On the relation between trust on input and reliability of output

Mauro Mazzieri
DEIT
Università Politecnica delle Marche
Ancona, Italy
mauro.mazzieri@gmail.com

Aldo Franco Dragoni
DEIT
Università Politecnica delle Marche
Ancona, Italy
a.f.dragoni@univpm.it

ABSTRACT

We are developing a framework to acquire knowledge from a peer-to-peer network and to evolve it maintaining local coherence and recoverability of previously acquired data.

In this context, we give trust to informations and information sources in a novel way. Informations acquired from the web is given a truth value, based on previous valuations. Each peer deal with them using a fuzzy extension of RDF semantics, apply extra entailment rules derived from an ontology, and deal with inconsistencies using belief-revision techniques. The result is a set of consistent and logically closed data that can be visualized and annotated by human users.

The knowledge shared with other peers has trust metadata on them. Trust values comes both from belief revision process outputs and from user annotations.

Keywords

belief revision, Semantic Web, trust

1. INTRODUCTION

Our reference scenario is that of RDF/S data acquired from a P2P network with a grow-only model: after quiescence, every peer has all the available informations about a shared and agreed topic of interest. An example is the RDFGrowth algorithm[8], where topics of interests are defined by an operator called GUED (Group URI Exposing Definition) and each community agree on a common ontology.

The grow-only model is somehow required by the monotonic nature of the semantic web. Problems arise when the data, according to the ontology, became inconsistent. In this case we locally use a belief-revision technique to restore consistency (section 3). The belief revision use a fuzzy extension of RDF semantics to represent trust values as "truth" values. The assignement of those trust values is explained in section 2.

Finally, section 4 summarize how new trust metadata are attached to information shared with other peers.

Demos and Posters of the *3rd European Semantic Web Conference (ESWC 2006)*, Budva, Montenegro, 11th - 14th June, 2006

2. ACQUIRED KNOWLEDGE

2.1 Trust values

We make a distinction between the **source** of a statement, i.e. the peer who produced the data, and the **peer** from which we actually downloaded the data.

The identity of a peer can be established with a digital signature, either on the document containing data (i.e. an RDF graph) or over singles data clusters (MSGs[9]): the former methos is used by peers, the latter by sources.

The attribution of a trust value to incoming data obeys the following criteria.

Explicit attribution An application-specific property can be used by a source to attach trust values to reified RDF statements, named graphs[2] or MSGs[9]. The actual value is internally weighted with the trust on the source inferred from previous experiences.

Source-based default The trust on statements with no explicit value is equal to the trust on their source.

Peer-based default If the source of a statement cannot be established, it is given a value equal to the trust of the peer from which we downloaded it.

2.2 Knowledge representation

Internally we represent data using a fuzzy semantic extension of RDF and RDF Schema[6, 7].

At the syntactic level, this extension adds a value (i.e. a number) to a triple ⟨subject, predicate, object⟩. The added element has a syntactic nature different from the others: it is not an element of the domain of the discourse, but a property related to the formalism used by the language to represent uncertainty and vagueness.

At the semantic level, the *extension* of a predicate (defined in [5] as a set of couples ⟨subject, object⟩) becomes a fuzzy set.

Analogously, *fuzzy RDF Schema* is based on the definition of a fuzzy *class extension* as a fuzzy subset of domain's elements. Not without some troubles, also domain, ranges, subproperties and subclasses are defined, allowing the representation of simple fuzzy ontologies.

3. THE BRIDGE BETWEEN MONOTONIC AND NONMONOTONIC WORLDS

The semantic web model has a strict monotonic discipline[5] and no global coherence requirement. Howevere, when we give RDF semantic meaning according to an ontology, we can have inconsistencies: for example, two different sources can give a different value to a functional property of a resource.

Our model is to deal with possible inconsistencies maintaining an external grow-only behavior, while internally using belief revision techniques to maintain local coherence.

