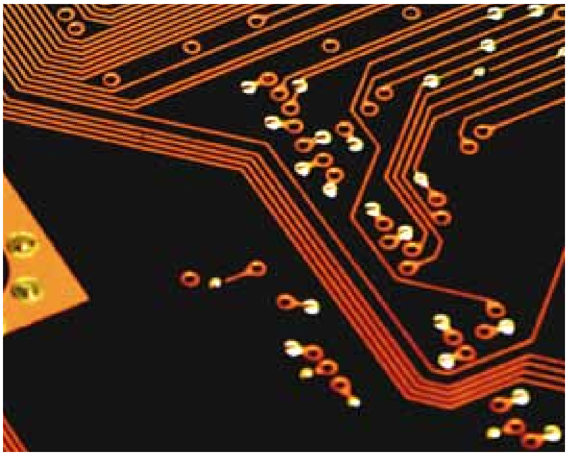

Shin: Generalized Trust Propagation with Limited Evidence

Chung-Wei Hang, Zhe Zhang, and Munindar P. Singh, *North Carolina State University*

Shin incorporates a probabilistic method for revising trust estimates in trustees, yielding higher prediction accuracy than traditional approaches that base trust exclusively on a series of referrals culminating with the trustee.

In e-commerce and social networks, trust—the truster's belief that interaction with the trustee will yield the expected outcome[1]—relies on interaction among autonomous agents. Ideally, a truster would base its trust in a trustee on evidence, or experience with the trustee, but prior interactions in dynamic settings, such as Amazon, tend to be rare.

To interact in these kinds of environments, an agent must rely on referrals to agents that pass on testimony of their experiences with the trustee. These agents might be product users that the prospective buyer trusts because of their stated views on other products that the buyer is familiar with. If its reviews coincide with the buyer's experience with those products, the agent is more likely to trust them.

Existing trust propagation approaches exploit this idea, using a series of referrals from one agent to the next,[2,3] with each such referral corresponding to a trust assessment. Most propagation approaches assume that trust propagates forward through a path of intermediary agents, with each agent trusting the next. These relationships naturally form a trust network—a weighted directed graph in which vertices represent agents, and edges represent directed trust relationships weighted by trust level. The outcomes of prior interactions affect each edge weight, and a truster can evaluate these assessments to decide whether or not to interact with a prospective trustee. In a forward path, the last intermediate agent is considered a witness for the trustee.

As an example of a forward-only propagation path, suppose that Charlie, Dennis, Elisa, and Flora are on the same social network. Charlie messages Dennis that he is looking for a real estate broker, and Dennis tells Charlie about his coworker, Elisa, who recently purchased a house in the area. Charlie contacts Elisa, who tells him about Flora, her real estate broker, whom she liked very much. On the basis of Elisa's referral, Charlie signs up with Flora. This propagation of trustworthiness estimation is a strictly forward path: that is, Charlie trusts Dennis, who trusts Elisa, who trusts Flora; thus, Elisa becomes the witness for Charlie, and he trusts her enough to use Flora's services as a real estate broker.

However, trust propagation does not always move in this forward direction, particularly in networks that involve rating a product or service. Rather, the decision to interact with a trustee often relies on a backward path from a rated entity to a rater. For example, in deciding to buy a monitor, John might look at what others have said about another product with which he is familiar, say a laptop. He would then look for common viewpoints on that laptop as a way to decide if he can trust those raters' reviews of the monitor he plans to buy. Existing propagation approaches cannot predict trustworthiness in these kinds of scenarios because in such cases trust assessment involves moving backward (from the laptop to the rater), not forward.

To address the need for trust, even when no suitable forward path exists, we created Shin (the Chinese word

 Published by the IEEE Computer Society

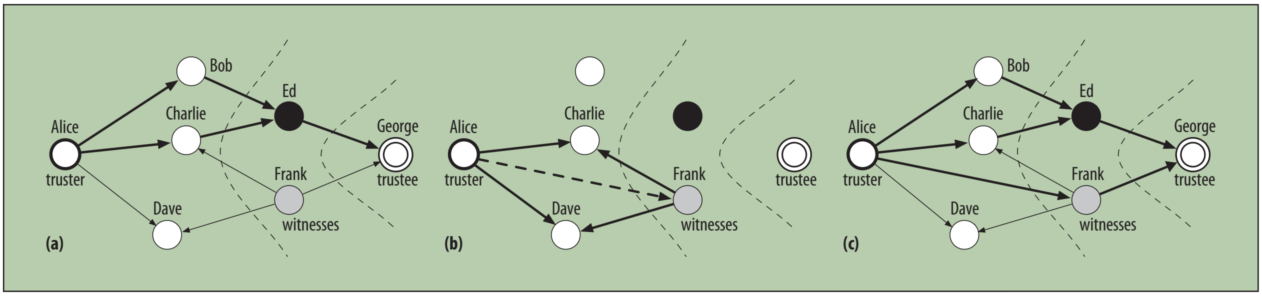

Figure 1. How trust propagation differs in Shin. The bold, black, and double-circled vertices represent the truster, witnesses, and trustee, respectively. Bold arrows indicate trust propagation. (a) Traditional trust propagation (bold edges) follows from truster (Alice) to trustee (George) through a reachable witness (Ed). (b) Shin estimates Alice's trust in an unreachable witness (Frank) by comparing Frank and Alice's ratings of common acquaintances Charlie and Dave. (c) For propagation, Shin treats witnesses Ed and Frank equally.

for trust), a generalized propagation technique that uses a probabilistic paradigm to estimate trust by comparing assessments from acquaintances that the truster and trustee have in common. In developing Shin, we included two of CertProp's three trust propagation operators, but we also extended CertProp to improve prediction accuracy for backward paths. Our evaluation of Shin's capabilities shows that it is superior to CertProp and other existing approaches when only a few trustworthy forward paths exist from the truster to the trustee.

