Semantic Analytics Visualization

Leonidas Deligiannidis[1,2], Amit P. Sheth[2] and Boanerges Aleman-Meza[2]

[1]Virtual Reality Lab and [2]LSDIS Lab, Computer Science,
The University of Georgia, Athens, GA 30602, USA
{ldeligia,amit,boanerg}@cs.uga.edu

Abstract. In this paper we present a new tool for semantic analytics through 3D visualization called "Semantic Analytics Visualization" (SAV). It has the capability for visualizing ontologies and meta-data including annotated web-documents, images, and digital media such as audio and video clips in a synthetic three-dimensional semi-immersive environment. More importantly, SAV supports visual semantic analytics, whereby an analyst can interactively investigate complex relationships between heterogeneous information. The tool is built using Virtual Reality technology which makes SAV a highly interactive system. The backend of SAV consists of a Semantic Analytics system that supports query processing and semantic association discovery. Using a virtual laser pointer, the user can select nodes in the scene and either play digital media, display images, or load annotated web documents. SAV can also display the ranking of web documents as well as the ranking of paths (sequences of links). SAV supports dynamic specification of sub-queries of a given graph and displays the results based on ranking information, which enables the users to find, analyze and comprehend the information presented quickly and accurately.

1 Introduction

National security applications, such as aviation security and terrorist threat assessments, represent significant challenges that are being addressed by information and security informatics research [25]. As the amount of information grows, it is becoming crucial to provide users with flexible and effective tools to retrieve, analyze, and comprehend large information sets. Existing tools for searching and retrieving information (such as search engines) typically focus on unstructured text and some may be adapted to support display of the results of text analytics. However, semantics is considered to be the best framework to deal with the heterogeneity and dynamic nature of the resources on the Web and within enterprise systems [35]. Issues pertaining to semantics have been addressed in many fields such as linguistics, knowledge representation, artificial intelligence, information systems and database management. Semantic Analytics involves the application of techniques that support and exploit the semantics of information (as opposed to just syntax and structure/schematic issues [32] and statistical patterns) to enhance existing information systems [30].

Semantic analytics techniques for national security applications have addressed a variety of issues such as aviation safety [33], provenance and trust of data sources [15], and the document-access problem of Insider Threat [3]. Ranking, or more specifically "context-aware semantic association ranking" [4], is very useful as it finds and presents to the end-user the most relevant information of his/her search. The presentation of these results is normally done via a list of paths (i.e., sequences of links). Sometimes these results also include documents ranked according to their relevance to the results. These interconnections of ranked links and documents can be viewed as a graph or a network. Visualization of large networks has always been challenging. There is an increasing need for tools to graphically and interactively visualize such modeling structures to enhance their clarification, verification and analysis [2, 38]. Effective presentation of such data plays a crucial role because it helps the end-user analyze and comprehend the data. As a result, data is transformed into information and then into knowledge [33]. Efficient understanding of semantic information leads to more actionable and timely decision making. Thus, without an effective visualization tool, analysis and understanding of the results of semantic analytics techniques is difficult, ineffective, and at times, impossible.

The fundamental goal of visualization is to present, transform and convert data into a visual representation. As a result, humans, with their great visual pattern recognition skills, can comprehend data tremendously faster and more effectively through visualization than by reading the numerical or textual representation of the data [12]. Interfaces in 2D have been designed for visualization results of queries in dynamic and interactive environments (e.g., InfoCrystal [6]). Even though the textual representation of data is easily implemented, it fails to capture conceptual relationships. Three-dimensional (3D) interactive graphical interfaces are capable of presenting multiple views to the user to examine local detail while maintaining a global representation of the data (e.g., SemNet [18]). Using virtual environments, the user is able to visualize the data and to apply powerful manipulation techniques. In addition, the user of such systems is able to view and listen to associated metadata for each subject of interest at a given location such as a suspicious phone call or an image of a handwritten message. Thus, we address the challenge of visualization in the context of semantic analytic techniques, which are increasingly relevant to national security applications. For example, semantic data allows rich representation of the movements of an individual such as a person P that traveled from city A to city B taking bus X, in which a terrorist Q was also traveling.

