

Dual-Channel Hybrid Graph Neural Network for Mobility Social Relationship Inference

Liangkun Chen^{2†}, Xiang Li^{2†}, Guiyuan Jiang^{1,2}, Zhongying Zhao³, Junyu Dong^{1,2}, Yanwei Yu^{1,2}

¹State Key Laboratory of Physical Oceanography, Ocean University of China

²Faculty of Information Science and Engineering, Ocean University of China

³College of Computer Science and Engineering, Shandong University of Science and Technology
 {clk, lixiang1202}@stu.ouc.edu.cn, jiangguiyuan@ouc.edu.cn, zyzhao@sdust.edu.cn,
 {dongjunyu, yuyanwei}@ouc.edu.cn

Abstract

Inferring latent social ties from large-scale spatiotemporal mobility traces is a foundational AI task with broad applicability. Existing hypergraph-based methods often model higher-order relations by treating hyperedges as static snapshots, thus failing to capture the temporal dynamics and co-evolution of user interactions. Meanwhile, many approaches still struggle to distinguish stable social gatherings from transient noisy co-occurrences. To address these challenges, we propose a novel **Dual-Channel Hybrid Graph Neural Network (HyGNN)** that jointly models temporal dynamics and high-order structural dependencies. The framework consists of two complementary components: (1) a temporal meeting graph channel that decomposes multi-user interactions into ordered snapshots and aligns trajectory evolution through message propagation, and (2) a gathering hypergraph channel that uses a structure-aware encoder with homogeneity-guided aggregation to filter noise in high-order co-occurrences. A mutual-information-aware fusion module integrates both views while preserving their distinct semantics. Extensive experiments across three real-world datasets show that HyGNN consistently outperforms state-of-the-art baselines in PRAUC and ROCAUC metrics, and robustly captures the nonlinear interplay between mobility and sociality. The code is available at <https://github.com/KunLiangChen/HyGNN>.

1 Introduction

The widespread adoption of Location-Based Services (LBS) (e.g., Foursquare, Yelp) continually generates large-scale spatiotemporal mobility data [Li *et al.*, 2023a], providing fine-grained traces of individual routines and co-presence. Inferring latent social relationships from such trajectories is a central problem in computational social science [Cho *et al.*, 2011; Yang *et al.*, 2019]. Mobility-derived interaction signals are known to be highly predictive of social ties [Wang *et al.*, 2014], enabling downstream applications including

friend suggestion [Yu *et al.*, 2020; Yu *et al.*, 2021], POI and session recommendation [Dai *et al.*, 2021; Dai *et al.*, 2022; Li *et al.*, 2024; Fu *et al.*, 2026], and context-aware advertising [Sharma *et al.*, 2021].

A substantial body of work leveraging advances in Network Representation Learning (NRL) [Cai *et al.*, 2023; Wu *et al.*, 2022; Dai *et al.*, 2023; Wang *et al.*, 2022; Antelmi *et al.*, 2023] has led to more sophisticated methods for social relationship inference. A common paradigm first constructs user interaction graphs from check-ins and then learns low-dimensional user embeddings whose similarities predict friendship [Backes *et al.*, 2017; Yu *et al.*, 2018; Wu *et al.*, 2019; Li *et al.*, 2023b]. To encode richer semantics, another line of research adopts heterogeneous or hypergraph formulations that treat check-ins as higher-order relations and preserve higher-order proximity between users and events [Xia *et al.*, 2022; Yang *et al.*, 2019; Li *et al.*, 2022]. More recently, MRGAN [Qin *et al.*, 2024] explicitly models multi-relational structures to capture both direct and indirect ties (e.g., personalized/common preferences) via relation attention and cross-relation attention, while H³GNN [Li *et al.*, 2025c] employs a heterogeneous hyperbolic hypergraph to alleviate geometric distortion in tree-like social structures and improve representation quality.

However, two fundamental challenges remain: (1) *Temporal dynamics are collapsed into static links*. Most approaches summarize mobility interactions as time-agnostic edges, implicitly assuming stationarity and overlooking how interactions emerge, strengthen, or decay over time. This static abstraction discards transient or context-dependent meetings that are crucial for understanding real social behavior [Wang *et al.*, 2014]. In practice, human mobility is *bursty* and *periodic* (e.g., weekday/weekend cycles), and link salience depends on *recency* and *ordering*: an edge formed years ago should not carry the same weight as a recent sequence of meetings. Moreover, aggregating across long windows introduces *edge staleness* and makes predictions highly sensitive to the chosen window size; temporal trends (strengthening, fading, seasonality) and short-lived co-presence are indistinguishable under static modeling. (2) *Pair-wise modeling is inherently limited and noisy*. Reducing complex group activities to isolated user pairs inflates spurious co-occurrences

and obscures high-order organization (e.g., gatherings, communities). At popular POIs with k participants, a pair-wise view explodes into $O(k^2)$ dyads, magnifying noise due to incidental co-location and popularity bias. Without an explicit group context, models cannot *localize* where/when interactions matter, nor can they quantify *within-event consistency* (homogeneity) or participants’ relative importance; repeated cohesive gatherings and one-off crowd encounters become indistinguishable. This decomposition also fails under overlapping groups and role asymmetries (leaders/companions), resulting in fragile predictions under structural noise.

To address these challenges, we propose HyGNN, a dual-channel representation learning framework that jointly models temporal pairwise dynamics and high-order group interactions from mobility data. Concretely, HyGNN builds a temporal meeting graph learning channel that segments interaction sequences into ordered snapshots and performs temporal message propagation, explicitly regularized to align both trajectory positions and evolution trends of friends; and a gathering hypergraph learning channel that treats multi-user co-occurrences as hyperedges and employs a structure-aware encoder with homogeneity-guided aggregation and adaptive augmentations. The two channels are then aligned and integrated via a dual-channel guided fusion module that performs semantic completion with auxiliary relational views (e.g., friendship, personalized/common preferences) and preserves cross-channel mutual information while preventing over-alignment. This design allows HyGNN to preserve and exploit temporal variability instead of collapsing it, and denoise pairwise artifacts by explicitly modeling high-order gatherings. Also, this design allows HyGNN to emphasize interaction contexts in both time (when) and structure (which gatherings/POIs). Extensive experiments on three real-life mobility datasets demonstrate that the proposed HyGNN significantly outperforms all state-of-the-art baselines, with an average improvement of 3.74% and 4.42% in terms of PRAUC and ROCAUC. This work contributes as follows:

- We propose HyGNN, a dual-channel architecture that couples the temporal meeting graph learning with the gathering hypergraph learning, enabling interaction localization (when/where) and *salience quantification* while mitigating pair-wise noise.
- We introduce a dual-channel guided fusion mechanism that completes semantics with auxiliary relations and aligns cross-channel information via mutual-information-aware objectives, preserving channel-specific signals without over-alignment.
- Extensive experiments on three large-scale mobility datasets show that HyGNN consistently outperforms state-of-the-art baselines, with average improvements of 3.74% in PRAUC and 4.42% in ROC-AUC.

