On Homeland Security and the Semantic Web: A Provenance and Trust Aware Inference Framework *

Li Ding , Pranam Kolari , Tim Finin , Anupam Joshi, Yun Peng, Yelena Yesha
University of Maryland Baltimore County, Baltimore MD 21250

Abstract

Discovering and evaluating interesting patterns and semantic associations in vast amount of information provided by many different sources is an important and time-consuming work for homeland security analysts. By publishing or converting such information in semantic web language, intelligent agents can automate the inference without compromising the semantics. This paper describes how trust and provenance can be represented/obtained in the Semantic Web and then be used to evaluate trustworthiness of discovered semantic associations and to make discovery process effective and efficient.

Introduction

The advent of the Semantic Web enables distributed publishing mechanisms on the Web as well as provides a large scale distributed knowledge store for computational agents with various purposes. Currently, significant amount of semantic web data (over 47,000,000 triples from over 130,000 web sites[1] is available directly in RDF documents using semantic web vocabularies like Friend Of A Friend (FOAF) [2] and RDF Site Summary(RSS)[3], and indirectly through information extraction tools such as Semagix Freedom(Hammond, Sheth, & Kochut 2002) and IBM's tools(Dill *et al.* 2003).

A promising application domain for the Semantic Web is homeland security, in which suspicious activities or associations among individuals and events must be recognized in millions of everyday reports from thousands of sources with varying trustworthiness, relevancy and consistency. When such reports are published in semantic web languages, human analysts can improve their decision quality and efficiency by using automated tools. Figure 1 shows a simple scenario in homeland security which discovers semantic associations (Sheth *et al.* 2004) among "Mr. X", "Terrorist Group" and "Osama Bin Laden" on the merged RDF graph from several sources. The atomic information unit is statement, e.g. *(Mr. X, isPresidentOf, Company A)*. Statements can be grouped by provenance, e.g. *(Company A, isLocatedIn, US)* has provenance *NASDAQ*. This paper focuses on one type of semantic association which is a simple path linking two RDF nodes in an given RDF graph, e.g. *Mr. X → Company A→ Organization B, → Mr. Y → Osama Bin Laden*.

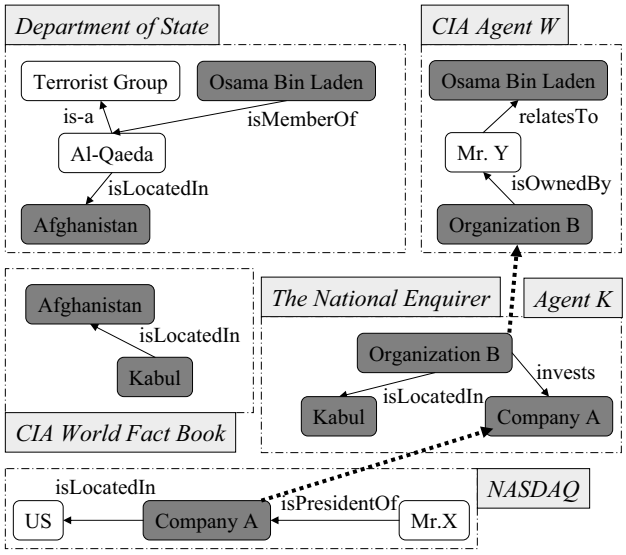

Figure 1: Discovering semantic associations from multiple documents

Automating semantic association discovery and evaluation requires (i) augmenting the Semantic Web by extracting more information from free text reports or databases; and (ii) discovering and evaluating semantic associations in large scale RDF graph. These are also the research objectives of SemDis project[4]. The first requirement has been addressed in (Hammond, Sheth, & Kochut 2002), and the second one has been addressed in (Sheth *et al.* 2004; Aleman-Meza *et al.* 2003) using domain-interest based

*Address correspondence to Li Ding, CSEE,UMBC, 1000 hilltop circle, Baltimore, MD 21250, (dingli1@umbc.edu). Partial support for this research was provided by DARPA contract F30602-00-0591 and by NSF awards NSF-ITR-IIS-0326460 and NSF-ITR-IDM-0219649.

[1]Reported by Swoogle (Ding *et al.* 2004a). More statistics are available at http://swoogle.umbc.edu/.

[2]http://www.foaf-project.org/

[3]http://purl.org/rss/1.0/

[4]http://semdis.umbc.edu

methods. This paper proposes an trust and provenance aware approach to the second requirement.

