Data Science and Engineering (DSE) Record, Volume 7, issue 1.

Regularized Matrix Factorization for Trust-Aware Recommender System

Paramate Phuengtrakul [1], Jakramate Bootkrajang [2] and Dussadee Praserttipong [2]

[1] Master's Degree Program in Data Science, Chiang Mai University, Chiang Mai, Thailand
paramate_ph@cmu.ac.th

[2] Department of Computer Science, Faculty of Science, Chiang Mai University,
Chiang Mai, Thailand
jakramate.b@cmu.ac.th, dussadee.p@cmu.ac.th

Abstract. Trust-aware recommender systems leverage social trust to mitigate rating sparsity and the cold-start problem, yet most public trust datasets represent trust as sparse binary links, which can underutilize structural information in the trust network. This paper proposes a Katz-based trust enrichment method that transforms binary directed trust into continuous multi-hop trust signals — capturing friend-of-friend propagation via truncated path counting — and integrates these signals into TrustSVD to improve recommendation accuracy while maintaining a practical accuracy–complexity trade-off. The proposed method is instantiated as four variants: Katz-2, Katz-3, and their corresponding boosted variants (Boosted Katz-2 and Boosted Katz-3), which differ in propagation depth and whether direct trust edges are re-emphasized after propagation. To characterize the value of multi-hop propagation, the proposed method is evaluated against three reference representations: the binary baseline used by the original TrustSVD and two local-overlap benchmarks (cosine and Jaccard similarity) that capture only direct neighborhood agreement without path propagation. Using FilmTrust and CiaoDVD, the study evaluates all seven trust representations under a unified training and hyperparameter tuning protocol, with performance reported via 5-fold cross-validation on RMSE and MAE for both all-user and cold-start user groups (rating < 10 and rating < 5). Results show that the proposed Katz-based method yields modest but consistent accuracy improvements over the binary baseline, with the clearest benefits in cold-start settings and in the sparser CiaoDVD dataset. Across settings, Katz-2 emerges as the most reliable variant of the proposed method, whereas the most extreme cold-start group in CiaoDVD (rating < 5) slightly favors the Jaccard benchmark. Given that training cost is dominated by repeated SGD updates and increases with enlarged effective trust neighborhoods, Katz-2 offers a strong default balance between accuracy gains and computational overhead.

Keywords: Trust-Aware Recommender System, Matrix Factorization, TrustSVD, Katz Index, Cold-Start Problem.

1 Introduction

Recommendation systems are widely utilized across a wide range of online applications, including online shopping, video streaming, news feeds, online advertising, and travel services. These systems are fundamental to personalizingcontent and supporting users as they navigate platforms with vast inventories. With too many available options, users often struggle to identify choices that best match their preferences. In this context, recommendation systems facilitate decision-making by filtering and presenting items that are most likely to be purchased or consumed. Such mechanisms not only improve product discovery and search efficiency but also help users discover new items, thereby enhancing their overall experience on the platform [Raza, 2024].

Recommendation systems have evolved through multiple paradigms, ranging from traditional algorithmic approaches to modern deep learning–based models. Deep learning–based recommender systems have demonstrated strong performance by modeling non-linear latent patterns and leveraging sequential behaviors and graph-structured relationships [He, 2017]. Despite their expressive power, deep learning approaches introduce practical challenges, including high computational cost, limited interpretability, and strong dependence on large-scale data and infrastructure. As a result, it remains pragmatic to revisit and enhance classical recommendation techniques that offer a more favorable balance between accuracy, scalability, and efficiency in realistic settings [Raza, 2024].

Traditional recommender systems have historically relied on two principal paradigms: Content-Based Filtering and Collaborative Filtering. In contrast to content-based approaches, Collaborative Filtering (CF) operates on the assumption that users who exhibited similar behavior in the past are likely to share preferences in the future. However, CF is particularly vulnerable to data sparsity, since users typically interact with only a small subset of items in large catalogs, leading to unreliable similarity estimates and unstable personalization. A closely related issue is the cold-start problem, where new users or new items lack sufficient interaction history to support accurate recommendations [Lam, 2008]. To mitigate these limitations, model-based methods such as matrix factorization learn compact latent representations that generalize better than direct neighborhood similarity under sparsity. Nevertheless, even strong matrix factorization models can remain constrained when rating evidence is extremely limited. Among auxiliary signals beyond ratings, social trust has received substantial attention because trust networks can provide additional evidence sources — grounded in homophily and social influence — for inferring preferences when rating histories are sparse [McPherson, 2001; Jamali, 2010; Ma, 2011].

Within the matrix factorization family, TrustSVD is a widely adopted baseline that extends SVD++-style recommendation by modeling both explicit and implicit influences from user trust relations alongside item-rating effects [Guo, 2015]. While TrustSVD demonstrates strong performance across trust datasets, the broader literature indicates that the benefit of trust integration depends critically on how trust is represented and utilized — especially when trust data itself is sparse and typically recorded as binary links [Guo, 2015; Fazeli, 2014]. Accordingly, the central challenge is not only whether trust should be integrated, but how binary and sparse trust information can be transformed into richer, more reliable signals that capture indirect propagation and heterogeneous influence. Among available transformation strategies, path-based similarity measures such as the Katz index [Katz, 1953] offer a principled mechanism to encode multi-hop connectivity while controlling propagation depth through exponential damping — properties that align closely with the known characteristics of trust transitivity and decay in social networks [Golbeck, 2005; Richters, 2011].

