http://linqs.cs.umd.edu

Supported by IARPA contract D12PC00337 and the Research Foundation Flanders (FWO-Vlaanderen)

Probabilistic Soft Logic for Trust Analysis in Social Networks

Bert Huang, Angelika Kimmig, Lise Getoor, and Jennifer Golbeck
Department of Computer Science
University of Maryland, College Park

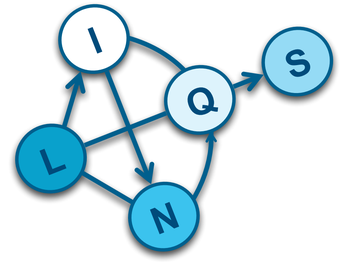

Trust Analysis

Trust analysis in social networks is a rich application for **statistical relational A.I. and learning** (SRL)

Modeling trust benefits many applications: recommendation engines, viral marketing, reputation management

Many existing methods for computational trust analysis, e.g., TidalTrust, EigenTrust, Advogato.

We analyze trust using the SRL tool **Probabilistic Soft Logic**

Probabilistic Soft Logic

Declarative language for probabilistic models using **first-order logic** (FOL) syntax

Truth values are relaxed to **soft truth** in [0,1]

Mechanisms for incorporating **similarity functions** and reasoning about **sets**

Logical operators relaxed via **Lukasiewicz t-norm**:

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

Continuous probability distribution over truth values:

$$\Pr(x) = \frac{1}{Z} \exp\left(- \sum_{r \in P} \sum_{g \in G(r)} w_r (1 - t_g(x)) \right)$$

P : the PSL program $G(r)$: set of all groundings of rule r
w_r : weight of rule r $t_g(x)$: truth value of grounding g

MAP inference is a **convex optimization**

Efficient sampling method for marginal inference

Weight learning via voted perceptron maximum likelihood

PSL trust models

We model trust with logical models inspired by ideas from social science

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

PSL-Triadic

models triadic closure, which implies the transitivity of trust, such that individuals determine whom to trust based on the opinions of those they trust.

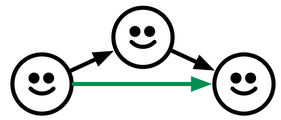

$$\begin{aligned}
\text{trusts}(A,B) &\Rightarrow \text{trusting}(A),\\
\neg\text{trusts}(A,B) &\Rightarrow \neg\text{trusting}(A),\\
\text{trusts}(A,B) &\Rightarrow \text{trustworthy}(B),\\
\neg\text{trusts}(A,B) &\Rightarrow \neg\text{trustworthy}(B),\\
\text{trusting}(A) \wedge \text{trustworthy}(B) &\Rightarrow \text{trusts}(A,B),\\
\neg\text{trusting}(A) \wedge \neg\text{trustworthy}(B) &\Rightarrow \neg\text{trusts}(A,B),\\
\text{trusting}(A) &\Rightarrow \text{trusts}(A,B),\\
\neg\text{trusting}(A) &\Rightarrow \neg\text{trusts}(A,B),\\
\text{trustworthy}(B) &\Rightarrow \text{trusts}(A,B),\\
\neg\text{trustworthy}(B) &\Rightarrow \neg\text{trusts}(A,B).
\end{aligned}$$

PSL-Personality

maintains predicates for how trusting and trustworthy users are. Trusting individuals are prone to trust others, while trustworthy individuals are likely to receive trust.

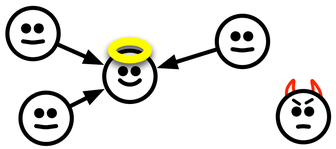

$$\begin{aligned}
\text{sameTraits}(A,B) &\Rightarrow \text{trusts}(A,B),\\
\neg\text{sameTraits}(A,B) &\Rightarrow \neg\text{trusts}(A,B),\\
\text{trusts}(A,B) \wedge \text{sameTraits}(B,C) &\Rightarrow \text{trusts}(A,C),\\
\neg\text{trusts}(A,B) \wedge \text{sameTraits}(B,C) &\Rightarrow \neg\text{trusts}(A,C),\\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

PSL-Similarity

models the correlation between feature similarity and trust, where the sameTraits predicate is computed as a normalized cosine similarity of users' local features.

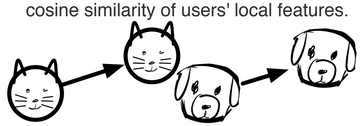

FilmTrust Experiment

Data description and experiment setup

FilmTrust: social movie recommendation service

Records users' trust scores for other users (1 to 10), users' ratings for movies (1 to 5)

500 users in largest connected component, 1574 user-user trust ratings

Four-fold cross validation for learning weights and parameters

Histograms of predicted trust vs. ground truth

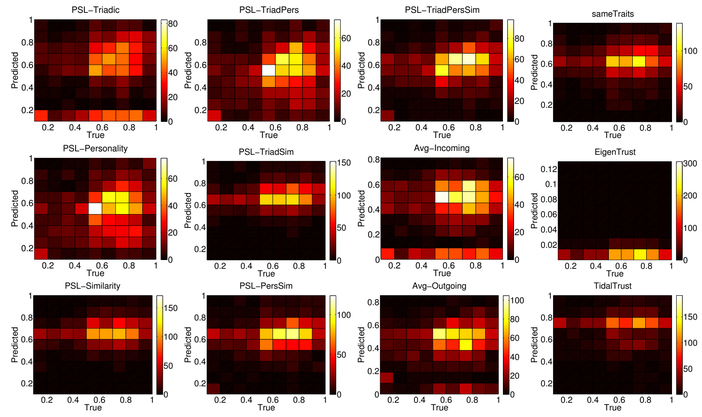

Average scores of trust predictions

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

(MAE: average error, τ: Kendall-tau, ρ: Spearman's rank correlation)

Discussion

PSL is a flexible tool for exploring social trust analysis

Soft first-order logic is convenient and effective for modeling the amount of trust between individuals

Future work: latent trust for social sentiment, larger-scale models of group trust and joint group membership