New Forms of Interaction With Hierarchically Structured Events

Sven Buschbeck, Anthony Jameson, and Tanja Schneeberger*
Sven.Buschbeck@dfki.de, jameson@dfki.de, tanjaschneeberger89@googlemail.com

German Research Center for Artificial Intelligence (DFKI), Saarbrücken, Germany

Abstract. Research in the semantic web community has given rise to some powerful methods for visualizing events and related media—mostly in terms of interactive timelines—and for enabling users to interact with these visualizations. The present paper aims to advance this line of research in two ways: (a) by developing interactive visualizations of event hierarchies of essentially arbitrary depth and size, which are more natural than timelines in the case of complex events that comprise subevents at various levels; and (b) by supporting forms of interaction that go beyond the usual activities of browsing and searching for events and related media, supporting additionally the sharing and annotation of media and the provision of interactive illustrations of narrative texts.

1 Introduction

1.1 Issues and Goals

This paper addresses the third of the three questions that were formulated in the call for papers for this workshop: "How can events be exploited for the provision of new or improved services?"

One focus of the paper is on new ways of visualizing events and associated media and user commentary, in particular in the case of a complex event (e.g., a soccer tournament) which consists of several levels of subevents (e.g., individual games in the tournament and events within the games). But we also show how the resulting interaction design makes possible novel forms of interaction with events and associated content.

Though the ideas presented have considerable generality, they are introduced here with reference to an implemented prototype that will be demonstrated interactively at the workshop. This prototype is being developed in the context of the European Integrating Project GLOCAL.[1] It is designed to enable professional and nonprofessional users to browse and search for media that are organized in terms of events; to comment

* The research described in this paper is being conducted in the context of the 7th Framework EU Integrating Project *GLOCAL: Event-based Retrieval of Networked Media* (http://www.glocal-project.eu/) under grant agreement 248984. We thank the anonymous reviewers for their perceptive suggestions, which have led to improvements in the final version of the paper.

[1] http://www.glocal-project.eu/

on such media and on the associated events; and to upload their own media, introducing them into the event-based organization and thereby in effect indexing them on-the-fly.

To understand the application context, consider a website that presents a large amount of media and information about a complex sports event, such as the soccer World Cup of 2010. Assume that a great many events and associated media have been detected and represented using methods such as those discussed in other papers for this workshop: How can users be enabled to explore these media and contribute media of their own?

One now-familiar way of organizing the media would involve associating each event in the tournament—and hence its associated media—with a given location and showing these locations on an interactive map (see Figure 1). This method can be seen with this example not to supply the most relevant context: It is usually less important to know where the game between Nigeria and the Korean Republic occurred in South Africa than to know where it occurred in the hierarchical structure of the whole tournament. Similarly, showing the game on a timeline would only vaguely and indirectly suggest its significance within the tournament. Consequently, although the GLOCAL interface supports visualizations in terms of geographical maps and timelines, we focus in this paper on the more innovative visualization in terms of subevent hierarchies.

