

Article

SeSKGC: A Semantic–Structural Fusion Framework for Knowledge Graph Completion

Ping Feng ^{1,2,3,*} , Siqi Xu ¹ , Xinping Du ¹ , Yan Chen ¹ and Yuyuan Dong ¹

¹ College of Computer Science and Technology, Changchun University, Changchun 130022, China; xusiqi29@gmail.com (S.X.); xinpindg82@gmail.com (X.D.); potato66669999@gmail.com (Y.C.); dyy030908@gmail.com (Y.D.)

² Jilin Provincial Key Laboratory of Manchurian Railway Data Engineering and Intelligent Decision-Making, Changchun 130022, China

³ Education Key Laboratory of Intelligent Rehabilitation and Barrier-Free Access for the Disabled, Changchun University, Changchun 130022, China

* Correspondence: pingfeng20188@gmail.com

Abstract

Knowledge graphs play a vital role in tasks such as recommendation systems, question-answering systems, and information retrieval. However, during practical construction, they commonly suffer from structural incompleteness and sparse relationships, which limit reasoning performance and downstream applications. Existing methods typically focus solely on either structural modeling or semantic modeling: embedding models relying solely on graph structures struggle to leverage textual information about entities and relationships. In contrast, semantic approaches relying solely on pre-trained language models struggle to accurately capture complex relationship patterns. To address this challenge, this paper proposes SeSKGC, a semantic–structural fusion knowledge graph completion model. At the semantic level, the model employs the DeBERTa pre-trained language model to encode entity and relation text. It incorporates a neighbor text augmentation mechanism to introduce local semantic context and enhance representation quality. At the structural level, it adopts complex-space rotation to model relationships using a RotatE-like approach, and aggregates local topological information through relative position attention to capture complex relationship patterns. At the scoring stage, the model employs a weighted fusion strategy to combine semantic and structural scores and utilizes InfoNCE contrastive loss for joint optimization. Experiments conducted on WN18RR and FB15k-237 datasets demonstrate that SeSKGC achieves overall superior performance on metrics including MRR and Hits@N compared to multiple representative baseline methods. Ablation studies and parameter sensitivity analysis of fusion weight λ further reveal that the semantic encoding and structural modeling modules exhibit distinct complementary roles, while the weighted fusion design in the scoring layer plays a crucial role in enhancing model performance and stability.

Keywords: knowledge graph completion; semantic–structural fusion; pre-trained language models; rotation in space



Received: 23 March 2026

Revised: 21 April 2026

Accepted: 23 April 2026

Published: 26 April 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Knowledge graphs (KGs) [1–3] represent entities and relationships in the form of triples. They can be viewed as directed graphs where nodes represent entities and edges represent the corresponding relationships between entities. As an external knowledge

foundation, knowledge graphs play a pivotal role in numerous natural language processing (NLP) tasks and serve as a core pillar of contemporary artificial intelligence systems. However, since most knowledge graphs are constructed through manual or semi-automated processes, a significant number of implicit entities and relationships remain undiscovered. Consequently, incompleteness has become a near-universal issue for knowledge graphs, severely limiting their application in tasks such as recommendation systems [4], question-answering systems [5], information retrieval [6], and natural language understanding. Consequently, in recent years, researchers have increasingly focused on knowledge graph completion (KGC), whose core idea is to predict potentially missing entities or relationships based on existing triples.

It is important to note that knowledge graphs are enormous in scale, often containing millions of entities, making most graph neural networks operating on the overall graph structure computationally unscalable. Therefore, existing knowledge graph completion methods typically model at the triple level and are mainly divided into two paradigms: structural embedding methods and semantic encoding methods.

Structural embedding methods model entities and relations through geometric transformations and are good at capturing structural patterns such as symmetry, antisymmetry, inversion, and combination. However, these methods rely on the graph structure itself and are difficult to extend to unseen entities or sparse scenarios. They can be divided into two subcategories. (1) Translation model-based methods: For example, based on the invariance of translation, Bordes et al. proposed the most representative classic translation model, TransE [7], and led a large number of researchers to study the Trans series models, among which the representative improved models include TransH [8], TransR [9], and TransD [10]. The main idea of TransE is that if the triple (h, r, t) is true, the sum of the vector representations of the head entity and the relation is close to the vector representation of the tail entity, and conversely, when the triple is true, $h + r \approx t$. From the perspective of inferring relation patterns, Sun et al. proposed the RotatE [11] model, which can model and infer various relation patterns, including symmetry, antisymmetry, inversion, and composite relations. The RotatE model characterizes each relation as a rotation operation from the source entity to the target entity in complex space. It not only has good scalability and linear complexity in terms of time and memory overhead, but also effectively captures a variety of relation patterns. It shows significantly better performance than existing methods in link prediction tasks. (2) Methods based on semantic matching models, such as DistMult [12], HolE [13], and QuatE [14], which measure the rationality of triples by designing scoring functions in a low-dimensional embedding space. DistMult uses a bilinear inner product to model the interaction between entities and relations. It is computationally simple but difficult to characterize antisymmetric relations. HolE improves the expressive power of complex relations by using circular convolution and the Fourier transform. QuatE extends the embedding to quaternion space and captures richer relation patterns through quaternion multiplication. Although these methods have certain advantages in terms of efficiency and relation pattern modeling, they generally have limitations such as limited expressive power and difficulty in utilizing textual semantic information.

Semantic encoding methods leverage textual information about entities and relations, and use pre-trained language models to learn contextual representations, demonstrating excellent performance in inductive scenarios. KG-BERT [15], a knowledge graph completion method based on a pre-trained language model, exemplifies this. It converts triples (h, r, t) into text sequences and uses BERT for sequence classification, thereby predicting the credibility of triples and completing missing knowledge. Unlike traditional embedding methods that rely solely on graph structures, KG-BERT fully utilizes textual descriptions of entities and relations to capture contextual semantic information, achieving superior

performance in triple classification, link prediction, and relation prediction tasks. However, this method suffers from significant computational overhead. Link prediction requires replacing almost all entities to generate candidate triples, which are then individually input into a deep Transformer model for computation, resulting in low efficiency. Furthermore, it does not consider the issues of structured knowledge and entity ambiguity.

Consequently, combining semantic information with structured knowledge has become a natural trend in the field of knowledge graph completion. Traditional structural embedding methods primarily rely on the topological structure of knowledge graphs for representation learning, and while they offer significant advantages in modeling relational patterns, they often struggle to fully capture the deep semantic features of entities and relationships when entity semantic information is underutilized. Furthermore, their generalization capabilities may be limited in scenarios involving sparse or unseen entities. However, semantic encoding methods based on pre-trained language models can extract contextual semantic information from entity and relationship texts, offering a new approach to address the aforementioned issues. Motivated by this, we introduce a semantic encoding module and integrate it with the structural modeling module to enhance representational capabilities and reasoning performance in knowledge graph completion tasks. It should be noted that the experiments in this paper are conducted under standard transductive knowledge graph completion settings, with the primary goal of verifying the effectiveness of semantic–structural integrated modeling.

The main contributions of this paper are as follows:

- We propose a unified semantic–structural fusion framework: we design a knowledge graph completion method that jointly models semantic and structural features within the same system, taking into account both the fine-grained modeling capabilities of semantic expression and the reasoning advantages of structural constraints.
- We propose a weighted fusion strategy for the scoring layer: achieving efficient unification of semantic and structural scoring layers, which not only preserves the complementarity of the two types of information, but also significantly improves the computational efficiency of the reasoning stage.
- We conduct experiments on two public datasets, WN18RR and FB15k-237. The results show that SeSKGC outperforms several representative baseline models, verifying the effectiveness and stability of the proposed method under standard transductive settings.

