# Media Aggregation via Events

Fausto Giunchiglia, Pierre Andrews, Gaia Trecarichi, and Ronald Chenu-Abente

University of Trento, Italy {fausto,andrews,gtrecari,chenu}@disi.unitn.it

**Abstract.** Events have been recognised as important metadata to fill the semantic gap between our experience of the world represented in media and its conceptualization. In this paper, we argue that, once event metadata can be extracted, there remains a gap between different users conceptualizations. We then show how a compositional event model can mitigate such a social semantic gap through higher level descriptions of events where an agreement can be reached. In turn, this enables semantic services which improve event-centric search and navigation of shared media.

## 1 Introduction

With the increase of information and media streams available to us, everyday's tasks such as searching and relating relevant data have become intractable. One of the recognised causes for this issue is the semantic gap existing between our *conceptualizations* of the world, usually expressed using language or other high level abstractions, and our *experience of the world*, whose most direct physical representation is kept in media. In fact, even if automatic image processing algorithms can help by extracting high level concepts from media (e.g., who is present in a photo), they still miss the general semantics of the experience memorized by media such as the context in which a photo was taken or what it means for the user (e.g., the feelings and impressions of what happened when a photo was taken). Such experiential aspect cannot yet be extracted automatically and thus a number of solutions and services are being proposed to tame the incoming streams of data. To this purpose, event models have been proposed to help in the extraction and indexing of event information within data streams.

However, once a high level representation of an event is extracted or manually provided in the local context of a single user, there is no guarantee that this conceptualization will be understood by other users. We believe that there is a second gap, the *social semantic gap* between a local user experience conceptualization and what other users might understand of this conceptualization.

In this paper, we propose a general event model that we believe helpful in aligning different local event representations and show how it can be applied to the issue of media management. In Section 2, we provide a motivating example for our event model. In Section 3, we introduce in more details the issue of social semantic gap, while in Sections 4, 5 and 6 we describe in more detail this new model. Section 7 shows how it can be applied to model experiences and media to help the sharing of such media with better metadata. Finally, Section 8 relates our work with state-of-the-art event models.

## 2 Motivating Example

Danda has just returned from a tour in the Italian region of Trentino with her friend[1]. She collected lots of material (e.g., digital photos, diaries, videos) and now wants to organise it digitally to revisit it later on and share her trip memories with her friends. With the current Web 2.0 technologies, she can rely on blogs, video sharing websites (e.g., YouTube) and online photo management software (e.g., Flickr, Virgilio Foto Album), to store and share the material. Since she likes writing, she opts to communicate the experience through her blog and thus dedicates some blog posts to describe the three days spent in Val di Non, the locality visited. In the first blog post titled "My journey diary - 11 Aug 2008", she describes the things that happened during the first day of the trip: the journey from Rimini to Trento by train, the one from Trento to Cles by the local railway Trento-Malé, the nice chats she had with the owner of the B&B during the journey from Cles to the B&B located near Revó, a small village close to the Santa Giustina lake, and so forth. The second day is described in a second blog post providing detailed descriptions of the breakfast she had in the B&B and of the itinerary followed to go to the Tovel lake. Finally, in two other blog posts, she describes the visits made to the Novella river park and the Tret waterfall and the concert of the Ramadas band attended during the third day. In all posts, images illustrate snippets of text to enhance the visual impact of the blog. However, the full gallery of pictures is maintained in a separate online photo management system which is pointed to by a link included in the blog post. An excerpt of the first blog post follows:

> "At Trento we wait for another train, this time on a local railway, the Trento-Malé. Our last stop is Cles from where buses depart for various villages, including Revó; however, the kind owner of the B&B waits for us there and gives us a lift by car..."

The way Danda discloses her trip experience allows her to fix her memory of the journey and to make friends and interested bloggers aware of it. It is important to note that Danda is ultimately interested in describing and sharing the events happened to her rather than just sharing a bunch of photos, these last being rather supporting material to give evidence of her experience and embellish the story. She describes the three-days trip by splitting it into days – or part of days – each corresponding to few significant events, which, in this case, are mostly visits to natural locations. The events range from small-scale ones (e.g., the breakfast, the move from Trento to Cles) to more large-scale, composite ones (e.g., the whole trip, the visit to a natural location, the concert) and span several types

[1] Our running example is based on a real "blog" story found at: `http://dandaworld.blogspot.com/2008/09/appunti-di-viaggio-my-journey-diary-11.html`.

(e.g., a visit, a move from one place to another, breakfast, walk, conversation, concert). Also, all their descriptions mention different entities such as locations (e.g., Trento, Cles, Tovel, Novella river park, Tret waterfall), time periods (e.g., 28th August 2008), people or group of people (e.g., trip companion, B&B owner, Ramadas band), and others (e.g., "Trento-Malé").

