

Scenarios and Challenges for the Sharing of Event-Indexed Media

Anthony Jameson
DFKI
Saarbrücken, Germany
<http://dfki.de/~jameson/>

ABSTRACT

The domain and goals of the new GLOCAL Integrating Project are summarized, and some likely usage contexts and scenarios are sketched. Several issues that are raised by GLOCAL which are relevant to VISSW 2010 are listed.

INTRODUCTION

This brief position paper was written shortly after the official start of an ambitious new project that is directly relevant to the VISSW 2010 workshop: The European integrating project GLOCAL (Event-based Retrieval of Networked Media). As a result, work is just beginning on the concrete elaboration of ideas that can be expected to be of interest to the workshop. By the time of the workshop itself, these ideas will have assumed a much more concrete form. Interested readers are referred to the workshop presentation (available from the author) for details.

THE GLOCAL STRUCTURE AND VISION

The Problem Addressed

Today's world is flooded with media, such as photos and videos, that depict all aspects of public and private life. Just about any individual user of computing technology can create such media, and people are often interested in sharing them with others. On a more institutional level, organizations whose responsibility is to inform the public, such as news agencies, not only collect and manage huge amounts of media themselves, but they also have an interest in leveraging the ability of private individuals to provide such media. For example, suppose there has been a demonstration surrounding a political meeting, in which several violent incidents occurred. Suppose further that a (fictitious) news agency "BCD" has captured some photos of the demonstration, and amateur photographers are likely to have captured some additional images that BCD would like to add to its archive. Many of these amateurs will be eager to share their photos with others, either directly in a peer-to-peer manner or by uploading them to a central repository where others can get them.

Although these forms of media sharing are already widespread and supported to a considerable extent by existing technology, the support is limited in various respects. The

typical metadata for an individual medium are geographical and temporal coordinates, supplemented to some extent with tags that are added by human annotators or through automatic analysis of the medium. Such metadata are adequate for some purposes, but GLOCAL aims to make it feasible to add richer metadata that facilitate more effective organization and sharing of media. Its approach is based on two key ideas:

Indexing of Media in Terms of Events Many (though not all) media are associated with events, which may range from important events reported around the world to personal events such as a family's walk in the park on a weekend. As has been suggested, for example, by [2], it is often natural to characterize a medium in terms of an event that is associated with; see Figures 1 and 2. Creating an organizational scheme for media in terms of events is certainly not trivial: In general, an event includes subevents, which may in turn include other subevents, etc. These subevents are related to each other in time, perhaps overlapping, and they may be causally related as well.

Although it has not yet been decided in GLOCAL to what extent the event-based representation of media will involve ontologies in the usual sense of the term, in this position paper the term *ontology* will be used for concreteness and clarity. An event-based ontology for organizing media will include, among other elements, (a) general concepts for types of events (e.g., "political demonstration"), (b) event instances (e.g., a demonstration on April 2, 2009 at the G-20 Summit in London), and (c) media instances (e.g., particular photos and videos from that demonstration).

Collaboration Between Global and Local Media Ontologies With the term *global media ontology*, we refer to a collection of media, represented as an ontology, that comprises a large number of media and makes them available to a large number of people. (An analogous collection in the domain of scientific documents would be the ACM classification of computer science publications, coupled with ACM's Digital Library.) In the domain of media management, we could imagine a press agency making available a large repository of photos and videos that it has organized and indexed in terms of an event-based ontology, as well as according to other criteria, such as the persons depicted, the place, and the time.

