

Manifold-Aligned Rectification Flow for Denoised Social Recommendation

Qing Meng¹, Zhufu Song¹, Huiyu Min^{2*}, Qian Huang¹ and Pengcheng Zhang¹

¹College of Computer Science and Software Engineering, Hohai University

²College of Information Science and Technology, Nanjing Forestry University

{qmeng, songzhufu}@hhu.edu.cn, hymin@njfu.edu.cn, {huangqian, pchzhang}@hhu.edu.cn

Abstract

Social recommendation leverages social relations to enhance user preference modeling. However, real-world social networks are often noisy and unreliable, where misleading social relationships introduce anisotropic perturbations into user representations. Existing denoising approaches, including heuristic filtering and generative reconstruction, struggle to produce social embeddings that are well aligned with user preferences, limiting their effectiveness in downstream recommendation tasks. To address this challenge, we propose Manifold-Aligned Rectification Flow (MARF), a flow matching based framework that explicitly rectifies noisy social representations toward preference-aligned embeddings. MARF jointly integrates social relations and user-item interactions to construct a preference-aware manifold as the target distribution, guiding the learning of a continuous vector field that captures preference-oriented transformations in the social domain. Through this learned transport process, MARF effectively bridges the gap between the social and preference domains, yielding robust and discriminative user representations. Extensive experiments demonstrate that our proposed model consistently outperforms state-of-the-art social recommendation methods, particularly under sparse and noisy conditions, validating its effectiveness and robustness.

1 Introduction

Recommender systems commonly suffer from data sparsity, as many users exhibit limited historical interactions with items, making it difficult to infer reliable preferences [Gunathilaka *et al.*, 2025]. Social recommendation addresses this challenge by incorporating social relationships as an auxiliary source of information, effectively performing data augmentation on sparse interaction data [Tang *et al.*, 2013; Sharma *et al.*, 2024]. The fundamental rationale behind social recommendation is rooted in theories of social homophily and social influence, which posit that socially connected users

tend to share similar interests and that user preferences can be shaped through social interactions, thereby enabling preference propagation across social networks.

Early social recommendation studies mainly extend traditional matrix factorization-based recommendation models by explicitly incorporating social information, such as co-factorization methods, ensemble methods, and regularization methods [Tang *et al.*, 2013]. With the emergence of deep learning, particularly graph neural networks (GNNs), social recommendation has gradually evolved toward graph-based modeling paradigms [Fan *et al.*, 2019; Sharma *et al.*, 2024]. By modeling interactions and social relationships as graphs, GNN-based methods enable multi-layer message passing and iterative information aggregation, allowing user representations to encode high-order neighborhood structures [Liu *et al.*, 2020; Xiong *et al.*, 2024], complex patterns of social influence [Wu *et al.*, 2019b; Wu *et al.*, 2019a; Ma *et al.*, 2024a], and rich cross-entity interactions [Fan *et al.*, 2019; Yu *et al.*, 2021; Wang *et al.*, 2025], leading to substantial performance improvements over traditional approaches.

However, social information and interactions often reside in different semantic spaces, and social influence does not always translate directly into preference similarity [Tang *et al.*, 2009; Jiang *et al.*, 2024]. Recent empirical studies [Sun *et al.*, 2024] also show that many socially connected users do not share common interests, indicating that social homophily and influence can be inconsistent in real world networks, which limits the benefit of naively aggregating social signals for recommendation. Consequently, directly using the observed social graph can introduce redundant or misleading information into user representations [Meng *et al.*, 2025].

Motivated by this observation, recent work has explored denoising techniques to improve social recommendation [Wang *et al.*, 2023]. Some exploratory studies have employed handcrafted features [Quan *et al.*, 2023], embedding-based approaches [Jiang *et al.*, 2024; Ma *et al.*, 2024b] to compute the similarity between social users using single or multiple networks [Chen *et al.*, 2025], and treat edges between users with low similarity as noise. With the success of denoising diffusion probabilistic models (DDPM) in image processing [Ho *et al.*, 2020], some studies have begun to explore the use of DDPM, leveraging the denoising capability of the reverse process to refine social network matrices or social user representations [Li *et al.*, 2024; Li and Wang, 2024;

*Corresponding Author.

Liu *et al.*, 2025b]. Nevertheless, noise in social networks often deviates significantly from the Gaussian assumption, making the accurate handling of non-Gaussian noise a key limitation that prevents the model from achieving an optimal solution. Therefore, RecFlow [Huang *et al.*, 2025] introduces flow matching to learn continuous velocity fields on the social graph, effectively mitigating non-Gaussian noise and enhancing the quality of user representations. Compared with DDPM, flow matching constructs continuous and deterministic transport paths with lower sampling variance and significantly fewer sampling steps, thereby attracting increasing attention in the recommendation domain [Liu *et al.*, 2025c; Liu *et al.*, 2025a; Huang *et al.*, 2025]. However, existing social recommendation models, regardless of the DDPM or flow-matching paradigms, denoise social embeddings mainly within the social domain. As the denoising process is not explicitly guided by user preference distributions, the resulting social user embeddings may be misaligned with preference space, limiting their utility for recommendation.