The contribution of this paper is a highly interactive tool for semantic analysis through 3D visualization (in a semi-immersive environment) built using Virtual Reality technology with the goal of enabling analysts to find and better comprehend the information presented. The main features of the "Semantic Analytics Visualization" (SAV) tool are three. First, SAV visualizes ontologies, metadata and heterogeneous information including annotated text, web-documents and digital media, such as audio and video clips and images. Second, interaction is facilitated by using a virtual laser pointer for selection of nodes in the scene and either play digital media, display images, or open annotated web documents. Third, SAV can display the results of semantic analytics techniques such as ranking of web documents as well as the ranking of paths (i.e., sequences of links).

2 Background

Industry and academia are both focusing their attention on information retrieval over semantic metadata extracted from the Web (i.e., collection of dots). In addition, it is increasingly possible to analyze such metadata to discover interesting relationships (i.e., connect the dots). However, just as data visualization is a critical component in today's text analytics tools, the visualization of complex relationships (i.e., the different ways the dots are connected) will be an important component in tomorrow's semantic analytics capabilities. For example, visualization is used in tools that support the development of ontologies such as ontology extraction tools (OntoLift [39], Text-to-Onto [28]) or ontology editors (Protégé (protege.stanford.edu), OntoLift). These tools employ visualization techniques that primarily focus on the structure of the ontology, or in other words, its concepts and their relationships.

The Cluster Map visualization technique [1] bridges the gap between complex semantic structures and their intuitive presentation. It presents semantic data without being burdened with the complexity of the underlying metadata. For end users, information about the population of the ontology (entities, instance data) is often more important than the structure of the ontology that is used to describe these instances. Accordingly, the Cluster Map technique focuses on visualizing instances and their classifications according to the concepts of the ontology. However, some knowledge bases and ontologies like WordNet [29], TAP [22] or SWETO (lsdis.cs.uga.edu/projects/semdis/sweto/) cannot be easily visualized as a graph, as they consist of a large number of nodes and edges. Similarly, there are many large bio-informatics ontologies like Gene ontology [20] and GlycO [21] which have several hundred classes at the schema level and a few thousand instances.

TouchGraph (www.touchgraph.com) is a spring-embedding algorithm and tool to implement customizable layouts. It is a nice tool but can be annoying to users because it keeps re-adjusting the graph to a layout it determines to be best. TGVizTab [2] is a visualization plug-in for Protégé based on TouchGraph. Large graphs can be cluttered and the user may need to manually move some nodes away from her/his point of interest to see clearly the occluded nodes or read a label on a node successfully. However, even while the user is trying to move a node out of her/his view, s/he ends up dragging the graph while TouchGraph is readjusting it. People are used to placing things where they want and coming back later to find them still there. This is quite difficult to do with TouchGraph. OWLViz, part of the CO-ODE project (www.co-ode.org), is a plug-in for Protégé to visualize ontologies but it only shows *is-a* (i.e., subsumption) relationships among concepts.

Many ontology based visualization tools are, at least partially, based on the Self-Organizing Map (SOM) [27] technique/algorithm. WEBCOM [26] is a tool that utilizes SOM to visualize document clusters where semantically similar documents are placed in a cluster. The order of document clustering helps in finding related documents. ET-Map [10] is also used to visualize a large number of documents and websites. It uses a variation of SOM called Multilayer SOM to provide a concept-based categorization for web servers.

Spectable [38] visualizes ontologies as taxonomic clusters. These clusters represent groups on instances of individual classes or multiple classes. Spectable displays class hierarchical relations while it hides any relations at the instance level. It presents each instance by placing it into a cluster based on its class membership; instances that are members of multiple classes are placed in overlapping clusters. This visualization provides a clear and intuitive depiction of the relationships between instances and classes.

In [37] a holistic "imaging" is produced of the ontology which contains a semantic layout of the ontology classes where instances and their relations are also depicted. Coloring is also employed to show which classes have more instances. However, this approach is not suitable to visualize instance overlap like in Spectable.