This work investigates a central question: can a Katz-based trust enrichment method — transforming binary trust into multi-hop weighted representations through truncated path propagation — improve the robustness of trust-aware matrix factorization beyond what local-overlap measures can achieve, without incurring prohibitive computational cost? The study pursues three objectives: (i) to determine whether the proposed Katz-based trust enrichment method improves TrustSVD's predictive accuracy (RMSE and MAE) compared with the binary trust baseline on FilmTrust and CiaoDVD, under both all-user and cold-start evaluation; (ii) to determine whether multi-hop trust propagation provides more reliable accuracy gains than local-overlap benchmarks (cosine and Jaccard similarity) across datasets and across cold-start severity levels (rating < 10 and rating < 5); and (iii) to characterize how propagation depth (Katz-2 vs Katz-3) and the boosting variant affect the accuracy–complexity trade-off, and which variant offers the best default balance between gains and training-time overhead.

2 Related Work

Recommender systems serve as information filtering tools that predict user preferences and suggest relevant items from large catalogs. The foundational architecture for collaborative recommendation was introduced by Resnick et al. (1994) through the GroupLens system, which demonstrated that aggregating user opinions could effectively filter Usenet articles and laid the groundwork for modern collaborative filtering [Resnick, 1994]. Traditional recommender systems are commonly classified into three principal paradigms: content-based filtering, collaborative filtering, and hybrid approaches. Burke (2002) provided a seminal taxonomy of hybrid methods, identifying several hybridization strategies that combine the strengths of content-based and collaborative filtering to mitigate issues such as the cold-start problem and overspecialization [Burke, 2002].

2.1 Matrix Factorization for Recommendation

Matrix factorization (MF) has become the dominant model-based collaborative filtering technique due to its ability to learn compact latent representations of users and items while maintaining computational efficiency [Bokde, 2015; Koren, 2009; Hu, 2008]. Koren, Bell, and Volinsky (2009) provided an influential tutorial on MF techniques for recommendation, while Funk (2006) introduced the foundational approach of learning latent factors via stochastic gradient descent on observed ratings [Koren, 2009; Funk, 2006]. Salakhutdinov and Mnih (2007) proposed Probabilistic Matrix Factorization (PMF), placing a probabilistic interpretation on the factorization process [Salakhutdinov, 2007], and Koren (2008) proposed SVD++, a model that integrates explicit and implicit feedback within a single factorization framework [Koren, 2008]. Despite these advances, classical MF models still depend heavily on observed interactions for learning stable user and item representations; in sparse settings this can lead to unstable latent representations and poor performance for users with few ratings, motivating extensions that incorporate auxiliary information such as social trust.

2.2 Trust in Social Networks and the Katz Index

The integration of social trust into recommender systems is grounded in well-established theories of trust in social networks. McPherson, Smith-Lovin, and Cook (2001) documented the principle of homophily — the tendency of individuals to associate with others who are similar to themselves [McPherson, 2001]. Golbeck (2005) formalized key properties of computational trust, including transitivity (if A trusts B and B trusts C, then A may extend some degree of trust to C) and composability [Golbeck, 2005]. Richters and Peixoto (2011) analyzed trust transitivity and showed that trust propagation is not universally reliable and can degrade over longer paths [Richters, 2011], implying that short-range propagation may yield useful signals while deeper propagation may introduce noise.

The Katz index, originally proposed by Katz (1953), computes similarity between nodes by counting the number of paths connecting them, with longer paths exponentially down-weighted by a damping factor [Katz, 1953]. Its application to trust networks is motivated by the observation that users connected through multiple short trust paths may share preferences even without a direct trust edge. Duricic et al. (2018) applied Katz-style reasoning to trust-based collaborative filtering to tackle cold-start settings by leveraging higher-order trust relationships rather than relying solely on direct trust links [Duricic, 2018].

2.3 Trust-Aware Recommender Systems

Trust-aware recommender systems augment traditional collaborative filtering with social trust information to improve recommendation quality, particularly in sparse and cold-start scenarios [Dong, 2022; Wu, 2016]. Early neighborhood-based methods — TidalTrust [Golbeck, 2005], MoleTrust [Massa, 2007], and TrustWalker [Jamali, 2009] — propagate trust directly through the social graph to identify reliable rating sources. MoleTrust showed that trust-based neighbor selection was particularly effective for cold-start users, establishing a key principle: trust is most valuable precisely when collaborative filtering is weakest [Massa, 2007].

Matrix factorization–based trust methods integrate social trust into the latent factor learning process. SoRec [Ma, 2008] jointly factorizes the rating and trust matrices through a co-factorization approach; SocialMF [Jamali, 2010] makes each user's latent factor dependent on the latent factors of their direct social neighbors, naturally enabling trust propagation through transitivity; SoReg [Ma, 2011] proposed a general social regularization framework with similarity-weighted regularization; and TrustMF [Yang, 2013] introduced a dual-role factorization that explicitly models the directional nature of trust. TrustSVD [Guo, 2015] extended the SVD++ framework by incorporating both explicit and implicit influence of user trust alongside implicit influence from rated items, establishing it as one of the strongest trust-aware MF baselines. Formally, TrustSVD predicts the rating of user u on item j as

$$\widehat{r_{u,j}} = \mu + b_u + b_j + q_j^T \left(p_u + |I_u|^{-1/2} \sum_{i \in I_u} y_i + |T_u|^{-1/2} \sum_{v \in T_u} w_v \right) \tag{1}$$

where μ is the global mean, b_u and b_j are user/item biases, p_u and q_j are user/item latent factors, y_i aggregates implicit feedback from items rated by u (the SVD++ component), and w_v aggregates implicit influence from users trusted by u (the trust component). Although TrustSVD demonstrated superior performance compared to ten baseline methods across four benchmark datasets, a structural limitation remains: it treats the trust input as a binary directed graph and weights all observed trust edges equally, assuming that every declared trust statement carries the same predictive value and ignoring indirect relational proximity in the trust network — a limitation that directly motivates the trust enrichment framework proposed here.

