
Probabilistic Soft Logic for Trust Analysis in Social Networks

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

Computer Science Dept.

University of Maryland

College Park, MD 20740

{bert, angelika, getoor, golbeck}@cs.umd.edu

Abstract

Trust plays a key role in social interactions. Explicitly modeling trust is therefore an important aspect of social network analysis in settings such as reputation management systems, recommendation systems, and viral marketing. Within the social sciences, trust is known to depend on network structure, context, individual actors' attributes, and group memberships and affiliations. Furthermore, trust is often measured quantitatively, according to degrees of trust, rather than as a binary indicator. In this paper, we propose trust modeling as a rich challenge for statistical relational learning (SRL). Additionally, we show that probabilistic soft logic (PSL) is particularly well-suited for this problem. PSL, like many SRL languages, provides an intuitive framework for capturing the relational aspects of trust modeling, while its soft truth values easily accommodate varying strengths of trust. We model various sociological theories of trust in PSL and experimentally compare the resulting PSL programs to existing trust prediction methods, demonstrating the ease of model development and showing that these interpretable first-order logic models produce results of competitive quality.

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 the social networks underlying these applications. While many trust models from the

social sciences can be expressed in terms of logical rules, such crisp formulations are usually too strict to be of practical use. Furthermore, trust relationships are often neither Boolean in nature nor known explicitly, thus adding another layer of uncertainty.

In this paper, we therefore highlight trust modeling as a challenging problem for statistical relational learning (SRL) and statistical relational artificial intelligence in general. We reinforce our point through a detailed discussion of trust modeling in *probabilistic soft logic* (Broecheler et al., 2010). We consider probabilistic soft logic (PSL) particularly effective for this problem domain, as it combines first-order rules as used in many SRL systems with soft truth values that naturally capture degrees of trust. PSL thus offers a natural, intuitive, and extensible framework for effective trust analysis.

We model three aspects of social trust theories in PSL and compare these models, as well as their combinations, against established prediction algorithms, showing that these interpretable first-order logic models produce results of competitive quality. In particular, we consider the problem of predicting trust in a partially labeled social network, in which each node represents a person and an edge represents some form of personal relationship (usually friendship). Given trust values for some person-person edges, the goal is to accurately predict the trust values for all edges whose trust values are unobserved.

1.1 Related Work

A large community of research focuses on computational modeling of social trust. Methods for analyzing trust include graph-based approaches (Golbeck, 2005; Kamvar et al., 2003; Richardson et al., 2003), probabilistic models (Kuter and Golbeck, 2007; Rettinger et al., 2011; Vydiswaran et al., 2011), as well as other logic-based approaches (Jøsang et al., 2006). These contributions tend to be fixed computational models based on particular theories of trust, whereas in this paper, we propose SRL as a general tool that provides the flexibility to explore various models with-

*Also at KU Leuven, Belgium

out the need for adapting inference algorithms.

The foundations for many computational approaches to trust stem from the vast sociological and psychological literature on human behavior. We review a number of these studies in Section 3. 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 (Salam et al., 2005).

Probabilistic soft logic (Broecheler et al., 2010) is one of various modern approaches that combine probabilistic reasoning with logic or programming languages, including *Markov logic networks* (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007), and FACTORIE (McCallum et al., 2009). Its main characteristic difference is the use of soft logic, which depends on continuous relaxations of Boolean algebra known as *triangular norms* (t-norms) (Klement et al., 2000).

2 Probabilistic Soft Logic

Probabilistic soft logic (PSL) (Broecheler et al., 2010) is a system for probabilistic modeling using first-order logic syntax. PSL uses soft truth values, relaxing truth to the interval $[0, 1]$ and adapting logical connectives accordingly. As a consequence of the soft logic formulation, inference in PSL is a convex optimization problem. Additionally, the soft truth values allow the natural integration of external functions ranging in the same interval, such as normalized similarity functions.¹ This section provides a short overview of PSL, its usage, and its internal representation.

