3-2008

Traveling the Semantic Web through Space, Time, and Theme

Amit P. Sheth

Wright State University - Main Campus, amit@sc.edu

Matthew Perry

Wright State University - Main Campus

Follow this and additional works at: <https://corescholar.libraries.wright.edu/knoesis>

Part of the [Bioinformatics Commons](#), [Communication Technology and New Media Commons](#), [Databases and Information Systems Commons](#), [OS and Networks Commons](#), and the [Science and Technology Studies Commons](#)

Repository Citation

Sheth, A. P., & Perry, M. (2008). Traveling the Semantic Web through Space, Time, and Theme. *IEEE Internet Computing*, 12 (2), 80-85.

<https://corescholar.libraries.wright.edu/knoesis/213>

This Article is brought to you for free and open access by the The Ohio Center of Excellence in Knowledge-Enabled Computing (Kno.e.sis) at CORE Scholar. It has been accepted for inclusion in Kno.e.sis Publications by an authorized administrator of CORE Scholar. For more information, please contact library-corescholar@wright.edu.

Traveling the Semantic Web through Space, Time, and Theme

Amit Sheth and Matthew Perry • *Wright State University*

Nearly all human activity is rooted in space and time, but we can in fact describe real-world entities and events along three dimensions: thematic, spatial, and temporal. As an example, consider the following event: “the Georgia Bulldogs defeated the Florida Gators 42 to 30 on Saturday, 27 October 2007, at Jacksonville Municipal Stadium.” The thematic dimension describes what occurred (a football game involving the Georgia Bulldogs and Florida Gators), the spatial dimension describes where the event occurred (Jacksonville Municipal Stadium in Jacksonville, Florida), and the temporal dimension describes when the event occurred (27 October 2007).

So far, Semantic Web researchers have focused most of their attention on the thematic dimension, but increasing amounts of spatial and temporal data are appearing on the Web. Examples include images taken with GPS-enabled cameras that automatically generate spatial coordinates and time-stamp metadata, time-stamped video of police cruisers posted on YouTube, and uploaded images in a Web-based photo album in which the user has provided location information. We’ve also seen increasing amounts of user-generated geospatial metadata created with geotagging vocabularies such as GeoRSS. The number of Web mashups created with public map services alone is a testament to the usefulness of maps and spatial data in a variety of applications. These real-world scenarios motivate us to argue that current tools for managing Semantic Web data must be extended to better handle spatial and temporal data. Better yet would be an extension and enrichment of the Web at the middleware and infrastructure level with spatial and temporal annotation, querying, and reasoning capabilities.

In this installment of Semantics and Services, we further develop the idea of spatial, temporal, and thematic (STT) processing of Semantic Web data and describe the Web infrastructure needed to support it. Starting from Ramesh Jain’s vision of the EventWeb¹ as a view of what’s possible with a Web that better accommodates all three dimensions of event-related information (thematic, spatial, and temporal), we outline the architecture needed to support it and current research that aims to realize it.

The Event Web Vision

Events are fundamental for relating entities in space and time.² Consider our college football game example: we can find substantial information about the game on the Web, from YouTube video clips to images on Flickr to stories from sports and news Web sites to audio clips from radio broadcasts to streaming of sensor-collected traffic and weather data. Relating all this data spatially and temporally around the sequence of thematic concepts of events – the plays – that make up the game will organize the data so that a vivid picture of the overall event – the game itself – emerges. Using temporal information, we can match video clips with audio commentary to get a better description of a given series of plays, for example, or we can incorporate spatial information to view images of the same play from different positions around the stadium.

Jain described vast collections of event data as the Web’s next evolution: “EventWeb organizes data in terms of events and experiences and allows natural access from users’ perspectives. For each event, EventWeb collects and organizes audio, visual, tactile, textual, and other data to provide people with an environment for experi-

encing the event from their perspective. EventWeb also easily reorganizes events to satisfy different viewpoints and naturally incorporates new data types – dynamic, temporal, and live. The current Web is document-centric hypertext. Unlike events, hypertext has no notion of time, space, or semantic structures other than often ad hoc hyperlinks.”¹

In our work, we envision a Web infrastructure that provides the means for realizing this web of interrelated events for traversal in any STT dimension. To illustrate this enhanced Web infrastructure, we draw an analogy to a GPS satellite system, which lets a GPS receiver automatically determine its location, speed, direction, and time. With such information, we can put a real-world event into its own spatial and temporal context. Similarly, the EventWeb provides an infrastructure for placing Web data and documents into their own spatial and temporal context via services that enhance Web data and documents with spatial and temporal metadata. We also envision the use of event registries in which users can upload other data about various events.

