A fuzzy model for context-dependent reputation

(Trust, Security and Reputation Workshop at ISWC 2004, Hiroshima, Japan.)

Victor S. Grishchenko

Ural State University, gritzko@ural.ru

Abstract. This paper extends the previous work on constructing reputation axiomatics [28] to define context-dependent reputation (e.g. occupation-specific). In short, to employ theoretic-set approach, both notions of reputation and recommendations are downcasted to the basic notion of responsibility. Notion of "a context" is understood as a synonym of "a universe of discourse".

1 Rationale

First, we have to understand the dimensionality of the problem. In the first approach we may expect the number of variables to be equal to the number of participants, n. Every variable represents a public reputation of some participant regarding some fixed topic. Numerous reputation services use this approach, such as e-mail blacklists and whitelists or extensively studied eBay[10] reputation system. There are reputation models which involve explicit or implicit recommendations to calculate n public reputation values, such as EigenTrust algorithm[17] and a family of algorithms based on Markov chains (e.g. PageRank[1]).

Still, there is no evidence that all participants share the same opinions or the same law. So, generally, reputation have to be recognized to be an opinion. This raises dimensionality of a perfect map of reputation to n^2. Straightforward aggregation is meaningless here: 1 billion of Chinese people think that something is good, but my relatives think it is bad. Opinion is meaningful in the context of the owner!

The task of n^2 data storage could hardly be solved with a central database or Distributed HashTable[21]. Considering context issues, it seems natural to let every participant host own opinion and experience, own "map of trust". At the same time we know that most participants can not support a map of size n (in other words, to trace every other participant). Many researchers considered an idea of Web-of-Trust, a graph formed by participant's expressed pairwise directed trust relationships [7, 13, 14, 16, 19, 20, 22, 27]. In this web a trust between distant entities is derived as a function of chains (paths) connecting those entities. M. Richardson et al [15] provide a good formal description of the approach. Another work in this area, Personalized PageRank [5, 4] is discussed in Sec. 2.3.

The purpose of this paper is to introduce a stable balanced reputation propagation scheme on arbitrary topologies. Section 2 introduces definitions and the basic event counting scheme. Section 3 extends the model to the case of interfering fuzzy contexts (the "reputable chef" case). For basic application considerations (i.e. deriving reputation of previously unknown entities from recommendations made by known ones, using coarsened reputation/opinion maps) see [28].

2 Measuring reputation; definitions

2.1 General considerations

Mui [13] describes reputation typology including the following aspects: context, personalization, individual or group, direct or indirect (the latter includes prior-derived, group-derived and propagated). This paper discusses personalized (n^2) reputation regarding some fixed context. Individual and group, direct and indirect flavours of reputation are defined via the basic uniform notion of *responsibility* for elementary events. A reputation context is represented as a set of all relevant (past) events $\mathbb{U}$ plus compliance requirements. Generally, we must consider $\mathbb{U}(t)$ but for the sake of simplicity I will focus on a static model. Propagation of reputation is performed by a social network of recommendations derived from the same notion of responsibility.

An irresponsible recommendation and imposed responsibility (both sound like definitions of a self-assured power) are of no interest to us. There are no distinction between groups and individuals; I use the same word "entities" to describe them.

2.2 Reputation is ...

> A reputation is an expectation about an agent's behavior based on information about or observations of its past behavior. [3]

So, a reputation is based on a responsibility, i.e. an association between events (behavior elements) and entities (agents). A reputation can not exist in anonymized environments. Speaking in terms of the formal model being explained a definition of a reputation is:

Definition 1. *A reputation is an expectation that a compliance of some future event will be near to an average compliance level of past events by the same responsible entities.*

Requirements for compliance are fixed. The simplest example is "the mail (event) is a spam (non-compliant)". So, an elementary (simple) event ε initiated by entity e, $\varepsilon \in \mathrm{E}_e$, may be valued by another entity v as $\rho_v(\varepsilon) \in [0, 1]$