PSL uses first-order logic (FOL) as its underlying modeling language. In a PSL program, relationships and attributes are modeled by different *predicates* (of arbitrary arity), and first order *rules* model dependencies or constraints on these predicates. For instance, $\text{TRUSTS}(A, B) \Rightarrow \text{KNOWS}(A, B)$ reads as “if A trusts B , then A knows B ”, where A and B are variables referring to arbitrary objects. Replacing these variables with constants from the domain of the program results in a *ground rule*. PSL extends the notion of rule to the soft context, i.e., rules can be assigned a weight from \mathbb{R}^+ indicating at what cost a grounding of the rule can be violated. For instance,

$$\text{TRUSTS}(A, B) \wedge \text{TRUSTS}(B, C) \stackrel{0.6}{\Rightarrow} \text{TRUSTS}(A, C)$$

models that the trust relation is not fully transitive and gets weaker along chains of links. Furthermore, a PSL program specifies known truth values for a subset of ground atoms. For instance, $\text{KNOWS}(\text{Alice}, \text{Bob}) = 1.0$ and $\text{TRUSTS}(\text{Alice}, \text{Bob}) = 0.6$ indicate that Alice knows Bob, but only trusts him somewhat above average. Throughout

¹For this reason, PSL was originally introduced as *probabilistic similarity logic*, emphasizing its ability to elegantly leverage similarity functions in its logic.

the text, we use the convention that predicates are written in small caps (e.g., TRUSTS) and variables are italicized capital letters (e.g., A).

To relax Boolean truth values to continuous variables, PSL uses the *Lukasiewicz t-norm* and its corresponding *co-norm* as the relaxation of the logical AND and OR, respectively. These relaxations are exact at the extremes, when variables are either true (1.0) or false (0.0), but provide a consistent mapping for values in-between. The formulas for the relaxation of the logical conjunction (\wedge), disjunction (\vee), and negation (\neg) are as follows:

$$\begin{aligned} a \tilde{\wedge} b &= \max\{0, a + b - 1\}, \\ a \tilde{\vee} b &= \min\{a + b, 1\}, \\ \tilde{\neg} a &= 1 - a, \end{aligned}$$

where we use $\tilde{\cdot}$ to indicate the relaxation from the Boolean domain. Rules are evaluated using the Lukasiewicz norms by converting the implication operator with the identity

$$X \Rightarrow Y \equiv \tilde{\neg} X \tilde{\vee} Y.$$

The probability distribution defined by a PSL program measures the overall distance to satisfaction, that is, the more groundings of rules have high truth values in an interpretation, the more likely it is. More formally, for a PSL program, let G be the set of all groundings for each rule. For any grounding $g \in G$, let w_g be the weight assigned to the rule, and $t_g(x) \in [0, 1]$ be the grounded rule’s truth-value under interpretation x . The probability distribution over interpretation x defined by the program is

$$\Pr(x; w) = \exp \left(- \sum_{g \in G} w_g (1 - t_g(x)) \right).$$

Considering each grounded rule a factor and each truth value a variable, this probability distribution becomes a log-linear Markov random field over continuous variables. Maximum likelihood inference for the unknown truth values corresponds to solving a linear program, where the truth-value variables are constrained to be consistent with respect to the t-norms and are weighted by rule potentials. Additional details, including a description of a learning algorithm for setting the weights, are provided by Broecheler et al. (2010).

3 Modeling Trust

The key role of trust in social interactions is mirrored by the vast body of work spanning many disciplines of science. On a high level, different types of factors influencing trust between two persons can be distinguished, relating to the type of relationship between them, the trusting person, the trusted person, and the context in which trust occurs (Levin, 2008). In this section, we review a narrow portion of this

literature, focusing on general principles that we model in PSL in Section 3.1.

Social network theory studies the structural balance of relationships. For example, social networks tend to exhibit *triadic closure*, which is loosely the concept that strong relationships are transitive (Granovetter, 1973). In the context of trust, this idea translates to how people determine whether to trust others by consulting with those they trust. For example, if Alice strongly trusts Bob, and Bob strongly trusts Chris, then triadic closure implies that Alice will likely trust Chris.