To this end, we propose a novel Manifold-Aligned Rectification Flow (MARF) framework for denoised social recommendation. Specifically, we introduce flow matching and explicitly construct a target distribution that captures user preference semantics from the interaction domain. Instead of denoising social representations within the social domain, we learn a continuous velocity field that transports representations toward a target distribution of social-aware preference. During inference, user social embeddings are directly mapped along the learned flow to the user preference domain, enabling effective alignment between social information and user interests. In summary, our contributions are threefold:

- We propose a flow matching based social recommendation model, named Manifold Aligned Rectification Flow (MARF), which utilizes user preference information from the interaction domain to guide social network denoising, explicitly aligning social user representations with user preferences.
- To bridge the semantic gap between social relations and user preferences, a Cross Domain Manifold Alignment (CDMA) module is deliberately designed. Within CDMA, a target distribution for flow matching is constructed from interaction graphs using the Social-aware Preference Manifold (SPM) construction, and social representations are progressively aligned with true user preferences through a flow-based transformation.
- Extensive experiments on three real-world datasets show that the proposed model significantly outperforms state-of-the-art baselines in both recommendation accuracy and robustness.

2 Related Work

2.1 Graph-Based Social Recommendation

Graph neural networks (GNNs) have become the dominant paradigm for social recommendation in recent years due to their ability to model high-order connectivity and structural influence in both social networks and interaction graphs [Wu *et al.*, 2019b; Yang *et al.*, 2023]. Early work such

as GraphRec [Fan *et al.*, 2019] incorporated heterogeneous graph structures to jointly model user-item interactions and social relations, showing significant improvement over classical matrix factorization approaches. DiffNet [Wu *et al.*, 2019a] and its extension DiffNet++ [Wu *et al.*, 2020] further enhance this perspective by recursively propagating social influence and interest representations through multiple graph layers. Yu *et al.* [Yu *et al.*, 2021] introduce hyper-graph neural network to capture high-order user relations beyond pairwise social links. Zhang *et al.* [Zhang *et al.*, 2023] propose to jointly leverage all relationships and construct a heterogeneous global graph learning framework to improve the performance of social recommendation. Later, Xiong *et al.* [Xiong *et al.*, 2024] propose a heterogeneous sampling and propagation mechanism to better capture high-order social neighborhood information while alleviating over-smoothing issues. Nevertheless, most GNN-based social recommendation methods implicitly presume that social neighbors provide trustworthy preference signals, and social information is typically fused through discrete neighborhood aggregation, making them vulnerable to noisy social relations.

2.2 Denoising Social Recommendation

Recent research has recognized that social recommendation can be negatively affected by noisy or weak social relations. To this end, several studies have explored denoising techniques to refine social signals and improve the reliability of user representations. Quan *et al.* [Quan *et al.*, 2023] propose a preference-guided graph that calculates the confidence of social relations and treats edges with low similarity as noisy. Later, a few studies [Jiang *et al.*, 2024; Ma *et al.*, 2024b] leverage user embeddings to calculate cosine similarities, and reconstruct the social network to retain informative relations. Yang *et al.* [Yang *et al.*, 2024] refine the observed social graph based on an information bottleneck principle, demonstrating that learned graph structures can enhance recommendation performance when noise is explicitly considered. Meanwhile, Ma *et al.* [Ma *et al.*, 2024a] use the contrastive learning to improve robustness under noisy social and interaction data. Chen *et al.* [Chen *et al.*, 2025] utilize preference signals from both social and interaction networks and iteratively remove noise from both domains. More recent approaches have shifted the denoising paradigm from explicit relation filtering to representation-level purification. For example, Yang *et al.* [Yang *et al.*, 2025] simulate multiple noisy social environments and propose a graph invariant learning to obtain robust user social embeddings. Building upon this perspective, an increasing number of studies further introduce generative models to perform denoising implicitly in the embedding space. Studies such as RecDiff [Li *et al.*, 2024] and ARD-SR [Sun *et al.*, 2025] introduce diffusion models to obtain denoised user embeddings, and Liu *et al.* [Liu *et al.*, 2025c] replace diffusion processes with flow matching, enabling direct learning of a continuous transport mapping for embedding rectification.

3 Preliminaries

Problem Formulation. Let $\mathcal{U} = \{u_1, \dots, u_{|\mathcal{U}|}\}$ and $\mathcal{V} = \{v_1, \dots, v_{|\mathcal{V}|}\}$ denote the sets of users and items. The user-




Figure 1: The overall architecture of MARF. The framework rectifies noisy social signals via: (1) Dual-View Manifold Encoding and Social-Aware Target Construction; (2) Cross Domain Manifold Alignment (CDMA) using flow matching to bridge the social-interest semantic gap; and (3) Adaptive Fusion to integrate rectified representations for recommendation.

item interaction graph is defined as $\mathcal{G}_I = (\mathcal{U} \cup \mathcal{V}, \mathcal{E}_I)$, where $\mathcal{E}_I$ represents user item interactions. The user social graph is defined as $\mathcal{G}_S = (\mathcal{U}, \mathcal{E}_S)$, and $\mathcal{E}_S$ is social relations between users. The objective is to learn a mapping function Φ that rectifies noisy social representations into reliable user embeddings, enabling accurate Top- K recommendation.

Flow Matching Paradigm. Flow Matching (FM) aims to learn a time-dependent vector field $\mathbf{v}_t(\mathbf{x})$ that generates a probability path $p_t(\mathbf{x})$ between distributions p_0 and p_1 . For any trajectory starting from an initial sample $\mathbf{x}_0 \sim p_0$, this transport process is governed by the Ordinary Differential Equation (ODE):

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_t(\mathbf{x}), \quad t \in [0, 1], \quad \text{s.t.} \quad \mathbf{x}(0) = \mathbf{x}_0. \quad (1)$$

Conditional Flow Matching (CFM) simplifies the learning objective by defining a conditional Optimal Transport (OT) path $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$ between samples $\mathbf{x}_0 \sim p_0$ and $\mathbf{x}_1 \sim p_1$. The vector field is then trained to regress the constant velocity $\mathbf{u}_t(\mathbf{x}_t|\mathbf{x}_0, \mathbf{x}_1) = \mathbf{x}_1 - \mathbf{x}_0$, guiding the transformation from the source manifold to the target manifold.