The Semantic Web has several characteristics which call for a trust and provenance based scheme: (i) data could be published throughout the Web by many sources; and (ii) data could be incomplete or semantically inconsistent. The open nature and space complexity will ultimately restrict us from reasoning on the entire knowledge from the Semantic Web; however, we can effectively prune parts of the graph by disregarding a subset of statements without compromising the final result. This is possible by leveraging the fact that not all of the information is required, much of it is irrelevant, and not all knowledge sources are trustworthy. Even after pruning, the coexistence of inconsistent statements will continue to inhibit logical inference, but trust and provenance support well-known heuristics that can resolve inconsistency through consensus based statements selection.

In the rest of this paper, we describe an inference framework to address the following issues: (i) how to capture provenance and trust information, (ii) how to use provenance and trust to evaluate the trustworthiness of a semantic association, and (iii) how to use trust and provenance to prune search on the Semantic Web.

Provenance and Trust Aware Inference

In order to deal with open and large-scale nature of the Semantic Web, we propose a provenance and trust aware inference framework based on the following assumptions:

- We assume that large amount of semantic web data are available and contains relevant information. We will not discuss how semantic web data is available on the Web through editors and information extraction tools.
- We assume the existence of domain filters that can exclude irrelevant information. Interested readers are encouraged to consult (Ding *et al.* 2004b; Sheth *et al.* 2004) for domain based information filtering. We also assume that domain filters only reason about relevance and are not sensitive to the trustworthiness or completeness of the information they evaluate.
- We assume sufficient information for deriving social and provenance knowledge, for example, social network, personal profile, social reputation, inter-personal trust and associations between web pages and social entities.

Capturing Provenance and Trust

Based on (Ding & Finin 2003), we argue that both provenance and trust are essentially associations among *foaf:Agent*, *rdfs:Statement*, and *wob:RDFDocument*, i.e. statements are created/used by agents and serialized in RDF documents. Although this model is not novel, it is needed but not yet used in the Semantic Web.

We focus on the provenance of statements and RDF documents and several well known provenance relations: (i) *where* – a statement can be physically serialized in one or more RDF documents on the Web, e.g. the statement *(Company A, isLocatedIn, US)* has sourceDocument *http://example.com/nasdaq.rdf* ; (ii) *whom* – a RDF document or a set of statements may be created/published by one or more agents, e.g. *http://example.com/nasdaq.rdf* has creator *NASDAQ*; and (iii) *why* – a set of statements is often copied/derived from several other sources. Provenance shows how statements are organized among agents on the Web; therefore, we may partition the semantic web data using *where-provenance* and *whom-provenance*. With *why-provenance*, users may further track a set of statements' "history of ownership" or "proof trace".

Here, we adopt the truthful and consistent semantics for trust, i.e. a set of statements are trusted by an agent only when they are consistent and believed to be true. We focus on two important associations in trust study: *belief*, which shows an agent's trust in a set of statements, and *trust* which shows an agent's trust in another agent's knowledge. The two associations shares many facets such as 'trustor', 'trust assertion', 'confidence' and only differs in the trusted object and the range of 'confidence' [5].

Belief can be derived from *whom-provenance*, e.g., *NASDAQ* fully believes *(Company A, isLocatedIn, US)* because it is the creator of *http://example.com/nasdaq.rdf* which contains the statement.

Inter-personal **trust** can often be derived in two steps. First, some common-sense rules are employed to extract social relations from raw semantic web data. For example, { *(X is_author_of P), (Y is_author_of P), (X is not Y)* } implies coauthor relation between X and Y. A more complicated social relation is "neighborOf(X,Y)" which we might deduce from { *(X livesIn C),(Y livesIn C), (C rdf:type Street),(X is not Y)*. Second, users may apply personalized rules to infer trust (including confidence value) from the extracted or existing social networks. Moreover, we may use well-known online reputation and ranking systems, such as Google and Citeseer, to derive the default trust in unfamiliar agents.

Trust based Belief Evaluation

Given a set of statements extracted from the Semantic Web, how much should we trust the model they describe? This can be viewed as a problem central to document analysis in which not all information sources are trusted at the same degree and it has obviously important applications in the homeland security domain. How much should we trust the semantic association from "Mr. X" to "Osama Bin Laden" given that "CIA World Fact Book" is highly trusted but "The National Enquirer" is somewhat less trusted. A more complex situation occurs when "Agent K" and "The National Enquirer" have conflicting beliefs over the statement *(Organization B, invests, Company A)* .

