Journal of King Saud University Computer and Information Sciences (2025) 37:245
https://doi.org/10.1007/s44443-025-00255-4

ORIGINAL PAPER

DHGNN: A dynamic heterogeneous graph neural network for interpretable inventor collaboration prediction

Xiaodong Xie[1] · Jie Wu[1] · Mengjia Xiang[2] · Jianting Tang[1] · Yongxiang Sheng[1]

Received: 5 June 2025 / Accepted: 23 August 2025 / Published online: 13 October 2025

Abstract
Inventor collaboration networks are both temporally evolving and semantically heterogeneous, involving diverse entity types and complex relational structures. These characteristics present significant challenges for accurately predicting future collaborations. We propose a Dynamic Heterogeneous Graph Neural Network (DHGNN) that jointly captures temporal dynamics and semantic dependencies while enabling interpretable prediction. DHGNN constructs cumulative heterogeneous graph snapshots and integrates relation-aware message passing with meta-path–guided multi-hop aggregation to capture both local and high-order collaboration patterns. A hierarchical attention mechanism combines continuous-time positional encoding with causal masking (preventing information leakage from future timestamps) to dynamically align temporal and semantic signals. For interpretability, DHGNN introduces a two-tier explanation framework: a meta-path attribution module estimates the contribution of each semantic path to the prediction, and a subgraph extraction module visualizes representative local structures aligned with the most influential paths. Experiments on a large-scale inventor–patent dataset demonstrate that DHGNN not only outperforms strong baselines but also generalizes well to cold-start scenarios, such as predicting collaborations between previously unconnected inventors.

Keywords Scientific collaboration prediction · Link prediction · Dynamic heterogeneous graphs · Dynamic graph neural networks · Meta-path attention · Inventor collaboration networks

1 Introduction

Scientific collaboration, particularly among inventors, serves as a pivotal mechanism for innovation, facilitating interdisciplinary knowledge exchange and technology advancement (Sun et al. 2015). As scientific and industrial ecosystems become increasingly complex and interconnected, forecasting such collaborations has gained increasing strategic attention from both academia and industry. This urgency is underscored by the rapid growth in global patent activity, with annual patent filings increasing from 1.98 million in 2010 to 3.55 million in 2023 (World Intellectual Property 2024). Accurate collaboration prediction supports strategic research planning, early identification of promising partnerships, and the discovery of emerging interdisciplinary linkages (Cho and Yu 2018). Yet, such prediction remains challenging due to the temporal dynamics, structural heterogeneity, and semantic complexity inherent in inventor collaboration networks.

Existing approaches encounter difficulties in effectively modeling the temporal and semantic complexities underlying inventor collaboration prediction (Fritz et al. 2023; Khoshraftar and An 2024). Specifically, static heterogeneous graph neural networks (GNNs) successfully represent multi-relational semantics but overlook temporal evolution, neglecting crucial changes in collaboration likelihood over time (Schlichtkrull et al. 2018; Pareja et al. 2020; Barros et al. 2021). Conversely, dynamic GNNs effectively model evolving relationships but generally assume homogeneous graph structures, failing to leverage the rich semantic signals available from diverse entity interactions (Hansheng et al. 2021; Xiao et al. 2024). Moreover, most current dynamic approaches discretize temporal data into coarse intervals or rely on sequential models, limiting their capacity to capture

✉ Jie Wu
2120172501@alu.nufe.edu.cn

[1] School of Economics and Management, Jiangsu University of Science and Technology, Zhenjiang 212100, China

[2] Department of Industrial & Systems Engineering, The Hong Kong Polytechnic University, Hong Kong 999077, China

nuanced continuous-time dynamics critical in real-world inventor networks (Xiong et al. 2025).

Additionally, current methods lack interpretability, hindering their practical application in scenarios where decision transparency is paramount. Such methods often inadequately leverage semantic information, like shared institutional affiliations or technical expertise, despite their potential to inform predictions in cases with sparse historical data.

Building on these observations, our study addresses the practical challenge of accurately forecasting inventor collaborations—a critical task that requires precise integration of temporal dynamics, semantic heterogeneity, and structural complexities inherent to innovation networks. While existing dynamic graph neural networks have demonstrated effectiveness in general settings, they typically lack specialized designs tailored for capturing nuanced multi-relational patterns characteristic of inventor–patent ecosystems. To address this critical gap, our proposed DHGNN not only offers a customized architecture that enhances predictive accuracy within this particular domain but also contributes methodological insights with broader applicability.

Unlike prior methods using discrete snapshots, DHGNN incorporates continuous-time positional encoding coupled with causal masking, explicitly designed to capture fine-grained temporal reasoning beyond conventional discrete-time or sequential models. Moreover, DHGNN explicitly leverages meta-path semantics—such as shared affiliations or technical domains—to inform predictions in both cold-start and structurally rich scenarios. For instance, semantic paths can reveal indirect collaboration signals, such as inventors sharing an institutional affiliation yet lacking previous direct co-inventions.

Specifically, our contributions are summarized as follows:

- **Unified Dynamic-Heterogeneous Modeling.** DHGNN introduces a comprehensive framework that jointly models the temporal evolution and semantic heterogeneity of inventor collaboration networks. It constructs cumulative heterogeneous snapshots to preserve temporal continuity, and integrates both relation-aware structural encoding and meta-path–guided aggregation to capture multi-relational semantics. This design enables accurate and context-aware collaboration prediction in dynamic settings.
- **Temporal–Semantic Attention Encoding.** To jointly capture structural semantics and temporal evolution in dynamic heterogeneous graphs, DHGNN integrates two levels of attention. At the structural level, we combine relation-aware message passing with meta-path–based multi-hop aggregation to encode both local interacons and long-range semantic patterns. At the temporal level, we inject continuous-time positional embeddings into historical representations and apply a causally masked attention mechanism to aggregate past collaboration signals. This unified design enables the model to learn fine-grained, time-aware semantic representations that adaptively reflect when and how inventor relations evolve.
- **Path-Level Interpretability with Subgraph Grounding.** DHGNN incorporates a two-tier interpretability framework that delivers transparent, post-hoc explanations for predicted collaborations. First, it computes path-level attribution weights by projecting final inventor representations through meta-path–specific bilinear decoders, revealing which semantic paths most influenced each prediction. Second, it extracts subgraphs aligned with the most influential meta-paths, providing a structural basis for interpreting the predicted link. This dual mechanism enables interpretable reasoning across diverse collaboration contexts—capturing semantic affiliation in cold-start cases and structural bridging in mature networks.
- **Cold-Start Robustness and Structure–Semantics Adaptivity.** DHGNN is designed to robustly predict collaborations under varying structural conditions, including cold-start scenarios and densely connected subgraphs. Through a hierarchical temporal–semantic attention mechanism, the model adaptively modulates its reliance on different relational signals according to the availability and reliability of contextual evidence. In sparse settings, semantic paths—such as shared affiliations or technical domains—are emphasized to compensate for missing structural cues. Conversely, when strong topological patterns exist, the model leverages direct and multi-hop structural paths to reinforce prediction confidence. This dynamic attention reallocation enables DHGNN to generalize effectively to unseen nodes while maintaining high accuracy in historically dense regions, offering a unified and context-aware framework for scientific collaboration prediction.

2 Related work

2.1 Dynamic heterogeneous graph representation learning

Real-world networks often exhibit both structural heterogeneity and temporal dynamics, with diverse node and relation types evolving over time. Jointly capturing these two aspects poses significant challenges due to their inherent complexity and interaction.

Existing dynamic graph models primarily focus on temporal evolution. For instance, EvolveGCN (Pareja et al. 2020), DynGEM (Goyal et al. 2018), and GCN-LSTM (Jinyin Chen et al. 2022) update node embeddings or network parameters recurrently over discrete snapshots.

Meanwhile, attention-based methods such as DySAT (Sankar et al. 2020) introduced dual attention along structural and temporal axes, while TGAT (Xu et al. 2020) proposed continuous-time encoding for inductive scenarios. Memory-based approaches such as TGN (Rossi et al. 2020) track historical node states, significantly enhancing temporal reasoning capabilities.

Conversely, heterogeneous graph neural networks (HGNNs) emphasize modeling semantic complexity through multi-type nodes and relations (Pengtao Wang et al. 2024). Classical methods such as metapath2vec (Dong et al. 2017) leverage meta-path-guided random walks to preserve semantic proximity. Advanced neural approaches, including R-GCN (Schlichtkrull et al. 2018) and HAN (Xiao Wang et al. 2019b), MAGNN (Yujie Fan et al. 2022) and HGT (Hu et al. 2020), use relation-specific transformations or hierarchical attention mechanisms for fine-grained semantic modeling. Nevertheless, these methods typically treat graphs as static, failing to represent dynamic changes in structures or relationships over time (Lei Chen et al. 2023).

Recently, several integrated models have emerged, attempting to unify temporal dynamics and semantic heterogeneity. Luo et al. (2020b) combined real-time relation-aware message passing with temporal embeddings. Xue et al. (2021) introduced hierarchical attention mechanisms that span both relation types and temporal intervals. Fan et al. (2021; 2022) proposed HTGNN, capturing type-aware and time-aware patterns simultaneously, and Zhang et al. (2023) developed DHGAS, which automatically searches for optimal dynamic message-passing strategies. Yet, these approaches still encounter limitations regarding interpretability and scalability, particularly crucial in practical, large-scale networks.

In summary, despite significant advances, integrating temporal and heterogeneous modeling remains challenging. Our proposed DHGNN aims addresses these issues by combining relation-aware propagation, meta-path–guided semantic aggregation, and continuous-time attention. This approach enables capturing dynamic, multi-semantic patterns necessary for accurate collaboration prediction.

2.2 Collaboration prediction on graphs

Collaboration prediction, essentially a graph-based link prediction task, has been extensively explored, particularly in social, academic, and innovation contexts (Ren 2024). Early methods predominantly relied on structural proximity metrics, such as common neighbors and preferential attachment, to infer potential links (Liben-Nowell and Kleinberg 2003). Pavlov and Ichise (2007) and Yu et al. (2014) extended these ideas using supervised learning techniques based on engineered network features. These initial approaches lacked flexibility, as they did not account for semantic or temporal variations in collaboration behaviors.

Recognizing the complexity of real-world collaborations, subsequent studies integrated multi-relational information. Wang and Sukthankar (2013) and Lopes et al. (2010) proposed multi-relational frameworks leveraging network topology and domain-specific metadata to predict collaborations. However, these methods typically did not model temporal dynamics explicitly.

The emergence of graph representation learning significantly advanced the field. DeepWalk (Perozzi et al. 2014) and node2vec (Grover and Leskovec 2016) enabled unsupervised node embeddings through random walks, but lacked the capacity to model complex semantic contexts or temporal evolution. GNN-based methods, such as SEAL (Muhan Zhang and Chen 2018), further improved accuracy by learning predictive structural motifs within enclosing subgraphs. In recommendation contexts, Graph Neural Collaborative Filtering (Xiang Wang et al. 2019a) and hybrid models combining GraphSAGE and TGN Zhu and Yaseen (2022) demonstrated temporal-aware structural modeling advantages. Yet, these models typically did not fully leverage high-order semantic interactions, critical for heterogeneous collaboration contexts.

Inventor collaboration prediction inherently involves complex semantic interactions, such as shared affiliations, technical domains, and indirect co-inventorship relationships evolving over time. Recent studies have thus emphasized meta-path-based semantics. Lathabai et al. (2022) demonstrated semantic paths as robust predictors of collaboration, motivating a stronger integration of semantic and temporal dynamics. Our DHGNN aims to addresses these distinct challenges by unifying temporal reasoning and multi-semantic path modeling, enabling context-sensitive predictions for dynamic inventor networks.

2.3 Explainability in graph neural network models

Despite their predictive power, graph neural networks (GNNs) are inherently opaque due to complex multi-layer structures and latent embeddings. This opacity limits trust and actionable insights in critical tasks such as collaboration forecasting, necessitating transparent model interpretations.

Post-hoc interpretability methods, such as GNNExplainer (Rex Ying et al. 2019a), identify minimal predictive subgraphs and critical features contributing to predictions. PGExplainer (Dongsheng Luo et al. 2020a) and GraphLIME (Huang et al. 2023), further enhance efficiency and generalization by learning parameterized explainers or local interpretable surrogates, respectively. These methods provide essential post-hoc transparency but do not inherently guide the model toward interpretable reasoning.

Alternatively, intrinsic interpretability approaches, notably attention-based methods such as HAN (Xiao Wang et al. 2019b), provide built-in attribution mechanisms. However, recent research indicates that attention weights alone may not reliably reflect true model decision rationales (Rex Ying et al. 2019a; Dongsheng Luo et al. 2020a).

To overcome these limitations, our DHGNN employs a hybrid interpretability framework, combining post-hoc identification of influential semantic paths and structured extraction of explanatory subgraphs. Specifically, our model first attributes predictions to interpretable semantic paths (meta-path attribution) and then selects a concise, structurally meaningful subgraph for visualizing and understanding predictions (subgraph extraction). This integrated strategy aims to mitigates the interpretability shortcomings associated with pure attention mechanisms, providing robust, comprehensible explanations aligned with both structural and semantic patterns.

3 Methodology

3.1 Framework overview

We propose DHGNN, a unified framework for forecasting inventor collaborations in dynamic and heterogeneous innovation networks. As illustrated in Fig. 1, the model integrates five components that collectively capture evolving relational semantics and support interpretable link prediction.

As illustrated in Fig. 1, DHGNN executes a five-stage pipeline. ① Build cumulative heterogeneous snapshots and initialize type-specific node features. ② Encode each snapshot with a two-stage heterogeneous encoder—relation-aware message passing followed by meta-path aggregation. ③ Temporally align snapshot embeddings via causally masked attention and sinusoidal positional encodings. ④ Predict future inventor collaborations with a trainable bilinear decoder. ⑤ Quantify meta-path contributions and extract explanatory sub-graphs, yielding transparent, instance-level explanations.

The framework begins by organizing historical patent records into a sequence of cumulative heterogeneous graph snapshots, comprising four node types and three timestamped relation types. Node features are initialized in a type-specific manner. Patent nodes are encoded using semantic embeddings derived from titles and abstracts. Inventor and assignee nodes, which lack structured attributes, are initialized with trainable vectors. Technical field nodes are one-hot encoded based on International Patent Classification (IPC) codes and transformed into dense embeddings. All node features are projected into a unified latent space to support joint representation learning.