FORWARD-ONLY SHORTCOMINGS

A witness is a party that can provide the truster with testimony about the trustee. A witness is reachable if a trustworthy (generally short) forward path links the truster to the witness.[2] Existing approaches fail in two important settings: when a truster cannot find a trustworthy path to reachable witness and when reachable witnesses do not exist.

No trustworthy path

A truster must be able to evaluate the witness's trustworthiness, which generally means having a forward path to that witness. Even if such a path exists, it might not be trustworthy, leading to inaccurate trust assessments. For example, suppose Mary is a recruiter looking for a potential software engineer. Her accountant, Jeff, refers her to a job candidate, William. Mary does not fully trust William because she does not trust Jeff's expertise in assessing software engineers. If she hires William, she is essentially following an untrustworthy path. Instead, she decides to interview William to gauge his expertise herself.

In this scenario, there is no quality forward path from Jeff to William, but Mary can still evaluate William's trustworthiness using an alternative method. However, Jeff's referral is crucial because it leads Mary to interview a reachable witness, William, whom she could not know without Jeff. Thus, a forward path from Mary to William exists, but it is not in and of itself trustworthy.

No reachable witnesses

In user-item rating networks, an edge corresponds to a user rating an item: no edges end at users or originate from items. A user can reach no other users and only a few items. For example, suppose Anne needs a recommendation for a smartphone, but she does not know anyone who has such a phone. Without reliable witnesses, Anne turns to reviews on smartphone forums. She does not know the reviewers directly, but begins to trust certain reviewers after reading their opinions on phones that she has also used. Thus, Anne propagates trust through reviews that share her opinions on reachable phone models (those she has experience with) and uses these reviews to identify unreachable new phone models.

Shin uses trust in common acquaintances to estimate the trustworthiness of unreachable witnesses. In Figure 1a, information about the subsequent agent's trustworthiness propagates in a strictly forward path (bold edges). In Figure 1b, the truster establishes trust in an unreachable witness by comparing the trust relationships to the acquaintances it has in common with the trustee. In Figure 1c, Shin propagates trust through all witnesses.

GENERALIZED TRUST PROPAGATION

Propagation networks have either a centralized reputation system, which consolidates the trust network, or a decentralized setting, in which each agent knows only its outer edges. Shin is decentralized in that each truster unilaterally pursues referrals and estimates its trust in others.

Shin is based on the idea that it is possible to compute the trust relationship between a truster and trustee using the known direct trust relationships between agent pairs in the network that are proximal to the truster and trustee.

Mathematically, a trust network $T(V,E,d)$ captures agents as vertices V and direct trust relationships as directed, weighted edges E, with the weight $d(a,b)$ of an edge from a to b expressing the amount of direct trust placed by truster a in trustee b. Shin measures direct trust as a value between zero and one, and assigns a trust network an edge if and

TRUST AS EVIDENCE AND BELIEF REPRESENTATIONS

From previous work,[1,2] we define trust in dual representations of evidence and belief. First, Alice's trust in Bob is modeled as the pair $\langle r,s\rangle$, where r represents Alice's positive experiences with Bob and s represents her negative ones. The probability of the next outcome being positive is

$$\alpha=\frac{r}{r+s}$$

and confidence in the probability is

$$c=\frac{1}{2}\int_0^1\left|f\left(x\middle|\langle r,s\rangle\right)-1\right|dx$$

where $f(x|\langle r,s\rangle)$ is the conditional probability of a positive outcome given $\langle r,s\rangle$ and is defined as

$$f\left(x\middle|\langle r,s\rangle\right)=\frac{x^r(1-x)^s}{\int_0^1 x^r(1-x)^s\,dx} \tag{1}$$

and the certainty of $\langle r,s\rangle$ is the probability mass above $\langle 0,0\rangle$, the no-evidence distribution.