2. Related Works

2.1. Structure Embedding Method

Several translation-based structural embedding methods have achieved significant results in link prediction tasks. The TransE [7] model pioneered an unsupervised knowledge graph embedding learning paradigm based on vector translation. Its core idea is to use the translation operation of relation vectors to make the true triples satisfy $h + r \approx t$. Building on this, the TransR [9] model introduces a multi-relation space mechanism to address the complex situation where the same entity participates in multiple facts in the knowledge graph. Specifically, this method represents relation vectors in the relation space R^k and models triple facts in the $R^{k \times d}$ space, thereby improving the model's relation discrimination ability, but also introducing additional computational overhead.

Furthermore, the TaRP [16] model, as a relation prediction framework, combines transformation-based embedding methods to generate entity and relation representations, and further introduces node type recognition and encoding mechanisms, thereby enhancing the accuracy of relation prediction. Another type of method attempts to generate node embeddings through neural networks. For example, the model proposed by Shallom [17]

utilizes a shallow neural network structure to model the relation prediction task, improving feature representation capabilities while maintaining computational efficiency.

Translation-based structural embedding methods have significant advantages in modeling graph topology, capturing multiple relation patterns, and maintaining inference efficiency, making them suitable for link prediction tasks on large-scale sparse knowledge graphs. However, these methods generally rely on pre-defined geometric assumptions and hand-designed scoring functions, limiting their ability to characterize high-order semantic information and long-range dependencies. When entity text descriptions are rich or new entities frequently appear, relying solely on structural signals often fails to fully extract latent semantics, easily leading to insufficient generalization in scenarios with unseen entities or long-tailed relationships.

2.2. Text Encoding Methods

Textual information in knowledge graphs has been utilized in various forms in recent years. The textual descriptions of nodes (especially lengthy texts) can provide rich semantic context for entities, helping models predict relationships when training on entities that have not yet been seen. However, retrieving the textual description of each node typically requires additional data preprocessing steps and relies on the support of external knowledge bases.

BERT [18], as a pre-trained model based on masked language modeling, has demonstrated outstanding performance in various Natural Language Processing (NLP) tasks. KG-BERT [15] was one of the earliest models to introduce BERT into knowledge graph completion tasks, primarily targeting link prediction (LP) tasks, while also proposing corresponding variant models for relation prediction (RP) and triple classification (TC). Inspired by this, BERTRL [19] further proposes to improve the model's inductive ability by linearizing the encoding of triples and fine-tuning them using a language model.

However, such language model-based approaches generally face problems such as high computational resource consumption and low inference efficiency, which limit their application in large-scale knowledge graphs. In contrast, lightweight pre-trained language models have a lower computational burden and can provide high-quality text embedding representations for new datasets without retraining. For example, GloVe [20], as a typical pre-trained word vector, can generate multi-dimensional embedding vectors for each word in the vocabulary (where n is the vocabulary size) in $O(n)$. However, GloVe only models at the word level and cannot capture contextual dependencies; while masked language models can model the context of text sequences, thus better representing the semantics of multi-word phrases or complete sentences. However, the inference cost of deep language models is still high, and end-to-end deployment on large-scale graphs faces practical constraints. Without effective alignment and constraints with the graph structure, relying solely on semantic representation can easily overlook explicit relationship patterns between entities, and the characterization of structural features such as symmetry, antisymmetry, and composition remains insufficient.

2.3. Fusion Paradigm

In recent years, scholars have proposed various knowledge graph completion models that attempt to fuse semantics and structure. The LASS [21] model fine-tunes BERT to obtain text encoding and uses the resulting loss function to construct semantic embeddings; a variant model suitable for triple classification tasks has also been proposed. The StAR [22] model aims to alleviate the problem of redundant language model parameters by reusing graph element embeddings, but its fusion mechanism is complex, computationally inefficient, and still has limitations in structural modeling. The KEPLER [23] model jointly trains

the semantic learning objective of the language model with the structural optimization objective of the link prediction task to achieve collaborative modeling of semantic and structural features. The KGLM [24] model retrain a RoBERTa-based language model and introduces natural language sentences transformed from knowledge graph triples as a supplementary corpus during training, while adding entity and relation type embedding layers to explicitly incorporate graph structure information during training. Furthermore, some models have achieved better results in link prediction tasks by introducing different negative sample sampling strategies during the text embedding learning stage. The SimKGC [25] model incorporates structured information into the candidate entity ranking process by using a re-ranking mechanism based on the path distance of target nodes in the graph; while the MoCoSA [26] model uses independently trained structured embeddings and merges them with text embeddings, thereby further improving the comprehensiveness and discriminative ability of knowledge representation.

In summary, the existing works have placed different emphases on semantic expressive power, structural constraint strength, and computational cost. Schemes that inject structural information during the pre-training stage achieve tighter integration but have higher training costs. Methods that explicitly align semantic and structural embeddings at the representation layer offer greater expressive flexibility but are prone to parameter inflation and training instability. The paradigm of fusion at the scoring layer is simple to implement and has low engineering implementation costs, but often results in limited coupling between semantic and structural modules and insufficient collaborative modeling capabilities. Therefore, how to strengthen the interaction and complementarity between semantic and structural elements while maintaining inference efficiency, and how to balance generalization and structural inference, remain key challenges in existing fusion paradigms, and also provide room for improvement for the SeSKGC model proposed in this paper.

3. Methodology

3.1. Problem Definition

Knowledge graph completion faces challenges such as sparsity, the diversity of entities and relations, and the integration of semantics and structure. Knowledge graphs typically contain a large number of entities and relations, but most relations are not fully expanded, making it difficult for models to effectively utilize existing information during reasoning. Secondly, the types of entities and relations in knowledge graphs are numerous, and complex relational patterns (such as symmetric, antisymmetric, and combinatorial relations) are difficult to capture accurately, especially when encountering new entities or relations; traditional methods often fail to perform effective reasoning in these cases, while existing methods typically rely on structural or semantic information, the effective integration of these two remains challenging.

This paper aims to solve the knowledge graph completion problem. We define a knowledge graph as a set of triples: $G = \{(h, r, t) \mid h, t \in \mathcal{E}, r \in \mathcal{R}\}$, where $\mathcal{E}$ represents the entity set, and $\mathcal{R}$ represents the relation set. In the triple (h, r, t) , h represents the head entity, r represents the relation, and t represents the tail entity. The knowledge graph may contain missing triples. The knowledge graph completion task aims to infer missing entities based on known facts, i.e., $(h, r, ?) \rightarrow t$ or $(?, r, t) \rightarrow h$. The model aims to learn a mapping function: $f : (\mathcal{E}, \mathcal{R}) \rightarrow \mathbb{R}^d$, such that the scores of true triples are higher than those of negative samples. The proposed SeSKGC method aims to enhance the model's reasoning ability by combining a semantic encoding module and a structural modeling module.