Each snapshot is encoded by a heterogeneous GNN that incorporates both relation-aware message passing and meta-path-guided semantic aggregation. To model the evolution of inventor representations over time, a temporal attention module aggregates historical embeddings through a causally masked attention mechanism, allowing the model to learn fine-grained temporal dependencies.

Finally, a bilinear decoder estimates the likelihood of future collaborations between inventor pairs. To enhance transparency, DHGNN includes a two-tier interpretability mechanism that performs path-level attribution and extracts subgraphs aligned with influential semantic paths to identify critical mediating entities.

3.2 Graph construction and initialization

Step ① To support temporal reasoning and semantic heterogeneity, DHGNN organizes historical patent data into a sequence of yearly heterogeneous graph snapshots. Each snapshot $\mathcal{G}^{(t)} = (\mathcal{V}^{(t)}, \mathcal{E}^{(t)})$ captures the inventor–patent landscape up to year t, where $V^{(t)}$ is the set of active nodes and $E^{(t)}$ denotes the set of directed, typed edges whose timestamps satisfy $\leq t$.

As shown in Fig. 2, each graph snapshot includes four node types—Inventors (V_{inv}), Patents (V_{pat}), Assignees (V_{assg}), and Technical Fields (V_{tech})—and three edge types representing authorship (I–P), ownership (P–A), and classification (P–F). These heterogeneous components evolve

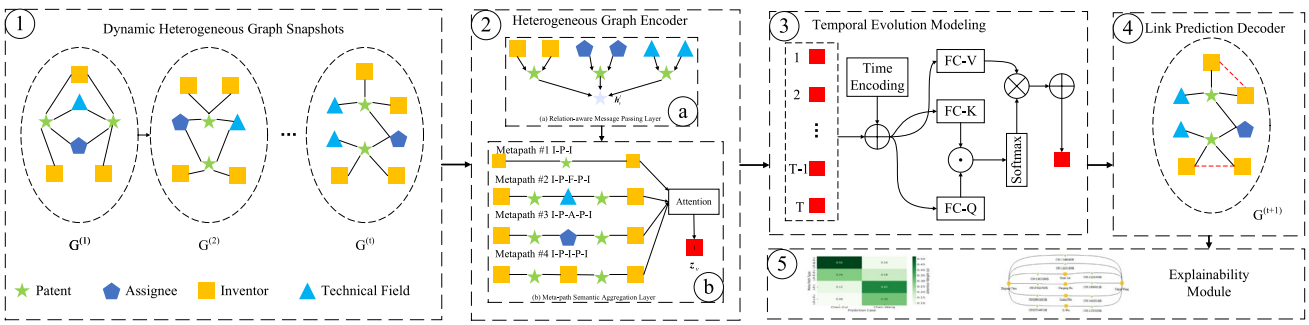

Fig. 1 DHGNN Framework Overview

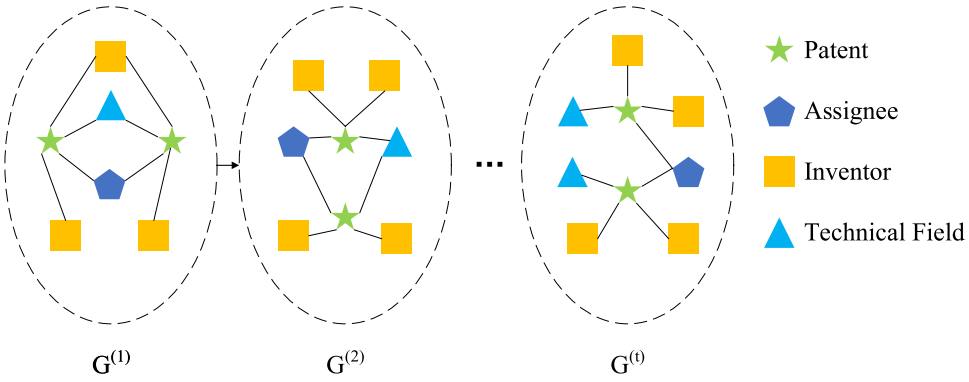

Fig. 2 Structure of the Dynamic Heterogeneous Graph

cumulatively across years, enabling the model to exploit historical collaboration and technological context.

Each edge $(u, v) \in \mathrm{E}^{(t)}$ is associated with a timestamp indicating the grant year of the corresponding patent. Only edges with timestamps $\leq t$ are included in $\mathrm{E}^{(t)}$, ensuring that each graph snapshot faithfully reflects the cumulative state of the innovation network at time t.

Node Feature Initialization. To handle heterogeneous nodes with diverse attributes, we adopt a type-aware initialization scheme combining semantic embeddings, learnable vectors, and categorical encodings. For each node $v \in \mathrm{V}^{(t)}$, we assign a raw feature vector $\mathbf{x}_v \in \mathrm{R}^{d'}$ based on its type $\phi(v)$.

Patent nodes use pretrained semantic embeddings from PatentSBERTa (Bekamiri et al. 2024), which captures rich contextual signals from titles and abstracts. Inventor and assignee nodes, which typically lack structured attributes, are initialized with randomly sampled, trainable embeddings. Technical field nodes (based on IPC codes) are represented by one-hot vectors that are mapped into a shared embedding space using a learnable linear projection.

To unify these heterogeneous features into a common latent space R^d, we apply a trainable type-specific projection matrix $W_{\phi(v)} \in \mathbb{R}^{d \times d'}$ to obtain projected features:

$$h_v^{(t,0)} = W_{\phi(v)} \cdot x_v \tag{1}$$

We denote the set of initial raw feature vectors at time t as $\{x_v \mid v \in \mathcal{V}^{(t)}\}$.All projected features $h_v^{(t,0)} \in \mathbb{R}^d$ are then used as input to the downstream heterogeneous graph encoder.

3.3 Heterogeneous graph encoder

Step ②. To effectively model both fine-grained local interactions and high-order semantic dependencies in dynamic heterogeneous graphs, we design a hierarchical encoder composed of two complementary layers: (a) relation-aware message passing for local structural encoding and (b) meta-path-based multi-hop aggregation for capturing semantic-level patterns across indirect collaborations. This encoder operates on each snapshot independently and produces node representations that are fed into the temporal modeling module in Step ③, thereby preserving both structural and semantic information before temporal dynamics are introduced.

(a) Relation-Aware Message Passing. We begin by encoding local structural patterns using a relation-aware message passing module. As illustrated in Fig. 3, each node aggregates information from its neighbors under different relation types.

Specifically, given a dynamic graph snapshot $\mathrm{G}^{(t)}$, the message from a neighbor $u \in \mathcal{N}_r(v)$ to node v is computed as $m_{u \to v}^{t,l,r} = W_r^{(l)} h_u^{t,l-1}$, where $\mathbf{W}_r^{(l)}$ is a learnable projection for relation r at layer l.

To reflect the relative importance of each neighbor, a relation-aware attention mechanism is applied:

$$e_{uv}^{t,l,r} = \mathrm{LeakyReLU}\left((a_r^{(l)})^\top [W_r^{(l)} h_u^{t,l-1} || W_r^{(l)} h_v^{t,l-1}] \right) \tag{2}$$

$$\alpha_{uv}^{t,l,r} = softmax_{u \in \mathcal{N}_r}(e_{uv}^{t,l,r}) \tag{3}$$

Messages are aggregated within each relation type as $m_v^{t,l,r} = \sum_{u \in N_r(v)} \alpha_{uv}^{t,l,r} m_{u \to v}^{t,l,r}$, and fused with a residual connection to yield the updated representation:

$$h_v^{t,l} = \sigma\left(\sum_{r \in \mathcal{R}} m_v^{t,l,r} + W_0^{(l)} h_v^{t,l-1} \right) \tag{4}$$

Here, $\sigma(\cdot)$ is the ReLU activation and $\mathbf{W}_0^{(l)}$ denotes the self-loop projection. The final output $h_v^t = h_v^{t,L}$ serves as input to the semantic aggregation layer.

(b) Meta-Path-Based Aggregation. While relation-aware aggregation captures direct structural signals, indirect and high-order collaborations—such as co-inventorship via shared assignees or technical domains—are often hidden in longer semantic paths. To enhance the model's capacity to represent such semantics, we introduce a meta-path-guided multi-hop aggregation module.

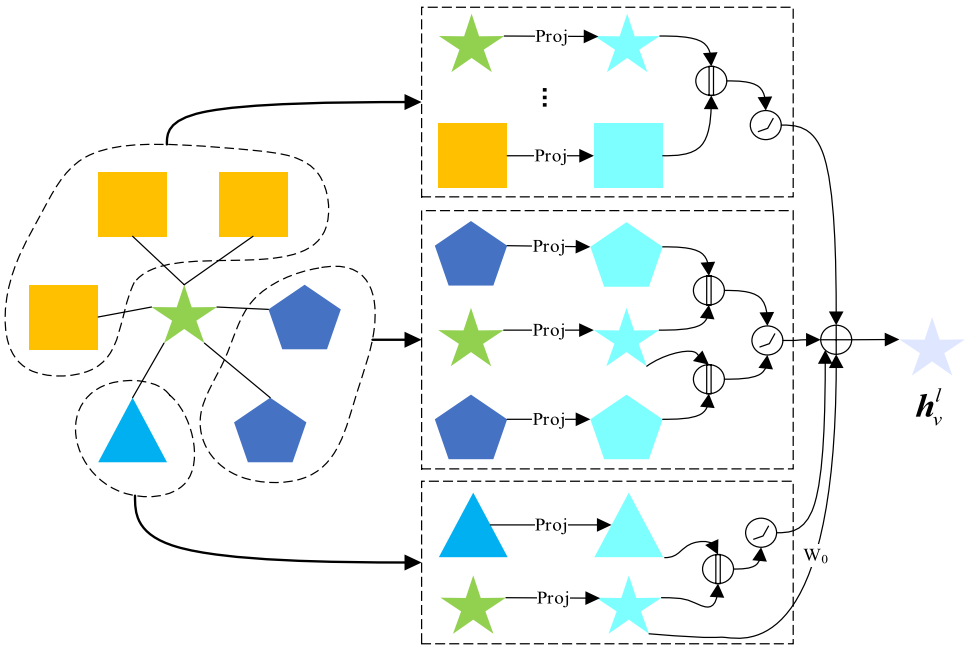

Fig. 3 Architecture of Relation-Aware Message Passing

Each meta-path instance defines a specific semantic route between two inventors via intermediate nodes and relations. We define a set of meta-paths $\mathcal{P} = \{P_1, P_2, P_3, P_4\}$, each encoding a distinct collaboration cue:

- P_1 = Inventor → Patent → Inventor (direct co-invention);
- P_2 = Inventor → Patent → Assignee → Patent → Inventor (shared affiliation);
- P_3 = Inventor → Patent → Technical Field → Patent → Inventor (shared domain expertise);
- P_4 = Inventor → Patent → Inventor → Patent → Inventor (indirect structural bridging via intermediate co-authors).

We select these domain-informed meta-paths to model key mechanisms that influence the formation of inventor collaborations and support partner prediction. Specifically, the I–P–I path captures direct co-invention and serves as a strong indicator of recurring collaborations. The I–P–A–P–I path reflects organizational proximity through shared assignees, consistent with empirical evidence that institutional ties and geographic co-location jointly facilitate collaboration (Crescenzi et al. 2016; Diemer and Regan 2022). I–P–F–P–I reflects topic similarity based on shared IPC technical fields, allowing the model to infer thematic overlap between inventors (Bar and Leiponen 2012; Jaffe and de Rassenfosse 2017). Finally, the I–P–I–P–I path represents structural bridging through mutual collaborators, thereby effectively capturing opportunities for indirect collaboration (Carayol et al. 2019).

These patterns reflect structural and semantic connections in the inventor–patent network, such as institutional ties, topic similarity, or redundant bridge-based pathways that capture indirect yet plausible co-invention routes.

As shown in Fig. 4, each path instance enables feature propagation across multi-hop semantic routes.

The meta-path module operates in two intuitive steps. First, for each meta-path, we aggregate the features of reachable neighbors into a compact summary vector. Second, we compute an importance weight for each path and combine the summaries into a unified embedding. This two-stage design—within-path aggregation followed by across-path weighting—allows the model to identify and prioritize the most relevant semantic cues for link prediction.

Specifically, given a path $P = (r_1, r_2, ..., r_L)$, the path-specific embedding for node v is initialized as $h_v^{t,(0),P} = h_v^t$, and updated recursively:

$$h_v^{t,(k),P} = \sum_{u \in \mathcal{N}_{r_k}(v)} \alpha_{uv}^{t,r_k} W_{r_k} h_u^{t,(k-1),P}, \quad k = 1, \ldots, L \tag{5}$$

The attention scores α_{uv}^{t,r_k} and transformation weights W_{r_k} are reused from the relation-aware encoder to ensure parameter consistency.