If an user-oriented personal media management system existed that allowed for an easy way to describe complex events, Danda would be supported in reliving the "trip event" by recalling salient events at the desired level of details, the persons met, and the places she visited in a more active, experiential manner. Moreover, she would experience pre- or post-trip visits by knowing more about the locations (e.g., facts and media), some co-located events and related stories. The framework we envision supports the well-understood need for event-centric media management systems.

In addition, by having a structured description of the media and event metadata, the events described by Diana can be matched with other users' events and with global definitions of events. We believe, this will reduce the social semantic gap between local descriptions and global understanding of events.

## 3 The Social Semantic Gap

There has been a long stream of research in media processing and media management to fill the semantic gap between what can be seen in a media (the person in a photo for instance) and what it actually represent as an experience for the person that created the media. This can be partly solved by adding event metadata to the media to help understand the context in which they were created. However, in many cases, this event information is created locally, either semi-automatically or manually, at a user level. The particular descriptions of the events at a local level, even if abstracting the content of the media, can still be different from their shared global conceptualization. As shown in Fig. 1, two different users might have a different high level description of a media, and thus there is still a gap between the different personal users' representations of their experiences.

In professional applications, like news media management for instance, there is a top-down agreement on the existing events (e.g., football championship, political conflicts) and thus the gap between the meaning of the events represented in the media and the ones described in the final products (e.g., newspaper articles) composed with these media is kept under control. For instance, everyone shares the understanding of what the "New Year Celebration" is and what experience the media used in the news for that event might represent. However, in the personal application sphere, any of the user's personal experiences can be transformed in an event and there is thus no widespread agreement about events and what they might represent for each single user. A user can create an event about her "Family Holiday" and understand the experiences that are represented in the media for this event, but, outside of her local context, such experiences and events are not meaningful anymore.

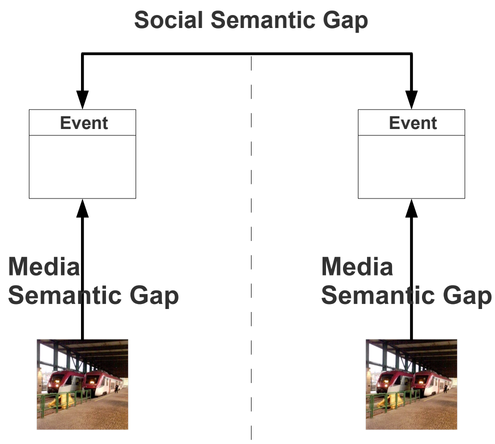


**Fig. 1.** Media and Social Semantic Gap

There can still be a social agreement between users that sit between the Global and Local contexts when the event is shared by a small group of participants. For instance, all the people attending a "Friends' night-out" event will understand the experiences in the media attached to this event, but this understanding will be lost to anyone outside of this group. Thus the events that can be extracted – or provided manually by the user – for the personal user's media are built bottom-up and are difficult to share: there is still a large *social semantic gap* between the local descriptions and their perception by the community.

We believe that the main cause of such a social semantic gap is due to the lack of aggregable metadata for these media. This makes it difficult to align single event descriptions from different users. We propose to solve this issue by modeling the events structural aspect [1] and provide a compositional event model where fine grained events can be composed into higher level events in order to provide a middle ground between top-down event agreement and bottom-up event creation. In fact, if the media are stored in a rooted structure of events, it will be easier to match between diverse local events into a global consensus. For instance, if a set of media is associated to a personal "Toast" event, it is impossible to know if this is a toast during a normal dinner, during a graduation ceremony or at a wedding reception. If this event is provided in the context of subsuming events, as illustrated in Fig. 2, it is then easier to know its semantics.

## 4 Entities

In our example, we can first notice that Danda refers to a number of "objects" in the real world, such as the places she visited or the people she met. There is a need to represent all these entities and their metadata into a model for managing the media relevant to the events she is describing. [2] proposes a unique entity space to store such resources and we follow a similar entity centric model to