Belief-based trust is a triple $\langle b,d,u\rangle$ of belief, disbelief, and uncertainty, where $b = c\alpha$, $d = c(1 - \alpha)$, and $u = 1 - c$, and evidence and belief map to each other.[2]

Consider a scenario that requires trust in unreachable witnesses. Charlie is a common acquaintance of Alice and Frank. Let Alice's trust in Charlie be $\langle r_{AC},s_{AC}\rangle$, and Frank's trust in Charlie be $\langle r_{FC},s_{FC}\rangle$. From these definitions, Alice's trust in Frank is $\langle 1 - q,q\rangle$, where q captures Frank's disagreement from Alice's perspective.[3]

From Equation 1, Frank's trust in Charlie exhibits the distribution $f(x|\langle r_{FC},s_{FC}\rangle)$. Alice's trust in Charlie carries the probability α_{AC}. From previous work,[3] we define the average error q as the integral over x of $f(x|\langle r_{FC},s_{FC}\rangle)(x - \alpha_{AC})^2$ and approximate it by

$$1-\sqrt{\left(\alpha_{AC}-\frac{(r_{FC}+1)}{(r_{FC}+s_{FC}+2)}\right)^2+\left(\frac{(r_{FC}+1)(s_{FC}+1)}{(r_{FC}+s_{FC}+2)^2(r_{FC}+s_{FC}+3)}\right)}$$

For each acquaintance that Alice and Frank have in common, we compare trust reports, computing a piece of evidence of Alice's trust in Frank as $\langle 1 - q,q\rangle$. With N common acquaintances, Alice's trust in Frank is the aggregation

$$\left\langle\sum_{n=1}^{N}(1-q_n),\sum_{n=1}^{N}q_n\right\rangle.$$

The complexity of calculating the trustworthiness of all unreachable witnesses approximates $O(uw \times cr)$, where uw is the average number of unreachable witnesses, and cr is the average number of common relations per witness. Complexity depends not only on evaluating the unreachable witnesses, but also on the choice of the forward propagation method and the propagation depth.

References

1. A. Jøsang, "A Logic for Uncertain Probabilities," *Int'l J. Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 9, no. 3, 2001, pp. 279-311.
2. Y. Wang and M.P. Singh, "Evidence-Based Trust," *ACM Trans. Autonomous and Adaptive Systems*, Nov. 2010, pp. 14:1-14:28.
3. Y. Wang, C.W. Hang, and M.P. Singh, "A Probabilistic Approach for Maintaining Trust Based on Evidence," *J. Artificial Intelligence Research*, vol. 40, 2011, pp. 221-267.

only if the corresponding direct trust is nonzero. In addition, Shin computes and uses the function t: $V \times V \rightarrow [0, 1]$ such that for $a,b \in V$, $t(a, b)$ is the amount of (direct or indirect) trust that truster a places in trustee b.

In simple terms, trust propagation is the problem of computing the amount of trust for a nonadjacent truster and trustee, or $t(a,b)$. As part of that computation, Shin uses CertProp's concatenation operator (⊗), which discounts trust values along a referral path, and its aggregation operator (⊕), which combines trust from referral paths.

The "Trust as Evidence and Belief Representations" sidebar describes the mathematical background of Shin's propagation approach in more detail.

Trust propagation through reachable witnesses

The scenario in Figure 1a is one context for illustrating how these operators work with reachable witnesses. Alice's trust in Ed through Bob is t_1(Alice, Ed) = t(Alice, Bob) ⊗ t(Bob, Ed). Alice's trust in Ed through Charlie is t_2(Alice, Ed) = t(Alice, Charlie) ⊗ t(Charlie, Ed). (Subscripts denote first and second trust paths.) To determine t(Alice, Ed), Alice uses Shin to aggregate t_1(Alice, Ed) and t_2(Alice, Ed) by t_1(Alice, Ed) ⊕ t_2(Alice, Ed). Then Alice's trust in George through Ed equals t_1(Alice, George) = t(Alice, Ed) ⊗ t(Ed, George).

Establishing trust in unreachable witnesses

Shin establishes trust in an unreachable witness by comparing the extent of trust the truster places in a common acquaintance with the trust the witness places in the same acquaintance. This idea is based on the trust updating method described in the sidebar and leads to an improved estimate of the witness's trustworthiness as a referrer.

The evaluation consists of comparing the agent's experience with a provided referral. The closer the referral is to the truster's own experience, the more trustworthy the referrer becomes in the truster's view. Similarly, the truster evaluates an unreachable witness's trustworthiness on the basis of how much trust the witness has in a common acquaintance. Shin then aggregates the truster's estimates about the witness's trustworthiness after considering all common acquaintances. For example, in Figure 1b, Alice uses Shin to determine Frank's trustworthiness t(Alice,

Table 1. Summary of the datasets in Shin evaluation.

Network	Type	Vertices	Edges	Weights (d)	Trust (t)
FilmTrust	Social: user to user	528	823	{1,2,3,...10}	$\langle d - 1, 10 - d \rangle$
Advogato	Social: user to user	5,406	51,839	{1,2,3,4}	$\langle d - 1, 4 - d \rangle$
Epinions	Bipartite: user to item	1,000 + 139,738	105,754	{1,2,3,4,5}	$\langle d - 1, 5 - d \rangle$

Table 2. Comparison of Shin with other trust propagation approaches in social networks.

	Prediction accuracy (root mean squared error)					
Dataset tested	Shin	CertProp	TidalTrust	Naïve	Average total witnesses	Average unreachable witnesses
FilmTrust	0.250	0.251	0.281	0.468	8.12	0.21
Advogato	0.1875	0.1875	0.245	0.385	67.17	1.71

Frank) by comparing t(Alice, Charlie) and t(Frank, Charlie), and comparing t(Alice, Dave) and t(Frank, Dave).

