

Social Network Analysis for Trust Prediction

Davide Ceolin¹ and Simone Potenza²

¹ d.ceolin@vu.nl

Computer Science Department, Vrije Universiteit Amsterdam
de Boelelaan, 1081a, 1081HV Amsterdam, The Netherlands

² simone@konnektid.nl

Konnektid

Herengracht 182, 1016 BR Amsterdam, The Netherlands

Abstract. From car rental to knowledge sharing, the connection between online and offline services is increasingly tightening. As a consequence, online trust management becomes crucial for the success of services run in the physical world. In this paper, we outline a framework for identifying social web users more inclined to trust others by looking at their profiles. We use user centrality measures as a proxy of trust, and we evaluate this framework on data from Konnektid, a knowledge-sharing social Web platform. We introduce five metrics for measuring trust. Performance achieved an accuracy between 43% and 99%.

1 Introduction

From renting a taxi for cheap to attending free online academic-level courses, the sharing economy has shown an incredible expansion in the recent times. Most - if not, all - of these successful sharing economy services rely on facilities offered through the Social Web. So, besides the “sharing” aspect, these companies are also a clear example of “online-to-offline” (O2O) services. In fact, these services exploit online social interaction to achieve the ultimate goal of offline exchanges (e.g., product or service sharing). In this scenario, online trust management plays a crucial role, because trust is a necessary precondition for users to rely on these services. Often times the actual achievement of the physical engagement depends on the trust that the user places in the online platforms and in the other peers that the user gets in touch with through the platforms themselves.

In this paper, we outline a framework for estimating user trust. This framework uses user features (extracted from her demographics, knowledge and network centrality) to predict trust by means of classification algorithms and, in general, of machine learning. We provide an overview of the framework, and we analyse the effectiveness of the network centrality part, evaluating it on a proprietary dataset provided by Konnektid. Konnektid is a knowledge-sharing online platform that allows users to share knowledge with peers. In Konnektid, users declare which subjects they wish to learn and to teach. Based on their own initiative, or through recommendations from the platform, users connect with their peers and, in case “teacher” and “student” match their needs and

wishes, they meet in person. This is the ultimate goal of the platform. In this scenario, we investigate the use of five network centrality measures (degree centrality, betweenness centrality, eigenvector centrality, closeness centrality and communicability) to estimate trust, which is represented by means of five different metrics (activity count, social activity count, average activity frequency and count of accepted requests and appointments), and computation is performed by means of classification algorithms (Support Vector Machines and Naïve Bayes). This provides us with a first evaluation of the framework. The rest of this paper is structured as follows. Section 2 describes related work. Section 3 describes our approach, which evaluation is presented in Section 4. Section 5 concludes.

2 Related Work

We refer to trust as 'Firm belief in the reliability, truth, or ability of someone or something' [13]. Sabater and Sierra [14], Artz and Gil [1], and Golbeck [6], provide extensive surveys of trust management models for trust in Computer Science, Social Web, and Web respectively. Sherchan et al. [15] present a survey of trust in social networks. Wu and Chiclana [17] make use of social network analysis to group decision-making problems. Despite the approach similarity, their ultimate goal is reaching consensus in decision-making, while we aim at identifying users that are more prone to trust.

We do not incentivize specific user behaviors, neither create any mechanism for handling and sharing user reputations, but the fact that we estimate trust based on user network centrality implicitly relates to reputation systems. Masum and Tovey [10] and of Golbeck [5] provide extensive analyses of this topic.

Kolaczek [8] proposes a method for estimating trust in social networks, but his focus is on autonomous multi-agent systems while we focus on human-based social networks. Similarly, Di Cagno and Sciubba [2] analyze the impact of trust in social networks, but they focus on lab-created networks instead of real-world data, as we do. Nepal et al. [11] consider an aspect of "social capital" built by users over time when estimating trust. We implicitly aim at fostering the creation of such capital with our work. Grabner-Kräuter and Bitter [7] propose a multi-faceted approach to trust analysis, distinguishing between individual and global aspects of trust. We make a distinction between global (e.g., network centrality) features and individual aspects (e.g., a user's decision to accept an appointment), and we will deep this separation in the overall framework (see Section 3). Lastly, this work can provide the basis for advanced uses of social network information in recommender systems [16] and quality assessment [3].

3 Approach

Our goal is to identify which user features correlate with user trust. First, we must identify useful user features. Second, we need to quantify (or estimate) trust. Lastly, we need to identify reasoning algorithms to link features and trust.