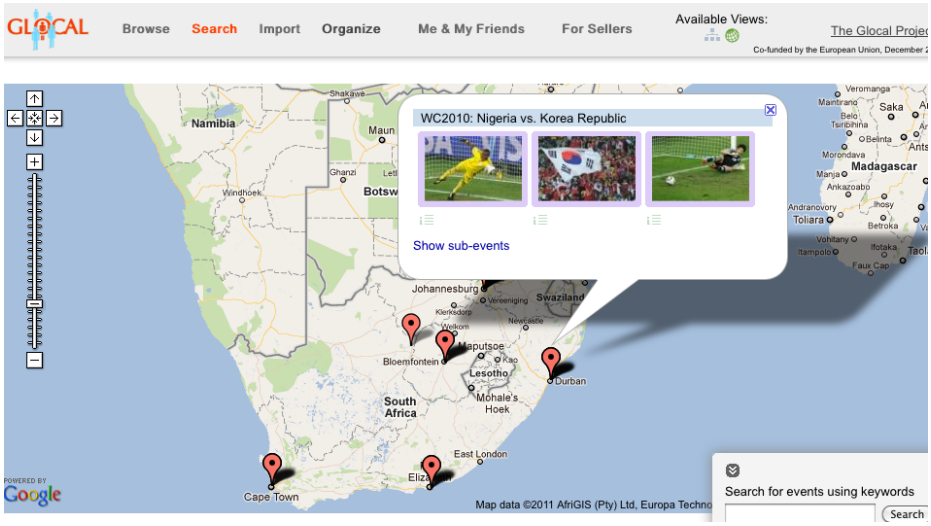

Fig. 1. Example of how events and associated media can be organized relatively conventionally in terms of a geographical map.

1.2 Background of Related Work

There is a fairly extensive tradition in the semantic web community of visualizing events in terms of timelines. Since the work up to 2007 has already been ably summa-

rized by André et al. ([1]), for reasons of space we will discuss just a few highlights, including the work of those authors.

Perhaps the most widely used tool of this sort is TIMELINE,[2] a web widget for visualizing temporal data. It was developed by David F. Hyunh as part of the SIMILE project. In a typical use of the widget, a number of related events (such as those involving the assassination of President Kennedy) are displayed on an interactive timeline (which may comprise a number of parallel lines in the vertical dimension to make it possible to represent events with high density). By clicking on an event in the timeline, the user can access a textual description, which includes further links, for example to a discussion page for that event.

TIMEMAP[3] builds on TIMELINE by integrating it with online mapping systems such as GOOGLE MAPS. Basically, a TIMEMAP display shows simultaneously (a) a timeline with events or intervals and (b) a map with corresponding locations. Clicking on an item in either the timeline or the map brings up additional information about the event in question. It is possible to include filtering functionality so as to show only the events that fulfill particular criteria.

The system CONTINUUM ([1]) likewise builds upon TIMELINE, advancing it in several ways. The way of interest for the present paper is CONTINUUM's ability to represent explicitly hierarchical relationships among events. For example, the lifetimes of classical composers can be shown as belonging to various eras, and each musical composition can be shown as a subevent in the life of its composer. Though this method may cover many types of event hierarchy elegantly, it cannot apparently deal with hierarchies of arbitrary depth and content. In the above example, the eras are represented as segments of the overall timeline; each composer's life is visualized in a box; and his or her compositions are listed in the box. It is not clear how further levels of hierarchy could be added to this visualization.

Outside of the semantic web area, some researchers developing new methods for media organization have explored the use of event hierarchies. For example, REMINISCING VIEW ([2]) enables users to organize photos in a way that is based on an underlying event hierarchy. But REMINISCING VIEW does not, as one might expect, present the media in a hierarchical layout; in fact, it does not visualize the events at all.

As we will see, an attempt to visualize event hierarchies (and the associated media, metadata, and user comments) explicitly requires a visual tree structure with various types of links. The solution presented below (in particular the use of *aggregator nodes*) was inspired in part by the work of Hirsch et al. ([3]) on visualization of large knowledge spaces.

Other aspects of the interface—such as expanding and collapsing subtrees, filtering, and focusing on subtrees— were inspired by functionality offered by mind-mapping tools (e.g., Mindjet's MINDMANAGER, which was used for a first semifunctional prototype). Although these systems have found widespread use on desktop computers, this functionality is seldom found in web-based systems (aside from web-based mind mapping tools such as MINDMEISTER[4].