3.2. Overall Framework

The proposed SeSKGC model aims to integrate semantic and structural information within a unified framework, thereby enhancing semantic representation and relational reasoning capabilities in knowledge graph completion tasks. The overall model architecture, as illustrated in Figure 1, comprises three primary modules: a semantic encoding module, a structural modeling module, and a fusion mechanism. The semantic encoding module extracts semantic features of entities and relations based on a pre-trained language model and captures local semantic context through a neighborhood text enhancement mechanism, enriching the semantic connotation of entity representations. The structural modeling module performs relation rotation modeling in complex space and introduces relative positional attention to perceive local topological structures. Finally, a weighted fusion mechanism unifies semantic and structural scores, achieving a unified semantic–structural knowledge graph completion. Through this design, the model can balance the fine-grained expression of semantic modeling with the reasoning capabilities of structural embedding, achieving higher completion accuracy and robustness in transductive settings.




Figure 1. SeSKGC Semantic–Structural Fusion Architecture Overview Diagram. This model comprises a semantic encoding module, a structural modeling module, and a weighted fusion mechanism. It achieves knowledge graph completion by synergistically optimizing semantic and structural scores.

3.2.1. Semantic Encoding Module

At the semantic level, SeSKGC uses a pre-trained language model to encode the textual descriptions of entities and relations to obtain context-aware semantic representations. Let the text sequences of the head entity, relation, and tail entity be X_h , X_r , and X_t , respectively. Their semantic embedding representation is as follows:

$$E_h = f(X_h), E_r = f(X_r), E_t = f(X_t) \quad (1)$$

Here, f represents the encoding function, and this paper uses DeBERTa as the semantic encoder. DeBERTa's decoupled attention mechanism and relative position encoding design can capture the relative dependencies in the text more precisely, thus maintaining consis-

tency with the relative position attention mechanism of the subsequent structural modeling module at the semantic level.

To further enhance the contextual richness of semantic representation, this module introduces a neighborhood text enhancement mechanism. For a target entity h , its k -hop neighbor set is denoted as $N(h)$. The model uses an attention mechanism to perform weighted aggregation of neighbor text features:

$$E'_h = \sum_{n \in N(h)} \alpha_{hn} E_n, \alpha_{hn} = \frac{\exp(E_h^T W E_n)}{\sum_{m \in N(h)} \exp(E_h^T W E_m)} \quad (2)$$

where E_n is the semantic vector of the neighboring entity, W is the learnable weight matrix, and α_{hn} is the neighbor attention weight, which measures the semantic contribution of entity n to the target entity h . By weighted aggregation of neighborhood semantic information, the model can incorporate the contextual semantics of related entities in the graph structure into the current entity representation, thereby alleviating the problem of semantic sparsity or insufficient entity description in knowledge graphs. The enhanced entity representation is as follows:

$$\tilde{E}_h = E_h + \gamma E'_h \quad (3)$$

Here, γ is the balance coefficient, which is used to control the contribution of semantic information from neighbors. It should be emphasized that the proposed semantic neighborhood enhancement does not perform uniform smoothing over all neighboring entities. Instead, the attention coefficient α_{hn} is used to selectively evaluate the semantic contribution of each neighbor to the target entity, while the residual formulation in Equation (3) preserves the original semantic representation of the target entity itself. Therefore, even in heterophilic knowledge graphs where connected entities may exhibit substantially different semantic attributes, the influence of less relevant neighbors can be suppressed through lower attention weights and the controllable contribution coefficient γ , rather than directly dominating the target representation.

To measure the plausibility of triples in the semantic space, the semantic score function uses cosine similarity calculation:

$$S_{sem}(h, r, t) = \cos(\tilde{E}_h + E_r, E_t) \quad (4)$$

The semantic encoding module obtains context-aware semantic representations through a pre-trained language model, and fully integrates local textual information of entities with global contextual features by leveraging neighborhood attention aggregation and semantic residual fusion mechanisms, providing semantic support for relational reasoning and fusion mechanisms in the subsequent structural modeling module.

3.2.2. Structural Modeling Module

The structural modeling module aims to capture the topological structure and complex relationship patterns in knowledge graphs, and its overview diagram is shown in Figure 2. This paper draws inspiration from RotatE, performing rotational modeling in complex space to characterize the directional and symmetric representation of relations. For a triple (h, r, t) , the rotational mapping is defined as:

$$h_r = h \circ r = (h_{Re} \circ \cos \theta_r - h_{Im} \circ \sin \theta_r) + i(h_{Re} \circ \sin \theta_r + h_{Im} \circ \cos \theta_r) \quad (5)$$

Here, h_{Re} and h_{Im} represent the real and imaginary part vectors of the head entity in the complex space, respectively. This operation is equivalent to performing a relation-dependent

phase rotation on the entity vectors in the complex plane, achieving an interpretable representation of the relational patterns from a geometric perspective. In this way, the model can characterize various complex relational patterns such as symmetric, antisymmetric, and combinatorial relationships.




Figure 2. Schematic Diagram of the Structural Modeling Module Based on Complex Plane Rotation and Relative Position Attention.

The structural score function is defined as the Euclidean distance between entities after rotation, and is used to measure the plausibility of triples in structural space.

$$S_{str}(h, r, t) = -\|h \circ r - t\|_2 \quad (6)$$

Furthermore, to enhance the local sensitivity of structural modeling, this module introduces a relative positional attention mechanism. For a target node i and a sampled neighbor j in its localized neighborhood, the attention weight with positional bias is defined as follows:

$$\alpha_{ij} = \frac{\exp(Q_i K_j^T + b_{d(i,j)})}{\sum_k \exp(Q_i K_k^T + b_{d(i,k)})} \quad (7)$$

Here, $d(i, j)$ denotes the truncated shortest-path distance between node i and node j , computed only within the localized k -hop subgraph centered at node i . In practice, this distance is obtained by limited-hop BFS on the local neighborhood, rather than by computing global all-pairs shortest paths over the full graph. Q_i and K_j represent the query vector and key vector of nodes i and j , respectively, obtained by linear transformation of structural features, and $b_{d(i,j)}$ is a distance-related learnable bias term used to explicitly model relative positional relationships in the local subgraph.

Based on the above attention weights, the final structure of node i is represented as:

$$H'_i = \sum_{j \in N(i)} \alpha_{ij} V_j \quad (8)$$

Here, V_j is the eigenvector of node j , and $N(i)$ represents the set of neighbors of node i . This weighted aggregation process enables the model to dynamically integrate local structural information based on complex space rotation modeling, so that the structural representation has both global relational reasoning ability and can reflect local topological differences.

Furthermore, the structural aggregation mechanism proposed in this paper does not rely on the homogeneity assumption that nearby entities should share similar features. It combines relational rotation modeling with local relative position attention, thereby performing neighborhood aggregation under structural constraints that account for relational features. Since attention weights are computed within truncated local subgraphs and

are jointly modulated by both structural features and distance-related biases, the model tends to assign higher weights to neighbors rich in structural information, rather than indiscriminately averaging all connected nodes. This design helps mitigate the risk of noise propagation in heterogeneous knowledge graphs caused by neighbors with structural mismatches or weak correlations.

The structural modeling module achieves an organic combination of geometric modeling and structural attention at the expression layer through the joint design of complex space rotation mapping and the relative position attention mechanism, providing structural support for the final semantic–structural fusion.

3.2.3. Integration Mechanism

After obtaining semantic and structural scores, SeSKGC achieves unified modeling of the two types of information through a weighted fusion mechanism at the score layer. The core idea of this mechanism is to dynamically and linearly combine the contextual features learned by the semantic encoding module in the semantic space with the relational patterns captured by the structural modeling module in the complex space, thereby enhancing overall reasoning consistency while maintaining the diversity of model representation. The fusion score function is defined as follows:

$$S(h, r, t) = \lambda \cdot S_{sem}(h, r, t) + (1 - \lambda) \cdot S_{str}(h, r, t) \quad (9)$$

Here, $\lambda \in [0, 1]$ is an adjustable fusion coefficient used to balance the importance of semantic and structural information. When λ takes a larger value, the model relies more on semantic representation in inference; conversely, it emphasizes the role of structural patterns. Through this linear weighting strategy, the model can adaptively allocate weights between the two representation spaces according to task characteristics and data distribution, achieving complementary fusion of information.