As Figure 1c shows, once Alice determines t(Alice, Frank) with Shin, she can calculate her trust in George through Frank as t_2(Alice, George) = t(Alice, Frank) $\otimes$ t(Frank, George). Finally, Alice can apply Shin to determine t(Alice, George) as t_1(Alice, George) $\oplus$ t_2(Alice, George).

FEASIBILITY STUDY

We evaluated Shin using scaled-down versions of the FilmTrust, Advogato, and Epinion networks. FilmTrust (http://trust.mindswap.org) is a social network of film buffs; its edges represent a user's estimation of the quality of another user's taste in movies. Advogato (www.advogato.org) is a free resource for software developers and a research testbed for group trust metrics and other social networking technologies.[4] The Epinions.com dataset (www.trustlet.org/wiki/Epinions_dataset) contains consumers' product reviews, including numerical ratings. In our bipartite representation, vertices are consumers and products, and edges are the reviews.

Table 1 shows the datasets and their corresponding weight-trust translations. To model datasets as trust networks, we used a linear transformation[3] to translate edge weights $d(u,v)$ (single integer) to our trust representation $t(u,v) = \langle r,s \rangle$. To keep our sample sizes manageable, we used random-walk sampling for Advogato and only the first 1,000 users in the Epinions dataset.

Social networks

To evaluate selected trust propagation approaches using the FilmTrust and Advogato datasets, we applied the leave-one-out cross-validation technique, viewing a single edge $e = (u,v)$ in a given dataset $G(V,E)$ as a test set in itself.

We first removed e from E to construct the corresponding training set and then used the training set to calculate the propagated trust between truster u and trustee v, $t(u,v)$. Finally, we measured each approach's effectiveness by comparing $t(u,v)$ with the actual trust value of e, the edge we had removed. When an approach failed to make a prediction, we treated the mean value of edge weights as the prediction (5.5 for FilmTrust and 2.5 for Advogato).

Table 2 shows how Shin compares with TidalTrust,[2] CertProp,[3] and Naïve on the FilmTrust and Advogato datasets. TidalTrust is a referral-based propagation approach that uses a recursive search with a weighted average of the trust assessments along the referral paths. The weighted average captures the notion that the trust assessments provided by more trustworthy referrers are weighted more heavily than those provided by untrustworthy referrers. Naïve takes the product of ratings along a referral path and computes propagated trust by averaging trust from multiple referral paths.

To measure accuracy, we treat the propagated trust $t(u,v)$ as a prediction and compare it with the known (but previously withheld) direct trust value $d(u,v)$. Each prediction thus yields an error value: its difference with the actual. Our accuracy metric is the well-known root mean squared error (RMSE).

As the table shows, Shin and CertProp significantly outperform TidalTrust and Naïve, but Shin yields similar results as CertProp because both social network datasets have relatively few unreachable witnesses on average: only 0.21 out of 8.12 witnesses in FilmTrust and 1.71 out of 67.17 in Advogato. Shin outperforms CertProp when there are few trustworthy reachable witnesses.

Bipartite network

To evaluate the Epinions dataset, we used threefold cross-validation. We randomly divided the edges in our dataset (first 1,000 users in the original dataset) into three

SHIN VERSUS COLLABORATIVE FILTERING

We compared Shin to collaborative filtering (CF) on Epinions, a bipartite dataset.

Figure A shows the prediction error when we varied the number of common trust relations considered in evaluating a witness. Shin generally outperforms CF when the evidence is limited, and it is robust against the number of common trust relations considered. In contrast, CF requires more evidence to improve accuracy.

Figure B shows the results with a varied number of witnesses considered, where witness evaluation is based on no more than two common trust relations. In this evaluation, Shin's prediction rate was more accurate with more considered witnesses, whereas CF was less sensitive to the number of considered witnesses. Shin is an evidence-based approach that aggregates evidence without modifying it, so more evidence produces more accurate predictions. CF adjusts evidence on the basis of correlation, so its prediction accuracy depends on the correlation's accuracy rather than the amount of evidence.

We further evaluated the accuracy of the trust established between trusters and witnesses. An Epinions.com user can explicitly state that he trusts the trustee's reviews. We used these explicit statements as the test set to verify the trust we computed from the ratings (as the training set).

For any two users *u* and *v* with at least one common rated acquaintance, we established a trust relationship from *u* to *v* on the basis of their respective trust relationships with their common acquaintances. We thus built a trust network where edges exist between any two users with common acquaintances. We then compared the built trust network with the test set. A description of the comparison method is available at http://research.csc.ncsu.edu/mas/code/trust/Shin/confusion-table.pdf.