[2] http://www.simile-widgets.org/timeline/
[3] http://code.google.com/p/timemap/
[4] http://www.mindmeister.com

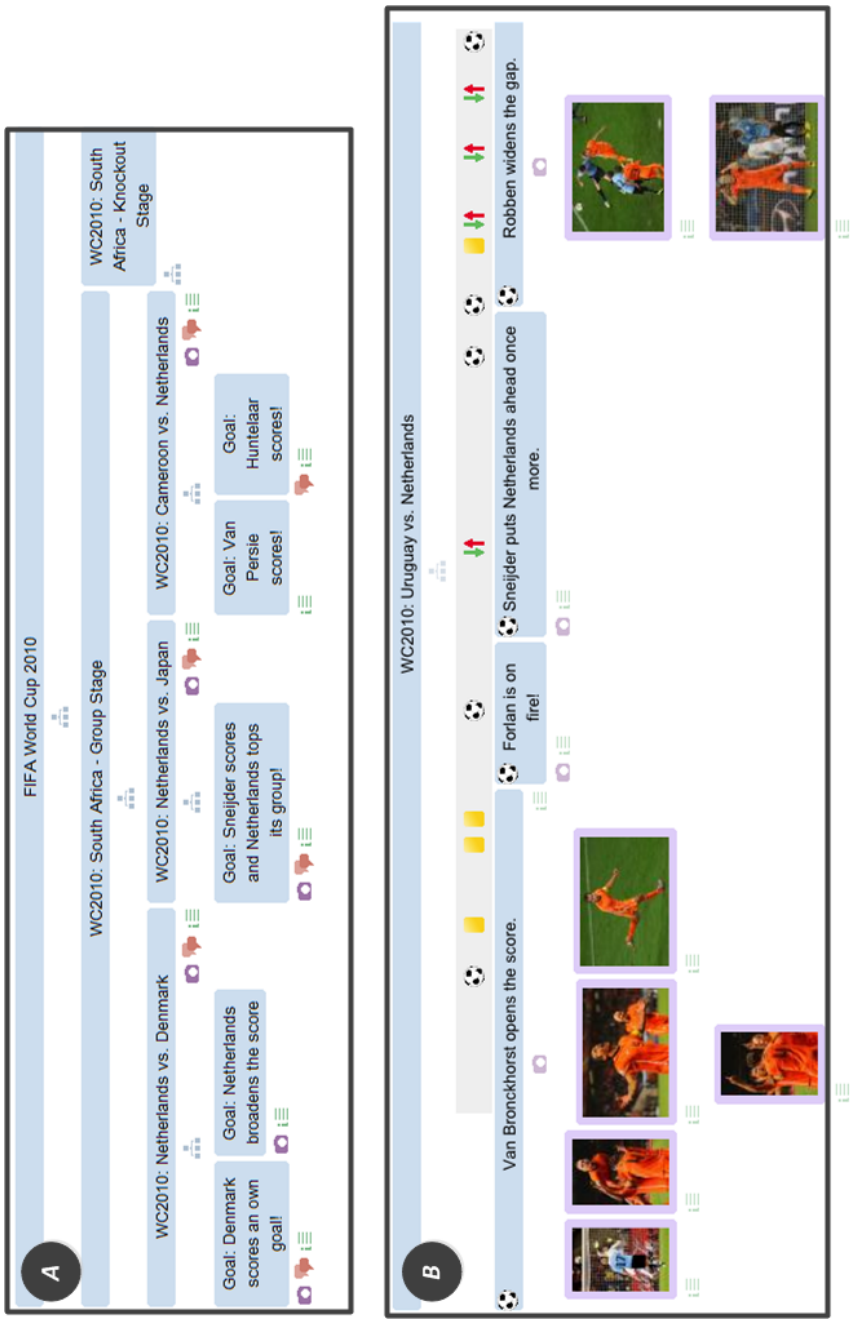

Fig. 2. A: A filtered and partly collapsed representation of the 2010 soccer World Cup as a hierarchy of events (explanation in text). B: The user has zoomed in on a single game and clicked on the "media" links for two goals, so as to be able to compare the associated media.

2 The GLOCAL User Interface

2.1 Hierarchical Browsing of Events

We will now introduce the GLOCAL interface, a web-based interface implemented with the GOOGLE WEB TOOLKIT[5]. It makes use of REST services to access media and event structures.