To further optimize the representation space distribution after semantic–structural fusion, this paper employs the InfoNCE contrastive learning loss function during the training phase. Its goal is to maximize the difference between positive and negative sample scores, as shown below:

$$\ell = -\log \frac{\exp(\text{sim}(S_i, S_i^+) / \tau)}{\sum_{j=1}^N \exp(\text{sim}(S_i, S_j^-) / \tau)} \quad (10)$$

Here, τ is the temperature parameter. This loss function, by narrowing the score distance between positive samples and widening the score distance between negative samples, enables the model to learn a consistent and discriminative embedding distribution in both semantic and structural spaces.

Furthermore, to improve the model's ability to distinguish negative samples, this module employs a structural negative sampling auxiliary loss:

$$\ell_{str} = -\log \sigma(S_{str}(h, r, t)) - \sum_{i=1}^k \log(1 - \sigma(S_{str}(h'_i, r, t'_i))) \quad (11)$$

Ultimately, the joint optimization objective of SeSKGC is defined as:

$$\ell_{total} = \ell + \mu \ell_{str} \quad (12)$$

where μ is the balance coefficient, used to adjust the relative weights of contrastive learning and structural learning.

Through the aforementioned fusion mechanism, the model achieves information complementarity and feature synergy at both the semantic and structural levels, maintaining contextual consistency in the semantic space while enhancing the geometric discernibility of the structural space, thereby effectively improving the model's prediction and relational reasoning capabilities.

3.2.4. Theoretical Analysis of Score Layer-Weighted Fusion

In Section 3.2.3, SeSKGC linearly weights the semantic and structural scores at the scoring layer using Formula (9) to obtain the final fused score, and combines the InfoNCE contrastive loss and structural negative sampling loss to form a joint optimization objective. This design is simple to implement in engineering, has good interpretability, and also has relatively good theoretical properties from the perspective of optimization convergence and generalization ability.

Convergence analysis. Under the conventional assumptions that the semantic and structural scores are continuously differentiable functions with respect to the parameters and that their gradients are bounded; that the InfoNCE loss and structural negative sampling loss are smooth with respect to the scores and satisfy the Lipschitz condition; and that the learning rate sequence used in training satisfies common decay conditions, then existing stochastic optimization theory indicates that the average gradient norm of the parameter sequence $\{\theta^{(t)}\}$ will converge to zero during the training process based on mini-batch stochastic gradient descent. That is:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \left\| \Pi_{\theta} L(\theta^{(t)}) \right\|^2 = 0 \quad (13)$$

This ensures that the model converges to a certain critical point in the desired sense, without unbounded divergence during training. It is important to emphasize that, since the fusion score $S(h, r, t)$ is simply a convex combination of the semantic and structural scores, its gradient with respect to the parameters can be written as:

$$\Pi_{\theta_{sem}} L = \lambda \cdot \Pi_{\theta_{sem}} L_{sem}, \Pi_{\theta_{str}} L = (1 - \lambda) \cdot \Pi_{\theta_{str}} L_{str} \quad (14)$$

In other words, based on the original semantic and structural channel gradients, only a scaling of the scalar weight $1 - \lambda$ is introduced, without changing the smoothness of the loss function or introducing additional unstable terms. Therefore, given that the semantic and structural modules are inherently convergent, introducing weighted fusion of the scoring layer does not compromise the convergence of the overall training objective. On the contrary, by jointly constraining the two representations under a unified objective, the model reaches a stable state more quickly within a finite number of training epochs. This is consistent with the smooth convergence observed in Sections 4.5 and 4.6 below.

Generalization error analysis. From the perspective of statistical learning theory, the semantic encoding module and the structural modeling module each correspond to a family of scoring functions (which can be understood as two types of “base models”). The weighted fusion of the scoring layer constructs a simple convex combination of these two types of functions. Given that the loss function is bounded and satisfies the Lipschitz condition, capacity measures such as Rademacher complexity can be used to give an upper bound between empirical risk and expected risk. Related conclusions show that the Rademacher complexity of the family of functions corresponding to the linear convex combination will not exceed the weighted sum of the complexities of each sub-family of functions, i.e.,

$$\zeta(F_{\lambda}) \leq \lambda \zeta(F_{sem}) + (1 - \lambda) \zeta(F_{str}) \quad (15)$$

The resulting upper bound on the generalization error is no worse than that obtained by using the semantic model or the structural model alone. Intuitively, this means that introducing a weighted fusion of the scoring layer without significantly increasing the model capacity does not lead to a significant risk of overfitting. On the contrary, by jointly constraining the two representations under a unified loss, it helps to find a more suitable trade-off between empirical risk and generalization error.

As seen in the ablation experiments in Section 4.5, removing the semantic module (w/o Sem) or the structural module (w/o Str) both lead to varying degrees of decrease in MRR and Hits@N. Even when retaining both scores but removing weighted fusion and performing only simple merging (w/o FW), the overall performance is still lower than the complete SeSKGC. This empirically confirms the above analysis: with a relatively constant model capacity, explicit score-layer weighted fusion leverages the complementarity of semantic and structural representations, and flexibly balances the two information streams across different datasets through adjustable fusion weights λ , thereby achieving better generalization performance and more stable inference performance in transductive settings.

3.3. Summary

This section elaborates on the overall structure of the SeSKGC model. The model centers on semantic–structural fusion, implementing a unified knowledge representation learning framework across three levels: semantic encoding, structural modeling, and information fusion. At the semantic level, the model employs the pre-trained language model DeBERTa for contextual semantic encoding and introduces local semantic information through a neighborhood text enhancement mechanism to enrich the semantic expression of entities and relationships. At the structural level, based on the idea of complex space rotation modeling, the model introduces a relative position attention mechanism to characterize complex relational patterns and local topological features in the knowledge graph. At the fusion level, the model achieves efficient unification of semantic and structural scores through a score-layer weighted fusion strategy, and combines InfoNCE contrastive learning and structural negative sampling for joint optimization, promoting collaborative learning between the semantic and structural spaces.

From an overall performance perspective, SeSKGC establishes a close connection between semantic representation and structural reasoning, balancing semantic richness and structural interpretability. Compared with existing representative semantic–structural fusion methods, SeSKGC emphasizes the “lightweight” and “adjustable” nature of the fusion approach. StAR [22] introduces multi-layer attention interactions between semantic and structural representations to deeply couple the two types of features, while this results in a tighter fusion, it incurs significant overhead in terms of parameter scale and implementation complexity. KEPLER [23] embeds knowledge into the language model during pre-training, requiring joint optimization of a large-scale PLM, leading to a longer training cycle and higher costs when transferring to new scenarios. In contrast, this paper decouples semantic encoding and structural modeling into two relatively independent sub-modules, while maintaining their respective advantages, it performs post-fusion only at the scoring layer using scalar weights λ . This avoids modifications to the pre-trained language model structure and the design of additional interactive networks, reducing the overall model complexity and engineering threshold. Furthermore, λ , as an easily adjustable hyperparameter, allows for flexible adjustment of the relative weights of semantics and structure under different datasets and task settings, resulting in a more balanced performance between generalization ability and inference efficiency.