We make use of a belief revision technique that drops the AGM[1] principle of priority to incoming information, as data come asynchronously from peers. Instead, we use a consistency-based approach[3] to extract a subset of **KB** that is maximal in term of fuzzy cardinality and minimal in term of inconsistence. This maximal subset may not contain the incoming data.

As there is a tradeoff between those requirements, a single parameter is maximazed: the *weighted cardinality value* fuzzy cardinality $*$ (1 $-$ fuzzy inconsistence).

The entire framework uses 3 correlated knowledge bases:

KB, a monotonic (and possibly inconsistent) knowledge base that grows with informations taken from the peer-to-peer network.

B, a maximal consistent subset of **KB**;

Th(B), the deductive closure of **B**, representing the current belief set of the agent, actually visualized by the user interface.

4. OUTPUT RELIABILITY VALUES

4.1 User annotations

The user can interact with the data in **Th(B)**, allowing a finer-grade control on which data is relevant and deserves to be visualized. In a way similar to that of [4], each statement can be either visualized, hidden because it is considered *unreliable*, or hidden because it is considered *irrelevant*. The inferred trust rating on statements is graphically visualized and can be overridden.

As explained in the following, each user interaction causes variations on **KB** statements truth values, and can eventually change the composition of **B** and **Th(B)**.

Moreover, the user can set configuration options to state a self-judgement about his knowledge of the topic and his ability to make judgements.

4.2 Shared outputs

The composition of **Th(B)** and the user annotation on its elements allows to give a feedback to **KB**, updating information metadata before they are shared.

If the user has made explicit his trust on a statement, this statement has in **KB** a truth value equal to user trust, and when it is shared it has the user's digital signature.

Statements with a source signature are transmitted as-is, thus preserving original truth value (unless the user overrided the value). A new value foa a source's reliability is calculated as the ratio between the fuzzy cardinality of statements in **Th(B)** and the cardinality of the set of statements from that source. However, this value is used only internally as a weight for trust value from that source.

The same happens for peers: statements are shared without an explicit trust value; an updated peer reliability value is calculated as the ratio between the fuzzy cardinality of statements in **Th(B)** and the cardinality of set of statement from the peer.

5. REFERENCES

[1] C. E. Alchourrón, P. Gärdenfors, and D. Mankinson. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50:510–530, 1985.

[2] J. Carroll, C. Bizer, P. Hayes, and P. Stickler. Named graphs, provenance and trust. Technical report, HP Labs, 2004.

[3] A. F. Dragoni. Belief revision: From theory to practice. *The Knowledge Engineering Review*, 12(2), 1997.

[4] Y. Gil and V. Ratnakar. Trusting information sources one citizen at a time. In *Proc. of International Semantic Web Conference (ISWC 2002)*, pages 162–176, 2002.

[5] P. Hayes. RDF Semantics. W3C recommendation, World Wide Web Consortium, 10 February 2004.

[6] M. Mazzieri. A fuzzy RDF semantics to represent trust metadata. In *1st Workshop on Semantic Web Applications and Perspectives (SWAP2004)*, Ancona, Italy, 10th December 2004.

[7] M. Mazzieri and A. F. Dragoni. A fuzzy semantics for semantic web languages. In *ISWC Workshop on Uncertainty Reasoning for the Semantic Web*, volume 173 of *CEUR Workshop Proceedings*, pages 12–22, Galway, Ireland, November 7 2005.

[8] G. Tummarello, C. Morbidoni, J. Petersson, F. Piazza, and P. Puliti. RDFGrowth, a P2P annotation exchange algorithm for scalable semantic web applications. In *Mobiquitous*, Boston, MA, August 23–25 2004.

[9] G. Tummarello, C. Morbidoni, P. Puliti, and F. Piazza. Signing individual fragments of an RDF graph. In *Proc. of the World Wide Web Conference (WWW2005)*, 2005.