Provenance as a Key Factor for Privacy-proof Trust

Davide Ceolin
VU University Amsterdam
de Boelelaan, 1081a
1081HV Amsterdam
The Netherlands
d.ceolin@vu.nl

ABSTRACT

Provenance is a key element to address some of the limitations of reputation systems, including privacy-related issues. In this paper, we outline these limitations and we describe the role of provenance in their mitigation.

CCS Concepts

•**Information systems** → **Personalization;**

Keywords

Provenance; Trust

1. INTRODUCTION

Trust and privacy are two tightly coupled concepts. The more we trust, the more we could be inclined to give away part of our privacy, in return for something. At the same time, to trust somebody, we may need to acquire information about her that violate her privacy. In fact, we might think that the more we know about a person, the better we can decide whether to trust her or not. The second scenario is the one we are tackling in this position paper. We propose to use provenance as a means to mitigate such conflict, thus preserving as much privacy as possible, while increasing the information available to form a trust judgment.

The rest of this paper is structured as follows. Section 2 introduces related work. Section 3 briefly describes reputation systems and their limitations. Section 4 describes how provenance can mitigate those limitations, and Section 5 concludes the paper.

2. RELATED WORK

Elsweiler [6] draws a parallel between provenance-based recommender systems and search systems. We can further extend such a parallel to provenance-based recommender systems: although their focus is different (handle user-provided content, rather than presenting content to the users), all these systems handle and model users based on behavior analysis through provenance. Another example of a behavior-based recommender system is presented by Hasler [7].

Reputation systems are widely studied and employed. Masum and Tovey [10] provide a comprehensive review of the field. The use of provenance as a basis for trust has been outlied in a seminal paper of Carroll et al. [1]. An example of implementation of those concepts is the work of Golbeck [8], who combines provenance and trust for semantic content filtering. Other examples of synergies of trust and provenance are represented by the work of Martinez et al. [9], who trace the provenance of crowdsourced works in the context of the CrowdTruth framework; by the work of Ebden et al. [5], who use provenance graphs network analysis as a basis for trust estimation; and by previous works of ours [4, 2], where we use so-called "provenance stereotypes" as a basis for trust prediction.

3. REPUTATION SYSTEMS

Reputation systems usually rely on a specific class of provenance information, that describes who created a given artifact. These systems collect evidence about user performance over time from one or more evaluators and summarize this evidence in a value representing the perceived (and expected) future user performance. This value is, in fact, the user reputation. User reputations are hence often represented by means of probabilistic values, although also categorical and ordinal values can be used. Evidence can be collected, for instance, in terms of counts of positive/negative past performance or in terms of judgments.

We identify the following limitations in this systems:

Cold start problem. This is a very well-known problem for user modeling systems and recommender systems. Whenever a new user joins the system, it is necessary to acquire a minimum amount of knowledge about her before being able to estimate her reputation. Before this amount of evidence is gathered, the user reputation is unknown or possibly unreliable. In fact, no or too few evidence can hardly be used to make estimates that are truly representative of the reputation.

Privacy intrusion. User reputations are often based on analyses of the user performance over time. However, when available, profile data could be used as well, in order to improve the accuracy of the reputation. Information like age, education level or gender could provide the basis for trust assessments. This information can be valuable as demonstrated in a previous work

Weaving Relations of Trust in Crowd Work: Transparency and Reputation across Platforms '16 Hannover, Germany

DOI: 10.1145/1235

of ours [3]. However, it is also potentially privacy-infringing and discriminatory, because it leads to trust assessments that are based on intrinsic user characteristics.

Inaccurate Point-wise Predictions Reputations are based on evidence sets that are (hopefully) representative of user performance. Hence, reputations are often implemented by means of expected probability to observe a positive performance by the user. This approach is reasonable (and often accurate) in the long run: probabilities represent asymptotic ratios between positive and negative evidence. However, when it comes to deciding whether to trust or not a given artifact on the basis of its author's reputation, then the chances of errors increase. For example, suppose that user x has a reputation of 0.9 (on a [0,1] scale). This means that we suppose that she is trustworthy 90% of the times. Suppose that x provides a new contribution as part of a crowdsourcing task. If we consider x reputation high enough to trust her contribution, we have 10% chance to be wrong. And, if we apply this policy for all x contributions, then, in the long run, we will have 10% of false positives. If we decide to discard 10% of x contributions, we need to clearly identify those untrustworthy: otherwise, we risk to add 10% of false negatives to the 10% of false positives (worst-case scenario).