Another common idea in analysis of trust is that of *reputation*, where people who are trusted gain a reputation of being trustworthy, thus garnering more trust (Cosmides and Tooby, 1992). Additionally, the qualities of the trustee (i.e., the person who is trusted) have been identified as important factors for determining trust. For example, whether Alice trusts Bob depends on Bob’s beliefs and goals, as well Alice’s notions of confidence in Bob. People also have person-specific innate tendencies for trust, which may stem from early-childhood experiences (Castelfranchi and Falcone, 2000).

Trust is also known to be affected by the similarity in traits of the involved people. In particular, trust exhibits the notion of *homophily*, a concept from social network theory which suggests that people connect to others with whom they are similar (Bhattacharya et al., 1998).

Finally, an important aspect of trust is its context-dependency. Trust determines how much individuals value information communicated from each other, so it is natural to consider the level of trust to be a function of the information’s topic area. Similarly, trust behavior varies significantly between different relationship types, such as trust between family members, co-workers, or religious group members (Glanville and Paxton, 2007).

3.1 Modeling Trust in PSL

We now present three sets of first-order logic rules, each modeling a different aspect from social theory. We view these as building blocks that could be used individually or in combination for trust models in many SRL formalisms. We demonstrate this principle in the context of PSL, allowing us to easily represent degrees of trust and rely on PSL’s parameter learning technique to estimate rule weights. We model trust relations with a binary predicate `TRUSTS`. A soft truth value for `TRUSTS(A, B) = 1.0` indicates that *A* fully trusts *B*, while `TRUSTS(A, B) = 0.5` indicates that *A* somewhat trusts *B*, and `TRUSTS(A, B) = 0.0` indicates that *A* does not trust *B*. In other probabilistic logic systems, such modeling can be cumbersome and may require discretization of the trust scale.

The first social phenomenon we model is triadic closure.

We encode the tendency for transitivity and reciprocity in trust using the rules listed in Figure 1, enumerating various triangle structures and their likely effect on trust. We refer to this model as PSL-Triadic.

The second social phenomenon we model is basic personality. More specifically, we consider additional predicates `TRUSTING` and `TRUSTWORTHY`, modeling whether a person is trusting or trustworthy, respectively. These predicates are not part of the input data, but they correspond to hidden variables that need to be inferred during prediction of trust values. The intuition is that a trusting person is likely to trust more, while a trustworthy person will earn more trust. The rules for this model, which we refer to as PSL-Personality, are listed in Figure 2.

The third social phenomenon we model is the effect of similarity on trust. Homophily is the tendency of individuals to associate with others who are similar. The trust ratings people have assigned to one another in our experiments are set in the context of movies (i.e., how much do users trust others’ opinions about movies). This makes the movie rating data especially relevant to understanding trust. Previous work has shown that trust in similar social network data is strongly correlated with similarity (Ziegler and Golbeck, 2007). In this PSL model, we consider an additional predicate `SAMETRAITS(A, B)`, which indicates the similarity of *A* and *B* according to their personal traits. For example, in our experiments, we measure the similarity of users’ survey responses on movie preferences. The intuition here is that people with similar traits tend to trust each other. We additionally consider the idea that people who trust (or do not trust) a particular individual will likely trust (or not trust) those similar to that individual. Conversely, similar people will trust (or not trust) similar sets of trustees. The rules encoding these intuitions, which form the model PSL-Similarity, are listed in Figure 3.

Finally, we also combine the models (into PSL-TriadPers, TriadSim, PersSim, and TriadPersSim) by simply creating PSL programs that include the rules from the component models. The next section reports on empirical experiments using these models.