Since TrustSVD's introduction, several studies have proposed extensions addressing its remaining limitations. Sun, Yan, and Ren (2022) incorporated trustworthiness weighting into TrustSVD's regularization, demonstrating that not all trust links carry equal information value [Sun, 2022]. Han et al. (2022) proposed a PMF-based model integrating user trust relationships, interest mining, and item correlation [Han, 2022], and Xu et al. (2021) addressed the asymmetric nature of trust by decomposing the trust matrix into separate truster and trustee feature matrices [Xu, 2021]. These findings imply that treating all trust edges equally is suboptimal, particularly under sparse trust conditions, and motivate a trust enrichment method that assigns graded weights through multi-hop path propagation.

3 Proposed Method: Katz-Based Trust Enrichment for TrustSVD

This work proposes a trust enrichment framework that transforms the binary directed trust input used by TrustSVD into continuous trust signals capturing local neighborhood overlap and multi-hop propagation. The proposed framework is designed as a drop-in replacement for the trust input of TrustSVD, allowing the same training procedure to be reused while isolating the effect of trust representation on recommendation accuracy and computational cost.

3.1 Trust Enrichment Framework

The starting point is the explicit directed trust network represented as a binary adjacency matrix A ∈ {0,1}^(|U|×|U|), where $A_{u,v} = 1$ denotes a truster-to-trustee statement u → v. Trust enrichment constructs a weighted trust matrix S from A using trust topology alone and then uses S as the trust input during TrustSVD training. The transformation step explicitly excludes ratings and co-rating information, ensuring that the enriched trust signal isolates the structural contribution of the trust graph. This design enables controlled comparison across trust representations because only the trust signal changes across model variants, so the same TrustSVD training procedure and hyperparameter selection protocol can be applied per representation while differences in performance can be attributed to how each transformation reshapes the effective trust input.

3.2 Katz Similarity (Katz-2, Katz-3)

Local-overlap measures such as cosine and Jaccard similarity capture trust similarity only through directly shared neighbors, limiting their reach to one-hop relational structure; they cannot represent indirect trust — the situation in which two users are connected through chains of intermediate trustees but share no direct neighbor. Path-based propagation addresses this limitation by counting paths of varying lengths between users. The Katz index [Katz, 1953] balances the two well-established properties of trust transitivity [Golbeck, 2005] and the empirical caution that trust propagation degrades over longer paths [Richters, 2011] by summing weighted path counts across multiple lengths and applying exponential damping to longer paths. Formally, the Katz similarity matrix is computed as the weighted sum of adjacency powers up to a maximum path length k:

$$\sigma^{(k_{\max}+1)} = \sum_{k=0}^{k_{\max}} (\alpha\, A)^{k} \qquad (2)$$

where A^k is the k-th power of the adjacency matrix (so $(A^2)_{uv}$ counts the number of 2-step trust paths from u to v, etc.), and $0 < \alpha < 1$ is a damping factor that down-weights longer paths. Intuitively, users connected via multiple short trust paths attain a higher Katz-based similarity, even without a direct link. Two truncated variants are considered: Katz-2, which includes paths of length 1 and 2 (direct trust and two-hop connections), and Katz-3, which includes up to length 3 paths. The damping factor α is set to ensure convergence and reasonable weighting (typically $\alpha < 1/\lambda\_max$, where λ_max is the largest eigenvalue of A). Since the trust network is directed, σ is asymmetric in general, and S_{uv} can be positive even if user u does not directly trust v when they are connected through intermediate trusted users.

3.3 Boosted Katz Similarity

Boosted Katz extends the Katz result by applying additional post-processing to preserve the primacy of direct trust links while still incorporating propagated similarity from longer paths as supplemental information. Specifically, it separates direct and indirect trust, normalizes the indirect component, and then adds the original direct trust links back as stronger signals. The transformation proceeds in four steps. First, the truncated Katz similarity matrix is computed by selecting k_max, propagating similarity through the trust network by accumulating weighted paths up to the selected path length. Second, degree normalization is applied using a degree matrix D to reduce the tendency of Katz similarity to assign higher values to high-degree users. Third, row normalization rescales each user's similarity vector using a maximum norm so that scores are comparable across users. Fourth, propagated similarities are boosted by separating indirect trust from direct trust — retaining similarity values only for user pairs without an explicit trust link, row-normalizing them, and finally constructing the boosted similarity matrix as

$$\widehat{\sigma}_{i,j} = \left(1 - A_{i,j}\right) \cdot \sigma_{\text{norm}}(i,j) + A_{i,j} \qquad (3)$$

This formulation ensures that direct trust links remain strong signals with value 1, while indirect trust relations obtained through propagation are retained as normalized supplemental signals. In this way, Boosted Katz differs from standard Katz by not only propagating trust through multi-hop paths, but also normalizing the propagated structure and explicitly boosting the role of direct trust in the final trust matrix.