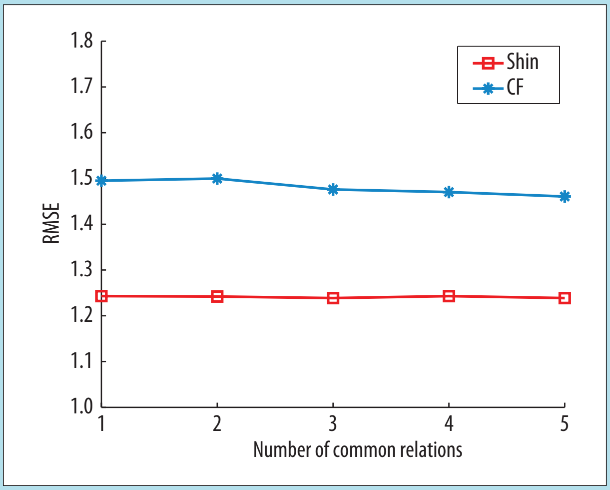

Figure A. ***Comparison of Shin and collaborative filtering (CF) on Epinions with limited evidence. Both Shin and CF make predictions by evaluating three witnesses on the basis of one to five common trust relations. RMSE: root mean squared error.***

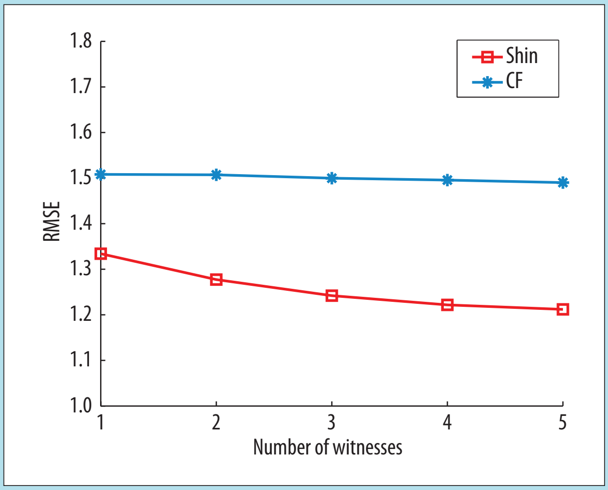

Figure B. ***Comparison of Shin and CF on Epinions with limited evidence. Both Shin and CF make predictions by evaluating only one to five witnesses. Evaluation is based on no more than two common trust relations. RMSE: root mean squared error.***

independent subsets of equal size and used two as the training set and the third as the test set. For each user-item pair in the test set, we considered either all or at most five random (to simulate limited-evidence scenarios) witnesses, each of which we evaluated on the basis of all or at most five random common trust relations.

We calculated the RMSE of trust predictions by propagating trust through all user-item pairs in each test set and then averaging RMSE over all the test sets. We then compared Shin with the user-oriented neighborhood-based collaborative filtering (CF) approach,[5] which uses Pearson correlation to measure rating predictions that are based on a weighted average of ratings from users with similar tastes. We chose this approach because it considers users individually.

Our results show that Shin yields more accurate predictions (RMSE = 1.12) than CF (RMSE = 1.35) when we considered all available witnesses and common relations. Shin also provides solid predictions with at most five random common trust relations, yielding an RMSE of 1.25, relative to CF's RMSE of 1.50.

As the "Shin versus Collaborative Filtering" sidebar describes, Shin effectively evaluates a witness's trustworthiness even with limited common trust relations, and its accuracy improves as the number of witnesses increases.

Sparse networks

To further demonstrate Shin's advantages, we modified CertProp and Shin so that each considered only a fraction, λ, of the reachable witnesses (treating the other witnesses as unreachable). Our aim was to consider settings in which social networks are sparse because they are still emerging or exist only in specialized domains. We conjecture that traditional methods will be less useful than Shin in such settings because their predictions rely on a witness-rich network.

As a simple example, assume we have 28 witnesses of trustee v, with eight being unreachable and 20 reachable from truster u. If λ is 0.55, CertProp and Shin propagate trust through only nine of the 20 reachable witnesses. In addition to forward propagation to the nine witnesses, Shin tries to establish trust from u to 19 witnesses, treating 11 reachable witnesses as unreachable plus considering all eight truly unreachable ones. If Shin successfully builds trust in five of the 19 witnesses, then it has successfully estimated trust using 14 (9+5) witnesses. CertProp, on the other hand, has calculated propagated trust with only nine witnesses.

In evaluating sparse networks, we varied λ from 0 to 100 percent. Figures 2 and 3 show the results for FilmTrust and Advogato. In Figures 2, 3, and 4, the lines refer to RMSE, and the regions refer to the number of witnesses. We compared Shin with TidalTrust as well as CertProp.

