SemanticSpy: Suspect Tracking Using Semantic Data in a Multimedia Environment

Amit Mathew, Amit Sheth, and Leonidas Deligiannidis

Department of Computer Science
University of Georgia, Athens, Georgia 30602, USA
{mathew,amit,ldeligia}@cs.uga.edu

Abstract. In this paper we describe SemanticSpy, a tool for tracking suspects and their activities around the globe. Using an RDF data store and a three-dimensional multimedia environment, SemanticSpy is able to bring semantic information to life. SemanticSpy utilizes a blend of semantic information and a media-rich environment that can display images, 3D models, and documents, and can play audio clips to give the user the ability to analyze information using the different types of media. We believe that this will help law enforcement officials gain insight into a suspect's behavior and give them the power to avert the next major crime or terrorist attack. Finally, we describe SemanticSpy's architecture and explain its advantages over other information environments.

1 Introduction

In this age of international terrorism and global drug smuggling rings, it is apparent that there is a need for law enforcement and counter-terrorism agencies to track the movement and activities of suspects around the world. The Semantic Web [1] can aid in this effort because it allows relationships among entities to be represented in a manner that is both human- and machine-processible. Relationships are a vital aspect of criminal and terrorist organization analysis because such organizations are composed of relationships among members, relationships among events, and so on. Furthermore, Semantic Web technologies, like RDF (Resource Description Framework) [7], provide a suitable format for metadata. Metadata is important whenever details matter: we don't just want to know that a meeting occurred, we also want to know who was involved, where it happened, and how it took place.

The number of relationships and amount of metadata generated from tracking a suspect can easily overwhelm security officials. SemanticSpy offers a way to harness this abundance of data and visualize it in a variety of ways. Phone conversations can be played, photographs of a suspect's meeting can be viewed, and a 3D model of the car a suspect is driving can be shown. With the help of semantic-enhanced data and interactivity, SemanticSpy can potentially allow an analyst to piece together seemingly trivial details to help forestall an act of terrorism.

Fig. 1. A screenshot of SemanticSpy

2 Motivations

The two major features of SemanticSpy that differentiate it from existing tools are its use of semantic data and its rich multimedia environment. The value of these two features and how they benefit the user are discussed next.

2.1 Advantages of Semantic Data

Semantic data has many advantages over data stored in a traditional database, but the two key aspects are its emphasis on relationships among entities and its use of metadata. Relationships are important in many aspects [9]. One is the idea of human networks and its relationship to organized crime, terrorism, and other illicit activities. Suspects participating in illegal activities often are part of social networks that need to be analyzed to determine leadership and level of involvement. Furthermore, relationships are important when trying to anticipate a suspect's actions. For example, a suspect may have a relationship to a chemical plant which may be legitimate (if the suspect is a chemical engineer) or not (if the suspect is investigating using hazardous materials for malevolent purposes). Explicit semantics afforded by named relationships in Semantic Web representations provide better quality of analysis and insight than what is possible by simple graphs with unnamed edges as used in many earlier approaches that study such networks.

The second important aspect of the Semantic Web is the utilization of metadata. Metadata is often treated as a side effect or unnecessary detail in traditional settings, but for intelligence agencies, metadata can make the difference in catching a criminal.

Metadata allows a user to detect relationships that would otherwise go unnoticed. For example, the fact that two seemingly unrelated people are traveling in a car may be innocuous. However, the fact they are traveling in the two cars with registrations belonging to the same person or persons in the same group may be important.

Use of such tools presupposes extensive capability for metadata extraction from heterogeneous data of the type discussed in [10].

2.2 Advantages of a Multimedia Environment

The main advantage that a multimedia environment offers over a traditional single media environment is that it allows the user to benefit from the strengths of each media type. Text allows the system to convey precise data to the user. In SemanticSpy, text is used mainly to show metadata about suspects and activities. Some types of information naturally occur in an audio format, such as phone calls or other recorded conversations. Although conversations can be transcribed to text, audio provides additional information through voice intonation and volume. Images are useful for showing photographs and frames from video (which may come from surveillance videos) and are a form of media where humans perform better than computers in comparing between two samples. Three-dimensional models are useful for conveying spatial information, such as the plan for the interior of a plane, and give security professionals more actionable information. For example, a 3D model of the car that suspected terrorists are driving may reveal a hidden compartment where a bomb may be placed.