Figure 2A illustrates several characteristics of the interface's visualizations and functionality. First, the hierarchical structure of the football tournament can be seen. The subevent link (⛶) connects an event with its subevents: By clicking on this link, the user can cause the subevents to toggle back and forth between being hidden and being displayed. In this figure, the entire "Knockout Stage" on the right has been collapsed to a single node, since the user wants to focus on the Group Stage.

Even the Group Stage contains much too many events to display at once. But the interface offers a filtering functionality, like that found in some mind-mapping tools, which enables the user to specify that only events that fulfill certain criteria are to be displayed—along with their superevents. In this figure, the user has chosen to focus on goals scored by the Netherlands team.

If the user wants to focus on one (complex) event in the hierarchy, he or she can click on an icon within the event node[6] to cause it to become the root node of the hierarchy. For example, in Figure 2B the game between Uruguay and the Netherlands has become the root node. It is now feasible to display subevents at a finer-grained level, where applicable using domain-specific symbols that correspond to the types of the events in question. Also, in the other direction, the user can step upwards in the hierarchy to include the parent node in the current view—and hence also the sibling nodes and their descendants, insofar as they match the current filters.

2.2 Organization of Media

Figure 2B also illustrates how, by clicking on the "media" link (📷), a user can cause thumbnails of the media associated with a given event to be displayed (or a selection of these thumbnails, if the number is large).[7] Figure 2B illustrates how this organization makes it possible, for example, to compare media associated with two of the events even if there are a number of events between them.

[5] http://code.google.com/p/google-web-toolkit/

[6] The icon that the user clicks on, along with several analogous icons, becomes visible only when the cursor hovers over the node. In this way, rich functionality can be offered with a minimum of clutter in the interface.

[7] As is usual with thumbnails, it is possible to click on a thumbnail to have a larger version of the media item displayed elsewhere on the screen. Currently the media are photos and videos, but the GLOCAL project is also working with other types of media item, such as automatically generated transcripts of audio files.

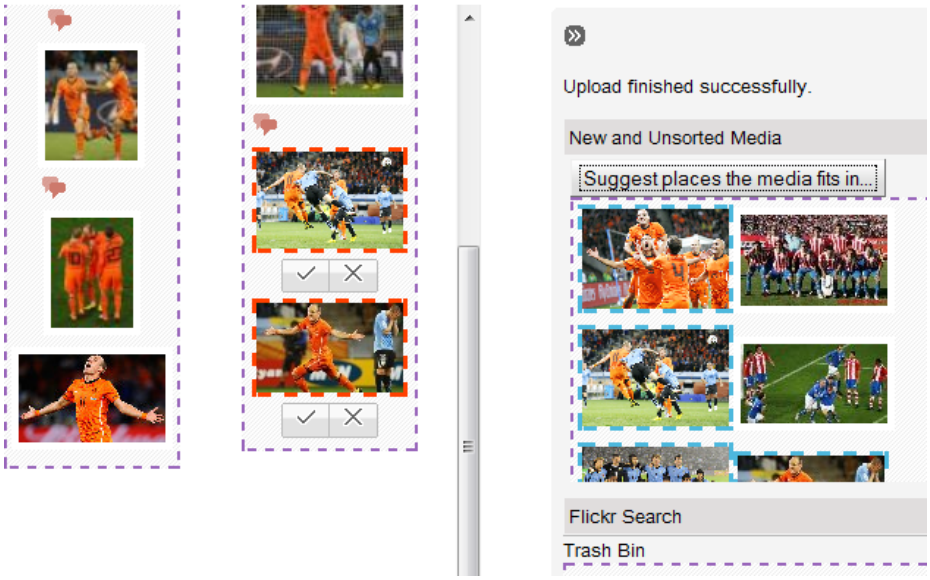

Fig. 3. The user has uploaded the media shown on the right and can now drag them into the appropriate boxes on the left to associate them with events.