Even when such a global ontology is available, it does not in general represent a suitable way for individuals to organize

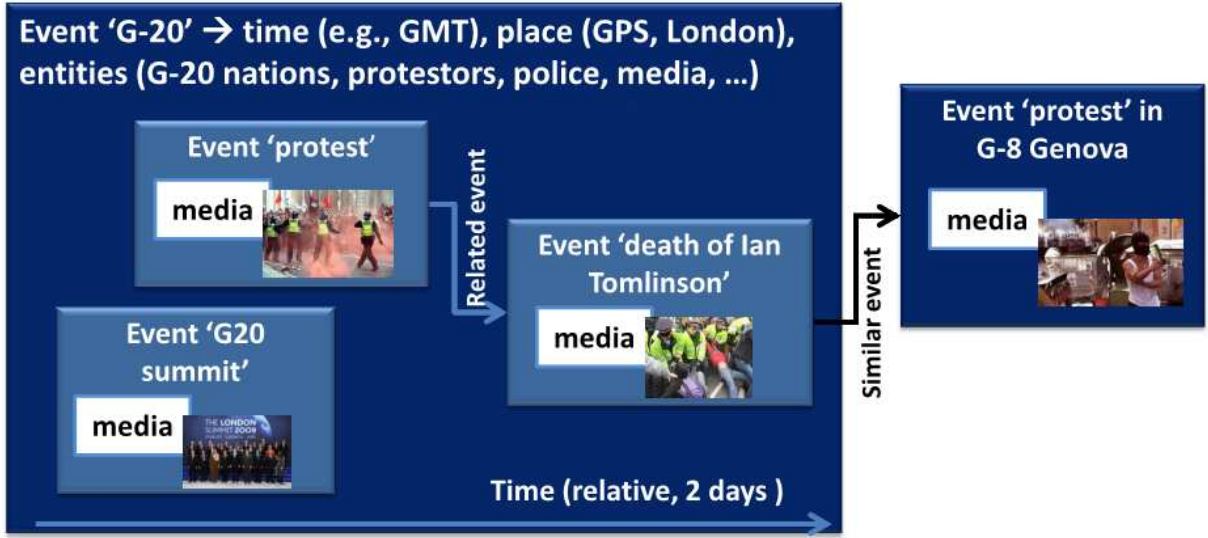


Figure 1: Illustration of the idea of indexing media (here: photos) in terms of events.

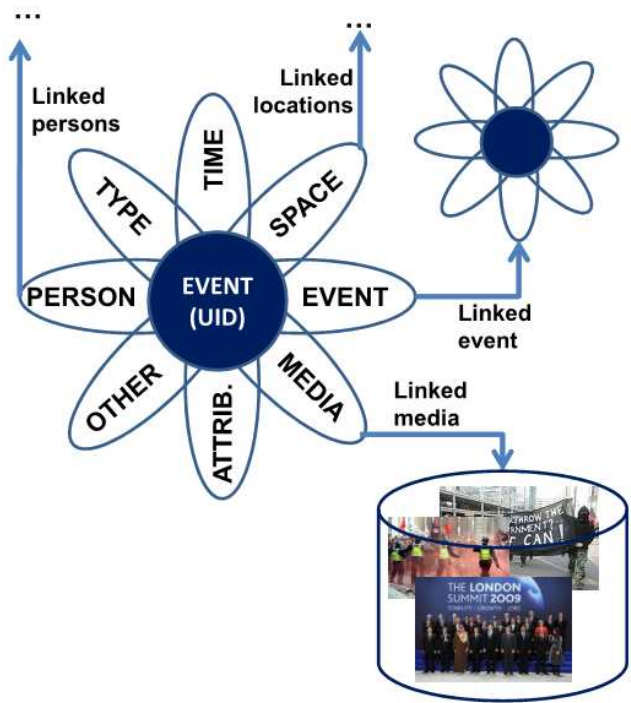


Figure 2: A more detailed view of the structure of an individual event.

their own media. An individual will in general be interested in only some small fraction of the global ontology, and they are likely to want to add concepts and events that are personally meaningful to them, such as concepts involving their own families. (Similarly, few computer scientists organize their own literature collection only in terms of the ACM taxonomy.) Therefore, the individual user should be helped to

create a local ontology that is meaningful to them.

A second major goal of GLOCAL is to provide ways of bridging the gap between global and local ontologies.

In the next section, we will illustrate these general themes of GLOCAL with reference to relatively concrete usage scenarios before listing, in the final section, some of the interaction design and research issues that the scenarios raise.