3.4 Benchmark Trust Representations

To characterize the relative value of the proposed Katz-based method, three benchmark trust representations are evaluated. Binary trust serves as the reference baseline by using the adjacency matrix itself as the trust signal (S_binary = A). Cosine similarity is included as a local-overlap benchmark that measures normalized overlap between outgoing trust vectors, providing a reference for what can be achieved using only one-hop neighborhood agreement. Jaccard similarity is a second local-overlap benchmark defined as the size of the intersection of u's and v's trust sets divided by the size of their union. Like cosine, Jaccard captures only direct neighborhood overlap and is included to assess whether the proposed multi-hop propagation provides gains beyond what local-overlap measures can achieve.

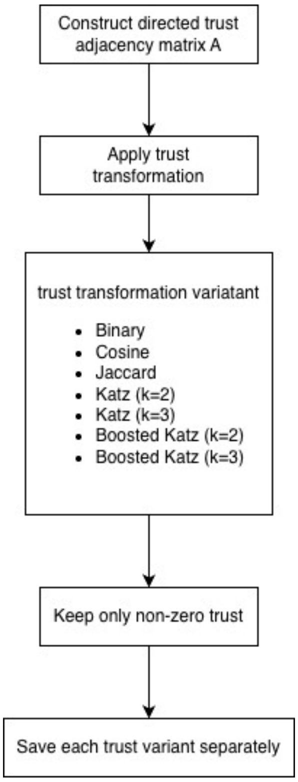

Fig. 1. Trust transformation framework: proposed Katz-based methods and local-overlap benchmarks.

3.5 Integration with TrustSVD and Complexity

The enrichment framework integrates with TrustSVD by replacing the binary trust matrix A with the enriched matrix S as the input to the trust-related components of training. Specifically, S enters the model through (i) the trustee aggregation term in the prediction function, $| T_u |^{-1/2} \sum_{v \in T_u} w_v$, where the effective trust neighborhood T_u is now derived from S rather than A, and (ii) the trust loss term, where the residual is computed against the weighted entries of S rather than binary 0/1 values. This integration preserves all SGD update rules of the original TrustSVD algorithm; the only change is the content of the trust input — its sparsity pattern, weight distribution, and effective neighborhood size — which is precisely the variable of interest in this study.

Training cost is dominated by repeated SGD optimization rather than the one-time trust preprocessing step. The trust loop iterates once per trust edge in $| T |$ and applies d-dimensional updates, so it scales directly with $| T |$, and any increase in the effective trust edges used for training increases trust-loop workload proportionally. To capture this effect, a trust-edge inflation factor is defined as $\rho_T^{(m)} = |T^{(m)}|/|T_0|$, where $|T_0|$ is the number of trust edges under the binary baseline and |T_m| the number used for training under transformation m. Under this definition the trust-loop component increases from $O(d\,|T_0|$ to $O(|T^{(m)}|) = O(dp_T^{(m)}|T_0|)$, making explicit that enrichment increases training time by enlarging the effective trust input optimized by SGD. The benefit and the overhead therefore grow together rather than independently, so the design problem is not whether to accept extra cost but how to obtain the most signal per added edge — which is exactly the balance the proposed Katz-2 variant is selected to strike.

4 Empirical Evaluation

4.1 Experimental Setup

The evaluation compares four proposed Katz variants (Katz-2, Katz-3, Boosted Katz-2, Boosted Katz-3) against three reference representations: the binary baseline (the original TrustSVD trust input) and two local-overlap benchmarks (cosine and Jaccard similarity). All seven trust representations are trained under the same TrustSVD pipeline so that performance differences can be attributed to the choice of trust representation rather than to differences in the learning algorithm. Trust-aware recommendation models are evaluated on two widely used benchmark datasets that provide both user–item ratings and an explicit directed trust network: FilmTrust [Guo, 2013] and CiaoDVD [Guo, 2014]. FilmTrust is comparatively dense (rating density ≈ 1.14%), whereas CiaoDVD is roughly two orders of magnitude sparser and contains a large proportion of cold-start users, making it a challenging testbed for trust-aware methods. Key dataset statistics are summarized in Table 1.

Table 1. Dataset statistics.

	FilmTrust	CiaoDVD
#Rating	35,497	72,345
#Users	1,508	17,615
#Items	2,071	16,121
Rating scale	0.5-4	1-5
Rating density	1.136%	0.025%
#Trust nodes	874	2,433
#Trust edges	1853	22,484
Trust density	0.0024	0.0038
Cold-users < 5	281	15,006
Cold-users < 10	501	16483

All experiments were executed on a containerized environment with an AMD Ryzen 9 9950X 16-core CPU and 32 GB of RAM. Although a GPU was available, all training was performed exclusively on the CPU, because the TrustSVD optimization involves stochastic gradient descent (SGD) updates applied sequentially to each training sample; this per-instance update process is inherently serial and does not benefit from batching. All models were trained using SGD with learning rate 0.001, latent dimensionality 10, and a maximum of 130 epochs. For robust reporting, 5-fold cross-validation was used: ratings were partitioned into five folds, each fold was used once as test data, and results were averaged across folds, with a fixed random seed (42) for reproducibility.