In both figures, as λ increases, the number of reachable witnesses decreases, and the number of evaluable witnesses increases. CertProp's accuracy

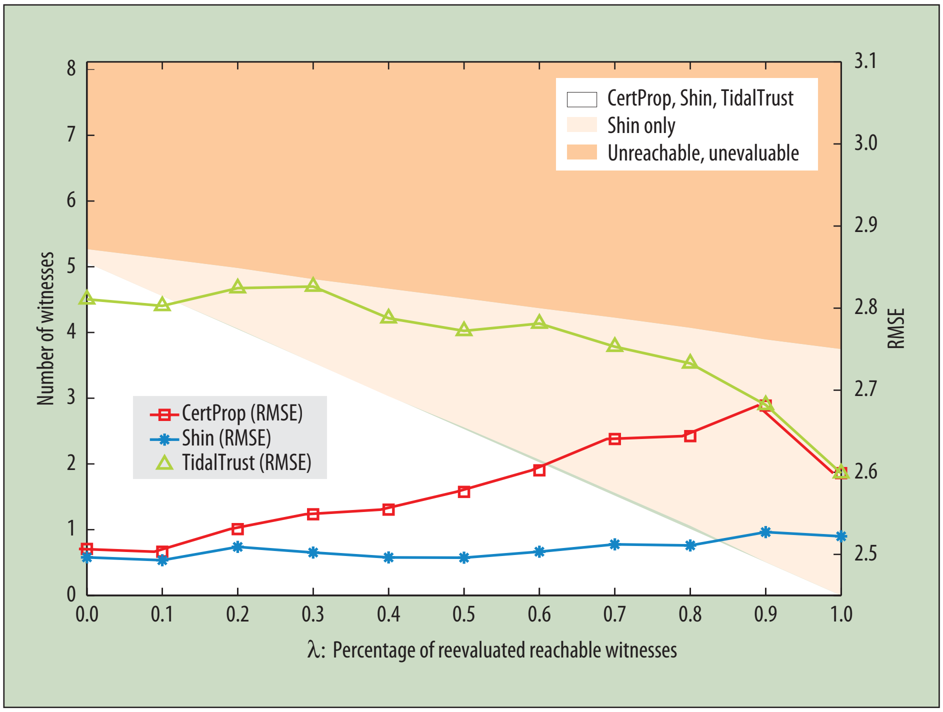

Figure 2. Comparison of Shin with CertProp and TidalTrust for the FilmTrust network when the percentage of unreachable witnesses varies. Unreachable witnesses are those for which the approach must consider backward paths only. Shin consistently outperforms CertProp and TidalTrust, which rely on networks with a rich set of reachable witnesses. RMSE: root mean squared error.

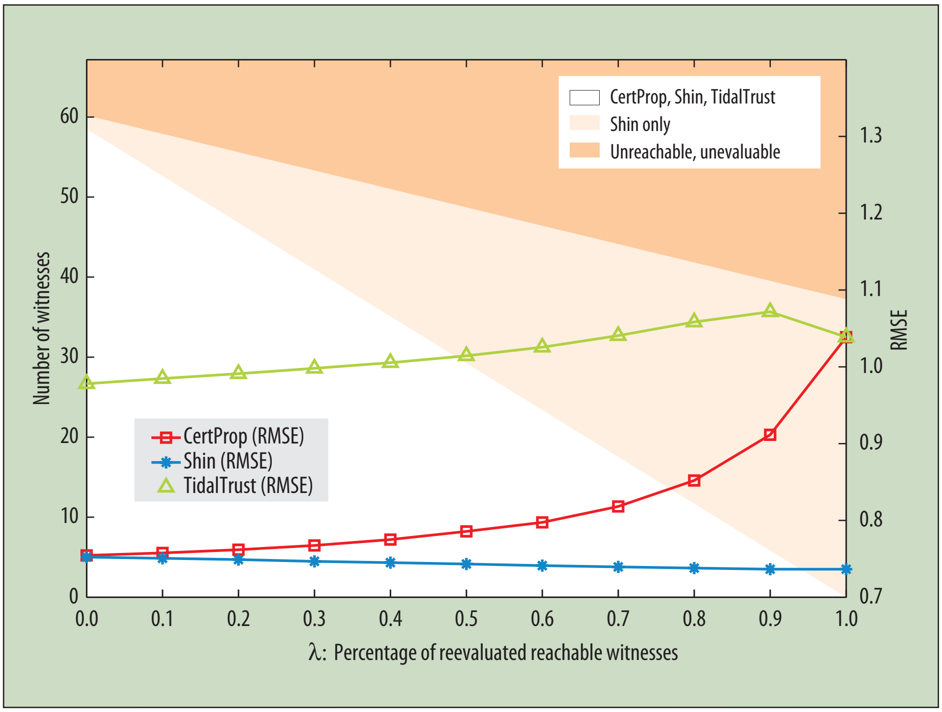

Figure 3. Comparison of Shin with CertProp and TidalTrust for the Advogato network when the percentage of unreachable witnesses varies. Despite the small number of reachable witnesses, Shin yields more accurate predictions by benefiting from additional evaluable witnesses. RMSE: root mean squared error.

