Transactions on Computational and Scientific Methods | Vo. 4, No. 9, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Entity-Aware Graph Neural Modeling for Structured Information Extraction in the Financial Domain

Yixian Wang
The University of Chicago, Chicago, USA
amywangyx@uchicago.edu

Abstract: This paper addresses the challenging problem of structured information extraction in the financial domain. It proposes a joint extraction model that integrates graph neural networks with an entity-aware mechanism. The goal is to improve the recognition of entities and relations in complex financial texts. The method constructs a Syntax-Semantics Hybrid Graph by incorporating dependency syntax and semantic co-occurrence relations. It models words, entities, and contextual information in a graph structure, effectively capturing long-range dependencies and hidden connections between non-contiguous entities. At the same time, the model introduces an Entity-Aware Representation Enhancement mechanism. Based on a pre-trained language model, it strengthens the explicit representation of entities. This guides the model to focus on core semantic areas during sequence encoding, improving the accuracy of entity boundary detection and relation extraction. The proposed method demonstrates strong performance across multiple evaluation tasks. It shows high stability and robustness, especially in handling high entity density, input noise perturbation, and adaptation to different pre-trained models. Comparative and ablation experiments confirm the effectiveness and complementarity of syntactic-semantic structure fusion and entity-aware enhancement in financial text extraction. The structured modeling and semantic representation strategy proposed in this paper provides a technical foundation for deeper understanding and high-quality knowledge construction in financial corpora. It also offers a new research perspective for extraction tasks in complex language scenarios.

Keywords: Information extraction, graph neural network, entity perception modeling, financial text processing

1. Introduction

In the era of information explosion, financial news, as a key information carrier, widely disseminates content related to corporate activities, policy trends, and market conditions. It significantly influences investor behavior and market fluctuations. Especially in capital markets, financial news is regarded as a critical source of signals that reveal potential risks and market opportunities[1]. As the volume of information continues to grow, manual methods can no longer extract valuable knowledge efficiently or accurately. This drives the trend toward automated analysis of financial news. Entity extraction and relation recognition, as core tasks in financial information extraction, play a fundamental role in building financial knowledge graphs, supporting decision-making, and enhancing intelligent investment systems[2].

However, financial texts are highly specialized, domain-specific, and ambiguous[3]. They contain numerous entities such as organizations, individuals, products, and indicators. The relationships among these entities are often embedded in complex semantic structures. Traditional natural language processing methods, which rely on lexical and syntactic rules, struggle to capture deep semantic associations and contextual interactions. Moreover, financial news features diverse writing styles, multiple reporting perspectives, fuzzy entity boundaries, and frequent ambiguities. These characteristics raise higher demands for entity recognition and relation modeling, urging researchers to explore more flexible and efficient modeling approaches[4].

In recent years, graph neural networks have emerged as powerful tools for modeling graph-structured data in natural language processing[5]. Unlike traditional sequence models, graph neural networks can naturally represent the structural relationships among entities. They can capture long-range dependencies that span across sentences or even documents. This structural advantage is particularly important for modeling complex interactions among multiple entities in financial news. Graph neural networks offer more reasonable representations for handling nested entities, implicit relations, and contextual dependencies. By transforming financial text into graph structures, they can improve the accuracy of entity extraction and uncover deeper semantic relationships, supporting financial event modeling and dynamic knowledge construction[6,7].

Furthermore, building an entity-relation recognition system tailored to financial news has far-reaching implications for the development of financial knowledge graphs. As key infrastructure for intelligent systems in the financial domain, knowledge graphs play a central role in information integration, knowledge inference, and risk warning. The quality of entity and relation recognition directly determines the accuracy and reliability of the graph[8]. Graph neural network-based modeling not only enhances the semantic expressiveness of information extraction but also enables unified knowledge structuring across diverse sources and formats of financial news. This improves the completeness and timeliness of the knowledge graph, boosting its application potential in financial technology, regulatory technology, and investment support[9].

From a broader perspective, advancing intelligent extraction of entities and relations in financial news is a key path for integrating natural language processing with financial engineering. It also represents a challenge and breakthrough in modeling complex semantic information. As information interaction in financial markets accelerates, building a system with high accuracy and robustness for information understanding helps financial institutions enhance their insight capabilities. It also provides more transparent and intelligent market information services for individual investors. In this context, research on entity extraction and relation recognition in financial news using graph neural networks holds both clear academic value and significant practical importance.

2. Related work

2.1 Graph Neural Networks

Graph neural networks, as a deep learning method for processing graph-structured data, have demonstrated remarkable performance in recent years across natural language processing, recommendation systems, and computer vision. The core idea is to transmit information between nodes in a graph through a message passing mechanism. This helps capture contextual structural features of the nodes. Compared with traditional sequence models, graph neural networks can effectively model high-order relationships between nodes and their neighbors in non-Euclidean space. This gives them stronger representational power. Especially when dealing with data that contain complex dependency structures, graph neural networks can enhance modeling depth and semantic capture due to their flexible topological representations[10].

In natural language processing, graph neural networks have been widely used in tasks such as text classification, relation extraction, named entity recognition, and event extraction[11,12]. Text, as a form of sequential data, can be transformed into graph structures through dependency parsing, coreference resolution, and entity linking. This enables a structured representation of semantic dependencies. In entity extraction and relation recognition tasks, words, phrases, and entities in text can be represented as nodes in a graph. Edges can express syntactic, semantic, or contextual dependencies. Graph neural networks can effectively model multi-hop relations and long-distance dependencies. This helps address the limitations of traditional models in capturing non-contiguous relationships[13,14]. Moreover, graph structures can integrate information from multiple heterogeneous sources, such as knowledge bases, context windows, and domain dictionaries. This further enhances the richness of representation and the accuracy of extraction[15].