We use the following notations: $S = \{s_1, s_2, ..s_n\}$ be a set of n statements, x, y, z be distinctive agents, s be any statement, $tv(x, y)$ be x's confidence over trustworthiness of y's knowledge, $H(x)$ be the set of highly trusted agents by x, $bel(x, S)$ be x's confidence on trustworthiness of S , $creator(s)$ be the set of creator of s, $B(s)$ be the set of agents who has belief over s.

[5] Trust confidence ranges in [0,1] where 0 for fully distrust, 0.5 for ignorance and 1 for fully trust(Ding *et al.* 2004b). Belief confidence ranges from [-1,1] where -1 for fully disbelief, 0 for nonbelief, and 1 for fully belief.

We first examine a simple situation satisfying the following assumptions: (i) statements are independent, semantically consistent, and fully believed by their creators, and (ii) the creators are independent and the investigator z has correct turst knowledge over the them. Equation 1 shows a Bayes model for computing z's overall confidence over S. With the first assumption, we can compute z's confidence over each member of S and then multiply them. With the second assumption, we compute z's confidence over a statement using "Noise-Or" Bayes model (Peng & Reggia 1990).

$$bel(z,S) = \prod_{s_i \in S} \left(1 - \prod_{x \in creator(s_i)} (1 - tv(z,x))\right) \quad (1)$$

For example, given *tv(z,NASDAQ)=0.99*, *tv(z,The National Enquirer=0.5*, *tv(Z, Agent K)=0.6*, *tv(z, Agent W)=0.8*), z's confidence on S0, i.e. the semantic path from "Mr.X" to "Osama Bin Laden", is $bel(z,S0) = 0.99 \times (1 - (1-0.5)(1-0.6)) \times 0.8 \cong 0.63$. This path is much more trustworthy than the cases that only one of "Agent K" and "The National Enquirer" is the creator.

However, statements sometimes could be semantically inconsistent and agents (even the creators) could assert beliefs with various confidence over statements, e.g. two sources believe in different names for the same person with different confidence. A straightforward approach is consensus model which is based on the intuition that trusted peers' beliefs are the only sources of reliable information. Equation 2 averages the discounted belief confidence from trusted agents.

$$bel(z,S) = \prod_{s_i \in S} \left(\sum_{x \in B(s_i) \cap H(z)} \frac{tv(z,x) * bel(x,\{s_i\})}{|B(s_i) \cap H(z)|}\right) \quad (2)$$

Let $s1$ be the statement "Organization B invests Company A". Suppose two agents, i.e. ne(The National Enquirer) and ak(Agent K), have conflicting beliefs, i.e. *bel(ne, {s1})=-1* and *bel(ak, {s1})=1*. The analyst z's final belief confidence depends on her trusts in the two agents: (i) when *tv(z,ne)=0.5* and *tv(z,ak)=0.9*, $bel(z,\{s1\}) = (-0.5 + 0.9)/2 = 0.2$ and (ii) when *tv(z,ne)=0.5* and *tv(z,ak)=0.5*, $bel(z,\{s1\}) = (-0.5 + 0.5)/2 = 0$ In both cases, z has two small confidence on that $s1$ and needs more field investigation reports. The absolute value of $bel(z,s1)$ also shows that the second case should be investigated first.

Trust based Knowledge Expansion

The size of the knowledge base determines the space complexity of inference over it. There are two well-known heuristics that controls knowledge base size: *domain interest heuristic*, which prioritizes knowledge based on domain of interest, and *trust and provenance based heuristic*, which prioritizes knowledge sources by trust. The primary difference is that the former requires services that categorizes and indexes all knowledge while the latter uses P2P trust network to navigate the Semantic Web and incrementally incorporates trusted external knowledge sources. With trust and provenance, an agent A runs semantic web inference (SWInference) as breath first search as shown below. The input includes: the agent A who conduct the inference, the $query$ of the inference, A's trust network TN_A, a trust threshold α and a social distance threshold β.