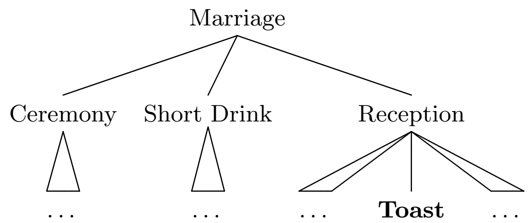


**Fig. 2.** Wedding Event Organisation

provide a uniform representation of objects in the real and virtual world. In our model, an entity $En$ is described by its metadata and associated services:

$$En =< id, type, \boldsymbol{Attr}, \boldsymbol{Rel}, \boldsymbol{S} >$$

Where: 1. *id*, is a unique identifier (e.g., an URI); 2. *type*, is the type of entity, that is, the category to which it belongs to (e.g., the entity "Danda" is of type `Person`); 3. ***Attr***, is a set of *attributes* composed of pairs $attr =< attr\_name, attr\_value >$ describing the properties (e.g., "date of birth") of that particular entity; 4. ***Rel***, is a set of *relational attributes* composed of pairs $rel =< rel\_name, rel\_value >$ describing the entity's relations (e.g., "friendOf") with other entities; 5. ***S***, is a set of *services* that can be leveraged on that specific entity; for example, a service "send email" can be enabled on the *`Person`* entity type (etype).

An important aspect of our model is that both attributes and relationships can be further defined by *meta-attributes*. For example, attributes like "job position" or relations such as "friendOf" are provided with metadata of their own, for example to describe the time period when these are valid or the circumstances (i.e., the events) that made them true.

Another interesting aspect is the *lattice* of etypes that is encoded in the *type* property of the entity. The specific type (e.g., *`Person`*, *`Location`*, *`Event`*) to which the entity belongs to is used to infer its possible attributes and services. Moreover, the hierarchy defined by the lattice allows to easily define new derived etypes by just inheriting the metadata and available services of parent etypes. For example, the new etype *`Author`* inherits both attributes (e.g., "name", "date of birth") and services (e.g., "send email") from the etype *`Person`* but extends them with more specific ones (e.g., "affiliation", "get h-index").

In our running example, we can identify a number of entities in Danda's recollection that can be represented in our model and used for future retrieval and reasoning: Danda participates in a "concert", which can be represented by an *`Event`* entity; "the owner of the B&B" is a *`Person`* entity; and specific *`Location`* entities are described, such as the "Trento" or the "Novella river park".

The following sections describe the less trivial entities used to construct a workable model of the events in Danda's blog post. The data structures discussed hereafter can be modeled as a set of attributes, relations and meta-attributes as formalised in the general Entity model but we provide a higher-level view for clearer reading.

## 5 Events

Events, unlike facts, are closely linked to their spatio-temporal collocation and, also, to the things constituting their subject (e.g., a sparrow in the event "a sparrow falls"). In addition, unlike objects, they have clear temporal boundaries but fuzzy spatial boundaries and they have a time-span [3]. Moreover, they are usually provided by descriptions and they may be composed of sub-events which are temporally, spatially and causally connected [4].

In our model, we assume that "local" events are created by users and their structuring and descriptions are thus subjective. An event entity $Ev$ is modeled as follows:

$$Ev =< evid, t, \boldsymbol{LEv}, Cx > \quad (1)$$

Where the elements of the tuple are:

- $evid$, the unique identifier of an event;
- $t$, the temporal collocation of an event, i.e., the time interval described by the event; it can be either a specific time interval marked off by the initial and final instants – e.g., "2009:01:14:10:00" to represent the 14th January 2009 at 10am – or a generic period of time where the temporal delimitations are not specified; also, the information on the date can be incomplete (e.g, "2009:01:::" ), or relative (e.g., "the day before Christmas"), and the time interval may not be continuous as, for instance, it would happen for a "Champions League" event.
- $\boldsymbol{LEv}$ represents a set of linked events and is described in more details in the following section.
- $Cx$, the *event context*, this being regarded as a distinguishing feature of event entities. As demonstrated in [5, 6], this context is useful for localised reasoning. The event context is represented as:

  $$Cx =< l, type, \boldsymbol{Pc} > \quad (2)$$

  Where:
  - $l$, defines the spatial collocation of an event. It identifies a "geographical entity" such as a geopolitical entity, a natural body (e.g., mountain, river, lake), or a man-made infrastructure (e.g., building, stretch of road). For example, consider our running example where Danda moved around the north of Italy, the location of this "transfer" event can be modeled by defining two geographical (point) locations (Rimini and Trento) that form the stretch of train track between Rimini and Trento.
    Note that, although the spatial collocation could be objectively defined, the participant's perception of it, that is represented in the context, is itself subjective, e.g, in terms of the actual extension of the location itself ("Trento" vs. "Trentino Region").

