) or if they are attending to an upcoming event (e.g. *I will go*). Another

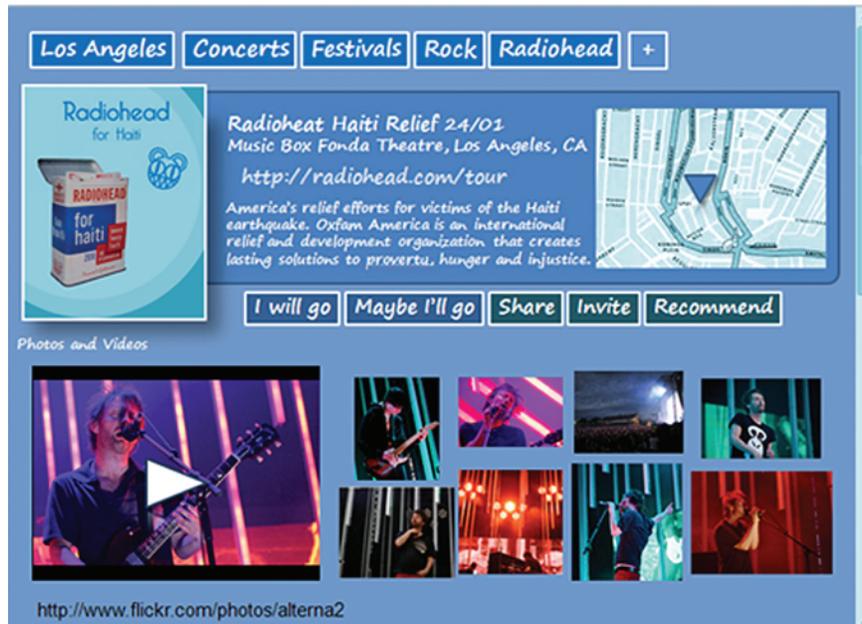


Fig. 4. Interface illustrating an event instance view for a Radiohead concert

prospective option is to allow the user to select events that he is unsure if he will attend (e.g. *I might go*). This will allow adding multiple events to a “maybe” list for future decision or even comparison.

Finally, since users are likely to revisit information they have viewed in the past [4], we will also support simple history mechanisms, by saving a list of recently viewed events. History mechanisms can also be incorporated into the facet search to allow users to undo query filtering and return to a specific query set.

5 Conclusion and Future Work

In this paper, we have described an event-based approach for users to explore, annotate and share media. We first conducted a user study where users were asked about real-world tasks they would like to carry out. We have then extracted requirements and described some scenarios for an event-based system for discovering and sharing media. We advocate the use of linked data technologies for integrating information contained in event directories and we described how event and media ontologies can be used. Finally, we present some sketches of user interfaces that we will develop in the coming months.

In following studies, we intend to use the scenarios we have written to understand how end-users interpret and fulfill associated goals. This will allow us

to identify patterns of interaction, information seeking strategies and information sources required to complete the described tasks. We will continuously increase our coverage of event directories by scraping more data sources and hence demonstrating how interoperability problems can be addressed using semantic web technologies.

6 Acknowledgments

The research leading to this paper was supported by the European Commission under contract FP7-216444, Petamedia Peer-to-peer Tagged Media, and contract FP7-215453, WeKnowIt. The authors would also like to thank Hyowon Lee from Dublin City University (DCU) for fruitful discussions on the design of the EventMedia interfaces.

References

1. A. K. Amin, M. Hildebrand, J. van Ossenbruggen, and L. Hardman. Designing A Thesaurus-Based Comparison Search Interface For Linked Cultural Heritage Sources. In *15th International Conference on Intelligent User Interfaces (IUI'10)*, pages 249–258, Hong Kong, China, 2010.
2. J. M. Carroll. *Making use: Scenario-based design of human-computer interactions*. MIT Press, 2000.
3. Aldo Gangemi and Peter Mika. Understanding the Semantic Web through Descriptions and Situations. In *2nd International Conference on Ontologies, Databases and Applications of SEMantics (ODBASE'03)*, pages 689–706, Catania, Italy, 2003.
4. Marti Hearst. *Search User Interfaces*. Cambridge University Press, 2009.
5. M. Hildebrand, J. van Ossenbruggen, L. Hardman, and G. Jacobs. Supporting Subject Matter Annotation Using Heterogeneous Thesauri, A User Study In Web Data Reuse. *International Journal of Human-Computer Studies*, 67(10):888–903, 2009.
6. Carsten Saathoff and Ansgar Scherp. Unlocking the Semantics of Multimedia Presentations in the Web with the Multimedia Metadata Ontology. In *19th World Wide Web Conference (WWW'10)*, Raleigh, USA, 2010.
7. Giovanni Maria Sacco and Yannis Tzitzikas. *Dynamic Taxonomies and Faceted Search: Theory, Practice, and Experience*, volume 25 of *The Information Retrieval Series*. Springer, 2009.
8. A. Scherp, T. Franz, C. Saathoff, and S. Staab. F—A Model of Events based on the Foundational Ontology DOLCE+ Ultra Light. In *5th International Conference on Knowledge Capture (K-CAP'09)*, Redondo Beach, California, USA, 2009.
9. R. Shaw, R. Troncy, and L. Hardman. LOD: Linking Open Descriptions Of Events. In *4th Asian Semantic Web Conference (ASWC'09)*, 2009.
10. Utz Westermann and Ramesh Jain. Toward a Common Event Model for Multimedia Applications. *IEEE MultiMedia*, 14(1):19–29, 2007.