SWInference(A**,** $query$**,** TN_A**,** α**,** β**):**
1. distance=0, KB={}, Agents ={}
2. if (distance $> \beta$) return;
3. newAgents = findAgents(TN_A, A, Agents, α, distance)
4. if (newAgents is empty) then return fail
5. Agents = Agents $\cup$ newAgents
6. KB = mergeKnowledge(TN_A, A, newAgents, KB)
7. doInference(KB, $query$)
8. if ($query$ is answered) then return with result
9. else distance++ and go to step 2

findAgents finds trusted agents from A's trust network with specified social distance from A. *mergeKnowledge* derives A's combined trust to all trusted agents according to the trust network, and merges the agents' knowledge using mechanisms described in section . We leave the details of these two functions to our previous work (Ding, Zhou, & Finin 2003; Ding *et al.* 2004b). *doInference* essentially runs a conventional inference. Semantic web inference terminates when (i) an answer is found at step 8, (ii) no more trust agents can be found at step 4, or (iii) the social distance limit is reached at step 2. Therefore, the space complexity, which is the union of trusted agents' knowledge within certain social distance, is bounded by α and β. In addition, "doInference" runs at most β times.

Another good feature of *SWInference* is that it assures completeness of high quality data (confidence above α and social distance below β)[6]. With such heuristics, semantic association discovery, which can be abstracted as finding subgraphs in RDF graph, will first use knowledge from the most trusted information sources like "Agent K" and then expand to less reliable sources like "The National Enquirer".

The performance of *SWInference* depends on the correctness of A's trust network. Existing works have shown effective mechanisms in evolving local trust by evaluating results using relevance utility function (Yu, Venkatraman, & Singh 2003) or using consensus based truthfulness utility function (Ding, Zhou, & Finin 2003). Moreover, by adding *domain-interest* facet in trust relation, the trust network may be constructed with respect to the domain interests and results in more efficient, for example, an analyst would not ask *NASDAQ* whether *Al-Qaeda* is a terrorist group despite of *NASDAQ* is highly trusted in stock domain.

Related Work

Provenance has been studied in digital library (e.g. Dublin Core[7]), database systems, (e.g. data provenance (Buneman, Khanna, & Tan 2001) and view maintains (Cui, Widom, & Wiener 2000)) and artificial intelligence (e.g. knowledge provenance (da Silva, McGuinness, & McCool 2003;

[6] Social distance and trust are both important quality factors. Social distance is a good heuristic to control risk, e.g., I may not believe any statement from my best friend's best friend's best friend.

[7] http://dublincore.org/

Fox & Huang 2003) and proof tracing (da Silva, McGuinness, & Fikes 2004)). Our work focuses on provenance in the Semantic Web context, and how it combines with trust network. Provenance network in the semantic web offers multiple-granularity way to group statements.

Trust based belief evaluation has been addressed from information assurance perspective(Hyvonen 2002). Recently, (Golbeck, Parsia, & Hendler 2003; Richardson, Agrawal, & Domingos 2003; R.Guha *et al.* 2004; Ding *et al.* 2004b) have remarked *trust network* as an social alternative. Our work focuses on utilizing trust and provenance in evaluating statements obtained from multiple sources.

The use of trust in reducing search complexity has been mentioned in (Marsh 1994) and realized in (Yu & Singh 2003; Ding, Zhou, & Finin 2003). Our approach focuses on how to create a trust network statically(offline) from the Semantic Web through rule based inference, and how to utilize trust and provenance to prune search space.

Conclusions and Future Work

We have described a provenance and trust aware inference framework including (i) an ontology for representing associations of trust and provenance; (ii) mechanisms to evaluate the trustworthiness of semantic association or any collection of statements obtained from multiple sources, and (iii) a trust based knowledge expansion mechanism that incrementally outsources knowledge from peers to bound the size of knowledge base for inference.

Representing belief is a complex issue since it requires explicit reference of RDF graph. Existing approaches are simply referring the entire RDF graph in an RDF document, RDF reification (Hayes 2004), and Named Graphs (Carroll *et al.* 2004). However, none of them is both efficient and expressive to reference an arbitrary RDF graph. Future work will build an RDF graph reference language.

The two trust based belief evaluation methods in this paper treat a statement as an atom. Future work will study the ontological dependency among statements.

Even with trust based knowledge expansion, serious scalability issues for local inference remain. Future work will find effective tools that support large scale local inference.

References

Aleman-Meza, B.; Halaschek, C.; Arpinar, I. B.; and Sheth, A. 2003. Context-aware semantic association ranking. In *First International Workshop on Semantic Web and Databases*.

Buneman, P.; Khanna, S.; and Tan, W.-C. 2001. Why and where: A characterization of data provenance. In *International Conference o Database Theory (ICDT)*, 316–330.

Carroll, J. J.; Bizer, C.; Hayes, P.; and Stickler, P. 2004. Named graphs, provenance and trust. Technical Report HPL-2004-57, HP Lab.

