

A Flexible Framework for Probabilistic Models of Social Trust

Bert Huang, Angelika Kimmig*, Lise Getoor, and Jennifer Golbeck

University of Maryland, College Park, MD 20742

Abstract. In social networks, notions such as trust, fondness, or respect between users can be expressed by associating a strength with each tie. This provides a view of social interaction as a weighted graph. Sociological models for such weighted networks can differ significantly in their basic motivations and intuitions. In this paper, we present a flexible framework for probabilistic modeling of social networks that allows one to represent these different models and more. The framework, probabilistic soft logic (PSL), is particularly well-suited for this domain, as it combines a declarative, first-order logic-based syntax for describing relational models with a soft-logic representation, which maps naturally to the non-discrete strength of social trust. We demonstrate the flexibility and effectiveness of PSL for trust prediction using two different approaches: a structural balance model based on social triangles, and a social status model based on a consistent status hierarchy. We test these models on real social network data and find that PSL is an effective tool for trust prediction.

1 Introduction

Trust is a complex social phenomenon and a critical component of human social interaction. Modeling trust therefore plays an important role in social network analysis, with applications including viral marketing, collaborative filtering, and security. Computational modeling of trust provides added insight into the communication patterns, information flow, and behavior of social networks underlying these applications. In this paper, we present a computational framework for relational probabilistic modeling that is particularly well-suited for trust analysis in social networks. This framework is based on *probabilistic soft logic* (PSL) [1], an analysis engine that combines first-order rules with soft truth-values. PSL allows one to naturally capture structural ideas about the strength of trust, making it a natural, intuitive, and extensible framework for effective trust analysis.

The role of trust in social interactions has led to a vast body of work spanning many disciplines of science. Different types of factors influencing trust between two persons can be distinguished, relating to the trusting person (or truster), the trusted person (or trustee), the type of relationship between them, and the context in which trust occurs [2]. *Structural balance theory* in the context of trust suggests that social structures of trust can be stable or unstable. For example, social networks tend to exhibit *triadic closure*, which is loosely the concept that strong relationships are transitive [3]. Figure 1(a)

* Also at KU Leuven, Belgium.




Fig. 1. Implied structures according to competing theories of *structural balance* and *status*. The positive trust relationships from A to B and B to C imply opposite relationships from C to A in the two models.

illustrates examples of such stable structures. If A strongly trusts B , and B strongly trusts C , then triadic closure implies that A will likely trust C (and vice versa). On the other hand, if A does not trust B , B does not trust C , and C does not trust A , this represents an unstable state that structural balance theory suggests should be less likely to occur, as the theory prefers triads with one or three strong trust links.

A competing idea is that these social systems are governed by *status* or reputation. This is related to ideas from social psychology on reputation [4], where individuals are trusted based on their expertise in a particular area. In a social status model, the notion of trust is that the trustee (i.e., the person being trusted) is of higher status than the truster (i.e., the person who is trusting). Thus, under a status model, individuals exist in a hierarchy from the most trustworthy to the least trustworthy, along which trust propagates in triangular structures. As for structural balance, if A strongly trusts B , and B strongly trusts C , then status also implies that A will likely trust C . However, as illustrated in Figure 1(b), in contrast to structural balance, status predicts that C will likely *not* trust A in this case. Similarly, if A does not trust B and B does not trust C , then status disagrees with structural balance and implies that A likely does not trust C .

1.1 Related Work

A large community of research focuses on computational modeling of social trust. Methods for analyzing trust include graph-based approaches [5,6,7], probabilistic models [8,9,10], as well as other logic-based approaches [11]. These contributions tend to be fixed computational models based on particular theories of trust, whereas in this paper, we propose PSL as a general tool that provides the flexibility to explore various models without the need to adapt and redesign inference algorithms.

The foundations for many of these computational approaches stem from the vast sociological and psychological literature on human behavior. Recent studies have analyzed some of these theories in the context of social media data, specifically comparing the structural balance- and status-based models we emulate in this work [12,13]. Trust is also an important topic in business analytics; for example, modeling of trust is a useful component for effective viral marketing and e-commerce [14].