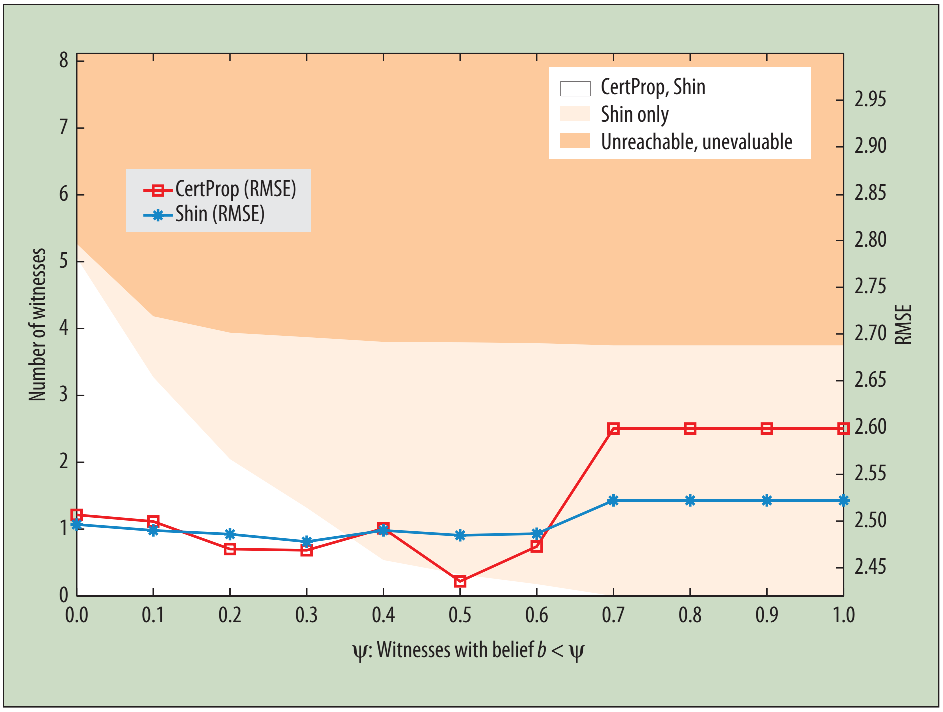

Figure 4. Comparison between Shin and CertProp with various values of Ψ in FilmTrust. By reevaluating the witnesses reached by untrustworthy referral paths, Shin yields more accurate prediction when reevaluating all witnesses ($\Psi \geq 0.7$). RMSE: root mean squared error.

decreases with fewer witnesses, but Shin benefits from additional evaluable witnesses and can thus provide stable predictions regardless of λ's value. TidalTrust produces the highest error rate. As λ increases, TidalTrust's accuracy improves because when TidalTrust fails to provide a prediction without a sufficient number of witnesses, simply guessing the mean value, 5.5, yields better accuracy.

We define an untrustworthy path as a referral path with belief b of concatenated trust lower than a threshold ψ. As described in the "Trust as Evidence and Belief Representations" sidebar, given a trust value, $\langle r,s \rangle$, $b = c\alpha$, where c is the trust value's certainty and α is $r/(r + s)$, belief incorporates both certainty and probability—it is the product of the two numbers.

To demonstrate how Shin improves trust predictions, we reevaluated the trustworthiness of witnesses that untrustworthy referral paths reach. Figure 4 shows how Shin and CertProp prediction compare with different ψ thresholds for the FilmTrust network. (We do not include TidalTrust in this comparison because its trust representation uses a different belief measurement.) Shin ignores the untrustworthy referral paths ($b < \psi$) and reevaluates the trustworthiness of the witnesses those paths reach by comparing the trust values to the common acquaintances between the truster and the witnesses.

When reevaluating all witnesses ($\psi \geq 0.7$), Shin's prediction accuracy is higher than CertProp's. As the white area in Figure 4 attests, only a few paths have a ψ threshold less than or equal to 0.7, where CertProp and Shin produce similar results (although Shin is less volatile with respect to ψ). Our results imply that in FilmTrust, trust propagation is more effective when witnesses are evaluated in terms of common relations with the truster rather than by the referral paths.

Figure 5 shows a similar comparison for Advogato. When $\psi \in [0.2, 1.0]$, reevaluating the witnesses reached by referral paths with $b < \psi$ yields better prediction. In Advogato, no paths have a b greater than 0.5. Consequently, in Advogato, the witnesses reached by trustworthy paths help, but the witnesses reached by untrustworthy paths can sometimes worsen the predictions, relative to reevaluating every reached witness.

Traditional trust propagation suffers when a trustee's witnesses are unreachable or are reachable only by untrustworthy referrals. Shin considers those witnesses by establishing trust between them and the truster agent. Shin is a decentralized approach in that a truster applies it to estimate the amount of trust to place in a trustee according to estimates from a small number of agents in proximity to the trustee.

We have evaluated Shin over both bipartite datasets in which all witnesses are unreachable and unipartite datasets in which some witnesses are reachable. In bipartite datasets, Shin builds on the trust relationships established in social networks, providing results competitive with collaborative filtering even with little evidence. In unipartite datasets, Shin provides more accurate results than CertProp and TidalTrust for a sparse network as the ratio of reachable to unreachable witnesses decreases.

One limitation of our work is that Shin models trustworthiness as two parameters whereas many real systems only capture one value. Translating between one-value and two-value representations could result in losing valuable information, which might degrade prediction accuracy. Despite this potential loss, we view the two-value representation as superior and recommend that future systems incorporate it. A two-value representation captures not only the trust level, but also the amount of evidence (the confidence) supporting that trust.