Realizing the EventWeb

Key components in the EventWeb architecture come from combining research about spatial and temporal data management in the geographic information systems (GIS) and database communities with current Semantic Web research and technologies (ontologies, representation languages, query languages, and so on). Let's first examine the architecture and then the various approaches for enabling its major components.

EventWeb Architecture

Figure 1 shows a system architecture for realizing the EventWeb. The major components include various services for processing spatial and temporal data and events, registries

Figure 1. EventWeb architecture. The main components are event registries and various services for managing event data.

Figure 2. Example instantiation of the EventWeb architecture. A custom metadata extraction service extracts event-related spatial, temporal, and thematic (STT) metadata about police incidents from dashboard video and corresponding incident reports and loads the resulting events into an event repository. A client uses a query service in combination with Google Maps to create a mashup displaying all accidents near a specific area on a map.

for storing event data, and shared STT ontologies. A shared understanding helps normalize data to a common frame of reference so that meaningful comparisons of events in space and time are possible.

The EventWeb needs five types of core services: catalog, spatial

and temporal metadata extraction, STT query, event notification, and event update services. Catalog services maintain a list of available event-related services and let providers register (and clients discover) their services. Metadata extraction services automatically

Figure 3. Ontology-based model of space, time, and theme. An upper-level ontology defining basic classes and relationships is shown in blue, and a sample military domain ontology is shown in magenta for illustration.

extract spatial and temporal meta-data from Web documents. The other three types of services are associated with event registries that store aggregated event data from various sources: STT query services let clients query and analyze data stored in event repositories, event notification services push relevant information about new events to associated clients, and event update services add to and edit event data stored in registries.

Figure 2 shows a possible interaction between information producers and consumers in this architecture.

Representing STT Data

The first requirement in this Web infrastructure is a representation of STT data. Our current approach uses standard data models and representation languages from the W3C – specifically, Resource Description Framework (RDF).

RDF represents metadata as triples in the form (subject, property, object), which denotes that a resource – the *subject* – has a *property* whose value is the *object*. We can view a set of RDF triples as a labeled graph in which a directed edge

labeled with the property name connects the subject to the object. RDF Schema (RDFS) provides a standard vocabulary for describing the classes and relationships used in RDF statements and consequently lets us define ontologies.

But to analyze the temporal properties of relationships in RDF graphs, we need a way to record the temporal properties of the statements in those graphs, and we must account for the effects of those temporal properties on RDFS inferencing rules. Claudio Gutierrez and his colleagues³ introduced the notion of temporal RDF graphs for this purpose.

Temporal RDF graphs model linear discrete absolute time and are defined as follows. Given a set of discrete, linearly ordered time points T , a temporal triple is an RDF triple with a temporal label $t \in T$ that represents its valid time; we use the notation $(s, p, o):[t]$ to denote this temporal triple. The expression $(s, p, o):[t_1, t_2]$ is a notation for $\{(s, p, o):[t] | t_1 \leq t \leq t_2\}$. A temporal RDF graph is a set of temporal triples. Let's consider a soldier s_1 assigned to the 1st armored division (1stAD) from 3 April 1942 until 14 June 1943 and then assigned

to the 3rd armored division (3rdAD) from 15 June 1943 until 18 October 1943. This would yield the following triples: $(s_1, assigned_to, 1stAD) : [04:03:1942, 06:14:1943]$, $(s_1, assigned_to, 3rdAD) : [06:15:1943, 10:18:1943]$. We can use any temporal ontology that defines a vocabulary of time units to precisely specify time intervals' start and end points.

To represent STT data using RDF, we defined a small upper-level ontology that defines the basic classes and relationships of the thematic and spatial domains (see Figure 3); we used temporal RDF to label relationship instances with their valid times.⁴ Our upper-level ontology distinguishes between *continuants*, which persist over time and maintain their identity through change, and *occurents*, which represent processes and events. *Spatial_Occurents* and *Named_Places* are spatial entities directly linked with *Spatial_Regions* that record their geographic location, and *Dynamic_Entities* represent those with dynamic spatial behavior. Temporal intervals on relationships denote when the relationship holds (valid time).

Automatic Semantic Metadata Extraction

Given the extensive research and rapidly growing set of capabilities in the field of automatic semantic metadata extraction,¹ our discussion on the topic only gives illustrative examples.