In the financial domain, text data often contain highly structured information and complex semantic relationships. Examples include trading relationships between financial institutions, acquisition behaviors between companies, and interactions between policies and markets. These types of information are naturally suited to graph-based modeling. The use of graph neural networks provides technical support for building more accurate and interpretable financial information extraction systems[16,17]. By constructing semantic or event graphs among entities, models can not only identify key entities but also understand dynamic semantic relations between them. This improves both the logical consistency and semantic completeness of information extraction. Therefore, applying graph neural networks to entity and relation modeling in financial news is not only a methodological innovation, but also offers a new direction for the intelligent development of financial information processing systems.

2.2 Entity Extraction

Entity extraction is one of the fundamental tasks in natural language processing[18]. It aims to identify and extract meaningful entity units from unstructured text, such as names of people, organizations, locations, time expressions, and numerical values. In financial texts, entity extraction is particularly important. It involves diverse financial concepts, including company names, stock codes, financial indicators, product names, counterparties, and legal terms. Financial texts are typically formal in style but rich in specialized terminology. Their semantic expressions are often complex[19,20]. Entities may be ambiguous or nested, which increases the difficulty of boundary detection and type classification. Therefore, entity extraction is considered a foundational step in financial information extraction systems. It directly affects the quality of subsequent tasks such as information linking, relation recognition, and knowledge graph construction[21].

Early methods for entity extraction relied on handcrafted rule templates, dictionary matching, or statistical learning models[22]. These approaches were effective in controlled scenarios. However, when applied to real-world financial texts with diverse language patterns and complex contextual relationships, they often showed poor adaptability and weak generalization[23]. With the development of deep learning, especially the wide adoption of pre-trained language models, entity extraction has gradually evolved from traditional sequence labeling to more semantically driven approaches. BiLSTM-CRF architectures were once among the mainstream methods. These models used sequential context to identify entities. However, they struggled with long texts, cross-sentence entities, and nested structures. This limited their effectiveness in complex financial settings[24].

In recent years, researchers have proposed more flexible and semantically aware models for entity extraction in the financial domain[25]. These include question-matching based models, hierarchical extraction frameworks, and graph neural network methods with structural awareness. These approaches overcome the limitations of linear sequence models. They retain contextual information while improving the model's ability to understand structural relationships among entities. For example, in a piece of financial news involving several companies, individuals, and transactions, the semantic boundaries of entities are not always clearly defined by specific words. Relying only on local context often fails to determine correct entity attribution. Structured modeling approaches introduce global information among entities. This helps improve the accuracy and consistency of extraction. Therefore, the continuous development of entity extraction techniques provides a solid foundation for information understanding in financial texts and offers essential technical support for the automatic construction and semantic analysis of financial knowledge.

3. Method

This study proposes a method for entity extraction and relation recognition in financial news by integrating Graph Neural Networks (GNN) and Pre-trained Language Models (PLM). The goal is to improve the model's ability to understand complex semantic structures and cross-sentence relations in financial texts. Compared with traditional approaches, this method introduces two key innovations. First, it constructs a Syntax-Semantics Hybrid Graph (SSHG) based on syntactic dependencies and semantic co-occurrence. This strategy builds high-quality semantic graph structures within the text. It allows the model to better capture non-contiguous entities and their contextual associations. Second, it introduces an Entity-Aware Representation Enhancement (EARE) mechanism. During encoding, the model is explicitly guided to focus on key domain-specific entities. This interaction at the entity level improves the accuracy of extraction and the robustness of relation recognition. These two improvements work together to strengthen the model's ability to capture implicit semantic and structural information in financial news. They provide a solid technical foundation for subsequent financial knowledge representation and event construction. The architecture of the overall model is illustrated in Figure 1.

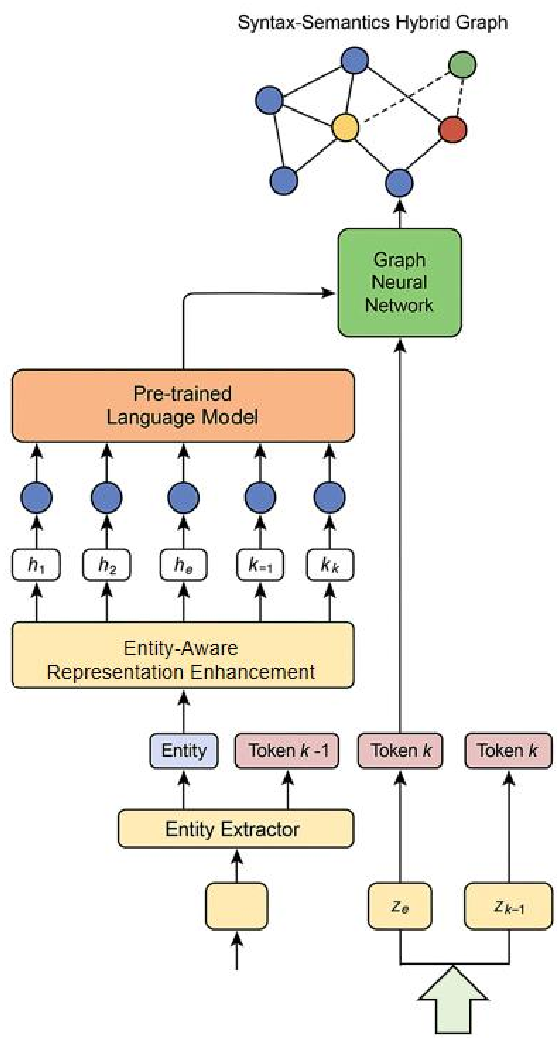

Figure 1. Overall model architecture diagram

3.1 Syntax-Semantics Hybrid Graph

In this study, we propose a hybrid graph structure named the Syntax-Semantics Hybrid Graph (SSHG), which combines syntactic structure and semantic co-occurrence information to enhance the model's ability to capture complex relationships between entities in financial text. In the constructed graph, both entities and contextual units are treated as nodes, while syntactic dependencies and semantic associations are introduced as edges. This dual-edge design explicitly encodes the intricate linguistic connections among entities in a structured form, thereby addressing the limitations of traditional sequence-based models in capturing long-distance and non-local dependencies.