Our previous research on Semantic Discovery [7] has focused on finding complex relationships including paths and relevant sub-graphs in graphs representing metadata and ontological terms and then present the information to the user as a list of paths (i.e., sequences of links). We have also implemented algorithms that perform ranking based on semantic associations [4, 8], and have studied applications in national/homeland security [3, 33], bioinformatics and detection of conflict of interest using social networks [5]. We experimented with a variety of visualization tools, including Protégé, TouchGraph and Jambalaya [36], to present the results to the user in a non-textual form. In this paper, we propose SAV as a virtual reality tool that better allows visualization of results from semantic analytics techniques.

3 The Virtual Reality Prototype

The interface presented in this paper is to be used in front of a group of people where only one person interacts with the tool. Images are drawn on a projection screen via a rear-projector system.

One of the major capabilities of SAV is to assist users in comprehending and analyzing complex relations of the semantic web. With SAV, the user can navigate, select and query the semantic information and select web documents. More interesting capabilities of SAV include the visualization of hundreds of web documents and other digital media and providing an environment where a user can easily and naturally interact with the environment such as finding related documents and at the end reading the documents. The documents in our experiments are annotated documents, such as web pages (the automatic semantic annotation system we used is described in [23]). The user (e.g., intelligence analyst) is interested not only in finding the documents in the Virtual Environment (VE) but also reading these documents. Hence, upon selecting a document, the document is loaded in a conventional browser where the user can read the annotated pages the same way s/he reads conventional web pages.

3.1 Software and Hardware

The prototype was designed using Java3D, the JWSU toolkit [14] and the GraphViz system [17, 19]. The application was running on a PC with 1GB of memory and an ATI Radeon 9800 video card. The user can interact with the VE using a PinchGlove –

for selection of objects and navigation. We also used a Polhemus FastRak to track the user's head and hand in space. One of the problems with such an interface, however, is that the user may need to simply turn her head to the left or right to see the rest of the environment and therefore, the user cannot possibly look at the display. For this reason, we used 3 sensors on our Polhemus Fastrak. The first sensor is used to track the position and orientation of the user's head. The second sensor is used to track the user's hand. In Figure 1(a) the user is looking straight and the third sensor held by her left hand does not provide any additional rotation. The third sensor provides additional rotation of the camera (the view) around the Y axis – the vertical axis. Thus, instead of the user turning to the left as shown in Figure 1(b), the user can rotate the sensor that is held by her second (non-dominant) hand as shown in Figure 1(c). Users were easily adapting to this mechanism almost immediately, within seconds.

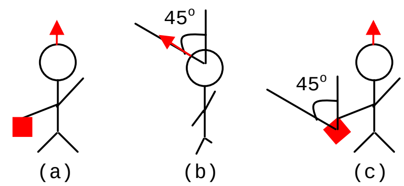

Fig. 1. Rotation sensor for the Virtual Environment.

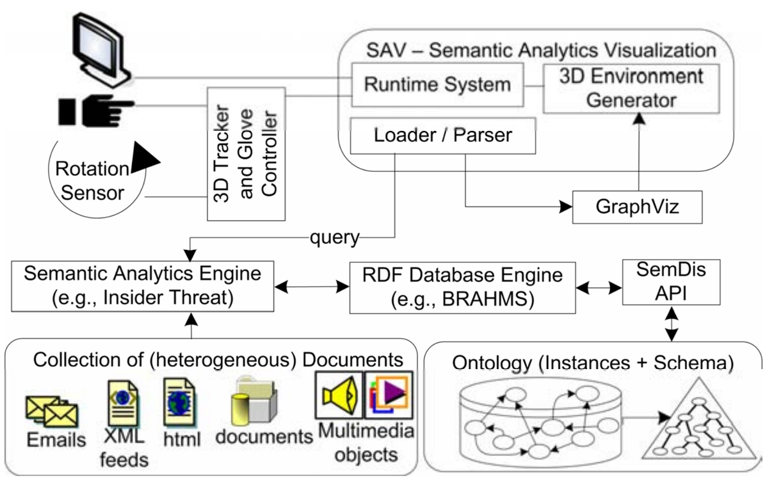

Fig. 2. System architecture of SAV.