4 Methodology

As illustrated in Figure 1, we propose a generative framework **Manifold-Aligned Rectification Flow (MARF)** to bridge the semantic gap between social information and user preferences. Specifically, the workflow begins with a dual-view manifold encoding module, which maps graph structures into latent spaces to obtain basic representations of users and items. Next, a Social-aware Preference Manifolds (SPM) is constructed from social relations and user-item interactions to provide a reliable target distribution for representation rectification in flow matching. Then, a cross-domain manifold

alignment (CDMA) module is then employed to learn a continuous transport mapping from the noisy social manifold to the preference manifold. Finally, the rectified social representations are fused with interaction-based user embeddings to form final user representations for recommendation.

4.1 Dual-View Manifold Encoding

To capture high-order connectivity in heterogeneous graphs, a dual-branch LightGCN [He *et al.*, 2020] is adopted as the backbone encoder, where independent message passing is performed on the interaction and social graphs.

In user-item interaction Graph $\mathcal{G}_I$, collaborative signals are modeled via simplified graph convolution with the layer-wise propagation rule defined as:

$$\mathbf{E}^{I,(l+1)} = (\mathbf{D}_I^{-\frac{1}{2}} \mathbf{A}_I \mathbf{D}_I^{-\frac{1}{2}}) \mathbf{E}^{I,(l)}, \quad (2)$$

where $\mathbf{A}_I$ is the adjacency matrix of $\mathcal{G}_I$ and $\mathbf{D}_I$ is the corresponding degree matrix. After propagating through L layers, we aggregate the embeddings from all layers via summation to form the final multi-scale representations:

$$\mathbf{E}^I = \sum_{l=0}^L \mathbf{E}^{I,(l)}. \quad (3)$$

where $\mathbf{E}^I$ is then partitioned into the user interest representation $\mathbf{E}_U^I \in \mathbb{R}^{N \times d}$ and the item representation $\mathbf{E}_V^I \in \mathbb{R}^{M \times d}$.

Similarly, in social graph $\mathcal{G}_S$, the user representation can be obtained by summing layer-wise social embeddings:

$$\mathbf{E}_U^S = \sum_{l=0}^L \mathbf{E}_U^{S,(l)}. \quad (4)$$

As user representations $\mathbf{E}_U^S$ are learned from the noisy social network, a subsequent module is needed to denoise and rectify the resulting social representations for downstream recommendation.

4.2 Cross Domain Manifold Alignment (CDMA)

Social connections often stem from offline acquaintances or social status, which do not necessarily translate to shared item preferences [Wang *et al.*, 2023; Li *et al.*, 2024]. To mitigate noise in social information, we introduce flow matching and design a cross domain manifold alignment (CDMA) module to learn a continuous transport mapping that align social representation with user preference representations.

Social-Aware Preference Manifolds (SPM) Construction.

User preference representations are derived via LightGCN from the user-item interaction graph in the previous section. However, directly adopting them as the target manifold for flow matching would cause social representations to collapse into preference-only embeddings, losing informative social signals. To further leverage social information, we refine the social adjacency matrix $\mathbf{A}_S$ by reweighting each entry based on the overlap of item interactions between users. Entries with shared interactions are assigned the Jaccard similarity of their item sets, while entries with no overlap are set to a small constant ϵ , and all other entries remain unchanged. From $\mathbf{A}_S$, the top- k neighbors for each user are selected to construct a refined matrix $\mathbf{A}_S^*$, which is then combined with the preference embeddings $\mathbf{E}_U^I$ to form the social-aware preference manifold $\mathbf{E}_{target} = \mathbf{A}_S^* \mathbf{E}_U^I$. Specifically, $\mathbf{A}_{uv}^* = 0$ indicates cases where two users have neither a social link nor any common interacted items; thus, no weight is assigned during manifold construction.

Learning the Rectification Field. Following the principle of flow matching, we aim to learn a continuous vector field whose induced flow transports samples between a simple reference distribution and the target distribution. In our setting, the target distribution is defined as the social-aware preference manifold $\mathbf{E}_{target}$. Accordingly, we formulate the rectification process as learning a flow from $\mathbf{E}_{target}$ at $t = 0$ to Gaussian noise $\mathbf{X}_n$ at $t = 1$. The forward process is constructed via the optimal transport interpolation:

$$\mathbf{Z}_t = (1 - t)\mathbf{E}_{target} + t\mathbf{X}_n, \quad (5)$$

which defines a time-dependent path connecting the target distribution and the reference noise distribution, where the time variable t is discretized according to the total sampling steps n , with a constant step size $\Delta t = 1/n$.

To avoid deterministic mappings and better model uncertainty in user interests, interaction embeddings ($\tilde{\mathbf{E}}_U^I$) are perturbed with noise, which is defined as:

$$\tilde{\mathbf{E}}_U^I = \mathbf{E}_U^I \odot (1 + \gamma\xi), \quad \xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (6)$$

where $\odot$ denotes element-wise product, and γ is a noise scale factor controlling the intensity of the perturbation ξ . Then, $\tilde{\mathbf{E}}_U^I$ is concatenated with the noisy flow state $\mathbf{Z}_t$ and a time embedding $\mathbf{e}_t$ to form a feature tensor $[\mathbf{Z}_t, \mathbf{e}_t, \tilde{\mathbf{E}}_U^I]$, where $\mathbf{e}_t$ is obtained by encoding the scalar time t into a sinusoidal embedding. Recognizing that the relative importance of these components evolves across the transport trajectory, attention weights are computed to achieve dynamic weight allocation. The resulting fused representation $\mathbf{h}_{fused}$ integrates information from all three sources, adaptively modulating the contri-