4 Experiments

In this section, we demonstrate the flexibility of trust modeling with probabilistic soft logic. For this task, we evaluate on the FilmTrust data set (Golbeck and Hendler, 2006). FilmTrust² is a web service designed to leverage user-to-user trust values and user-to-movie ratings for movie recommendation. The dataset consists of a set of anonymized users, their trust values for other users, and their ratings for a set of movies. Since the user trust values are rather sparse, we prune the data to only include the largest con-

²<http://trust.mindswap.org/FilmTrust/>

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

Figure 1: Rules for PSL model of triadic closure (PSL-Triadic). Triadic closure implies the transitivity of trust, such that individuals tend to determine whom to trust based on the opinions of those they trust.

$$\begin{aligned}
& \text{TRUSTS}(A, B) \Rightarrow \text{TRUSTING}(A), \\
& \neg \text{TRUSTS}(A, B) \Rightarrow \neg \text{TRUSTING}(A), \\
& \text{TRUSTS}(A, B) \Rightarrow \text{TRUSTWORTHY}(B), \\
& \neg \text{TRUSTS}(A, B) \Rightarrow \neg \text{TRUSTWORTHY}(B), \\
& \text{TRUSTING}(A) \wedge \text{TRUSTWORTHY}(B) \Rightarrow \text{TRUSTS}(A, B), \\
& \neg \text{TRUSTING}(A) \wedge \neg \text{TRUSTWORTHY}(B) \Rightarrow \neg \text{TRUSTS}(A, B), \\
& \text{TRUSTING}(A) \Rightarrow \text{TRUSTS}(A, B), \\
& \neg \text{TRUSTING}(A) \Rightarrow \neg \text{TRUSTS}(A, B), \\
& \text{TRUSTWORTHY}(B) \Rightarrow \text{TRUSTS}(A, B), \\
& \neg \text{TRUSTWORTHY}(B) \Rightarrow \neg \text{TRUSTS}(A, B).
\end{aligned}$$

Figure 2: Rules for PSL model of basic personality (PSL-Personality). This model maintains predicates for whether users are trusting or trustworthy, and uses these predicates to determine each pairwise trust. Trusting individuals are more prone to offer trust, while trustworthy individuals are more prone to receive trust.

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

Figure 3: Rules for trust via similarity and homophily (PSL-Similarity). These rules model the correlation between feature similarity and trust, where the predicate SAMETRAITS indicates whether two individuals have similar traits. The two phenomena modeled here are homophily, which implies that people tend to trust those similar to themselves, and the idea that similar people trust and are trusted in similar ways.




Figure 4: Histograms of predicted trust values over true trust annotations. The brightness of each grid cell indicates the number of edges with the corresponding true trust (horizontal axis) and predicted trust (vertical axis). More mass along the diagonal indicates predictions consistent with the ground truth.

ected component of users. 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. Similarly, users rate movies with a recommendation rating between 1 and 5. There are 500 users in the largest connected component, among which there are 1574 total user-to-user trust values. The trust values are directed and thus not symmetric. For each pair of users within a two-hop distance, we compute their soft similarity SAMETRAITS via a normalized inner product of their overlapping rated-entry vectors. Let $r(A, M)$ denote the rating given by user A for movie M , scaled to the interval $[0, 1]$ (by dividing by 5). If a user did not rate a movie, the score for that user-movie pair is 0. We compute SAMETRAITS with the formula

$$\text{SAMETRAITS}(A, B) = \frac{\sum_{M \in \text{Movies}} r(A, M)r(B, M)}{\sum_{M \in \text{Movies}} I(r(A, M)r(B, M) > 0)}.$$

The task we consider is collective prediction of trust values. We generate four folds where, in each fold, 1/4 of the trust values are hidden at random. The prediction algorithm can use the remaining 3/4 of the trust values to learn parameters for a model and perform inference of the unknown trust values. PSL learns weights for the rules in each given model from these observed trust values. We

consider a transductive prediction setting, in which the inference algorithm is given which pairs of users rated each other (i.e., the full network structure), but is not given the actual trust values on the held-out 1/4.

4.1 Baselines

We now discuss a range of baselines, including two popular approaches from the literature, to which we compare our PSL models from Section 3.1.

As a simple baseline, we consider predicting the average trust across all observed trust values for every prediction (denoted in tables as Avg-Global). Clearly, the global average is not very informative, so we additionally consider a node-centric metric where we compute average trust values for each node (and only use the global average if no trust values are available for a node). We use two variants of this: the average of incoming trust values (Avg-Incoming) and the average of outgoing trust values (Avg-Outgoing).