3.1.1 Software Architecture

The application loads the ontology result using BRAHMS [24] and the SemDis API [31]. The result is fed to GraphViz's layout algorithm(s) and then we generate the

Virtual Environment (VE). The system architecture of SAV is shown in figure 2. After we generate the 3D environment based on the output of GraphViz and the weight (ranking) information of the documents, the user can interact with the VE using a 3D tracker and PinchGloves.

3.2 Retrieving and Loading the Data

The visualization of sequences of links and relevant documents for a query is built upon our previous work on Insider Threat analysis [3]. The ontology of this application captures the domain of National Security and Terrorism. It contains relevant metadata extracted from different information sources such as government watch-lists, sanction-lists, gazetteers, and lists of organizations. These sources were selected for their semi-structured format, information richness and their relationship to the domain of terrorism. The ontology population contains over 32,000 instances and over 35,000 relationships among them. For ontology design and population, the Semagix Freedom toolkit (www.semagix.com) was used. Freedom is a commercial software toolkit based on a technology developed at and licensed from the LSDIS lab [34].

A query can be a selection of a concept (i.e., class) or a particular entity from the ontology. The results of a search are related entities which are discovered by means of traversing links to other entities (in a graph representation of the ontology). Different documents are relevant to a query depending upon which entities appear within the content of the documents. The final result is a group of entities (or nodes) interconnected by different named relationships (such as 'located in'), together with a set of documents that are related to the entities. The interconnections of the entities can be viewed as sequences of links where each sequence has a ranking score. The documents also have a rank score based on their relevance to the query. For visualization purposes, a cached version of each document is placed in a separate web server to provide quick and persistent access from SAV. These documents are often web pages but can also include images, audio, or video clips whereby meta-data annotations provide the connections to entities in the ontology.

SAV accesses data to be visualized by loading the results from the query module of the application for detection of Insider Threat mentioned earlier. Note that this application exemplifies a semantic analytics technique. The architecture of SAV considers data independence of any particular semantic analytics technique. The relevant components of results fed into and visualized by SAV are: (i) each entity represents a node in the results graph; (ii) nodes are connected by named relationships; (iii) each path contains a ranking score; (iv) each document contains a relevance score and is related to one or more nodes.

3.3 Semantic Visualization

One of the more complex ontology searches we visualized is presented in the remainder of this paper as our main illustrative example. The data to be visualized (i.e., search results) includes information about entities and relationships among them, relevant annotated documents associated with entities, media associated with each entity, ranking of web pages and path ranking. To visualize the results, we partition the space into two volumes, the foreground and background volumes. In the foreground we visualize the entities and their relationships and in the background we visualize the documents. Figure 3 illustrates the foreground of SAV.

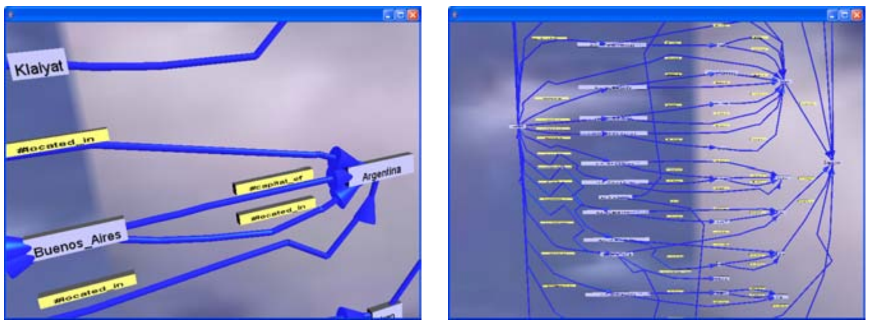

Fig. 3. (Left) Entities (bluish rectangles) and relationships (arrows between entities – a yellow rectangle above the arrow is the relationship's label) visualized in SAV. (Right) the entities and their relationships from a distance; we removed the document nodes to make the figure more readable.