By incorporating both sentence-level syntax and document-level semantic co-occurrence, SSHG enables the model to represent non-continuous and cross-sentence entity associations within a unified graph framework. This not only improves the expressive capacity of the graph neural network but also enhances its semantic coverage and contextual awareness. As a result, the model becomes more effective in identifying latent relations and subtle interactions that are otherwise difficult to capture using linear architectures. The detailed architecture of the SSHG module is illustrated in Figure 2.

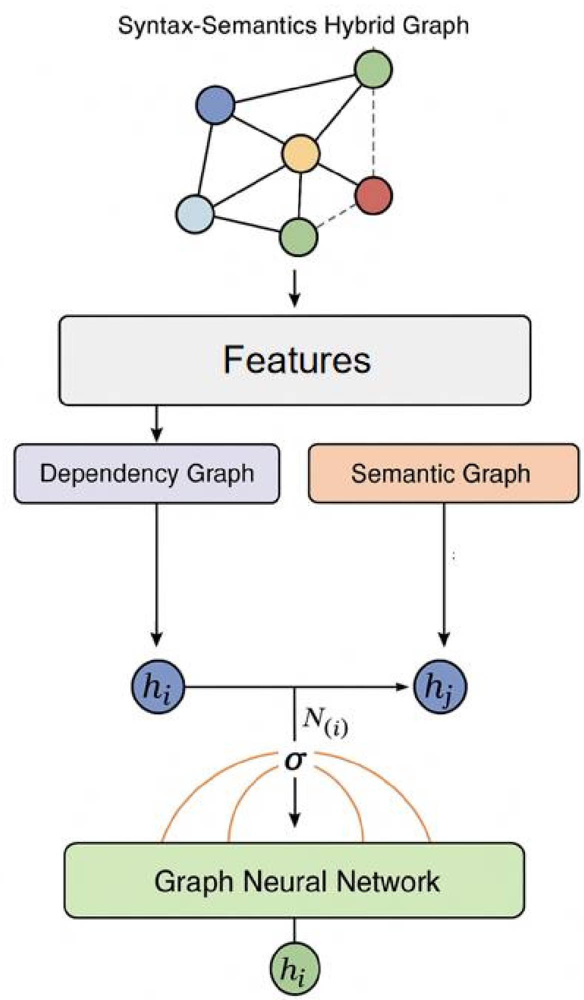

Figure 2. SSHG module architecture

In the graph construction process, given an input text sequence $X = \{x_1, x_2, ..., x_n\}$, we first use the syntactic parsing tool to obtain the dependency graph $G_{syn} = (V, E_{syn})$, where V represents the word node set and

E_{syn} represents the edge set based on the syntactic dependency relationship. At the same time, we introduce the semantic co-occurrence graph $G_{syn} = (V, E_{syn})$, where the edge E_{syn} is obtained by calculating the co-occurrence frequency of word pairs in the local window, TF-IDF similarity, or the similarity between pre-trained vectors. The final hybrid graph $G = (V, E_{syn} \cup E_{sem})$ contains both syntactic edges and semantic edges, thus forming a more comprehensive semantic structure.

To enhance the distinguishability of nodes in the graph, we introduce an entity-aware labeling function $\delta(v_i)$ to identify whether a node is an entity, which is defined as follows:

$$\delta(v_i) = \begin{cases} 1, & \text{if } v_i \in \text{Entity Set} \\ 0, & \text{otherwise} \end{cases}$$

In addition, the edge weight function $w(e_{ij})$ is defined as:

$$w(e_{ij}) = \begin{cases} \lambda_1, & \text{if } e_{ij} \in \mathrm{E}_{syn} \\ \lambda_2 \cdot \cos(e_i, e_j), & \text{if } e_{ij} \in \mathrm{E}_{sem} \end{cases}$$

e_i represents the word vector, and λ_1 and λ_2 are hyperparameters. Such a weighting mechanism helps to improve the semantic selectivity of edges.

In the GNN propagation phase, we update the representation of each node v_i in the form of message aggregation:

$$h_i^{(l+1)} = \sigma\Big(\sum_{j \in N(i)} w(e_{ij}) \cdot W^{(l)} h_j^{(l)}\Big)$$

Where $N(i)$ represents the neighbor set of node v_i, $W^{(l)}$ is the parameter matrix of the l-th layer, and σ is the activation function. Through multi-layer propagation, each entity node can gradually aggregate its context and structural information to achieve cross-sentence and cross-relation knowledge modeling. The introduction of this graph structure lays a structured semantic foundation for subsequent entity extraction and relationship recognition.

3.2 Entity-Aware Representation Enhancement

In natural language processing tasks, entities usually carry the most core semantic information in a sentence. Especially in financial texts, entities such as company names, amounts, and event objects have a critical impact on task completion. In order to make full use of these important semantic units, this study introduces the Entity-Aware Representation Enhancement mechanism, which guides the model to focus on key semantic nodes by modeling and strengthening entity information in the encoding stage, thereby improving the overall representation and expression capabilities. Based on the pre-trained language model, this module explicitly injects entity features, which helps to improve the accuracy and consistency of subsequent graph modeling and relationship recognition stages. Its module architecture is shown in Figure 3.