To adaptively fuse path-specific semantics, we apply a semantic-level attention mechanism. For each path P, its contribution is scored via:

$$s_v^{t,P} = q^\top \tanh(W_P h_v^{t,(L),P} + b_P) \tag{6}$$

$$\beta_v^{t,P} = \frac{\exp(s_v^{t,P})}{\sum_{P\prime \in \mathcal{P}} \exp(s_v^{t,P\prime})} \tag{7}$$

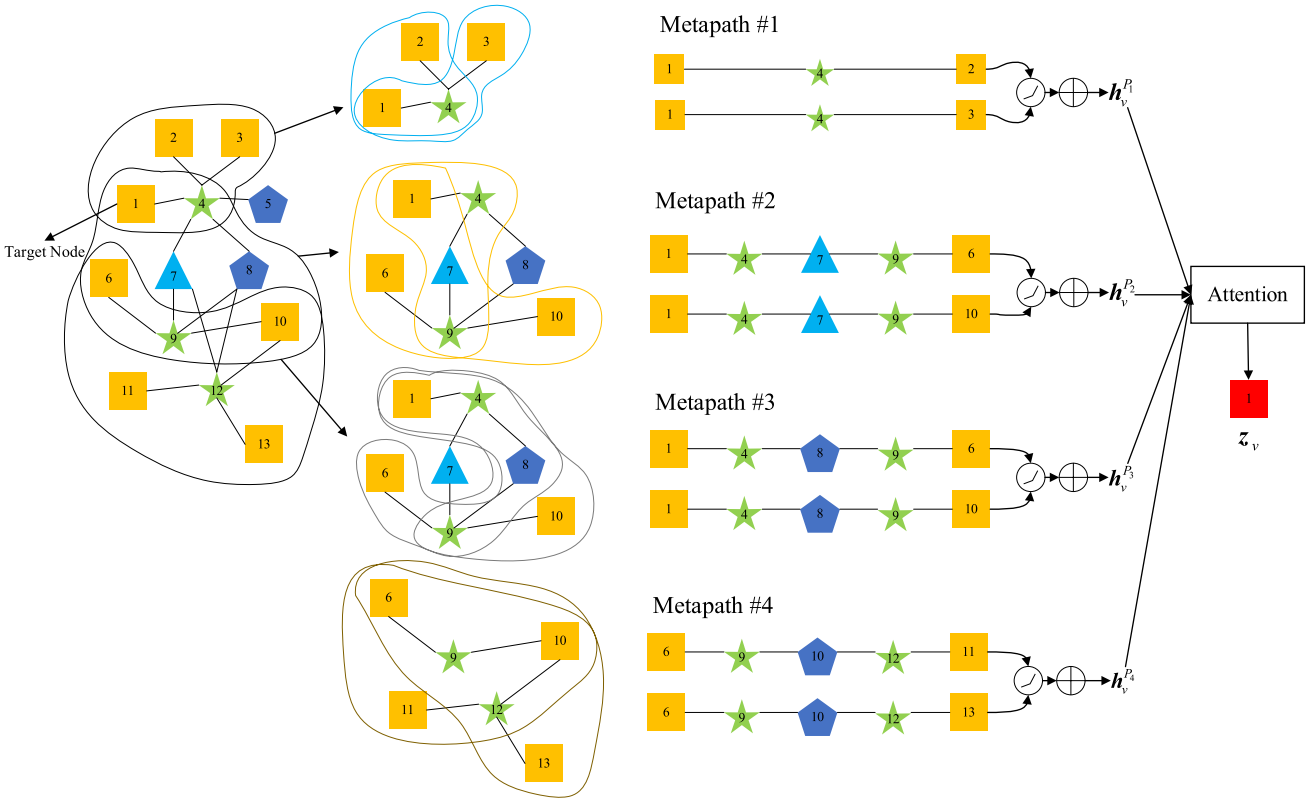

Fig. 4 Meta-Path Aggregation with Instance- and Semantic-Level Attention

The final node embedding aggregates information across all meta-paths:

$$z_v^t = \sum_{P \in \mathcal{P}} \beta_v^{t,P} h_v^{t,(L),P} \tag{8}$$

The proposed encoder jointly captures both local heterogeneous interactions and high-order semantic dependencies via a hierarchical representation learning process. By integrating the outputs of relation-aware message passing and meta-path-level attention, it enables flexible, interpretable, and context-sensitive modeling of potential inventor collaborations in dynamic patent graphs.

3.4 Modeling temporal evolution

Step ③. In dynamic heterogeneous collaboration networks, the structural environment around an inventor continuously evolves—driven by emerging patents, institutional transitions, and shifting technological domains. To accurately model such temporal dynamics, we employ an attention-based aggregation framework that tracks the evolution of node representations across historical timestamps.

Our temporal module adopts a hybrid scheme: we segment the dynamic heterogeneous graph into cumulative yearly snapshots and apply causally masked self-attention with continuous positional encodings over the historical embeddings. This design is appropriate for the patent collaboration prediction setting for two main reasons. First, the delay between patent application and grant usually spans one to two years and varies significantly across technologies and assignees. Under these conditions, using finer temporal resolutions such as quarterly snapshots may impose spurious chronological orders and amplify temporal noise. Second, finer temporal segmentation considerably increases the number of snapshots and the overall computational burden, which offers limited practical benefits given the sparse and delayed nature of patent collaboration events. In contrast, yearly snapshots provide a balanced temporal abstraction that aligns with the natural pace of patent-related activities while maintaining both model robustness and computational efficiency.

Let $z_v^{t\prime} \in \mathrm{R}^d$ denote the heterogeneous structural embedding of inventor v at time $t\prime$, as generated by the encoder in §3.3. We define the historical sequence up to timestamp t as:

$$\mathcal{Z}_v^{(t)} = \left\{ z_v^{t\prime} | 1 \le t\prime \le t \right\} \tag{9}$$

To inject temporal context into each representation, we define a temporal encoding function. We add the

timestamp-specific positional embedding $\gamma_{t\prime}$ to $\mathbf{z}_v^{t\prime}$ so that temporal cues are injected in a parameter-free, dimension-consistent manner. This choice follows the practice of prior temporal GNN models such as HTGNN (Yujie Fan et al. 2022), which adopt similar additive strategies and report favorable performance.

$$l(z_v^{t\prime}) = z_v^{t\prime} + \gamma_{t\prime} \tag{10}$$

where $\gamma_{t\prime} \in \mathrm{R}^d$ is a timestamp-specific positional embedding based on sine and cosine signals:

$$\gamma_{t\prime}^{(2k)} = \sin\left(\frac{t\prime}{10000^{2k/d}}\right) \tag{11}$$

$$\gamma_{t\prime}^{(2k+1)} = \cos\left(\frac{t\prime}{10000^{2k/d}}\right) \tag{12}$$

where k is the index of each element.

To ensure causality and properly model evolving semantics, we apply a scaled dot-product attention mechanism over past time steps only, this attention setup is similar to the mechanism used in Transformer models (Vaswani et al. 2017). The query vector is obtained from the current (most recent) representation via

$$q_v^{(t)} = W_q \cdot z_v^{(t)} \tag{13}$$

And each temporally encoded historical vector is projected into key and value vectors:

$$k_v^{(t\prime)} = W_k \cdot l(z_v^{t\prime}) \tag{14}$$

$$v_v^{(t\prime)} = W_v \cdot l(z_v^{t\prime}) \tag{15}$$

where $W_q, W_k, W_v \in \mathbb{R}^{d\times d}$ are learnable matrices.

The attention score between the current step t and a past step $t\prime \le t$ is computed as:

$$e_v^{(t,t\prime)} = \frac{q_v^{(t)} \cdot k_v^{(t\prime)}}{\sqrt{d}} \tag{16}$$

$$\alpha_v^{(t,t\prime)} = \frac{\exp(e_v^{(t,t\prime)})}{\sum_{\tau=1}^{t}\exp(e_v^{(t,\tau)})} \tag{17}$$

We aggregate the temporal context using a weighted sum over all past values:

$$h_v^{(t)} = \sum_{t\prime=1}^{t} \alpha_v^{(t,t\prime)} \cdot v_v^{(t\prime)} \tag{18}$$

To stabilize learning, we apply residual connection and layer normalization:

$$h_v^{(t)} \leftarrow \mathrm{LayerNorm}(z_v^{(t)} + h_v^{(t)}) \tag{19}$$

The resulting temporally-aware embedding $h_v^{(T)}$ at the last timestamp T is passed to the link prediction decoder described in §3.5.

3.5 Link prediction and learning objective

Step ④. After computing the dynamic temporal embeddings $h_v^{(T)}$ for all inventor nodes at the final observation year T, we formulate the prediction task as a binary link prediction problem. Specifically, the model aims to estimate whether a collaboration link will form between any inventor pair (u, v) in the next time step $T + 1$.

Link Prediction Decoder. Given two inventor nodes u and v with temporal embeddings $\mathbf{h}_u^{(T)}$ and $\mathbf{h}_v^{(T)}$ produced by the temporal evolution module, we compute a link prediction score s_{uv} using a parameterized bilinear decoder (Yue Wang et al. 2019c). We adopt the bilinear form as it introduces trainable relation weights, offering a balance between model flexibility and parameter efficiency. This structure is suitable for heterogeneous collaboration graphs with potentially asymmetric interactions.

$$s_{uv} = \left(h_u^{(T)}\right)^T W h_v^{(T)} \tag{20}$$

This bilinear form models the interaction between latent representations with a learnable matrix $W \in \mathbb{R}^{d\times d}$, allowing the model to flexibly capture latent relation patterns. A higher score indicates a stronger semantic alignment between the inventor pair.

The prediction probability is transformed into a probability using a logistic sigmoid function:

$$\hat{y}_{uv} = \sigma(s_{uv}) = \frac{1}{1+\exp(-s_{uv})} \tag{21}$$

where $\hat{y}_{uv} \in (0,1)$ denotes the predicted probability that inventors u and v will co-author a patent at time $T + 1$.

Supervision Strategy and Sampling. To train the model, we construct a labeled dataset of inventor pairs from historical data. Positive samples are inventor pairs (u, v) that collaborated during the prediction window $T + 1$. Negative samples are inventor pairs that did not collaborate at time $T + 1$.

Given the extreme class imbalance, we adopt a negative sampling strategy (Mikolov et al. 2013). For each positive pair, we randomly sample K negative pairs from the candidate space. This ensures the negatives are balanced and comparable, improving training stability.

Loss Function Design. The model is trained by minimizing binary cross-entropy (BCE) over all labeled examples. For an inventor pair with ground truth label $y_{uv} \in \{0,1\}$ and predicted probability $\hat{y}_{uv}$, the loss is:

$$\mathcal{L}_{uv} = -(y_{uv}\log(\hat{y}_{uv}) + (1 - y_{uv})\log(1 - \hat{y}_{uv})) \tag{22}$$

The total loss $\mathcal{L}$ across the mini-batch $\mathcal{B}$ of inventor pairs is:

$$\mathcal{L} = \frac{1}{|\mathcal{B}|} \sum_{(u,v)\in\mathcal{B}} \mathcal{L}_{uv} \tag{23}$$

To prevent overfitting caused by link sparsity, we additionally apply dropout in graph layers and weight-decay regularization.

3.6 Model interpretability

Step ⑤. In scientific collaboration forecasting, interpretability is essential for fostering user trust and promoting practical adoption. Beyond achieving accurate predictions, models must clearly articulate the underlying rationale behind their forecasts. To this end, we propose a two-tier explainability framework comprising (1) post-hoc meta-path-level attribution and (2) subgraph-based local explanation, providing transparent and actionable insights into the model's decision-making processes. The entire interpretability process is strictly post hoc and does not interfere with model training. Meta-path attribution is conducted on frozen temporal embeddings without backpropagation, and subgraph extraction is a deterministic, non-parametric procedure. No gradients are computed, and no model parameters are updated during explanation.

Meta-Path Attribution. Although our encoder employs meta-path attention during the semantic aggregation stage (§3.3), these attention weights primarily reflect local semantic preferences during node representation learning, rather than directly corresponding to the final prediction outcomes. To accurately quantify the actual contribution of each semantic path to prediction results, we introduce a post-hoc meta-path attribution module, which is applied after temporal embedding fusion without requiring modifications to the existing encoder structure.

Specifically, after obtaining the final inventor representations h_u and h_v via temporal fusion (§3.4), we define a learnable path-specific projection matrix W_p for each predefined meta-path p. Using these matrices, we calculate the contribution score of each semantic path as follows:

$$score_p^{(u,v)} = h_u^\top W_p h_v \tag{24}$$

Next, these scores are normalized using the softmax function to yield the meta-path-level attribution weights:

$$\beta_p^{(u,v)} = \frac{\exp(score_p^{(u,v)})}{\sum_{p'\in P}\exp(score_{p'}^{(u,v)})} \tag{25}$$

Here, $\beta_p^{(u,v)}$ explicitly quantifies the relative contribution of meta-path p to the predicted collaboration likelihood between inventors u and v. This attribution mechanism transparently decomposes the final prediction score into semantic contributions (e.g., shared affiliations, similar technical domains, or structural bridging), significantly enhancing the interpretability of the prediction outcomes.

Subgraph-Level Explanation. To further ground the predictions in explicit graph-based evidence, we extract a local subgraph from the meta-path identified as having the highest attribution weight:

$$p^* = \operatorname{argmax}_{p\in\mathcal{P}} \beta_p^{(u,v)} \tag{26}$$

Under the selected influential meta-path p^*, we rank neighbor nodes by their node-level attention $\alpha_{ij}^{(p^*)}$ and select the smallest set whose cumulative attention mass reaches a coverage target η; in case of extreme sparsity we keep top-1 to guarantee a non-empty explanation, and we cap the subgraph by a size budget K to maintain readability. The hyperparameters η and K are chosen on the validation set (Rex Ying et al. 2019a; Dongsheng Luo et al. 2020a; Funke et al. 2023).

3.7 Algorithm of DHGNN

In this subsection we present the training and interpretation workflow of DHGNN through a concise pseudocode narrative. Algorithm 1 is organized into four functional blocks. Steps 1–3 initialize type-specific node features and the full set of trainable parameters. Steps 4–13 constitute the dynamic-snapshot encoder: a relation-aware message-passing layer followed by meta-path semantic aggregation is executed for every yearly snapshot (Steps 5–9), after which the resulting per-snapshot embeddings are fed into a causal temporal-attention module that slides over the latest L years (Steps 10–13). Steps 14–22 implement link-prediction supervision—binary cross-entropy is accumulated for each positive–negative pair drawn with ratio K_{neg} (Steps 14–19), averaged over the mini-batch (Step 20), and the parameters are optimized with Adam under L2 weight decay (Step 21). Finally, Steps 23–31 describe the optional post-hoc explanation phase: given a query inventor pair, the frozen encoder produces their temporal embeddings (Step 24); path-specific scores are computed and normalized to obtain meta-path attribution weights (Steps 25–27); the most influential path guides a subgraph extraction procedure that selects the smallest neighborhood reaching coverage η and capped by K (Step 28). Step 31 returns the trained parameters together with the attribution weights and the extracted explanatory subgraph. After several epochs (or early-stopping rounds), DHGNN converges and can both predict future collaborations and provide interpretable meta-path-level explanations.