3.1 User features

We identify three classes of user features useful to this aim:

User demographics. The propensity of users to engage in socializing and in other cooperation and interaction activities might be affected by their demographic profile. For example, younger users might be more inclined to participate, or this inclination could be influenced by cultural factors which could be, in turn, correlated with the nationality of the user.

User knowledge profile. Demographics characterize the user with respect to the population she belongs to. These characteristics are often not decided by the user (e.g., age), and are either immutable or subject to slow changes. A useful user profile can be built also based on the knowledge that the user demonstrates, her tastes, and the knowledge that the user wishes to acquire, thus inducing more dynamics (e.g., user tastes need to be updated periodically). Also, this profile depends on the platform: in some, skills are more important (e.g., knowledge-sharing platforms), in others (e.g., media-sharing platforms), users tastes are more relevant.

User network centrality Social network users interact with other peers. Their network centrality can be measured in diverse manners: degree centrality, betweenness centrality, etc. These measures provide an indication of with how many users a given user interacts, whether a given user links different parts of the whole social network that would be disjoint otherwise. Intuitively, we suppose that the higher the network centrality of a user is, the higher is her tendency to trust and interact. However, which centrality measures better indicate trust and how strong such a correlation is, needs to be investigated.

3.2 Trust Measures

Trust is a belief that somebody shows with respect to something or somebody in a given context [12]. Since we situate in the realm of Social Web apps, we identify two main subjects of trust, namely the app itself and other users (which the app allows getting in touch with).

Trust in the App Trust in the app and in the service provider are a necessary precondition for the user to join a Social Web app. This implies trust in how user personal information is dealt with, and trust in the app behavior and its functionality. This prerequisite is the basis for building user engagement. Users do hardly engage with Social Web platforms they do not trust, especially when these are aimed at creating contacts in the real world (as in the case of O2O). We use engagement indicators like the number of user accesses as trust proxies.

Trust in Other Users Trust in other users is the key aspect of Social Web apps. While trust in the app is a precondition for the user to utilize it, trust in other users is the requirement for users to join the app. Social Web apps

are meant to enhance and facilitate user interaction, thus relying on trust to be established. Measuring trust in users is important, for example, to identify users whose engagement needs to be fostered by means of recommendations or other actions. Depending on the platform, user trust can be estimated based on the number or frequency of interactions that a user has with others. In O2O apps, user trust can be measured by user acceptance of real-world transactions.

3.3 Reasoning Algorithms

Having identified the possibly relevant user features, and having identified proxies for trust, then we will use machine learning algorithms for identifying correlations between them. We prefer classification algorithms since we treat trust metrics as qualifying classes. So, we employ the Support Vector Machines and Naïve Bayes algorithms. Alternative approaches (e.g., to improve computational performance) will be evaluated when we will extend our framework.

4 Evaluation

4.1 Dataset Description

We perform a preliminary evaluation of our approach on a dataset of user interactions by Konnektid [9] consisting of the logs of 37,423 user actions performed between September 2012 and August 2015. The only personal information present in this dataset is anonymous user identifiers. Actions are classified as:

ProfileRegistered, ProfileUpdate To access the platform, users register their profile (which contains both demographics and indications about what they wish to learn and to teach). Profiles can be updated by users anytime.

DirectRequest, NeighbourRequest, GroupRequest. Users can issue requests to learn particular skills. These requests can be directed to selected users, or broadcasted, also to the neighboring users (geolocated).

DirectMessageSent Users can exchange textual messages.

AppointmentCreated, AppointmentUpdated, AppointmentAccepted.
The goal of the app is to facilitate user encounter, in person, to let them teach something each other.

Graph Description We model the social graph of Konnektid as follows. Each node of the graph is represented by a user. Each edge represents any possible kind of interaction occurred among users. In this manner, we model user interaction, without focusing on its quality or frequency, but merely from the “social” point of view. We will consider different kinds of graphs in the future.

4.2 Network Centrality Features

On the graph described above, we calculate the following five network centrality measures to be used as features for trust prediction in this setting.

Degree Centrality The degree centrality of a node is equal to the degree of that node, i.e., to the number of edges that connect that node.

Closeness Centrality The closeness centrality of a node is the reciprocal of the sum of the distances between that node and all the other nodes.

Betweenness Centrality The betweenness centrality of a node counts how many times it acts as part of the shortest path between two nodes.

Communicability Centrality This is the sum of closed walks of all lengths starting and ending at a given node. This is defined as: $CC(i) = \sum_{j=1}^N C_{i,j} = [e^A]_{i,j}$, where i, j are nodes and A is the adjacency matrix [4].

