Journal of King Saud University Computer and Information Sciences (2026) 38:66
https://doi.org/10.1007/s44443-025-00461-0

ORIGINAL PAPER

QLGAN: a quantum-lineage graph attention network for temporal knowledge graph entity alignment

Jia Li[1] · Yuxi Ma[2] · Lingzhong Meng[2]

Received: 28 August 2025 / Accepted: 30 December 2025

Abstract
Temporal entity alignment, the task of identifying equivalent entities across evolving knowledge graphs (KGs), is a critical yet challenging problem. Existing methods often struggle to holistically model the complex interplay between structural topology and temporal dynamics while also failing to capture long-range dependencies encoded in multi-hop relational paths. To address these limitations, we propose the **Quantum-Lineage Graph Attention Network (QLGAN)**, a novel quantum-classical hybrid model for temporal entity alignment. QLGAN uniquely integrates a path encoder directly within its attention mechanism, which enriches neighbor representations with multi-hop relational semantics before aggregation. Furthermore, it leverages a Variational Quantum Circuit (VQC) to perform powerful non-linear feature transformations, enabling a more effective fusion of complex spatiotemporal features. The model's alignment process is based on a hybrid similarity framework that combines quantum-inspired metrics with structurally aware temporal propagation, followed by an optimal transport alignment algorithm. Extensive experiments on benchmark datasets validate the effectiveness of our approach in learning robust and contextually-rich representations for temporal knowledge graphs.

Keywords Temporal entity alignment · Graph attention network · Variational quantum circuit

1 Introduction

Entity Alignment (EA) (Wang et al. 2018) serves as a critical bridge for knowledge integration, facilitating the consolidation of disparate Knowledge Graphs (KGs) by identifying entities that represent the same real-world concept. Traditionally, EA research has focused on static relational structures, aiming to align entities based on their invariant neighborhood topologies. However, as knowledge-driven systems are increasingly operationalized in dynamic, real-world environments, the limitations of static snapshots have become evident. Real-world facts are inherently transient and context-dependent; they evolve, expire, and transition over time, necessitating a paradigm shift from static KGs to Temporal Knowledge Graphs (TKGs).

✉ Lingzhong Meng
lingzhong@iscas.ac.cn

1 College of Computer Science and Technology, Zhengzhou University of Light Industry, Zhengzhou 450000, China

2 Institute of Software Chinese Academy of Sciences, Beijing 100190, China

In the context of TKGs, the alignment task—Temporal Entity Alignment (TEA) (Xu et al. 2021)—demands a sophisticated synthesis of both spatial connectivity and temporal dynamics. Unlike its static counterpart, TEA requires models to reconcile entities by deciphering their historical trajectories and evolutionary patterns. While existing methods, such as TEA-GNN (Xu et al. 2021), TREA (Xu et al. 2022), STEA (Cai et al. 2022), and DualMatch (Liu et al. 2023), have made significant progress by incorporating time-aware mechanisms, they still face two core challenges that limit their efficacy in complex, real-world scenarios.

First, the exploration of temporal information remains insufficient. Most current models treat timestamps merely as discrete features or independent time intervals, failing to fully leverage the rich, underlying patterns within temporal data, such as continuity, periodicity, and event sequences. Consequently, these models struggle to capture complex evolutionary patterns across time steps, especially when an entity's local graph structure is disconnected from known alignment seeds.

Second, existing models struggle to capture complex, long-range dependencies and multi-hop interactions. Traditional

Published online: 15 January 2026

graph neural networks (GNNs) primarily rely on local neighborhood aggregation, which overlooks the rich semantic information embedded in multi-hop relational lineages. This limitation becomes particularly severe in large, sparse knowledge graphs where critical contextual clues for disambiguation often reside in entities and relations far beyond the immediate neighborhood. Furthermore, these methods treat different reasoning lineages as isolated paths, failing to model their complex interplay. They cannot "entangle" or fuse information from seemingly disparate lineages, which may, in fact, provide mutually reinforcing or contradictory evidence for a potential alignment.

To illustrate these challenges more concretely, we consider a classic scenario within temporal entity alignment. Suppose we have an entity "Dr. Evelyn Reed" in TKG_1 and another entity "E. Reed" in TKG_2. Both entities are linked to similar static attributes, such as being affiliated with "BioSynth Labs" and publishing on "Quantum Biology", suggesting a potential alignment based on a local, static view. However, existing methods often fail to adequately leverage temporal information and long-range dependencies, leading to two critical limitations.

First, the insufficient exploration of temporal information can lead to erroneous alignments. For instance, TKG_1 might contain the temporal fact **(Dr. Evelyn Reed, published_in, Nature Physics, 2023)**, while TKG_2 holds a conflicting fact **(E. Reed, left_academia, 2021)**. A model that treats timestamps as mere discrete features would struggle to process this contradictory temporal sequence, potentially leading to an incorrect match.

Second, the inability to capture multi-hop interactions poses a significant challenge, especially in large, sparse knowledge graphs. The correct match for "Dr. Evelyn Reed" might be "Evelyn R." in TKG_2, but this link may only be discoverable through a long-range relational lineage, such as a shared co-author at MIT during a specific period. Traditional GNNs, which rely on local neighborhood aggregation, are likely to miss this critical, distant connection. Furthermore, these methods treat different reasoning lineages as isolated paths and fail to "entangle" or fuse information from seemingly disparate paths that may, in fact, provide mutually reinforcing evidence for a potential alignment.

To address these challenges, we propose a novel quantum-classical hybrid framework designed specifically for temporal entity alignment: the Quantum-Lineage Graph Attention Network (QLGAN). QLGAN fundamentally enhances the model's ability to holistically represent spatiotemporal features and capture long-range dependencies. To tackle the insufficient exploration of temporal information, QLGAN employs a multi-view encoder architecture combined with a Variational Quantum Circuit (VQC). The VQC leverages the high-dimensional Hilbert space to perform powerful non-linear feature transformations, enabling it to learn highly complex, entangled spatiotemporal features that are often inaccessible to classical models. To resolve the issue of long-range dependencies, QLGAN uniquely integrates a Lineage encoder directly within its attention mechanism. This design explicitly encodes relational paths between entities and fuses this path-level semantic information with neighbor features before the attention-based aggregation step. This "path-aware" aggregation ensures that the model can capture context beyond the immediate neighborhood, generating more comprehensive and accurate entity representations. The alignment is performed through a hybrid similarity framework that combines quantum-inspired metrics with structurally-aware temporal propagation, followed by an optimal transport alignment algorithm. Through these innovations, QLGAN can more deeply mine the intrinsic patterns of TKGs, effectively reduce the reliance on pre-aligned data, and demonstrate superior performance in complex alignment scenarios.

Our main contributions are as follows:

1. We present QLGAN, a novel quantum-classical hybrid model that addresses the core challenges of temporal entity alignment. It cohesively integrates a path encoder into a quantum-enhanced attention network, explicitly designed to handle the complexities of evolving knowledge graphs.
2. To tackle the **insufficient exploration of temporal information**, QLGAN employs a multi-view encoder architecture combined with a **Variational Quantum Circuit (VQC)**. The VQC leverages the high-dimensional Hilbert space to perform powerful non-linear feature transformations, enabling the model to learn complex, entangled spatiotemporal features that are often inaccessible to classical models.
3. To resolve the issue of **long-range dependencies**, QLGAN integrates a unique attention mechanism. This design explicitly encodes relational paths between entities, fusing this multi-hop semantic information with local neighbor features before the aggregation step. This "path-aware" aggregation ensures that the model captures context beyond the immediate neighborhood, generating more comprehensive and accurate entity representations.
4. We design a hybrid similarity framework that combines quantum-inspired metrics with structurally-aware temporal propagation, followed by an optimal transport alignment algorithm. This framework enables the model to effectively fuse deep-learned semantics with explicit structural-temporal information.

2 Related work

2.1 Entity alignment

Entity Alignment (EA) serves as a cornerstone task in knowledge fusion, aimed at identifying entities that represent the same real-world object across disparate Knowledge Graphs (KGs). The field evolved primarily along two axes regarding the nature of the underlying data: static and temporal alignment.

Static entity alignment The most extensively studied paradigm assumed that KG facts remained fixed over time. Early pioneering works, inspired by multilingual word embeddings, introduced translation-based models such as **MTransE** (Chen et al. 2017) and its variant **ITransE** (Zhu et al. 2017). These methods embedded entities into a shared vector space and treated relations as translation vectors. However, standard translation approaches often struggled to capture complex graph topologies and geometric consistencies. To mitigate this, **RREA** (Mao et al. 2020) proposed relation-specific reflection matrices, while **HyperKA** (Sun et al. 2020a) introduced hyperbolic embeddings to better capture hierarchical structures.

The advent of Graph Neural Networks (GNNs) marked a significant paradigm shift. **GCN-Align** (Wang et al. 2018) pioneered the use of GNNs to aggregate neighborhood features. This concept was subsequently refined by models focusing on wider structural contexts, such as **AliNet** (Sun et al. 2020b), which aggregated multi-hop neighbors, and **Dual-AMN** (Mao et al. 2021), which modeled intra- and cross-graph interactions. Further efforts focused on enriching entity representations with complex semantics: **MRAEA** (Mao et al. 2020) and **MuGNN** (Cao et al. 2019) utilized meta-relation semantics and multi-channel encoders, respectively, to learn alignment-oriented embeddings from diverse perspectives. Similarly, **KE-GCN** (Yu et al. 2021) combined GCNs with knowledge graph completion objectives to enhance task-specific representations.

Recognizing that structural information alone was often insufficient, subsequent research enriched representations with attributes (**AttrGNN** (Liu et al. 2020), **AttrE** (Trsedya et al. 2019)), textual descriptions (**HMAN** (Yang et al. 2020)), and visual data (**EVA** (Liu et al. 2021), **MSNEA** (Chen et al. 2022)). Finally, addressing scalability, **ClusterEA** (Gao et al. 2022) and the non-neural framework **LightEA** (Mao et al. 2022) were developed to handle large-scale graph matching efficiently. Despite these advancements, the static assumption fundamentally limited these models when applied to dynamic, real-world KGs where knowledge evolved.

