Semantic Network Analysis for Evidence Evaluation:

The Threat Anticipation Initiative

Robin Burk, Mark Davis, Michele Morara, Steve Rust
Battelle Memorial Institute
505 King Ave.
Columbus, OH USA 43201

Alan Chappell, Michelle Gregory, Liam McGrath, Cliff Joslyn
Pacific Northwest National Laboratory
P.O. Box 999
Richland, WA USA 99352

***Abstract*—Semantic network analysis offers a computational method for discovery, pattern matching, and reasoning with large amounts of unstructured, semi-structured and structured information. The Threat Anticipation Platform replaces more cumbersome and computationally complex forms of semantic inference with metrics on graph representations of labeled, directed semantic networked data to identify the degree of evidence within multiple data sources for specified hypotheses about potential events.**

Keywords: semantic; network analysis; distance metrics; event representation

I. Introduction

Semantic networks are of growing interest as information analysis goes beyond traditional structured databases and text-based search queries. These networks take several forms. WordNet [1] and FrameNet [2] are qualitative descriptive linguistic networks that establish the usage of and relationships among word senses in a natural languages such as English. Conceptual graphs [3] and formal ontologies [4] represent knowledge so as to allow computational inferences to be made. Ontologies typically describe semantic hierarchies in terms of concept subsumption and can be extended through the use of description logics [5] which allow relationships beyond subsumption to form the basis for inferencing. Description logics permit both terminological (concept/class relationship) and assertion (specific fact) statements. In the emerging semantic database paradigm [6], such ontologies provide the typing information for another underlying semantic network recording ontologically-compliant predicates. Finally, formal concept analysis offers rigorous methods for deriving ontologies through the use of partially-ordered sets and, specifically, mathematical lattice representations [7].

Although ontologies and description logics offer improved efficiency of knowledge representation and processing compared to propositional and predicate calculus-based knowledgebases, inadequacies in visualization and query yield them unwieldy as the basis for discovery within and reasoning across large amounts of information. This forms the motivation for our application of graph-theoretical analysis of semantic networks.

II. THE THREAT ANTICIPATION PLATFORM

Among potential applications of semantic technologies, intelligence analysis presents some of the most significant challenges. Intelligence analysts must evaluate large amounts of information in order to anticipate potential threats to national security, military operations and other critical capabilities. This information is collected in a variety of formats such as unstructured text documents (Web postings, intelligence reports etc.), structured databases and semi-structured data. Typically it is generated without adherence to a standardized vocabulary or data model. In addition, this information is often of varying provenance and from sources of varying reliability. A key task is to analyze this information, evaluate its (often tacit) meaning and anticipate potential threats as a result of the relationships among entities (people, places and things), and the events in which they participated or might participate.

To support the analysis process, the Threat Anticipation Platform (TAP) brings together a number of semantic capabilities, including innovative approaches to semantic network analysis. TAP offers a broad set of capabilities designed to enable generation of specific hypotheses within a defined class of hypotheses, and their ranking with respect to their evidential support within a variety of sources of prior knowledge for any domain of interest. TAP includes:

- A background knowledge model of the domain of application being assessed, including an ontological description
- A simple, computationally tractable, and human interpretable representation of hypotheses and events
- The incorporation of multiple sources of information, including semantic databases derived from both textual documents and structured databases, domain specific predictive models, and expert opinion

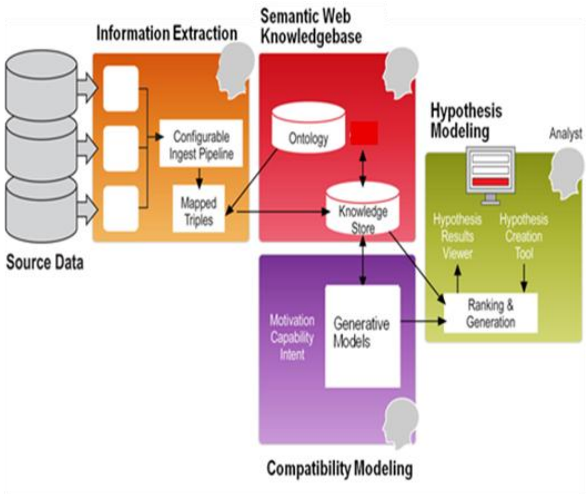

Figure 1. TAP Flow

- The ability to tolerate incomplete, inaccurate, and contradictory data

- The ability to rank alternative specific hypotheses within a hypothesis class based on their level of support within the available information

TAP processing flow consists of several stages (Fig. 1). First, information about entities and events are extracted from sources using natural language and other techniques, as guided by a domain model in the form of an ontology. This information is stored in a knowledgebase, along with meta-information regarding the source of each assertion. The entity and event information is used to generate a semantic network which captures the relationships among entities and between entities and events.