butions of topological priors and behavioral signals at different stages of the rectification process. This mechanism effectively resolves the semantic misalignment between the social and interest manifolds, facilitating a more precise transition to the target space. Then, the fused representation $\mathbf{h}_{fused}$ is fed into a flow model f_θ to predict the target, namely social-aware user preference manifold. Following prior work [Liu *et al.*, 2025c], the model is trained to directly regress $\mathbf{E}_{target}$ instead of the instantaneous velocity field. This approach alleviates the challenge of fitting noisy vector fields, improving numerical stability and prediction accuracy, particularly on sparse data. Finally, the training loss is defined as the Mean Squared Error (MSE) between the predicted target and the ground-truth social-aware preference manifold:

$$\mathcal{L}_{FM} = \mathbb{E}_{t, \mathbf{X}_n} [\|\hat{\mathbf{E}}_{target} - \mathbf{E}_{target}\|^2]. \quad (7)$$

where $\hat{\mathbf{E}}_{target} = f_\theta(\mathbf{h}_{fused})$, and f_θ serves as the flow model capable of modeling the transformation toward the social-aware user preference manifold, and is implemented by a two layer MLP.

Cross Domain Rectification. Starting from the social representation ($\mathbf{E}_U^S$) as the initial state, we set $\mathbf{Z}_1 = \mathbf{E}_U^S$ and apply the flow model f_θ by integrating backward in time from $t = 1$ to $t = 0$, progressively transforming social representations toward the social-aware user preference manifold. However, Rahaman *et al.* [Rahaman *et al.*, 2019] show that standard MLPs are biased toward low-frequency signals, producing over-smoothed vector fields that are insufficient for modeling large semantic discrepancies. Inspired by Fourier feature theory [Tancik *et al.*, 2020], we scale the input of the flow model to enable the network to capture high-frequency signals.

When predicting the target $\hat{\mathbf{E}}_{target}$, the flow model takes inputs $\mathbf{Z}_t$, $\tilde{\mathbf{E}}_U^I$, and $\mathbf{e}_t$. Here, $\mathbf{e}_t = \text{TimeEmb}(\lambda \cdot t)$ is a standard sinusoidal encoding for time awareness. To capture high-frequency signals, only the scalar time t is scaled by λ while other semantic inputs remain fixed:

$$\hat{\mathbf{E}}_{target} = f_\theta(\mathbf{Z}_t, \mathbf{e}_t, \tilde{\mathbf{E}}_U^I). \quad (8)$$

Then, the rectified velocity is computed as:

$$\hat{v}_\theta(\mathbf{Z}_t, t) = \hat{\mathbf{E}}_{target} - \mathbf{Z}_t \quad (9)$$

The discrete update is performed using a Euler solver over n steps with $\Delta t = 1/n$:

$$\mathbf{Z}_{t-\Delta t} = \mathbf{Z}_t - \hat{v}_\theta(\mathbf{Z}_t, t) \cdot \Delta t. \quad (10)$$

At $t = 0$, the final rectified embedding $\mathbf{Z}_0$ is obtained as the social-aware user embedding $\hat{\mathbf{E}}_U^S$.

4.3 Adaptive Fusion and Optimization

Adaptive Fusion (AF). To integrate rectified social information with collaborative signals, we design a user-specific adaptive gating mechanism that balances the contribution of the interaction user embedding $\mathbf{E}_U^I$ and the rectified social-aware user embedding $\hat{\mathbf{E}}_U^S$. Formally, we concatenate the two

representations and feed them into a two-layer MLP to compute a scalar gating weight $\alpha_u \in (0, 1)$:

$$\alpha_u = \sigma \left(\mathbf{w}_2^\top \tanh \left(\mathbf{W}_1 [\mathbf{E}_U^I; \hat{\mathbf{E}}_U^S] + \mathbf{b}_1 \right) + \mathbf{b}_2 \right), \quad (11)$$

where $[\cdot; \cdot]$ denotes concatenation, $\sigma(\cdot)$ is the sigmoid function, and $\{\mathbf{W}_1, \mathbf{b}_1, \mathbf{w}_2, \mathbf{b}_2\}$ are learnable parameters. The final user representation $\mathbf{E}_U$ is obtained via a weighted summation: $\mathbf{E}_U = \alpha_u \mathbf{E}_U^I + (1 - \alpha_u) \hat{\mathbf{E}}_U^S$. This allows the model to selectively leverage social priors only when they are informative, reducing the impact of residual noise.

Optimization Objective. We employ the Bayesian Personalized Ranking (BPR) loss to optimize the recommendation task. BPR assumes that a user prefers an observed interaction item i over an unobserved negative item j . The predicted preference score is defined as the inner product between the final user embedding and the item embedding: $\hat{y}_{ui} = \mathbf{E}_u^\top \mathbf{E}_i^I$. The BPR objective minimizes the following pairwise ranking loss:

$$\mathcal{L}_{BPR} = - \sum_{(u,i,j) \in \mathcal{D}_{train}} \ln \sigma(\hat{y}_{ui} - \hat{y}_{uj}), \quad (12)$$

where $\mathcal{D}_{train} = \{(u, i, j) | i \in \mathcal{V}_u, j \notin \mathcal{V}_u\}$ denotes the training triplets. A multi-task loss is employed to jointly optimize recommendation performance and flow matching rectification:

$$\mathcal{L} = \mathcal{L}_{BPR} + w \mathcal{L}_{FM} + \lambda_{reg} \|\Theta\|_2^2, \quad (13)$$

where Θ denotes the set of all trainable parameters, and w and λ_{reg} are hyperparameters controlling the weight of the flow matching loss and 2-norm regularization, respectively.