Temporal entity alignment To address the limitations of static models, Temporal Entity Alignment (TEA) emerged as a critical frontier. Directly applying static models to Temporal Knowledge Graphs (TKGs) typically yielded suboptimal results, as they failed to account for the evolution of entities and relations, particularly in the presence of noisy temporal data (Cai et al. 2022). Consequently, specialized models were developed to explicitly model the temporal dimension.

Early works such as **TEA-GNN** (Xu et al. 2021) incorporated time-aware attention mechanisms to differentiate neighbors based on temporal relevance. **TREA** (Xu et al. 2022) advanced this by combining temporal relational attention with a sequential regularizer to better model event sequences. In a different vein, **STEA** (Cai et al. 2022) posited that a unified temporal representation existed in TKGs and proposed a simplified GNN with an efficient temporal matching mechanism. More recently, researchers focused on disentangling temporal and structural features. **DualMatch** (Liu et al. 2023) introduced a dual-encoder architecture to separately capture temporal and relational information, while **TSM-GNN** (Jia et al. 2024) utilized relation-entity graphs to approximate embeddings before refining them with temporal information.

A recent surge in interest concerned Large Language Model (LLM)-based methods, which leveraged the generative and reasoning capabilities of pre-trained models. **ChatEA** (Jiang et al. 2024a) introduced a KG-code module to translate KG structures into LLM-understandable formats, adopting a two-stage strategy of retrieval and multi-step reasoning. Similarly, **MM-ChatAlign** (Jiang et al. 2024b) leveraged the visual reasoning capabilities of multimodal LLMs to enhance semantic alignment. While these models highlighted the importance of dedicated temporal and semantic modeling, effectively fusing complex spatiotemporal features with long-range dependencies remained an open challenge.

2.2 Quantum graph neural networks

The integration of quantum computing principles with machine learning gave rise to Quantum Machine Learning (QML) (Biamonte et al. 2017), a field leveraging phenomena such as superposition and entanglement to address problems intractable for classical computation. This approach was particularly promising for graph data (Beer et al. 2023), where the combinatorial complexity of connections necessitated robust modeling, leading to the development of Quantum Graph Neural Networks (QGNNs) (Innan et al. 2024).

A primary challenge lay in the effective mapping of quantum principles (OZHIGOV 2022) onto classical or hybrid architectures. Research largely focused on determin-

ing which components of a traditional GNN benefited most from "quantization". Two dominant strategies emerged:

- **Hybrid Quantum-Classical Models** (Liu et al. 2025): Several studies (Ray et al. 2024; Tüysüz et al. 2021) successfully replaced the terminal layers of GNNs—typically responsible for classification or prediction—with Variational Quantum Circuits (VQC). In this paradigm, the quantum component functioned as a highly expressive, non-linear transformation layer within a classical backbone.
- **Quantum Data Encoding** (Rath and Date 2024): An alternative approach (Qi et al. 2024) involved encoding the graph structure itself (e.g., the adjacency matrix) directly into a quantum state. Techniques such as "data re-uploading" were employed to leverage the high-dimensional Hilbert space for clearer structural representation. Additionally, methods utilizing quantum walks (Xu et al. 2024) were explored to enhance the capability of GNNs in capturing long-range dependencies (Pomarico et al. 2021).

Despite these advancements, significant limitations persisted in current QGNN research. Existing methods largely focused on static graphs or node classification tasks, often neglecting the complex interplay between static structure and dynamic temporal evolution. Furthermore, most quantum-inspired graph models struggled to explicitly capture multi-hop dependencies in sparse temporal settings. This necessitated the development of a unified framework capable of integrating quantum expressivity with temporal lineage modeling.

3 Methodology: Quantum-Lineage Graph Attention Network (QLGAN)

We propose the **Quantum-Lineage Graph Attention Network (QLGAN)**, an end-to-end framework designed to learn robust entity representations for Temporal Knowledge Graph (TKG) alignment. As illustrated in Fig. 1, QLGAN addresses the expressivity bottlenecks of standard GNNs by synergizing classical structural encoding with quantum-inspired feature transformations. The architecture processes entity features $h_i^{(l)}$ and time-relation embeddings (v_t, v_r) through a multi-stage pipeline comprising three core modules: (1) a Lineage-Entangled Encoder, (2) a Quantum-Enhanced Mapping module, and (3) a Global Gating fusion mechanism.

3.1 Problem formulation

Definition 1 (Temporal Knowledge Graph) A Temporal Knowledge Graph is defined as $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{Q}\}$, where $\mathcal{E}$ is the set of entities, $\mathcal{R}$ is the set of relations, and $\mathcal{T}$ is the set of timestamps. The core of the TKG is the set of temporal quadruples $\mathcal{Q} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E} \times \mathcal{T}$. A quadruple $(h, r, t, \tau) \in \mathcal{Q}$ represents a relational fact existing between head entity h and tail entity t via relation r at time τ.

Definition 2 (Temporal Entity Alignment Task (Xu et al. 2021)) Given two heterogeneous TKGs, $\mathcal{G}_1 = (\mathcal{E}_1, \mathcal{R}_1, \mathcal{T}_1, \mathcal{Q}_1)$ and $\mathcal{G}_2 = (\mathcal{E}_2, \mathcal{R}_2, \mathcal{T}_2, \mathcal{Q}_2)$, the objective is to discover a set of equivalent entity pairs $\mathcal{A} = \{(e_{i,1}, e_{i,2}) \mid e_{i,1} \in \mathcal{E}_1, e_{i,2} \in \mathcal{E}_2, e_{i,1} \equiv e_{i,2}\}$ that refer to the same entity in different sources.

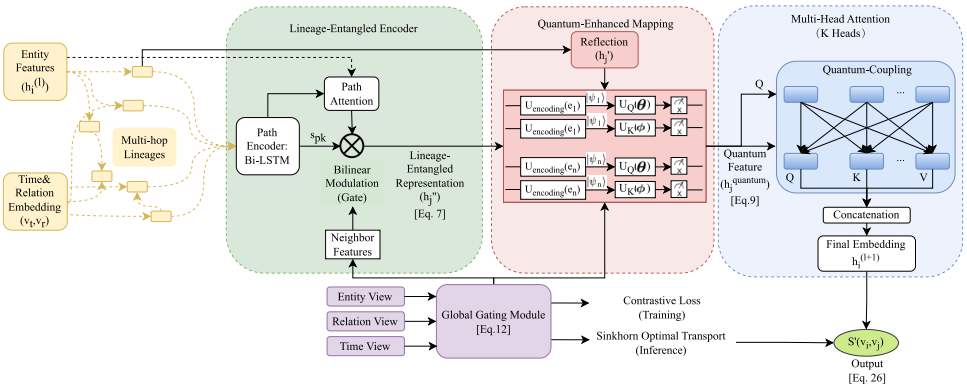

Fig. 1 Overview of the QLGAN Layer. The architecture integrates a Lineage-Entangled Encoder (left) to capture multi-hop structural history, a Quantum-Enhanced Mapping module (center) to project features into a Hilbert space for non-linear transformation, and a Quantum-Coupling Attention mechanism (right) for final aggregation. The Global Gating Module (bottom) fuses multi-view information for the final alignment objective

3.2 Lineage-entangled entity encoding

The Quantum-Enhanced Graph Learning Attentional Network (QLGAN) operates on an entity graph $\mathcal{G} = (\mathcal{E}, \mathcal{R})$. A single QLGAN layer aims to compute a superior representation $\mathbf{h}_i^{(l+1)}$ for a central entity e_i by aggregating information from its structural neighborhood $\mathcal{N}(i)$. Our design fundamentally addresses the limitations of conventional GNNs by explicitly encoding multi-hop relational semantics directly into the neighbor features, moving beyond simple single-hop message passing.

3.2.1 Lineage-entangled feature modulation

Unlike multi-hop GNNs that typically perform message passing over an expanded graph or use paths solely for weighting, the QLGAN introduces Lineage-Entangled Feature Modulation. This mechanism treats the path semantics as an operator that transforms the neighbor's feature vector, creating a dependency (or "entanglement") between the entity representation and its relational history.

Step 1: ***Contextual Path Encoding***

For a central entity e_i and a neighbor $e_j \in \mathcal{N}(i)$, we first identify all distinct multi-hop **lineages** (paths) $\mathcal{P}_{ij} = \{p_1, p_2, \ldots\}$ connecting them. Each lineage $p_k = (r_1, \ldots, r_L)$ is a sequence of relations. We utilize a dedicated Path Encoder (e.g., an attentive Bi-LSTM) to capture the ordered semantic context of p_k, yielding a path embedding $\mathbf{s}_{p_k}$.

We then consolidate the multiple path embeddings into a single cohesive lineage feature vector $\mathbf{f}_{\text{lineage}}(i, j)$ using a dedicated **path attention mechanism** $\text{Attn}_{\text{path}}$, which is conditioned on the central entity e_i's current state $\mathbf{h}_i^{(l)}$:

$$\beta_k = \frac{\exp(\mathbf{a}_{\text{path}}^\top \cdot [\mathbf{s}_{p_k} \| \mathbf{h}_i^{(l)}])}{\sum_{p_m \in \mathcal{P}_{ij}} \exp(\mathbf{a}_{\text{path}}^\top \cdot [\mathbf{s}_{p_m} \| \mathbf{h}_i^{(l)}])} \tag{1}$$

The aggregated lineage feature is:

$$\mathbf{f}_{\text{lineage}}(i, j) = \sum_{p_k \in \mathcal{P}_{ij}} \beta_k \mathbf{s}_{p_k} \tag{2}$$

Step 2: ***Entangled Feature Fusion***

Crucially, the aggregated lineage feature $\mathbf{f}_{\text{lineage}}(i, j)$ is then used to **modulate** the neighbor's current feature $\mathbf{h}_j^{(l)}$ through a **Gated Fusion Unit**, rather than a simple concatenation and projection. This mechanism enforces the 'entanglement'—the representation of e_j becomes inherently dependent on the specific lineage connecting it to e_i.