Our compliance expectation on a future event is based on compliance of past events by the same responsible entities, $\rho(\varepsilon) = \rho(\bigcup E_{e_i})$. A reputation of an entity is a compliance expectation on events initiated by that entity: $\rho(e) = \rho(\varepsilon)$. Considering the initiator only (i.e. one fully responsible entity) and assuming events to be of equal value (not distinctively priced) we have

$$\rho_v(e) = \rho_v(\mathrm{E}_e) = \frac{\sum_{\varepsilon \in \mathrm{E}_e} \rho_v(\varepsilon)}{|\mathrm{E_e}|} \tag{1}$$

where $|\mathrm{E_e}|$ is the number of elements (events), $\mathrm{E_e}$ is generally a set of events which affect reputation of e (it is equal to the set of events initiated by e here). We will distinct E_e as a set of known events and $\mathbb{E}_e$ as a set of all such events whether known or unknown to us. (This is the last time I mention $\mathbb{E}$ in this paper.)

2.3 Recommendation: responsibility for other's events

Definition 2. *A recommendation is an expressed opinion of an entity that some another entity is reputable which opinion the recommender is responsible for.*

Full responsibility for an event mean that the event will be included into the entity's relevant event set, thus affecting the reputation of the entity. A reputation of a recommending entity will be affected by any event that affects a reputation of recommended one.

It is useful, if recommendation could be of different certainty ("cautious", fuzzy, $0 < c < 1$), so a weight of a recommended event will be lesser than weights of events initiated by the entity itself (or, another way, the recommended event belongs to the event set of the recommender in a fuzzy way having membership degree $\mu_E = c$). To migrate to fuzzy sets a compliance-of-a-set function (Eq. 1) have to be generalized; it will be equal to a weighted mean (centroid):

$$\rho(\mathrm{E}) = \frac{\sum_{\varepsilon \in \mathrm{E}} c_\varepsilon \rho(\varepsilon)}{|\mathrm{E}|}, \text{ where } |\mathrm{E}| = \sum_{\varepsilon \in E} c_\varepsilon \tag{2}$$

Discounted inclusion $\subset_c$ is an operator further used to express recommendation and, therefore, fuzzy inclusion of a recommended event set into the recommender's set of responsibility. $\mathrm{E}_e \subset_c \mathrm{E}_r$ if entity r recommends entity e with certainty c, so $\forall \epsilon \in \mathrm{E}_e : \mu_{\mathrm{E}_r}(\epsilon) \geq c \cdot \mu_{\mathrm{E}_e}(\epsilon)$. An operation of set discounting $c\mathrm{E}$ is defined as follows: $\mu_{c\mathrm{E}}(\epsilon) = c\mu_{\mathrm{E}}(\epsilon)$. So, $\mathrm{E}_e \subset_c \mathrm{E}_r \Leftrightarrow c\mathrm{E}_e \subset \mathrm{E}_r$ where the subsethood on the right side is the original fuzzy containment by Zadeh: $A \subset B$ is true if $\mu_A(\epsilon) \leq \mu_B(\epsilon)$ for every ϵ. (Discounted inclusion correlates with Goguen implication; generally, it may be understood as a less common denominator statement on a class of entities permitting useful implications for any given entity inside the class.)

Note: this way we define a closure model using multiplication as a concatenation function and maximum as an aggregation function. This combination has a feature of *strong global invariance*[15]. It is different from Personalized PageRank which uses sum for aggregation. This difference may be described as "recommendation" (this model) vs. "voting" (PPR). Practical consequences of this feature still have to be evaluated.

So, what the entity is responsible for? A set of all events that affect a reputation of an entity e was denoted as E_e. It contains events initiated by e (membership degree 1.0) as well as events initiated by recommended entities including those recommended by recommended ones, transitively (membership degree is equal to certainty c or $c_1 c_2 ... c_n$ for transitive cases). Recursive formulae:

$$E_e = O_e \cup \underline{R}_e = O_e \ \cup \ \bigcup c_{er_i} E_{r_i} \tag{3}$$

where
c_{er_i} is a certainty of recommendation of r_i by e;
O_e - a set of events, initiated by e;
$\underline{R}_e$ - appropriately discounted events by entities recommended by e.