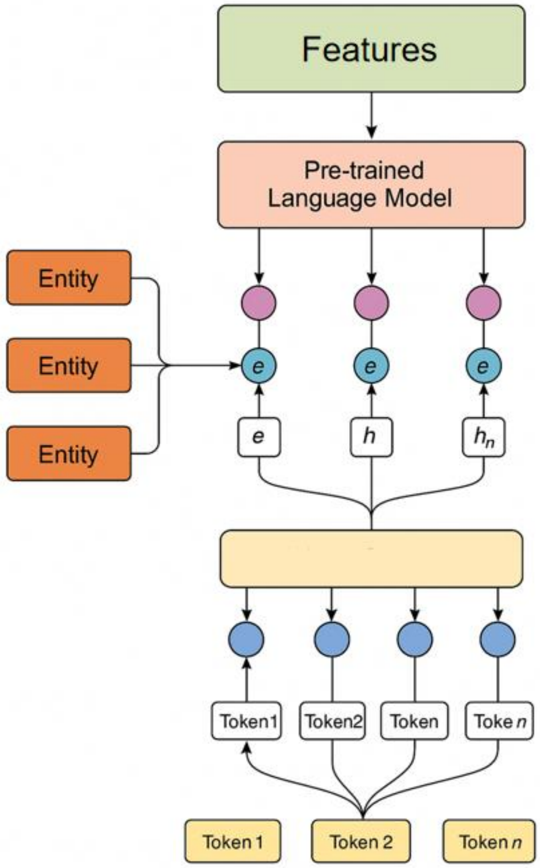

Figure 3. EARA module architecture

In the specific implementation, assume that the input text is $X = \{x_1, x_2, ..., x_n\}$ and its corresponding initial representation is $H = \{h_1, h_2, ..., h_n\}$. We first use the entity recognizer to annotate the entity set $E = \{e_1, e_2, ..., e_m\}$, and each entity is represented by its position in the text (s_i, t_i). Then, the extracted entity representation e_i can be defined as the pooling operation of its contained word vectors:

$$e_i = Pooling(h_{s_i}, ..., h_{t_i})$$

To achieve entity-guided contextual enhancement, we introduce an entity attention mechanism to fuse the entity representation with the original sequence representation, defined as:

$$h_j = h_j + \sum_{i=1}^{m} a_{ij} \cdot e_i$$

The attention weight a_{ij} represents the correlation between the jth word and the i-th entity, and is calculated as:

$$a_{ij} = \frac{\exp(h_j^T W_a e_i)}{\sum_{k=1}^{m} \exp(h_j^T W_a e_k)}$$

In addition, to further enhance the semantic influence of entities, we introduced an entity gating mechanism in the hidden layer to dynamically adjust the representation of each position, which is defined as follows:

$$\hat{h}_j = \sigma(W_g[h_j; e^*]) \otimes h_j + (1 - \sigma(W_g[h_j; e^*])) \otimes e^*$$

Where e^* represents the aggregation of global entity vectors, σ is the Sigmoid function, and $\otimes$ is the element-wise multiplication operation. This mechanism ensures that when encoding sentences, entity information can affect the global semantic expression in an explicit and controllable way, which is especially important for modeling key entities in long texts.

Finally, all enhanced representations $\{\hat{h}_1, \hat{h}_2, ..., \hat{h}_n\}$ will be passed as input to the graph neural network module to participate in information propagation on the semantic graph. The entity enhancement module not only improves the model's sensitivity to key semantics, but also provides a more semantically rich representation input for subsequent structured modeling, thus laying the foundation for complex relationship extraction and event construction.

4. Experimental Results

4.1 Dataset

This study uses the FinRE dataset as the primary source for entity recognition and relation extraction tasks. FinRE is a financial-domain information extraction dataset collected from a large number of financial news reports. It covers typical financial events such as corporate activities, investment and financing, mergers and acquisitions, and executive changes. The dataset is designed for joint modeling of entities and relations. It features strong domain specificity and structural complexity, making it suitable for evaluating a model's performance in real financial contexts.

The FinRE dataset includes annotations for various types of entities, such as companies, individuals, job titles, amounts, and time expressions. It also provides detailed classifications of semantic relations between entities, including "holds position", "investment event", "acquisition target", and "funding round". Each sample contains the original news text along with its corresponding structured annotations. This allows the dataset to be used for named entity recognition, relation classification, and event modeling tasks.

The dataset has high annotation quality and is based on reliable sources. It is widely used in financial natural language processing research. Its complex nested entities, diverse relation types, and highly specialized financial terminology present challenges for domain-adaptive modeling. At the same time, it offers an effective benchmark for financial text understanding tasks.

4.2 Experimental Results

1) *Comparative experimental results*

This paper first gives the comparative experimental results, as shown in Table 1.

Table 1: Comparative experimental results

Method	**Accuracy**	**Precision**	**Recall**
BERT-CRF[26]	83.1%	79.6%	81.3%

PL-Marker[27]	84.7%	82.5%	83.6%
PURE[28]	85.4%	83.1%	84.2%
UniRE[29]	86.0%	84.2%	85.1%
Ours	88.5%	86.7%	87.6%

As shown in Table 1, the proposed model demonstrates clear advantages in entity extraction and relation recognition tasks within the financial domain. It achieves the best overall performance in terms of accuracy and extraction completeness. Compared with the traditional BERT-CRF model, our approach improves accuracy by 5.4 percentage points. It also increases precision and recall by 7.1% and 6.3%, respectively. This indicates that the model not only identifies more correct entities and relations but also shows stronger robustness and generalization.

Further analysis shows that the proposed method still outperforms strong baseline models recently introduced in top conferences, such as PL-Marker, PURE, and UniRE. Methods like PURE and UniRE have adopted joint modeling strategies and some context-aware mechanisms. However, they still face challenges in capturing long-range dependencies and deep semantics in financial texts, which are often syntactically complex and densely packed with overlapping entities. The Syntax-Semantics Hybrid Graph introduced in this study addresses this gap. It enables the model to capture non-contiguous entities and their latent relations, thus improving overall extraction performance.

In addition, the Entity-Aware Representation Enhancement mechanism plays a key role in performance improvement. In financial texts, entities often carry the core semantic content of the sentence. By explicitly introducing entity information during the encoding phase of the pre-trained model, the mechanism helps guide model attention toward critical areas. This leads to gains in both precision and recall. The simultaneous improvement of these two metrics confirms that the mechanism enhances both entity boundary detection and relation classification.

Overall, the superior performance of the proposed method results from a dual-focus strategy that integrates structural information and semantic awareness. By incorporating enhanced modeling of entities and context at both the encoding layer and the graph neural propagation layer, the model improves its understanding of financial language details. It also strengthens its ability to reason about higher-level structural relationships. This design proves to be more adaptive and accurate when dealing with diverse and tightly coupled entity relations in financial scenarios.