We propose a **Bilinear Modulation and Gating** mechanism for the Lineage-Entangled Representation $\mathbf{h}_j''$:

$$\mathbf{g}_{ij} = \sigma(\mathbf{W}_g[\mathbf{h}_j^{(l)} \| \mathbf{f}_{\text{lineage}}(i, j)]) \tag{3}$$

$$\mathbf{z}_{ij} = \mathbf{h}_j^{(l)} \odot (\mathbf{W}_z \mathbf{f}_{\text{lineage}}(i, j)) \tag{4}$$

$$\mathbf{h}_j'' = \mathbf{g}_{ij} \odot \mathbf{z}_{ij} + (\mathbf{1} - \mathbf{g}_{ij}) \odot \mathbf{h}_j^{(l)} \tag{5}$$

where $\mathbf{g}_{ij}$ is a lineage-specific gate vector, σ is the sigmoid function, $\odot$ denotes the Hadamard product, and $\mathbf{W}_g$, $\mathbf{W}_z$ are trainable projection matrices. The term $\mathbf{z}_{ij}$ captures a **bilinear interaction** between the neighbor's feature and the path feature, allowing the lineage to selectively transform $\mathbf{h}_j^{(l)}$. The gate $\mathbf{g}_{ij}$ then determines the extent to which this transformation is applied, resulting in the final **lineage-entangled representation** $\mathbf{h}_j''$.

This design fundamentally differs from prior work: $\mathbf{h}_j''$ is not merely an updated feature, but a composite vector where the neighbor's identity ($\mathbf{h}_j^{(l)}$) is intrinsically **conditioned** and **transformed** by the relational context ($\mathbf{f}_{\text{lineage}}$) relative to the central node e_i.

3.2.2 Quantum-inspired feature mapping

The QLGAN employs a two-stage **Quantum-Enhanced Feature Mapping** on the input feature $\mathbf{h}_j''$ to significantly boost feature expressiveness and encode **relational context** into a high-dimensional quantum state space.

Stage 1: ***Relational Reflection***

We first apply a **Relational Reflection Operation** to incorporate local structural information. This operation, inspired by geometric embedding models (Mao et al. 2020) in Knowledge Graphs, transforms $\mathbf{h}_j''$ by reflecting it across the hyperplane defined by the embedding of the direct connecting relation $\mathbf{v}_{r_{ij}}$. This ensures that the transformed feature $\mathbf{h}_j'$ is sensitive to the type of relationship r_{ij} connecting the entities:

$$\mathbf{h}_j' = \mathbf{h}_j'' - 2 \cdot (\mathbf{h}_j'' \cdot \mathbf{v}_{r_{ij}}) \cdot \mathbf{v}_{r_{ij}} \tag{6}$$

The result $\mathbf{h}_j'$ is a relation-aware vector, ready for quantum encoding

Stage 2: ***Variational Quantum Circuit (VQC) Mapping***

Second, the transformed feature $\mathbf{h}_j'$ is mapped by a **Variational Quantum Circuit (VQC)** (Verdon et al. 2019), which serves as a highly complex, nonlinear feature map $\Phi_{\text{VQC}} : \mathbb{R}^n \to \mathbb{R}^m$. This stage is critical for introducing **quantum advantages**:

- **High-Dimensional Quantum State Space**: The VQC implicitly projects the classical data into an exponentially

large Hilbert space (2^N for N qubits). This projection allows for separation of data points that are intractable in classical Euclidean space, leveraging the concept of a **quantum kernel** for enhanced separability.

- **Entanglement Encoding**: The unitary operations within $U(\boldsymbol{\theta})$ (particularly the entangling gates) generate complex **quantum entanglement** between the qubit states. This entanglement fundamentally encodes rich, non-local correlations and higher-order feature interactions, which are difficult to model efficiently with standard classical neural networks.
- **Universal Expressiveness**: Under certain conditions, VQCs can be shown to approximate arbitrary unitary transformations, offering a degree of **universal expressiveness** potentially surpassing bounded-depth classical circuits.

The final **quantum-enhanced representation** $\mathbf{h}_j^{\text{quantum}}$ is extracted by measuring the expectation values of a set of Pauli-Z operators ($\{Z_k\}_{k=1}^m$):

$$\mathbf{h}_j^{\text{quantum}} = \Phi_{\text{VQC}}(\mathbf{h}'_j) = (\langle\psi_L|Z_1|\psi_L\rangle, \ldots, \langle\psi_L|Z_m|\psi_L\rangle)^T \tag{7}$$

where $|\psi_L\rangle = U(\boldsymbol{\theta})|\psi_0\rangle$ is the state after evolution through the trainable circuit $U(\boldsymbol{\theta})$ from the initial encoded state $|\psi_0\rangle = \text{AngleEmbedding}(\mathbf{h}'_j)$. Here, $\boldsymbol{\theta}$ represents the trainable circuit parameters.

The resulting $\mathbf{h}_j^{\text{quantum}}$ thus carries features enriched not only by the local relational context but also by complex, high-order quantum correlations.

3.2.3 Quantum-inspired entanglement-aware attention mechanism

To effectively fuse the classical representation $\mathbf{h}_i^{(l)}$ of entity e_i with the **quantum-enhanced features** $\mathbf{h}_j^{\text{quantum}}$ from its neighbors, we design the **Quantum-Inspired Entanglement-Aware Attention Mechanism**. This layer adopts a multi-head attention mechanism, where the resulting attention coefficient ϵ_{ij} is reinterpreted as a **Quantum-Correlation Coupling Weight**, designed to capture the propagation of potential non-classical associations (e.g., quantum entanglement) across the graph.

The raw attention score e_{ij} between node e_i and its neighbor e_j is computed by applying an inner product between a learnable attention vector $\mathbf{a}_{\text{att}}$ and the concatenated, linearly transformed features.

Specifically, we first linearly map the input features from dimension D to D' using a shared weight matrix $\mathbf{W} \in \mathbb{R}^{D' \times D}$. The raw attention score e_{ij} is then calculated as:

$$e_{ij} = \mathbf{a}_{\text{att}}^{\top} \cdot \text{LeakyReLU}\left(\mathbf{W}\mathbf{h}_i^{(l)} \| \mathbf{W}\mathbf{h}_j^{\text{quantum}}\right) \quad \text{for } e_j \in \mathcal{N}(i) \tag{8}$$

where $\mathbf{a}_{\text{att}}$ is a learnable attention vector, $\mathcal{N}(i)$ denotes the set of immediate neighbors of entity e_i, and $\|$ denotes the vector concatenation operation. The LeakyReLU activation is introduced to enhance non-linearity, following standard Graph Attention Network (GAT) practices.

The raw scores e_{ij} are subsequently normalized to obtain the final **Quantum-Correlation Coupling Weight** ϵ_{ij}. We employ a **Sparse Softmax** function to ensure the attention distribution is computed exclusively over the set of neighbors $\mathcal{N}(i)$:

$$\epsilon_{ij} = \text{softmax}_j(e_{ij}) = \frac{\exp(e_{ij})}{\sum_{k \in \mathcal{N}(i)} \exp(e_{ik})} \tag{9}$$

This weight ϵ_{ij} quantifies the contribution of neighbor e_j's quantum-enhanced information to the new representation of the target node e_i.

The intermediate aggregated feature $\mathbf{h}_i^{\text{QLGAN}}$ for node e_i is obtained by accumulating the quantum-enhanced features $\mathbf{h}_j^{\text{quantum}}$ from all neighbors $e_j \in \mathcal{N}(i)$, weighted by the normalized coupling coefficients ϵ_{ij}:

$$\mathbf{h}_i^{\text{QLGAN}} = \sum_{j \in \mathcal{N}(i)} \epsilon_{ij} \mathbf{h}_j^{\text{quantum}} \tag{10}$$

To improve the stability and expressive capacity of the model, we utilize a **Multi-Head Attention Mechanism**. Assuming K independent attention heads, each executes the aggregation process described above, yielding K independent aggregated features $\mathbf{h}_i^{\text{QLGAN},k}$. The final layer output $\mathbf{h}_i^{(l+1)}$ is achieved by concatenating the results of all attention heads, followed by a non-linear activation σ:

$$\mathbf{h}_i^{(l+1)} = \sigma\left(\|_{k=1}^{K} \mathbf{h}_i^{\text{QLGAN},k}\right) \tag{11}$$

where σ is an activation function, such as the Exponential Linear Unit (ELU). With concatenation, the input dimension for the subsequent layer will be $K \cdot D'$.

3.3 Multi-view feature fusion and gating

We integrate entity-centric ($\mathbf{H}^{\text{ent}}$), relation-centric ($\mathbf{H}^{\text{rel}}$), and temporal ($\mathbf{H}^{\text{time}}$) views, each processed by a separate QLGAN encoder.

The **local neighborhood feature** $\mathbf{h}_i^{\text{local}}$ for entity e_i:

$$\mathbf{h}_i^{\text{local}} = \left[\mathbf{H}^{\text{ent}}[i] \| \left(\frac{\mathbf{H}^{\text{rel}}[i] + \mathbf{H}^{\text{time}}[i]}{2}\right)\right] \tag{12}$$

We introduce a **learnable global proxy vector** $\mathbf{v}_{\text{proxy}}$ to derive a global-aware feature $\mathbf{h}_{\text{proxy}}$:

$$\mathbf{h}_{\text{proxy}} = \text{MLP}(\mathbf{h}_i^{\text{local}} \odot \mathbf{v}_{\text{proxy}}) \tag{13}$$

A **Gated Fusion Mechanism** adaptively combines the local and global features:

$$\mathbf{g} = \sigma(\mathbf{W}_{\text{gate}}\mathbf{h}_{\text{proxy}} + \mathbf{b}_{\text{gate}}) \tag{14}$$

$$\mathbf{h}_{\text{final},i} = \mathbf{g} \odot \mathbf{h}_i^{\text{local}} + (\mathbf{1} - \mathbf{g}) \odot \mathbf{h}_{\text{proxy}} \tag{15}$$

The vector $\mathbf{h}_{\text{final},i}$ is the final entity embedding.