ASSUMPTIONS ABOUT THE OVERALL SITUATION

Suppose that a press agency “BCD” maintains a public website that offers, among other things, media that are indexed according to events. BCD would like the public to contribute their own media to the site. In return, it provides services to the public that help them to organize and extend their own personal collections of similar media. Thus both BCD and the public benefit. This relationship is somewhat similar to that between Google (with its YouTube for videos and Web Albums for photos) or Flickr (with its photo sharing site) on the one hand and individual users on the other hand, except that we assume that BCD exploits technology, developed in GLOCAL, for representing events and indexing media in terms of them.

The Global Representation of Media

We assume that BCD supplies each user with web pages on which they can organize their own media, making use of both (a) the ontology used for BCD’s own central collection and (b) some of the software tools for classifying media that BCD itself uses. For example, BCD may make available a tool that takes as input the space and time metadata of a medium and outputs one or more events that the medium seems likely to belong to, making use of BCD’s large event repository. For instance, a video clip taken at a certain date and time in a certain part of London may be tentatively classified as depicting a particular part of a demonstration (cf. Figure 1). This situation is comparable to the way in which Google and Flickr

provide tools that make it easy to geolocalize a new photo supplied by a user. This web environment should represent a fairly attractive place to organize media for many users, even though they are unlikely to find BCD's global media ontology entirely suitable for organizing their own media.

The Local Representation of Media

Since only a tiny part of the central ontology of BCD is of interest to any given individual user, having the rest available for the user's local collection would in general just be a distraction. One way of enabling users to benefit from the central ontology is to allow them to import selected parts of the BCD ontology (without the media instances, at first) to serve as an initial organizing ontology for their own collection. For example, if the BCD scheme covers all of professional sports, the user may want to adopt the part that covers only their favorite sport, in their own country, on the level of play that interests them.

The user may then want to expand this small ontology by adding personally relevant concepts, for example concerning a regional amateur sports league that is not included in BCD's ontology.

This combination of selection and expansion is a fairly frequently occurring way of bridging a gap between a general organizational scheme and one that is appropriate for an individual user. For example, the desktop search system used by the present author comes preconfigured with 11 categories of items that the user might search for (e.g., email, attachments, music, tasks, . . .), each type being associated with an appropriate query interface screen. A typical user may decide that only three or four of these categories are relevant to them and thus remove the other categories from view; and the user may go on to change or expand the interfaces for these categories or to define new categories of their own—though good interface design is required if these latter tasks are to be performed easily.

Even though the ontology for the local collection is quite different from that of the global collection, there is a strong enough relationship between them to facilitate sharing between the collections. In particular, the situation is more favorable than that which occurs when ontology matching is to be done between two ontologies that arose independently. We will now consider two typical use cases to see how this sharing might work.

TWO USE CASES

Use Case 1: A User Wants to Upload Media to BCD's Collection

Suppose that a user would like to donate some media that they have created themselves to BCD's collection, for example because they were able to take a particularly interesting video clip at a demonstration.

One challenge here is to figure out where each medium belongs in BCD's ontology. The user can in effect collaborate with BCD's software to solve this problem, since each of them has strengths that the other lacks: The user probably has pretty good information about where the medium was created and what it refers to; the system has a better knowl-

edge of BCD's ontology, and it can probably infer something about the medium by analyzing the metadata or the medium itself.

If the user hasn't yet classified the medium in terms of their own ontology, the system may start by inferring what it can from the medium's content and metadata and indicating one or more possible classifications in terms of the global ontology, asking the user to help make the final decision.

If the user has already inserted the medium into their own collection (maybe with the help of the system), finding the correct location for it in the global ontology should be relatively easy, if the local and global schemes are related in the ways sketched above.

In a more sophisticated variant of this use case, the user may have a set of new media and in effect offer them to the BCD system, asking the system which ones it is "interested in". In this case, the system's basic strategy could be to see where each of the offered media would fit into the central system and to judge whether it would fill a gap there; if so, the medium is inserted there and the user is asked for any desired additional information. This idea of having a central system indicate what media it would like to have could be an important way of motivating users to contribute. People are more willing to make an effort to contribute something if they know that their contribution will fill a gap and would not be likely to be made by someone else (cf. the Wikipedia SUGGESTBOT of [1]).