2) *Ablation Experiment Results*

This paper also further gives the results of the ablation experiment, and the experimental results are shown in Table 2.

Table 2: Ablation Experiment Results

Method	**Accuracy**	**Precision**	**Recall**
BaseLine	84.2%	81.0%	82.1%
+SSHG	86.1%	83.5%	84.3%
+EARA	85.6%	83.1%	83.8%
Ours	88.5%	86.7%	87.6%

As shown in the ablation results in Table 2, the two core modules proposed in this study—Syntax-Semantics Hybrid Graph (SSHG) and Entity-Aware Representation Enhancement (EARE)—both contribute significantly to the overall performance of the model. The baseline model, without structured modeling or entity-awareness mechanisms, achieves an accuracy of 84.2%. Its precision and recall are 81.0% and 82.1%, respectively. This indicates that it has certain limitations in identifying complex semantic relations.

When only the SSHG module is added, the model's accuracy increases to 86.1%. Precision and recall also rise to 83.5% and 84.3%. This change confirms the effectiveness of structured semantic graphs in modeling syntactic dependencies and semantic co-occurrence. By introducing graph structures, the model captures deeper semantic associations among entities. It performs better in handling long-range dependencies and non-contiguous relations. This enhances contextual support for entity and relation decisions.

On the other hand, adding the EARE module also improves model performance. The accuracy reaches 85.6%, showing that the entity-awareness mechanism effectively guides the model to focus on key information. This module enhances entity representations explicitly. It helps the model recognize important semantic units during sequence encoding. As a result, overall extraction precision improves, especially in sentences with multiple co-occurring entities.

The full model, which integrates both SSHG and EARE, achieves an accuracy of 88.5%. Both precision and recall reach their highest levels. This indicates that the two mechanisms are functionally complementary. Structured modeling provides semantic-level graph support. The entity enhancement mechanism improves semantic focus at the representation level. Together, they enable the model to generalize better and understand deeper semantic structures in complex financial texts during entity extraction and relation recognition tasks.

3) Experiment on the adaptability of different pre-trained language models to downstream extraction tasks

This paper also gives an experiment on the adaptability of different pre-trained language models to downstream extraction tasks, and the experimental results are shown in Figure 4.

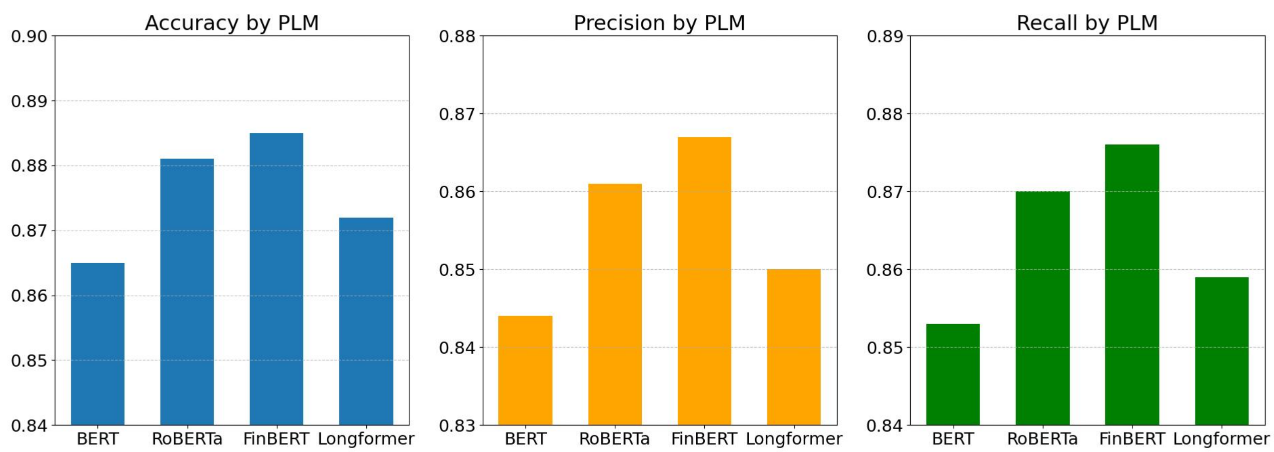

Figure 4. Experiment on the adaptability of different pre-trained language models to downstream extraction tasks

As shown in the results of Figure 4, different pre-trained language models (PLMs) exhibit significant performance differences in downstream entity extraction and relation recognition tasks. Overall, FinBERT achieves the best or near-best results in terms of accuracy, precision, and recall. This indicates that domain-specific pre-trained models have clear advantages in modeling financial texts. Although BERT has a certain level of representational power as a base model, it shows limitations in adapting to financial semantic contexts.

RoBERTa performs consistently and outperforms BERT across all three evaluation metrics, including accuracy, precision, and recall. Its enhanced ability to model deeper contextual dependencies plays a key role in improving the understanding of financial texts. Although RoBERTa is not specifically pre-trained on financial-domain corpora, its superior capacity for semantic representation enables it to maintain high levels of accuracy and recall, even when dealing with complex and syntactically varied language structures. This result indicates that general-purpose enhanced language models, when equipped with more advanced architectural designs, can still exhibit strong adaptability and performance, even in specialized domains without explicit domain-specific pretraining.

The strength of FinBERT further confirms the compatibility between entity-aware mechanisms and domain-specific language models. In the architecture of this study, the Entity-Aware Representation Enhancement module depends on accurate identification of dense entities such as financial terms, company names, and numerical expressions. FinBERT benefits from higher coverage of such entities in its pre-training data, making it a better base encoder for this task. Its leading performance in precision and recall also suggests that the model is more effective in capturing financial entities with clear semantic boundaries and explicit pragmatic targets.

In contrast, Longformer, despite its architectural advantage in handling long-range dependencies and extended input sequences, demonstrates slightly weaker performance in this particular task. One possible explanation is that financial news and announcements, although not excessively long, rely more heavily on precise semantic interpretation of entity relationships rather than on modeling paragraph-level continuity. As a result, the ability to process long texts alone is not sufficient to compensate for the lack of adaptation to domain-specific semantics.