5 Experiments

5.1 Experimental Settings

Datasets and Evaluation Metrics. We evaluate our model on three widely-used benchmark datasets: Ciao, Epinions, and Yelp, using the pre-processed versions provided by RecDiff [Li *et al.*, 2024] for fair comparison. Ciao and Epinions are consumer review platforms where users establish trust relationships, whereas Yelp is a large-scale local business recommendation platform containing undirected friendship links. The detailed statistics are summarized in Table 1.

Dataset	#Users	#Items	#Interactions	#Social Ties
Ciao	1,925	15,053	23,223	65,084
Epinions	14,680	233,261	447,312	632,144
Yelp	99,262	105,142	672,513	1,298,522

Table 1: Statistics of experimental datasets.

Evaluation Protocols. Following the standard setting, we randomly split each dataset into training, validation, and testing sets with a ratio of 7:1:2. Following common practices in recommendation research, we employ two widely used metrics to assess recommendation performance: Recall@N and NDCG@N. Unless otherwise specified, Recall and NDCG are evaluated at N = 20 throughout the experiments.

Baselines. We compare MARF with a comprehensive set of baselines, covering: (1) Matrix Factorization methods: TrustMF [Yang *et al.*, 2016]; (2) GNN-based Social Recommendation: DiffNet [Wu *et al.*, 2019a], GraphRec [Fan *et al.*, 2019], DGRRec [Song *et al.*, 2019], NGCF [Wang *et al.*, 2019], and KCGN [Huang *et al.*, 2021]; (3) Self-Supervised Learning methods: MHCN [Yu *et al.*, 2021], SMIN [Long *et al.*, 2021], and DSL [Wang *et al.*, 2023]; (4) Generative Models: GDMSR [Quan *et al.*, 2023], RecDiff [Li *et al.*, 2024], and RecFlow [Huang *et al.*, 2025].

Implementation Details. Our proposed model MARF is optimized with the Adam optimizer, and all experiments are conducted on NVIDIA RTX 3090 GPUs. Hyperparameters are tuned based on validation performance. The embedding dimension d is set to 64 for Ciao and 128 for Epinions and Yelp. The batch size is set to 2048 for Ciao and 4096 for the larger datasets. The learning rate is selected from $\{1e^{-2}, 5e^{-3}, 1e^{-3}, 1e^{-4}\}$ and the λ_{reg} in 2-norm regularization in $\{1e^{-3}, 2e^{-3}, 3e^{-3}\}$. The amplification factor $\lambda = 1000$ as the default configuration based on hyperparameter analysis. Moreover, a grid search is conducted over the neighbor size k (5–50, step 5), the number of inference steps n (10–50, step 10), and the multi-task weight w for the flow matching loss (0.01–1.0, step 0.05). For a fair comparison, all baselines are initialized and tuned according to the original papers.

5.2 Overall Performance

To evaluate the effectiveness of the proposed MARF model, we compare it against the baseline methods. The results, summarized in Table 2.

Method	Ciao		Yelp		Epinions	
	Recall	NDCG	Recall	NDCG	Recall	NDCG
TrustMF	0.0539	0.0343	0.0371	0.0193	0.0265	0.0195
SAMN	0.0604	0.0384	0.0403	0.0208	0.0329	0.0226
DiffNet	0.0528	0.0328	0.0557	0.0292	0.0384	0.0273
GraphRec	0.0540	0.0335	0.0419	0.0201	0.0334	0.0246
DGRRec	0.0517	0.0319	0.0410	0.0209	0.0326	0.0236
NGCF	0.0559	0.0363	0.0450	0.0230	0.0353	0.0243
MHCN	0.0621	0.0378	0.0567	0.0292	0.0438	0.0321
KCGN	0.0602	0.0350	0.0460	0.0234	0.0220	0.0146
SMIN	0.0588	0.0354	0.0485	0.0251	0.0333	0.0228
GDMSR	0.0560	0.0355	0.0513	0.0246	0.0368	0.0241
DSL	0.0606	0.0389	0.0504	0.0259	0.0365	0.0267
RecDiff	0.0712	0.0419	0.0597	0.0308	0.0460	0.0336
RecFlow	0.0725	0.0438	0.0618	0.0341	0.0486	0.0341
MARF	0.0742	0.0453	0.0642	0.0335	0.0513	0.0360
Improv.	2.34%	3.42%	3.88%	-	5.56%	5.57%

Table 2: Performance comparison with baselines. The last row indicates the relative improvement over the best baseline.

The results in Table 2 indicate that MARF consistently achieves superior performance across all datasets. Detailed analyses are provided as follows:

Superiority over State-of-the-Art Baselines. MARF outperforms all baselines on the majority of metrics. Compared to the strongest baseline, RecFlow, our model achieves relative improvements in Recall of 2.34%, 3.88%, and 5.56% on Ciao, Yelp, and Epinions, respectively. On the large-scale Yelp dataset, MARF significantly improves Recall (0.0642 vs. 0.0618) while maintaining comparable NDCG to RecFlow. These results suggest that the superior performance of MARF may be due to the explicit alignment of social and preference representations, which allows effective integration of social influence and individual preferences. Consequently, our proposed model can retrieve relevant items and improve both recall and ranking quality.

Robustness against Structural Noise. Traditional GNN-based methods (e.g., DiffNet, KCGN) and SSL-based methods (e.g., MHCN) exhibit limited performance, especially on the noisy Epinions dataset. These methods rely on message passing over raw social structures, making them vulnerable to noise propagation and over-smoothing. In contrast, generative models such as RecDiff, RecFlow, and MARF show better robustness. MARF further improves performance by employing Social aware Preference Manifold (SPM) construction to filter irrelevant relations before the flow matching process, ensuring that model optimization is guided by reliable behavioral signals rather than biased social structures.