- *type*, is the type of event (e.g., conference, trip, visit, concert);
- ***Pc***, the event's participants is a set of relations to other entities. The corresponding entity values could belong to ***Person***, ***Organization*** as well as to non-agentive etypes. Each relation to these participating entities is annotated with *meta-attributes* describing aspects of each entity which are only relevant in the current event's context. This includes the role of the entity in such an event, i.e., the modality of its participation in the event: for instance, a person can be a professor in a graduation event but is then a mother in the event describing the birthday of her daughter. Our vision of role is in line with the one given in [7], where one of the key features of a role is that of being linked to the notion of context – the event's context in our model.
  In addition, an entity participating to a given event could be described by properties valid only within the event context: for example, a temporal attribute such as "jacketColour" is only attached to the relation between a particular event and the entity. Note that, since we regard events as subjectively perceived entities, the above mentioned properties are meant to be objective (e.g., "jacket colour") as well as subjective (e.g., "personality"). As for the attributes, relations defining the metadata of an event's participants can also be relevant only to the event's context; for instance, relations such as "girlfriendOf" or "near".

## 6 Event Compositionality

Danda's trip has several events that can be identified and captured through the previously given event definition. For example, the transfer from Rimini to Trento can be represented as:

$$\begin{aligned} Transfer(\text{Rimini}, \text{Trento}) =& < evid01, \text{"11/08/2008 on the early morning"}, \\ & < \text{"Railroad between Rimini and Trento"}, \text{"train transfer"}, \{\text{Danda}, \\ & \text{Danda's friend}\} >, \emptyset > \end{aligned}$$

while the transfer from Trento to Revó can be represented as:

$$\begin{aligned} Transfer(\text{Trento}, \text{Revó}) =& < evid02, \text{"11/08/2008 on the late morning"}, \\ & < \text{"Railroad between Trento and Revó"}, \text{"train transfer"}, \{\text{Danda}, \\ & \text{Danda's friend}\} >, \emptyset > \end{aligned}$$

These two events can be aggregated to define a more general *Departure* event representing the journey that would include both *Transfer(Rimini,Trento)* and *Transfer(Trento,Revó)*:

$$\begin{aligned} Departure =& < evid03, \text{"11/08/2008 morning"}, \\ & < \text{"Railroad between Rimini and Revó"}, \text{"train transfer"}, \{\text{Danda}, \\ & \text{Danda's friend}\} >, \{evid01, evid02\} > \end{aligned}$$

The $\boldsymbol{LEv}$ from the event *Departure* contains references to the events it is aggregating. The running example, as a whole, can be represented as a single complex entity *Trip* as illustrated in Fig. 3.

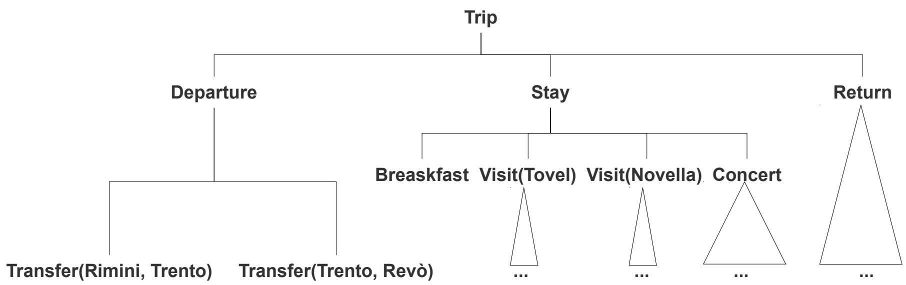


**Fig. 3.** Event Structure for the Danda's trip

Note that, in Fig. 3, the event that represents the whole *Trip*, is subdivided in *Departure*, *Stay* and *Return* (which refer to the going journey, the stay period and the return journey respectively). Furthermore, each of these sub-events is, in turn, subdivided in other sub-events. This tree-like compositional structure is enabled by the use of the $\boldsymbol{LEv}$ component from the event definition in (1).

As explained before, $\boldsymbol{LEv}$ represents a set of linked events that are parts of the event to which that $\boldsymbol{LEv}$ belongs. However, to keep $\boldsymbol{LEv}$ as useful as possible for its complex event modeling purposes, the following restrictions are applied to it:

**Restrictions on time** $t$, the time duration defined for a complex entity $CEv$, must subsume the time duration of all of its sub-events pointed by $\boldsymbol{LEv}$. That is, if we have a function $\texttt{time}(Ev) \rightarrow t$:

$$\text{Given } CEv =< evid, t, Cx, \boldsymbol{LEv} >, \forall e \in \boldsymbol{LEv}, \texttt{time}(e) \sqsubseteq t$$

By enforcing the previous restriction, all the individual time periods involved in the children events, are guaranteed to be subsumed in the time period of the parent complex event. This enables the representation of a complex event, like the one from our running example, in a timeline as shown in Fig. 4.
Note how in Fig. 4 all events comply with this rule. For example, the event *Stay* spans from the first to the third day and is within this period that its children events (*Breakfast*, *Visit(Tovel)*, *Visit(Novella)* and *Concert*) take place. Furthermore, note that the whole time period of *Stay* is not entirely covered by sub-events (e.g., a small period of time exists between the end of the event *Visit(Novella)* and the start of the event *Concert*). These blanks correspond to unspecified events in Danda's trip such as, for example, the transfer between visits where she had no memorable experiences.