4. Experiments and Analysis

4.1. Datasets

To verify the effectiveness of the proposed SeSKGC model in the semantic–structural fusion knowledge graph completion task, this paper selects two classic public datasets: WN18RR [27] and FB15k-237 [28]. Both datasets are widely used in knowledge graph completion research under transductive settings and can comprehensively evaluate the model’s inference performance under different relational patterns.

WN18RR: Derived from the WordNet semantic vocabulary database, it is a deduplicated version of WN18 [7]. This dataset mainly contains hierarchical and inverse relations (such as hypernym/hyponym), suitable for testing the model’s ability to model symmetric and antisymmetric relations and inverse relational reasoning.

FB15k-237: Derived from the Freebase knowledge base, it is a version of the original FB15k [7] after removing inverse relation pairs, retaining more complex relational patterns (such as combinatorial relations, hierarchical relations, etc.). This dataset is closer to real-world application scenarios and can test the model’s performance in complex relational modeling and semantic interaction fusion. Detailed information about the datasets is shown in Table 1.

It should be noted that although the semantic encoding module used in this paper is based on a pre-trained language model—which may offer semantic advantages due to its use of a pre-trained language model—we have not yet specifically validated the model’s ability to infer unseen entities on standard inductive datasets (such as Wikidata5M [23] and NELL-One [29]). Therefore, the experimental results in this paper are primarily intended to illustrate the complementary roles of semantic and structural information in the task of completing transductive knowledge graphs, rather than to draw empirical conclusions regarding inductive generalization capabilities. Further evaluations in inductive scenarios will be conducted in future work.

Table 1. Statistics of WN18RR and FB15k-237 datasets.

Dataset	Entities	Relations	Train	Valid	Test
WN18RR	40,943	11	86,835	3034	3134
FB15k-237	14,541	237	272,115	17,535	20,466

4.2. Evaluation Metrics

To comprehensively evaluate the model’s performance in the knowledge graph completion task, this paper employs five commonly used evaluation metrics, including mean rank (MR), mean reciprocal rank (MRR), and Hits@N ($N \in \{1, 3, 10\}$). These metrics have been widely used in link prediction research in recent years to measure the overall ranking quality and prediction accuracy of the model. MR and MRR reflect the model’s overall ranking performance in the entity ranking task. MR represents the average rank of the correct entity among all candidate entities; a smaller value indicates better model performance. MRR measures the model’s predictive consistency and stability by calculating the reciprocal average of the correct entity rankings for each test sample. Hits@N represents the proportion of correct entities ranked in the top N across all test samples, used to measure the model’s high-ranking hit capability. All metrics are calculated under a filtering setting; that is, during the testing phase, for each triplet to be predicted (h, r, t) , if there exists a true triplet with the same head or tail entities that has already appeared in the training, validation, or test sets, its score is ignored to ensure the fairness and accuracy of the ranking evaluation.

4.3. Experimental Setup

In the experimental implementation, this paper adopts a unified data partitioning strategy for both datasets, dividing them into non-overlapping training, validation, and test sets. The training set is used for learning model parameters, the validation set for hyperparameter tuning and model selection, and the test set for final evaluation of the model's generalization ability. All experiments were conducted on an NVIDIA RTX 4090 GPU environment, and the experimental platform was implemented using Python 3.6.5 and the PyTorch 1.10.1 deep learning framework.

Model training employed the Adam [30] optimizer for iterative parameter updates. To obtain reliable experimental results, we performed hyperparameter grid searches on two datasets. On the WN18RR dataset, the learning rate was set to 5×10^{-5} , and the fusion weight λ to 0.4; on the FB15k-237 dataset, the learning rate was set to 1×10^{-5} , and the fusion weight λ to 0.6. The batch size was consistently set to 1024, and the epochs to 200. This configuration ensured stable validation set performance on both datasets. For the text encoding part, semantic embeddings of entities and relations were cached offline to reduce the training and inference overhead caused by repeated computations during iteration. All embedding dimensions were determined through grid search on the validation set.

Regarding the loss function, the model jointly optimized InfoNCE contrastive loss and structural negative sampling loss to simultaneously improve discriminative ability in the semantic space and relation modeling quality in the structural space. The specific values of the temperature coefficient τ and the semantic–structural balance coefficient μ were also determined based on validation set performance. In implementation, we do not compute shortest-path distances globally over the entire knowledge graph. Instead, for each target node, we construct a localized truncated neighborhood and obtain hop distances to sampled neighbors using limited-hop BFS. These local distances are then used to instantiate the relative positional bias in Equation (7). Therefore, the positional bias computation is restricted to local subgraphs and does not involve global all-pairs shortest-path preprocessing.

4.4. Comparative Experiment

Our proposed SeSKGC model is compared with several SOTA models using different technical approaches. These models cover text-based methods, structure-based methods, and methods that fuse text and structure. To ensure the representativeness and comprehensiveness of the comparison, we selected several models with typical characteristics as comparison objects.

KG-BERT [15]: The first method to introduce a pre-trained language model into the knowledge graph completion task. KG-BERT linearizes the triples h, r, t into natural language sequences, inputs them into a Transformer-based BERT model, and judges the validity of the triples through a binary classification task.

MTL-KGC [31]: This method proposes a multi-task learning framework that jointly models three tasks: relevance ranking, relation prediction, and link prediction, thereby effectively improving the model's ranking performance.

StAR [22]: A representative semantic–structural fusion method. StAR divides each triple into two parts (head-relation, relation-tail), encodes the two parts using a text encoder, and achieves the fusion of semantic and structural information through spatial metrics and a deterministic scoring mechanism.

RotatE [11]: Introduces rotation operations in complex space to characterize various relational patterns (symmetric, antisymmetric, inverse, combinatorial, etc.), significantly enhancing the model's structural reasoning ability. RotatE demonstrates excellent rela-

tional modeling capabilities in link prediction tasks and is a representative of structural embedding methods.

AttH [32]: This model, through rotation, reflection, and attention mechanisms in hyperbolic space, simultaneously captures the hierarchy and complex logical relationships of knowledge graphs in low-dimensional embeddings, significantly improving link prediction performance.

TuckER [33]: A knowledge graph link prediction model based on Tucker tensor decomposition. By decomposing the triple tensor into the product of the core tensor and the entity and relation embedding matrices, it achieves efficient and expressive linear modeling. All benchmark models are implemented using publicly available code or by the authors, and the proposed model is compared under the same embedding dimension and number of training epochs. The comparison results with the above six benchmark models are shown in Table 2.

Table 2. Comparison of SeSKGC and Representative Baseline Models in Link Prediction Performance on WN18RR and FB15k-237 Datasets (MR, MRR, Hits@1/3/10).