Input: Dynamic snapshots $\{G(1)\ldots G(T)\}$, meta- path set $P=\{P_1\ldots P_m\}$, labelled inventor pairs $D_{train}=\{(u,v,y)\}$, history length L, Epochs E, Negatives-per-positive K_{neg}, query pair (u^*,v^*) for interpretation

Output: Trained parameters Θ, Meta-path attribution weights $\beta_p(u^*,v^*)$, Explanation subgraph G_{expl}

// INITIALISATION

1: **for** each node $v\in G(1)$:

2: $h_v=W_{\phi(v)}x_v$ // Type-specific linear projection Eq. (1)

3: Initialize parameters $\Theta=\left\{\{W_r^{(l)}\},\{W_0^{(l)}\},a,W_L,q,b,W_q,W_k,W_v,W,\{W_p\},\{W_{\phi(\cdot)}\}\right\}$

// TRAINING PHASE

4: **for** epoch = 1 … E:

5: // snapshot encoding

6: **for** t = 1 … T:

7: $h_v^t=RelMsgPass(G(t),h)$ // Eq. (2- 4)

8: $z_v^t=MetaPathAgg(h,P)$ // Eq. (5- 8)

9: **end for**

10: **// temporal evolution**

11: **for** inventor v:

12: $h_v^T=CausalTemporalAttn\left(\{z_v^{T-L},\ldots,z_v^T\}\right)$ // Eq. (9- 19)

13: **end for**

14: **// link prediction & loss**

15: loss = 0

16: **for** $u,v,y\in SamplePairs(D_{\text{train}},K_{neg})$:

17: $\hat{y}=BilinearDecoder(h_u^T,h_v^T)$ // Eq. (20- 21)

18: $loss=loss+BCE(\hat{y},y)$ // Eq. (22)

19: **end for**

20: $loss=loss/|\mathrm{B}|$ // Eq. (23)

21: $\Theta=Adam(\Theta,\nabla loss)$

22: **end for**

// EXPLANATION PHASE

23: **if** (u^*,v^*) is specified:

24: $(h_{u^*}^T,h_{v^*}^T)=DHGNN.Encode(G(T-L),\ldots,G(T))$ // forward only

25: **for** $p\in P$: $score[p]=Path\,\mathrm{Proj}(h_{u^*},h_{v^*},P)$ // Eq. (24)

26: $\beta_p(u^*,v^*)=Softmax(score[p])$ // Eq. (25)

27: $p^*=argmax\beta_p(u^*,v^*)$ // Eq. (26)

28: $G_{expl}=ExtractSubgraph(p^*,\eta,K)$

29: **else:**

30: $\beta_p(u^*,v^*)=\varnothing$; $G_{expl}=\varnothing$

31: **return** Θ, $\beta_p(u^*,v^*)$, G_{expl}

4 Empirical evaluation

4.1 Dataset construction and preprocessing

We construct a large-scale dynamic heterogeneous graph dataset derived from real-world patent records in the artificial intelligence (AI) domain to facilitate systematic evaluation of our proposed framework. As a representative field of high innovation intensity, AI has undergone rapid structural transformations and increasing cross-domain collaboration, which are prominently reflected in national strategies such as China's "New Generation AI Development Plan." These characteristics—frequent team reconfigurations, evolving relational patterns, and semantic complexity—render AI patent data a suitable empirical ground for validating graph-based modeling approaches under temporal and structural heterogeneity.

In this context, we selected a large-scale dataset of AI patents from China as the primary experimental testbed. This dataset captures the high structural and temporal complexity inherent in AI innovation networks. Although geographically specific, it exhibits key structural characteristics that are broadly representative, providing a reliable basis for evaluating model performance and supporting the transferability and applicability of our framework to other national or regional contexts.

Patent retrieval and filtering. We collected the data from Incopat, a comprehensive commercial patent analytics platform with full access to structured inventor, assignee, classification, and legal status metadata. The retrieval was performed in May 2025 using the following structured query: (INDUSTRY2 = (1.5)) AND (PNC = CN AND PT = ("4")) AND (AP-COUNTRY = ("CN")) NOT ((INTT = ("不公告发明人")) OR (INTT = ("请求不公布姓名"))) NOT (NO-IN = (1)) AND (ADY = ("2015" TO "2024")).

This query targets invention patents (PT = 4) granted in China between 2015 and 2024 and classified in the AI sector under SEIC code "1.5." To ensure data integrity and analytical relevance, we retain only patents with at least two named inventors and exclude records where inventor identities are undisclosed. A total of 188,090 patent records were retrieved. After deduplication at the patent family level, we retained 187,660 unique invention patents for downstream analysis.

Inventor Name Disambiguation. We resolved inventor-name ambiguity through a four-stage, constraint-driven pipeline (Tran et al. 2020).

(1) Blocking. Inventor records were grouped by exact matches of normalised simplified-Chinese full names. Names were standardized with OpenCC and stripped of whitespace and punctuation. For high-frequency names (e.g., "张伟"), we introduced sub-blocking that concatenated the inventor name with a three-letter assignee code (e.g., "张伟–qhd"), thereby isolating same-name inventors across institutions and sharply reducing the comparison space while maintaining high within-block precision.

(2) Pairwise Similarity Scoring. Within each name–assignee block, pairwise similarity was estimated by a gradient-boosted Siamese classifier trained on 10 000 manually labelled inventor pairs. Four feature groups were concatenated: (i) the mean of 768-dimensional PatentSBERTa embeddings over all patent titles and abstracts associated with each inventor; (ii) one-hot encodings of assignee affiliations and top-level IPC classes; (iii) Jaccard similarities between the two inventors' co-inventor sets and between their co-assignee sets, capturing structural and organizational overlap; (iv) an exponential temporal kernel, $\exp(-\Delta t/\tau)$, reflecting recency-weighted collaboration proximity. The model returned a match probability for each pair, serving as a soft similarity score.

(3) Constrained Clustering. An approximate k-nearest-neighbor graph was built with FAISS, where edges were weighted by the match probabilities. Hard constraints were imposed: pairs with probability ≥ 0.90 were given must-link labels, whereas pairs ≤ 0.15 received cannot-link labels; intermediate edges remained unlabeled. A constraint-aware variant of HDBSCAN then performed density-based clustering, balancing precision with boundary flexibility. The 0.90/0.15 cut-offs and HDBSCAN parameters (min_samples = 5, min_cluster_size = 15) were selected through grid-search on a manually labelled validation set. The resulting clustering achieves a silhouette coefficient of 0.712, indicating strong intra-cluster cohesion and inter-cluster separation.

(4) Graph refinement. To enhance cluster coherence, we exploited the inventor–co-inventor bipartite graph. Modularity-guided label propagation allowed a node to adopt a neighbor's label only if this increased global modularity ($\Delta Q > 0.02$). Iteration ceased once no further gain was achievable, ensuring that the final clusters respected both pairwise similarity and the latent collaboration topology.

Evaluation against a manually curated benchmark of 1 500 inventor profiles yielded precision 0.948, recall 0.932, and F1-score 0.940. The 1500 inventor profiles were randomly sampled from our empirical dataset to ensure broad coverage of different institutions and time periods. Labeling was independently conducted by two doctoral students with expertise in Chinese patent data, and any disagreements were resolved through discussion. To mitigate long-tail sparsity and stabilize downstream modelling, we retained inventors with at least ten AI-related invention patents granted between 2015 and 2024 and removed patents authored solely by a single high-productivity inventor. The resulting graph comprises 7 653 prolific inventors and 35 537 unique patents.

4.2 DHG Construction and temporal settings

Graph Schema and Temporal Granularity. We construct a dynamic heterogeneous graph from the cleaned AI patent corpus, designed to support temporal graph neural network modeling. Inventor nodes denote disambiguated high-productivity individuals—those with at least ten granted AI invention patents during the observation period—resulting in 7,653 unique inventor entities. Patent nodes represent deduplicated invention records, restricted to patents co-authored by at least two high-productivity inventors, yielding 35,537 nodes. Assignee nodes are created by normalizing applicant institution names, enabling consistent aggregation across affiliated patents. Technical field nodes, corresponding to four-character IPC subclass codes (e.g., G06F, H04L), are used to denote the technological domains of patents.

We focus on the Chinese AI patent corpus due to its rich metadata, institutional diversity, and fast temporal evolution—characteristics that make it particularly suitable for evaluating dynamic heterogeneous models. While region-specific, our approach is generalizable to other collaboration networks with comparable multi-entity structure and timestamped relationships.

Typed edges capture key relational semantics within the inventor–patent network: (i) I → P edges denote inventorship links from inventors to their associated patents. Thus, a patent with multiple inventors yields multiple I → P edges; (ii) P → A edges link patents to their assignee institutions; and (iii) P → F edges connect patents to their IPC-based technical fields. All edges are timestamped using the grant year of each patent, embedding fine-grained temporal dynamics into the evolving network structure. Annual statistics across nodes and edges are reported in Table 1.

Table 1 reveals a general upward trend in network connectivity from 2015 to 2022, suggesting that collaboration intensity in the inventor–patent network increased during this period. The slight decline in 2023 and 2024 indicates a possible stabilization in structural density. This trend reflects the dynamic evolution of collaboration intensity in the AI patent landscape.

We construct input features for four node types using a type-specific embedding scheme. Patent nodes are represented by 768-dimensional semantic vectors generated from concatenated titles and abstracts using PatentSBERTa. Inventor and assignee nodes lack structured attributes and are initialized with randomly sampled, trainable embeddings. Technical field nodes are derived from four-character IPC subclass codes and encoded as one-hot vectors. The one-hot vector length corresponds to the 274 IPC subclasses observed in the dataset, with unused subclasses omitted to reduce sparsity. All node features are then mapped into a shared 128-dimensional embedding space via learnable, type-specific linear projections. A summary of the node feature construction is provided in Table 2 below.

Hyperparameter Configuration. The DHGNN model uses a specific set of hyperparameter settings selected for this study. All inventor nodes are embedded into a

Table 1 Yearly statistics of the heterogeneous inventor–patent network (2015–2024)

Year	Inventors	Patents	Assignees	Fields	I → P	P → A	P → F
2015	529	415	100	45	1233	437	578
2016	798	615	157	50	1892	685	853
2017	1159	1065	224	62	3617	1200	1517
2018	1471	1355	267	80	4630	1546	2054
2019	1991	2009	339	92	6619	2223	3204
2020	3156	3659	523	125	11,358	4011	6544
2021	4292	6168	781	158	18,993	6899	11,758
2022	4755	8400	932	166	25,528	9614	19,239
2023	4354	6541	902	137	19,872	7433	12,903
2024	3636	5310	803	118	15,739	6279	11,008

Table 2 Type-Aware Initialization of Node Features

Node Type	Feature Initialization	Feature Dimension
Patent	Pretrained PatentSBERTa embedding from titles and abstracts	768
Inventor	Learnable embedding (random initialization)	128
Assignee	Learnable embedding (random initialization)	128
Technical field	One-hot IPC subclass encoding (274 subclasses used)	274

128-dimensional latent space. Both the structure-level and temporal attention modules employ 8 attention heads, resulting in 16-dimensional subspaces for each head. The input graph is organized as annual snapshots, and the latest $L = 3$ snapshots are used to capture temporal dependencies. Structural information is encoded using a two-layer relation-aware graph encoder before temporal modeling.

To prevent overfitting caused by sparse collaboration links, two regularization techniques are applied. A dropout rate of 0.10 is introduced after each relation-level message-passing operation and each meta-path aggregation step. Dropout is not used in the temporal attention and bilinear decoding components, as preliminary tests indicated that additional randomness in these stages impairs convergence. Additionally, L2 regularization is applied to all trainable parameters with a weight decay coefficient of 1×10^{-4}. All regularization settings are selected based on validation performance and kept fixed throughout all experiments. For subgraph extraction, we tune the coverage threshold $\eta \in \{0.85, 0.90, 0.95\}$ and the size limit $K \in \{10, 12, 16\}$ based on validation performance. The optimal setting $\eta = 0.90$, $K = 12$ is used consistently across all experiments to ensure faithful yet concise explanations.

Model optimization is performed using the Adam algorithm with an initial learning rate of 10^{-3}, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The maximum number of training epochs is capped at 80. No learning rate decay schedule is applied, as early stopping based on validation AUC—with a patience of 20 epochs—proved sufficient for convergence. Each training mini-batch contains 1024 inventor pairs, consisting of one observed (positive) and one randomly sampled unobserved (negative) pair. This 1:1 ratio maintains a balanced learning objective in the presence of class imbalance. All experiments are conducted using PyTorch 2.1 with automatic mixed precision (AMP) enabled, on a Windows 11 system with an NVIDIA RTX 4060 GPU (8 GB VRAM), an Intel Core i7-13650HX CPU, and 32 GB of RAM. Random seeds are fixed across runs to ensure consistent results.

4.3 Baselines and evaluation metrics

We benchmark DHGNN comprehensively against a variety of baseline methods spanning heuristic measures, static and dynamic graph neural networks (GNNs), as well as homogeneous and heterogeneous modeling approaches. Classical heuristic methods, such as Common Neighbors (CN) and Adamic–Adar (AA), provide baseline performance solely derived from structural similarities. Static homogeneous models, including GCN and GraphSAGE, are employed to evaluate the effectiveness of node representation learning without incorporating temporal or semantic aspects. By comparison, static heterogeneous models like R-GCN, HAN, and HGT integrate relational semantics explicitly.

To isolate the impact of temporal modeling, dynamic homogeneous models such as EvolveGCN-H, DySAT, and TGN are also included, capturing temporal dynamics but neglecting semantic heterogeneity. Advanced baselines consisting of dynamic heterogeneous GNNs—such as DHGAS, HTGNN, and DynHEN—are further considered, as these models integrate both temporal evolution and multi-type relational semantics. Collectively, these diverse approaches serve as comprehensive benchmarks for evaluating the robustness and effectiveness of DHGNN.

To ensure temporal validity and prevent label leakage, we adopt a strictly chronological data splitting strategy that aligns with established practices in dynamic graph learning (Yujie Fan et al. 2021; Xing et al. 2022). For each prediction year T, we construct the cumulative heterogeneous graph snapshot $G(T)$, which includes all inventor–patent–affiliation–field relations up to year T. The model aggregates temporal context using a three-year sliding window, incorporating snapshots from years $T-2$, $T-1$, and T to produce time-sensitive inventor representations. These representations are used to predict inventor–inventor collaborations occurring in the future window $(T, T+1]$, in accordance with the causal masking design introduced in Sect. 3.4.

The dataset is chronologically partitioned as follows. The training set spans the supervision intervals from 2018 to 2021. For example, when $T = 2018$, the model uses 2016–2018 historical snapshots to predict links formed in 2019. The validation set targets links in 2022 based on history from 2019–2021 ($T = 2021$), while the test set predicts 2023 collaborations using snapshots from 2020–2022 ($T = 2022$). The most recent year, 2024, is held out exclusively for qualitative case studies and is not involved in any parameter tuning or evaluation.

This configuration yields an 8:1:1 split of the overall dataset into training, validation, and test subsets, ensuring strict non-overlap of target links across different stages. For each prediction year, positive labels consist of inventor pairs that collaborate at least once within $(T, T+1]$. Negative pairs are uniformly sampled from non-collaborating inventor pairs in the same period. All baseline models are retrained under the same supervision windows and candidate sets to ensure fair and reproducible comparisons. Table 3 provides an overview of all baseline models used in this study.