We include the SAMETRAITS predicate itself as a baseline, since PSL-Similarity heavily depends on this similarity function.

EigenTrust (Kamvar et al., 2003) is a global metric analogous to PageRank (Page et al., 1999) that computes a trust value for each node by finding the left principle eigenvec-

Table 1: Average scores of trust predictions using mean average error (MAE), Kendall-tau statistic τ , and Spearman’s rank correlation ρ . Each statistic is computed separately on each fold, and the average over all folds is listed here.

Method	MAE	τ	ρ
PSL-Triadic	0.2985	0.0717	0.0944
PSL-Personality	0.2586	0.1681	0.2265
PSL-Similarity	0.2198	0.1089	0.1539
PSL-TriadPers	0.2509	0.1801	0.2417
PSL-TriadSim	0.2146	0.1197	0.1688
PSL-PersSim	0.2154	0.1771	0.2444
PSL-TriadPersSim	0.2246	0.1907	0.2598
SAMETRAITS	0.2461	0.0531	0.0739
Avg-Incoming	0.3751	0.0120	0.0167
Avg-Outgoing	0.3327	0.1088	0.1463
Avg-Global	0.2086	–	–
EigenTrust	0.6729	-0.0229	-0.0291
TidalTrust	0.2387	0.0478	0.0649

tor 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 (Golbeck, 2005) 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 for another. In contrast to EigenTrust, 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.

4.2 Results

For each algorithm, we measure the average score over the four folds for three metrics: mean average error (MAE), Kendall’s τ statistic, and Spearman rank correlation ρ . MAE measures the absolute error on the soft truth values, while τ and ρ measure ranking performance. For true trust values t_1, \dots, t_n and predicted trust values p_1, \dots, p_n , where the orderings of these values from greatest to least are t' and p' respectively, these metrics are defined as fol-

lows:

$$\begin{aligned} \text{MAE} &= \frac{1}{n} \sum_{i=1}^n |t_i - p_i|, \\ \tau &= \left(\sum_{i=1}^n \sum_{j=i+1}^n \text{sgn}(t'_i - t'_j) \text{sgn}(p'_i - p'_j) \right) / \binom{n}{2}, \\ \rho &= 1 - 6 \sum_{i=1}^n \frac{(t'_i - p'_i)^2}{n(n^2 - 1)}. \end{aligned}$$

The average scores are listed in Table 1. Perhaps surprisingly, simply predicting the global average produces the lowest (best) MAE. However, this can be explained by the fact that users tend to rate approximately the same value for all their friends, which makes predicting the average a reasonable guess for minimizing the error over the full range of predictions. This indicates that the mean error is not very informative here. Indeed, more variation is visible in the ranking metrics, which better capture the overall range of predictions. This is further illustrated by two-dimensional histograms of the trust value agreements in Figure 4. These histograms plot the distribution of predictions in a 10×10 grid, where true and predicted values correspond to the horizontal and vertical axes, respectively. Perfect predictions would have all the mass along the diagonal.

Both EigenTrust and TidalTrust are somewhat crippled by the problem setup, since removing many trust values creates gaps in the network, which strongly affect these methods. In particular, TidalTrust depends on the existence of alternate paths between nodes, and, despite the initial network being a connected component, the removal of a full quarter of the trust edges 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.

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. Each individual PSL model produces different predictions, as visualized in Figure 4, and, while PSL-Personality produces the best ranking among the individual models, each combination of models produces better rankings than its component models alone. For example, PSL-TriadPers outperforms both PSL-Triadic and PSL-Personality. The combination of all models, PSL-TriadPersSim produces the overall best-scoring ranking. Additionally, the benefits of collective inference via the logical rules are evident from the improvement of PSL-Similarity over the raw SAMETRAITS score,

since the PSL model uses the SAMETRAITS values with added relational logic to produce a better ranking.

