

Trustworthy-constraint Deep Graph Learning For Enterprise Financial Risk Prediction

Wenting Ma

School of Finance and Economic, Zhengzhou University of Science and Technology, Zhengzhou, Henan, China 450064

Corresponding author. E-mail: wentingmazst@163.com

Received: March 14, 2025; Accepted: April 05, 2025

Financial risk prediction for enterprises is a hot topic in the field of financial technology. Deep learning-Based methods achieve encouraging the financial risk prediction performance due to the power ability of the feature learning. However, there exists two issues in deep learning-Based methods. (1) Current methods fail to accurately model the complex relationships among listed companies in the real financial market. (2) Current methods lack credible estimates in the decision-making process, which leads to the questionable reliability of the decision-making results. To this end, a trustworthy-constraint deep graph learning network (TDGL-net) is proposed to achieve the above goal, which includes the multi-view feature encoding, the heterogeneous graph information aggregation, the trustworthy decision-making mechanism. Specifically, TDGL-net integrates a multi-dimensional bilinear neural tensor and Transformer into a unified multi-view feature encoding to learn comprehensive representation. Then, TDGL-net models heterogeneous graph information aggregation via the cross-category association and the intra-category association, to capture complex inter-enterprise relationships with momentum spillover effects for enhancing the representation discrimination. Additionally, TDGL-net incorporates a trustworthy decision-making mechanism to adaptively integrate information from deep embedding representations and graph embedding representations according to enterprise-specific contexts, improving decision reliability and accuracy. Ultimately, extensive experimental evaluations on the real-world dataset reveal that TDGL-net delivers state-of-the-art performance in predicting enterprise financial risk.

Keywords: Financial risk prediction; Heterogeneous graph information aggregation; Trustworthy decision-making mechanism

© The Author(s). This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are cited.

[http://dx.doi.org/10.6180/jase.202601_29\(1\).0011](http://dx.doi.org/10.6180/jase.202601_29(1).0011)

1. Introduction

Financial risk prediction for enterprises has always been a crucial research direction in the field of financial technology [1, 2]. For instance, in stock fluctuation prediction, investors are constantly striving to forecast the future trend of a listed company's stock, aiming to maximize profits in the volatile financial market. Over the years, this task has continuously inspired researchers' strong interest in developing more effective forecasting models. Machine learning methods have shown great potential in financial risk tasks.

In the research fields of finance and artificial intelligence, existing studies typically rely on the historical prices of the enterprise itself (such as opening price, closing price, volume) and apply various time - series analysis techniques to predict the trend of stocks [3-5]. According to the efficient market hypothesis, which posits that financial markets are informationally efficient, some researchers also mine more indicative features from market-outside data, such as online media, including news information and social media [6, 7].

However, when focusing on market-outside data, many

studies often overlook the impact of stock fluctuations from related companies, known as the momentum spillover effect in finance [8]. Although recent research has started to model the stock momentum spillover effect using graph neural networks, most of these studies only consider the simple explicit relationships between related companies. This simplified approach inevitably fails to accurately model the complex relationships among listed companies in the real financial market. In reality, the relationships between listed companies are influenced by diverse factors such as industry competition patterns and supply chain upstream and downstream relationships. Modeling based solely on simple explicit relationships is like viewing through a tube, only capturing a part of the problem and failing to grasp the whole picture [9].

Moreover, the current financial risk decision-making process also has certain limitations, mainly reflected in the lack of credible estimates in the decision-making process, which leads to the questionable reliability of the decision-making results [10, 11]. In financial risk prediction, a credible estimate is crucial for decision-makers as it provides a more accurate risk assessment and a more reliable basis for decision-making. Without credible estimates in the decision-making process, it may lead decision-makers to misjudge the risks and make unreasonable decisions. In such cases, financial institutions may suffer huge losses due to decision-making failures when facing the complex and changing financial market, and may even trigger systemic financial risks, posing a threat to the stable operation of the entire financial market.

To address these challenges, a trustworthy-constraint deep graph learning network (TDGL-net) is proposed to achieve the above goal, which includes the multi-view feature encoding, the heterogeneous graph information aggregation, the trustworthy decision-making mechanism. Specifically, TDGL-net integrates a multi-dimensional bilinear neural tensor and Transformer into a unified multi-view feature encoding framework to capture multi-dimensional interactions between corporate financial information and media sentiment, thereby providing a more comprehensive representation of historical information. Through cross-category association learning and intra-category association learning, TDGL-net can aggregate heterogeneous graph information, learn complex inter-enterprise relationships, and capture momentum spillover effects to enhance the discriminative power of financial information representations. Additionally, TDGL-net incorporates a trustworthy decision-making mechanism based on subjective logic and uncertainty estimation, which adaptively integrates information from deep embedding representations and graph

embedding representations according to enterprise-specific contexts. This approach improves decision reliability and accuracy, mitigates over-confidence issues, and delivers more dependable financial risk predictions. Ultimately, extensive experimental evaluations on the real-world dataset reveal that TDGL-net delivers state-of-the-art performance in predicting enterprise financial risk.

The key contributions of TDGL-net include:

- A unified multi-view feature encoding framework that captures multi-dimensional interactions between corporate financial information and media sentiment, providing a more comprehensive representation of historical information.
- Cross-category and intra-category association learning mechanisms that effectively aggregate heterogeneous graph information, learn complex inter-enterprise relationships, and capture momentum spillover effects.
- A trustworthy decision-making mechanism that adaptively integrates information from deep embedding representations and graph embedding representations according to enterprise-specific contexts, improving decision reliability and accuracy.