We evaluate model performance using four complementary metrics: area under the ROC curve (AUC), average precision (AP), Hits@K (K = 5, 10, 20), and mean reciprocal rank (MRR). AUC assesses the model's ability to distinguish positive links under class imbalance, while AP reflects

Table 3 Overview of Baseline Models

Category	Model	Reference
Heuristic	CN, AA	—
Static Homogeneous	GCN, GraphSAGE	(Kipf and Welling 2016; Hamilton et al. 2017)
Static Heterogeneous	R-GCN, HAN, HGT	(Schlichtkrull et al. 2018; Xiao Wang et al. 2019b; Hu et al. 2020)
Dynamic Homogeneous	EvolveGCN-H, DySAT, TGN	(Pareja et al. 2020; Sankar et al. 2020; Rossi et al. 2020)
Dynamic Heterogeneous	DHGAS, HTGNN, DynHEN	(Zeyang Zhang et al. 2023; Yujie Fan et al. 2022; Xing et al. 2022)
Ours	DHGNN	—

ranking quality across the full recall spectrum. Hits@K measures top-K accuracy by checking whether the ground-truth link appears among the top-K predicted scores. MRR captures how early a correct prediction appears in the ranked list. All metrics are computed per test instance and averaged over five independent runs under the same temporal train/validation/test split to ensure robustness.

4.4 Experimental results

Table 4 compares the performance of our proposed DHGNN model against a comprehensive set of baseline methods on the 2023 test slice.

The experimental results highlight the following key observations:

Table 4 Link Prediction Results on the 2023 Test Slice (Mean of 5 Runs)

Model	AUC	AP	Hits@5	Hits@10	Hits@20	MRR
CN	0.646	0.212	2.8	4.3	7.9	0.105
	±0.014	±0.013	±0.9 pp	±1.2 pp	±1.6 pp	±0.008
AA	0.667	0.229	3.2	5.1	8.6	0.109
	±0.011	±0.012	±1.0 pp	±1.3 pp	±1.7 pp	±0.008
GCN	0.742	0.318	6.4	10.5	16.7	0.162
	±0.010	±0.011	±0.8 pp	±1.0 pp	±1.3 pp	±0.007
GraphSAGE	0.763	0.337	7.0	11.6	18.1	0.174
	±0.009	±0.011	±0.9 pp	±1.1 pp	±1.2 pp	±0.007
R-GCN	0.775	0.352	7.8	12.4	19	0.18
	±0.008	±0.010	±0.8 pp	±1.0 pp	±1.3 pp	±0.006
HAN	0.792	0.371	8.6	13.5	20.4	0.193
	±0.007	±0.009	±0.9 pp	±1.2 pp	±1.4 pp	±0.005
HGT	0.823	0.402	9.8	15.2	23.1	0.212
	±0.007	±0.009	±0.9 pp	±1.1 pp	±1.4 pp	±0.006
EvolveGCN-H	0.784	0.357	8.2	12.7	19.7	0.188
	±0.009	±0.010	±0.8 pp	±1.0 pp	±1.5 pp	±0.007
DySAT	0.803	0.376	8.9	13.8	21	0.198
	±0.008	±0.010	±1.0 pp	±1.0 pp	±1.3 pp	±0.004
TGN	0.834	0.411	10.3	15.9	24.5	0.218
	±0.007	±0.009	±0.9 pp	±1.1 pp	±1.4 pp	±0.006
DHGAS	0.853	0.435	11.3	17.6	26.9	0.236
	±0.006	±0.008	±0.8 pp	±0.9 pp	±1.1 pp	±0.005
HTGNN	0.861	0.447	11.9	18.4	28	0.243
	±0.008	±0.009	±0.9 pp	±1.2 pp	±1.4 pp	±0.006
DynHEN	0.844	0.427	10.8	16.8	25.8	0.229
	±0.006	±0.008	±0.8 pp	±1.1 pp	±1.3 pp	±0.005
DHGNN (ours)	**0.892**	**0.486**	**13.8**	**21.9**	**31.6**	**0.269**
	±0.005	±0.007	±0.7 pp	±0.9 pp	±1.1 pp	±0.005

(1) **Graph Neural Networks Significantly Outperform Heuristics.** All GNN-based methods substantially exceed heuristic baselines (CN and AA), with gains exceeding 10 AUC points. This demonstrates the effectiveness of representation learning even in structurally sparse graphs.

(2) **Heterogeneous Modeling Enhances Predictive Power.** Static heterogeneous models (R-GCN, HAN, HGT) consistently outperform static homogeneous models (GCN, GraphSAGE), with HGT achieving 0.823 AUC and 23.1% Hits@20—showcasing the value of semantic type distinctions in complex inventor–patent networks.

(3) **Temporal Dynamics Introduce Orthogonal Gains.** Incorporating time through dynamic homogeneous models (DySAT, TGN) yields measurable improvements over their static counterparts. For instance, TGN outperforms HAN by + 4.1% Hits@20 and + 0.025 MRR, validating the benefit of time-aware message passing.

(4) **Dynamic Heterogeneous Architectures Achieve Strong Performance.** Models that combine temporal dynamics and semantic heterogeneity—such as DHGAS, HTGNN, and DynHEN—consistently perform well across all metrics. HTGNN reaches 0.861 AUC and 0.243 MRR, making it the strongest baseline before DHGNN.

(5) **DHGNN Sets a New State-of-the-Art.** Our proposed model DHGNN consistently achieves the highest scores on all metrics, with particularly large gains in early retrieval precision (Hits@5: 13.8%) and MRR (0.269). Compared to the strongest prior baseline HTGNN, DHGNN yields **+ 3.1 points AUC, + 4.0 points AP, + 3.6 percentage points Hits@20, and + 2.6 points MRR**, clearly demonstrating the advantage of jointly modeling temporal dynamics, heterogeneous semantics, and path-level relational structure within a unified framework.

To more clearly articulate the architectural distinctions of DHGNN, Table 5 offers a comparative overview with three representative dynamic heterogeneous GNN baselines—DHGAS, HTGNN, and DynHEN. The comparison spans four fundamental design aspects (message passing, temporal modeling, heterogeneous relation handling, and interpretability), along with corresponding performance metrics and training time.

As summarized in Table 5, DHGNN provides a carefully designed architecture tailored for dynamic inventor collaboration prediction. Compared to prior models, DHGNN integrates relation-aware local propagation with meta-path–guided multi-hop semantic aggregation, allowing effective modeling of both local neighbor information and high-order semantic interactions, such as co-affiliation relationships captured by specific meta-paths.

In contrast to DHGAS and HTGNN, which utilize unified or hierarchical attention mechanisms for temporal and heterogeneous relation aggregation, DHGNN employs discrete cumulative snapshots to explicitly represent temporal evolution. It further enhances temporal modeling by incorporating continuous-time positional encodings along with causally masked self-attention to maintain proper chronological ordering during representation learning. This approach helps DHGNN effectively capture temporal dynamics without imposing unrealistic assumptions of uniform temporal effects.

Moreover, DHGNN explicitly incorporates interpretability through a two-tier explanation mechanism—meta-path attribution followed by subgraph extraction—whereas none of the three baselines provide comparable explanatory capabilities. This design significantly enhances transparency and interpretability, which are highly valuable in scientific collaboration forecasting scenarios.

Together, these architectural improvements translate into superior predictive performance across multiple evaluation metrics, as evidenced by the results presented in Table 5. However, these advances also incur slightly higher computational costs, leading to a moderate increase in training time relative to simpler baselines such as DynHEN and HTGNN. Therefore, future research can explore strategies to optimize model complexity, potentially developing more efficient implementations that maintain DHGNN's interpretability and predictive performance.

4.5 Ablation studies

To quantitatively assess the contribution of each architectural component within DHGNN, we design a set of controlled ablation experiments by selectively removing key modules: the temporal attention mechanism, the meta-path-based semantic propagation layer, and the dual-level attention framework. We further implement a TGAT-style ablated variant of DHGNN by replacing the original snapshot-level temporal aggregation module with a single continuous-time attention layer that incorporates functional time encoding and causal masking (Xu et al. 2020). In parallel, to assess whether a finer temporal granularity yields additional predictive gains, we introduce an ablation variant that segments the dynamic graph into cumulative quarterly snapshots instead of yearly ones, while maintaining the same data partitioning protocol to ensure fair comparison. All other components are kept identical. Table 6 presents the performance of each ablated variant on the 2023 test set.

As shown in Table 6, each component plays an indispensable role in achieving the final performance. When the temporal module is removed, the model—reduced to a static heterogeneous GNN—suffers a notable 2.2-point drop in AUC and a 3.6-point reduction in Hits@10. This confirms that temporal dynamics encode non-redundant structural patterns that static topology alone cannot capture. Similarly, disabling the meta-path propagation layer leads to a 2.5-point loss in AUC and the second-largest decrease in Hits@10

Table 5 Model Architecture Comparison

Dimension	DHGNN	DHGAS	HTGNN	DynHEN
Message passing	Relation-aware and meta-path–guided multi-hop propagation capturing both local and high-order semantics	Unified type- and time-aware attention; the NAS procedure decides which relations/timestamps to attend	Hierarchical: intra-relation, inter-relation, and across-time aggregation in sequence	Two-stage flow: per-snapshot heterogeneous encoding, followed by a temporal attention block
Temporal modeling	Discrete cumulative snapshots combined with continuous-time positional embeddings and causally masked self-attention over history	Discrete snapshots; attention scores are explicitly parameterized by relation type and the time difference, integrating space and time in one layer	Discrete snapshots; each layer includes an across-time aggregation that links adjacent slices	Discrete snapshots; a temporal self-attention module is applied after spatial encoding
Heterogeneous relations	Relation-aware parameters and predefined meta-path–guided multi-hop channels with meta-path-level attention	Type-aware mappings with relation/time-specific projections; no predefined meta-paths	Relation-wise attentions only; no meta-path semantics	Bipartite user–item setting with a small set of predefined meta-paths inside HGCN (e.g., u–i–u, i–u–i)
Interpretability	Two-tier explanation: meta-path attribution plus subgraph extraction for each prediction	No explicit explanation module	No explicit explanation module	No explicit explanation module
AUC/ AP/ Hits@20/ MRR)	0.892/ 0.486/ 31.6%/ 0.269	0.853/ 0.435/ 26.9%/ 0.236	0.861/ 0.447/ 28.0%/ 0.243	0.844/ 0.427/ 25.8%/ 0.229
Training Time (mins)	37.5	33.3	27.1	21.3

Table 6 Performance of DHGNN and Ablated Variants (2023 Test Slice)

Variant	Temporal Module	Meta-paths	Dual-level Attn	AUC	AP	Hits@10
– w/o Temporal module	✗	✓	✓	0.870 ±0.006	0.456 ±0.008	18.3 ±1.0 pp
– w/o Meta-paths	✓	✗	✓	0.867 ±0.009	0.452 ±0.007	17.7 ±1.1 pp
– w/o Path-level attention	✓	✓	relation-level only*	0.878 ±0.005	0.468 ±0.006	19.5 ±0.8 pp
– w/o Any attention	✓	✓	✗	0.855 ±0.007	0.435 ±0.009	16.6 ±1.2 pp
–TGAT-style temporal	✓	✓	✓	0.883 ±0.010	0.482 ±0.011	21.2 ±1.5 pp
–Quarterly snapshots	✓	✓	✓	0.879 ±0.009	0.474 ±0.010	20.3 ±1.3 pp
DHGNN (full)	✓	✓	✓	0.892 ±0.005	0.486 ±0.007	21.9 ±0.9 pp

* "relation-level only" retains relation-aware neighbor attention and meta-path propagation (Eq. (5)) but removes meta-path–level attention (Eqs. (6)–(7)) by replacing the learned path-level weights with uniform averaging. The temporal module remains unchanged

among all variants (−4.2 pp), emphasizing the importance of modeling higher-order type-aware interactions between inventors, patents, and assignees.

Disabling meta-path–level semantic attention while preserving relation-aware neighborhood aggregation, results in a moderate performance drop (−0.014 AUC, −2.4 pp Hits@10). This suggests that relation-type aggregation alone lacks the expressiveness to capture fine-grained collaboration cues without explicit path-level weighting. When all attention mechanisms are removed and replaced with uniform mean aggregation, the model suffers the most substantial degradation (−0.037 AUC, −5.3 pp Hits@10). This highlights that attention in DHGNN not only enhances interpretability but also serves as a vital signal amplifier, particularly under structural sparsity.

The TGAT-style temporal variant delivers performance closely matching that of DHGNN in terms of AUC, AP, and Hits@10. However, it requires over 50% longer training time. Although fine-grained timestamps are available, the effective timing of collaboration is difficult to determine due to long and highly variable delays between patent application and grant. Our proposed scheme—segmenting the graph into cumulative yearly snapshots and applying masked self-attention with continuous positional encodings—better aligns with the inherent temporal granularity of inventor–patent dynamics and offers a more favorable balance among accuracy, efficiency, and robustness.

The quarterly snapshot variant registers performance drops of 0.013 in AUC, 0.012 in AP, and 1.6 pp in Hits@10 relative to the default yearly setting. Increasing temporal granularity brings a finer chronological view, yet each quarterly snapshot contains substantially fewer inventor–collaboration events than a yearly snapshot. This lower interaction density per snapshot weakens the contextual signal available to temporal attention and offsets the potential benefit of finer resolution. Furthermore, the larger number of snapshots introduces a considerable computational burden, with training time increasing by more than twofold. These observations indicate that yearly snapshots offer a more favorable balance between temporal expressiveness and computational efficiency, which is crucial for modeling long-horizon and inherently sparse inventor–patent collaboration networks.

These ablation results collectively validate the necessity of integrating temporal evolution, heterogeneous semantics, and hierarchical attention. Each component contributes complementary inductive biases, and only their joint application enables DHGNN to fully realize its predictive capacity.

4.6 Sensitivity Analysis

To evaluate how the predictive performance of our model responds to variations in key hyperparameter configurations, we perform a targeted sensitivity analysis on six influential parameters. These include architectural choices such as the number of GNN layers, as well as optimization and regularization factors like learning rate, dropout rate, and weight decay. The resulting AUC trends under each setting are summarized in Fig. 5, offering insights into the stability and tuning robustness of DHGNN across a wide operational range.