3 Functionality

SemanticSpy offers several different predefined views that aid in the tracking of suspects. These include the suspect view, the activity view, and the global view. The suspect view, shown in the upper right corner of Figure 1, shows a list of suspects that can be selected for tracking. In this figure, the suspects are taken from the list of the FBI's ten most wanted criminals [3]. Notice that the suspects' names are accompanied by their photographs, which allows security officials to see a suspect and possibly recognize the suspect in another picture. The suspect view also shows additional information about the suspect, like their physical features. The list of suspects can be scrolled through, with the currently tracked suspect shown on the top. Once a suspect is selected, the user can choose a date and see the suspect's activities on that date. In this figure, March 5, 2005, is chosen. When a day on which a suspect has been involved in some activity is selected, information for that day is shown in the activity view on the left side.

The activity view shows the type of activity that occurred and its associated metadata. Currently, SemanticSpy supports three types of activities: travel, meetings, and phone calls; more activities are under development. The different types of activities can have different types of metadata. For example, a travel activity can detail the type of travel used, whether it be by air, sea, or land. If the suspect traveled by air, the airline's flight number, and an image or model of the airplane they traveled in can

be shown. In Figure 1, a 3D model of a plane is shown with the flight the suspect took on that day. If the suspect traveled by a rental car, information about the rental car company can be shown. Each travel activity is associated with two geographic points, the point of departure and point of arrival. These points are shown in the global view, shown in the center of Figure 1.

The global view is made of a three-dimensional sphere, textured with an image of the earth displayed with terrain. The global view is helpful in several ways. It allows users to absorb information "at a glance." For example, in Figure 1, a user can find the details of each of the suspect's trips, but from the image in the global view it is immediately clear that a suspect traveled first from somewhere in North America to somewhere in South America, and later back to North America. This easy access to general information prevents the user from being overloaded by information, and allows them to more likely notice the information that can stop an international criminal. The currently active points are highlighted in red, while points of previous travel are highlighted in yellow. The point-plotting feature is not restricted to travel. It can be used to communicate any relationship between two points, such as a phone call between parties in two different locations. Another helpful feature of the map is that it shows the topographical information of the Earth. This is helpful because a user may notice a suspect traveling to suspicious terrain, such as a desert.

The other types of activities supported by SemanticSpy are meetings and phone calls. Meetings are displayed with a single geographic point, an optional photograph of the meeting area if one is available, and relevant metadata, like who was present at a particular meeting. The other type of supported activity is phone conversations. Phone calls can be accompanied by an audio recording of the conversation and information about the phone call, like who was called and from what number. All these elements can be used to gain insight into the motivations of a suspect's actions.

A typical use case starts with a user selecting a suspect that may be involved in planning a terrorist attack. The user may then choose a date from a month ago to see the suspect's previous travels and activities. The user may choose to drill down on one of those activities, like playing a recorded phone conversation between the current suspect of interest and another unknown person, or viewing the people involved in a meeting with the suspect of interest.

4 Architecture

SemanticSpy's architecture, as seen in Fig. 2, is separated into two major components: the front-end visualization system and the back-end RDF data store. The visualization system is powered by MAGE [8], a high-performance game and virtual reality engine that supports simulation and visualization tasks. MAGE supports the viewing of 3D models, 2D images, and audio, while still maintaining interactive framerates. MAGE also has a scripted material system, making it easy to update suspects' photographs or 3D models on-the-fly.

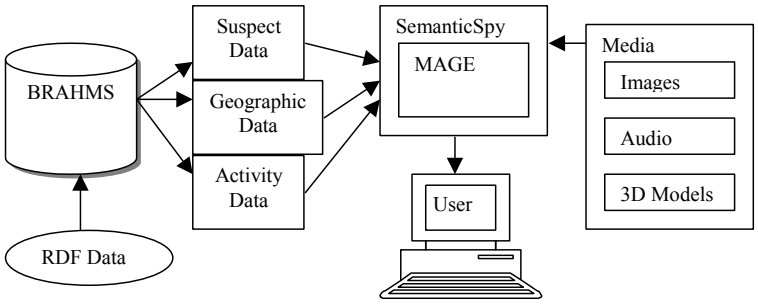

Fig. 2. A diagram of SemanticSpy's architecture

BRAHMS [5], developed in the LSDIS lab, is used as SemanticSpy's RDF data store. BRAHMS is a main-memory system that is very fast, making it an appropriate choice for a real-time visualization system. The points on the globe that mark the suspects' travel are determined by GPS coordinates in the RDF input file. The GPS coordinates are converted to Cartesian coordinates and plotted on the globe. The lines connecting the points show the different legs of a suspect's journey.