The two TrustSVD regularization terms — λ (rating-related regularization) and λ_t (trust-related regularization) — were tuned via exhaustive grid search over $\lambda \in \{0.5, 0.8, 1.0, 1.2\}$ and $\lambda_t \in \{0.1, 0.5, 0.8, 1.0, 1.2, 1.5, 2.0\}$, using a single 80/20 train–test split for validation. Tuning was conducted separately for each trust transformation variant, since the effective scale and sparsity of the trust signal differ across binary and continuous trust matrices. The combination minimizing validation RMSE was chosen, and models were retrained under the selected hyperparameters using the full 5-fold protocol. In all cases, moderate regularization values (λ around 1.0 and λ_t around 1.0) worked well. Two cold-start groups were defined for evaluation: users with fewer than 10 ratings (Cold<10) and users with fewer than 5 ratings (Cold<5). Models were evaluated using Mean Absolute Error (MAE) and Root Mean Square Error (RMSE).

4.2 Results: All-User and Cold-Start Performance

FilmTrust (all users) results show small but consistent metric differences across trust representations, with the strongest RMSE achieved by Katz-2 and the strongest MAE achieved by Katz-3. Relative to the Binary baseline (RMSE 0.797661; MAE 0.620701), Katz-2 attains the lowest mean RMSE (0.796498), while Katz-3 attains the lowest mean MAE (0.619843), and cosine similarity yields comparable but slightly weaker improvements across both metrics (Table 2). CiaoDVD (all users) results show more uniformly favorable changes from trust enrichment: compared with Binary (RMSE 0.941847; MAE 0.727965), Katz-3 achieves the lowest mean RMSE (0.938368) and the lowest mean MAE (0.725368), with Katz-2 and Jaccard closely following (Table 3).

Table 2. FilmTrust (all users, 5-fold CV; baseline = Binary).

Algorithm	Mean RMSE	Mean MAE	ΔRMSE (%)	ΔMAE (%)
Binary (baseline)	0.797661	0.620701	0.000	0.000
Cosine similarity	0.796721	0.619983	-0.118	-0.116
Jaccard similarity	0.796982	0.620297	-0.085	-0.065
Katz-2	**0.796498**	0.619999	**-0.146**	-0.113
Katz-3	0.797712	**0.619843**	+0.006	**-0.138**
Katz-2 (boost)	0.797006	0.620322	-0.082	-0.061
Katz-3 (boost)	0.796561	0.620093	-0.138	-0.098

Table 3. CiaoDVD (all users, 5-fold CV; baseline = Binary).

Algorithm	Mean RMSE	Mean MAE	ΔRMSE (%)	ΔMAE (%)
Binary (baseline)	0.941847	0.727965	0.000	0.000
Cosine similarity	0.939103	0.725928	-0.291	-0.280
Jaccard similarity	0.938709	0.725620	-0.333	-0.322
Katz-2	0.938496	0.725472	-0.356	-0.342
Katz-3	**0.938368**	**0.725368**	**-0.369**	**-0.357**
Katz-2 (boost)	0.940137	0.726698	-0.182	-0.174
Katz-3 (boost)	0.938952	0.725740	-0.307	-0.306

For FilmTrust cold-start users, Katz-2 provides the best RMSE and MAE. For users with rating < 10, relative to Binary (RMSE 0.811614; MAE 0.640427), Katz-2 reduces both RMSE (0.808890) and MAE (0.638731), while Katz-3 does not improve over the baseline in this group (Table 4). For users with rating < 5, Katz-2 remains the top-performing method and is the only approach that improves both RMSE and MAE relative to Binary, attaining the lowest RMSE (0.859762) and MAE (0.685933), whereas cosine and Jaccard increase MAE despite near-baseline RMSE (Table 5). For CiaoDVD users with rating < 10, Katz-2 achieves the lowest RMSE (0.946817) and MAE (0.726236), with Katz-3 and Jaccard nearly indistinguishable (Table 6). For CiaoDVD users with rating < 5, the best performance shifts slightly toward a local-overlap signal: Jaccard attains the lowest RMSE (0.949587) and MAE (0.732082), while Katz-2 and Katz-3 remain competitive but marginally weaker (Table 7).

Table 4. FilmTrust (cold-start: rating < 10; baseline = Binary).

Algorithm	Mean RMSE	Mean MAE	ΔRMSE (%)	ΔMAE (%)
Binary (baseline)	0.811614	0.640427	0.000	0.000
Cosine similarity	0.810019	0.639455	-0.197	-0.152
Jaccard similarity	0.810302	0.640025	-0.162	-0.063

Table 4. FilmTrust (cold-start: rating < 10; baseline = Binary) (continued).

Algorithm	Mean RMSE	Mean MAE	ΔRMSE (%)	ΔMAE (%)
Katz-2	**0.808890**	**0.638731**	**-0.336**	**-0.265**
Katz-3	0.811916	0.640806	0.037	0.059
Katz-2 (boost)	0.810366	0.639982	-0.154	-0.069
Katz-3 (boost)	0.809682	0.639683	-0.238	-0.116

Table 5. FilmTrust (cold-start: rating < 5; baseline = Binary).

Algorithm	Mean RMSE	Mean MAE	ΔRMSE (%)	ΔMAE (%)
Binary (baseline)	0.862363	0.686805	0.000	0.000
Cosine similarity	0.863087	0.688308	0.084	0.219
Jaccard similarity	0.862405	0.688921	0.005	0.308
Katz-2	**0.859762**	**0.685933**	**-0.302**	**-0.127**
Katz-3	0.863324	0.688750	0.111	0.283
Katz-2 (boost)	0.860891	0.687029	-0.171	0.033
Katz-3 (boost)	0.861222	0.687475	-0.132	0.097

Table 6. CiaoDVD (cold-start: rating < 10; baseline = Binary).