2 Related Work

Traditional Social Relationship Mining. Early work analyzes raw check-ins or mobile traces to link mobility patterns with social ties. *However, these approaches are largely descriptive and aggregate; they seldom treat group co-presence*

as a primary structure, and incorporate time only as coarse covariates rather than explicit dynamics—limiting interaction localization and temporal resolution.

Graph Learning for Social Relationship Inference. With the rise of representation learning, research shifted toward graph abstractions of latent relations [Cai *et al.*, 2023; Wu *et al.*, 2022; Wang *et al.*, 2022; Antelmi *et al.*, 2023; Dai *et al.*, 2023]. Random-walk methods (e.g., walk2friends [Backes *et al.*, 2017]) embed user-location bipartite graphs; Yu *et al.* [Yu *et al.*, 2018] further build meeting graphs with hierarchical walk sampling. GCN-based methods (e.g., Heter-GCN [Wu *et al.*, 2019]) perform message passing on heterogeneous mobility graphs, modeling co-occurrence as pairwise edges and inferring ties via embedding similarity. These methods set strong baselines but typically summarize interactions as *static* pairwise links, making them sensitive to popularity bias and noise. *Consequently, they offer limited ability to pinpoint where/when interactions matter or to weight event homogeneity in crowded venues.*

Multi-Relational and Structure-Learning Perspectives. Recent work adopts multiple relational views and structure learning to reduce noise and enrich semantics [Yan *et al.*, 2025; Qin *et al.*, 2023]. Closest to our baselines, MRGAN [Qin *et al.*, 2024] models *direct* and *indirect* relations (e.g., personalized and common preferences) and integrates their impacts with relation-specific and cross-relation attention. Although strong, MRGAN remains largely pairwise and time-agnostic in meeting modeling, underusing temporal evolution and group context. *Without explicit group co-presence and temporal ordering, multi-relation aggregation may confuse incidental co-location with meaningful gatherings and fail to localize salient events.*

Hypergraph Learning for Social Relationship Inference. Hypergraphs naturally encode multi-user co-occurrences, capturing high-order interactions that pairwise graphs miss. Recent studies leverage hypergraphs to model LBSN complexities, including spatio-temporal semantics [Yang *et al.*, 2019; Yang *et al.*, 2020], contrastive learning for friend recommendation [Li *et al.*, 2022], and high-order dynamic contexts [Trung *et al.*, 2022]. HHL-Traj [Cao *et al.*, 2024] uses hypergraph neural networks to capture spatial-behavioral patterns via trajectory-based hyperedges. Modeling spatial-behavioral patterns via trajectory-based hypergraph neural networks. H³GNN [Li *et al.*, 2025c] advances this line by coupling hypergraphs with *hyperbolic* geometry to reduce distortion in tree-like social structures, offering a strong hypergraph baseline for LBSNs. However, most hypergraph methods, including H³ GNN, do not explicitly model the fine-grained temporal evolution of meetings; they may insufficiently distinguish repeated cohesive gatherings from transient crowd effects without such a mechanism. *Thus, high-order structure alone cannot resolve when a group interaction is salient versus incidental, nor can it mitigate temporal staleness.*

3 Problem Definition

We set $U = \{u_1, u_2, \dots, u_N\}$ as the set of users and the set of POIs $P = \{p_1, p_2, \dots, p_M\}$.