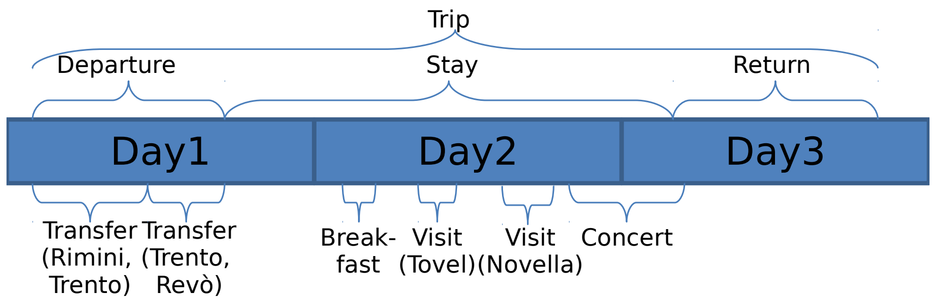


**Fig. 4.** Running example as a timeline.

**Restriction on context** the context $Cx$ defined for the complex entity $CEv$ must subsume the context metadata of all of its sub-events pointed by $\boldsymbol{LEv}$. That is, given (1) and (2), if we have the functions:

$$\texttt{location}(Cs) \rightarrow l$$
$$\texttt{location}(Ev) \rightarrow \texttt{location}(Cx)$$
$$\texttt{participants}(Cx) \rightarrow \boldsymbol{Pc}$$
$$\texttt{participants}(Ev) \rightarrow \texttt{participants}(Cx)$$

Given $CEv =< evid, t, \boldsymbol{LEv} >, Cx$

$$\forall e \in \boldsymbol{LEv}, \texttt{location}(e) \sqsubseteq \texttt{location}(Cx)$$
$$\forall e \in \boldsymbol{LEv}, \texttt{participants}(e) \subset \texttt{participants}(Cx)$$

By enforcing the previous restriction, all the geographic locations from the children events are guaranteed to be subsumed by the parent's location and all the participants from the sub-events are guaranteed to be included in their parent's set of participants.

From our previous examples, it is clear that the location for the event *Transfer(Rimini,Trento)* is the railroad between the cities Rimini and Trento; furthermore, the location for the event *Transfer(Trento,Revó)* is the railroad between the cities Trento and Revó. Applying the location subsumption restriction, the location for the parent event *Departure* would be the train road between Rimini and Revó or more generally the entity for North of Italy (both of which subsume the locations from the sub-events). Likewise, the participants from *Breakfast* and *Visit(Tovel)* would also be included in the set of participants of its parent event *Stay*.

**Restriction on linked events** an event cannot be included in $\boldsymbol{LEv}$ if doing so would cause the creation of a loop in the events structure. Let $\texttt{connection}(Start, End)$ be a function that returns a sequence of events that, through their $\boldsymbol{LEv}$ components, define a directed path from the $Start$ event to the $End$ event (or $\emptyset$ if no such sequence exists). Then, the restriction can be expressed as:

$$CEv = < evid, t, \boldsymbol{LEv}, Cx >,\ \texttt{connection}(CEv, CEv) = \emptyset$$

This restriction is introduced to avoid the conceptual problems that would arise from an event being its own predecessor, directly or through other intermediate events.

The composition of events presented in this section has the following advantages:

- *Avoid repetition of information*: if there is a particular information that applies to all of the children of a complex event, instead of repeating the content on each of the sub-events, this information can be included directly at the event that is aggregating them. For example, thanks to this, it is not necessary to describe the weather in both the *Transfer(Rimini,Trento)* and the *Transfer(Trento,Revó)* events. If the weather did not change between these two events, the details of the weather can be included in their aggregating event *Departure* and, through compositionality, this information will apply to all of its children.
- *Capture information emerging from the aggregation*: there may exist information that emerges from the composition itself and is not part of any of its individual sub-components. For example, Danda could describe the *Concert* event as being "long and tiring" but each of its sub-events may not have these properties individually. The "long and tiring" description would then only apply to the aggregation of these individual events into the *Concert* event.
- *Capture information from unspecified sub-events*: as seen in Fig. 4, there may exist some blanks between events at high granularity levels. A lower granularity or parent event can then be used to capture information belonging to these blanks: for example, suppose Danda wanted to add a photo she took right after having her breakfast on the 2nd Day (*Breakfast* event) but before her visit to Tovel (*Visit(Tovel)* event); instead of adding a new event only for that photo, she could just include it in the *Stay* event.