Algorithm	Mean RMSE	Mean MAE	ΔRMSE (%)	ΔMAE (%)
Binary (baseline)	0.948619	0.727552	0.000	0.000
Cosine similarity	0.947074	0.726429	-0.163	-0.154
Jaccard similarity	0.946911	0.726289	-0.180	-0.173
Katz-2	**0.946817**	**0.726236**	**-0.190**	**-0.181**
Katz-3	0.946836	0.726267	-0.188	-0.177
Katz-2 (boost)	0.947667	0.726826	-0.101	-0.100
Katz-3 (boost)	0.947514	0.726695	-0.117	-0.118

Table 7. CiaoDVD (cold-start: rating < 5; baseline = Binary).

Algorithm	Mean	Mean	ΔRMSE	ΔMAE
Binary (baseline)	0.950960	0.732999	0.000	0.000
Cosine similarity	0.949700	0.732169	-0.132	-0.113
Jaccard	**0.949587**	**0.732082**	**-0.144**	**-0.125**
Katz-2	0.949607	0.732139	-0.142	-0.117
Katz-3	0.949636	0.732200	-0.139	-0.109
Katz-2 (boost)	0.950282	0.732569	-0.071	-0.059
Katz-3 (boost)	0.950163	0.732545	-0.084	-0.062

4.3 Discussion of Findings

The proposed Katz-based trust enrichment method outperforms the binary baseline across all evaluated conditions, with the clearest gains appearing in cold-start segments and in the sparser CiaoDVD dataset. When compared against the local-overlap benchmarks (cosine and Jaccard), the proposed method demonstrates competitive or superior performance in most settings. Specifically, Katz-2 emerges as the most reliable variant of the proposed method, achieving the best RMSE on FilmTrust (all users), the best results on FilmTrust cold-start groups (rating < 10 and rating < 5), and the best RMSE on CiaoDVD cold-start users with rating < 10. On CiaoDVD all-users, Katz-3 achieves the strongest overall performance, followed closely by Katz-2 and the Jaccard benchmark. The only condition in which a benchmark slightly outperforms the proposed method is the most extreme CiaoDVD cold-start subgroup (rating < 5), where Jaccard achieves marginally better RMSE and MAE than the best Katz variant — though the margin is small and not statistically significant after Holm correction.

Going from Katz-2 to Katz-3 does not uniformly improve results, despite increasing indirect connectivity: Katz-3 helps in some all-user settings (notably CiaoDVD) but does not consistently outperform Katz-2 under cold-start, indicating that deeper propagation can introduce over-smoothing or dilute useful local trust signals. The boosted Katz variants provide consistent improvements over the binary baseline but do not surpass the standard Katz variants in most conditions, suggesting that the additional post-processing step of re-emphasizing direct trust does not yield further gains once multi-hop structure is already incorporated. An explainability analysis using Spearman correlations between social-network attributes and rating error clarifies why: on CiaoDVD, the proposed Katz variants benefit from creating many soft propagated paths among already-covered users rather than from expanding user coverage; on FilmTrust, moderate graph expansion by the proposed Katz-2 improves signal quality without over-smoothing, while deeper propagation becomes brittle in cold-start regimes. Taken together, the results support the conclusion that multi-hop trust propagation provides a more robust enrichment strategy than direct neighborhood overlap across most conditions, while also confirming that the optimal propagation depth and variant choice are dataset- and segment-dependent.

4.4 Correlation Analysis of Social-Network Attributes

To explain why the proposed Katz-based method wins under most settings while specific benchmark variants become more competitive in extreme cold-start regimes, a correlation analysis relates rating error to the structural characteristics of each trust representation. Spearman correlations are computed between rating error (RMSE, MAE across all-user, Cold<10, and Cold<5 groups) and nine social-network attributes derived from each trust representation. A negative correlation indicates that higher feature values are associated with lower error (i.e., the attribute is helpful for prediction), whereas a positive correlation indicates that higher feature values are associated with higher error.

The nine attributes are defined as follows. n_users is the number of users that appear in the trust graph (i.e., have at least one outgoing or incoming trust edge). n_edges is the total number of directed edges in the (possibly weighted) trust graph; for continuous representations such as cosine, Jaccard, and Katz, this counts the number of non-zero entries after the transformation, so higher values indicate denser trust connectivity. density is the ratio of observed edges to the maximum possible number of edges in the trust graph (i.e., n_edges / [n_users × (n_users − 1)]), where higher density indicates a more strongly connected trust network at the structural level. avg_in_degree is the average number of incoming trust edges per user, reflecting how many other users on average place trust in a given user, so higher values indicate the existence of more "trusted" users (i.e., trustees) on average. avg_out_degree is the average number of outgoing trust edges per user, reflecting how actively users on average extend trust to others, so higher values indicate broader trust-giving behavior across the population. w_mean is the mean weight of non-zero edges in the trust matrix; for the binary baseline this is always 1.0, whereas for enriched representations lower values indicate that trust strength is distributed across many soft (sub-unit) signals rather than concentrated in unit-weight edges, capturing how "soft" or "graded" the trust weighting is. w_median is the median weight of non-zero edges, providing a robust complement to w_mean; together they characterize the distribution shape of trust weights — for example, a low w_median with a higher w_mean suggests a heavy-tailed distribution dominated by a few strong edges. pct_unit_weight is the percentage of trust edges with weight equal to 1.0, where higher values indicate that the representation retains many binary-like (full-strength) edges while lower values indicate that propagation or normalization has softened most edge weights, distinguishing representations that preserve direct trust primacy (e.g., boosted Katz variants) from those that fully replace it with graded propagated signals (e.g., standard Katz). Finally, trust_rating_coverage is the percentage of rating-active users that also appear in the trust graph, where higher values indicate that a larger share of users whose ratings the model must predict can benefit from trust-based regularization; this attribute is particularly important for understanding whether enrichment changes which users receive trust signal versus how trust signal is distributed among already-covered users.