This observation suggests that structural capability must be matched with domain relevance to achieve optimal performance. These experimental results underscore the importance of selecting pre-trained language models that are well aligned with the target semantic domain. They also provide further evidence supporting the effectiveness of the proposed approach, which combines domain-specific pre-trained representations with structure-enhanced strategies, allowing the model to better capture fine-grained financial semantics and complex entity interactions.

4) Analysis of the influence of entity density on information extraction accuracy

This paper also provides a detailed analysis of the impact of entity density on the accuracy of information extraction. By examining how varying levels of entity concentration affect model performance, the study highlights the challenges posed by densely populated semantic structures. The corresponding experimental results, which illustrate this relationship across different density levels, are presented in Figure 5.

As shown in Figure 5, entity density has a significant impact on the performance of information extraction tasks. When the entity density in the text is low to moderate, the model can stably extract key entities and their semantic relations. The F1 scores are 0.873 and 0.882, respectively, indicating high extraction accuracy and completeness. This shows that in contexts with moderate semantic load, the model has strong semantic discrimination and contextual understanding. It can effectively perform joint modeling of entities and relations.

However, as entity density increases, model performance declines. In the "Very High" density setting, the F1 score drops to 0.835. This result suggests that in texts with dense entities, the model faces greater ambiguity.

When nested, overlapping, or co-referent entities are present, it becomes harder for the model to accurately distinguish boundaries and types. This affects the final recognition results. High density may also cause mismatches in relation recognition, which reduces overall extraction accuracy.

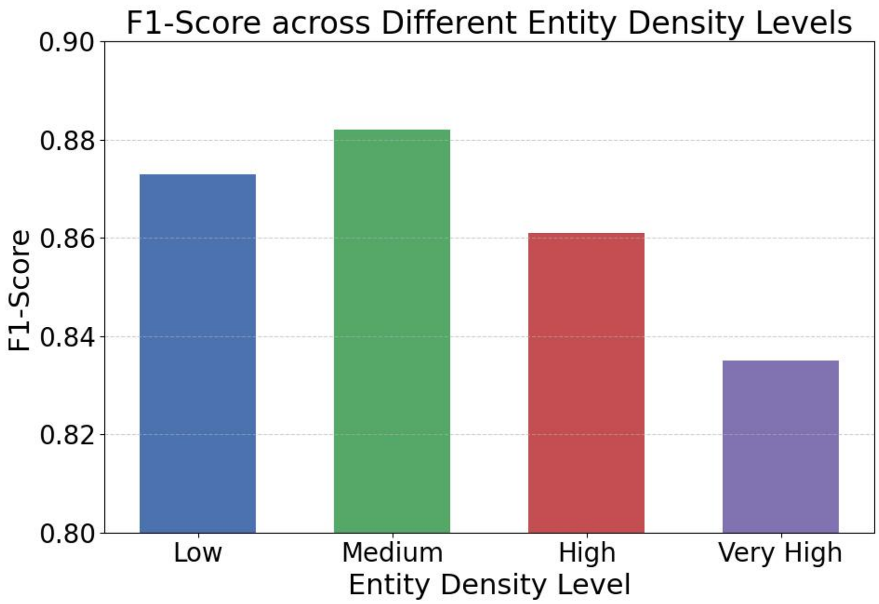

Figure 5. Analysis of the influence of entity density on information extraction accuracy

This trend further confirms the necessity of introducing the Syntax-Semantics Hybrid Graph and Entity-Aware mechanisms in this study. Although complex entity structures pose challenges, structured graph modeling and entity-aware mechanisms help reduce the interference caused by high density. They guide the model to focus on the structural and semantic links among high-weight entities. This helps maintain performance stability. These findings also show that the proposed model has strong capabilities in handling complex texts.

In summary, higher entity density results in more complex and intertwined semantic structures, which significantly increase the demands placed on a model's ability to abstract meaningful patterns and accurately identify entity boundaries. In financial texts and other domains where entities frequently overlap or co-occur, traditional sequence-based modeling approaches are often inadequate for capturing such intricacies. To effectively handle these challenges, multi-level methods that integrate structural awareness with explicit entity modeling provide a more robust solution. These approaches offer greater practical value and broader application potential, particularly in real-world information extraction tasks that involve dense and semantically rich content.

5) Model robustness evaluation under input noise perturbations

This paper also provides an evaluation of the model's robustness under varying levels of input noise perturbations, aiming to assess its stability in less controlled or noisy environments. The experimental results, which demonstrate how performance changes with increasing noise intensity, are presented in Figure 6.

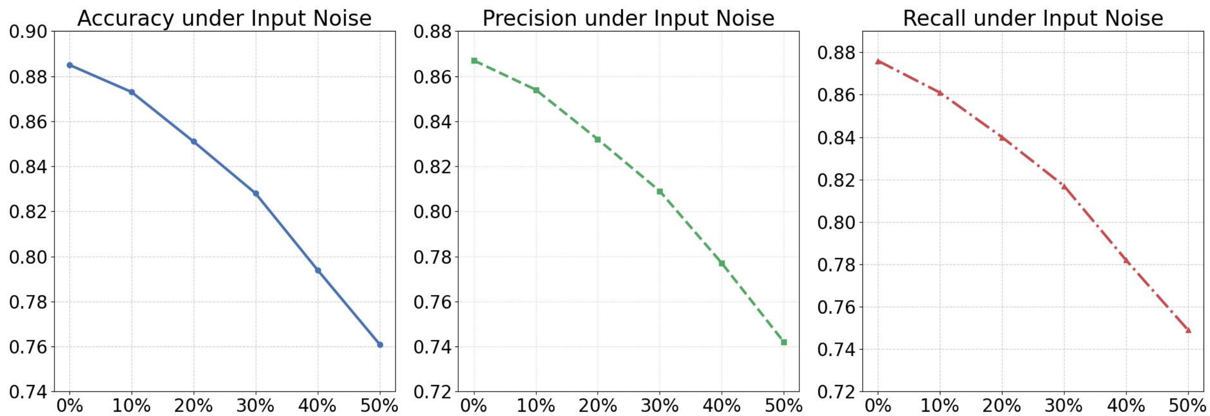

Figure 6. Model robustness evaluation under input noise perturbations

As shown in the results of Figure 6, the model's performance in terms of accuracy, precision, and recall steadily declines as input noise increases. This indicates that the model is sensitive to the completeness of input information. When no noise is added (0%), the model maintains a high level of performance. As the noise level reaches 50%, accuracy drops by more than 12 percentage points. This suggests that when input structures are disrupted, the model's ability to identify key information is significantly affected.