2 Probabilistic Soft Logic

Probabilistic soft logic (PSL) [1] is a general purpose system for probabilistic modeling and reasoning in relational domains.¹ PSL uses a first-order language to specify features of graphical models over ground atoms with soft truth-values from the interval $[0, 1]$. In this setting, finding the most likely truth-value assignment (*most probable explanation* inference) can be done efficiently. We refer to Broecheler et al. [1] and Bach et al. [15] for the technical details of the formalism and latest advances in efficient inference, respectively, and restrict the discussion here to an illustration of the key ideas by example. We start from a social network given as ground facts of the form $\text{KNOWS}(a, b)$, indicating that user a is acquainted with user b , where we are interested in the truth-values of ground facts of the form $\text{TRUSTS}(a, b)$, indicating that user a trusts user b . The following two PSL rules model general constraints that we might expect to hold in this domain, namely, trust being mutual and transitive. They are part of the structural balance model discussed in full detail in Section 2.1.

$$\begin{aligned} \text{TRUSTS}(A, B) &\stackrel{0.8}{\Rightarrow} \text{TRUSTS}(B, A) \\ \text{TRUSTS}(A, B) \wedge \text{TRUSTS}(B, C) \wedge \text{KNOWS}(A, C) &\stackrel{0.5}{\Rightarrow} \text{TRUSTS}(A, C) \end{aligned}$$

Rules have nonnegative weights (written above the implication operator \Rightarrow in the previous example) indicating their relative importance. The probability of a truth-value assignment to all ground atoms is defined as a function of the weighted *distance to satisfaction* of each ground rule. Generalizing the notion of rule satisfaction from the Boolean case to continuous truth-values, a rule is satisfied if the truth-value of its head (i.e., consequence) is at least that of the body (i.e., antecedent). The distance to satisfaction of an unsatisfied rule is the difference between the truth-values of the body and head. In our example program above, we thus prefer trust networks where many links are mutual and respect transitivity if users know each other, where the (hypothetical) weights indicate that the mutuality is considered more important.

2.1 Modeling Trust in PSL

We now expand the sketch above into models of competing theories for social trust in PSL. As before, we reason about two predicates KNOWS and TRUSTS , representing an observed social network and trust relationships between individuals, respectively. For any two individuals A and B in the social network, we set $\text{KNOWS}(A, B) = 1.0$ if A is acquainted with B , and 0.0 otherwise. Soft truth-values for TRUSTS atoms represent degrees of trust. For instance, $\text{TRUSTS}(A, B) = 1.0$ indicates that A fully trusts B , while $\text{TRUSTS}(A, B) = 0.5$ indicates that A somewhat trusts B , and $\text{TRUSTS}(A, B) = 0.0$ indicates that A does not trust B . We assume that if A trusts B , A knows B , but not necessarily vice versa. We use these models to predict unobserved truth-values of $\text{TRUSTS}(A, B)$ for pairs of individuals for whom $\text{KNOWS}(A, B)$ is true. In all our models, we include a prior for the truth-value of an atom $\text{TRUSTS}(A, B)$ centered around the global average of all observed trust scores.

¹ PSL is available as an open-source software package at <http://psl.umiacs.umd.edu>

In order to predict the degree of trust between two individuals, structural balance considers sixteen possible stable triangular structures involving the two individuals and a third individual. For example, an individual is likely to trust people his or her friends trust; this tendency is encoded as the first rule below. Simplified versions of the rules for each of these structures are:

$$\begin{aligned}
& \text{TR}(A, B) \wedge \text{TR}(B, C) \Rightarrow \text{TR}(A, C), & \text{TR}(B, A) \wedge \text{TR}(B, C) \Rightarrow \text{TR}(A, C), \\
& \text{TR}(A, B) \wedge \neg\text{TR}(B, C) \Rightarrow \neg\text{TR}(A, C), & \text{TR}(B, A) \wedge \neg\text{TR}(B, C) \Rightarrow \neg\text{TR}(A, C), \\
& \neg\text{TR}(A, B) \wedge \text{TR}(B, C) \Rightarrow \neg\text{TR}(A, C), & \neg\text{TR}(B, A) \wedge \text{TR}(B, C) \Rightarrow \neg\text{TR}(A, C), \\
& \neg\text{TR}(A, B) \wedge \neg\text{TR}(B, C) \Rightarrow \text{TR}(A, C), & \neg\text{TR}(B, A) \wedge \neg\text{TR}(B, C) \Rightarrow \text{TR}(A, C), \\
& \text{TR}(A, B) \wedge \text{TR}(C, B) \Rightarrow \text{TR}(A, C), & \text{TR}(B, A) \wedge \text{TR}(C, B) \Rightarrow \text{TR}(A, C), \\
& \text{TR}(A, B) \wedge \neg\text{TR}(C, B) \Rightarrow \neg\text{TR}(A, C), & \text{TR}(B, A) \wedge \neg\text{TR}(C, B) \Rightarrow \neg\text{TR}(A, C), \\
& \neg\text{TR}(A, B) \wedge \text{TR}(C, B) \Rightarrow \neg\text{TR}(A, C), & \neg\text{TR}(B, A) \wedge \text{TR}(C, B) \Rightarrow \neg\text{TR}(A, C), \\
& \neg\text{TR}(A, B) \wedge \neg\text{TR}(C, B) \Rightarrow \text{TR}(A, C), & \neg\text{TR}(B, A) \wedge \neg\text{TR}(C, B) \Rightarrow \text{TR}(A, C),
\end{aligned} \tag{1}$$

where we write TR as shorthand for TRUSTS to save space, and the full version of each rule is of the form,

$$\begin{aligned}
& \text{KNOWS}(A, B) \wedge \text{KNOWS}(B, C) \wedge \text{KNOWS}(A, C) \wedge \\
& \text{TRUSTS}(A, B) \wedge \text{TRUSTS}(B, C) \Rightarrow \text{TRUSTS}(A, C).
\end{aligned}$$

In these full versions of the rules, a parallel, positive KNOWS atom is added for each TRUSTS atom, which ensures that the groundings for A , B , and C are relevant entities representing acquaintance triangles in the social network.

In addition to the triangle rules, a natural extension of the structural balance model may include reciprocation of trust, which is captured using the rules

$$\begin{aligned}
& \text{TRUSTS}(A, B) \Rightarrow \text{TRUSTS}(B, A), \\
& \neg\text{TRUSTS}(A, B) \Rightarrow \neg\text{TRUSTS}(B, A).
\end{aligned} \tag{2}$$

The status model only makes predictions in the eight cases represented by the following simplified rules, where it agrees with structural balance on four triangular structures, but makes opposite predictions on the other four:

$$\begin{aligned}
& \text{TR}(X, Y) \wedge \text{TR}(Y, Z) \Rightarrow \text{TR}(X, Z), & \text{TR}(Y, X) \wedge \neg\text{TR}(Y, Z) \Rightarrow \neg\text{TR}(X, Z), \\
& \neg\text{TR}(X, Y) \wedge \neg\text{TR}(Y, Z) \Rightarrow \neg\text{TR}(X, Z), & \neg\text{TR}(Y, X) \wedge \text{TR}(Y, Z) \Rightarrow \text{TR}(X, Z), \\
& \text{TR}(X, Y) \wedge \neg\text{TR}(Z, Y) \Rightarrow \text{TR}(X, Z), & \text{TR}(Y, X) \wedge \text{TR}(Z, Y) \Rightarrow \neg\text{TR}(X, Z), \\
& \neg\text{TR}(X, Y) \wedge \text{TR}(Z, Y) \Rightarrow \neg\text{TR}(X, Z), & \neg\text{TR}(Y, X) \wedge \neg\text{TR}(Z, Y) \Rightarrow \text{TR}(X, Z),
\end{aligned} \tag{3}$$

where again we use shorthand for space, and in our full implementation, we include positive KNOWS atoms mirroring each TRUSTS atom that appears in a rule.