SAV runs the "dot" filter of the GraphViz product to generate the coordinates of every visible node in the scene along with the splines' control points. GraphViz uses splines to connect entities, not straight lines, in order to minimize edge crossing, generate a symmetric layout and make the graph more readable. Generating splines in space requires the creation of objects with complex geometry. Instead, we create multiple cylinder objects and we connect them together to form a line that connects the control points of each spline.

GraphViz is invoked twice, once for the foreground and once for the background layout. Then, the generated output of GraphViz is loaded and 3D objects are created and attached in the scene. As illustrated in figure 3, we represent entities as rectangles and their relationships to other entities as directional splines in space; each relationship has a label attached on it. Each entity and relationship has a textual label. We dynamically generate PNG images (using Java2D) out of the textual label (with anti-aliasing enabled) of each node and then we apply the un-scaled generated images as textures on the nodes (entities and relationships). We found that by doing this the labels are readable even from a distance. Text anti-aliasing must be enabled while the images are generated and when applying the textures no scaling should be done – the images should have dimensions of power of two.

The documents are represented as red spheres in the background as shown in figure 5. The position of a document changes depending on its ranking. The higher the ranking, or else the more important a document is, the closer to the foreground it is placed. The intensity of the red color also becomes higher when the ranking is higher (redder spheres indicate higher ranking). Additionally, depending on a document's ranking, the width of the sphere representing the document becomes bigger, in which case a sphere becomes a 3D oval shape (wider 3D ovals indicate higher ranking).

Since the most relevant documents are closer to the foreground, the user can select them more easily because they are closer to the user and the objects (the spheres) are bigger in size.

The number of documents relative to a search could range from one hundred to several thousand. As a result, we visualized the entities and their relationships in the foreground while the document information stays in the background without cluttering the user's view.

4 User Interaction

The user can interact and navigate in the VE using her hand while wearing a Pinch-Glove to inform the system of her actions. The user can change the speed of travel using 3D menus similar to the ones in [13]. The state diagram that describes the functionality of SAV is shown in figure 4.

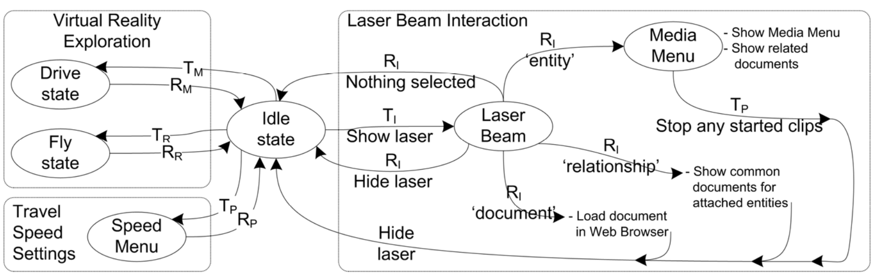

Fig. 4. State machine of SAV.

There are six states in the state machine. Initially the user is in the *Idle* state. The user can choose the "drive" or "fly" technique to explore the VE. "Drive" means that the user can travel in the VE at the direction of her hand while staying at the ground level where the ground level is defined by the horizontal plane of where the user is at the initiation of the traveling technique. "Fly", on the other hand, means that the user can freely travel in space at the direction of her hand.

On the state transition edges we place a label for which "pinch" (the contact of two fingers) moves the user to different states. A "T" indicates that the thumb and another finger are touching and an "R" indicates the release of the two fingers. The subscript indicates the second finger; the finger that touches the thumb ("I" for index, "M" for middle, "R" for ring, and "P" for pinky). For example, $\mathbf{T_R}$ means that the thumb and the "ring" fingers are touching.

When the user moves in the *Speed* menu state, s/he can increase or decrease the speed of travel. While in the *Idle* state, the user can activate the laser beam to select an entity, a relationship, or a document. Upon selecting a document, we load the document in a web browser and the user can read the annotated web page, click on links, and so on. Upon selecting an entity, all documents become semi-transparent with the exception of the documents that are related to the selected entity. Edges that connect documents to other entities become fully transparent. Additionally, all entities and relationships become semi-transparent except for the selected entity. This helps the user to easily find the related documents associated with the selected entity, and also focus on the detail presented currently while keeping the overview of the entire scene. Thus, the user's attention is on the selection while s/he is still capable of looking around to see the rest of the environment that is transparent but still visible. The transparency provides a solution to the "detail" and "overview" problem where the query result stays in focus while the rest of the VE becomes transparent but still visible. Additionally, by means of head rotations, the user can still get an overview of the environment while also being capable of seeing the detail of the selection.