In the following section we show an example of how this structured metadata of the events and the media that can be attached to them can be used to fill the semantic gap.

## 7 Event-Centric Media Management

In the previous sections, we presented a general event model where an event is independent from media. In fact, an event can exist totally independently from media in many applications. However, in this paper we are interested in how such an event model can be used to move from a media-centric management to a different metaphor where events are of importance for organizing media, which is the goal of a number of projects such as GLOCAL[2], PRONTO[3] and EventMedia[4]. This is supported by [1] that prompts for a common event model for media management. Our approach is to separate clearly the event metadata and the experiential aspect of this event, which depends on the user describing the event and the intended audience. Thus, we introduce a new entity in our model to store a particular description of an experience as the relation between an event and the media describing it:

$$Exp =< Ev, \boldsymbol{M} >$$

where $Ev$ is a relation to the event in question and $\boldsymbol{M}$ is a set of relations to media describing the user's experience of this event.

In this context, the creation and structuring into complex events proves useful when a "story" about that event's experience has to be told. For example, in our example, Danda goes into the higher granularity events for the part of the story that her audience will favor. Conversely, she also chose to stay at a lower level of details when describing events that might be of less importance for that particular audience. However, for her own use when searching for media or when telling stories to other audiences about her experience, she is still interested in keeping as most details as possible in the events structure and metadata.

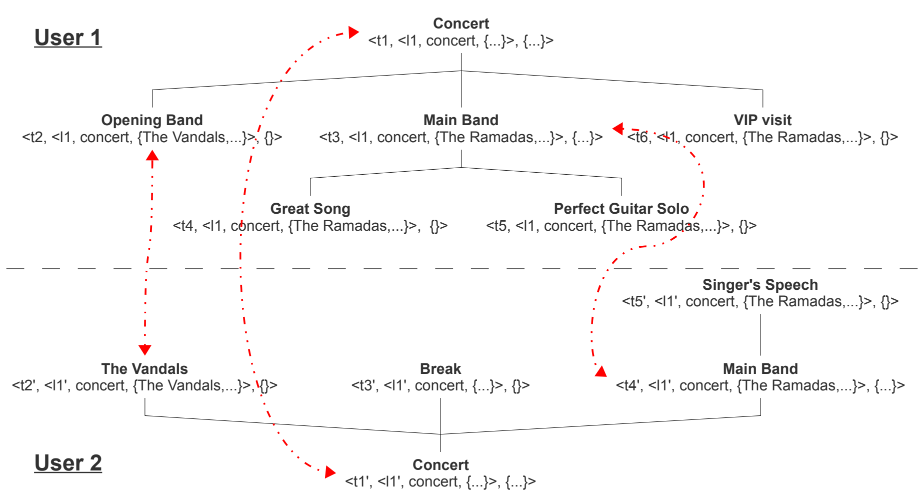


**Fig. 5.** Two Local Structures for the Same Concert Event

A second application of this granular description of events in media management is the support provided when sharing an experience with a community

[2] `http://www.glocal-project.eu/`
[3] `http://www.ict-pronto.org/`
[4] `http://eventmedia.cwi.nl/`

(i.e., going from a local context to a shared or global context as discussed in Section 3). If we consider the "concert" event in our running example, more than one user will have attended the concert and annotated media of this concert with subjective event information; we could thus have the two event sub-structures shown in Fig. 5. The two users have a different experience of the concert and have represented it as a set of different sub-events. If only the lower level of granularity was available, it would be difficult to say that the media describing the "Singer's Speech" and the "Great Song" are related. However, by providing information about the "concert" higher level event, general metadata is made available to help match between users local events as the subjective description of the time ($t1$ and $t1'$) and location ($l1$ and $l1'$) are more likely to overlap. The media of both users can be related by matching as much events as possible in a top-down – general to specific – fashion to reduce the social semantic gap. In particular, the sub-structure of events is slightly different, but by comparing the metadata of each event, the events "Opening Band" and "The Vandals", as well as "Main Band" and "The Ramadas", can be matched for instance with an ontology matching algorithm [8]. In addition, once the matching is done, the management system can propose to the *User 1* to add the "Break" sub-event to her own event structure and thus organise her local media in a more refined manner. After that, the media can be reclassified automatically within the sub-event, for instance with the get-specific algorithm [9]. Thus, the users will be able to share, search and navigate new media of the events they have experienced more easily.