In contrast to the structural balance model, a natural addition to enforce a consistent status hierarchy suggests the inversion of trust between pairs of individuals. We can represent this with the rules

$$\begin{aligned} \text{TRUSTS}(X, Y) &\Rightarrow \neg\text{TRUSTS}(Y, X), \\ \neg\text{TRUSTS}(X, Y) &\Rightarrow \text{TRUSTS}(Y, X). \end{aligned} \quad (4)$$

3 Experiments

We now demonstrate the flexibility of trust modeling with probabilistic soft logic by evaluating different models on real social trust data.² We consider a structural balance model (referred to in our discussion below as PSL-Balance) comprised of the rules in (1), and a structural balance model with reciprocation (PSL-Balance-Recip), comprised of (1) and (2), as well as a status model (PSL-Status), comprised of (3), and a status model with inversion (PSL-Status-Inv), comprised of (3) and (4). We use the FilmTrust data set [17]³ as well as data from Epinions.com [7]. The FilmTrust data consists of a set of anonymized users, their trust values for other users, and their ratings for a set of movies (which we omit from this study). Users rate each other on a discrete scale of whole numbers from 1 to 10, which we normalize to $[0, 1]$, making each trust value interpretable as a soft truth-value. There are 1,754 users in the data set, among which there are 2,055 total user-to-user trust values. The trust values are directed and thus not symmetric. We sample via *snowball sampling* a network of 2,000 users from the Epinions data, which contains 8,675 discrete $\{-1, 1\}$ trust scores between users, which we treat as false and true TRUSTS predicate values.

The task we consider is collective prediction of trust values given the fully-observed social network. We generate eight folds where, in each fold, $1/8$ of the trust values are hidden at random. The prediction algorithm can use the remaining $7/8$ of the trust values and the full structure of the social network to learn parameters for a model and perform inference of the unknown trust values. For example, PSL learns weights for the rules in each given model from these observed trust values.

3.1 Baselines

We compare our PSL models to a range of baselines, including two popular approaches for computational trust modeling. As a simple baseline, we predict the average trust across all observed trust values for every prediction. EigenTrust [6] is a global metric that computes a trust value for each node by finding the left principle eigenvector of a normalized trust matrix. The trust matrix is normalized such that each row sums to 1.0, making the normalized trust matrix stochastic. EigenTrust’s prediction is then the stationary distribution of the stochastic process described by the normalized trust matrix, or equivalently the limit on the probability of landing on each node as a random walk approaches infinity, where the probability of walking to a neighbor is proportional to how much the current node trusts the neighbor.

TidalTrust [5] is a graph-based algorithm that propagates trust values through neighbors by recursively using the weighted average of neighbor trust to decide a node’s trust

² Early versions of these experiments appeared in [16].

³ FilmTrust is a web service designed to leverage user-to-user trust values and user-to-movie ratings for movie recommendation. <http://trust.mindswap.org/FilmTrust>

Table 1. Average scores of FilmTrust trust predictions using mean average error (MAE), Kendall-tau statistic τ , and Spearman’s rank correlation ρ for the full test set and the non-default predictions (MAE*, τ^* , and ρ^*). Each statistic is computed separately on each fold, and the average over all folds is listed here. Scores that are statistically equivalent to the best score according to a two-sample t-test with rejection threshold 0.05 in each metric are typed in bold.

Method	MAE	τ	ρ	MAE*	τ^*	ρ^*
Average	0.210	n/a	n/a	n/a	n/a	n/a
EigenTrust	0.339	-0.054	-0.074	0.339	-0.054	-0.074
TidalTrust	0.229	0.059	0.078	0.236	0.089	0.117
PSL-Balance	0.207	0.136	0.176	0.193	0.235	0.314
PSL-Balance-Recip	0.207	0.139	0.188	0.193	0.241	0.318
PSL-Status	0.224	0.112	0.144	0.230	0.205	0.277
PSL-Status-Inv	0.224	0.065	0.085	0.238	0.143	0.189

for another. TidalTrust predicts distinct trust values per link, rather than a single global trust value per node. To predict an unknown trust value from a source node to a sink node, the algorithm uses a breadth-first search to determine the set of minimum length paths from the source to the sink. TidalTrust then recursively computes the neighbor-weighted trust for the sink node along these paths, starting from the sink node until finally reaching the source, at which point it outputs the final weighted trust.