3.4 Training objective: margin-based alignment loss

We employ a bidirectional, margin-based contrastive loss (Liu et al. 2023) over the set of pre-aligned entity pairs $\mathcal{L}_{\text{train}} = \{(e_i, e_j)\}$. The distance metric is $d(e_i, e_j) = \|\mathbf{h}_{\text{final},i} - \mathbf{h}_{\text{final},j}\|_2^2$.

Bidirectional contrastive loss The loss aligns e_i from the source KG ($\mathcal{K}_S$) towards $\mathcal{K}_T$, and e_j from $\mathcal{K}_T$ towards $\mathcal{K}_S$. $\gamma > 0$ is the margin and $\lambda > 0$ is the scaling factor.

1. **Source-to-Target Alignment**:

$$\mathcal{L}_{i\to\cdot} = \log\left(1 + \sum_{e'_k \in \mathcal{E}_T \setminus \{e_j\}} \exp\left(\lambda(d(e_i, e_j) - d(e_i, e'_k) + \gamma)\right)\right) \tag{16}$$

2. **Target-to-Source Alignment**:

$$\mathcal{L}_{\cdot\to j} = \log\left(1 + \sum_{e'_l \in \mathcal{E}_S \setminus \{e_i\}} \exp\left(\lambda(d(e_j, e_i) - d(e_j, e'_l) + \gamma)\right)\right) \tag{17}$$

Final objective function The final training objective $\mathcal{L}$ is the average over all training pairs:

$$\mathcal{L} = \frac{1}{|\mathcal{L}_{\text{train}}|} \sum_{(e_i, e_j) \in \mathcal{L}_{\text{train}}} (\mathcal{L}_{i\to\cdot} + \mathcal{L}_{\cdot\to j}) \tag{18}$$

where $\mathcal{L}_{\text{train}} = \{(e_s, e_t) \mid e_s \in \mathcal{E}_S, e_t \in \mathcal{E}_T, e_s \equiv e_t\}$ denotes the set of pre-aligned seed anchor pairs.

4 Quantum-enhanced entity alignment

To rigorously validate the representational quality and expressive power of the learned embeddings, we apply our model to the critical downstream task of **entity alignment (EA) in Temporal Knowledge Graphs (TKGs)**. This task aims to establish verifiable correspondences between entities across two disparate TKGs, $\mathcal{G}_S$ (Source) and $\mathcal{G}_T$ (Target), that refer to the same real-world object at comparable time points.

A central challenge in TKG-EA is the robust measurement of entity equivalence, which must account for both their **structural position** and **temporal evolution**. Our alignment mechanism is driven by the **Hybrid Similarity Framework**, a core component of the QLGAN architecture. This framework fundamentally differs from prior KG similarity metrics by integrating two novel concepts:

1. **Quantum-Inspired Metrics**: Utilizing a tensor-based, non-linear metric to capture complex, high-dimensional relational entanglement.
2. **Structurally-Aware Temporal Feature Fusion**: A dedicated mechanism for dynamically weighing and combining structural proximity and time-series feature alignment.

The complete, end-to-end entity alignment procedure, which utilizes the Hybrid Similarity score for bipartite graph matching, is formally detailed in Algorithm 1.

Algorithm 1 Quantum-enhanced entity alignment.

Input: Source entity embeddings $\mathbf{V}_S$, target embeddings $\mathbf{V}_T$; Source initial features $\mathbf{F}_S$, target initial features $\mathbf{F}_T$; Source normalized adjacency $\tilde{\mathbf{A}}_S$, target normalized adjacency $\tilde{\mathbf{A}}_T$; Hyperparameters w, K, w_1, w_2, w_3.
Output: Alignment matrix $\mathbf{P}$.
{Step 1: Non-linear Embedding Similarity (S_{embed})}
$\mathbf{S}_{\text{fid}} \leftarrow \text{QuantumFidelity}(\mathbf{V}_S, \mathbf{V}_T)$
$\mathbf{S}_{\text{ent}} \leftarrow \text{QuantumEntanglement}(\mathbf{V}_S, \mathbf{V}_T)$
$\mathbf{S}_{\text{sup}} \leftarrow \text{QuantumSuperposition}(\mathbf{V}_S, \mathbf{V}_T)$
$\mathbf{S}_{\text{embed}} \leftarrow w_1\mathbf{S}_{\text{fid}} + w_2\mathbf{S}_{\text{ent}} + w_3\mathbf{S}_{\text{sup}}$
{Step 2: Structurally-Aware Feature Diffusion (S_{struct})}
$\mathbf{F}_S^{(0)} \leftarrow \mathbf{F}_S, \mathbf{F}_T^{(0)} \leftarrow \mathbf{F}_T$
for $k = 0$ to $K - 1$ **do**
  $\mathbf{F}_S^{(k+1)} \leftarrow \text{L2-normalize}(\tilde{\mathbf{A}}_S\mathbf{F}_S^{(k)})$
  $\mathbf{F}_T^{(k+1)} \leftarrow \text{L2-normalize}(\tilde{\mathbf{A}}_T\mathbf{F}_T^{(k)})$
end for
$\mathbf{S}_{\text{struct}} \leftarrow \text{AverageSim}(\{\mathbf{F}_S^{(k)}\}_{k=0}^K, \{\mathbf{F}_T^{(k)}\}_{k=0}^K)$
{Step 3: Similarity Fusion}
$\mathbf{S}_{\text{final}} \leftarrow (\mathbf{S}_{\text{embed}} + w \cdot \mathbf{S}_{\text{struct}})/(1 + w)$
{Step 4: Optimal Transport Alignment}
$\bar{S}_{\text{final}} \leftarrow \text{mean}(\mathbf{S}_{\text{final}})$
$\mathbf{S}' \leftarrow \tanh(\mathbf{S}_{\text{final}} - \bar{S}_{\text{final}})$
$\mathbf{K} \leftarrow \exp(\gamma\mathbf{S}')$
$\mathbf{P} \leftarrow \text{Sinkhorn}(\mathbf{K})$
return P

4.1 Hybrid similarity framework

Our alignment framework estimates the correspondence between entities by combining two complementary, robust similarity scores: a non-linear embedding-based similarity (S_{embed}) and a structurally-aware feature similarity (S_{struct}). The fusion of these two provides a comprehensive measure of entity equivalence, integrating both latent semantic information and explicit structural context.

4.1.1 Non-linear embedding similarity (S_{embed})

This component measures the proximity of entities based on their learned low-dimensional vector representations $\mathbf{v}_i$ and $\mathbf{v}_j$. Diverging from standard L_1/L_2 norms or cosine similarity, we employ a suite of non-linear measures inspired by concepts from quantum information theory, which introduce high-order feature interactions to model complex semantic relationships. All entity vectors are assumed to be L_2-normalized.

1. **Quantum Fidelity-Inspired Metric** (S_{fid}): This metric is inspired by quantum fidelity, which quantifies the "closeness" of two quantum states. By using the squared inner product, it captures the probabilistic overlap (transition probability) between the two semantic embeddings:

$$S_{\text{fid}}(\mathbf{v}_i, \mathbf{v}_j) = |\mathbf{v}_i^T \mathbf{v}_j|^2 \tag{19}$$

2. **Quantum Entanglement-Inspired Metric** (S_{ent}): To model latent, non-local correlations within the feature space, we introduce a non-linear similarity that incorporates a feature permutation.

$$S_{\text{ent}}(\mathbf{v}_i, \mathbf{v}_j) = \tanh(\mathbf{v}_i^T \mathbf{v}_j + \lambda(\mathbf{v}'_i)^T \mathbf{v}'_j) \tag{20}$$

   where the **permuted representations** are defined as $\mathbf{v}'_i = \mathbf{v}_i + P\mathbf{v}_i$ and $\mathbf{v}'_j = \mathbf{v}_j - P\mathbf{v}_j$. P is the **circular shift operator** which permutes the elements of a vector ($\mathbf{v}[k] \rightarrow \mathbf{v}[k-1] \bmod D$). λ is a hyperparameter that controls the influence of the permuted component, and the tanh function normalizes the score to $[-1, 1]$.
3. **Quantum Superposition-Inspired Metric** (S_{sup}): This metric aims to capture complex interactions by measuring the similarity between two "superposed" states, which are linear mixtures of the original vectors.

$$S_{\text{sup}}(\mathbf{v}_i, \mathbf{v}_j) = (\alpha\mathbf{v}_i + \beta\mathbf{v}_j)^T(\alpha\mathbf{v}_j + \beta\mathbf{v}_i) \tag{21}$$

   where α and $\beta = \sqrt{1-\alpha^2}$ are hyperparameters that control the mixing ratio of the superposition.
4. **Final Embedding Similarity** (S_{embed}): To fully exploit the complementary strengths of all three non-linear measures (probabilistic overlap, non-local correlations, and feature mixing interactions), we propose a **robust Hybrid Metric** by combining them through a weighted average:

$$S_{\text{embed}} = w_1 S_{\text{fid}} + w_2 S_{\text{ent}} + w_3 S_{\text{sup}} \tag{22}$$

   where w_1, w_2, and w_3 are weights that sum to one ($\sum_{k=1}^{3} w_k = 1$). By learning these weights, the model adaptively determines which quantum-inspired mechanism contributes most to the semantic similarity for a given task, ensuring all components are leveraged in the final score.

4.1.2 Structurally-aware feature similarity (S_{struct})

This component incorporates explicit temporal features ($\mathbf{F}_t$) by enhancing them with structural context from the local graph topology via a **Graph Diffusion** process:

$$\mathbf{F}_t^{(k+1)} = \text{L2-normalize}(\tilde{\mathbf{A}}\mathbf{F}_t^{(k)}) \tag{23}$$

where $\mathbf{F}_t^{(k)}$ is the feature matrix at the k-th hop, and $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2}(\mathbf{A} + \mathbf{I})\mathbf{D}^{-1/2}$ is the symmetrically normalized adjacency matrix.