Analysts then specify a hypothesis template, i.e. a set of partially specified hypotheses for which they would like to know the fully specified hypotheses for which the source information provides the best evidence. TAP applies semantic network analysis, along with generative domain models if desired, to identify and rank specific hypothesis instances.

A. *Event and Hypothesis Representation*

TAP includes a novel method for representing hypotheses in a structure based upon the grammar of natural language. In TAP an *hypothesis* is an expression over a set of *events*. An *event* is a verb with a set of syntactic *roles* dependent upon the verb. Each *role* in a specific event or hypothesis is filled by exactly one entity. An example of an event is <u>Attacker</u> attacks <u>Target</u> using <u>Weapon</u>, where the roles are underlined. An hypothesis expression then combines events using Boolean connectors (e.g., And, Or), temporal constraints derived from Allen's interval algebra [8], and constraints on which entities can participate in which roles among events.

An *event class* consists of a verb and a set of *role classes*. A *role class* is a structured description of the set of entities that may participate in that *role*. A participant constraint requires that the participants in a subset of roles either be the same entity or be different entities (e.g., the object in Event 1 must be the same as the Subject in Event 2). A participant constraint may cover roles within one event or may cover roles across multiple events. A *completely specified event* is one where some specific entity is associated with every role.

The TAP event representation has a number of advantages:

- It is sufficiently structured to enable computation.
- It is simple yet expressive.
- It can be translated to relatively natural prose so that users can understand what they have specified

This hypothesis representation scheme is also remarkably flexible. For example, hypotheses about technology proliferation might include event classes such as:

- **Event 1**: ***Nation A*** *acquires* <u>*some capability*</u> *from* <u>*some nation*</u>.
- **Event 2:** ***Nation A*** *develops* <u>*some capability*</u> *using* <u>*some resources*</u>.
- **Event 3:** ***Nation A*** *develops* <u>*some technology*</u> *using* <u>*some capabilities*</u>.

Such events can then be combined to form a hypothesis class: (**Event 1** AND **Event 2**) occur before **Event 3.**

B. *Ontology Management*

TAP uses a formal ontology to guide both information extraction and generation of the semantic network. Although attempts have been made to develop standardized ontologies at the highest levels of abstraction [9], no universal ontologies do (or many would argue, should) exist. However, ontologies may be extended by adding facets for new domain areas. Two issues that arise frequently within knowledge management are ontology evolution and ontology mapping. Over time, an ontology will be modified to capture additional data or to represent existing data more accurately. As it evolves, an ontology may be corrupted as duplicate or overlapping concepts are introduced. Future plans for TAP include algorithmic methods for evaluating an ontology and detecting potential ambiguities and inconsistencies. In addition, because there is no single, universal ontology, related knowledge may exist in separate repositories but be represented using a different ontology. TAP includes methods for mapping one ontology to another, and assessing the results [10].

Finally, TAP includes innovative display and graphical user interface methods to address the severe difficulties associated with user interaction with large ontologies, creating intuitively accessible representations of the ontology for analyst reference [11].

C. *Information Extraction*

TAP includes an automated ability to extract information, including entities and relationships, from libraries of text documents as well as from other data sources, and record them in an RDF semantic graph database. Our novel ontology mapping techniques allow for the association of entities and relations to ontology categories, and thus the construction of an ontologically-compliant semantic graph from such input information. In addition, TAP allows users to correct errors in automated information extraction and to resolve references to the same entity across documents if they are not resolved during the extraction process. Metadata about the provenance and reliability of the information is maintained and can be modified by analysts during specific hypothesis activity.