3.2 Results

On the FilmTrust data, since the ground truth is continuous-valued, we measure for each algorithm the average score over the eight folds for three metrics: mean average error (MAE), Kendall’s τ statistic, and Spearman’s rank correlation ρ . MAE measures the absolute error on the soft truth-values, while τ and ρ measure ranking performance. The average scores are listed in Table 1. Three PSL models, all but PSL-Status-Inv, are statistically tied for the best-performing method on all three metrics, according to a two-sample t-test with rejection threshold 0.05. This suggests that the inversion rules do not help in this setting. Both EigenTrust and TidalTrust do not do as well here as in their natural problem setup. Here, the prediction algorithms must do joint inference over many unknown trust values, where a significant fraction of the values are unknown. This can disrupt network-based methods that depend on the connectivity of observed information. For example, TidalTrust depends on the existence of alternate paths between nodes, and the removal of a full eighth of this already sparse network significantly increases the number of pairs for which a directed path does not exist. In these cases, we set TidalTrust to predict the global average of all trust values. Since EigenTrust returns a probability distribution over the nodes, its predictions are not on the same scale as the true values, thus making it difficult to directly compare the raw error. Nevertheless, the disconnected state of the network causes the spectral prediction to seemingly fail at recovering any signal from the data when measuring rank correlations. In contrast, PSL takes advantage of the edges with unobserved trust values to propagate information across the network during collective inference, and is thus more robust to the disconnections from the sampling process.




Fig. 2. Small example subgraphs of FilmTrust network. We plot the one-hop neighborhood around two triangles where the two primary PSL models disagree. The top row contains the network surrounding a triangle whose ground truth behaves consistently with the structural balance model, and the bottom row contains one surrounding a triangle that is more consistent with the status model. Edges are colored by trust scores, ranging from blue (no trust) to red (maximal trust). The left column (a,d) contains the ground truth trust network, the middle column (b,e) contains the predictions by the PSL-Balance model, and the right column (c,f) contains those by the PSL-Status model.

Because of the network’s sparsity, especially after subsampling for testing, methods that propagate trust values suffer. In our experiments, we allow each method to predict the global average when no information is available to propagate, due to the query edge being disconnected from any observed edge. In PSL, the prior has a similar effect. To isolate performance on nontrivial predictions, we also measure accuracy statistics only on the edges for which the method predicts non-default values. We list these in Table 1 as MAE^* , τ^* , and ρ^* . On the non-default predictions, the PSL models show a clear advantage over others, suggesting that their joint inference effectively propagates trust throughout the network.

Comparing the competing PSL models, of all 2055 edges, PSL-Balance and PSL-Status predict 514 trust scores that differ by at least 0.1. In Figure 2, we visualize a few cases where these predictions differ. Since both models produce similarly high accuracy, we suspect the trust behavior in the FilmTrust network follows a combination of the two models.

Table 2. Average area under precision-recall curves of various methods predicting Epinions trust relationships. We compute precision and recall with respect to finding the rarer non-trust links. Scores statistically equivalent to the best score are displayed in bold, where statistical significance is measured via a two-sample t-test with rejection threshold 0.05.

Method	AUC
Average	0.069762
PSL-Balance	0.316843
PSL-Balance-Recip	0.343011
PSL-Status	0.296563
PSL-Status-Inv	0.279580
EigenTrust	0.131159
TidalTrust	0.129785

On the Epinions data, since the ground truth is discrete-valued, we measure for each algorithm the precision and recall on the non-trust links. We compute these retrieval metrics on the non-trust links because the majority of links in the data set are trust links, with 7,974 positive links and only 701 negative links. Thus, we expect a better prediction to retrieve the rarer link type. Table 2 lists the average area under the precision-recall curve for each method. The PSL-Balance-Recip model again produces the best-scoring prediction, while PSL-Balance and PSL-Status produce statistically equivalent results, according to two-sample t-tests with rejection threshold 0.05.