5 Discussion

In this paper, we highlight social trust analysis as a challenging problem for statistical relational artificial intelligence. The dynamics of social trust can naturally be modeled by first-order rules, while statistical techniques address the inherent uncertainty of trust assessments. We demonstrate the promise of statistical relational models for trust analysis by testing various probabilistic soft logic (PSL) models in a small, collective trust prediction problem. The soft truth values of PSL make it particularly well suited to model degrees of trust. A comparison of our PSL models with existing trust prediction methods shows that these interpretable first-order logic models produce results of competitive quality.

As illustrated by our PSL models, the use of a generic SRL framework allows for easy exploration of the space of possible trust models, as extensions and variations can readily be incorporated. As this paper only examines a small portion of the literature on trust, such an exploration 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 offers an exciting application for statistical relational learning and the potential to transfer models between specific domains. We are actively exploring the application of PSL for these problems as well.

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/NBA, or the U.S. Government. A. Kimmig is a postdoctoral fellow of the Research Foundation Flanders (FWO Vlaanderen).

References

- R. Bhattacharya, T. Devinney, and M. Pillutla. A formal model of trust based on outcomes. *The Academy of Management Review*, 23(3):459–472, 1998.
- M. Broecheler, L. Mihalkova, and L. Getoor. Probabilistic similarity logic. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010.
- C. Castelfranchi and R. Falcone. Trust is much more than subjective probability: Mental components and sources of trust. In *Hawaii International Conference on System Sciences (HICSS)*, 2000.
- L. Cosmides and J. Tooby. *Cognitive Adaptions for social exchange*. The adapted mind: Evolutionary psychology and the generation of culture. Oxford University Press, 1992.
- L. De Raedt, A. Kimmig, and H. Toivonen. ProbLog: A probabilistic Prolog and its application in link discovery. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- J. Glanville and P. Paxton. How do we learn to trust? a confirmatory tetrad analysis of the sources of generalized trust. *Social Psychology Quarterly*, 70(3):230–242, 2007.
- J. Golbeck. *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, University of Maryland, College Park, College Park, MD, USA, 2005.
- J. Golbeck and J. Hendler. FilmTrust: Movie recommendations using trust in web-based social networks. *ACM Transactions on Internet Technology*, 6(4):497–529, 2006.
- M. Granovetter. The Strength of Weak Ties. *The American Journal of Sociology*, 78(6):1360–1380, 1973.
- A. Jøsang, R. Hayward, and S. Pope. Trust network analysis with subjective logic. In *Australasian Computer Science Conference (ACSC)*, 2006.
- S. Kamvar, M. Schlosser, and H. Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. In *International Conference on World Wide Web (WWW)*, 2003.
- E. Klement, R. Mesiar, and E. Pap. *Triangular Norms*. Kluwer Academic Publishers, 2000.
- U. Kuter and J. Golbeck. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In *National Conference on Artificial Intelligence (AAAI)*, 2007.
- D. Levin. Trust. In S. Clegg and J. Bailey, editors, *International Encyclopedia of Organization Studies*, pages 1573–1579. Sage, 2008.
- A. McCallum, K. Schultz, and S. Singh. FACTORIE: Probabilistic programming via imperatively defined fac-

- tor graphs. In *Neural Information Processing Systems (NIPS)*, 2009.
- L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, 1999.
- A. Rettinger, M. Nickles, and V. Tresp. Statistical relational learning of trust. *Machine Learning*, 82(2):191–209, 2011.
- M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62(1-2):107–136, 2006.
- M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *International Semantic Web Conference*, volume 2870 of *Lecture Notes in Computer Science*, pages 351–368. Springer, 2003.
- A. Salam, L. Iyer, P. Palvia, and R. Singh. Trust in e-commerce. *Commun. ACM*, 48(2):72–77, 2005.
- V. Vydiswaran, C. Zhai, and D. Roth. Content-driven trust propagation framework. In *Knowledge Discovery and Data Mining (KDD)*, pages 974–982, 8 2011.
- C.-N. Ziegler and J. Golbeck. Investigating interactions of trust and interest similarity. *Decis. Support Syst.*, 43(2): 460–475, 2007.