So, according to Definitions 1 and 2, we expect events initiated by some known entity e to have compliance level of $\rho(E_e)$ as of Eq. 3.

What about recommended entities? Due to Definition 2 a recommender entity is responsible for events initiated by recommended ones. So, according to Definition 1, a reputation of a recommended entity depends on reputations of recommenders (i.e. events by recommended entities belong to wider sets than own event set of the initiator). Thus, for recommended entities we have:

$$E_e \ = \ O_e \cup \underline{R}_e \cup \overline{R}_e \ = \ O_e \ \cup \ \bigcup c_{er_i} E_{r_i} \ \cup \ \bigcup c_{m_j e} E_{m_j} \ ^{1} \tag{4}$$

where m_j are recommender entities. As a result, "an echo" of an event traverses edges in either direction, because everything that affects a reputation of a recommender, also affects reputations of recommended entities and vice-versa.

3 Contexts

The model explained in Section 2 assumes some fixed context. Taking all possible contexts as a numbered set we may provide corresponding event universes and reputations for every context. Practical applications may require more flexible approach to contexts: an event may relate to a given context at some degree, also different contexts may be semantically close. This section aims to extend the model to fuzzy compliance (relevance, reputation) contexts using the same mathematical apparatus.

First, we have to extend the universe of discourse to all events, independent of context. This extended universe will be denoted as $\mathbb{A}$. Any fuzzy set $\mathbb{U}$ in universe $\mathbb{A}$ is a context.

Taking $\mathbb{U}$ as a universe of discourse, every entity defines a fuzzy subset of compliant events $\mathbb{U}_v^+$ defined by a membership function $\rho_v(\varepsilon)$. As before, E_e is a fuzzy set containing events entity e is responsible for.

[1] As far as I see, this way of discounting of recommenders' event sets does not follow from definitions, e.g. $c^2 E$ does not contradict them also. cE is chosen for the sake of symmetry and balance, to prevent reputation "xeroxes" and "laundries".

Thus, if moving to the universe of discourse $\mathbb{A}$, Eq. 2 changes as follows:

$$\rho_v(E)|_{\mathbb{U}} = \frac{\sum \mu_E(\varepsilon)\rho_v(\varepsilon)}{\sum \mu_E(\varepsilon)} = \frac{\sum c_\varepsilon \mu_{\mathbb{U}}(\varepsilon)\rho_v(\varepsilon)}{\sum c_\varepsilon \mu_{\mathbb{U}}(\varepsilon)} \tag{5}$$

where c_ε is the degree of responsibility, as before; $\mu_{\mathbb{U}}$ is the degree of relevance to the context. So, a weight of an event becomes proportional to its relevance to a context in the case of interfering fuzzy contexts.

The rationale behind Eq. 5 and migration to the narrower universe may be explained as follows: whether we are looking for “reputable chef” (i.e. cooks well) or “chef AND (generally) reputable man” (also pays taxes, etc). For the latter case logical operaions (AND/OR) are enough, while the former needs a move to the universe of cooking.

Understanding context as a semantic domain marked by a character sequence we may model it as a fuzzy set containing relevant events and so, to easily express theoretic-set relations between different domains, i.e. inclusion, equality, disjointless and other, using the relation of discounted inclusion $\subset_c$ and similar tools. E.g. (a simple crisp example) “merinos are sheep”, so a reputable sheep breeder will handle merinos well. A net of such interconnected domains can hardly be called a taxonomy or topological space. So, I suggest the term “indranet” because of some Buddhist connotations. Important feature of indranets (quasi-topological spaces formed by discounted inclusion, union and intersection) is a possibility for any subset to contain all other subsets, at some degree (e.g. due to recommendation graph closure any given entity may be responsible for any event on the planet, *at some vanishing degree*). The previous example of “reputable chef” may be understood as a seeking in the indranet formed by two orthogonal dimensions: occupation taxonomy and binary notion of reputability. General search algorithms for indranets is the author’s future work.