Benefit of Cross-Domain Manifold Alignment (CDMA). A key distinction between MARF and recent generative baselines (RecDiff, RecFlow) lies in the learning objective. RecDiff and RecFlow primarily focus on restorative denoising—reconstructing a clean social graph from the noisy input. Consequently, the learned representations are often limited to capturing social structures, making it difficult to jointly encode social information and user preferences into a unified representation space. In contrast, MARF formulates social rectification task as a transport process that map user representations from the social manifold to a preference-aligned manifold. By effectively bridging the semantic gap between social relationships and user preferences, MARF learns a more discriminative transformation that is not constrained by the inherent geometry of the social graph.

5.3 Ablation Study

To evaluate the contribution of each core component in MARF and validate our manifold rectification assumptions, we conduct ablation experiments on the Ciao and Epinions datasets. The full MARF model is compared against three variants, each removing one key module of the proposed model:

- **w/o CDMA:** This variant removes the Cross-Domain Manifold Alignment (CDMA) module and directly uses the social-aware preference manifold for prediction.
- **w/o SPM:** This variant replaces the Social-aware Preference Manifold (SPM) construction with standard Jaccard-based neighbor selection, where the top- k social neighbors with the highest Jaccard similarity are retained, ignoring the original social relations.

- **w/o AF:** This variant substitutes the Adaptive Fusion (AF) module with simple average pooling to integrate rectified social information and collaborative signals.

Method Variant	Ciao		Epinions	
	Recall	NDCG	Recall	NDCG
w/o CDMA	0.0623	0.0392	0.0491	0.0344
w/o AF	0.0672	0.0399	0.0435	0.0299
w/o SPM	0.0642	0.0386	0.0500	0.0355
MARF (Full)	0.0742	0.0453	0.0513	0.0360

Table 3: Ablation study of MARF components.

The results are summarized in Table 3. Several key observations can be drawn as follows:

Effectiveness of CDMA and SPM. Ablation results (Table 3) reveal that while both modules are essential, CDMA serves as the primary denoising component. The w/o CDMA variant exhibits a substantial performance drop (e.g., 16.0% Recall@20 on Ciao), confirming its necessity in bridging the cross-domain semantic gap through flow transport. While SPM (w/o SPM) effectively mitigates social noise by prioritizing behaviorally consistent neighbors, the more significant decline in w/o CDMA underscores that flow-based rectification is the core denoising stage. Essentially, SPM provides a reliable structural prior that guides CDMA’s continuous transport to resolve semantic misalignments that static filtering alone cannot bridge.

Necessity of Adaptive Fusion. Removing the adaptive fusion (AF) module leads to a significant performance drop, particularly on Epinions, where the NDCG decreases by 17.2%. Given that the reliability of social signals varies substantially across users, simple average pooling can introduce feature interference. The gating mechanism in AF module successfully resolves the view conflict by adaptively weighting social information according to the complementarity to the interaction view.

5.4 Mechanism Analysis and Visualization

To empirically validate the effectiveness of MARF, we analyze the semantic alignment between the rectified social representations and the behavioral ground truth. We adopt the social-aware preference manifold $\mathbf{E}_{target}$ as the target manifold. Figure 2 illustrates the Kernel Density Estimation (KDE) of the cosine similarity between user social embeddings and this target, focusing specifically on the top 20% of hard samples with the largest initial semantic gap.

As shown in Figure 2, the Initial social embeddings (gray) are heavily concentrated in the low-similarity region, reflecting the intrinsic semantic misalignment in raw social data. In contrast, the rectified distribution (red) shows a clear shift towards higher similarity values. This confirms that MARF does goes beyond local smoothing within the original graph, performing active cross manifold transport that effectively realigns social information with the interaction-based preference manifold.




Figure 2: Geometric Rectification Analysis.

5.5 Robustness and Sensitivity Analysis

In this section, we analyze the robustness of MARF under varying data conditions and hyperparameter configurations.

Robustness across User Activity Levels. To validate the critical role of the Social-aware Preference Manifold (SPM) in handling varying data sparsity, we analyze the performance gain over the *w/o SPM* variant across three user groups: *Cold-Start* (< 5), *Mid-Activity* (5-15), and *High-Activity* (> 15). As illustrated in Figure 3, SPM functions as a dynamic structural rectifier via a dual-mechanism advantage:




Figure 3: Impact of SPM across Activity Levels on the Ciao dataset.

- **SPM Rescues Cold-Start Users:** For the Cold-Start group, the *w/o SPM* variant struggles due to the lack of co-interacted items. Our proposed SPM construction strategy addresses this by automatically triggering the structural fallback coefficient ϵ , which activates the supplementary set to provide essential topological priors. This enables a steady recovery in performance, confirming that SPM successfully extracts valuable signals from noise social relations when behavioral data is sparse.
- **SPM Denoises High-Activity Users:** In the high-activity group, standard heuristics accumulate severe noise from complex social circles. By strictly following the priority logic, SPM effectively acts as a semantic filter, ensuring that item ranking is guided by reliable signals rather than being distorted by social noise.




Figure 4: Hyperparameter analysis on Ciao dataset.

Hyperparameter Analysis. We evaluate the robustness of MARF regarding hyperparameters k , n , w and λ . The results are illustrated in Figure 4.

Neighbor Size k determines the density of the constructed social-aware preference manifold. Performance initially improves as k increases but degrades beyond a threshold due to over-smoothing and the introduction of social noise.

Inference Steps n follows an inverted-U trend, as enough steps are needed to reduce Euler solver truncation error and accurately traverse complex manifold trajectories.