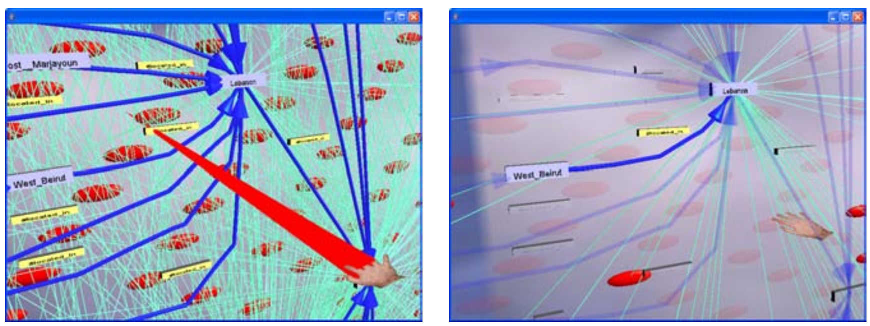

Fig. 5. (Left) Remote object selection using ray casting. A laser bean extends from the user's hand to infinity. The first object that is penetrated by the laser is selected. (Right) After a selection of a relationship, all entities and relationships become semi-transparent but the selected relationship and the attached entities. Additionally, all documents become semi-transparent but the common documents attached to the entities.

Entities, relationships and documents can be selected using the user's virtual laser pointer. The user can activate the virtual laser pointer by touching her/his thumb and index finger when in the *Idle* state. We implemented a ray casting technique for remote object selection (selectable objects are away for the user's hand reach), since this is one of the best techniques for selecting remote objects [9]. The laser beam stays activated until the user initiates a release of the two fingers at which time a selection is performed.

Entities at the leaf level may contain digital media while some entities represent concepts in the ontology and do not have any associated media attached to them. That is, after the selection of an entity, we display the digital media menu (the current state becomes *Media*). The user can then play video and audio clips, and load in images in the VE. To implement this aspect of the application we used the Java Media Framework (JMF).

Selecting a relationship (e.g. an edge) is similar to selecting an entity. However, when selecting a relationship, all nodes in the scene become semi-transparent except the selected relationship, its attached entities and the common documents of the attached entities as shown in figure 5 (right). Selection of a relationship is performed by selecting the (yellowish) label of the line that connects two entities.

By making the non-relevant entities, relationships and documents semi-transparent, we un-clutter the visualization space in order to assist the user in the exploration of related documents and help her/him focus on the detail of a sub-query [16]. Since the non-relevant portion of the environment is still visible but semi-transparent, the user can keep a view of the whole data available, while pursuing detailed analysis of a sub-selection.

5 Conclusions and Future Work

We presented a highly interactive tool, SAV, for visualizing ontologies, metadata and documents or digital media. SAV is based on VR technology and its primary goal is to display the information to the user in a simple and intuitive way within a Virtual Environment. SAV builds a dynamic environment where users are able to focus on different sub-queries (detail) while still keeping them in context with the rest of the environment (overview).

SAV provides an environment where users can select a small set of objects to examine, dynamically and in real-time, providing better contextual information. The users can select and manipulate objects directly and naturally via skills gained in real life as it is described in [11]. Such a system will be of increasing importance for interactive activities in semantic analytics.

Throughout our preliminary studies and demonstrations, users and observers showed a high interest in SAV. Their comments and suggestions will be considered in future releases. Usability studies and experiments are planned to observe how SAV is used and how we can improve upon its interactivity. This will help us discover any further requirements needed to improve the functionality and use of this tool.

6 Acknowledgments

This research was principally supported by "SemDis: Discovering Complex Relationships in Semantic Web" funded by the National Science Foundation (NSF) grant IIS-0325464, projects funded by ARDA and the Office of Instructional Support and Development at UGA. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF, ARDA or UGA.