Eigenvector centrality This computes the centrality for a node based on the centrality of its neighbors. The eigenvector centrality for node i is $\mathbf{Ax} = \lambda \mathbf{x}$ where A is the adjacency matrix of the graph G with eigenvalue λ .

4.3 Trust in the Platform

Trust in the platform is estimated based on the user activity. Trust is necessarily tangled with other user attitudes, like user engagement, and user preferences. Even if it is not possible to discern the influence of trust on user activities, trust is necessarily their prerequisite: users interact with the platform because, consciously or not, they trust it. Trust is present in any other interaction of the users with any other platform. However, in this case, the platform is a means to interact with strangers that users will decide whether to encounter or not. Hence, trust in the platform implies trust in its ability to preserve privacy and in its ability to identify potentially interesting encounters. We propose the following measures as proxies for this type of trust:

Count of Activities The first measure of user interaction is given by overall the count of user activities. This measure corresponds to the degree centrality computed on a graph representing all the interactions performed among users, while our graph of interest is unweighed and undirected, and represents any kind of interaction among users, regardless of their frequency or type.

Results We run the Support Vector Machine (SVM) algorithm with Stochastic Gradient Descent (SGD) preprocessing to predict the number of activities of each user, treating this problem as a classification problem (so to predict the “1-activity users”, the “2-activities users”, etc.). We evaluated SGD-SVM with 10-fold cross validation, obtaining 43% accuracy (there are 67 different classes in total, i.e., users have 67 different numbers of actions performed). Accuracy is computed as the percentage of correctly classified items. Accuracy rises to 84% when we group actions in groups of 5 (i.e., users who performed between 0 and 4 actions fall into the same class, etc.)

Count of “social activities” Users can perform different activities on the Social Web app. Besides the fact that all these activities are meant to facilitate social interaction, only some of them actually involve other users. For example, a

user might decide to update her own profile in order to be more easily contacted, but this action does not directly involve other users. This measure is equivalent to the degree centrality computed on the network reporting only the following activities: message sending, offer sending, requests sending, appointment making and updating.

Results We employed SGD-SVM also in this case, and we evaluated it by running 10-fold cross-validation also in this case. We obtain an accuracy of 67%, which reaches 92% by grouping the counts of social actions in classes modulo 5.

Weighed Activity Frequency The count of activities is a possible indicator of user interaction with the platform. However, this indicator does not take into account the time span of this interaction: a user might perform a high number of activities in a limited period of time, and then disappear. Or, she could demonstrate trust and engagement in the platform by participating frequently. So, as another measure of trust, we propose a weighed measure of user frequency. On the one hand, in fact, we value frequent user activities. On the other hand, we ‘penalize’ users who do not return to the platform. We define the measure in such a manner that it ranges from 0 (no trust) to ∞ (full trust). Also, we define this measure so to take a specific point of view that corresponds to a specific time instant t : to decide whether a user u ‘disappeared’ for a long period of time, we must be sure that a long period of time occurred between our observational point and her last appearance. The resulting metric is defined as:

$$weighed_freq(u, t) = e^{-\frac{t(u)_{last} - t(u)_{first}}{\#activities(u)}} * e^{-(t - t(u)_{last})}$$

Results SVM with 10-fold cross validation, reaches 94% accuracy.

4.4 Trust in other Users

Here we define metrics for estimating the trust users express in other users. These metrics are computed from the logs of user activities in the platform.

Count of accepted requests Users receive requests from other users. We count how many times each user reacts to a request with an offer. This count is affected by the user “good-will”, by the fact that she is interested in the content of the offers received, as well as by the intention of the user to trust the requester: ultimately these offers should lead to meetings in person.

Results SGD-SVM with 10-fold cross-classification achieves 99% accuracy in this case. All the users receive requests because, besides those issued by other users, the system itself periodically sends requests, in an attempt to facilitate encounters. However, the longevity of users is likely to be linked to the number of requests received, and thus it could make sense to analyze also the ratio between the number of requests received and the number of offers made.

Count of accepted appointments The second measure of trust in other users that we propose to adopt is the number of appointments a given user accepted.

Results SGD-SVM with 10-fold cross-validation achieves 99% of accuracy in this case, but this is due also to the sparsity of appointments accepted with respect to the total counts of activities (indeed, these correspond to about 1% of the activities). More interesting, in this case, is the recall. Given the sparsity of the data targeted, we can sacrifice part of the precision of the results in order to identify a large enough set of candidate users that comprises most of the users who actually accepted an appointment. Recall of SGD-SVM is, in fact, 14%. In this case, we run also Naïve Bayes as an alternative classification algorithm. This allows us still achieving 99% but with 55% recall.