Named entity recognition is the problem of identifying occurrences of known entities in a document — for example, recognizing the entity “Wright State University” in an HTML document and explicitly asserting that this string refers to an instance of the concept “University” identified on the Web by a specific URI. This model reference to the URI links the document with knowledge stored in the ontology. Our previous work with the Semantic Enhancement Engine² represents an example of commercial-grade named entity recognition. In addition to textual data, extraction of multimedia data must be supported, which could involve linkage of low-level features in an image or video frame with high-level concepts from an ontology.³ Identifying spatial entities and dates is necessary for extracting spatial and temporal information — for example, the Spatially-aware Information Retrieval on the Internet (SPIRIT) project⁴ recognized named places (such as park names) and associated the corresponding low-level spatial features (such as points, lines, and polygons) with documents to create spatial metadata. Additionally, our recent work⁵ recognizes onscreen time-stamp information from police videos to associate explicit temporal metadata with those videos.

Relationship extraction is the process of identifying instances of named relationships in documents, and it’s critical for ex-

tracting event data. Such extraction lets us identify interactions between entities that indicate events as well as the relations that indicate an event’s spatial and temporal properties, such as “occurred near location x” or “happened before 3:00 pm.” In our recent work,⁶ we used natural language processing techniques to identify instances of Unified Medical Language System (UMLS) relationships in documents from the PubMed repository.

References

1. A. McCallum, “Information Extraction: Distilling Structured Data from Unstructured Text,” *ACM Queue*, vol. 3, no. 9, 2005, pp. 48–57.
2. B. Hammond, A. Sheth, and K. Kochut, “Semantic Enhancement Engine: A Modular Document Enhancement Platform for Semantic Applications over Heterogeneous Content,” *Real World Semantic Web Applications*, V. Kashyap and L. Shklar, eds., Jos Press, 2002, pp. 29–49.
3. Y. Jin, L. Wang, and L. Khan, “Improving Image Annotations Using WordNet,” *Proc. 11th Int’l Workshop on Advances in Multimedia Information Systems*, Springer, 2005, pp. 115–130.
4. C.B. Jones et al., “The SPIRIT Spatial Search Engine: Architecture, Ontologies, and Spatial Indexing,” *Proc. 3rd Int’l Conf. Geographic Information Science*, Springer, 2004, pp. 125–139.
5. C. Henson et al., “Video on the Semantic Sensor Web,” *Proc. W3C Video on the Web Workshop*, 2007, www.w3.org/2007/08/video/positions/Wright.pdf.
6. C. Ramakrishnan, K. Kochut, and A. Sheth, “A Framework for Schema-Driven Relationship Discovery from Unstructured Text,” *Proc. 5th Int’l Semantic Web Conf.*, Springer, 2006, pp. 583–596.

Metadata Extraction

A fundamental task needed for analyzing events on the Web is semantic metadata extraction. Consequently, our architecture’s metadata extraction component is responsible for creating the semantic data sets that underpin the EventWeb. The architecture will require the ability to extract named entities and relationships as well as spatial and temporal information from both textual and multimedia data. We envision large collections of specialized extraction services for various types of data and extraction tasks (see the “Automatic Semantic Metadata Extraction” sidebar).

Event Notification

Event notification services let information consumers specify events of interest and then notify them when such events occur. Realizing event

notification services therefore requires a mechanism for consumers to identify and subscribe to events and an infrastructure to respond to those subscriptions.

One option for event specification could be a form of semantic template⁵ in which users identify concepts of interest in domain ontologies (event types, specific entities, and so forth) along with spatial and temporal regions to focus event requests in space and time. The system could then judge relevance based on the semantic proximity of the events and the concepts of interest. Clearly, the event’s spatial and temporal proximity to the regions specified in the template will be very important for determining relevance. Another option would be to formulate an STT query as an event request.

At the infrastructure level, we can use research in publish-sub-

scribe systems to manage collections of information requests. Research in datastream management systems and continuous queries are also relevant at the event repository level for efficient processing of notification requests as the repository is updated.

Querying STT Data

To search and analyze objects and events on the Web in STT dimensions, we need better support for STT data queries. We presented a prototype implementation of a basic set of spatial and temporal query operators for RDF graphs.⁶ These operators represent a solid first step toward a framework for querying in the EventWeb. Their implementation allowed graph pattern queries (involving spatial variables) over temporal triples and supported filtering results based on spatial and temporal predicates.