The final similarity score $\mathbf{S}_{\text{struct}}$ is calculated as the average of the similarities computed at **every** propagation step (from $k = 0$ to K hops) to ensure a robust, multi-scale structural context:

$$\mathbf{S}_{\text{struct}} = \frac{1}{K+1} \sum_{k=0}^{K} \mathbf{S}^{(k)} \tag{24}$$

where $\mathbf{S}^{(k)}$ is the cosine similarity matrix computed from the features $\mathbf{F}_t^{(k)}$.

4.2 Similarity fusion and optimal transport alignment

4.2.1 Similarity fusion

The two complementary similarity measures are combined linearly to produce the final, fused similarity matrix $\mathbf{S}_{\text{final}}$:

$$\mathbf{S}_{\text{final}} = \frac{\mathbf{S}_{\text{embed}} + w \cdot \mathbf{S}_{\text{struct}}}{1 + w} \tag{25}$$

where $w \geq 0$ is a hyperparameter balancing the contribution of the two components.

4.2.2 Optimal Transport (OT) alignment

The final alignment matrix $\mathbf{P}$ is computed using an ϵ-regularized **Optimal Transport (OT)** framework, solved via the differentiable **Sinkhorn-Knopp algorithm** (Knight 2007).

1. **Preprocessing:** The fused similarity matrix is centered and scaled to stabilize the optimization process:

$$\mathbf{S}' = \tanh(\mathbf{S}_{\text{final}} - \bar{S}_{\text{final}}) \tag{26}$$

   where $\bar{S}_{\text{final}}$ is the mean of all elements in $\mathbf{S}_{\text{final}}$.
2. **Transport Kernel and Iteration:** The transport kernel is formed as $\mathbf{K} = \exp(\gamma \mathbf{S}')$, where γ is an adaptive scaling factor. The Sinkhorn algorithm then iteratively normalizes $\mathbf{K}$ using row and column marginal constraints until convergence to a doubly stochastic matrix $\mathbf{P}$, which represents the optimal soft alignment plan.

5 Experiments

5.1 Datasets

Statistics for all datasets are listed in Table 1.

We validate the performance of QLGAN on two distinct and challenging Temporal Knowledge Graph (TKG) entity alignment benchmarks. The selection of these datasets (Xu et al. 2021)—DICEWS and YAGO-WIKI50K—is designed to rigorously test QLGAN's generalization ability across different graph scales, structural complexities, and temporal dynamics.

1. **DICEWS-1K/200:** This dataset, established based on the time knowledge database 1, consists of entity alignment tasks focusing on political events with specific time annotations 2, such as the quadruple (Barack Obama, Access, Ukraine, 2014-07-08)3.The dataset is characterized by a high volume of time-point events, offering a focused challenge on precise temporal sequence modeling. The variation between DICEWS-1K and DICEWS-200 lies solely in the number of seed alignments (S), with 1, 000 and 200 entity pairs known in advance, respectively4. This split allows us to test the model's **robustness to data sparsity** and its effectiveness in low-resource settings.
2. **YAGO-WIKI50K:** This is a large-scale, heterogeneous knowledge graph linking YAGO with similar entities in Wikidata. YAGO and Wikidata are critical benchmarks as they contain temporal information in various forms, including time points, beginning or end times, and time intervals. This diversity poses a challenge for models that only process discrete timestamps. By succeeding on YAGO-WIKI50K, QLGAN demonstrates its capability to handle the complexities of biographical and encyclopedic knowledge at scale, validating its **scalability and generalization across diverse temporal data formats**.

Table 1 Statistics of original datasets

Dataset	DICEWS-1K/200	YAGO-WIKI50K
$\lvert\mathcal{E}_1\rvert$ (Source Entities)	9,517	49,629
$\lvert\mathcal{E}_2\rvert$ (Target Entities)	9,537	49,222
$\lvert\mathcal{R}_1\rvert$ (Source Relations)	247	11
$\lvert\mathcal{R}_2\rvert$ (Target Relations)	246	30
$\lvert\mathcal{T}\rvert$ (Timestamps)	4,017	245
$\lvert\mathcal{Q}_1\rvert$ (Source Quadruples)	307,552	221,050
$\lvert\mathcal{Q}_2\rvert$ (Target Quadruples)	307,553	317,814
$\lvert\mathcal{P}\rvert$ (Total Alignments)	8,566	49,172
$\lvert\mathcal{S}\rvert$ (Training Seeds)	1,000 / 200	5,000

5.2 Evaluation metrics

Consistent with previous state-of-the-art work (Xu et al. 2021, 2022) in entity alignment, we primarily use **Hits@k** (Hk) and the **Mean Reciprocal Rank (MRR)** to evaluate model performance.

We choose these ranking metrics because Entity Alignment is fundamentally a ranking problem, where the objective is to place the true counterpart entity as high as possible within a list of candidates, rather than a binary classification task. Specifically:

- **Hits@k (Hk):** Measures the proportion of correctly ranked knowledge graph entity pairs found in the top k positions. This directly reflects the model's ability to correctly prioritize the true match among candidates.
- **MRR:** Calculates the mean of the reciprocals of these ranks. The MRR metric places a heavy emphasis on obtaining a high rank (e.g., Hit@1), which is crucial for distinguishing superior models, especially in highly competitive alignment tasks.

Hits@k and MRR are the established standards in this domain precisely because they capture the required ranking performance effectively.

5.3 Implementation details

To guarantee the transparency and facilitate the full reproducibility of our results, we provide a comprehensive record of the quantum circuit specifications, key parameter

initialization strategies, and the computational environment used. The complete source code for QLGAN is publicly available at https://github.com/AnonymousforKing/QLGAN.

5.3.1 Hyperparameter configuration

The primary hyperparameters used for training QLGAN on the DICEWS and YAGO-WIKI50K datasets are summarized in Table 2. These settings reflect the optimal configuration found through validation on the DICEWS-1K and YAGO-WIKI50K dataset.

5.3.2 Variational Quantum Circuit (VQC) specification

The core of our model's feature transformation lies in the VQC. Its exact specifications, derived from the model architecture, are essential for replication:

- **Quantum Platform**: The VQC is implemented and simulated using the PennyLane framework (version 0.33) with the "default.qubit" backend.
- **Quantum Width** (N): The circuit width is set to $N = 128$ qubits, matching the embedding dimension D.
- **Circuit Depth** (L): The VQC employs a shallow architecture with a total of $L = 1$ trainable layer.
- **Architecture**: The single layer consists of: an initial Angle Embedding layer, a trainable R_Y rotation layer, a fixed CNOT entangling layer (circular connectivity), and a trainable R_Z rotation layer.
- **Parameter Initialization**: All $2 \times N = 256$ trainable parameters (rotation angles) are initialized by uniform random sampling within the range $[-\pi, \pi]$.

Table 2 Key hyperparameters for QLGAN training

Category	Parameter	Value
Model Dimension	Embedding Dimension (D)	128
	Maximum Quantum Dimension ($D_{\text{max-Q}}$)	128
	Attention Heads (K)	4
Graph/Lineage	Lineage Depth ($L_{\max}$)	4
	Graph Diffusion Steps (K_{struct})	3
	Batch Size (Training)	1024
	Negative Samples ($\mathcal{L}_{\cdot \to j}$)	256
Training	Optimizer	Adam
	Learning Rate (η)	1×10^{-3}
	Margin (γ)	3.0
Alignment Fusion	Structural Weight (w)	0.8
	Sinkhorn Scaling (γ_{sinkhorn})	50.0

5.3.3 Computational environment

All experiments were conducted on a single machine utilizing a high-performance Apple M2 Max chip. The specific environment details are as follows:

- **Hardware**: Apple M2 Max Chip (12-Core CPU, 30-Core GPU), 96 GB Unified Memory.
- **Software Environment**: macOS Sonoma, Python 3.9.
- **Key Libraries**: PyTorch (version 2.1), PyTorch Geometric (PyG), PennyLane (version 0.33).

The training time analysis in §5.8.5 was recorded in this environment, confirming the competitive computational feasibility of QLGAN compared to classical baselines.

5.4 Experimental setup

In our experiments, we compare some traditional and state-of-the-art entity alignment models to evaluate the entity alignment performance between the temporal knowledge graphs of our proposed model. We only add pure structures (quadruples) when completing the task of entity alignment between temporal knowledge graphs. Our proposed model does not learn iteratively and does not add additional information, such as attribute information, name information, and description information. Therefore, we only compare pure entity alignment models.

Three aspects of the performance of our proposed model: (1) Main results. (2) Ablation study. (3) Auxiliary Experiments.

5.5 Baselines and comparative setup

To strictly evaluate QLGAN, we benchmarked it against a comprehensive range of state-of-the-art methods. To ensure reproducibility, all baselines were retrained using their official hyperparameters. These methods are categorized into two groups:

(1) Static entity alignment models, which rely solely on structural triples (h, r, t);
(2) Temporal entity alignment models, which explicitly model temporal dynamics (h, r, t, τ).
Following standard protocols, we exclusively utilized structural and temporal information (quadruples) and excluded auxiliary data (e.g., attributes or images) to ensure a fair comparison of core topological and temporal learning capabilities.

5.6 Main results

The primary experimental results, comparing our proposed QLGAN against a comprehensive suite of static and temporal

entity alignment models, are presented in Table 3. The evaluation was conducted on the DICEWS-200 and DICEWS-1K datasets, which differ in the number of available seed alignments. The results lead to several key conclusions.

First, there is a clear and substantial performance gap between specialized Temporal Entity Alignment Models and the Static Entity Alignment Models. For instance, the best-performing static model, **RREA**, achieves a Hit@1 of 0.722 on DICEWS-1K, which is significantly lower than all temporal models. This observation validates the fundamental premise that explicitly modeling the temporal dimension is crucial for accurate alignment in evolving knowledge graphs.