As shown in Fig. 5, DHGNN attains its highest AUC under moderately tuned hyper-parameters. Specifically, a three-year historical window ($L = 3$) yields the highest AUC, effectively balancing recency and temporal depth. An embedding size of 128 achieves the best trade-off between expressiveness and overfitting, while both smaller and larger dimensions lead to slight performance drops. The model is

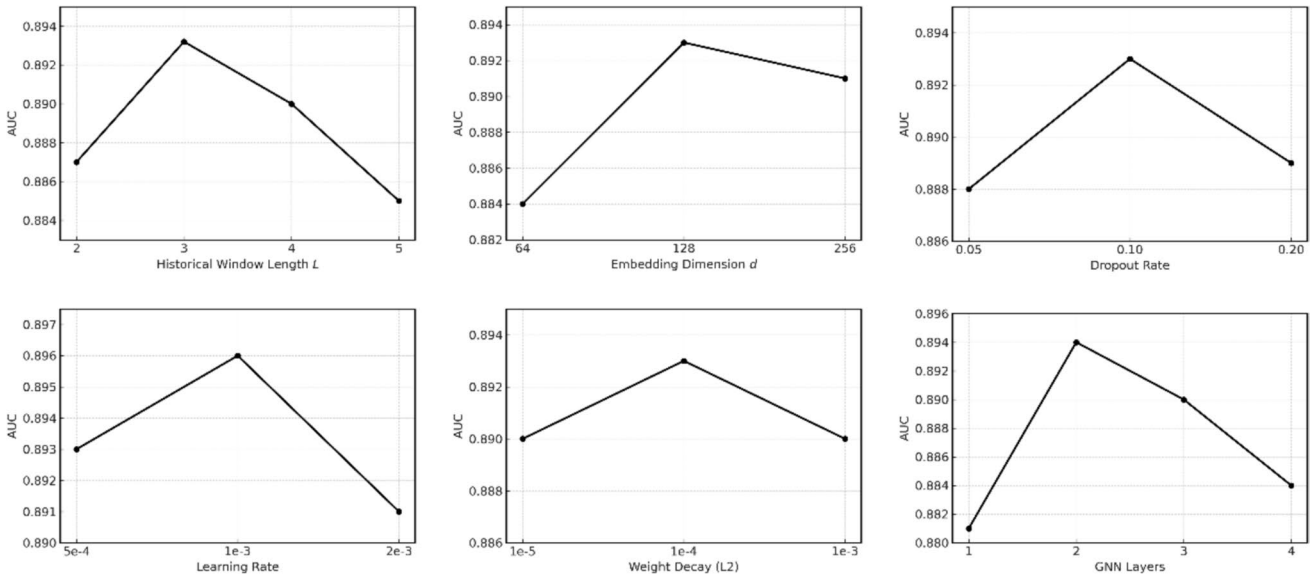

Fig. 5 Sensitivity Analysis Results

most robust when applying a 0.1 dropout rate, indicating sufficient regularization without disrupting representation learning. A learning rate of 1×10^{-3} results in the most reliable convergence, and weight decay performs best at 1×10^{-4}, ensuring generalization without under-regularization. Lastly, using two GNN layers consistently outperforms shallower and deeper variants. Overall, these results demonstrate that DHGNN maintains stable performance across a wide hyperparameter range, while emphasizing the impact of well-calibrated configurations on model effectiveness.

Table 7 Quantitative interpretability benchmarks

Method	Faithfulness↑	Infidelity ↓	Sparsity ↓	Stability ↑
DHGNN	0.88 ± 0.02	0.07 ± 0.01	0.10	0.95 ± 0.01
GNNExplainer	0.90 ± 0.01	0.05 ± 0.00	0.15	0.85 ± 0.03
PGExplainer	0.86 ± 0.02	0.09 ± 0.01	0.12	0.90 ± 0.02
GraphLIME	0.72 ± 0.03	0.15 ± 0.02	0.00*	0.88 ± 0.02

*GraphLIME produces feature-level explanations and does not select edges, thus its edge-sparsity is effectively 0.00

4.7 Evaluation of explanation quality

We quantify the effectiveness of DHGNN's two-tier explanation module using four standard metrics from graph explainability literature. These include faithfulness, infidelity, sparsity, and stability, which collectively assess whether the explanations are causally relevant, parsimonious, and consistent.

We conduct the evaluation on the 2023 test slice of the Chinese AI patent dynamic heterogeneous graph. Specifically, we select the top 1,000 inventor pairs with the highest predicted collaboration scores in 2023. For each pair (u, v), we construct a 2-hop typed ego-graph centered at (u, v) on the cumulative graph snapshot up to 2023. This instance-aligned subgraph preserves heterogeneous edge types and respects causal time, ensuring consistency across methods.

We compare DHGNN with three representative model-agnostic explanation baselines: GNNExplainer (R. Ying et al. 2019b), PGExplainer (Dongsheng Luo et al. 2020a), and GraphLIME (Huang et al. 2023). All methods operate post hoc on frozen node representations and share the same ego-graph structure. DHGNN explanations are derived by ablating the most influential meta-path instance and its associated edges. For baselines, we follow each method's original masking or perturbation protocol. Table 7 reports the quantitative comparison.

Table 7 presents a quantitative comparison of explanation quality across four representative models, revealing three key observations:

(1) DHGNN achieves a balanced performance across faithfulness, stability, and sparsity. Specifically, as reported in Table 7, DHGNN attains faithfulness of 0.88 and stability of 0.95, with a sparsity value of 0.10. Compared to GNNExplainer, DHGNN shows marginally lower faithfulness (-0.02) but noticeably higher stability ($+0.10$) and more concise explanations (-0.05 sparsity). These results indicate that DHGNN's explanations generally require fewer edges to support the predictions, thus achieving clearer and more consistent rationales.

(2) Structure-aware explanation methods provide greater explanatory effectiveness than purely feature-based methods. GraphLIME, which generates feature-level explanations

without explicitly selecting edges, inherently yields zero edge sparsity. However, it obtains the lowest faithfulness (0.72) and highest infidelity (0.15). This indicates a clear limitation in capturing structural dependencies, confirming that neglecting structural relationships substantially reduces explanatory power in dynamic heterogeneous graphs.

(3) DHGNN shows improved stability compared to PGExplainer with similar levels of faithfulness. PGExplainer achieves reasonably good performance (faithfulness: 0.86, stability: 0.90); however, DHGNN slightly surpasses it in both stability (+0.05) and sparsity (0.10 vs. 0.12). This observation suggests that DHGNN's meta-path-based explanations tend to be more stable and concise.

4.8 Cross-national generalization of collaboration prediction

To assess the cross-national generalizability of DHGNN, we extend the experiment to include inventor–patent networks from the United States and Republic of Korea, alongside the original Chinese dataset. All patent data were retrieved from the IncoPat database, covering granted patents in the artificial intelligence sector from 2015 to 2024. Given language divergence and institutional naming inconsistencies—particularly in Korean—the inventor and assignee fields across all datasets were unified using DWPI-standardized names, thereby mitigating ambiguities during entity disambiguation and graph construction.

Each dataset is modeled as a dynamic heterogeneous graph with four node types (Inventor, Patent, Assignee, IPC Field) and three timestamped relation types. To ensure comparability, all networks adopt a three-year cumulative snapshot window for temporal segmentation. Table 8 summarizes DHGNN's predictive performance across the three national datasets.

As shown in Table 8, the U.S. graph yields the strongest prediction performance among the three, which can be attributed to several structural characteristics. First, U.S. AI patents demonstrate a higher average number of inventors per patent, expanding the semantic neighborhood available for meta-path aggregation. Second, institutional diversity is more pronounced, with frequent inter-organizational co-invention, which enhances both the depth and breadth of high-order semantic paths. Third, the name standardization is particularly effective in English-language datasets, where DWPI alignment captures institutional hierarchies more precisely.

The Korean graph, while substantial in size, presents greater challenges due to denser name duplication, smaller inventor teams, and higher intra-organization clustering. These characteristics reduce the diversity of available semantic paths and limit long-range structural signals. Despite these constraints, the model still maintains competitive performance, demonstrating its robustness under sparse or topologically constrained conditions.

In summary, the observed differences validate DHGNN's capacity to generalize across structurally divergent inventor–patent networks, and suggest that performance is positively correlated with the semantic richness, disambiguation clarity, and inter-institutional link density of the underlying data.

4.9 Cold-start evaluation

We define a cold-start scenario as a prediction task involving inventor pairs with no prior co-invention records in the training data but who form new collaborations in the test year. This setting requires the model to rely on semantic path cues (e.g., shared assignees, research domains, or indirect co-inventors) and temporal evolution patterns to infer latent

Table 8 Performance of DHGNN on National Inventor–Patent Networks

Country	AUC	AP	Hits@5	Hits@10	Hits@20	MRR
China	0.892 ±0.005	0.486 ±0.007	13.8 ±0.7 pp	21.9 ±0.9 pp	31.6 ±1.1 pp	0.269 ±0.005
United States	0.916 ±0.009	0.507 ±0.010	15.9 ±0.9 pp	25.1 ±1.0 pp	35.3 ±1.2 pp	0.287 ±0.009
Korea	0.862 ±0.008	0.452 ±0.012	14.2 ±1.1 pp	20.3 ±1.2 pp	28.2 ±1.5 pp	0.257 ±0.011

Table 9 Performance on Cold-Start Collaboration Prediction

Model	AUC	AP	MRR	Hits@10
CN	0.620 ± 0.015	0.270 ± 0.016	0.041 ± 0.005	4.7 ± 1.3pp
TGN	0.758 ± 0.012	0.365 ± 0.014	0.148 ± 0.007	11.1 ± 1.1pp
HTGNN	0.792 ± 0.010	0.392 ± 0.012	0.168 ± 0.006	12.5 ± 1.1pp
DHGNN	0.836 ± 0.009	0.420 ± 0.010	0.192 ± 0.006	14.2 ± 0.9pp

collaboration potential. We compare DHGNN against three representative baselines selected from different modeling families: CN (structural heuristic), TGN (temporal GNN), and HTGNN (heterogeneous temporal GNN). All other experimental configurations follow the protocol described in Sects. 4.2–4.3, ensuring strict comparability. Table 9 summarizes the performance across four ranking-based metrics.

As shown in Table 9, DHGNN consistently outperforms all baselines under cold-start conditions, achieving a 4.4-point gain in AUC and a 1.7-point improvement in Hits@10 over HTGNN. Notably, the largest improvements appear in ranking-sensitive metrics such as MRR and Hits@10, which are critical in practical recommendation settings. These results suggest that DHGNN can effectively identify top collaboration candidates even in the absence of historical links, by leveraging semantic context and temporal alignment. This finding is consistent with our interpretability analysis, which highlights the importance of high-order semantic paths in cold-start predictions.

We note that the extreme case of inventors without any prior patents was not included in this evaluation. This is primarily due to limitations in available data and the infeasibility of deriving input representations without patent history. Addressing such cases would likely require external auxiliary information or alternative feature designs, which can be considered as a promising direction for future work.

4.10 Case studies

To investigate DHGNN's predictive reasoning in identifying inventor collaborations under complex semantic and structural conditions, we present an empirical case analysis centered on Zhigang Chen, a prolific inventor affiliated with iFLYTEK Corporation. The primary goal is to determine the model's capability to accurately forecast plausible future collaborations by integrating semantic signals and indirect structural cues alongside explicit historical collaboration data.

We construct cumulative co-invention subgraphs reflecting Chen's collaborative interactions from 2021 to 2023, as depicted in Fig. 6. These aggregated network structures effectively illustrate the evolving dynamics of Chen's collaboration environment, with nodes representing inventors and edges indicating established collaborations through co-authored granted patents. Notably, the network reveals a stable core group of collaborators accompanied by peripheral inventors, some of whom directly collaborated with Chen during this period, while others remained indirectly connected or unconnected, indicating potential latent relationships accessible through semantic or structural inference.

To systematically assess DHGNN's predictive accuracy, we generated a ranked list of the top-20 potential collaborators for Chen in 2024. Table 10 reports these predictions along with actual collaborations recorded in the 2024 patent data. Impressively, three of the top-five predicted collaborators confirmed co-authorships in 2024. The remaining predictions, though not realized, represent semantically plausible collaborations driven by shared institutional affiliations (e.g., iFLYTEK) and overlapping technological domains (e.g., IPC codes G06F and G06N). These results highlight DHGNN's capability to leverage multi-dimensional semantic cues and indirect structural paths rather than solely relying on immediate structural proximity.

Case I: Cold-Start Collaboration Based on Semantic Affiliation. A central objective of this study is to forecast scientific collaborations in cold-start scenarios—where traditional structural signals such as co-authorship are unavailable. DHGNN addresses this challenge by leveraging heterogeneous meta-path semantics to infer potential collaborations. To evaluate its effectiveness, we analyze a top-5 prediction between Zhigang Chen and Yiming Cui, who had no recorded co-invention activities from 2021 to 2023, yet

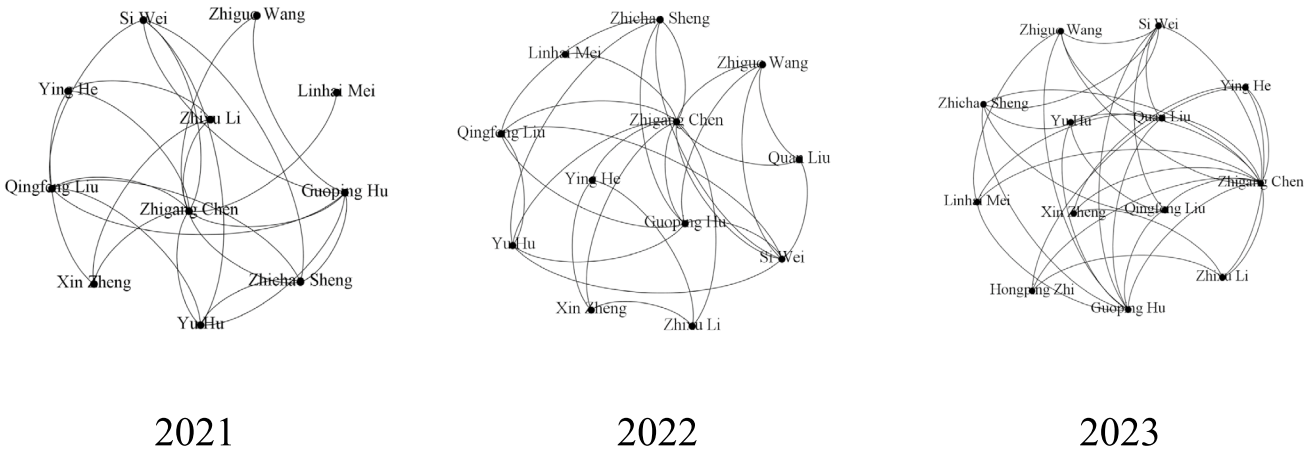

Fig. 6 Temporal Evolution of Zhigang Chen's Co-Invention Network

Table 10 Top-20 Predicted Collaborators for Zhigang Chen (2024)

Rank	Inventor Name	Real Collaboration (2024)
1	Weitai Zhang	No
2	Zhixu Li	Yes
3	Qian Zhao	No
4	Junhua Liu	Yes
5	Yiming Cui	Yes
6	Xudong Dai	Yes
7	Jiaxue Hu	No
8	Guoping Hu	Yes
9	Si Wei	Yes
10	Zhiguo Wang	Yes
11	Bing Qin	No
12	Bing Yin	No
13	Wei Song	No
14	Jianqing Gao	No
15	Ting Liu	No
16	Ruiji Fu	No
17	Xin Zheng	Yes
18	Genshun Wan	No
19	Jia Pan	No
20	Linhai Mei	Yes

Table 11 Meta-Path Attention Distribution for Case I (Chen–Cui)

Meta-Path	Semantic Type	Attention Weight (α)
I-P-A-P-I	Shared Assignee (iFLYTEK)	0.52
I-P-F-P-I	Shared IPC fields (G06F, G06N)	0.28
I-P-I	Weak structural co-invention	0.12
I-P-I-P-I	Multi-Hop Structural Bridge	0.08

were forecast to collaborate in 2024—a prediction subsequently confirmed by actual patent data.