(The media on the left that have checkboxes are ones that the system has tentatively aligned in response to the user's clicking on the suggestion button on the right.)

2.3 Importing and Aligning Media

Consider a user who has taken some photos and videos of the game between the Netherlands and Uruguay and wants to introduce them into the platform. Figure 3 shows how the user can upload media into a sort of inbox on the right-hand side of the screen. The system then displays a dotted box for each of the existing events, so that the user can drag a thumbnail into a box to indicate that it belongs to the event in question.

Since the system will often be able to make a good guess about the event that a given media item depicts, the interface also allows the user to ask the system to suggest an alignment of media to events (by clicking on the link "Suggest places ..."). [8] In the cases where the system has a recommendation, the system places a thumbnail tentatively in the dotted box for one of the events on the left. The user can then accept or reject the system's suggestion by clicking on one of two icons associated with the thumbnail.

Instead of a uploading media from a local computer, the user can import media from another site such as FLICKR by formulating a search query which is passed to FLICKR's

[8] Partners in the GLOCAL project are exploring various methods for suggesting alignments of media items to events, including image analysis and the use of spatial and temporal metadata. Whatever method is used, its accuracy is likely to be imperfect, making it natural to put users in the loop in some way.

API. The media retrieved in this way are placed in the user's inbox, where they can be subjected to the treatment just described.

Media contributed in this way become available for sharing with other users; they also provide the system with more information about the events and media that it already has.

2.4 Narrative Plus Visualization

Figure 4 illustrates a novel use of event representations that was suggested by our collaboration with the AGORA project.[9] Instead of merely commenting on individual events or media, a (professional or amateur) user can create a textual *narrative* and then provide illustrations of it by supplying links to relevant *views* of the event representation. Each view shows a subset of the events and the associated media and metadata which the author of the narrative can (a) specify interactively by applying filters and clicking on aggregator nodes and then (b) save with a bookmark (much as users of GOOGLE MAPS can save a view of a map with a bookmark, which can be embedded in a web page or emailed to another user). The reader of the narrative can interact with each view in the same way, in particular exposing information that was not visible in the view as specified by the author. In the example in Figure 4, a Netherlands fan might choose to have the cautions incurred by the Spanish team displayed.

This style of interaction is reminiscent of the increasingly popular trend of *media curation*, which is supported by platforms like STORIFY[10] and OURSTORY,[11] which supports the creation of timelines. A difference is that each view consists not of arbitrary content (e.g., photos or TWITTER feeds) that has been acquired from somewhere in the internet but rather of a specified view of a very large content repository. An important consequence of this difference is that the reader of a narrative is not restricted to contemplating the provided illustrations but can interact with the visualization.

We believe that this approach will (a) enable news agencies quickly to create interactive illustrations of their news stories (in particular, longer stories that cover a number of related events) and (b) enable amateur users to provide richer forms of user-generated content (including, for example, personal essays on complex events such as political campaigns and wars, supported by interactive visualizations). Once a large number of illustrated narratives of this sort exist, it should be possible to mine them in various ways to support new forms of searching and browsing.

3 Lessons Learned

3.1 Ongoing Evaluations

User testing of the GLOCAL interface has so far been formative rather than summative. Early mockups created with the MINDMANAGER software were tested with 2–3 users at a time, each test yielding feedback about the perceived usefulness of the