Cui, Y.; Widom, J.; and Wiener, J. L. 2000. Tracing the lineage of view data in a warehousing environment. *ACM Trans. on Database Systems* 25(2):179–227.

da Silva, P. P.; McGuinness, D. L.; and Fikes, R. 2004. A proof markup language for semantic web services. Technical Report KSL-04-01, Stanford.

da Silva, P. P.; McGuinness, D. L.; and McCool, R. 2003. Knowledge provenance infrastructure. *Data Engineering Bulletin* 26(4):26–32.

Dill, S.; Eiron, N.; Gibson, D.; Gruhl, D.; Guha, R.; Jhingran, A.; Kanungo, T.; Rajagopalan, S.; Tomkins, A.; Tomlin, J. A.; ; and Zien, J. Y. 2003. Semtag and seeker: Bootstrapping the semantic web via automated semantic annotation. In *The Twelfth International World Wide Web Conference (WWW2003)*.

Ding, L., and Finin, T. 2003. Weaving the web of belief into the semantic web.

Ding, L.; Finin, T.; Joshi, A.; Pan, R.; Cost, R. S.; Peng, Y.; Reddivari, P.; Doshi, V. C.; ; and Sachs, J. 2004a. Swoogle: A search and metadata engine for the semantic web. In *Proceedings of the Thirteenth ACM Conference on Information and Knowledge Management*.

Ding, L.; Kolari, P.; Ganjugunte, S.; Finin, T.; ; and Joshi, A. 2004b. Modeling and evaluating trust network inference. In *Seventh International Workshop on Trust in Agent Societies at AAMAS 2004*.

Ding, L.; Zhou, L.; and Finin, T. 2003. Trust based knowledge outsourcing for semantic web agents. In *Proceedings of IEEE/WIC International Conference on Web Intelligence*.

Fox, M., and Huang, J. 2003. Knowledge provenance: An approach to modeling and maintaining the evolution and validity of knowledge. Technical report, University of Toronto.

Golbeck, J.; Parsia, B.; and Hendler, J. 2003. Trust networks on the semantic web. In *Proceedings of Cooperative Intelligent Agents*.

Hammond, B.; Sheth, A.; and Kochut, K. 2002. Semantic enhancement engine: A modular document enhancement platform for semantic applications over heterogeneous content. In *Real World Semantic Web Applications*. IOS Press. 29–49.

Hayes, P. 2004. Rdf semantics (w3c recommendation, 10 february 2004). http://www.w3.org/TR/2004/REC-rdf-mt-20040210/.

Hyvonen, E. 2002. The semantic web – the new internet of meanings. In *Semantic Web Kick-Off in Finland: Vision, Technologies,Research, and Applications*.

Marsh, S. P. 1994. *Formalising trust as a computational Concept*. Ph.D. Dissertation, University of Stirling.

Peng, Y., and Reggia, J. 1990. *Abductive Inference Models for Diagnostic Problem Solving*. Springer-Verlag.

R.Guha; Kumar, R.; Raghavan, P.; and Tomkins, A. 2004. Propagation of trust and distrust. In *Proceedings of the 1st Workshop on Friend of a Friend, Social Networking and the Semantic Web*.

Richardson, M.; Agrawal, R.; and Domingos, P. 2003. Trust management for the semantic web. In *Proceedings of the Second International Semantic Web Conference*.

Sheth, A.; Aleman-Meza, B.; Arpinar, I. B.; Halaschek, C.; Ramakrishnan, C.; Bertram, C.; Warke, Y.; Avant, D.; Arpinar, F. S.; Anyanwu, K.; and Kochut, K. 2004. Semantic association identification and knowledge discovery for national security applications. *Special Issue of Journal of Database Management on Database Technology for Enhancing National Security*.

Yu, B., and Singh, M. P. 2003. Searching social networks. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.

Yu, B.; Venkatraman, M.; and Singh, M. P. 2003. An adaptive social network for information access: Theoretical and experimental results. *Journal of the Applied Artificial Intelligence* 17(1).

Semantic Web researchers, in contrast, accept that paradoxes and unanswerable questions are a price that must be paid to achieve versatility. We make the language for the rules as expressive as needed to allow the Web to reason as widely as desired.Â Two important technologies for developing the Semantic Web are already in place: eXtensible Markup Language (XML) and the Resource Description Framework (RDF). XML lets everyone create their own tagsâ€”hidden labels such as or that annotate Web pages or sections of text on a page. Scripts, or programs, can make use of these tags in sophisticated ways, but the script writer has to know what the page writer uses each tag for.