To interpret the model's decision, we examine the semantic attention distribution over four predefined meta-paths. As shown in Table 11, the I–P–A–P–I path accounts for the largest share of attention ($\alpha = 0.52$), highlighting that shared affiliation—through repeated co-assignment to iFLYTEK—played a dominant role in the prediction. The I–P–F–P–I path, capturing domain-level proximity via shared IPC classifications such as G06F and G06N, ranks second ($\alpha = 0.28$), indicating that technical relevance provided meaningful, albeit supplementary, context.

In contrast, structural paths receive minimal weight: I–P–I (direct weak ties) and I–P–I–P–I (bridge-based weak ties) score only $\alpha = 0.12$ and 0.08, respectively. This distribution is consistent with the cold-start setting, where no historical structural connectivity exists to guide inference—reinforcing the importance of non-structural semantics in such cases.

To further explain the prediction, we extract representative subgraph instances under the two most influential meta-paths—I–P–A–P–I and I–P–F–P–I. As visualized in Fig. 7, the first five instances connect Zhigang Chen and Yiming Cui through the assignee node iFLYTEK, demonstrating repeated co-assignment across multiple shared patents. Additionally, the lower part of the figure presents two instances of the I–P–F–P–I path, which link the inventors via common IPC fields such as G06F, indicating overlapping expertise in AI-related domains.

In this cold-start scenario, the model's prediction is predominantly informed by institutional affiliation. With an attention weight of $\alpha = 0.52$, the I–P–A–P–I path—capturing repeated co-assignment to the assignee iFLYTEK—emerges as the primary semantic driver, consistent with prior findings on organizational proximity in scientific collaboration.

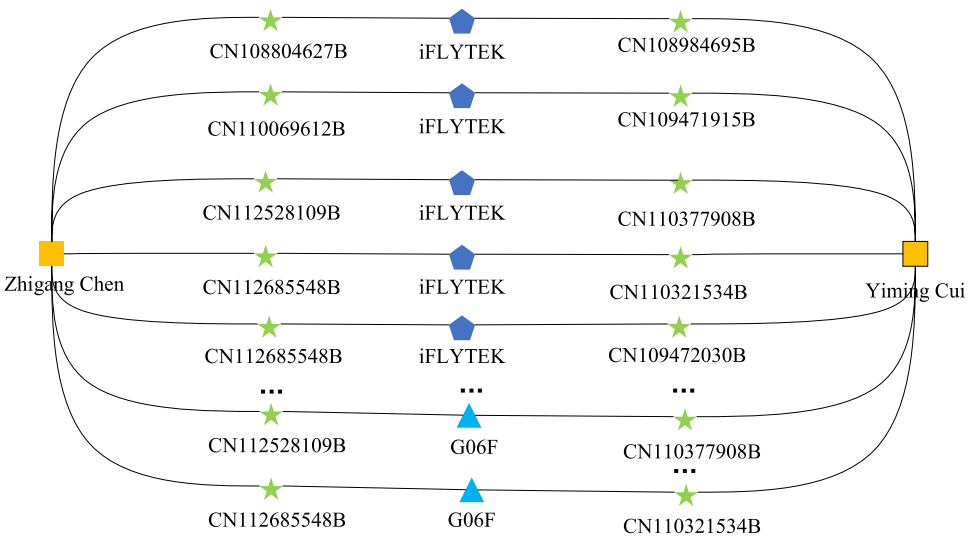

Fig. 7 Meta-Path Visualization for Case I (Chen–Cui)

The I–P–F–P–I path, with $\alpha = 0.28$, provides complementary domain-level context through shared IPC fields such as G06F and G06N. In contrast, both structural paths—I–P–I and I–P–I–P–I—are significantly downweighted ($\alpha = 0.12$ and 0.08), reinforcing the model's ability to suppress unreliable structural signals when historical co-invention is absent.

These results exemplify DHGNN's capacity to identify latent collaborative potential by leveraging semantic signals embedded in heterogeneous innovation networks. Through interpretable meta-path attention, the model delivers forward-looking predictions even under structurally sparse conditions.

Case II: Structure-Led Collaboration with Bridge-Based Reasoning. Beyond its ability to identify latent collaborations through semantic alignment, DHGNN is particularly effective when structural signals are strong and temporally relevant. This case analyzes a high-confidence prediction between Zhigang Chen and Shijin Wang, who had already co-invented two patents during 2021–2023 and were correctly forecast to collaborate again in 2024.

As shown in Table 12, the I–P–I path receives the highest attention weight ($\alpha = 0.42$), indicating that recent co-invention was the most decisive factor in the model's reasoning. In parallel, the model also leverages I–P–I–P–I paths ($\alpha = 0.30$) to capture multi-hop bridge structures, where intermediate inventors such as Quan Liu and Guoping Hu had previously collaborated with both target inventors. These indirect paths act as redundant yet informative structural reinforcements, enhancing model robustness.

Semantic meta-paths provide contextual signals that complement structural reasoning. The I–P–F–P–I path ($\alpha = 0.18$) captures shared expertise in AI-related domains (e.g., G06N), while I–P–A–P–I ($\alpha = 0.10$) reflects partial institutional overlap. Though secondary in weight, these channels support the plausibility of future collaboration through domain and organizational alignment.

Figure 8 visualizes a composite subgraph constructed from all four meta-paths used in the model's prediction. Two I–P–I instances—based on shared patents CN111400489B and CN112231458B—form the direct co-invention backbone between Zhigang Chen and Shijin Wang. Additionally, I–P–I–P–I chains such as Chen → Quan Liu → Wang and Chen → Guoping Hu → Wang capture multi-hop bridge structures, offering complementary structural support.

This case illustrates how DHGNN fuses temporal structural continuity with bridge-based reasoning to make interpretable collaboration predictions. While the model prioritizes high-frequency co-invention via I–P–I, it simultaneously integrates multi-hop bridging and semantic cues for enhanced stability and contextual validity. These results affirm DHGNN's ability to perform structure-aware, semantically supported reasoning in dynamic heterogeneous graphs—precisely addressing the challenge of interpretable forecasting in innovation networks.

The two cases highlight DHGNN's adaptive reasoning: in the cold-start case (Chen–Cui), the model relies on semantic paths—I–P–A–P–I ($\alpha = 0.52$) and I–P–F–P–I ($\alpha = 0.28$)—due to the absence of structural links; in contrast, the structure-led case (Chen–Wang) prioritizes I–P–I ($\alpha = 0.42$) and I–P–I–P–I ($\alpha = 0.30$), leveraging direct and bridge-based ties. As shown in Fig. 9, DHGNN flexibly shifts attention between structure and semantics, enabling transparent prediction across varying collaboration contexts.

Table 12 Meta-Path Attention Distribution for Case II (Chen–Wang)

Meta-Path	Semantic Channel	Attention Weight (α)
I-P-I	Recent direct co-invention	0.42
I-P-I-P-I	Multi-hop bridges via shared inventors	0.30
I-P-F-P-I	Common IPC domains (e.g., G06N)	0.18
I-P-A-P-I	Shared assignee	0.10

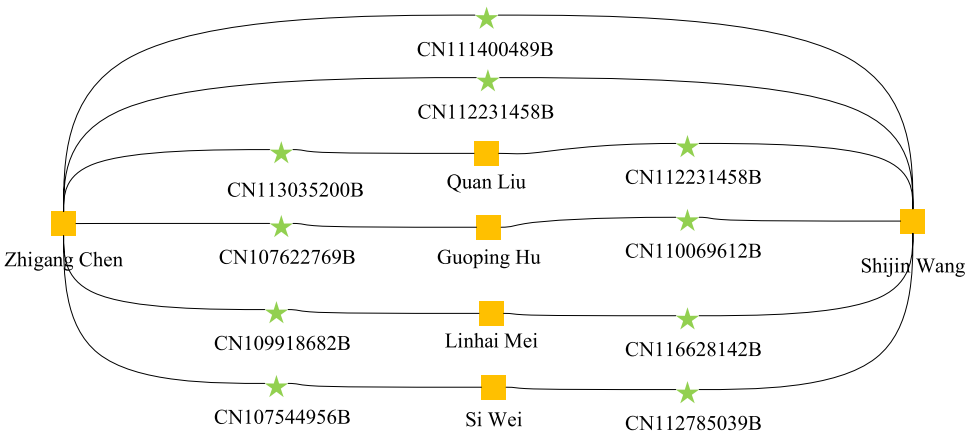

Fig. 8 Meta-Path Visualization for Case II (Chen–Wang)

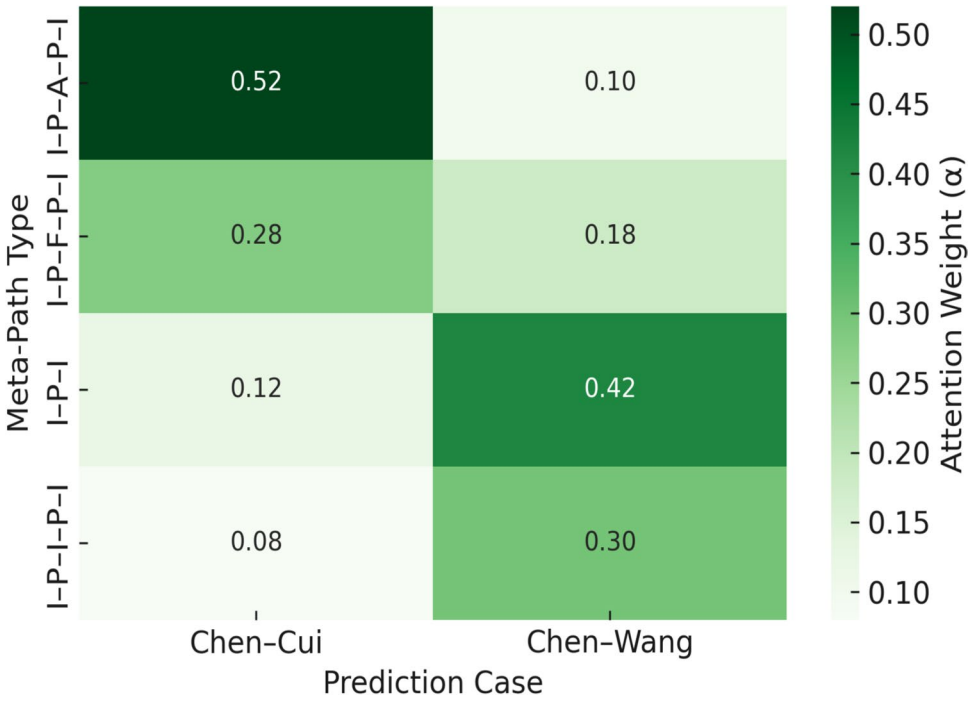

Fig. 9 Meta-Path Attention Distribution for Case I and Case II

5 Discussion

The experimental results clearly demonstrate that DHGNN consistently outperforms a broad range of baseline methods, validating the importance of its dynamic heterogeneous design. Specifically, DHGNN achieves new state-of-the-art accuracy across all evaluation metrics, surpassing both static and dynamic graph baselines. For example, it improves AUC by over 3 points and Hits@10 by 3.5% compared to the strongest prior model, HTGNN. These improvements underscore the complementary benefits of jointly modeling temporal evolution and heterogeneous semantics.

Conventional static heterogeneous GNNs like HGT effectively encode multi-relational semantics but ignore temporal dynamics. Conversely, dynamic homogeneous models such as TGN emphasize continuous-time link dynamics but lack relational diversity. Even recent dynamic heterogeneous models (e.g., DynHEN) incorporate time and type-specific aggregation yet lack DHGNN's hierarchical temporal-semantic attention and meta-path–based interpretability. DHGNN's superior performance highlights the value of integrating temporal dependencies with rich semantic meta-path structures, effectively capturing collaboration patterns that previous models miss.

Beyond accuracy, our analysis reveals how each component contributes critically to DHGNN's overall effectiveness. The hierarchical temporal-semantic attention mechanism is not merely an interpretability feature but also functionally essential. Ablation studies show significant performance drops when removing either temporal attention or meta-path aggregation, demonstrating the necessity of these modules. For instance, removing temporal attention results in approximately a 2.2-point AUC drop and 3–4% decrease in Hits@10. Eliminating meta-path aggregation similarly reduces Hits@10 by over 4%. Such empirical findings affirm that the full model design is synergistically beneficial.

Critically, DHGNN provides robust interpretability, yielding actionable insights into collaboration formation. The hierarchical attention weights transparently highlight influential semantic paths and historical interactions. This transparency is particularly valuable in innovation networks, where understanding predictive rationale is as important as prediction accuracy itself. Our case studies illustrate that DHGNN can clearly explain predictions under both cold-start and structure-dominated scenarios by adaptively prioritizing semantic and structural signals. Under cold-start conditions, DHGNN leverages semantic paths—such as shared institutional affiliations or research domains—to predict new collaborations. In structure-rich contexts, the model appropriately shifts focus to direct co-invention paths and structural bridging through mutual collaborators. These intuitive explanations enhance trust and practical utility.