The precision curve shows a clear drop in the accuracy of entity boundary detection under noise. This implies that noise not only reduces the recognizability of entities but also increases the risk of incorrect predictions. Since the proposed method relies heavily on entity representation enhancement and structured graph construction, any form of input variation may damage the integrity of the graph. This, in turn, disrupts the semantic propagation of entities within the graph.

The recall curve also shows a downward trend, indicating that noise causes the model to miss key entities or relations. It fails to capture the original semantic content effectively. In financial texts, semantic understanding often depends on accurate expression of specific entities such as names, dates, and monetary values. Once these are disturbed, the model struggles to build correct contextual associations, leading to lower recall.

Overall, the experimental results confirm that the proposed method has strong extraction ability but depends on input stability. In scenarios with frequent noise or low-quality text, model performance may decline. This finding highlights the need to enhance robustness in future model design. Possible directions include adversarial training, data augmentation, or robust graph structure optimization, to improve system reliability in real-world financial text processing.

6) Loss function changes with epoch

Finally, this paper gives a graph of the loss function changing with epoch, and the experimental results are shown in Figure 7.

As shown in the results of Figure 7, the proposed method demonstrates good convergence and stability during training. Both train loss and validation loss drop rapidly within the first 10 epochs. This indicates that the model quickly captures the main semantic structures between entities and relations. It also shows strong feature fitting ability in the early stages. These results suggest that the model's structural modeling and entity enhancement mechanisms in the financial domain begin to contribute positively from the early training phase, improving extraction efficiency.

As training progresses, both losses gradually stabilize after epoch 12. The gap between validation loss and train loss remains small, with no obvious sign of overfitting. This indicates that the model has good generalization ability in structurally complex and semantically dense financial texts. It maintains consistent predictions on unseen data. The results reflect the robustness of the graph neural network mechanisms and pre-trained representations in aligning semantic features.

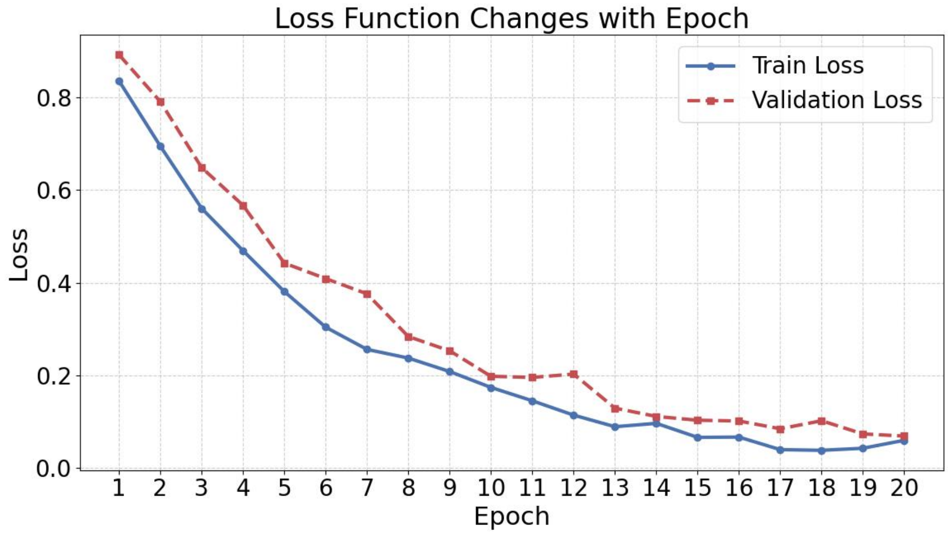

Figure 7. Loss function changes with epoch

It is worth noting that during the later training phase (epoch 15 to 20), the validation loss shows slight fluctuations. However, the overall trend remains stable, and no significant performance degradation occurs. Given the model's integration of graph structure and entity information, these fluctuations may be related to the complexity of entity distributions or the presence of long-tail relation types in the input. Such variations fall within an acceptable range of local disturbances.

Overall, the loss curves confirm the reliability and training efficiency of the proposed method in financial information extraction tasks. The strong convergence trend and consistent validation results further demonstrate the practical value of structural modeling and entity enhancement. These components help improve model robustness and representation power, providing effective support for entity and relation recognition in complex financial language scenarios.

5. Conclusion

This study addresses the problem of structured information extraction from complex texts in the financial domain. It proposes a deep learning model that integrates a Syntax-Semantics Hybrid Graph with an Entity-Aware Representation Enhancement mechanism. The method combines syntactic structures with semantic co-occurrence to build a context-aware graph neural network. It also introduces an entity enhancement strategy to improve the modeling of semantic dependencies between entities and relations. In various experimental tasks, the model shows strong extraction performance and robustness. It performs especially well in scenarios with overlapping entities, long-range dependencies, and semantic ambiguity. This research advances joint extraction modeling from both methodological and technical perspectives. It explores how graph structures and pre-trained language models can work together, enriching the study of financial natural language processing. The proposed architecture is both generalizable and scalable. It applies not only to financial news but also has the potential to transfer to other sub-domains such as announcements, financial reports, and regulatory documents. This method supports the development of more accurate and real-time financial knowledge graphs, which provide key support for tasks such as investment analysis, regulatory monitoring, and intelligent search.