5 Discussion

This paper introduces a framework for predicting trust in Social Web apps. In particular, in this framework, we analyze the use of network centrality measures to predict trust that users show in the platform and in other users. Our analyses show that interpersonal trust is well-captured by user centrality: the more central a user is, the more prone to trust others he will be. This is useful, for instance, to identify users to recommend to newcomers, in order to increase the likelihood of positive outcomes of interactions. Also, there is a clear link between trust in the platform (and, hence, engagement), and user centrality. This link is weaker than interpersonal trust, but still identifies in the number of diverse network links one possible motivation for user engagement. In the platform that we analyze, user engagement and interpersonal trust are tightly bound because of the nature of the task performed: users interact with the platform in order to interact with other users. These results could hence be expected. However, they show that, besides the fact that users are motivated to use the platform because of already-established acquaintances, the creation of new links and their diversification are important factors to consider to foster quality interaction. In fact, diverse centrality measures focus on different aspects of connectivity, from the mere number of connection (like in the case of degree centrality) to the ability to connect diverse groups of users (like in the case of betweenness centrality).

In the future, we will develop further this framework in all its three main components: features, trust measures, and reasoning algorithms. We will expand the set of centrality measures considered and add in the computation also demographics, knowledge features (as defined in Section 3), and possibly other classes of features, as the current selection is heavily driven by the case study at our disposal. Also, we aim at investigating further the trust metrics proposed, in order to extend them, as well as to identify relations between them (e.g., one trust metric might be highly correlated with others; this kind of information is useful to increase computation performance). Lastly, we will consider other prediction algorithms. For example, besides the classification angle taken in this paper, given that user actions occur sequentially, it might be useful to model them in

terms of Markov chains, to predict whether the sequence of actions (rather than the set of action) performed by a user provides indications for trust.

Acknowledgements This work has been partially funded by the Dutch national program COMMIT under the Big Data Veracity project and by a research voucher awarded by the Network Institute of the Vrije Universiteit Amsterdam.

References

1. D. Artz and Y. Gil. A survey of trust in computer science and the semantic web. *Journal of Semantic Web*, 2007.
2. D. D. Cagno and E. Sciubba. Trust, trustworthiness and social networks: Playing a trust game when networks are formed in the lab. *Journal of Economic Behavior & Organization*, 75(2):156 – 167, 2010.
3. D. Ceolin, J. Noordegraaf, L. Aroyo, and C. van Son. Towards web documents quality assessment for digital humanities scholars. In *ACM WebSci 2016*, 2016.
4. E. Estrada and N. Hatano. Communicability in complex networks. *Phys. Rev. E*, 77:036111, Mar 2008.
5. J. A. Golbeck. *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, 2005. AAI3178583.
6. J. A. Golbeck. Trust on the World Wide Web: A Survey. *Foundations and Trends in Web Science*, 1(2):131–197, 2006.
7. S. Grabner-Kräuter and S. Bitter. Trust in online social networks: A multifaceted perspective. *Forum for Social Economics*, 44(1):48–68, 2015.
8. G. Kołaczek. *Agent and Multi-agent Technology for Internet and Enterprise Systems*, chapter Social Network Analysis Based Approach to Trust Modeling for Autonomous Multi-agent Systems, pages 137–156. Springer, 2010.
9. Konnektid. Konnektid. <http://www.konnektid.com>.
10. H. Masum and M. Tovey, editors. *The Reputation Society*. MIT Press, Boston, MA, USA, Feb. 2012.
11. S. Nepal, W. Sherchan, and C. Paris. Strust: A trust model for social networks. In *TrustCom*, pages 841–846, 2011.
12. K. O’Hara. A General Definition of Trust. Technical report, University of Southampton, 2012.
13. Oxford English Dictionary. Trust. <https://en.oxforddictionaries.com/definition/trust> (accessed March 27, 2017).
14. J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24:33–60, 2005.
15. W. Sherchan, S. Nepal, and C. Paris. A survey of trust in social networks. *ACM Comput. Surv.*, 45(4):1–33, 2013.
16. J. Wu, L. Chen, Q. Yu, P. Han, and Z. Wu. Trust-aware media recommendation in heterogeneous social networks. *World Wide Web*, 18(1):139–157, 2015.
17. J. Wu and F. Chiclana. A social network analysis trust–consensus based approach to group decision-making problems with interval-valued fuzzy reciprocal preference relations. *Knowledge-Based Systems*, 59:97 – 107, 2014.