The analysis indicates that different attributes carry predictive potential in different user segments. For the all-user and moderate cold-start (Cold<10) segments, the weight-softness attributes — w_mean, w_median, and pct_unit_weight — are the most influential: they show strong positive correlations with RMSE and MAE, meaning that more binary-like weighting is associated with worse error, and that softening trust into many graded sub-unit signals has the greatest potential to improve performance. This is the regime in which the proposed Katz variants help most, because they replace hard unit-weight edges with soft propagated paths. Connectivity attributes such as n_edges and avg_in_degree also carry positive potential in these segments — for FilmTrust all-users, higher n_edges and avg_in_degree are associated with lower MAE — indicating that expanding effective trust connectivity improves prediction up to a point. Notably, trust_rating_coverage is not the driver of these gains: it remains unchanged between the binary baseline and the Katz variants (13.81% on CiaoDVD, 49.07% on FilmTrust), so improvement comes from how trust signal is distributed among already-covered users rather than from covering more users.

For the extreme cold-start segment (Cold<5), the picture shifts and the weight-softness attributes lose their dominance, producing a more mixed correlation pattern that is less well explained by broad propagation alone. In this regime, density rather than weight softness becomes the salient attribute, but its potential is bounded: for FilmTrust Cold<5 RMSE, density is positively associated with error, signalling that excessive densification becomes brittle when user evidence is extremely scarce. This explains why a highly dense but narrowly scoped local-overlap signal (the Jaccard benchmark, density 0.3733 over only 4.96% trust_rating_coverage) can match or marginally outperform the proposed method for the very coldest users, while still leaving the proposed Katz method best overall. The practical reading is that the features with the greatest potential to raise performance are segment-dependent: graded, soft trust weighting and moderate connectivity expansion benefit the all-user and Cold<10 groups most, whereas in the Cold<5 group strong local overlap within a small coherent subset of users becomes comparatively more valuable than deeper propagation, and aggressive densification should be avoided.

5 Conclusion

This paper proposed a Katz-based trust enrichment method that transforms binary directed trust into continuous multi-hop trust signals through truncated path propagation, and integrates these signals into TrustSVD. When compared against the binary TrustSVD baseline and against two local-overlap benchmarks (cosine and Jaccard similarity), the proposed method demonstrates competitive or superior accuracy in most conditions, with the clearest gains appearing in cold-start regimes. The improvements are generally modest, and the best variant within the proposed method family varies by dataset and by cold-start severity, indicating that multi-hop trust enrichment is beneficial but that variant selection should be guided by the target dataset and user segment. The correlation analysis further clarifies where the gains come from: graded, soft trust weighting (low w_mean and pct_unit_weight) and moderate connectivity expansion are the attributes most associated with lower error in the all-user and Cold<10 segments, whereas for the most extreme cold-start users (Cold<5) strong local overlap within a small coherent subset becomes comparatively more valuable than deeper propagation — with trust-rating coverage held constant, confirming that the improvement stems from how trust signal is distributed among already-covered users rather than from covering more users.

Because training cost is driven by repeated SGD optimization, the proposed method increases training workload primarily by enlarging the effective trust input optimized during learning and by expanding trustee neighborhoods that contribute to user aggregation. On FilmTrust, the binary trust graph contains 1,853 edges; Katz-2 expands this to 11,362 edges ($\rho \approx 6.1\times$), whereas Katz-3 expands it further to 38,000 edges ($\rho \approx 20.5\times$). On the larger and sparser CiaoDVD network, Katz-3 inflates the 22,484 binary edges to 996,109 edges ($\rho \approx 44.3\times$) — which is precisely why Katz-2 is identified as the balanced default: it captures most of the available multi-hop signal while keeping the trust-edge inflation, and therefore the trust-loop workload, an order of magnitude smaller than Katz-3. The additional cost is incurred entirely during an offline training phase and does not affect serving-time latency, and the inflation factor is a controllable design parameter that can be bounded through thresholding or top-K retention per user.

The proposed Katz-based enrichment is most valuable as a lightweight, drop-in enhancement for trust-equipped, rating-sparse, cold-start-heavy systems that are retrained offline, with Katz-2 serving as the recommended default within that operating regime. Practitioners are advised to adopt Katz-2 as the default choice for robust improvement across both all-user and cold-start segments; to consider Katz-3 only when the target dataset is highly sparse and all-user performance is the primary concern; to validate against the Jaccard benchmark on the target cold-start population, since this local-overlap measure may match the proposed method in extreme cold-start regimes (rating < 5) but does not generalize as well across user segments; and to treat boosted Katz variants as alternatives rather than upgrades. Future work includes adaptive calibration of the Katz damping factor, trust network embeddings as a complementary transformation, sparsity-preserving enrichment strategies, and evaluation under neural recommender architectures such as graph neural network–based trust models.

References

[Bokde, 2015] Bokde, D., Girase, S. and Mukhopadhyay, D., Matrix Factorization Model in Collaborative Filtering Algorithms: A Survey, Procedia Computer Science, Vol.49, 2015, pp.136–146.

[Burke, 2002] Burke, R., Hybrid Recommender Systems: Survey and Experiments, User Modeling and User-Adapted Interaction, Vol.12, No.4, 2002, pp.331–370.