Multi-task Weight w controls the contribution of the flow matching loss in the joint optimization. Noisy datasets (e.g., Epinions) require conservative values to reduce the impact of intrinsic structural noise, while cleaner datasets benefit from stronger flow matching guidance to align representations.

Amplification Factor λ controls the trade-off between generative smoothness and ranking discriminability. Performance peaks around $\lambda \approx 1000$, where high-frequency signals help decouple representations from irrelevant ties, while larger values cause numerical oscillation and disrupt vector field continuity.

6 Conclusion

In this paper, we address the challenge of data sparsity and noise in social recommendation. While social links provide useful signals, the semantic gap between social and preference domains often introduces misleading information. To tackle this, we propose Manifold Aligned Rectification Flow (MARF), which employs flow matching to transport noisy social representations toward a preference-aware target distribution. The Cross Domain Manifold Alignment (CDMA) module explicitly bridges social relations and user preferences, ensuring that refined embeddings align with users' true preferences. Experiments on three real-world datasets show that MARF outperforms state-of-the-art baselines, achieving higher recommendation accuracy and robust performance under sparse and noisy conditions. This work highlights the value of cross-domain alignment and flow matching for reliable social recommendation.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grant No.62302150. China Postdoctoral Science Foundation No.2024M750731. Jiangsu Province Higher Education Basic Science (Natural Science) Foundation No.24KJB520014.

References

- [Chen *et al.*, 2025] Wenjie Chen, Yi Zhang, Honghao Li, Lei Sang, and Yiwen Zhang. Dual-domain collaborative denoising for social recommendation. *IEEE Transactions on Computational Social Systems*, 2025.
- [Fan *et al.*, 2019] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *The world wide web conference*, pages 417–426, 2019.
- [Gunathilaka *et al.*, 2025] Thennakoon Mudiyanseelage Anupama Udayangani Gunathilaka, Prabhashrini Dhanushika Manage, Jinglan Zhang, Yuefeng Li, and Wayne Kelly. Addressing sparse data challenges in recommendation systems: A systematic review of rating estimation using sparse rating data and profile enrichment techniques. *Intelligent Systems with Applications*, page 200474, 2025.
- [He *et al.*, 2020] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 639–648, 2020.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [Huang *et al.*, 2021] Chao Huang, Huance Xu, Yong Xu, Peng Dai, Lianghao Xia, Mengyin Lu, Liefeng Bo, Hao Xing, Xiaoping Lai, and Yanfa Tang. Knowledge-aware coupled graph neural network for social recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4115–4122, 2021.
- [Huang *et al.*, 2025] Yinxuan Huang, Ke Liang, Zhuofan Dong, Xiaodong Qu, Wang Tianxiang, Yue Han, Jingao Xu, Bin Zhou, and Ye Wang. Flow matching for denoised social recommendation. In *Forty-second International Conference on Machine Learning, ICML2025, Vancouver, BC, Canada, July 13-19, 2025*, 2025.
- [Jiang *et al.*, 2024] Wei Jiang, Xinyi Gao, Guandong Xu, Tong Chen, and Hongzhi Yin. Challenging low homophily in social recommendation. In *Proceedings of the ACM Web Conference 2024*, pages 3476–3484, 2024.
- [Li and Wang, 2024] Jiuqiang Li and Hongjun Wang. Graph diffusive self-supervised learning for social recommendation. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, pages 2442–2446, 2024.
- [Li *et al.*, 2024] Zongwei Li, Lianghao Xia, and Chao Huang. Recdiff: Diffusion model for social recommendation. In *Proceedings of the 33rd ACM international conference on information and knowledge management*, pages 1346–1355, 2024.
- [Liu *et al.*, 2020] Yang Liu, Liang Chen, Xiangnan He, Jiaying Peng, Zibin Zheng, and Jie Tang. Modelling high-order social relations for item recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 34(9):4385–4397, 2020.
- [Liu *et al.*, 2025a] Chengkai Liu, Yangtian Zhang, Jianling Wang, Rex Ying, and James Caverlee. Flow matching for collaborative filtering. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 1765–1775, 2025.
- [Liu *et al.*, 2025b] Chengyi Liu, Jiahao Zhang, Shijie Wang, Wenqi Fan, and Qing Li. Score-based generative diffusion models for social recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 2025.
- [Liu *et al.*, 2025c] Feng Liu, Lixin Zou, Xiangyu Zhao, Min Tang, Liming Dong, Dan Luo, Xiangyang Luo, and Chenliang Li. Flow matching based sequential recommender model. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 3108–3116. International Joint Conferences on Artificial Intelligence Organization, 8 2025.
- [Long *et al.*, 2021] Xiaoling Long, Chao Huang, Yong Xu, Huance Xu, Peng Dai, Lianghao Xia, and Liefeng Bo. Social recommendation with self-supervised metagraph in-formax network. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, pages 1160–1169, 2021.
- [Ma *et al.*, 2024a] Gang-Feng Ma, Xu-Hua Yang, Haixia Long, Yanbo Zhou, and Xin-Li Xu. Robust social recommendation based on contrastive learning and dual-stage graph neural network. *Neurocomputing*, 584:127597, 2024.
- [Ma *et al.*, 2024b] Wenzhe Ma, Yuexian Wang, Yanmin Zhu, Zhaobo Wang, Mengyuan Jing, Xuhao Zhao, Jiadi Yu, and Feilong Tang. Madm: A model-agnostic denoising module for graph-based social recommendation. In *Proceedings of the 17th ACM international conference on web search and data mining*, pages 501–509, 2024.
- [Meng *et al.*, 2025] Qing Meng, Huiyu Min, Ming Shan Hee, Roy Ka-Wei Lee, Bing Tian Dai, and Shuai Xu. Usefulness and diminishing returns: Evaluating social information in recommender systems. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 2105–2114, 2025.
- [Quan *et al.*, 2023] Yuhan Quan, Jingtao Ding, Chen Gao, Lingling Yi, Depeng Jin, and Yong Li. Robust preference-guided denoising for graph based social recommendation. In *Proceedings of the ACM web conference 2023*, pages 1097–1108, 2023.