Second, among the temporal models, our proposed **QLGAN** consistently achieves new state-of-the-art performance across all metrics (H1, H10, and MRR) on both datasets. Notably, **QLGAN** outperforms the strongest baseline, **Dual-Match**. On the more challenging DICEWS-1K dataset, **QLGAN** achieves a Hit@1 score of **0.959**, surpassing **Dual-Match** by 0.6 percentage points and demonstrating its robustness even with fewer seed alignments.

The superior performance of **QLGAN** can be attributed to its novel architecture which holistically addresses the core challenges in temporal entity alignment. Unlike baselines that rely on local neighborhood aggregation, our **Lineage**-aware encoder captures long-range, multi-hop dependencies, providing richer contextual information. Furthermore, the use of a Variational Quantum Circuit (VQC) for feature transformation enables a more powerful and effective fusion of complex spatiotemporal features, leading to more discriminative entity representations.

In summary, the results in Table 3 confirm the effectiveness of our quantum-classical hybrid approach. By synergistically combining deep lineage-based context, quantum-enhanced feature transformations, and an adaptive training-inference framework, **QLGAN** sets a new benchmark for temporal entity alignment.

5.7 Ablation study: validating component contributions

To rigorously validate the contribution and necessity of each key component within the proposed QLGAN framework, we conducted a comprehensive ablation study. We compare the performance of the full QLGAN model against six distinct variants, each disabling a crucial module. The results, measured by Hit@1 (H1), Hit@10 (H10), and Mean Reciprocal Rank (MRR), are presented in Table 4.

The impact of each component confirms the synergistic design of QLGAN, highlighting that the highest performance is achieved only when complex feature engineering, advanced encoding, and specialized metric learning are combined.

- **w/o $\mathbf{h}^{\text{quantum}}$ (Removal of VQC Mapping)**: This variant, which replaces the Variational Quantum Circuit (VQC) mapping with a simple classical non-linear projection (e.g., standard MLP), shows an MRR reduction to 0.956. The $\mathbf{h}^{\text{quantum}}$ feature is critical for high-dimensional correlation encoding.
  Mechanism Explained: Although the performance drop is numerically small compared to removing the temporal signal, this confirms that the VQC provides a superior, higher-expressive feature transformation than classical alternatives. By projecting features into the exponentially

Table 3 Main results

Models	DICEWS-200			DICEWS-1K		
	H1	H10	MRR	H1	H10	MRR
Static Entity Alignment Models						
GCN-Align (Wang et al. 2018)	0.067	0.104	0.175	0.101	0.241	0.150
MuGNN (Cao et al. 2019)	0.367	0.583	0.412	0.525	0.794	0.617
HyperKA (Sun et al. 2020a)	0.383	0.653	0.474	0.588	0.842	0.669
MRAEA (Mao et al. 2020)	0.476	0.733	0.564	0.675	0.870	0.745
RREA (Mao et al. 2020)	0.659	0.824	0.719	0.722	0.883	0.780
KE-GCN (Yu et al. 2021)	0.373	0.625	0.451	0.549	0.827	0.650
Temporal Entity Alignment Models						
TSM-GNN (Jia et al. 2024)	0.756	0.878	0.801	0.821	0.914	0.855
TEA-GNN (Xu et al. 2021)	0.876	0.941	0.902	0.897	0.947	0.911
ChatEA (Jiang et al. 2024a)	0.868	0.906	0.887	0.871	0.912	0.911
MM-ChatAlign (Jiang et al. 2024b)	0.870	0.910	0.891	0.883	0.922	0.901
TREA (Xu et al. 2022)	0.910	0.960	0.927	0.914	0.966	0.933
Dual-Match (Liu et al. 2023)	0.953	0.974	0.961	0.953	0.973	0.961
QLGAN	0.954	0.976	0.964	0.959	0.979	0.965

Table 4 Ablation study results on the entity alignment task, demonstrating the necessity of each component. MRR provides the primary metric for ranking quality assessment

Models	H1	H10	MRR
QLGAN (Full)	0.959	0.979	0.965
w/o $\mathbf{h}^{\text{quantum}}$ (Removal VQC Mapping)	0.948	0.971	0.956
w/o $\mathbf{f}_{\text{lineage}}$ (Removal Lineage Modulation)	0.945	0.968	0.953
w/o S_{hybrid} (Removal Q-Inspired Metric)	0.911	0.956	0.947
w/o S_{entity} (Removal Entity Information)	0.883	0.944	0.913
w/o S_{temporal} (Removal Temporal Signal)	0.793	0.911	0.843

large Hilbert space and encoding entanglement, $\mathbf{h}^{\text{quantum}}$ captures more nuanced, non-linear dependencies, leading to demonstrably sharper alignment boundaries.

- **w/o $\mathbf{f}_{\text{lineage}}$ (Removal of Lineage Modulation)**: Removing the lineage-entangled feature modulation results in a noticeable drop in MRR to 0.953. The $\mathbf{f}_{\text{lineage}}$ component provides multi-hop relational context.
  Mechanism Explained: When removed, the QLGAN layer reverts to aggregating information based only on immediate, single-hop neighbors, ignoring the rich path semantics that connect distant entities. This loss of deep, multi-hop context reduces the quality of the neighbor features $\mathbf{h}''_j$, leading to less precise final entity embeddings.
- **w/o $\mathbf{S}_{\text{hybrid}}$ (Removal of Quantum-Inspired Metric)**: This variant shows a significant drop in MRR to 0.947. The $\mathbf{S}_{\text{hybrid}}$ component is a specialized, learned metric designed to leverage the quantum-enhanced features for distance calculation.
  Mechanism Explained: This result validates that simply using standard distance metrics (e.g., Euclidean or Cosine) on the final embeddings is insufficient. The specialized $\mathbf{S}_{\text{hybrid}}$ metric acts as a **learned comparator**, efficiently translating the complex, high-dimensional correlations encoded by the VQC into an effective similarity score.
- **w/o $\mathbf{S}_{\text{entity}}$ (Removal Entity Information)**: Ablating the core model-based entity similarity component results in the second largest performance drop (MRR decreases from 0.965 to 0.913). The $\mathbf{S}_{\text{entity}}$ metric is derived directly from the deep, contextualized representations learned by the QLGAN encoder.
  Mechanism Explained: This severe decline confirms that the QLGAN encoder's learned deep features are fundamentally crucial for distinguishing between entities. The model cannot perform accurate alignment without its primary source of semantic and structural knowledge.
- **w/o $\mathbf{S}_{\text{temporal}}$ (Removal of Temporal Signal)**: This variant exhibits the most substantial performance degradation, with a dramatic drop in MRR from 0.965 to 0.843 (a decrease of 12.2 percentage points). The component $\mathbf{S}_{\text{temporal}}$ explicitly models structurally-aware temporal consistency between entities.
  Mechanism Explained: Without this explicit temporal guidance, the model relies solely on static feature vectors, failing to capture the dynamic, time-sensitive nature of entity relationships. This confirms that for alignment tasks in evolving graphs, the **temporal dimension is an indispensable, non-substitutable signal**.

In conclusion, this comprehensive ablation study validates that the QLGAN achieves peak performance through a powerful synergy between its components. The model's success is predicated on the integration of:

1) Indispensable Temporal Guidance ($\mathbf{S}_{\text{temporal}}$);
2) High-Quality Semantic Features ($\mathbf{S}_{\text{entity}}$);
3) Novel Feature Refinement Mechanisms ($\mathbf{f}_{\text{lineage}}$ and $\mathbf{h}^{\text{quantum}}$) coupled with a Specialized Similarity Metric ($\mathbf{S}_{\text{hybrid}}$).

5.8 Auxiliary experiments and parameter sensitivity analysis

In addition to the main results and the comprehensive ablation study, we conduct several auxiliary experiments to deeply investigate the robustness, scalability, and impact of key design choices within QLGAN.

5.8.1 Performance on large-scale graph (Scalability)

To rigorously evaluate the effectiveness and scalability of QLGAN in processing large-scale datasets, we conducted experiments on the challenging YAGO-WIKI50K graph. We compared QLGAN's performance against two leading temporal entity alignment baselines: TREA and AGN.

The results in Table 5 confirm QLGAN achieves state-of-the-art performance across all evaluation metrics, significantly outperforming both competitive baselines on this large-scale dataset. QLGAN surpasses the strong AGN model, achieving an absolute improvement of 0.4 percentage points on Hit@1 and 0.002 on MRR. The performance gain over TREA is more pronounced, highlighting QLGAN's robustness. The superior performance can be attributed to the novel integration of the **Lineage-Entangled Modulation**

Table 5 Scalability results on the YAGO-WIKI50K dataset

Model	Hit@1	Hit@10	MRR
TREA	0.940	0.989	0.958
AGN	0.960	0.990	0.972
QLGAN	0.964	0.991	0.974

and the **VQC-based feature enhancement**. Unlike baselines that struggle to aggregate complex features efficiently at scale, QLGAN's ability to enrich neighbor representations with contextually aware, multi-hop semantics prior to the final aggregation leads to more precise and powerful entity embeddings, a capability critical for accurately capturing the intricate relationships in large-scale dynamic knowledge graphs.

5.8.2 Influence of embedding dimension

To investigate the representational capacity of QLGAN, we conducted a sensitivity analysis on the embedding dimension (D) by varying it across {8, 16, 32, 64, 128}. The results, measured by Hit@1, are shown in Fig. 2.

The figure clearly demonstrates a strong positive correlation between the embedding dimension and model performance. As D increases from 8 to 128, the Hit@1 score steadily improves from approximately 92.5% to 95.9%. This trend is expected: higher dimensions mitigate the representational bottleneck, enabling the model to effectively capture the complex structural, semantic, and temporal features. Notably, performance continues to improve significantly even when D increases from 64 to 128. This suggests that the **quantum-enhanced features** generated by QLGAN can effectively utilize the representational power of high-dimensional space without suffering from the typical saturation or severe overfitting issues often seen in classical models within this range. Considering the trade-off between performance and computational overhead, we selected $D = 128$ as the optimal dimension for subsequent experiments.