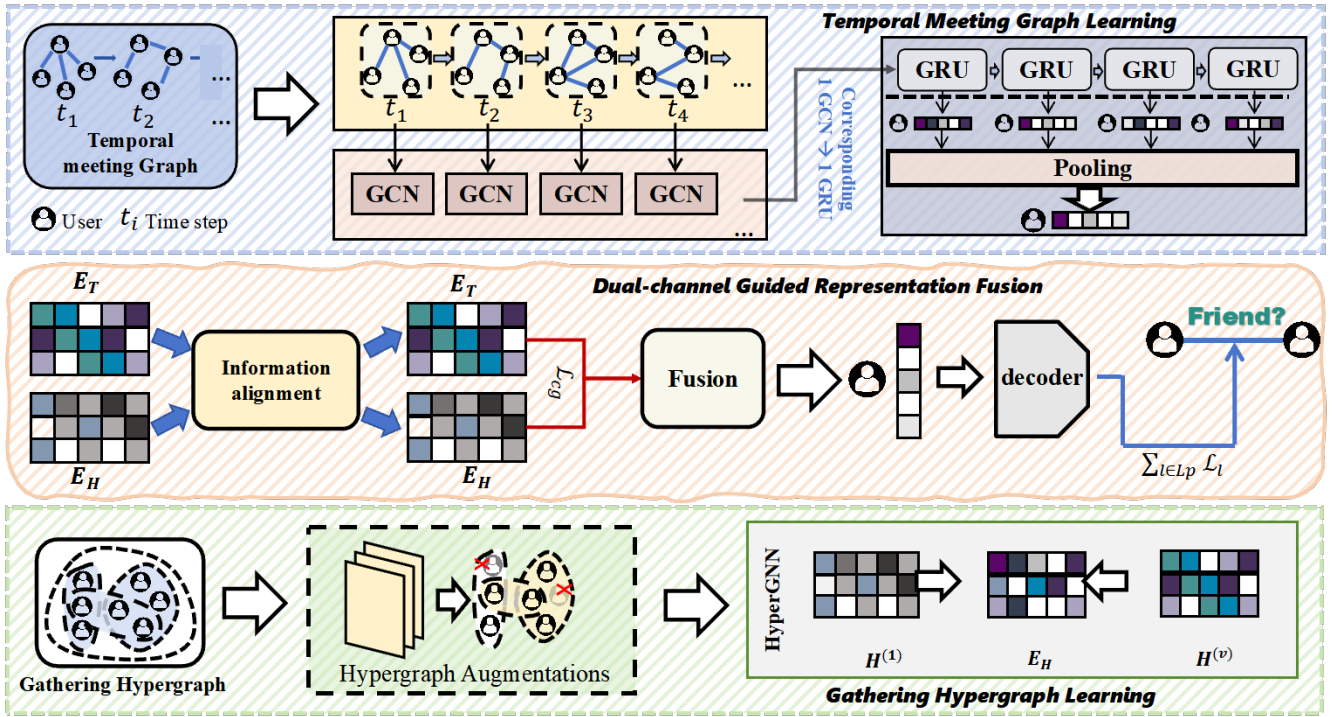


Figure 1: The overview of the proposed HyGNN.

Definition 1 (Check-in Record). We denote a check-in as $\langle u, t, p \rangle$, indicating that **user** u performs a check-in at **POI** p at **time** t . The POI p is represented by $\langle l, p_{id} \rangle$, where l is the geographical location (latitude and longitude) and p_{id} is the unique POI index.

Definition 2 (Check-in Sequence). For user u , the check-in sequence C_{seq}^u is a chronological sequence of check-ins, denoted as $(\langle u, t_1, p_1 \rangle, \dots, \langle u, t_i, p_i \rangle)$.

In human mobility analysis, social ties are often evidenced by users' spatiotemporal co-occurrences, commonly referred to as **Meeting Events** in the literature [Yu et al., 2018; Qin et al., 2023; Qin et al., 2024].

Definition 3 (Meeting Event). A meeting event occurs between two users u_i and u_j if they have check-ins at the same POI within a temporal threshold δ . Formally, let $\langle u_i, t, p \rangle \in C_{seq}^{u_i}$ and $\langle u_j, t', p' \rangle \in C_{seq}^{u_j}$ be two check-ins. A meeting event is established if $p = p'$ and $|t - t'| \leq \delta$.

Pairwise meetings over-fragment group co-presence and inflate spurious dyads. We therefore define a **Gathering Event** that aggregates users co-present at a POI within Δt into a high-order relation, preserving group context and aiding interaction localization and denoising.

Definition 4 (Gathering Event). We define a **Gathering Event** as a user set $\mathcal{G}_{p,t} \subseteq U$ at POI p , where $\mathcal{G}_{p,t} = \{u_i \mid \langle u_i, t_i, p_i \rangle, p_i = p, t \leq t_i \leq t + \Delta t\}$. Each $\mathcal{G}_{p,t}$ represents users who checked in at the same POI within a temporal window Δt .

Based on the definitions above, the problem addressed in this paper is formulated as follows:

Problem (Social Relationship Inference from Mobility Data). Given a user set U , the objective is to learn an inference model $\mathcal{F}(u_i, u_j) \rightarrow \hat{y}$ for each user pair based on their check-in sequences and partially observed friendship labels. The prediction $\hat{y}$ is binary, where 1 indicates friendship and 0 indicates no friendship.

4 Methodology

This section presents HyGNN, a dual-channel framework that integrates *temporal pairwise* signals and *high-order gathering* patterns for social relationship inference. As illustrated in Figure 1, the overall architecture comprises three components: (1) *Temporal Meeting Graph Learning*, (2) *Gathering Hypergraph Learning*, and (3) *Dual-channel Guided Representation Fusion* module.

To address structural noise in high-order modeling and to capture temporal dynamics, we construct a hypergraph and a temporal graph from raw check-ins and design two independent channels to learn their respective representations. A knowledge-completion module further enhances representation quality by leveraging auxiliary relational signals. Finally, both channels are fused through a lightweight pipeline. Positive/negative user pairs are sampled from the friendship graph to supervise learning, and we specify dedicated objectives for each channel as well as for the fusion stage.

4.1 Temporal Meeting Graph Learning

We introduce a temporal graph channel that explicitly encodes *when* user meetings occur and *how* such interactions evolve. The proposed module organizes meeting events into

a sequence of monthly snapshots and applies a time-aware encoder to derive temporal user embeddings.

Temporal Meeting Graph Construction

Given the check-ins $\langle u, t, p \rangle$ (user $u \in U$, timestamp t , POI p), we group records by month and POI. For each POI, we collect all users who visited it within the same month and extract monthly meeting edges E_M^t from meeting events:

$$E_M^t = \{(u_i, u_j) \mid \exists \langle u_i, t_i, p_i \rangle, \langle u_j, t_j, p_j \rangle, p_i = p_j, |t_i - t_j| \leq \delta, t_i, t_j \in \text{month } t\}. \quad (1)$$

Here, δ is the temporal co-occurrence threshold, $t \in \{1, \dots, T\}$ indexes months. We define the *Temporal Meeting Graph* as $\mathcal{G}_M^T = \{G_M^1, G_M^2, \dots, G_M^T\}$, where $G_M^t = (U, E_M^t)$ encodes monthly meeting relationships among the same node set U . Segmenting by months balances sparsity and temporal resolution, providing stable snapshots while preserving meaningful interaction dynamics.

Temporal Message Propagation

To learn time-evolving patterns, we propagate messages over the sequence $\mathcal{G}_M^T$ so that HyGNN captures both intra-month structure and inter-month evolution.

(a) Node Feature Initialization. For month t , let $\tilde{\mathbf{A}}_t \in \mathbb{R}^{N \times N}$ be the normalized adjacency of G_M^t (row/degree normalization as in LightGCN). The input feature of node u_i at t , denoted $\mathbf{X}_i^t$, is set to the normalized degree of u_i in G_M^t , serving as a lightweight activity indicator. We collect features into $\mathbf{X}^t \in \mathbb{R}^{N \times 1}$ and project them during propagation.