From an application perspective, this research holds practical value for both financial technology (FinTech) and regulatory technology (RegTech). In financial markets, automated semantic understanding and information structuring are fundamental for high-frequency trading, risk assessment, and market analysis. In regulatory contexts, the model can help identify high-risk expressions, track corporate behavior, and strengthen compliance checks. As financial text volume continues to grow, manual processing is no longer sufficient to meet the dual demands of efficiency and accuracy. Intelligent semantic parsing systems are becoming a core part of future infrastructure. Looking ahead, financial texts will continue to become more diverse in structure and more complex in semantics. This raises new challenges for the flexibility and generalization of extraction models. Future work may incorporate cross-document graph construction, enhanced temporal modeling, or multimodal information such as charts, tables, and images to improve understanding of heterogeneous data. There is also room for deeper exploration of model robustness and security. Research on adversarial perturbations and domain transfer can further advance the development of mature and practical financial language processing systems.

References

[1] Repke, Tim, and Ralf Krestel. "Extraction and representation of financial entities from text." Data Science for Economics and Finance: Methodologies and Applications. Cham: Springer International Publishing, 2021. 241-263.

[2] Hillebrand, Lars, et al. "Kpi-bert: A joint named entity recognition and relation extraction model for financial reports." 2022 26th International Conference on Pattern Recognition (ICPR). IEEE, 2022.

[3] Zhao, Lingyun, et al. "A BERT based sentiment analysis and key entity detection approach for online financial texts." 2021 IEEE 24th International conference on computer supported cooperative work in design (CSCWD). IEEE, 2021.

[4] Loukas, Lefteris, et al. "FiNER: Financial numeric entity recognition for XBRL tagging." arXiv preprint arXiv:2203.06482 (2022).

[5] Aguda, T., Siddagangappa, S., Kochkina, E., Kaur, S., Wang, D., Smiley, C., & Shah, S. (2024). Large language models as financial data annotators: A study on effectiveness and efficiency. arXiv preprint arXiv:2403.18152.

[6] Sinha, Ankur, et al. "SEntFiN 1.0: Entity-aware sentiment analysis for financial news." Journal of the Association for Information Science and Technology 73.9 (2022): 1314-1335.

[7] Huang, Allen H., Hui Wang, and Yi Yang. "FinBERT: A large language model for extracting information from financial text." Contemporary Accounting Research 40.2 (2023): 806-841.

[8] Jacobs, Gilles, and Véronique Hoste. "SENTiVENT: enabling supervised information extraction of company-specific events in economic and financial news." Language Resources and Evaluation 56.1 (2022): 225-257.

[9] Mao, Xuting, et al. "Financial fraud detection using the related-party transaction knowledge graph." Procedia Computer Science 199 (2022): 733-740.

[10] Corso, Gabriele, et al. "Graph neural networks." Nature Reviews Methods Primers 4.1 (2024): 17.

[11] Veličković, Petar. "Everything is connected: Graph neural networks." Current Opinion in Structural Biology 79 (2023): 102538.

[12] Zhou, Yu, et al. "Graph neural networks: Taxonomy, advances, and trends." ACM Transactions on Intelligent Systems and Technology (TIST) 13.1 (2022): 1-54.

[13]Wu, Lingfei, et al. "Graph neural networks: foundation, frontiers and applications." Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining. 2022.

[14]Gao, Chen, et al. "A survey of graph neural networks for recommender systems: Challenges, methods, and directions." ACM Transactions on Recommender Systems 1.1 (2023): 1-51.

[15]Reiser, Patrick, et al. "Graph neural networks for materials science and chemistry." Communications Materials 3.1 (2022): 93.

[16]Liu, Rui, et al. "Federated graph neural networks: Overview, techniques, and challenges." IEEE transactions on neural networks and learning systems (2024).

[17]Nasar, Zara, Syed Waqar Jaffry, and Muhammad Kamran Malik. "Named entity recognition and relation extraction: State-of-the-art." ACM Computing Surveys (CSUR) 54.1 (2021): 1-39.

[18]Bai, Lu, et al. "Clinical entity extraction: comparison between MetaMap, cTAKES, CLAMP and Amazon Comprehend Medical." 2021 32nd Irish Signals and Systems Conference (ISSC). IEEE, 2021.

[19]Tuo, Meimei, and Wenzhong Yang. "Review of entity relation extraction." Journal of Intelligent & Fuzzy Systems 44.5 (2023): 7391-7405.

[20]Whittaker, Edward, and Ikuo Kitagishi. "Large language models for simultaneous named entity extraction and spelling correction." arXiv preprint arXiv:2403.00528 (2024).

[21]Bose, Priyankar, et al. "A survey on recent named entity recognition and relationship extraction techniques on clinical texts." Applied Sciences 11.18 (2021): 8319.

[22]Shah-Mohammadi, Fatemeh, Wanting Cui, and Joseph Finkelstein. "Comparison of ACM and CLAMP for Entity Extraction in Clinical Notes." 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). IEEE, 2021.

[23]Zheng, Hengyi, et al. "PRGC: Potential relation and global correspondence based joint relational triple extraction." arXiv preprint arXiv:2106.09895 (2021).

[24]Xu, Tiefeng, et al. "Power grid fault event reasoning based on entity extraction model of knowledge map." Journal of Intelligent & Fuzzy Systems 45.5 (2023): 8479-8488.

[25]Qiao, Bo, et al. "A joint model for entity and relation extraction based on BERT." Neural Computing and Applications (2022): 1-11.

[26]Souza, Fábio, Rodrigo Nogueira, and Roberto Lotufo. "Portuguese named entity recognition using BERT-CRF." arXiv preprint arXiv:1909.10649 (2019).

[27]Ye, Deming, et al. "Packed levitated marker for entity and relation extraction." arXiv preprint arXiv:2109.06067 (2021).

[28]Hang, Tingting, et al. "Joint extraction of entities and overlapping relations using source-target entity labeling." Expert Systems with Applications 177 (2021): 114853.

[29]Wang, Yijun, et al. "UniRE: A unified label space for entity relation extraction." arXiv preprint arXiv:2107.04292 (2021).