[9] http://agora.cs.vu.nl/
[10] http://storify.com
[11] http://www.ourstory.com/

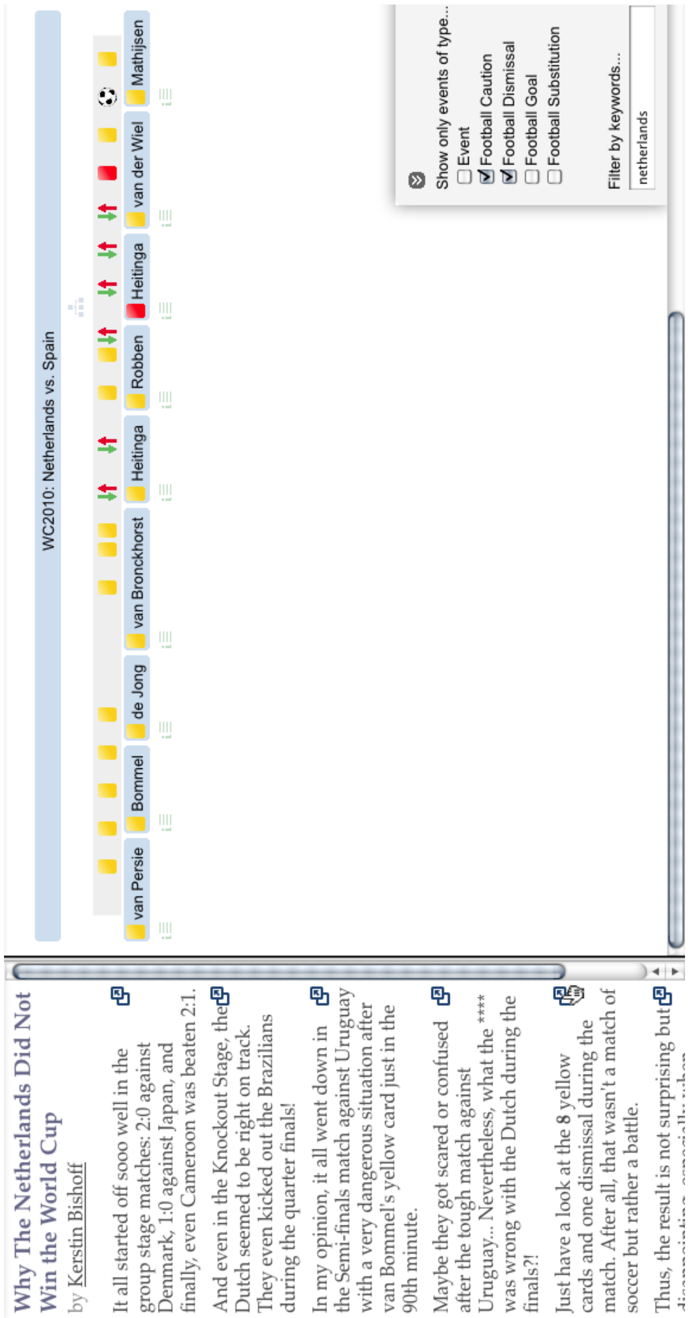

Fig. 4. Example of the use of the "narrative plus visualization" functionality offered by the interface.

(Each time the reader clicks on one of the hyperlinks to the right of the text, a previously specified view of the events and associated media is shown in the right-hand panel—in this example, a visualization of the "8 yellow cards and one dismissal" incurred by the Dutch team.)

functionality and ideas about how to improve the functionality and the visual representations. A more recent test was conducted with a version of the prototype much like the version described above. It took the basic form of a contextual inquiry ([4]): Four student-age users with typical amounts of experience with media exchange systems and social networks were observed and unobtrusively questioned as they performed various tasks with the interface. They were then interviewed about their experience.

One general result was that the representation of numerous events within a single hierarchy was seen to have advantages over the distribution of event representations over numerous hierarchically hyperlinked pages: Since the hierarchy remains basically visible at all times while the user is interacting with the complex event in question, the participants found it easier to remain oriented within the event structure. The participants stated that they could imagine using a similar interface for dealing with media and information concerning sports events, upcoming and past cultural events, politics and current events, and private events such as weddings.

The way in which controls for some functions remain invisible until they are hovered over implies that it takes a few minutes for users to become aware of all of the available functionality. But from then on, users know that they need only hover over a node of interest to be reminded of the available functions. Making these controls more conspicuous appears undesirable; in fact most of the suggestions made by the users of this and earlier versions concerned ways of reducing clutter in the interface.