[Dong, 2022] Dong, M. et al., A Survey for Trust-Aware Recommender Systems: A Deep Learning Perspective, Knowledge-Based Systems, Vol.249, 2022.

[Duricic, 2018] Duricic, T., Lacic, E., Kowald, D. and Lex, E., Trust-Based Collaborative Filtering: Tackling the Cold Start Problem Using Regular Equivalence, Proceedings of the ACM Conference on Recommender Systems, 2018, pp.446–450.

[Fazeli, 2014] Fazeli, S. et al., Trust-Aware Recommender Systems: Incorporating Implicit Trust in Matrix Factorization, RecSys Workshop, 2014.

[Funk, 2006] Funk, S., Netflix Update: Try This at Home, Available at http://sifter.org/~simon/journal/20061211.html, 2006.

[Golbeck, 2005] Golbeck, J., Computing and Applying Trust in Web-Based Social Networks, PhD Dissertation, University of Maryland, 2005.

[Guo, 2013] Guo, G., Zhang, J. and Thalmann, D., A Simple but Effective Method to Incorporate Trusted Neighbors in Recommender Systems, Proceedings of UMAP, 2013.

[Guo, 2014] Guo, G., Zhang, J., Thalmann, D. and Yorke-Smith, N., ETAF: An Extended Trust Antecedents Framework for Trust Prediction, Proceedings of ASONAM, 2014.

[Guo, 2015] Guo, G., Zhang, J. and Yorke-Smith, N., TrustSVD: Collaborative Filtering with Both the Explicit and Implicit Influence of User Trust and Item Ratings, Proceedings of AAAI Conference on Artificial Intelligence, 2015.

[Han, 2022] Han, X. et al., Recommendation Model Based on Probabilistic Matrix Factorization Integrating User Trust Relationship, Interest Mining, and Item Correlation, IEEE Access, 2022.

[He, 2017] He, X., Liao, L., Zhang, H., Nie, L., Hu, X. and Chua, T.S., Neural Collaborative Filtering, Proceedings of the International World Wide Web Conference, 2017, pp.173–182.

[Hu, 2008] Hu, Y., Koren, Y. and Volinsky, C., Collaborative Filtering for Implicit Feedback Datasets, IEEE International Conference on Data Mining (ICDM), 2008, pp.263–272.

[Jamali, 2009] Jamali, M. and Ester, M., TrustWalker: A Random Walk Model for Combining Trust-Based and Item-Based Recommendation, Proceedings of ACM KDD, 2009, pp.397–406.

[Jamali, 2010] Jamali, M. and Ester, M., A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks, Proceedings of ACM RecSys, 2010, pp.135–142.

[Katz, 1953] Katz, L., A New Status Index Derived from Sociometric Analysis, Psychometrika, Vol.18, No.1, 1953, pp.39–43.

[Koren, 2008] Koren, Y., Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model, Proceedings of ACM SIGKDD, 2008, pp.426–434.

[Koren, 2009] Koren, Y., Bell, R. and Volinsky, C., Matrix Factorization Techniques for Recommender Systems, IEEE Computer, Vol.42, No.8, 2009, pp.30–37.

[Lam, 2008] Lam, X.N., Vu, T., Le, T.D. and Duong, A.D., Addressing Cold-Start Problem in Recommendation Systems, International Conference on Ubiquitous Information Management and Communication, 2008, pp.208–211.

[Ma, 2008] Ma, H., Yang, H., Lyu, M.R. and King, I., SoRec: Social Recommendation Using Probabilistic Matrix Factorization, Proceedings of ACM CIKM, 2008, pp.931–940.

[Ma, 2011] Ma, H., Zhou, D., Liu, C., Lyu, M.R. and King, I., Recommender Systems with Social Regularization, Proceedings of ACM WSDM, 2011, pp.287–296.

[Massa, 2007] Massa, P. and Avesani, P., Trust-Aware Recommender Systems, Proceedings of ACM RecSys, 2007, pp.17–24.

[McPherson, 2001] McPherson, M., Smith-Lovin, L. and Cook, J.M., Birds of a Feather: Homophily in Social Networks, Annual Review of Sociology, Vol.27, 2001, pp.415–444.

[Resnick, 1994] Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P. and Riedl, J., GroupLens: An Open Architecture for Collaborative Filtering of Netnews, Proceedings of ACM CSCW, 1994, pp.175–186.

[Richters, 2011] Richters, O. and Peixoto, T.P., Trust Transitivity in Social Networks, PLoS ONE, Vol.6, No.4, 2011.

[Salakhutdinov, 2007] Salakhutdinov, R. and Mnih, A., Probabilistic Matrix Factorization, Advances in Neural Information Processing Systems (NIPS), 2007.

[Sun, 2022] Sun, R., Yan, J. and Ren, F., Improving the Recommendation Accuracy of TrustSVD via Trustworthy Analysis in the Social Network Environment, Journal of Information Science, 2022.

[Wu, 2016] Wu, L. et al., Understanding Graph-Based Trust Evaluation in Online Social Networks, ACM Computing Surveys, Vol.49, No.1, 2016.

[Xu, 2021] Xu et al., Recommendation Algorithm of Probabilistic Matrix Factorization Based on Directed Trust, Computers & Electrical Engineering, Vol.93, 2021.

[Yang, 2013] Yang, B., Lei, Y., Liu, D. and Liu, J., Social Collaborative Filtering by Trust, Proceedings of IJCAI, 2013.