5.8.3 Influence of quantum-inspired similarity metric

To validate the necessity of the non-linear, quantum-inspired metrics defined in S_{embed}, we performed a direct comparison against a strong classical baseline similarity metric, $S_{\text{embed}}^{\text{classical}}$ (e.g., CSLS similarity). We fixed the QLGAN architecture and embeddings, replacing only the calculation module in the final alignment stage.

The results in Fig. 3 clearly demonstrate that our proposed S_{embed} offers a significant performance advantage, achieving a Hit@1 score of 95.9% compared to 94.8% for $S_{\text{embed}}^{\text{classical}}$. This performance gap is primarily due to the unique design of S_{embed}, which incorporates non-linear measures inspired by quantum fidelity, entanglement, and superposition. These measures are specifically tailored to capture the complex, high-order feature interactions encoded in the VQC-generated embedding space. Unlike traditional L_p norms or CSLS, S_{embed} acts as a sophisticated, **learned comparator** that is naturally adapted to the geometry of QLGAN's embedding space, resulting in demonstrably more accurate alignment decisions.

5.8.4 Influence of lineage depth

To determine the optimal influence range for the **Lineage-Entangled Modulation**, we conducted a parameter sensitivity analysis on the maximum path depth ($L_{\max}$) within the Lineage module, evaluating Hit@1 for depths {2, 4, 6, 8}.

The results in Fig. 4 show a distinct "rise-then-fall" trend, peaking at a path depth of $L_{\max} = 4$. Performance significantly improves from 0.942 at $L_{\max} = 2$ to the peak of **0.959**. However, increasing the depth further to 6 and 8 causes performance degradation to 0.951 and 0.946, respectively. This finding confirms two crucial points: First, overly short paths ($L_{\max} = 2$) result in a lack of rich, multi-hop contextual information. Second, excessively long paths ($L_{\max} \geq 6$) introduce relational noise or lead to a phenomenon analogous to "over-smoothing" in GNNs, where the distinct local context is diluted by irrelevant, distant information, interfering with the final embedding quality. $L_{\max} = 4$ provides the

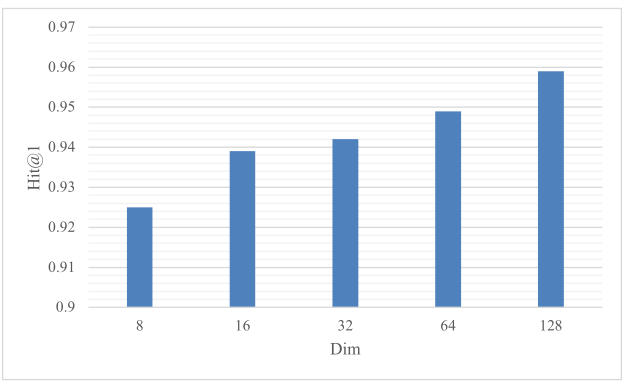

Fig. 2 The impact of different embedding dimensions on Hit@1 performance

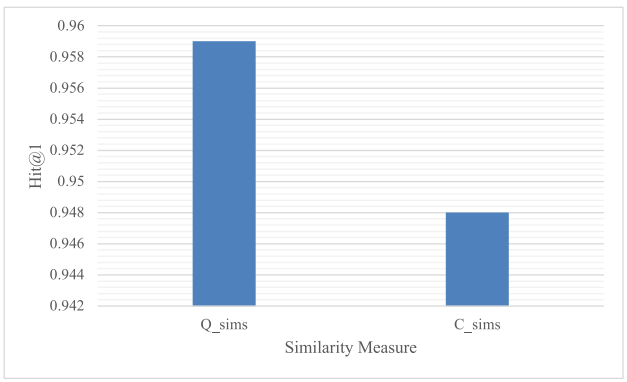

Fig. 3 Performance comparison between our proposed quantum-inspired similarity (S_{embed}) and a classical similarity ($S_{\text{embed}}^{\text{classical}}$)

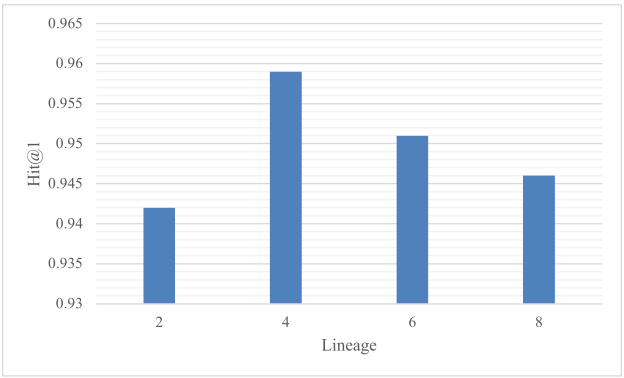

Fig. 4 The impact of different Lineage path depths ($L_{\max}$) on Hit@1 performance

optimal balance between comprehensive context enrichment and noise suppression.

5.8.5 Computational efficiency analysis

To demonstrate the practical feasibility of QLGAN, which integrates complex mechanisms such as the Variational Quantum Circuit (VQC) and multi-hop lineage encoding, we conducted a direct comparative analysis of the training efficiency against the strong classical baseline, DualMatch. The comparison focuses on the total time required for training and alignment on a standard benchmark dataset, with results presented in Table 6.

Quantitative results demonstrate that QLGAN consistently outperforms the strong baseline, DualMatch, across all alignment metrics (e.g., 0.965 MRR versus 0.961). While this performance advantage is accompanied by an increase in computational cost—specifically, a rise in per-epoch runtime from 109 seconds to 125 seconds (approximately +14.7%)—we observe that this overhead remains within a manageable range for practical deployment. The trade-off indicates that the model effectively leverages additional computational resources to secure higher alignment precision, validating the efficiency of the proposed hybrid architecture even under classical simulation constraints.

The observed latency is primarily attributed to two architecturally necessary components designed to enhance feature expressivity. First, the simulation of the Variational Quantum Circuit (VQC) during forward feature mapping projects data into a high-dimensional Hilbert space, enabling the capture of complex non-linear correlations that standard classical layers often fail to distinguish. Second, the Lineage Path Encoder, which constructs and modulates multi-hop neighbor features, introduces explicit computational steps but is critical for resolving structural ambiguities inherent in sparse temporal graphs. These mechanisms are not redundant parameters but essential structures for modeling high-order dependencies.

To strictly contextualize these findings, we offer a two-fold justification. First, in the current saturated landscape where baselines already exceed 0.96 MRR, the observed improvements (+0.4% MRR) signify a meaningful capability to resolve "hard" entities that simpler models cannot address. Second, and more importantly, this work validates the feasibility of hybrid Quantum-Classical learning frameworks. While currently simulated on classical hardware (hence the overhead), QLGAN provides a blueprint for leveraging future quantum computing resources. The additional complexity is thus a strategic investment: it secures immediate robustness gains while establishing a "quantum-ready" paradigm for the next generation of computing hardware.

Table 6 Comparison of computational efficiency (total runtime per epoch) and performance

Models	Time	Hit@1	Hit@10	MRR
QLGAN	125 s	0.959	0.979	0.965
DualMatch	109 s	0.953	0.974	0.961

5.8.6 Intrinsic clustering quality evaluation

We assess the quality of the final entity embeddings $\mathbf{h}^{(L)}$ generated by QLGAN and the baseline DualMatch using two established intrinsic clustering metrics: the Silhouette Score (measuring separability and cohesion) and Inertia (measuring tightness within clusters). These analyses objectively quantify the structural superiority of the representation space learned by our quantum-enhanced method.

As illustrated in Fig. 5 (Left), QLGAN consistently achieves a significantly higher Silhouette Score than DualMatch across a varying number of clusters (K). This sustained superiority demonstrates that QLGAN's features are inherently more separable and cohesive in the learned space. This structural advantage directly stems from the VQC's ability to project features into an exponentially large Hilbert space, effectively resolving non-linear boundaries intractable for classical models.

Furthermore, Fig. 5 (Right) shows that QLGAN exhibits consistently higher Inertia values across all K. When coupled

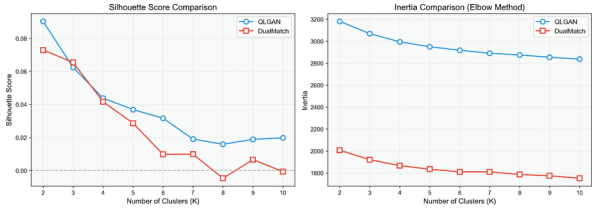

Fig. 5 Comparison of intrinsic clustering metrics. Left: Silhouette Score vs. K, Right: Inertia (Elbow Method) vs. K. QLGAN exhibits higher Silhouette Scores, indicating superior separability and cohesion

Fig. 6 t-SNE visualization of the final entity embeddings. Left: Overall space comparison. Middle/Right: Internal structure of the QLGAN and DualMatch spaces, demonstrating QLGAN's superior feature compactness and coherence

with the high Silhouette Scores, this implies that QLGAN successfully encodes richer and more distinct semantic information into its representations, utilizing a broader, yet better-organized, feature space.

5.8.7 Embedding space visualization via t-SNE

To provide a crucial qualitative understanding, we visualize the final embeddings using t-SNE in Fig. 6.

Figure 6 (Left) shows the QLGAN data (blue) to be tighter and more distinctly clustered than the DualMatch data (red), directly supporting the quantitative Silhouette Score findings. Figure 6 (Middle and Right) reveal the internal structure. The QLGAN space (Middle) forms a relatively compact and well-defined spherical cluster, indicating high internal cohesion and a low presence of outliers. Conversely, the DualMatch embeddings (Right) appear more scattered and less coherent, exhibiting significant noise and less defined structure.

Qualitative proof of quantum enhancement The superior structure and coherence of the QLGAN embedding space are the direct results of the Lineage-Entangled Modulation followed by the VQC Mapping. The VQC, acting as a complex quantum feature map, effectively organizes and purifies the data in the high-dimensional quantum space, leading to a measured classical vector $\mathbf{h}^{\text{quantum}}$ that is inherently better suited for alignment. This visualization serves as strong qualitative evidence that QLGAN successfully leverages quantum-inspired mechanisms to learn a representation space with superior geometric properties.