Model	WN18RR					FB15k-237				
	MR	MRR	Hits@1	Hits@3	Hits@10	MR	MRR	Hits@1	Hits@3	Hits@10
KG-BERT	97	0.452	0.422	0.464	0.512	153	—	—	—	0.420
MTL-KGC	89	0.331	0.203	0.383	0.597	132	0.267	0.172	0.298	0.458
StAR	51	0.401	0.243	0.491	0.709	117	0.296	0.205	0.322	0.482
RotatE	3340	0.476	0.428	0.492	0.571	117	0.338	0.241	0.375	0.533
AttH	—	0.486	0.443	0.499	0.573	—	0.348	0.252	0.384	0.540
TuckER	—	0.470	0.443	0.482	0.526	—	0.358	0.266	0.394	0.544
ours	57	0.546	0.457	0.621	0.723	115	0.361	0.232	0.396	0.557

Experimental results show that the proposed model outperforms current mainstream baseline models on both datasets. On WN18RR, SeSKGC achieves significant improvements across all ranking metrics. Compared to the structural baseline AttH, SeSKGC's MRR increases from 0.486 to 0.546, a relative improvement of approximately 12%; Hits@1 increases from 0.443 to 0.457, a relative improvement of approximately 3%; and on Hits@3, which measures overall ranking quality, it increases from 0.499 to 0.621, an improvement of over 24%, while Hits@10 also improves by approximately 2% compared to the best baseline. Although it lags slightly behind some structural models in terms of MRR, SeSKGC's advantages in MRR and Hits at all levels are more significant, indicating that semantic–structural joint modeling can significantly improve prediction accuracy without significantly sacrificing overall ranking costs.

On FB15k-237, SeSKGC also demonstrates stable competitiveness. Compared to the current best-performing TuckER, SeSKGC improves its MRR from 0.358 to 0.361, and its Hits@3 and Hits@10 scores from 0.394 and 0.544 to 0.396 and 0.557, respectively, achieving a gain of approximately 0.5–2.4% in mid-to-high rank accuracy. Simultaneously, the MR decreases from 117 to 115, indicating an overall improvement in ranking quality. It's worth noting that SeSKGC's Hits@1 is slightly lower than TuckER on this dataset, suggesting that in scenarios with more long-tail relationships and more complex structures, this model is more advantageous in improving overall recall and mid-to-high rank accuracy, while there is still room for further optimization in extremely high rank accuracy.

The results from both datasets show that SeSKGC, by introducing a pre-trained language model and neighborhood text enhancement mechanisms in the semantic encoding stage, utilizing complex space rotation and relative position attention in the structural modeling stage, and employing a weighted fusion strategy in the scoring layer, creates a strong complementary relationship between semantic and structural information in the

entity ranking task. This design not only achieves significant improvements on the traditional WN18RR with strong structure and rich inverse relations, but also maintains a stable advantage on the FB15k-237 with complex relation types, verifying the applicability and effectiveness of the proposed method in different knowledge graph scenarios.

4.5. Ablation Experiment

To analyze the specific contribution of each component module to the overall performance, this paper designed ablation experiments, constructing the following three model variants: (1) w/o Sem: removing the semantic encoding module and retaining only the structural modeling module; (2) w/o Str: removing the structural modeling module and using only the semantic encoding module for triple scoring; (3) w/o FW: retaining both semantic and structural scoring channels, but removing the weighted fusion strategy and simply merging the two scores. The complete model is denoted as SeSKGC. The performance of each model on WN18RR and FB15k-237 is shown in Table 3.

Table 3. Ablation experiment results and performance contribution analysis (MRR, Hits@1/3/10) of SeSKGC components on the WN18RR and FB15k-237 datasets.

Model	WN18RR				FB15k-237			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
w/o Sem	0.535	0.439	0.606	0.708	0.321	0.185	0.370	0.532
w/o Str	0.496	0.428	0.590	0.663	0.313	0.206	0.361	0.498
w/o FW	0.531	0.447	0.617	0.715	0.349	0.227	0.392	0.550
ours	0.546	0.457	0.621	0.723	0.361	0.232	0.396	0.557

Based on the WN18RR results, SeSKGC achieves the best performance across all metrics compared to the three variants. Comparing the complete model with w/o Sem, the MRR improves from 0.535 to 0.546, and Hits@1 and Hits@10 improve from 0.439 and 0.708 to 0.457 and 0.723, respectively, indicating that the semantic encoding module still provides stable gains on datasets with many inverse relations. However, the degeneration is more pronounced with w/o Str after removing the structural module, with an MRR of only 0.496 and Hits@10 dropping to 0.663, indicating that complex space rotation and relative position attention play a crucial role in capturing complex relational patterns. In contrast, the performance of w/o FW falls between the two, suggesting that introducing adjustable fusion weights can further improve the MRR by about 1% while preserving both semantic and structural information, making the model more balanced across all metrics.

On the FB15k-237 dataset, the ablation results show a slightly different, but generally consistent, trend compared to WN18RR. Compared to the complete model, the MRR of w/o Sem decreased from 0.361 to 0.321, and Hits@1 decreased from 0.232 to 0.185, showing a significant decline. The degradation of w/o Str was even more pronounced in Hits@3 and Hits@10 (0.361 and 0.498, respectively), indicating that the semantic module contributes more to high precision (Hits@1, MRR), while the structural module is more important for mid-to-high recall (Hits@3, Hits@10). Although w/o FW utilizes both types of information, its overall performance is still slightly lower than SeSKGC, with MRR and Hits@10 of 0.349 and 0.550, respectively. This shows that simple addition is insufficient to fully leverage the complementary advantages of semantic and structural information, while weighted fusion at the scoring layer is more effective in balancing the two types of information. To more intuitively illustrate the performance differences between different variants, the results in Table 3 are further plotted as bar charts, as shown in Figures 3 and 4.

Experimental results show that the semantic encoding module and the structural modeling module each have their own emphasis on different datasets, but neither can be dispensed with. The weighted fusion mechanism, though seemingly simple, plays a crucial role in steadily improving the overall performance. Ablation experiments indirectly verify the necessity and effectiveness of the semantic–structural joint modeling framework and fusion strategy proposed by SeSKGC.




Figure 3. Ablation Study Results for Each Component Module of SeSKGC on the WN18RR Dataset.




Figure 4. Ablation Study Results for Each Component Module of SeSKGC on the FB15k-237 Dataset.

4.6. Parameter Sensitivity Experiment

To further examine the impact of the fusion weight λ in the weighted fusion strategy of the scoring layer, this paper designs parameter sensitivity experiments on the WN18RR and FB15k-237 datasets, while maintaining other hyperparameter settings consistent with Section 4.3, only the fusion coefficient λ between semantic and structural scores in Formula (9) is adjusted. In the actual parameter tuning process, we initially tried more extreme values of λ (close to 0 or 1), and found that the performance on the validation set fluctuated significantly. Therefore, we ultimately controlled λ within the relatively stable [0, 1] range shown in the figure. For each value, the MRR and Hits@1,3,10 indices were calculated under the filtering settings, and the performance curves as a function of λ were plotted, as shown in Figures 5 and 6.

From the experimental results, it can be observed that the model performance changes relatively smoothly with the fusion weight λ . When λ takes extreme values, i.e., excessively biased towards the semantic or structural channels, the MRR and Hits indices at each level

decrease to varying degrees; while in the range where the contributions of semantic and structural aspects are more balanced, the model performs better overall on both datasets. This indicates that the semantic encoding module and the structural modeling module are significantly complementary in the link prediction task, and relying solely on either channel is insufficient to fully realize the model's potential. On the other hand, the curve fluctuations are relatively limited under different λ values, demonstrating that SeSKGC is robust to the fusion weights: the model maintains stable MRR and Hits@N levels across a wide parameter range without significant oscillations or performance collapse. This means that reliable performance can be obtained in practical applications without overly fine-tuning λ , further validating the practicality and effectiveness of the proposed score-layer weighted fusion strategy.




Figure 5. Parameter Sensitivity Analysis of λ on the WN18RR Dataset with Weight Fusion.




Figure 6. Parameter Sensitivity Analysis of λ on the FB15k-237 Dataset with Weight Fusion.