While DHGNN achieves promising results, certain limitations still constrain the broader applicability of our findings. First, the architecture depends on a manually curated set of four meta-paths; this design choice secures interpretability but risks missing latent relational patterns that lie outside domain expertise. Second, our empirical evidence—while cross-national—remains restricted to AI-related patent corpora from China, the United States and Korea; extrapolating

to other innovation regimes or to non-patent collaboration graphs therefore warrants caution. Third, the cold-start analysis excluded inventors without any patent history, leaving the model's behavior under extreme information sparsity untested. Finally, the richer temporal-semantic encoding introduces a moderate training-time overhead relative to lighter baselines, which may deter deployment in resource-constrained settings.

6 Conclusion

In this study, we introduced DHGNN, a Dynamic Heterogeneous Graph Neural Network designed specifically for scientific collaboration prediction. DHGNN jointly models structural heterogeneity and temporal evolution while providing interpretable predictions grounded in semantic structures.

The proposed approach constructs cumulative heterogeneous snapshots and employs hierarchical attention mechanisms to capture meaningful patterns from heterogeneous neighbors and leverage multi-step semantic connections. By integrating continuous-time positional encodings with causally masked self-attention, DHGNN adaptively emphasizes informative historical interactions. Experimental evaluations indicate competitive performance across multiple metrics, highlighting the benefits of unified temporal-semantic modeling.

Critically, DHGNN provides interpretability through meta-path–level attribution scores and explanatory subgraphs. This capability enables transparent reasoning across diverse contexts—from predicting collaborations in cold-start scenarios relying on semantic affiliations, to structurally dense scenarios prioritizing direct and multi-hop connections.

Despite these strengths, DHGNN still has several limitations that suggest directions for future research. First, DHGNN currently relies on a fixed repertoire of domain-specified meta-paths; incorporating automated or learnable path discovery could expose hidden semantics and broaden applicability. Second, the framework is confined to dyadic links; extending it to group-level or hypergraph formulations would allow richer modelling of multi-inventor collaborations. Third, generality must still be demonstrated across heterogeneous scientific and industrial networks and future work will explore more standardized, multi-domain benchmarks to evaluate its broader applicability.

Acknowledgements This work was supported by the National Natural Science Foundation of China under Grants 72171122.

Authors' contributions Xiaodong Xie: Responsible for research design, conducting experiments, and writing the paper. Jie Wu: Responsible for research proposal design and providing guidance on the paper. Mengjia Xiang: Contributed to writing the paper. Jianting Tang: Contributed to the research proposal design. Yongxiang Sheng: Contributed to the research proposal design and provided guidance on the paper.

Funding National Natural Science Foundation of China,72171122

Data availability Data will be made available on reasonable request.

Declarations

Competing interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Bar T, Leiponen A (2012) A measure of technological distance. Econ Lett 116(3):457–459. https://doi.org/10.1016/j.econlet.2012.04.030

Barros CDT, Mendonça MRF, Vieira AB, Ziviani A (2021) A survey on embedding dynamic graphs. ACM Comput Surv 55(1):1. https://doi.org/10.1145/3483595

Bekamiri H, Hain DS, Jurowetzki R (2024) PatentSBERTa: a deep NLP based hybrid model for patent distance and classification using augmented SBERT. Technol Forecast Soc Change 206:123536. https://doi.org/10.1016/j.techfore.2024.123536

Carayol N, Bergé L, Cassi L, Roux P (2019) Unintended triadic closure in social networks: the strategic formation of research collaborations between French inventors. J Econ Behav Organ 163:218–238. https://doi.org/10.1016/j.jebo.2018.10.009

Chen J, Wang X, Xu X (2022) Gc-lstm: graph convolution embedded LSTM for dynamic network link prediction. Appl Intell 52(7):7513–7528. https://doi.org/10.1007/s10489-021-02518-9

Chen L, Li Y, Deng X (2023) Multi-view learning-based heterogeneous network representation learning. Journal of King Saud University - Computer and Information Sciences 35(10):101855. https://doi.org/10.1016/j.jksuci.2023.101855

Cho H, Yu Y (2018) Link prediction for interdisciplinary collaboration via co-authorship network. Soc Netw Anal Min 8(1):25. https://doi.org/10.1007/s13278-018-0501-6

Crescenzi R, Nathan M, Rodríguez-Pose A (2016) Do inventors talk to strangers? On proximity and collaborative knowledge creation. Res Policy 45(1):177–194. https://doi.org/10.1016/j.respol.2015.07.003

Diemer A, Regan T (2022) No inventor is an island: social connectedness and the geography of knowledge flows in the US. Res Policy 51(2):104416. https://doi.org/10.1016/j.respol.2021.104416

Dong, Y., Chawla, N. V., & Swami, A. (2017). metapath2vec: Scalable Representation Learning for Heterogeneous Networks. (Paper presented at the Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada)

Fan, Y., Ju, M., Hou, S., Ye, Y., Wan, W., Wang, K., et al. (2021). Heterogeneous Temporal Graph Transformer: An Intelligent System for Evolving Android Malware Detection. (Paper presented at the Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, Virtual Event, Singapore)

Fan, Y., Ju, M., Zhang, C., & Ye, Y. (2022). Heterogeneous Temporal Graph Neural Network. Proceedings of the 2022 SIAM International Conference on Data Mining (SDM) (pp. 657–665). Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9781611977172.74.

Fritz C, De Nicola G, Kevork S, Harhoff D, Kauermann G (2023) Modelling the large and dynamically growing bipartite network of German patents and inventors. J R Stat Soc Ser A Stat Soc 186(3):557–576. https://doi.org/10.1093/jrsssa/qnad009

Funke T, Khosla M, Rathee M, Anand A (2023) Zorro: valid, sparse, and stable explanations in graph neural networks. IEEE Trans Knowl Data Eng 35(8):8687–8698. https://doi.org/10.1109/TKDE.2022.3201170

Goyal, P., Kamra, N., He, X., & Liu, Y. (2018). DynGEM: Deep Embedding Method for Dynamic Graphs. ArXiv, abs/1805.11273.

Grover, A., & Leskovec, J. (2016). node2vec: Scalable Feature Learning for Networks. (Paper presented at the Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, California, USA)

Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. (Paper presented at the Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, California, USA)

Hansheng, X., Yang, L., Jiang, W., Wei, Y., Hu, Y., & Lin, Y. (2021) 'Modeling Dynamic Heterogeneous Network for Link Prediction Using Hierarchical Attention with Temporal RNN' F. Hutter, K. Kersting, J. Lijffijt, & I. Valera Machine Learning and Knowledge Discovery in Databases. 2021//. Cham: Springer International Publishing, pp. 282–298.

Hu, Z., Dong, Y., Wang, K., & Sun, Y. (2020). Heterogeneous Graph Transformer. (Paper presented at the Proceedings of The Web Conference 2020, Taipei, Taiwan)

Huang Q, Yamada M, Tian Y, Singh D, Chang Y (2023) GraphLIME: local interpretable model explanations for graph neural networks. IEEE Trans Knowl Data Eng 35(7):6968–6972. https://doi.org/10.1109/TKDE.2022.3187455

Jaffe AB, de Rassenfosse G (2017) Patent citation data in social science research: overview and best practices. J Assoc Inf Sci Technol 68(6):1360–1374. https://doi.org/10.1002/asi.23731

Khoshraftar S, An A (2024) A survey on graph representation learning methods. ACM Trans Intell Syst Technol 15(1):Article 19. https://doi.org/10.1145/3633518

Kipf, T., & Welling, M. (2016). Semi-Supervised Classification with Graph Convolutional Networks. ArXiv, abs/1609.02907.

Lathabai HH, Nandy A, Singh VK (2022) Institutional collaboration recommendation: an expertise-based framework using NLP and network analysis. Expert Syst Appl 209:118317. https://doi.org/10.1016/j.eswa.2022.118317

Liben-Nowell, D., & Kleinberg, J. (2003). The link prediction problem for social networks. (Paper presented at the Proceedings of the twelfth international conference on Information and knowledge management, New Orleans, LA, USA)

Lopes, G. R., Moro, M. M., Wives, L. K., & de Oliveira, J. P. M. (2010) 'Collaboration Recommendation on Academic Social Networks'. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 190–199.

Luo, D., Cheng, W., Xu, D., Yu, W., Zong, B., Chen, H., et al. (2020a). Parameterized explainer for graph neural network. (Paper presented at the Proceedings of the 34th International Conference on Neural Information Processing Systems, Vancouver, BC, Canada)

Luo, W., Zhang, H., Yang, X., Bo, L., Yang, X., Li, Z., et al. (2020b). Dynamic Heterogeneous Graph Neural Network for Real-time Event Prediction. (Paper presented at the Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Virtual Event, CA, USA)

Mikolov, T., Chen, K., Corrado, G. S., & Dean, J. (2013) 'Efficient Estimation of Word Representations in Vector Space' International Conference on Learning Representations.

Pareja A, Domeniconi G, Chen J, Ma T, Suzumura T, Kanezashi H et al (2020) EvolveGCN: evolving graph convolutional networks for dynamic graphs. Proceedings of the AAAI Conference on Artificial Intelligence. https://doi.org/10.1609/aaai.v34i04.5984

Pavlov, M., & Ichise, R. (2007). Finding experts by link prediction in co-authorship networks. (Paper presented at the Proceedings of the 2nd International Conference on Finding Experts on the Web with Semantics - Volume 290, Busan, Korea)

Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). DeepWalk: online learning of social representations. (Paper presented at the Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, New York, New York, USA)

Ren X (2024) Link prediction using extended neighborhood based local random walk in multilayer social networks. Journal of King Saud University - Computer and Information Sciences 36(2):101931. https://doi.org/10.1016/j.jksuci.2024.101931

Rossi, E., Chamberlain, B. P., Frasca, F., Eynard, D., Monti, F., & Bronstein, M. M. (2020). Temporal Graph Networks for Deep Learning on Dynamic Graphs. ArXiv, abs/2006.10637.

Sankar, A., Wu, Y., Gou, L., Zhang, W., & Yang, H. (2020). DySAT: Deep Neural Representation Learning on Dynamic Graphs via Self-Attention Networks. (Paper presented at the Proceedings of the 13th International Conference on Web Search and Data Mining, Houston, TX, USA)

Schlichtkrull, M., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., & Welling, M. (2018) 'Modeling Relational Data with Graph Convolutional Networks' A. Gangemi, R. Navigli, M.-E. Vidal, P. Hitzler, R. Troncy, L. Hollink, et al. The Semantic Web. 2018//. Cham: Springer International Publishing, pp. 593–607.

Sun X, Lin H, Xu K, Ding K (2015) How we collaborate: characterizing, modeling and predicting scientific collaborations. Scientometrics 104(1):43–60. https://doi.org/10.1007/s11192-015-1597-3

Tran, T. Q., Fukuchi, K., Akimoto, Y., & Sakuma, J. (2020). Statistically Significant Pattern Mining with Ordinal Utility. (Paper presented at the Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Virtual Event, CA, USA)

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. Advances in neural information processing systems, 30.

Wang Y, Sun Y, Liu Z, Sarma SE, Bronstein MM, Solomon JM (2019c) Dynamic graph CNN for learning on point clouds. ACM Trans Graph 38(5):Article 146. https://doi.org/10.1145/3326362

Wang P, Shu J, Liu L (2024) Heterogeneous network link prediction based on network schema and cross-neighborhood attention. Journal of King Saud University - Computer and Information Sciences 36(7):102154. https://doi.org/10.1016/j.jksuci.2024.102154

Wang, X., & Sukthankar, G. (2013). Link prediction in multi-relational collaboration networks. (Paper presented at the Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, Niagara, Ontario, Canada)

Wang, X., He, X., Wang, M., Feng, F., & Chua, T.-S. (2019a). Neural Graph Collaborative Filtering. (Paper presented at the Proceedings

of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, Paris, France)

Wang, X., Ji, H., Shi, C., Wang, B., Ye, Y., Cui, P., et al. (2019b). Heterogeneous Graph Attention Network. (Paper presented at the The World Wide Web Conference, San Francisco, CA, USA)

World Intellectual Property, O. (2024) 'World Intellectual Property Indicators 2024'. Geneva: WIPO. Available at: https://www.wipo.int/publications/en/details.jsp?id=4759.

Xiao W, Zhao H, Zheng VW, Song Y (2024) Meta-path based proximity learning in heterogeneous information networks. Data Min Knowl Discov 39(1):8. https://doi.org/10.1007/s10618-024-01076-1

Xing Z, Song R, Teng Y, Xu H (2022) DynHEN: a heterogeneous network model for dynamic bipartite graph representation learning. Neurocomputing 508:47–57. https://doi.org/10.1016/j.neucom.2022.08.050

Xiong, J., Zareie, A., & Sakellariou, R. (2025). A Survey of Link Prediction in Temporal Networks. ArXiv, abs/2502.21185.

Xu, D., Ruan, C., Körpeoglu, E., Kumar, S., & Achan, K. (2020). Inductive Representation Learning on Temporal Graphs. ArXiv, abs/2002.07962.

Xue HS, Yang LW, Jiang W, Wei Y, Hu Y, Lin Y (2021) Modeling dynamic heterogeneous network for link prediction using hierarchical attention with temporal RNN. (Paper presented at the Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases—ECML PKDD 2020, Ghent, Belgium.) https://doi.org/10.1007/978-3-030-67658-2_17

Ying R, Bourgeois D, You J, Zitnik M, Leskovec J (2019b) GNNexplainer: generating explanations for graph neural networks. Adv Neural Inf Process Syst 32:9240–9251

Ying, R., Bourgeois, D., You, J., Zitnik, M., & Leskovec, J. (2019a). GNNExplainer: Generating Explanations for Graph Neural Networks. (Paper presented at the Advances in Neural Information Processing Systems (NeurIPS))

Yu Q, Long C, Lv Y, Shao H, He P, Duan Z (2014) Predicting co-author relationship in medical co-authorship networks. PLoS ONE 9(7):e101214. https://doi.org/10.1371/journal.pone.0101214

Zhang Z, Zhang Z, Wang X, Qin Y, Qin Z, Zhu W (2023) Dynamic heterogeneous graph attention neural architecture search. Proceedings of the AAAI Conference on Artificial Intelligence 37(9):11307–11315. https://doi.org/10.1609/aaai.v37i9.26338

Zhang, M., & Chen, Y. (2018). Link prediction based on graph neural networks. (Paper presented at the Proceedings of the 32nd International Conference on Neural Information Processing Systems, Montréal, Canada)

Zhu, J., & Yaseen, A. (2022). A Recommender for Research Collaborators Using Graph Neural Networks (Original Research). Frontiers in Artificial Intelligence, Volume 5 - 2022. https://doi.org/10.3389/frai.2022.881704.