III. NETWORK ANALYSIS

A portion of the facts within a semantic knowledgebase can be represented as a network or graph, where the entities become vertices or nodes, and the specific relationships between entities (e.g., object properties) become directed edges or links. Such a graph is, in reality, multi-modal or labeled in the sense that each vertex has a type (e.g., person, organization, etc.) and each edge has a type (e.g., employee-of, has-experience-with). While analysis of simple, unlabeled, and undirected networks is a fairly developed field, analysis of labeled, directed graphs is not nearly as mature.

A core TAP capability is to provide a score for the amount of evidential support of events and hypotheses within the semantic graph database. In turn, a major factor of this score is the computation of a semantic distance between entities. Most simple measures of distance only consider the shortest path between pairs of nodes. TAP has multiple distances available. The primary distance used assesses all paths between those nodes, while another allows users to specify the path types (including link type and direction) they are interested in. In this way, TAP can not only factor in the length of the paths, but also their type, number, and topology. TAP uses this semantic distance as a proxy for the likelihood of there being an existing but unknown association or a possible future association between the entities.

TAP supports weighting both the entities (vertices) and the relationships (edges) in the network. This can be used to make the entities or relationships more or less important based upon other criteria such as their semantic type. When such weights are set to zero, it has the effect of removing that entity or relationship from consideration, while retaining it in the knowledgebase.

In addition TAP includes a semantic path type weighting mechanism which effectively replaces a path consisting of a sequence of directed relationships of a certain type with a single, weighted, direct relationship. Performing this replacement throughout the knowledgebase allows the semantic distance calculation to incorporate the semantic information more directly and accurately.

Most simple methods of network analysis are only able to compute the distance between pairs of individual nodes. TAP is able to compute the distance between two sets of nodes and the variance of a set – a measure of how far the entities in a set (e.g., event) are from each other. These abilities enable the scoring of events and hypotheses.

TAP uses these concepts of distance, coupled with spectral analysis of matrices, to cluster the nodes in a network. These methods have proven highly effective on both canonical network analysis datasets and on the networks derived from semantic knowledge bases of terrorist activity. TAP combines clustering with other transformations such as filtering of data by type to create techniques for identifying clusters of activity within graphs with some level of noise (i.e., irrelevant data). Clustering can be used to reduce the size of the problem, allowing other, more computationally intensive algorithms to proceed or allowing manual intervention on a much smaller dataset.

IV. GENERATING AND RANKING HYPOTHESES

As stated earlier, the main purpose of TAP is to rank and generate specific hypotheses within a defined hypothesis class. To do so, TAP relies predominantly upon statistical inference rather than logical inference to select candidate entities to fill specified roles. While logical inference is the traditional method of processing semantic knowledge, logical inference relies heavily upon the accuracy, self-consistency, and, to some degree, the completeness of the knowledge. In particular, in the absence of self-consistency, inference via two-valued propositional logic fails entirely, enabling either no or any conclusions to be drawn. In addition, two-valued propositional logic provides no means to evaluate one hypothesis against another – a hypothesis may only be shown to be possible or impossible. TAP does not treat facts as simply true or false. Rather, facts can be weighted in terms of various desired dimensions including confidence and relevance. TAP scores hypotheses based on a statistical preponderance of evidence. This ability to handle inaccurate, incomplete, and inconsistent data is critical in many domains, including intelligence analysis.

The Hypothesis Ranking and Generation framework within TAP is designed to incorporate multiple sources of prior knowledge (Fig. 2). These include, in addition to the ontology and the semantic network analytic algorithms, several ways in which expert domain understanding can be applied.

A. *Compatibility Assessment*

When humans evaluate hypotheses, they draw on their knowledge of the domain and of the entities and events in questions. Including such knowledge in a computational platform is important but requires that a number of issues be

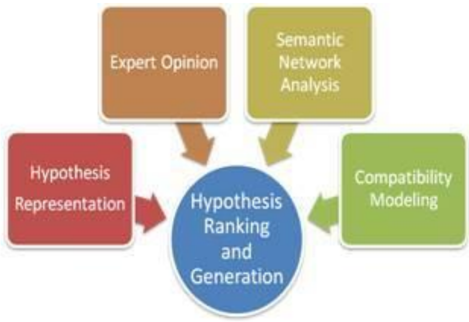

Figure 2. Sources of information incorporated into hypothesis ranking and generation.

addressed. These include:

- *What is the most appropriate representation for specific types of domain knowledge?*
- *How will those representations be populated for a given domain?*
- *How will the resulting formalized knowledge be applied computationally?*

In TAP important domain knowledge is captured in the ontology and embedded in the extraction that populates the knowledgebase. TAP's innovative approach to distance metrics in the resulting semantic graph forms the basis for identifying and ranking potential entity-role assignments, thereby contributing to hypothesis generation and evaluation. TAP allows the definition of participant constraints for event roles. A participant constraint requires that the participants in a subset of roles either be the same entity or be different entities. A participant constraint may cover roles within one event or may cover roles across multiple events. In addition, the TAP semantic network can carry weights on both the vertices (entities) and edges (relationships) to encode analyst judgments regarding the relative importance of the information thus represented, which are in turn factored into hypothesis ranking.

However it is sometimes the case that domain knowledge important for hypothesis evaluation does not fit well into ontological and semantic graph representation. TAP augments semantic graph analysis with domain-specific methods for assessing entity-role compatibility.

Compatibility is a measure of how likely an entity (e.g., a person) is to play a certain role (e.g., attacker) in a certain type of event (e.g., suicide bombing) *relative* to all other entities in the knowledgebase. For instance, compatibility assessment in TAP as applied to the domain of non-state sponsored terrorism provides models that assess motivation, capability, and target value of certain entities. Motivation quantifies the relative likelihood of an actor (e.g., a person or organization) to play any role in an event. Capability quantifies the relative likelihood of an actor to perform a certain action (e.g., how capable is a certain person of constructing a certain type of weapon). Target value quantifies how desirable a target is to a certain class of attackers. An automated, quantitative model for motivation based upon facts in the semantic knowledgebase and upon input from socio-cultural experts specialized in terrorism assigns compatibility values to all of the actors present in the knowledgebase. This compatibility assessment is factored with the semantic network analysis metrics when assigning entities to roles and ranking the resulting specific hypotheses.

Because compatibility assessment draws on domain-specific insights, TAP does not constrain the computational method used for this purpose, requiring only that the method provide a numeric compatibility score. TAP applications can therefore be tailored as closely as desired to a given domain and problem set. Methods for assessing compatibility include

- Manual assignment. TAP allows an analyst or other user to manually assign numeric compatibility values to entities for given events.
- Elicited expert knowledge. TAP includes Bayesian models of the factors associated with means, intent and motivation of non-state actors for various kinds of terror attacks.
- Domain rules. TAP's use of semantic graph metrics for hypothesis ranking provides computational efficiency when compared to the use of description logic within the ontology to infer relationships. However, a limited number of logical rules that capture key domain-specific and commonsense knowledge can improve the accuracy of compatibility assessment while bounding the amount of knowledge required for this purpose.
- Learned models. Where a sufficient volume of representative data exists, probabilistic models of entity-role compatibility can be constructed using machine learning techniques.
- Hybrid approaches. Multiple compatibility assessment methods can be combined in an ensemble.

B. *Other Methods for Incorporating Expert Opinion*

TAP hypothesis scoring algorithms incorporate expert opinion though several direct and indirect paths. In addition to generative compatibility models, expert opinions can be entered as facts into the knowledge base. Analyst expert opinion can also be included by manually setting compatibility and confidence values.

Two features of TAP are critical in the context of expert opinion. First, multiple opinions can be included in the models, facts, and the compatibility and confidence values. Second, the TAP scoring algorithms are robust to contradictory information, evaluating the alternatives based on preponderance of evidence, including expert opinion.

C. Scoring Hypotheses

A TAP hypothesis consists of an expression over a set of events. This expression can include Boolean connectors, temporal constraints, and participant constraints. To score a hypothesis, TAP scores the individual events using a combination of semantic distance and compatibility values to generate a relative likelihood for an event within its event class. TAP then scores the hypothesis by combining the event level scores as dictated by the hypothesis expression.