Our next step will be to formalize network properties that can serve as useful indicators of Shin's effectiveness relative to traditional approaches. For example, in FilmTrust, most witnesses are reachable through referrals, so one of our aims is to explore how FilmTrust differs from networks where witnesses are mostly unreachable. Retrieving frequency patterns from a network is a possible investigation path. For example, a witness in a network with a more frequent path-of-three pattern could be more reachable than a witness in a network where such a pattern occurs less frequently. Understanding such patterns could be the basis for selecting the most appropriate trust propagation approach in a particular scenario.

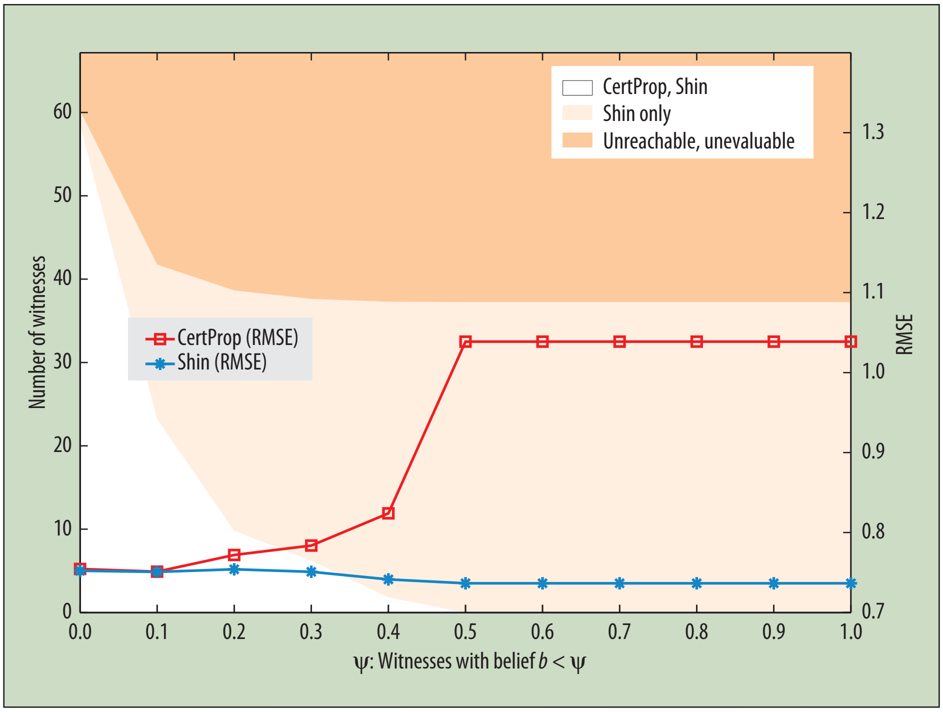

Figure 5. Comparison between Shin and CertProp with various values of Ψ in Advogato. By reevaluating the witnesses reached by paths with $b < \Psi$, Shin yields more accurate prediction when $\Psi > 0.1$. RMSE: root mean squared error.

Acknowledgments

We thank Victor Lesser for useful discussion and the reviewers for their comments.

This work was supported in part by the US Army Research Laboratory through its Network Sciences Collaborative Technology Alliance (NS-CTA) under Cooperative Agreement W911NF-09-2-0053 and by the US Army Research Office (ARO) under grant W911NF-08-1-0105 managed by NCSU Secure Open Systems Initiative (SOSI).

References

1. S. Paradesi, P. Doshi, and S. Swaika. "Integrating Behavioral Trust in Web Service Compositions," Proc. 7th Int'l Conf. Web Services (ICWS 09) IEEE, 2009, pp. 453-460.
2. J. Golbeck, "Computing and Applying Trust in Web-Based Social Networks," PhD thesis, Univ. of Maryland, College Park, 2005.
3. C.W. Hang, Y. Wang, and M.P. Singh, "Operators for Propagating Trust and Their Evaluation in Social Networks," Proc. 8th Int'l Conf. Autonomous Agents and Multiagent Systems (AAMAS 09), ACM, vol. 2, 2009, pp. 1025-1032.
4. J. Leskovec and C. Faloutsos, "Sampling from Large Graphs," Proc. 12th Int'l Conf. Knowledge Discovery and Data Mining (KDDM 06), ACM, 2006, pp. 631-636.
5. J.L. Herlocker et al., "An Algorithmic Framework for Performing Collaborative Filtering," Proc. 22nd Int'l ACM Conf. Research and Development in Information Retrieval (SIGIR 99), ACM, 1999, pp. 230-237.

***Chung-Wei Hang** is a postdoctoral researcher in computer science at North Carolina State University. His research interests include multiagent systems, social computing, and service-oriented computing. Hang received a PhD in computer science from North Carolina State University. Contact him at chang@ncsu.edu.*

***Zhe Zhang** is a doctoral student in computer science at North Carolina State University. His research interests include social networks, text mining, and multiagent systems. Zhang received an MS in computer science from North Carolina State University. Contact him at zzhang13@ncsu.edu.*

***Munindar P. Singh** is a professor of computer science at North Carolina State University. His research interests include multiagent systems and service-oriented computing, emphasizing the study of interactions among autonomous parties. Singh received a PhD in computer science from the University of Texas at Austin. He is an IEEE Fellow. Contact him at singh@ncsu.edu.*