(b) T-GCN Cell. Following the time-aware design, the core cell proceeds from $t = 1$ to T . It blends the current topology signal with the previous hidden state to update $\mathbf{H}^t \in \mathbb{R}^{N \times D}$, where the i -th row is the embedding of user u_i at time t . Specifically, we apply LightGCN [He *et al.*, 2020] to both current inputs and past states with $\tilde{\mathbf{A}}_t$:

$$\mathbf{Z}_X^t = \text{LightGCN}(\mathbf{X}^t, \tilde{\mathbf{A}}_t), \quad (2)$$

$$\mathbf{Z}_H^{t-1} = \text{LightGCN}(\mathbf{H}^{t-1}, \tilde{\mathbf{A}}_t), \quad (3)$$

and then update with a GRU as in T-GCN:

$$\mathbf{H}^t = \text{GRU}(\mathbf{Z}_X^t, \mathbf{Z}_H^{t-1}). \quad (4)$$

Here, $\mathbf{Z}_X^t, \mathbf{Z}_H^{t-1} \in \mathbb{R}^{N \times D}$ are topology-refined signals; the GRU operates row-wise over users, carrying temporal memory of each user's interaction trajectory.

(c) Temporal Meeting Embedding. The cell outputs $\mathbf{H}^1, \dots, \mathbf{H}^T$ are aggregated by mean pooling along the time axis to obtain a fixed-size temporal representation:

$$\mathbf{E}_T = \frac{1}{T} \sum_{t=1}^T \mathbf{H}^t. \quad (5)$$

The matrix $\mathbf{E}_T \in \mathbb{R}^{N \times D}$ serves as the temporal embedding passed to the fusion module. Intuitively, $\mathbf{E}_T$ summarizes both the level (*position*) and variability (*trend*, captured in the recurrent state) of users' meeting behaviors across months.

4.2 Gathering Hypergraph Learning

We detail the Gathering Hypergraph Channel, which captures high-order co-occurrence via hyperedges while incorporating per-user overlap and per-hyperedge homogeneity. Notations: let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote the hypergraph with $\mathcal{V} = U$ (nodes are users) and $\mathcal{E}$ the set of gathering events (hyperedges). For hyperedge $e \in \mathcal{E}$, $\mathcal{U}(e) \subseteq \mathcal{V}$ is its incident node set; $\mathcal{P}(e)$ the set of unordered node pairs within e .

Gathering Hypergraph Construction

We instantiate a hypergraph whose nodes are users and hyperedges correspond to gathering events (as defined in the problem setting). The user *overlap* feature measures participation intensity across hyperedges: $\sum_{e \in \mathcal{E}} \mathbb{I}(u_i \in \mathcal{U}(e))$, normalized by the maximum user count for comparability (yielding a scalar per user). To characterize within-event consistency, we define the *homogeneity* of a hyperedge e via pairwise co-occurrence similarity $S(u_i, u_j)$ (the number of hyperedges containing both u_i and u_j), and take its average over all pairs in e as follows:

$$\text{Homo}(e) = \frac{1}{|\mathcal{P}(e)|} \sum_{u_i, u_j \in \mathcal{P}(e)} S(u_i, u_j), \quad (6)$$

followed by max-normalization across $\mathcal{E}$. Larger $\text{Homo}(e)$ indicates stronger, more consistent co-occurrence among participants and will be used to weight edge2node aggregation.

Gathering-aware Hypergraph Encoder

We build augmented views to train a hypergraph-aware encoder that is robust to structural noise.

(a) Hypergraph Augmentation. Refer to CHGNN [Song *et al.*, 2024], for each hyperedge, we consider three operations: (1) *Identity* (no change); (2) *Drop Hyperedge*; (3) *Random Node Masking*. We use HyperConv layers [Bai *et al.*, 2021] to produce a hyperedge embedding matrix $E_{he} \in \mathbb{R}^{E \times op}$ by iteratively aggregating nodes and hyperedges (with $op = 3$). We treat E_{he} as per-hyperedge operation logits and apply Gumbel-softmax to sample the augmentation per hyperedge and construct an augmented view. Masked nodes are sampled with an overlap-aware strategy that more likely masks high-overlap nodes, preserving key semantics while encouraging invariance to redundant co-occurrences.

(b) HGNN with Structure-aware Aggregation. After projecting node features, we perform mean aggregation to obtain hyperedge representations:

$$\mathbf{Z}_e = \frac{1}{|\mathcal{V}(e)|} \sum_{v \in \mathcal{V}(e)} \mathbf{H}_v^{(1)}, \quad (7)$$

and compute homogeneity-weighted attention from hyperedges to nodes:

$$\alpha_{ev} = \frac{\text{Homo}(e)}{\sum_{e' \in \mathcal{E}(v)} \text{Homo}(e')}. \quad (8)$$

Weighted edge-to-node aggregation then yields:

$$\mathbf{H}_v^{(v)} = \sum_{e \in \mathcal{E}(v)} \alpha_{ev} \cdot \mathbf{Z}_e. \quad (9)$$

We apply residual and normalization:

$$\mathbf{H}^{(2)} = \mathbf{H}^{(1)} + \mathbf{H}^{(v)}, \mathbf{X}_{hy} = \frac{\mathbf{H}^{(2)}}{\|\mathbf{H}^{(2)}\|_2}. \quad (10)$$

After activation and dropout, HGNN outputs the hypergraph embedding $\mathbf{E}_H \in \mathbb{R}^{N \times D}$ for the current view; if two augmented views are used, they are averaged to obtain a single $\mathbf{E}_H$ for downstream fusion.

4.3 Dual-channel Guided Representation Fusion Representation Fusion

The design of the dual-channel fusion module is motivated by the distinct yet complementary semantic roles of the two channels. [Li *et al.*, 2025a; Li *et al.*, 2025b] To integrate the dynamic interactions from the temporal graph and high-order correlations from the hypergraph, we propose a gated cross-attention fusion module. Indirect relations have proven effective in MRGAN [Qin *et al.*, 2024], where a *semantic fender* aids cross-relation completion and alignment. Since the two channels deliver embeddings of different richness, we adopt a similar idea to calibrate and align them before fusion.