References

1. Aduna Cluster Map Library version 2005.1 (Integration Guide), 2005.
2. Alani, H., TGVizTab: An Ontology Visualization Extension for Protégé. In *Knowledge Capture, Workshop on Visualization Information in Knowledge Engineering*, (2003).

3. Aleman-Meza, B., Burns, P., Eavenson, M., Palaniswami, D. and Sheth, A.P., An Ontological Approach to the Document Access Problem of Insider Threat. In *IEEE International Conference on Intelligence and Security Informatics (ISI-2005)*, (Atlanta, Georgia, USA, 2005), Springer, 486-491.
4. Aleman-Meza, B., Halaschek, C., Arpinar, I.B. and Sheth, A., Context-Aware Semantic Association Ranking. In *First International Workshop on Semantic Web and Databases*, (Berlin, Germany, 2003), 33-50.
5. Aleman-Meza, B., Nagarajan, M., Ramakrishnan, C., Ding, L., Kolari, P., Sheth, A.P., Arpinar, I.B., Joshi, A. and Finin, T., Semantic Analytics on Social Networks: Experiences in Addressing the Problem of Conflict of Interest Detection. In *15th International World Wide Web Conference*, (Edinburgh, Scotland, 2006).
6. Anselm, S., InfoCrystal: A Visual Tool for Information Retrieval. In *IEEE Visualization Conference*, (San Jose, California, USA, 1993), 150-157.
7. Anyanwu, K. and Sheth, A.P., r-Queries: Enabling Querying for Semantic Associations on the Semantic Web. In *Twelfth International World Wide Web Conference*, (Budapest, Hungary, 2003), 690-699.
8. Anyanwu, K., Sheth, A.P. and Maduko, A., SemRank: Ranking Complex Relationship Search Results on the Semantic Web. In *14th International World Wide Web Conference*, (Chiba Japan, 2005), 117-127.
9. Bowman, D. and Hodges, L., An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments. In *1997 Symposium on Interactive 3D Graphics*, (1997), 35-38.
10. Chen, H., Schuffels, C. and Orwig, R. Internet Categorization and Search: A Self-Orginizing Approach. *Journal of Visual Communication and Image Representation, 7* (1). 88-102.
11. Chuah, M.C., Roth, S.F., Mattis, J. and Kolojejchick, J., SDM: Selective Dynamic Manipulation of Visualizations. In *ACM Symposium on User Interface Software and Technology*, (1995), 61-70.
12. DeFanti, T.A., Brown, M.D. and McCormick, B.H. Visualization: Expanding Scientific and Engineering Research Opportunities. *IEEE Computer, 22* (8). 12-25.
13. Deligiannidis, L. and Jacob, R.J.K., The London Walkthrough in an Immersive Digital Library Environment. In *2005 International Conference on Modeling, Simulation and Visualization Methods*, (2005), 179-185.
14. Deligiannidis, L., Weheba, G., Krishnan, K. and Jorgensen, M., JWSU: A Java3D Framework for Virtual Reality. In *International Conference on Imaging Science, Systems, and Technology*, (2003).
15. Ding, L., Kolari, P., Finin, T., Joshi, A., Peng, Y. and Yesha, Y., On Homeland Security and the Semantic Web: A Provenance and Trust Aware Inference Framework. In *AAAI Spring Symposium on AI Technologies for Homeland Security*, (Stanford University, California, USA, 2005).
16. Eick, S.G. and Wills, G.J., Navigating Large Networks with Hierarchies. In *IEEE Visualization Conference*, (San Jose, California, USA, 1993), 204-210.
17. Ellson, J., Gansner, E.R., Koutsofios, E., North, S.C. and Woodhull, G., Graphviz - Open Source Graph Drawing Tools. In *Graph Drawing, 9th International Symposium*, (Vienna, Austria, 2001), 483-484.
18. Fairchild, K.M., Poltrock, S.E. and Furnas, G.W., SemNet: Three-dimensional graphic representation of large knowledge bases. In *Cognitive Science and its Application for Human-Computer Interface*, (Erlbaum, Hillsdale, NJ, USA, 1988), 201-233.