Let’s look at an example from the

Desired search in natural language

Find all soldiers with symptoms indicative of exposure to chemical X who fought in battles within 2 miles of sightings of any members of enemy group Y

Query specified using the *spatial_eval* SQL table function

```
SELECT a FROM TABLE (spatial_eval (
  '(?a has_symptom ?b) (Chemical_X induces ?b) (?a fought_in ?c)', ?c,
  '(?d member_of Enemy_Group_Y) (?d spotted_at ?e)', ?e,
  'geo_distance(distance=2 units=mile)'));
```

Figure 4. Example spatial, temporal, and thematic (STT) query over an RDF graph. The SQL query uses the *spatial_eval* operator to search for specific types of thematic relationships and filter the found relationships based on their spatial properties.

battlefield intelligence domain: suppose an analyst is assigned to monitor the health of soldiers to detect exposure to a chemical or biological agent that might imply a biochemical attack. The analyst could search for connections among soldiers, chemicals, enemy groups, and battlefield events; Figure 4 illustrates how to specify such a search in our system.

With this query, we use the *spatial_eval* operator to specify a relationship among a soldier, a chemical agent, and a battle location as well as a relationship between members of an enemy organization and their known locations. We then limit the results by the spatial proximity of the battles and enemy sightings. The *spatial_eval* operator is one of the implemented functions. In addition, a *spatial_extent* operator allows users to retrieve the spatial geometry associated with the spatial entities composing a thematic relationship and optionally filter the results using a spatial predicate – for example, “find all soldiers participating in military events that take place within an input bounding box.” For temporal aspects, an analogous *temporal_extent* operator returns a given relationship’s temporal properties and allows optional filtering – for example, “return all soldiers exhib-

iting a given symptom during a specific time period.” A *temporal_eval* operator can also answer queries such as “find soldiers who exhibited symptoms after participating in a given military event.” With Web 2.0-based semantic interfaces, the power of such STT query capability transfers to the hands of casual Web users, letting them ask questions such as “show all event photos and videos taken in Central Park on New Year’s Eve,” or “create a montage of multimedia content on cultural attractions in Vienna created in March.” A preliminary step toward such capability appears in our Semantic Sensor Web project at <http://knoesis.wright.edu/projects/sensorweb/>.

We see great potential for realizing the EventWeb in the sensor networks domain. The Open Geospatial Consortium’s (OGC) sensor Web enablement initiative proposes a suite of specifications related to sensors, sensor data models, and sensor Web services. These standards were intended to allow discovery, exchange, and processing of sensor data, but it’s clear that purely syntactic standards specifications aren’t sufficient for realizing this goal. Adding semantics through domain ontologies

and spatial and temporal ontologies would allow the extra machine processing capabilities required to realize the sensor Web’s goal and yield a Web of events in the sensor networks domain. As initial steps in this direction, we’re working on semantic extensions to the OGC standards.⁷

The result of the enhanced infrastructure presented here will be an organization of information on the Web that’s closer to a human’s perspective than a machine’s. We naturally conceptualize our interactions as events, and the STT relations between events are crucial to our understanding of the world. The EventWeb will consequently lead to better understanding and use of the vast amounts of data currently on the Web and surely to come. □

References

1. R. Jain, “EventWeb: Developing a Human-Centered Computing System,” *Computer*, vol. 41, no. 2, 2008, pp. xx-xx.
2. U. Westermann and R. Jain, “Events in Multimedia Electronic Chronicles (EChronicles),” *Int’l J. Semantic Web and Information Systems*, vol. 2, no. 2, 2006, pp. 1-23.
3. C. Gutierrez, C. Hurtado, and A. Vaisman, “Temporal RDF,” *Proc. European Conf. Semantic Web*, Springer, 2005, pp. 93-107.
4. M. Perry, F. Hakimpour, and A. Sheth,

- “Analyzing Theme, Space and Time: An Ontology-Based Approach,” *Proc. 14th ACM Int’l Symp. Geographic Information Systems*, ACM Press, 2006, pp. 147–154.
5. K. Gomadam et al., “A Semantic Framework for Identifying Events in a Service Oriented Architecture,” *Proc. IEEE Int’l Conf. Web Services*, IEEE CS Press, 2007, pp. 545–552.
 6. M. Perry et al., “Supporting Complex Thematic, Spatial and Temporal Queries over Semantic Web Data,” *Proc. 2nd Int’l Conf. Geospatial Semantics*, Springer, 2007, pp. 228–246.
 7. C. Henson et al., “Video on the Semantic Sensor Web,” *Proc. W3C Video on the Web Workshop*, 2007; www.w3.org/2007/08/video/positions/Wright.pdf.

Amit Sheth is an IEEE fellow, LexisNexis Ohio Eminent Scholar, and director of the Kno.e.sis Center at Wright State University. Contact him via <http://knoesis.wright.edu>.

Matthew Perry is a researcher at the Kno.e.sis Center and a PhD candidate in computer science at Wright State University. His research focuses on spatial-temporal-thematic query processing. Contact him via <http://knoesis.wright.edu/students/mperry/>.