To keep complexity low, we convolve the friendship, common-preference, and personalized-preference graphs with LightGCN, and use the resulting embeddings $\mathbf{H}_F, \mathbf{H}_{PP}, \mathbf{H}_{CP} \in \mathbb{R}^{N \times D}$ to enrich both channels via semantic completion. Concretely, the temporal embedding $\mathbf{E}_T$ and the hypergraph embedding $\mathbf{E}_H$ are refined through interaction with $\mathbf{H}_F, \mathbf{H}_{PP}$, and $\mathbf{H}_{CP}$, providing a unified space for subsequent fusion. Then we project both channels into a shared latent space via lightweight MLP heads:

$$\mathbf{P}_H = f_H(\mathbf{E}_H), \quad \mathbf{P}_T = f_T(\mathbf{E}_T), \quad (11)$$

and apply single-head cross-attention (queries from $\mathbf{P}_H$, keys/values from $\mathbf{P}_T$):

$$\mathbf{R} = \text{softmax}\left(\frac{\mathbf{P}_H W_q (\mathbf{P}_T W_k)^\top}{\sqrt{d}}\right) (\mathbf{P}_T W_v), \quad (12)$$

where $W_q, W_k, W_v \in \mathbb{R}^{D \times d}$ and d is the attention width. Finally, we form the fused representation via a gated residual and a small feed-forward network:

$$\mathbf{H} = \text{LayerNorm}(\mathbf{E}_H + \text{MLP}([\mathbf{P}_H + \mathbf{g} \odot \mathbf{R} \parallel \mathbf{P}_T])), \quad (13)$$

where $\mathbf{g} \in \mathbb{R}^d$ is a learnable gate controlling temporal signal flow, $\mathbf{H} \in \mathbb{R}^{N \times D}$ is the unified user embedding used for friendship prediction; its i -th row $\mathbf{h}_i$ corresponds to user u_i .

Model Optimization

HyGNN is trained under a dual-channel objective that couples the main prediction with mutual information-preserving regularizers and the channel-specific losses.

For the main task, given positive pairs Θ^+ and negative pairs Θ^- sampled from the friendship graph (with a fixed negative ratio per mini-batch), let $\mathbf{h}_i \in \mathbb{R}^D$ be the fused embedding of u_i in $\mathbf{H}$. We define the pairwise affinity

$$\phi(u_i, u_j) = \mathbf{h}_i^\top \mathbf{h}_j, \quad (14)$$

and the binary cross-entropy loss per channel

$$\mathcal{L}_{lp}^{(\cdot)} = - \sum_{(u_i, u_j) \in \Theta^+} \log \sigma(\phi(u_i, u_j)) - \sum_{(u_i, u_j) \in \Theta^-} \log(1 - \sigma(\phi(u_i, u_j))), \quad (15)$$

with $\sigma(\cdot)$ the sigmoid. The total link-prediction loss is

$$\mathcal{L}_{lp} = \mathcal{L}_{lp}^T + \mathcal{L}_{lp}^H. \quad (16)$$

We further include auxiliary completion tasks on CP/PP graphs (binary cross-entropy), denoted by $\mathcal{L}_{mr}$.

To align channels while preserving uniqueness, we introduce a contrastive guidance loss. Let

$$\mathbf{S}_{ij} = \frac{(h_i^{(h)})^\top h_j^{(t)}}{\tau}, \quad (17)$$

with temperature τ and superscripts (h) , (t) indicating hypergraph and temporal-channel embeddings. The bi-directional bootstrap loss is as follows:

$$\mathcal{L}_{cg} = -\frac{1}{2N} \sum_{i=1}^N \left[\log \frac{\exp(\mathbf{S}_{ii})}{\sum_j \exp(\mathbf{S}_{ij})} + \log \frac{\exp(\mathbf{S}_{ii})}{\sum_j \exp(\mathbf{S}_{ji})} \right]. \quad (18)$$

Now we formally set the final objective:

$$\mathcal{L}_{all} = \mathcal{L}_{lp} + \lambda_{mr} \mathcal{L}_{mr} + \lambda_{cg} \mathcal{L}_{cg}, \quad (19)$$

where λ_{mr} and λ_{cg} are non-negative scalars that balance auxiliary tasks and cross-channel alignment. In practice, we tune $(\lambda_{mr}, \lambda_{cg})$ on a validation split; model selection uses PRAUC/ROCAUC for the social link prediction task.

5 Experiments

Dataset	Gowalla	Foursquare	Foursquare
City	AUSTIN	TKY	NEW YORK
#users	7,897	8,443	9,088
#check-ins	336,309	954,835	251,323
#POIs	19,095	84,820	25,395
#friends	76,212	53,890	21,431

Table 1: Statistics of datasets

This section is dedicated to model evaluation, wherein we assess the performance of HyGNN against established baseline methods using three real-world datasets. Specifically, our experiments are structured to address the following research questions:

- **RQ1:** How does HyGNN perform compare with existing state-of-the-art baselines for social relationship inference?
- **RQ2:** How do key components contribute to the performance?
- **RQ3:** How do the key hyperparameters influence the performance of HyGNN?

Method	Austin		TKY		NYC	
	ROCAUC	PRAUC	ROCAUC	PRAUC	ROCAUC	PRAUC
DeepWalk (<i>KDD</i> , 2014)	0.6035	0.6172	0.5963	0.6130	0.6149	0.6227
node2vec (<i>KDD</i> , 2016)	0.6212	0.6159	0.5967	0.5900	0.6276	0.6359
GCN (<i>ICLR</i> , 2017)	0.6275	0.6318	0.6001	0.6192	0.6253	0.6191
R-GCN (<i>ESWC</i> , 2018)	0.6928	0.7013	0.6683	0.6777	0.6850	0.7051
HAN (<i>WWW</i> , 2019)	0.7685	0.7863	0.7243	0.7531	0.7554	0.7645
HGMAE (<i>AAAI</i> , 2023)	0.8265	0.8364	0.7886	0.8100	0.8088	0.8262
emb-cat (<i>Ubicomp</i> , 2018)	0.6405	0.6337	0.6353	0.6334	0.6399	0.6581
Heter-GCN (<i>IJCAI</i> , 2019)	0.7753	0.7867	0.7404	0.7399	0.7501	0.7673
MVMN (<i>TNNLS</i> , 2020)	0.7612	0.7831	0.7271	0.7323	0.7643	0.7723
LBSN2Vec++ (<i>TKDE</i> , 2022)	0.7521	0.7472	0.7116	0.7069	0.7337	0.7405
MSC-LBSN (<i>TKDE</i> , 2022)	0.7892	0.8054	0.7522	0.7516	0.7807	0.7811
HHGNN (<i>CIKM</i> , 2022)	0.7963	0.8016	0.7845	0.8018	0.7824	0.8071
HMGCL (<i>WWWJ</i> , 2023)	0.8356	0.8409	0.7954	0.8142	0.8337	0.8370
SRINet (<i>AAAI</i> , 2023)	0.8315	0.8693	0.7996	0.8241	0.8305	0.8414
MORGAN (<i>IJCAI</i> , 2024)	<u>0.9204</u>	<u>0.9354</u>	0.8747	0.8912	<u>0.8543</u>	<u>0.8695</u>
H ³ GNN (<i>TKDD</i> , 2025)	0.9189	0.9308	<u>0.8791</u>	<u>0.8972</u>	0.8258	0.8543
Ours	0.9478	0.9618	0.9246	0.9430	0.8760	0.9098
Improvement	2.98% ↑	2.82% ↑	5.70% ↑	5.81% ↑	2.54% ↑	4.63% ↑