At the time of this writing, a much larger-scale evaluation is being prepared in collaboration with GLOCAL partner AFP. It will involve interaction with representations of media about a set of current events: the recent uprisings in northern Africa. At the time of the workshop, it will be possible to provide some information about this evaluation study.

3.2 Overview of Contributions

This paper has aimed to contribute (at least to some extent) to each of the four questions of this workshop that concern the exploitation of events for the provision of new or improved services:

1. How can event representations be better exploited in support of activities like semantic annotation, semantic search, and semantically enhanced browsing?

We have illustrated how, when media are closely related to events, organizing the media within an interface in terms of events opens up new and improved possibilities for search, browsing, and annotation. Essentially, the benefits are analogous to those that come from organizing media in terms of geographical maps and/or timelines. The novel contribution of this paper is to show how additional functionality such as the support for interaction with event hierarchies and flexible filtering supports these activities.

2. What application areas for semantic technologies can benefit from an increased use of event representations?

To date, two application areas that are illustrated by the GLOCAL project are (a) the provision of news and media by news agencies (GLOCAL's partners include AFP and CITIZENSIDE); and (b) the exchange of media among nonprofessional users, as well as contributions by such users to media offerings of the type just mentioned. The way in which the GLOCAL interface encourages users to contribute media and comment on

existing media distinguishes it from previous methods used for exploiting event representations for interaction in the semantic web.

3. How can we improve existing methods for visualizing event representations and enabling users to interact with them in semantic web user interfaces?

Though the semantic web community has already produced impressive and useful techniques for visualizing events and supporting interaction with them, the GLOCAL user interface augments these approaches in several ways: The novel use of functionality typical of mind mapping applications introduces new ways of interacting with event hierarchies. The idea of enabling users to illustrate narratives with interactive event representations is a novel approach to event visualization that shifts some of the representational burden from the graphical visualization to natural language text. It is true that, with enough imagination and effort, it may be possible to visualize just about any relationship between two events, even if they are temporally and spatially far apart (see, e.g., the visualizations of this sort offered by [1]). But if statements about such relationships are subjective and intended only for consumption by a human user—not for automatic processing and inference—it may be most natural to have them expressed in natural language, reserving formal representation for the relations that form the backbone of the system's internal representations and external visualizations.

4. What requirements for event detection and representation methods are implied by advances in methods for exploiting events?

The main requirement introduced by the GLOCAL interface is the need to detect and represent the *subevent* relation. Representation is straightforward, and in some domains characterized by clearly structured complex events (e.g., sports tournaments, conferences), detection may also be easily automated. In domains where any hierarchical structure is not defined in advance but rather emerges as events evolve—for example, political and military uprisings such as those that have occurred in northern Africa in 2011—the identification of events and their subevents is likely to require sophisticated automatic analysis and/or human intervention. But this conclusion need not be discouraging, given that there exist many professionals and amateurs (e.g., journalists and historians) who are more than willing to apply their knowledge and skill to the analysis and interpretation of events.

References

1. André, P., Wilson, M.L., Russell, A., Smith, D.A., schraefel, m.: Continuum: Designing timelines for hierarchies, relationships and scale. In: UIST '07 Proceedings of the 20th Annual ACM symposium on User interface Software and Technology. (2007) 101–110
2. Chen, J., Hibino, S.: Reminiscing View: Event-based browsing of consumers photo and videoclip collections. In: Tenth IEEE International Symposium on Multimedia. (2008) 23–30
3. Hirsch, C., Hosking, J., Grundy, J.: Interactive visualization tools for exploring the semantic graph of large knowledge spaces. In: Proceedings of the Workshop on Visual Interfaces to the Social and Semantic Web (VISSW 2009), in conjunction with IUI 2009, Sanibel Island, FL (2009)
4. Beyer, H., Holtzblatt, K.: Contextual Design: Defining Customer-Centered Systems. Morgan Kaufmann, San Francisco (1998)