4. PROVENANCE-BASED SOLUTIONS

Provenance-based reputation systems aim at detaching from a user-centered view when deciding whether to trust something or not. These systems use all the provenance information available to take behavior-based decisions. Provenance is usually recorded in terms of traces that describe the sequence of actions that led to a given artifact (e.g., an annotation) together with metadata like timestamps.

Provenance-based reputation systems allow to mitigate the **cold start problem** because, when provenance tracing is put in place, the availability of provenance traces will increase as soon as the system will be used by any user. This means that the information on which the reputation system is based has a higher availability than in the user-centered case. Therefore, the chances of being short of evidence are lower. For the same reason, this approach helps also tackling the problem of **privacy intrusion**. In fact, higher availability of provenance traces reduces the need for additional data, including profile data. Moreover, we note that in a case study analysed in a previous work of ours [4, 2], user characteristics correlate with provenance stereotypes. In other words, on statistical bases, similar users tended to follow regular patterns in their behaviors. This means that, when this is the case, by relying on provenance stereotypes, we implicitly rely on user-based reputation. However, since we do not know which class of users adopts a given provenance stereotype (i.e., we do not know who are the users who behave in a given matter), user privacy is preserved. Lastly, since provenance is a multi-dimensional class of information (since it captures 'who', 'when', 'how', 'where') data came to be, and since provenance-based evidence has a higher availability than user-centered evidence, **point-wise predictions** have higher chances to succeed. In fact, by combining evidence from multiple sources, we can come up with accurate predictions, as we showed in the cited works.

5. CONCLUSION

Using provenance for trust assessment is a challenging task: provenance is fine-grained, complex, multi-dimensional. However, provenance allows addressing some of the limitations of reputation systems, like the cold-start problem, privacy intrusion, and inaccurate point-wise predictions. Therefore, we see in the efforts for properly handling provenance (e.g., via stereotyping or clustering) one pivotal step to benefit from crowdsourcing without compromising privacy. Using standardized provenance modeling strategies (e.g., PROV [11]), in the future, we aim also at identifying a common ground of platform-independent provenance-based trust insights.

6. ACKNOWLEDGMENTS

This work was supported by the Amsterdam Academic Alliance Data Science (AAA-DS) Program Award to the UvA and VU Universities.

7. REFERENCES

[1] J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler. Named graphs, provenance and trust. In *WWW '05*, pages 613–622. ACM, 2005.

[2] D. Ceolin, P. Groth, V. Maccatrozzo, W. Fokkink, W. R. van Hage, and A. Nottamkandath. Combining user reputation and provenance analysis for trust assessment.

[3] D. Ceolin, P. Groth, A. Nottamkandath, W. Fokkink, and W. R. Hage. *Uncertainty Reasoning for the Semantic Web III*, chapter Analyzing User Demographics and User Behavior for Trust Assessment, pages 219–241. Springer, 2014.

[4] D. Ceolin, A. Nottamkandath, and W. Fokkink. Efficient Semi-automated Assessment of Annotation Trustworthiness. *Journal of Trust Management*, 1:1–31, May 2014.

[5] M. Ebden, T. D. Huynh, L. Moreau, S. Ramchurn, and S. Roberts. Network analysis on provenance graphs from a crowdsourcing application. In *IPAW 2012*, pages 168–182. Springer, June 2012.

[6] D. Elsweiler. Behaviour with search and recommender systems: What can it tell us? In *CARR '14*, pages 1–1. ACM, 2014.

[7] A. Geyer-Schulz, M. Hahsler, A. Neumann, and A. Thede. Behavior-based recommender systems as value-added services for scientific libraries. In *Statistical Data Mining & Knowledge Discovery*, pages 433–454. Chapman & Hall / CRC, July 2003.

[8] J. Golbeck. *IPAW 2006*, chapter Combining Provenance with Trust in Social Networks for Semantic Web Content Filtering, pages 101–108. Springer, 2006.

[9] C. Martinez-Ortiz, L. Aroyo, O. Inel, S. Champilomatis, A. Dumitrache, and B. Timmermans. Provenance-driven representation of crowdsourcing data for efficient data analysis. In *e-Science*, pages 300–303. IEEE, 2015.

[10] H. Masum and M. Tovey, editors. *The Reputation Society*. MIT Press, Boston, MA, USA, Feb. 2012.

[11] W3C. PROV-O. http://www.w3.org/TR/prov-o/, 2013.