5 Related Work

Much of the work done in Semantic Web visualization today is in the area of graph visualizations, especially of large ontologies. However, there exists some work dealing with Semantic information in a spatial context. Themescape [11] provides topographical maps that show the relations among different types of documents. Ghinea et al [4] use spatial data in the context of the human body to study back pain. SemanticSpy builds upon these research works by combining geospatial visualization and Semantic Web technologies to support suspect tracking. GIS-type visualization applications that predate the Semantic Web have also been developed. For example, Eick [2] developed a method for visualizing communication networks using a three-dimensional globe that contains nodes connected with lines. Eick's work interestingly gives meaning to the thickness and color of the lines that connect the nodes.

6 Future Work and Conclusions

The development of SemanticSpy has opened up the possibility for productive future research. SemanticSpy can be extended with more types of multimedia (e.g., video), and more types of activities, such as job interviews, emails and other documents, and social gatherings. Also, the global visualization system can be extended to support zooming and satellite views. Another direction we wish to explore is to move SemanticSpy to a virtual environment to see if an immersive environment improves the user's ability to process information. Finally, we wish to harness Semantic Web tech-

nologies to intelligently extract related suspects from a large dataset for viewing in SemanticSpy, using connections based on financial transactions and social networks [6].

SemanticSpy allows law enforcement officials to view information about suspects in new and powerful ways by using a multimedia-rich environment to capitalize on humans' powerful senses and pattern recognition abilities. Furthermore, SemanticSpy is an experimental tool for visualizing semantic information in new ways, by creating domain-specific visualizations that empower analysts to detect important relationships, which facilitate their investigative tasks.

Acknowledgements. The authors would like to thank Maciej Janik for his assistance with the BRAHMS system. We would also like to thank Boanerges Aleman-Meza, Angela Maduko, and other members of the Large Scale Distributed Information Systems Lab at the University of Georgia who provided valuable feedback during the development of SemanticSpy.

References

1. T. Berners-Lee, J. Hendler, and O. Lassila. "The Semantic Web," *Scientific American*, May 2001.
2. S.G. Eick. "Aspects of Network Visualization," *Computer Graphics and Applications* 16(2), 69-72, 1996.
3. FBI's Ten Most Wanted, http://www.fbi.gov/mostwant/topten/fugitives/fugitives.htm
4. G. Ghinea and D. Gill. "Using Scalable Vector Graphics and Geographical Information Systems in a Back Pain Study," Visualizing the Semantic Web, Springer 2002, pp. 180 – 197.
5. M. Janik and K. Kochut. "BRAHMS: A WorkBench RDF Store And High Performance Memory System for Semantic Association Discovery," *Fourth International Semantic Web Conference ISWC 2005*, Galway, Ireland, 6-10 November 2005.
6. V. Krebs, "Mapping Networks of Terrorist Cells," *Connections*, 24(3): 43-52, 2002.
7. O. Lassila and R. R. Swick. "Resource Description Framework (RDF) model and syntax specification," W3C Recommendation, Feb. 1999. http://www.w3.org/TR/REC-rdf-syntax/
8. A. Mathew and L. Deligiannidis. "MAGE-VR: A Software Framework for Virtual Reality Application Development," *Proc. of The 2005 International Conference on Modeling, Simulation and Visualization Methods*, pp. 191-197, June 2005.
9. A. Sheth, I. Budak Arpinar, and V. Kashyap, "Relationships at the Heart of Semantic Web: Modeling, Discovering, and Exploiting Complex Semantic Relationships," Enhancing the Power of the Internet Studies in Fuzziness and Soft Computing, M. Nikravesh, B. Azvin, R. Yager and L. Zadeh, Springer-Verlag, 2003, pp. 63-94
10. A. Sheth, C. Bertram, D. Avant, B. Hammond, K. Kochut, Y. Warke, Managing Semantic Content for the Web, IEEE Internet Computing, July/August 2002, pp. 80-87.
11. ThemeScape Product Suite, Cartia Inc, http://www.cartia.com/