Table 2: Performance comparison of all models on three real-world datasets. Improvement denotes the improvement of the best results compared with the second-best results. The best results are shown in bold, and the best among baselines is underlined.

5.1 Experiment Settings

Datasets. In the experiment, we used two publicly available real-world check-in datasets. Gowalla [Cho *et al.*, 2011], Foursquare [Yang *et al.*, 2019]. Like MORGAN [Qin *et al.*, 2024], we do not discard inactive users with few check-ins. We present the data statistics in Table 1.

Baselines. To address the need for a structured comparison, we evaluate HyGNN against three categories of models:

- **General Graph Models:** DeepWalk [Perozzi *et al.*, 2014], node2vec [Grover and Leskovec, 2016], and GCN [Kipf and Welling, 2016]. These methods focus on basic pairwise connectivity.
- **Heterogeneous Learning:** R-GCN [Schlichtkrull *et al.*, 2018], HAN [Wang *et al.*, 2019], and HGMAE [Tian *et al.*, 2023]. These capture semantic diversity but often neglect high-order group dynamics.
- **Social Inference SOTA methods:** (1) Walk-based: emb-cat [Yu *et al.*, 2018], LBSN2Vec++ [Yang *et al.*, 2020]; (2) GNN-based: Heter-GCN [Wu *et al.*, 2019], MVMN [Zhang *et al.*, 2020], MSC-LBSN [Trung *et al.*, 2022], SRINet [Qin *et al.*, 2023]; (3) High-order/Hypergraph: HHGNN [Li *et al.*, 2022], HMGCL [Li *et al.*, 2023b], MORGAN [Qin *et al.*, 2024], and the hyperbolic model H³GNN [Li *et al.*, 2025c].

Implementation Details. For our experiments, we allocated the social network friendships as follows: a 25% sample was used for the training set (which informed the social relations within the heterogeneous mobility graph), 5% was randomly selected for validation, and the remaining 70% comprised the testing set. HyGNN is implemented in PyTorch and opti-

mized by Adam [Kingma and Ba, 2014] for parameter learning. We repeat each experiment 10 times with different random seeds and report the mean results. We report mean results across 10 runs with different random seeds. Statistical significance ($p < 0.05$) is determined via two-tailed paired t-tests between our model and the best baseline (underlined).

5.2 Evaluation Metrics

Derived from the former paper [Qin *et al.*, 2024] in our evaluation, we employ the widely used Area Under the ROC Curve (ROCAUC) and the Area Under the Precision-Recall Curve (PRAUC) to evaluate the performance of various methods. In the context of social relationship inference, we calculate the pairwise cosine similarity of learned user representations in each method to derive a score indicating the likelihood of two users being friends.

5.3 Overall Performance (RQ1)

Table 2 presents the comprehensive experimental results on three real-world public datasets, evaluated using the ROCAUC and PRAUC metrics; as shown, our proposed model consistently achieves the SOTA performance across all three datasets and both metrics, with the best scores highlighted in bold, demonstrating superior effectiveness and robust generalization. Quantitatively, our model significantly outperforms all baselines, yielding an average improvement of 3.74% on ROCAUC and 4.42% on PRAUC across the three datasets compared to the next best SOTA method (underlined results); notably, on the Austin dataset MORGAN achieved strong results, yet our HyGNN still managed to achieve further gains of 2.98% and 2.82%, respectively, and on the more chal-

lenging TKY dataset our HyGNN secures the largest relative improvements of 5.70% on ROCAUC and 5.81% on PRAUC over the second-best result (H³GNN). Advanced heterogeneous SOTA methods (MRGAN, H³GNN, and SRINet) achieve the best baseline results by incorporating deeper structure or temporal factors, but contain inherent flaws that limit performance—MRGAN’s static pair-wise modeling of encounter events causes information loss by ignoring dynamics, sequence, and duration; H³GNN, while time-aware, fails to capture fine-grained temporal interaction information; and SRINet, though effective at handling mobility noise, neglects high-order gathering information in multi-user meeting events, limiting its ability to identify strong social ties.

5.4 Ablation Study (RQ2)




Figure 2: Performance (PRAUC) of different model variants on Austin, TKY, and NYC datasets.

To quantify the efficacy of our proposed architectural components, we conduct a rigorous ablation study to systematically evaluate their individual contributions to the overall performance. We ablate four parts as follows:

- **w/o T** removes temporal meeting graph learning module.
- **w/o H** removes gathering hypergraph learning module.
- **w/o A** removes the information alignment
- **w/o F** removes dual-channel guided representation fusion.

The results of ablation studies are summarized in Figure 2. We have reached the following conclusions. (1) All model variants underperform the complete model across every dataset, validating the necessity of the four proposed modules. (2) **w/o F** suffers the largest overall performance drop (avg drop of 3.26%), falling substantially below other variants on all datasets. This indicates that simply concatenating dual-channel embeddings leads to information scattering and noise interference. Thus, the fusion module acts as a crucial feature selector and adaptive weighting mechanism. (3) **w/o T** shows an average drop of 2.67%, most severe on Austin (0.9551→0.9156). This highlights that dynamic temporal modeling is essential, especially where mobility patterns rely on sequential semantics—merely capturing the final encounter structure is insufficient for accurate inference.