- [Rahaman *et al.*, 2019] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019.
- [Sharma *et al.*, 2024] Kartik Sharma, Yeon-Chang Lee, Sivagami Nambi, Aditya Salian, Shlok Shah, Sang-Wook Kim, and Srijan Kumar. A survey of graph neural networks for social recommender systems. *ACM Computing Surveys*, 56(10):1–34, 2024.
- [Song *et al.*, 2019] Weiping Song, Zhiping Xiao, Yifan Wang, Laurent Charlin, Ming Zhang, and Jian Tang. Session-based social recommendation via dynamic graph attention networks. In *Proceedings of the Twelfth ACM international conference on web search and data mining*, pages 555–563, 2019.
- [Sun *et al.*, 2024] Youchen Sun, Zhu Sun, Yingpeng Du, Jie Zhang, and Yew Soon Ong. Self-supervised denoising through independent cascade graph augmentation for robust social recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2806–2817, 2024.
- [Sun *et al.*, 2025] Youchen Sun, Zhu Sun, Yingpeng Du, Jie Zhang, and Yew Soon Ong. Model-agnostic social network refinement with diffusion models for robust social recommendation. In *Proceedings of the ACM on Web Conference 2025*, pages 370–378, 2025.
- [Tancik *et al.*, 2020] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Ravi Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, volume 33, pages 7537–7547, 2020.
- [Tang *et al.*, 2009] Jie Tang, Jimeng Sun, Chi Wang, and Zi Yang. Social influence analysis in large-scale networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 807–816, 2009.
- [Tang *et al.*, 2013] Jiliang Tang, Xia Hu, and Huan Liu. Social recommendation: a review. *Social network analysis and mining*, 3(4):1113–1133, 2013.
- [Wang *et al.*, 2019] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. Neural graph collaborative filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 165–174, 2019.
- [Wang *et al.*, 2023] Tianle Wang, Lianghao Xia, and Chao Huang. Denoised self-augmented learning for social recommendation. *arXiv preprint arXiv:2305.12685*, 2023.
- [Wang *et al.*, 2025] Zheng Wang, Wanwan Wang, Yimin Huang, Zhaopeng Peng, Ziqi Yang, Ming Yao, Cheng Wang, and Xiaoliang Fan. P4gcn: Vertical federated social recommendation with privacy-preserving two-party graph convolution network. In *Proceedings of the ACM on Web Conference 2025*, pages 2924–2934, 2025.
- [Wu *et al.*, 2019a] Le Wu, Peijie Sun, Yanjie Fu, Richang Hong, Xiting Wang, and Meng Wang. A neural influence diffusion model for social recommendation. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 235–244, 2019.
- [Wu *et al.*, 2019b] Qitian Wu, Hengrui Zhang, Xiaofeng Gao, Peng He, Paul Weng, Han Gao, and Guihai Chen. Dual graph attention networks for deep latent representation of multifaceted social effects in recommender systems. In *The world wide web conference*, pages 2091–2102, 2019.
- [Wu *et al.*, 2020] Le Wu, Junwei Li, Peijie Sun, Richang Hong, Yong Ge, and Meng Wang. Diffnet++: A neural influence and interest diffusion network for social recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 34(10):4753–4766, 2020.
- [Xiong *et al.*, 2024] Fei Xiong, Haoran Sun, Guixun Luo, Shirui Pan, Meikang Qiu, and Liang Wang. Graph attention network with high-order neighbor information propagation for social recommendation. In *IJCAI-24: Thirty-Third International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence, 2024.
- [Yang *et al.*, 2016] Bo Yang, Yu Lei, Jiming Liu, and Wenjie Li. Social collaborative filtering by trust. *IEEE transactions on pattern analysis and machine intelligence*, 39(8):1633–1647, 2016.
- [Yang *et al.*, 2023] Yonghui Yang, Le Wu, Kun Zhang, Richang Hong, Hailin Zhou, Zhiqiang Zhang, Jun Zhou, and Meng Wang. Hyperbolic graph learning for social recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 36(12):8488–8501, 2023.
- [Yang *et al.*, 2024] Yonghui Yang, Le Wu, Zihan Wang, Zhuangzhuang He, Richang Hong, and Meng Wang. Graph bottlenecked social recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3853–3862, 2024.
- [Yang *et al.*, 2025] Yonghui Yang, Le Wu, Yuxin Liao, Zhuangzhuang He, Pengyang Shao, Richang Hong, and Meng Wang. Invariance matters: Empowering social recommendation via graph invariant learning. *arXiv preprint arXiv:2504.10432*, 2025.
- [Yu *et al.*, 2021] Junliang Yu, Hongzhi Yin, Jundong Li, Qinyong Wang, Nguyen Quoc Viet Hung, and Xiangliang Zhang. Self-supervised multi-channel hypergraph convolutional network for social recommendation. In *Proceedings of the web conference 2021*, pages 413–424, 2021.
- [Zhang *et al.*, 2023] Yiming Zhang, Lingfei Wu, Qi Shen, Yitong Pang, Zhihua Wei, Fangli Xu, Ethan Chang, and Bo Long. Graph learning augmented heterogeneous graph neural network for social recommendation. *ACM Transactions on Recommender Systems*, 1(4):1–22, 2023.