6 Conclusion and impact

In this paper, we introduced QLGAN (Quantum-Enhanced Graph Learning Attentional Network), a pioneering hybrid model designed to overcome the limitations of classical Graph Neural Networks in handling complex temporal entity alignment within dynamic knowledge graphs. Our work represents a significant step forward in integrating quantum-inspired mechanisms for high-level representation learning.

The key technical and scientific contributions of QLGAN can be summarized as follows:

1. **Lineage-Entangled Feature Modulation:** We developed a novel, multi-hop context propagation mechanism that uses relational path semantics ($\mathbf{f}_{\text{lineage}}$) to actively modulate neighbor features ($\mathbf{h}_j^{(l)}$). This entanglement mechanism ensures that message passing is highly conditional and structurally informed, capturing dependencies that traditional GNNs ignore.
2. **Quantum-Inspired Feature Encoding:** We utilized a **Variational Quantum Circuit (VQC)** as a highly expressive, non-linear feature mapping Φ_{VQC}. The VQC projects features into an exponentially large Hilbert space, leveraging the principles of quantum entanglement to encode richer, high-order correlations into the entity representations ($\mathbf{h}_j^{\text{quantum}}$).
3. **Adaptive Hybrid Similarity Framework:** We introduced a robust framework for alignment that synergizes the semantic depth of quantum-inspired non-linear metrics (S_{embed}) with the explicit structural context derived from graph diffusion (S_{struct}), optimized via Optimal Transport.

Empirical results demonstrate QLGAN's state-of-the-art performance across various benchmark entity alignment tasks. Crucially, our auxiliary analysis—confirmed by superior Silhouette Scores and coherent t-SNE visualizations—provides qualitative evidence that the QLGAN's quantum-enhanced features lead to a representation space with fundamentally better separability and structure compared to classical baselines. The comprehensive ablation study further validates that this success is a direct result of the powerful synergy between explicit temporal modeling and our novel quantum-classical feature refinement components.

QLGAN establishes a compelling and effective paradigm for the emerging field of Hybrid Quantum-Classical Graph Learning. For future work, we intend to formally analyze the quantum advantage provided by the VQC in terms of feature separability bounds. Furthermore, exploring the deployment of QLGAN on near-term quantum hardware and adapting its lineage-entangled principles to dynamic link prediction problems will be key areas of research, solidifying the role of quantum computing in advancing complex relational AI.

Author Contributions Jia Li: Conceptualization, Methodology, Software, Writing – original draft; Yuxi Ma:Investigation and Data Curation; Lingzhong Meng: Supervision.

Funding This work was supported by the Doctoral Research Start-up Fund of Zhengzhou University of Light Industry (2024BSJJ030).

Data Availability https://pan.baidu.com/s/1pbYaqEmKByffYyOCZqWXJQ?pwd=eq8z

Declarations

Competing interests 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

Beer K, Khosla M, Köhler J, Osborne TJ, Zhao T (2023) Quantum machine learning of graph-structured data. Phys Rev A 108(1):012410

Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S (2017) Quantum machine learning. Nature 549(7671):195–202

Cai L, Mao X, Ma M, Yuan H, Zhu J, Lan M (2022) A simple temporal information matching mechanism for entity alignment between temporal knowledge graphs. In: Proceedings of the 29th international conference on computational linguistics, pp 2075–2086

Cao Y, Liu Z, Li C, Li J, Chua T-S (2019) Multi-channel graph neural network for entity alignment. The 57th Annual Meeting of the Association for Computational Linguistics (ACL)

Chen L, Li Z, Xu T, Wu H, Wang Z, Yuan NJ, Chen E (2022) Multimodal siamese network for entity alignment. In: Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining, pp 118–126

Chen M, Tian Y, Yang M, Zaniolo C (2017) Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. the 26th International Joint Conference on Artificial Intelligence(IJCAI)

Gao Y, Liu X, Wu J, Li T, Wang P, Chen L (2022) Clusterea: Scalable entity alignment with stochastic training and normalized mini-batch similarities. In: Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining, pp 421–431

Innan N, Sawaika A, Dhor A, Dutta S, Thota S, Gokal H, Patel N, Khan MA-Z, Theodonis I, Bennai M (2024) Financial fraud detection using quantum graph neural networks. Quantum Mach Intell 6(1):7

Jia W, Ma R, Yan L, Niu W, Ma Z (2024) Time-aware structure matching for temporal knowledge graph alignment. Data & Knowl Eng 151:102300

Jiang X, Shen Y, Shi Z, Xu C, Li W, Li Z, Guo J, Shen H, Wang Y (2024) Unlocking the power of large language models for entity alignment. In: Proceedings of the 62nd annual meeting of the association for computational linguistics (Volume 1: Long Papers), pp 7566–7583

Jiang X, Shen Y, Shi Z, Xu C, Li W, Zihe H, Guo J, Wang Y (2024) Mm-chatalign: A novel multimodal reasoning framework based on large language models for entity alignment. In: Findings of the association for computational linguistics EMNLP 2024, pp 2637–2654

Knight PA (2007) The sinkhorn–knopp algorithm: Convergence and applications. In: Web information retrieval and linear algebra algorithms, 11.02. - 16.02.2007

Liu C-Y, Kuo E-J, Abraham Lin C-H, Gemsun Young J, Chang Y-J, Hsieh M-H, Goan H-S (2025) Quantum-train: Rethinking hybrid quantum-classical machine learning in the model compression perspective. Quantum Mach Intell 7(2):80

Liu Z, Cao Y, Pan L, Li J, Chua T-S (2020) Exploring and evaluating attributes, values, and structures for entity alignment. In EMNLP

Liu F, Chen M, Roth D, Collier N (2021) Visual pivoting for (unsupervised) entity alignment. In: AAAI, 2021

Liu X, Wu J, Li T, Chen L, Gao Y (2023) Unsupervised entity alignment for temporal knowledge graphs. In: Proceedings of the ACM web conference 2023, pp 2528–2538

Mao X, Wang W, Wu Y, Lan M (2021) Boosting the speed of entity alignment 10*: Dual attention matching network with normalized hard sample mining. In WWW,2021

Mao X, Wang W, Wu Y, Lan M (2022) Lightea: A scalable, robust, and interpretable entity alignment framework via three-view label propagation. In: Proceedings of the 2022 conference on empirical methods in natural language processing, pp 825–838

Mao X, Wang W, Xu H, Lan M, Wu Y (2020) Mraea: an efficient and robust entity alignment approach for cross-lingual knowledge graph. In: Proceedings of the 13th international conference on web search and data mining, pp 420–428

Mao X, Wang W, Xu H, Wu Y, Lan M (2020) Relational reflection entity alignment. In: Proceedings of the 29th ACM international conference on information & knowledge management, pp 1095–1104

OZHIGOV Y (2022) Three principles of quantum computing. Quantum Inf Comput 22(15&16):1280–1288

Pomarico D, Fanizzi A, Amoroso N, Bellotti R, Biafora A, Bove S, Didonna V, Forgia DL, Pastena MI, Tamborra P et al (2021) A proposal of quantum-inspired machine learning for medical purposes: An application case. Mathematics 9(4):410

Qi H, Wang L, Gong C, Gani A (2024) A survey on quantum data mining algorithms: challenges, advances and future directions. Quantum Information Processing 23(3)

Rath M, Date H (2024) Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy. EPJ Quantum Technology 11(1):72

Ray A, Madan D, Patil S, Pati P, Rapsomaniki M, Kohlakala A, Dlamini TR, Muller SJ, Rhrissorrakrai K, Utro F, et al (2024) Hybrid quantum-classical graph neural networks for tumor classification in digital pathology. In: 2024 IEEE international conference on quantum computing and engineering (QCE), vol 1, pp 1611–1616 . IEEE

Sun Z, Chen M, Hu W, Wang C, Dai J, Zhang W (2020) Knowledge association with hyperbolic knowledge graph embeddings. In EMNLP

Sun Z, Wang C, Hu W, Chen M, Dai J, Zhang W, Qu Y (2020) Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In: Proceedings of the AAAI conference on artificial intelligence, pp 222–229

Trsedya BD, Qi J, Rui Z (2019) Entity alignment between knowledge graphs using attribute embeddings. Proceedings of the AAAI Conference On Artificial Intelligence 33:297–304

Tüysüz C, Rieger C, Novotny K, Demirköz B, Dobos D, Potamianos K, Vallecorsa S, Vlimant J-R, Forster R (2021) Hybrid quantum classical graph neural networks for particle track reconstruction. Quantum Mach Intell 3(2):29

Verdon G, Mccourt T, Luzhnica E, Singh V, Leichenauer S, Hidary J (2019) Quantum graph neural networks

Wang Z, Lv Q, Lan X, Zhang Y (2018) Cross-lingual knowledge graph alignment via graph convolutional networks. In: Proceedings of the 2018 conference on empirical methods in natural language processing, pp 349–357

Xu Y, Huang H, State R (2024) Ctqw-graphsage: Trainabel continuous-time quantum walk on graph. In: International conference on artificial neural networks, pp 79–92 . Springer

Xu C, Su F, et al (2021) Time-aware graph neural network for entity alignment between temporal knowledge graphs. In: EMNLP

Xu C, Su F, Xiong B, Lehmann J (2022) Time-aware entity alignment using temporal relational attention. In: Proceedings of the ACM web conference 2022, pp 788–797

Yang H-W, Zou Y, Shi P, Lu W, Lin J, Sun X (2020) Aligning cross-lingual entities with multi-aspect information. In: Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)

Yu D, Yang Y, Zhang R, Wu Y (2021) Knowledge embedding based graph convolutional network. In: Proceedings of the web conference 2021, pp 1619–1628

Zhu H, Xie R, Liu Z, Sun M (2017) Iterative entity alignment via joint knowledge embeddings. In: IJCAI, vol 17, pp 4258–4264