4 Discussion

This paper proposes the use of probabilistic soft logic (PSL) as a natural framework for modeling trust in social networks. Such a generic framework allows for easy exploration of trust models based on different assumptions about social phenomena. To demonstrate the effectiveness of PSL for this task, we apply competing trust models based on structural balance and status to predict user trust data. Further exploration of the literature on trust within this framework is a promising direction for future work. For instance, one could model multiple relationship types and trust topics, capturing the intuition that a person may trust a sibling more than a co-worker about family issues, while trusting the co-worker more about career advice. Similarly, different people have varying degrees of expertise on particular topics, earning them different levels of trust dependent on the context. Finally, the structure of social trust is similar in form to various other phenomena in social networks, such as opinion, social influence, and complex contagion modeling. Each of these problems may benefit from the power and flexibility of a system such as PSL.

Acknowledgments. This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI / NBC) contract number D12PC00337. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or

endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government. A. Kimmig is a postdoctoral fellow of the Research Foundation Flanders (FWO Vlaanderen).

References

1. Broecheler, M., Mihalkova, L., Getoor, L.: Probabilistic similarity logic. In: Conference on Uncertainty in Artificial Intelligence, UAI (2010)
2. Levin, D.: Trust. In: Clegg, S., Bailey, J. (eds.) *International Encyclopedia of Organization Studies*, pp. 1573–1579. Sage (2008)
3. Granovetter, M.: The Strength of Weak Ties. *The American Journal of Sociology* 78(6), 1360–1380 (1973)
4. Cosmides, L., Tooby, J.: *Cognitive Adaptions for social exchange. The adapted mind: Evolutionary psychology and the generation of culture*. Oxford University Press (1992)
5. Golbeck, J.: *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, University of Maryland, College Park, College Park, MD, USA (2005)
6. Kamvar, S., Schlosser, M., Garcia-Molina, H.: The Eigentrust algorithm for reputation management in P2P networks. In: *International Conference on World Wide Web, WWW* (2003)
7. Richardson, M., Agrawal, R., Domingos, P.: Trust Management for the Semantic Web. In: Fensel, D., Sycara, K., Mylopoulos, J. (eds.) *ISWC 2003. LNCS*, vol. 2870, pp. 351–368. Springer, Heidelberg (2003)
8. Kuter, U., Golbeck, J.: Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In: *National Conf. on Artif. Intelligence, AAAI* (2007)
9. Rettinger, A., Nickles, M., Tresp, V.: Statistical relational learning of trust. *Machine Learning* 82(2), 191–209 (2011)
10. Vydiswaran, V., Zhai, C., Roth, D.: Content-driven trust propagation framework. In: *Knowledge Discovery and Data Mining, KDD* (2011)
11. Jøsang, A., Hayward, R., Pope, S.: Trust network analysis with subjective logic. In: *Australasian Computer Science Conference, ACSC* (2006)
12. Leskovec, J., Huttenlocher, D., Kleinberg, J.: Signed networks in social media. In: *Conference on Human Factors in Computing Systems, CHI* (2010)
13. Guha, R., Kumar, R., Raghavan, P., Tomkins, A.: Propagation of trust and distrust. In: *International Conference on World Wide Web, WWW* (2004)
14. Salam, A., Iyer, L., Palvia, P., Singh, R.: Trust in e-commerce. *Commun. ACM* 48(2), 72–77 (2005)
15. Bach, S., Broecheler, M., Getoor, L., O’Leary, D.: Scaling constrained continuous Markov random fields with consensus optimization. In: *Adv. in Neur. Information Proc. Sys. (NIPS)* (2012)
16. Huang, B., Kimmig, A., Getoor, L., Golbeck, J.: Probabilistic soft logic for trust analysis in social networks. In: *International Workshop on Stat. Relational Artif. Intelligence* (2012)
17. Golbeck, J., Hendler, J.: FilmTrust: Movie recommendations using trust in web-based social networks. *ACM Transactions on Internet Technology* 6(4), 497–529 (2006)