4.7. Time Complexity Analysis

To evaluate the computational efficiency of the SeSKGC model within the semantic-structural fusion framework, this section analyzes its time complexity from three aspects: semantic encoding, structural modeling, and fusion computation. Let the number of entities be $|E|$, the number of relations and embedding dimension be $|R|$, the embedding dimension be d , the average number of neighbors be k , and the average length of the text sequence be L .

At the semantic level, the model uses the pre-trained language model DeBERTa to encode the context of entity and relational text. The complexity of each forward computation is mainly determined by the Transformer's self-attention mechanism, which is approxi-

mately $O(L^2d)$. If there are $|E| + |R|$ text sequences in the graph, the overall complexity can be expressed as:

$$O((|E| + |R|) \cdot L^2d) \quad (16)$$

Since DeBERTa is a parameter-sharing PLM, the actual computation can significantly reduce inference overhead through batch input and caching mechanisms. Therefore, this part can be considered a moderately controllable fixed cost in the overall training process.

At the structural level, complex-space rotation modeling and relative positional attention are the main sources of computation. It should be noted that the relative positional bias in Equation (7) relies only on truncated shortest-path distances within localized neighborhoods, rather than on global all-pairs shortest-path computation over the full graph. For each target node, these local hop distances are obtained by limited-hop BFS within the sampled neighborhood, so the additional overhead is linear in the neighborhood size. Under bounded neighborhood sampling, this cost is dominated by the subsequent local attention and aggregation operations. Therefore, the structural component can be approximately expressed as:

$$O(|E| \cdot (1 + k) \cdot d) \quad (17)$$

Since $K \ll |E|$ and the neighborhood size can be controlled through local sampling, this stage has good linear scalability and is suitable for large-scale knowledge graph completion tasks.

The score-layer weighted fusion (Equation (9)) and InfoNCE contrastive loss (Equations (10)–(12)) only involve vector similarity computation and scalar operations, with complexity approximately $O(B \cdot d)$, where B is the batch size. During training, the semantic and structural modules can be executed independently, which further improves efficiency. Therefore, the single-step training complexity of SeSKGC can be approximately written as:

$$O((|E| + |R|)L^2d + |E|kd) \quad (18)$$

This expression assumes that the positional distances used in the structural attention are computed locally within truncated neighborhoods, and thus no global $O(|V|^3)$ shortest-path computation is introduced.

Comprehensive analysis reveals that SeSKGC's computational overhead is primarily concentrated in the semantic encoding portion, while structural modeling and fusion computation both exhibit a linear growth trend. During the inference phase, because DeBERTa can cache text embeddings offline, the overall prediction complexity can be reduced to $O(|E|kd)$, achieving high inference efficiency. Compared to end-to-end solutions that directly rely on large language models (such as KG-BERT or KGLM), SeSKGC has significant advantages in terms of time and memory consumption, while maintaining strong semantic–structural synergy.

5. Discussion

From a semantic–structural fusion perspective, SeSKGC is closer to a lightweight, adjustable scoring layer fusion paradigm. On the semantic side, the model utilizes DeBERTa and neighborhood text enhancement mechanisms to construct contextual semantic representations of entities and relationships; on the structural side, it combines RotatE-style complex space rotation and relative position attention to characterize various higher-order relationship patterns. The two channels are linearly integrated at the scoring layer using scalar weights λ . This maintains the clarity of the overall structure and controllable implementation costs, while also enabling the model to achieve competitive performance on

WN18RR and FB15k-237 datasets. Nevertheless, the above design should be regarded as a mitigation strategy rather than a complete solution to heterophily. Although the semantic attention mechanism, residual fusion, localized structural aggregation, and score-level decoupling collectively improve the robustness of the model to heterogeneous neighbors, the current framework does not explicitly identify or suppress malignant heterophily, where connected entities may exhibit highly conflicting semantic characteristics or structural roles. In such cases, local aggregation may still introduce noisy signals into representation learning or score estimation. Therefore, incorporating heterophily-aware neighbor filtering, relation-conditional suppression, or more adaptive aggregation mechanisms remains an important direction for future research.

It is important to note that the experiments in this paper are conducted under the standard transductive setting, where the training and test sets share the same entity set, and the model is evaluated only on the two public datasets WN18RR and FB15k-237. This setting is suitable for examining the complementary effects of semantic enhancement and structural modeling in standard link prediction tasks. However, it does not support empirical conclusions regarding the model's ability to generalize to unseen entities. Although the semantic encoding module, built upon a pre-trained language model and neighborhood text augmentation, may provide a basis for extending the model to unseen-entity settings, this capability has not been empirically validated in the present study. Its effectiveness in standard inductive benchmarks still requires further investigation in future work.

At the application level, SeSKGC's semantic-structural collaborative modeling approach aligns well with the construction of digital humanities and industry knowledge graphs. Taking the ongoing construction of the South Manchuria Railway historical archives and the knowledge graph of cities along the railway line as an example, with the continuous organization of new archives, place names, and relationships, the sets of entities and relationships in the graph exhibit dynamic expansion characteristics. Traditional embedding models that rely solely on structure struggle to absorb incremental knowledge in a timely manner. With the support of the semantic encoding module, the model may help initialize semantic representations for newly introduced entities based on archival text, place name descriptions, or biographical materials, while the structural modeling module may further adjust their positions under existing topological constraints. Through weighted fusion scoring, these two modules could provide a potential technical path for incremental completion and auxiliary analysis in the South Manchuria Railway historical knowledge graph.

This research is also one of the interim results of the "South Manchuria Railway" project. Within this project framework, future work will further expand and validate the applicability of the proposed method in real-world business scenarios: on the one hand, rigorous unseen entity experimental settings will be constructed on the South Manchuria Railway historical knowledge graph and other open-world datasets to systematically evaluate the performance of SeSKGC in incremental entity access and cross-graph migration; on the other hand, the overhead of the semantic encoding part will be further compressed without significantly sacrificing performance, for example, by attempting knowledge distillation on the existing DeBERTa encoder or by integrating smaller pre-trained language models to enable deployment in larger-scale graphs and resource-constrained environments. Simultaneously, plans are in place to extend the current "structure + text" unimodal setting to multimodal knowledge graphs, introducing heterogeneous information such as images, tables, or event logs to explore more complex reasoning and decision-making tasks in vertical fields such as healthcare and finance, thereby more fully leveraging the application value of the semantic-structural fusion framework.

6. Conclusions

This paper addresses the knowledge graph completion problem under conditions of sparse structure and complex relational patterns, proposing the SeSKGC semantic–structural co-modeling framework. On the semantic side, the model utilizes a pre-trained language model to encode entities and relational texts, and incorporates a neighborhood text enhancement mechanism to alleviate the problem of insufficient entity description. On the structural side, it employs complex space rotation modeling and relative position attention to effectively characterize various high-order relational patterns such as symmetry, antisymmetry, and combinations. The outputs of the two channels are unified at the scoring layer through an adjustable weighted fusion strategy, and jointly optimized using InfoNCE contrastive loss and structural negative sampling, enabling the semantic and structural spaces to converge collaboratively under the same training objective. Experimental results on the WN18RR and FB15k-237 benchmark datasets show that SeSKGC outperforms many mainstream baseline models in terms of MRR and Hits@N. Ablation experiments verify the key roles of the semantic encoding module, structural modeling module, and weighted fusion mechanism in improving performance. Parameter sensitivity analysis further demonstrates the model’s robustness to the fusion weight λ . Overall, SeSKGC significantly improves semantic expression and structural reasoning capabilities in knowledge graph completion tasks while maintaining high inference efficiency, providing a concise and effective implementation path for the design and application of subsequent semantic–structural fusion methods.