References

1. L. Page, S. Brin, et al: The PageRank Citation Ranking: Bringing Order to the Web, 1998, http://www-db.stanford.edu/ backrub/pageranksub.ps
2. Francisco Botana: Deriving fuzzy subsethood measures from violations of the implication between elements, LNAI 1415 (1998) 234-243
3. A. Abdul-Rahman, S. Hailes: Supporting trust in virtual communities, in Proceedings 3rd Ann. Hawaii Int’l Conf. System Sciences, 2000, vol 6, p. 6007
4. P.-A. Chirita, W. Nejdl, Oana Scurtu: Knowing Where to Search: Personalized Search Strategies for Peers in P2P Networks, SIGIR’04
5. G. Jeh, J. Widom: Scaling personalized web search. In proc. of WWW’2003
6. P. Resnick, R. Zeckhauser, E. Friedman, K. Kuwabara: Reputation systems, Communications of the ACM, Volume 43, Issue 12, 2000
7. Bin Yu, M.P. Singh: A social mechanism of reputation management in electronic communities. in Proc. of CIA’2000, 154-165
8. J. Klensin, RFC 3467 “Role of the Domain Name System (DNS)”, 2001
9. K. Aberer, Z. Despotovic: Managing trust in a P2P information system, in proc. of CIKM’01

10. P. Resnick, R. Zeckhauser: Trust among strangers in internet transactions: empirical analysis of eBay's reputation system, Technical report, University of Michigan, 2001
11. R. Cox, A. Muthitacharoen, R.T. Morris: Serving DNS using a peer-2-peer lookup service, in IPTPS, Mar. 2002
12. L. Mui, M. Mohtashemi, A. Halberstadt: A computational model of trust and reputation, HICSS'02
13. L. Mui, PhD thesis: Computational models of trust and reputation: agents, evolutionary games and social networks, MIT, 2002, http://medg.lcs.mit.edu/ftp/lmui/
14. Bin Yu, M.P. Singh: Detecting deception in reputation management, in Proceedings of AAMAS'03
15. M. Richardson, R. Agrawal, P. Domingos: Trust management for the Semantic Web, in Proc. of ISWC'2003
16. Jennifer Golbeck, Bijan Parsia, James Hendler: Trust networks on the Semantic Web, in Proc. of CIA'2003.
17. S.D. Kamvar, M.T. Schlosser, H. Garcia-Molina: The EigenTrust algorithm for reputation management in P2P networks, in Proceedings of the WWW'2003
18. "Lightweight MTA Authentication Protocol (LMAP) Discussion and Comparison", John Levine, Alan DeKok. Internet Draft, 2004
19. Jennifer Golbeck, James Hendler: Reputation network analysis for email filtering, for the 1st Conf. on Email and Anti-Spam, 2004
20. Jennifer Golbeck, James Hendler: Inferring reputation on the semantic web, for EKAW'04.
21. Michal Feldman, Antonio Garcia-Martinez, John Chuang: Reputation management in peer-to-peer distributed hash tables, http://www.sims.berkeley.edu/∼mfeldman/research/
22. R. Guha, R. Kumar, P. Raghavan, A. Tomkins: Propagation of trust and distrust, in Proc. of WWW'2004
23. A. Fernandes, E. Kotsovinos, S. Ostring, B. Dragovic: Pinocchio: incentives for honest participation in distributed trust management, in Proceedings of iTrust'2004
24. Philipp Obreiter: A case for evidence-aware distributed reputation systems, in Proceedings of iTrust'2004
25. Paolo Massa, Bobby Bhattacharjee: Using trust in recommender systems: an experimental analysis, in Proceedings of iTrust'2004
26. Radu Jurca, Boi Faltings: Truthful reputation information in electronic markets without independent verification, icwww.epfl.ch/publications/documents/IC_TECH_REPORT_200408.pdf
27. Web-o-Trust effort by Russell Nelson, http://www.web-o-trust.org/
28. Viktor S. Grishchenko: Redefining Web-of-trust, may be found at http://blogs.plotinka.ru/gritzko/accounting2.pdf
29. W3C recommendation: OWL Web Ontology Language Overview, http://www.w3.org/TR/owl-features/