## 8 Related Work

Research fields spanning from Semantic Web [10, 2] to Information Extraction [11, 12] to Digital Libraries [13] have recognised the key role played by the entities and their linking relationships.

In [10], entities are seen as atomic objects of reference and reasoning for Semantic Web applications which are described by a general conceptual model and categorized into types derived from a user-study. This work is part of the OKKAM project[5], which proposes a framework where entities are assigned with a globally unique identifier to ease data integration and the development of innovative "entity-centric" applications [2]. Categories for entities are also defined in [11, 12], together with guidelines for the accomplishment of named entity recognition and relation extraction tasks. The issue of linking together different kinds of entities in the Digital Library domain is addressed by Buckland [13], who advocates the need for a metadata infrastructure able to interconnect place-name gazetteers, biographical dictionaries, time period and subject indexes.

Furthermore, the concept of the "Web of Things[6]" (see [14]) has recently emerged and efforts such as the Linking Open Data Initiative[7] entered the scene

[5] `http://fp7.okkam.org`
[6] `http://ercim-news.ercim.eu/content/view/343/536/`
[7] `http://linkeddata.org/`

to address its principles. The role played by linked data for supporting users in tasks beyond the simple fact-finding and question answering (e.g., finding connections among people, places and events) is analyzed by looking at specific prototypes in [15].

In regards to the conceptual representation of events, research work exists on generic models as well as models tailored to domain-specific events (e.g., journalistic, historical, cultural-heritage, multimedia events): the generic E-model [16] is extended to enable "event-centric" rather than "media-centric" media management systems [17]; on the same line, the Eventory media repository [18], the Medi*AE*ther system [19] and a multimodal event browsing tool [20] are proposed. The F Event Model [21] proposes a formal model that, together with standard properties, supports mereological, causal, correlation relationships and interpretations of arbitrary events. The Event Ontology[8], developed as part of a music ontology framework, supports music events (e.g., compositions, recordings, performances) but is not tied to such domain. With a journalistic perspective, EventML[9] provides an XML schema for exchanging news events among news agencies. The CIDOC ontology [22] and the CultureSampo approach (see [23]) aims to solve interoperability issues between metadata standards for heritage resources. In [24], most of the above-mentioned models are compared and an event ontology is proposed to fulfill the vision of a Linked Data Event Model[10]. The model is purposely kept minimal to capture the well-understood "when", "where", "what" and "who" facets of historical non-composite events; aspects defined as more interpretative dimensions are excluded. In this respect, our model differs from others since user-driven contextual metadata make events as always being subjective entities. Furthermore, by means of the *experience* entity we define relationships between events and media. This is also allowed in Eventory [18] where such relationships are made explicit through ad-hoc interfaces for events, media and connection creation; however, we foresee a system where these types of relationships are deduced from the way the user describes her experience rather than being explicitly created.

For what specifically concerns the event's structural relationships – i.e., their compositionality – Rafatirad et al. [25] design "subevent-of" relationships forming directed acyclic graphs over events and implement composition operators allowing to aggregate the spatial and temporal attributes of composite events from atomic events for which these attributes are known; Singh et al. [26] regard the information needed to model compound events as semi-structured (dynamic) data and hence use XML to manage them: for example, a "group meeting" event is divided into compound events (e.g., "Introduction", "Presentation", "Discussion") and simple events (e.g., "ask_question", "answer_question"). The XML schema used to describe a compound event includes an "how" element which model its process, i.e., the simple events composing it. A hierarchical model to represent events at different granularities has been recently proposed in the domain of

[8] http://motools.sf.net/event

[9] http://iptc.org/

[10] http://linkedevents.org/ontology/

multimedia observation systems [27]; here, "transient" events are detected from sensors datastreams and clustered into "atomic" events, these last being in turn composed into "compound" event. The model is however application-specific and aimed at reducing the semantic gap rather than what we defined as social semantic gap. On the same line, van Hage et al. [28] present the Simple Event Model (SEM), applied in a Maritime Safety and Security use case to deduce simple behaviour events from sensor data. The model provides a minimal set of classes describing all the event's aspects, and add the notion of roles and the possibility to associate types to these classes thus to maintain the compatibility with external resources.

## 9 Conclusion

In this paper we present a general Event model to store the key metadata of an event. In particular, our model allows for the storage of complex subjective information relevant to the event's context.

We then show how such general model can be applied to the media management issue by introducing the *experience* entity that links an event with the media representing the user's experience of such an event. We believe that by clearly separating the event model from the experiential metadata of the event, this metadata is easier to use in heterogeneous applications.