Author Contributions: Conceptualization, P.F. and S.X.; methodology, S.X.; data curation, S.X., X.D. and Y.C.; investigation, Y.D.; project administration, P.F.; supervision, P.F.; writing—original draft preparation, S.X.; writing—review and editing, P.F., X.D., Y.C. and Y.D. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Jilin Provincial Department of Science and Technology Project, Research and Development of an Integrated Platform for Intelligent Mining of Jilin Manchurian Railway Historical Materials and Public Cultural Services Driven by Multimodal Semantics (Project No. 20260203051SF).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Vrandečić, D.; Krötzsch, M. Wikidata: A free collaborative knowledgebase. *Commun. Acm* **2014**, *57*, 78–85. [\[CrossRef\]](#)
2. Bollacker, K.; Evans, C.; Paritosh, P.; Sturge, T.; Taylor, J. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*; Association for Computing Machinery: New York, NY, USA, 2008; pp. 1247–1250.
3. Suchanek, F.M.; Kasneci, G.; Weikum, G. Yago: A core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web*; Association for Computing Machinery: New York, NY, USA, 2007; pp. 697–706.
4. Gao, J.; Peng, P.; Lu, F.; Claramunt, C.; Xu, Y. Towards travel recommendation interpretability: Disentangling tourist decision-making process via knowledge graph. *Inf. Process. Manag.* **2023**, *60*, 103369. [\[CrossRef\]](#)
5. Sun, H.; Bedrax-Weiss, T.; Cohen, W. Pullnet: Open domain question answering with iterative retrieval on knowledge bases and text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 2380–2390.
6. Xiong, C.; Power, R.; Callan, J. Explicit semantic ranking for academic search via knowledge graph embedding. In *Proceedings of the 26th International Conference on World Wide Web*; International World Wide Web Conferences Steering Committee: Geneva, Switzerland, 2017; pp. 1271–1279.

7. Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; Yakhnenko, O. Translating embeddings for modeling multi-relational data. *Adv. Neural Inf. Process. Syst.* **2013**, *26*, 2787–2795.
8. Wang, Z.; Zhang, J.; Feng, J.; Chen, Z. Knowledge graph embedding by translating on hyperplanes. *Proc. AAAI Conf. Artif. Intell.* **2014**, *28*, 1112–1119. [[CrossRef](#)]
9. Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; Zhu, X. Learning entity and relation embeddings for knowledge graph completion. *Proc. AAAI Conf. Artif. Intell.* **2015**, *29*, 2181–2187. [[CrossRef](#)]
10. Ji, G.; He, S.; Xu, L.; Liu, K.; Zhao, J. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2015; pp. 687–696.
11. Sun, Z.; Deng, Z.H.; Nie, J.Y.; Tang, J. Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv* **2019**, arXiv:1902.10197. [[CrossRef](#)]
12. Yang, B.; Yih, W.; He, X.; Gao, J.; Deng, L. Embedding entities and relations for learning and inference in knowledge bases. *arXiv* **2014**, arXiv:1412.6575.
13. Nickel, M.; Rosasco, L.; Poggio, T. Holographic embeddings of knowledge graphs. *Proc. AAAI Conf. Artif. Intell.* **2016**, *30*, 1955–1961. [[CrossRef](#)]
14. Zhang, S.; Tay, Y.; Yao, L.; Liu, Q. Quaternion knowledge graph embeddings. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 2731–2741.
15. Yao, L.; Mao, C.; Luo, Y. KG-BERT: BERT for knowledge graph completion. *arXiv* **2019**, arXiv:1909.03193. [[CrossRef](#)]
16. Cui, Z.; Kapanipathi, P.; Talamadupula, K.; Gao, T.; Ji, Q. Type-augmented relation prediction in knowledge graphs. *Proc. AAAI Conf. Artif. Intell.* **2021**, *35*, 7151–7159. [[CrossRef](#)]
17. Demir, C.; Moussallem, D.; Ngomo, A.C.N. A shallow neural model for relation prediction. In *2021 IEEE 15th International Conference on Semantic Computing (ICSC)*; IEEE: New York, NY, USA, 2021; pp. 179–182.
18. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 4171–4186.
19. Zha, H.; Chen, Z.; Yan, X. Inductive relation prediction by BERT. *Proc. AAAI Conf. Artif. Intell.* **2022**, *36*, 5923–5931. [[CrossRef](#)]
20. Pennington, J.; Socher, R.; Manning, C.D. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2014; pp. 1532–1543.
21. Shen, J.; Wang, C.; Gong, L.; Song, D. Joint language semantic and structure embedding for knowledge graph completion. *arXiv* **2022**, arXiv:2209.08721. [[CrossRef](#)]
22. Wang, B.; Shen, T.; Long, G.; Zhou, T.; Wang, Y.; Chang, Y. Structure-augmented text representation learning for efficient knowledge graph completion. *Proc. Web Conf.* **2021**, *2021*, 1737–1748.
23. Wang, X.; Gao, T.; Zhu, Z.; Liu, Z.; Li, J.; Kepler, J.T. A unified model for knowledge embedding and pre-trained language representation. *Trans. Assoc. Comput. Linguist.* **2021**, *9*, 176–194. [[CrossRef](#)]
24. Youn, J.; Tagkopoulos, I. Kglm: Integrating knowledge graph structure in language models for link prediction. In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2023; pp. 217–224.
25. Wang, L.; Zhao, W.; Wei, Z.; Liu, J. Simkgc: Simple contrastive knowledge graph completion with pre-trained language models. *arXiv* **2022**, arXiv:2203.02167.
26. He, J.; Liu, J.; Wang, L.; Li, X.; Xu, X. Mocosa: Momentum contrast for knowledge graph completion with structure-augmented pre-trained language models. In *2024 IEEE International Conference on Multimedia and Expo (ICME)*; IEEE: New York, NY, USA, 2024; pp. 1–6.
27. Dettmers, T.; Minervini, P.; Stenetorp, P.; Riedel, S. Convolutional 2d knowledge graph embeddings. *Proc. AAAI Conf. Artif. Intell.* **2018**, *32*, 1811–1818. [[CrossRef](#)]
28. Toutanova, K.; Chen, D. Observed versus latent features for knowledge base and text inference. In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and Their Compositionality*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2015; pp. 57–66.
29. Xiong, W.; Yu, M.; Chang, S.; Guo, X.; Wang, W.Y. One-shot relational learning for knowledge graphs. *arXiv* **2018**, arXiv:1808.09040. [[CrossRef](#)]
30. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
31. Kim, B.; Hong, T.; Ko, Y.; Seo, J. Multi-task learning for knowledge graph completion with pre-trained language models. In *Proceedings of the 28th International Conference on Computational Linguistics*; International Committee on Computational Linguistics: Barcelona, Spain, 2020; pp. 1737–1743. [[CrossRef](#)]

32. Chami, I.; Wolf, A.; Juan, D.C.; Sala, F.; Ravi, S.; Ré, C. Low-dimensional hyperbolic knowledge graph embeddings. *arXiv* **2020**, arXiv:2005.00545. [[CrossRef](#)]
33. Balažević, I.; Allen, C.; Hospedales, T. Tucker: Tensor factorization for knowledge graph completion. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 5185–5194.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.