Use Case 2: The User Wants to Retrieve Media of Interest From the Central Repository

In addition to contributing to the BCD repository, a typical user is likely to want to benefit from this repository by browsing it or by issuing queries to find particular items.

To start with the case of querying: The indexing of media in terms of events opens up new possibilities for querying: If the user manages to specify a particular event in the central ontology, they can retrieve all media that are indexed as being related to that event. Specifying a particular event in the central ontology may not be trivial, but there are various alternative approaches, one or more of which may be applicable in any given situation. For example, the user can specify some combination of keywords, spatiotemporal coordinates, or examples of media from their own collection; and if this query does not uniquely identify one event in the central ontology, the system can ask for disambiguation.

The user may want to insert a successfully retrieved item into their own collection. If (as we assume) the ontology of the personal collection (aside from the specific media themselves) is based on a subset of the larger ontology, the position of the medium in the central ontology will often be a good guide as to where it belongs in the user's ontology, though the user may need to provide input with regard to the medium's classification in terms of events and other concepts that they have added to their local ontology themselves.

CHALLENGES FOR VISUAL INTERFACES TO THE SOCIAL AND SEMANTIC WEB

The scenarios sketched above are relevant to the VISSW workshop in that they involve sharing of resources with the help of semantic technologies. The novel aspects of the scenarios—which are due to GLOCAL’s characteristic features—imply that the interaction designs and user interfaces that support these scenarios will have to solve some fairly novel problems:

1. How can a user be enabled to browse a media collection that is indexed and organized in terms of events (as well as in terms of individual persons and other concepts and entities)? This central ontology, however it may be formalized exactly, will include a large number of concepts, relations, and instances, most of which are of no interest to the current user at a given time. A browsing interface will have to visualize this content in a comprehensible manner and filter it in appropriate ways.
2. How can a user be helped to create a personal ontology that is based on a subset of a much larger central ontology?
Although the basic idea here may be simple, it is not obvious how the details of the interaction should be designed.
3. How can a user be helped to formulate queries that specify a desired event in a central event-based ontology?
As was suggested above, there are various types of specification that may go into such a query, and which specifications are needed may become evident to the user only during the querying process itself. So there should be some support for iterative refinement of such queries, as well as interface elements that make it convenient to specify keywords, geographical and temporal regions, or example media as part of the query.
4. How can a system visualize a suggestion to the user as to where a given medium should be inserted in the user’s local collection?
One can imagine solutions such as providing a high-level view of the user’s local ontology and highlighting places where the system thinks that a given medium might belong. But then it must also be made easy for the user to evaluate these suggestions, perhaps with the help of some explanation of the suggestion by the system or the ability to compare the new medium in question with those already at the suggested location.

Because of the very early stage in which this work finds itself, the contribution of this position paper has been merely to point to some promising possible ways of supporting the sharing of media between global and local repositories with the help of event-based indexing. The issues raised have some generality that goes beyond the specific use cases and usage scenarios. By the time of the workshop itself, it will be possible to present proposed solutions that have been worked out for these and related questions, with the aim of stimulating discussion at the workshop.

Acknowledgments

The research described in this position paper is being conducted in the context of the 7th Framework EU Integrating Project *GLOCAL: Event-based Retrieval of Networked Media* under grant agreement 248984. Thanks are due to the anonymous reviewers, whose comments will substantially affect the presentation at the workshop.

REFERENCES

1. Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. SuggestBot: Using intelligent task routing to find work in Wikipedia. In Tessa Lau and Angel R. Puerta, editors, *IUI 2007: International Conference on Intelligent User Interfaces*, pages 32–41. ACM, New York, 2007.
2. Ramesh Jain. EventWeb: Developing a human-centered computing system. *IEEE Computer*, 2008(2):42–50, 2008.