19. Gansner, E.R. and North, S.C. An Open Graph Visualization System and Its Applications to Software Engineering. *Software - Practice and Experience, 30* (11). 1203-1233.
20. Gene-Ontology. http://www.geneontology.org.
21. Glyco-Ontology. http://lsdis.cs.uga.edu/Projects/Glycomics.
22. Guha, R.V. and McCool, R. TAP: A Semantic Web Test-bed. *Journal of Web Semantics, 1* (1). 81-87.
23. Hammond, B., Sheth, A. and Kochut, K. Semantic Enhancement Engine: A Modular Document Enhancement Platform for Semantic Applications over Heterogeneous Content. In Kashyap, V. and Shklar, L. eds. *Real World Semantic Web Applications*, Ios Press Inc, 2002, 29-49.
24. Janik, M. and Kochut, K., BRAHMS: A WorkBench RDF Store and High Performance Memory System for Semantic Association Discovery. In *Fourth International Semantic Web Conference*, (Galway, Ireland, 2005).
25. Kantor, P.B., Muresan, G., Roberts, F., Zeng, D.D., Wang, F.-Y., Chen, H. and Merkle, R.C. (eds.). *Intelligence and Security Informatics, IEEE International Conference on Intelligence and Security Informatics (ISI 2005)*. Springer, Atlanta, GA, USA, 2005.
26. Kohonen, T., Self-organization of very large document collections: State of the art. In *8th International Conference on Artificial Neural Networks*, (1998), 65-74.
27. Kohonen, T. *Self Organizing Maps*. Springer, Espoo, Finland, 1994.
28. Maedche, A. and Staab, S., The Text-To-Onto Ontology Learning Environment (Software Demonstration at ICCS-2000). In *Eight International Conference on Conceptual Structures*, (Darmstadt, Germany, 2000).
29. Miller, G.A. WordNet: A Lexical Database for English. *Communications of the ACM, 38* (11). 39-41.
30. Polikoff, I. and Allemang, D. Semantic Technology - TopQuadrant Technology Briefing v1.1, 2003.
31. SemDis. http://lsdis.cs.uga.edu/projects/semdis/.
32. Sheth, A.P. Semantic Meta Data for Enterprise Information Integration, *DM Review*, 2003.
33. Sheth, A.P., Aleman-Meza, B., Arpinar, I.B., Halaschek, C., Ramakrishnan, C., Bertram, C., Warke, Y., Avant, D., Arpinar, F.S., Anyanwu, K. and Kochut, K. Semantic Association Identification and Knowledge Discovery for National Security Applications. *Journal of Database Management, 16* (1). 33-53.
34. Sheth, A.P., Bertram, C., Avant, D., Hammond, B., Kochut, K. and Warke, Y. Managing Semantic Content for the Web. *IEEE Internet Computing, 6* (4). 80-87.
35. Sheth, A.P. and Ramakrishnan, C. Semantic (Web) Technology in Action: Ontology Driven Information Systems for Search, Integration and Analysis. *IEEE Data Engineering Bulletin, 26* (4). 40-47.
36. Storey, M.A.D., Noy, N.F., Musen, M.A., Best, C., Fergerson, R.W. and Ernst, N., Jambalaya: An Interactive Environment for Exploring Ontologies. In *International Conference on Intelligent User Interfaces*, (2002).
37. Tu, K., Xiong, M., Zhang, L., Zhu, H., Zhang, J. and Yu, Y., Towards Imaging Large-Scale Ontologies for Quick Understanding and Analysis. In *Fourth International Semantic Web Conference*, (Galway, Ireland, 2005), 702-715.
38. van Harmelen, F., Broekstra, J., Fluit, C., ter Horst, H., Kampman, A., van der Meer, J. and Sabou, M., Ontology-based Information Visualization. In *International Conference on Information Visualisation*, (London, England, UK, 2001), 546-554.
39. Volz, R., Oberle, D., Staab, S. and Studer, R. OntoLiFT Prototype - WonderWeb: Ontology Infrastructure for the Semantic Web, *WonderWeb Deliverable*.