5.5 Parameter Sensitivity Analysis (RQ3)

The effect of the hyperparameter λ_{cg}




Figure 3: The effect of the hyperparameter.

Fig. 3a shows the sensitivity of performance (PRAUC) to the fused mutual information loss weight λ_{cg} , tested between 0.1 and 0.5. The optimal λ_{cg} is region-dependent: TKY and Austin favor $\lambda_{cg} \geq 0.4$, while NYC peaks at $\lambda_{cg} = 0.1$. This indicates that while the fused information is beneficial for TKY and Austin, it introduces noise for NYC, minimizing its contribution yields the best result.

The effect of the hyperparameter λ_{mr}

Fig. 3b shows the sensitivity of PRAUC to λ_{mr} (tested from 0.1 to 1.0), which controls cross-relation alignment. Performance first rises then declines as λ_{mr} increases. The best results occur at $\lambda_{mr} = 0.5$, with PRAUCs of 0.9414 (TKY), 0.9070 (NYC), and 0.9545 (Austin). This indicates that moderate alignment effectively balances relation-specific learning with semantic fusion, while excessive weighting causes over-regularization and performance drop.

The effect of hyperparameters τ

Fig. 3c and 3d show the PRAUC sensitivity to the temperature parameter τ , varied from 0.005 by doubling. Both datasets exhibit non-linear responses. TKY reaches its best performance at $\tau = 0.32$ (PRAUC=0.9406), showing relatively stable robustness, while Austin peaks at $\tau = 0.04$ (PRAUC=0.9618) and drops sharply with larger τ , indicating higher sensitivity on this kind of dataset. Overall, a moderate τ range of 0.04–0.32 provides the best balance, with dataset-specific tuning still required.

6 Conclusion

This paper proposes HyGNN, a dual-channel hybrid graph neural network for mobility social relationship inference. HyGNN jointly models temporal graphs and gathering hypergraph structures from mobility data to capture user preferences, learn expressive user representations, and discriminate between the varying influences of direct and indirect relations. Experiments on three real-world datasets show that HyGNN significantly outperforms SOTA baselines.

Acknowledgments

This work is partially supported by the National Natural Science Foundation of China under Grant Nos. 62176243, 62472263 and 62302120, and Key R&D Program of Shandong Province, China under Grant No. 2025CXPT185.

Contribution Statement

Yanwei Yu is the corresponding author. Liangkun Chen and Xiang Li contribute equally.

References

- [Antelmi *et al.*, 2023] Alessia Antelmi, Gennaro Cordasco, Mirko Polato, Vittorio Scarano, Carmine Spagnuolo, and Dingqi Yang. A survey on hypergraph representation learning. *ACM Computing Surveys*, 56(1):1–38, 2023.
- [Backes *et al.*, 2017] Michael Backes, Mathias Humbert, Jun Pang, and Yang Zhang. walk2friends: Inferring social links from mobility profiles. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1943–1957, 2017.
- [Bai *et al.*, 2021] Song Bai, Feihu Zhang, and Philip HS Torr. Hypergraph convolution and hypergraph attention. *Pattern Recognition*, 110:107637, 2021.
- [Cai *et al.*, 2023] Xuheng Cai, Lianghao Xia, Xubin Ren, and Chao Huang. How expressive are graph neural networks in recommendation? In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 173–182, 2023.
- [Cao *et al.*, 2024] Yuan Cao, Lei Li, Xiangru Chen, Xue Xu, Zuojin Huang, and Yanwei Yu. Hypergraph hash learning for efficient trajectory similarity computation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 175–186, 2024.
- [Cho *et al.*, 2011] Eunjoon Cho, Seth A Myers, and Jure Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090, 2011.
- [Dai *et al.*, 2021] Shaojie Dai, Yanwei Yu, Hao Fan, and Junyu Dong. Personalized poi recommendation: spatio-temporal representation learning with social tie. In *International Conference on Database Systems for Advanced Applications*, pages 558–574. Springer, 2021.
- [Dai *et al.*, 2022] Shaojie Dai, Yanwei Yu, Hao Fan, and Junyu Dong. Spatio-temporal representation learning with social tie for personalized poi recommendation. *Data Science and Engineering*, 7(1):44–56, 2022.
- [Dai *et al.*, 2023] Shaojie Dai, Jinshuai Wang, Chao Huang, Yanwei Yu, and Junyu Dong. Dynamic multi-view graph neural networks for citywide traffic inference. *ACM Transactions on Knowledge Discovery from Data*, 17(4):1–22, 2023.
- [Fu *et al.*, 2026] Haoyan Fu, Zhida Qin, Wenhao Xue, Qixian Wang, Xufeng Liang, Tianyu Huang, and John CS Lui. Multi-relation enhanced dynamic hypergraph for session-based recommendation. *ACM Transactions on Information Systems*, 44(2):1–28, 2026.
- [Grover and Leskovec, 2016] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016.
- [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.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [Li *et al.*, 2022] Yongkang Li, Zipei Fan, Jixiao Zhang, Dengheng Shi, Tianqi Xu, Du Yin, Jinliang Deng, and Xuan Song. Heterogeneous hypergraph neural network for friend recommendation with human mobility. In *Proceedings of the 31st ACM international conference on information & knowledge management*, pages 4209–4213, 2022.
- [Li *et al.*, 2023a] Shuzhe Li, Wei Chen, Bingqi Yan, Zhen Li, Shunzhi Zhu, and Yanwei Yu. Self-supervised contrastive representation learning for large-scale trajectories. *Future Generation Computer Systems*, 148:357–366, 2023.
- [Li *et al.*, 2023b] Yongkang Li, Zipei Fan, Du Yin, Renhe Jiang, Jinliang Deng, and Xuan Song. Hmgcl: Heterogeneous multigraph contrastive learning for lbsn friend recommendation. *World Wide Web*, 26(4):1625–1648, 2023.
- [Li *et al.*, 2024] Shuzhe Li, Wei Chen, Bin Wang, Chao Huang, Yanwei Yu, and Junyu Dong. Mcn4rec: Multi-level collaborative neural network for next location recommendation. *ACM Transactions on Information Systems*, 42(4):1–26, 2024.
- [Li *et al.*, 2025a] Xiang Li, Chaofan Fu, Zhongying Zhao, Guangjie Zheng, Chao Huang, Yanwei Yu, and Junyu Dong. Dual-channel multiplex graph neural networks for recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 37(6):3327–3341, 2025.
- [Li *et al.*, 2025b] Xiang Li, Changsheng Shui, Zhongying Zhao, Junyu Dong, and Yanwei Yu. Multi-channel hypergraph contrastive learning for matrix completion. *ACM Transactions on Information Systems*, 44(1):1–27, 2025.
- [Li *et al.*, 2025c] Yongkang Li, Zipei Fan, and Xuan Song. Heterogeneous hyperbolic hypergraph neural network for friend recommendation in location-based social networks. *ACM Transactions on Knowledge Discovery from Data*, 19(3):1–29, 2025.