In addition, TAP computes the marginal relative likelihoods for partially specified events and hypotheses (i.e., where only some roles are associated with specific entities). This includes the ability to determine the relative likelihood of each entity in the knowledgebase to participate in each role in each event in the hypothesis. This expands the types of questions that may be asked to include:

- *Some nation will help **Nation A** to develop **Technology B**. Which nations are most likely to do this?*
- ***Nation A** will develop some technology. Which technologies are most likely?*

The ability to compute marginal relative likelihoods for each entity enables the final, powerful ability to generate lists of the most likely events and hypotheses by assembling combinations of the most likely entities. In this manner, TAP can answer a question such as "***Nation A** will help some nation to develop some technology. What are the most likely combinations of some nation and some technology?*"

V. Summary

The Threat Anticipation Platform provides powerful general hypothesis management and analysis capabilities which can be tailored to specific domains and problems of interest. Hypothesis classes can be easily constructed manually by analysts through a graphical user interface. Alternately, pre-defined hypothesis classes can be applied automatically.

Central to TAP is the construction and analysis of semantic networks of assertions extracted from a variety of sources and formats in conformance with an ontology. TAP is architected to allow the selection of specific semantic network metrics and compatibility assessment measures which best balance computational complexity with accuracy for a given application. Metrics can range from simple characteristics such as path distance or connectedness to more sophisticated and innovative metrics developed for semantic graphs representing highly complex relationships. Similarly, TAP allows a choice among a variety of methods for assessing entity-role compatibility.

TAP's highly modular software architecture facilitates the evaluation of alternative metrics and compatibility assessment methods for a given domain and problem space. Alternatives can be measured for accuracy, completeness and computational efficiency against representative source information. TAP applications can thereby be scaled appropriately to the using organization's needs. In addition, formal evaluation of scoring metrics and compatibility assessment provides confidence in the hypothesis rankings provided by an operational system.

The Threat Anticipation Platform provides a flexible framework for ranking and automated generation of role-based hypotheses using disparate sources of information. In addition, TAP includes components that may help solve constituent problems such as handling of imprecise, incomplete, and contradictory data; construction and management of large knowledge bases; representation of hypotheses; analysis of semantic networks; integration of expert opinion; and quantitative, goal-based modeling of compatibility, including capability, motivation and intent. TAP's modular architecture offers a choice among semantic graph metrics and compatibility assessment methods, ensuring that a tailored and carefully evaluated analytic capability for domains of interest.

References

[1] Project website at http://wordnet.princeton.edu/. Accessed 28 March 2011.

[2] Project website at http://framenet.icsi.berkeley.edu/. Accessed 28 March 2011.

[3] J. Sowa, Knowledge Representation: Logical, Philosophical, and Computational Foundations, Pacific Grove, CA: Brooks/Cole, 2000

[4] T. Gruber, L. Liu and M. Özsu, eds. Ontology. Springer-Verlag, 2008

[5] P. Hitzler, M. Krötzsch, B. Parsia, P. Patel-Schneider and S. Rudolph. "OWL 2 Web Ontology Language Primer". OWL 2 Web Ontology Language. http://www.w3.org/TR/2009/REC-owl2- primer-20091027/. Retrieved 28 March 2011.

[6] M Schmidt, T Hornung, N Kuchlin, G Lausen, C Pinkelal, "An Experimental Comparison of RDF Data Management Approaches in a SPARQL Benchmark Scenario", Proc. 7th Int. Conf. Semantic Web (ISWC 08), pp. 82-97, doi>10.1007/978-3-540-88564-1_6, 2008

[7] B. Ganter, G. Stumme, R.Wille, eds. Formal Concept Analysis: Foundations and Applications, Lecture Notes in Artificial Intelligence, no. 3626, Springer-Verlag, 2005

[8] J. Allen, "Maintaining Knowledge about temporal intervals", Communications of the ACM, v26:11, pp. 832-843, 1983.

[9] Project website for the Suggested Upper Merged Ontology at http://www.ontologyportal.org/. Accessed 28 March 2011.

[10] CA Joslyn, P Paulson, A White: (2009) "Measuring the Structural Preservation of Semantic Hierarchy Alignments", *Proc. 4th Int. Wshop. on Ontology Matching (OM-2009)*, CEUR, v. 551.

[11] A Chappell, A Bladek, CA Joslyn, E Marshall, L McGrath, P Paulson, S Stolberg, A White: (2009) "Supporting the Analytic Knowledge Manager: Formal Methods for Ontology Display and Management", 2009 Conf. on Ontologies for the Intelligence Communty (OIC 2009), CEUR Vol. 555, http://ceur-ws.org/Vol-555/paper3.pdf