In addition, by proposing a compositional model to represent the structural aspect of events, we allow for an easier alignment between users' personal descriptions of events and thus bridge the "social semantic gap" between different local representations of shared experiences.

## Acknowledgement

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n248984 GLOCAL: Event-based Retrieval of Networked Media.

## References

1. Westermann, U., Jain, R.: Toward a common event model for multimedia applications. IEEE MultiMedia **14** (2007) 19–29
2. Bouquet, P., Stoermer, H., Niederee, C., Mana, A.: Entity name system: The backbone of an open and scalable web of data. In: In Proc. of IEEE-ICSC. (2008) 554–561
3. Casati, R., Varzi, A.: Events. Stanford Encyclopedia of Philosophy. (2006)
4. Scheffler, U.: Events as shadowy entities. Konstanzer Berichte: Logik und Wissenschaftstheorie **28** (1992)
5. Giunchiglia, F.: Contextual reasoning. Epistemologia **16**(I Linguaggi e le Macchine) (1993) 345–364

6. Ghidini, C., Giunchiglia, F.: Local Models Semantics, or Contextual Reasoning = Locality + Compatibility. Artificial Intelligence **127** (2001) 221–259
7. Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Ferrario, R., Gangemi, A., Guarino, N.: Social roles and their descriptions. In: Proc. of Knowledge Representation Workshop, AAAI Press (2004) 267–277
8. Giunchiglia, F., Yatskevich, M., Shvaiko, P.: Semantic matching: algorithms and implementation. In: JoDS, IX. (2007)
9. Giunchiglia, F., Zaihrayeu, I., Kharkevich, U.: Formalizing the get-specific document classification algorithm. In: ECDL. (2007) 26–37
10. Bazzanella, B., Chaudhry, J., Palpanas, T.: Towards a general entity representation model. SWAP (2008)
11. : Automatic content extraction english annotation guidelines for entities, version 6.6 2008.06.13. Technical report, Linguistic Data Consortium (2008)
12. : Automatic content extraction english annotation guidelines for relations, version 6.2 - 2008.04.28. Technical report, Linguistic Data Consortium (2008)
13. Buckland, M.K.: Description and search: metadata as infrastructure. Brazilian Journal of Information Science (2006)
14. Bizer, C., Heath, T., Berners-Lee, T.: Linked data - the story so far. IJSWIS, Special Issue on Linked Data (2009)
15. Hardman, L., Ossenbruggen, J.v., Troncy, R., Amin, A., Hildebrand, H.: Interactive information access on the web of data. In: WebSci'09: Society On-Line. (2009)
16. Westermann, U., Jain, R.: E - a generic event model for event-centric multimedia data management in echronicle applications. In: ICDEW, IEEE Computer Society (2006) 106
17. Scherp, A., Agaram, S., Jain, R.: Event-centric media management. In: SPIE. (2008)
18. Wang, X., Mamadgi, S., Thekdi, A., Kelliher, A., Sundaram, H.: Eventory - an event based media repository. In Semantic Computing. IEEE (2007)
19. Boll, S., Westermann, U.: Medither: an event space for context-aware multimedia experiences. In: ACM SIGMM workshop on Experiential telepresence. (2003) 21–30
20. Appan, P., Sundaram, H.: Networked multimedia event exploration. In: ACM Multimedia. (2004) 40–47
21. Scherp, A., Franz, T., Saathoff, C., Staab, S.: F–a model of events based on the foundational ontology dolce+dns ultralight. In: K-CAP'09. (2009) 137–144
22. Doerr, M., Ore, C., Stead, S.: The cidoc conceptual reference model - a new standard for knowledge sharing. In Conceptual modeling (2007) 51–56
23. Ruotsalo, T., Hyvnen, E.: An event-based approach for semantic metadata interoperability. In: ISWC. (2007) 407–420
24. Shaw, R., Troncy, R., Hardman, L.: Lode: Linking open descriptions of events. In: ASWC. (2009)
25. Rafatirad, S., Gupta, A., Jain, R.: Event composition operators: Eco. In: EiMM'09, ACM (2009) 65–72
26. Singh, R., Li, Z., Kim, P., Pack, D., Jain, R.: Event-based modeling and processing of digital media. In: CVDB'04. (2004)
27. Atrey, P.K.: A hierarchical model for representation of events in multimedia observation systems. In: EiMM'09, ACM (2009) 57–64
28. van Hage, W.R., Malaisé, V., de Vries, G., Schreiber, G., van Someren, M.: Combining ship trajectories and semantics with the simple event model (sem). In: EiMM'09. (2009)