- [Perozzi *et al.*, 2014] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710, 2014.
- [Qin *et al.*, 2023] Guangming Qin, Lexue Song, Yanwei Yu, Chao Huang, Wenzhe Jia, Yuan Cao, and Junyu Dong. Graph structure learning on user mobility data for social relationship inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 4578–4586, 2023.
- [Qin *et al.*, 2024] Guangming Qin, Jianpeng Qi, Bin Wang, Guiyuan Jiang, Yanwei Yu, and Junyu Dong. Multi-relational graph attention network for social relationship inference from human mobility data. In *33rd International Joint Conference on Artificial Intelligence (IJCAI’24)*, pages 2315–2323. International Joint Conferences on Artificial Intelligence Organization, 2024.
- [Schlichtkrull *et al.*, 2018] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018.
- [Sharma *et al.*, 2021] Anshuman Sharma, Yogesh K Dwivedi, Vikas Arya, and Muhammad Qutubuddin Siddiqui. Does sms advertising still have relevance to increase consumer purchase intention? a hybrid pls-sem-neural network modelling approach. *Computers in Human Behavior*, 124:106919, 2021.
- [Song *et al.*, 2024] Yumeng Song, Yu Gu, Tianyi Li, Jianzhong Qi, Zhenghao Liu, Christian S Jensen, and Ge Yu. Chgnn: A semi-supervised contrastive hypergraph learning network. *IEEE Transactions on Knowledge and Data Engineering*, 36(9):4515–4530, 2024.
- [Tian *et al.*, 2023] Yijun Tian, Kaiwen Dong, Chunhui Zhang, Chuxu Zhang, and Nitesh V Chawla. Heterogeneous graph masked autoencoders. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 9997–10005, 2023.
- [Trung *et al.*, 2022] Huynh Thanh Trung, Tong Van Vinh, Nguyen Thanh Tam, Jun Jo, Hongzhi Yin, and Nguyen Quoc Viet Hung. Learning holistic interactions in lbsns with high-order, dynamic, and multi-role contexts. *IEEE Transactions on Knowledge and Data Engineering*, 35(5):5002–5016, 2022.
- [Wang *et al.*, 2014] Hongjian Wang, Zhenhui Li, and Wang-Chien Lee. Pgt: Measuring mobility relationship using personal, global and temporal factors. In *2014 IEEE International Conference on Data Mining*, pages 570–579, 2014.
- [Wang *et al.*, 2019] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The world wide web conference*, pages 2022–2032, 2019.
- [Wang *et al.*, 2022] Xiao Wang, Deyu Bo, Chuan Shi, Shao-hua Fan, Yanfang Ye, and Philip S Yu. A survey on heterogeneous graph embedding: methods, techniques, applications and sources. *IEEE transactions on big data*, 9(2):415–436, 2022.
- [Wu *et al.*, 2019] Yongji Wu, Defu Lian, Shuowei Jin, and Enhong Chen. Graph convolutional networks on user mobility heterogeneous graphs for social relationship inference. In *IJCAI*, pages 3898–3904, 2019.
- [Wu *et al.*, 2022] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. Graph neural networks in recommender systems: a survey. *ACM Computing Surveys*, 55(5):1–37, 2022.
- [Xia *et al.*, 2022] Lianghao Xia, Chao Huang, Yong Xu, Jia-shu Zhao, Dawei Yin, and Jimmy Huang. Hypergraph contrastive collaborative filtering. In *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*, pages 70–79, 2022.
- [Yan *et al.*, 2025] Ziang Yan, Xingyu Zhao, Hanqing Ma, Wei Chen, Jianpeng Qi, Yanwei Yu, and Junyu Dong. Correlation-attention masked temporal transformer for user identity linkage using heterogeneous mobility data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 39, pages 12999–13007, 2025.
- [Yang *et al.*, 2019] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudre-Mauroux. Revisiting user mobility and social relationships in lbsns: a hypergraph embedding approach. In *The world wide web conference*, pages 2147–2157, 2019.
- [Yang *et al.*, 2020] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudré-Mauroux. Lbsn2vec++: Heterogeneous hypergraph embedding for location-based social networks. *IEEE Transactions on Knowledge and Data Engineering*, 34(4):1843–1855, 2020.
- [Yu *et al.*, 2018] Yanwei Yu, Hongjian Wang, and Zhenhui Li. Inferring mobility relationship via graph embedding. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–21, 2018.
- [Yu *et al.*, 2020] Junliang Yu, Hongzhi Yin, Jundong Li, Min Gao, Zi Huang, and Lizhen Cui. Enhancing social recommendation with adversarial graph convolutional networks. *IEEE Transactions on knowledge and data engineering*, 34(8):3727–3739, 2020.
- [Yu *et al.*, 2021] Junliang Yu, Hongzhi Yin, Min Gao, Xin Xia, Xiangliang Zhang, and Nguyen Quoc Viet Hung. Socially-aware self-supervised tri-training for recommendation. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 2084–2092, 2021.
- [Zhang *et al.*, 2020] Wei Zhang, Xin Lai, and Jianyong Wang. Social link inference via multiview matching network from spatiotemporal trajectories. *IEEE transactions on neural networks and learning systems*, 34(4):1720–1731, 2020.