The organization of TDGL-net is structured as follows: Section 2 provides a comprehensive analysis of the proposed TDGL-net, detailing its fundamental elements and functional processes. Section 3 presents a thorough empirical evaluation of TDGL-net's effectiveness using benchmark datasets, alongside an in-depth comparison with leading-edge techniques. Section 4 summarizes the key contributions of TDGL-net and outlines potential directions for future research and enhancement.

2. Method

Let (G, X, S) denote financial information dataset of N enterprises, where G is the Enterprise-CEO dual relationship heterogeneous graph, $X = \{x_i^t = \{op_i^t, cp_i^t, hp_i^t, lp_i^t, tv_i^t\}_i^N\}$ denotes history information of on the t -th day of the i -th enterprise including opening price, closing price, highest price, lowest price, and trading volume (tv), and $S = \{s_i^t = \{pos_i^t, neg_i^t, dif_i^t\}_i^N\}$ denotes the positive media sentiment, the negative media sentiment, and media sentiment difference on the t -th day of the i -th enterprise. The goal of this paper is to utilize the financial information dataset (G, X, S) of enterprises to achieve the enterprise financial risk prediction. To this end, a trustworthy-constraint deep graph learning network (TDGL-net) is proposed to achieve the above goal, which includes the multi-view feature encoding, the heterogeneous graph information aggregation,

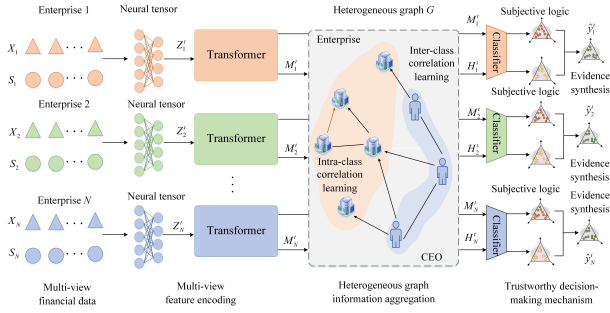


Fig. 1. The architecture of TDGL-net. The multi-view feature encoding module integrates multi-dimensional bilinear neural tensors and Transformer networks to capture dynamic interactions between financial indicators and media sentiment. The heterogeneous graph information aggregation module employs dual-channel attention networks to model cross-category entity relationships and intra-category market dynamics. Finally, the trustworthy decision-making mechanism leverages Dirichlet distribution-based uncertainty quantification to adaptively fuse deep semantic embeddings with graph structural representations, ensuring robust and reliable financial risk predictions.

the trustworthy decision-making mechanism, as shown in Fig. 1.

2.1. Multi-view feature encoding

Generally, financial information is influenced by multi-view time-series market signals. Considering the strong temporal dynamics of corporate financial information, historical states are useful for predicting future trends. Therefore, when forecasting the financial risk of company i on date t , it should account for market signals from the past several days. Specifically, TDGL-net first captures multi-view interactions between historical information and media sentiment. Then, TDGL-net feeds the fusion features into a Transformer block and takes the final hidden state as the sequential embedding of enterprise i to preserve spatio-temporal dependencies.

More specifically, TDGL-net introduce a multi-dimensional bilinear neural tensor layer to correlate the company's financial historical information with the media sentiment towards the company across multiple dimensions:

$$z_i^t = \text{Relu}(x_i^t W_T s_i + \mathcal{V} \begin{bmatrix} x_i \\ s_i \end{bmatrix} + b_T), \quad (1)$$

where Relu is the activation function, $W_T \in \mathbb{R}^{5 \times 3 \times 4}$ is the trainable tensor, $\mathcal{V} \in \mathbb{R}^{8 \times 4}$ is the learnable parameter, b_T represents the bias term. Then, TDGL-net inputs multi-

view fusion feature into Transformer to capture spatio-temporal information over the past T days:

$$m_i^t = \text{Transformer}(z_i^{t-T}, z_i^{t-T+1}, \dots, z_i^t), \quad (2)$$

where m_i^t denotes the sequential embedding of enterprise i .

2.2. Heterogeneous graph information aggregation

In real markets, the fluctuations in a company's financial information are often influenced by related companies, a phenomenon known as the momentum spillover effect in finance. To address this challenge, TDGL-net leverages a Deep Graph Learning architecture, motivated by the need to accurately model the complex interdependencies and momentum spillover effects among enterprises. Traditional methods often fail to capture these intricate relationships, leading to suboptimal predictions. Deep Graph Learning, with its ability to handle heterogeneous and dynamic graph structures, provides a powerful framework for uncovering hidden patterns and interactions in financial data. Consequently, based on the Enterprise-CEO dual relationship heterogeneous graph, TDGL-net employs a dual-attention strategy to model both inter- and intra-class correlations between enterprises and CEOs. This approach not only captures the direct relationships within the same category (e.g., interactions between companies) but also the cross-category relationships (e.g., affiliations between companies and their CEOs). By learning relational embeddings that incorporate momentum spillover signals, TDGL-net enhances the discriminative power of financial information representations, leading to more accurate and reliable financial risk predictions.

Inter-class correlation learning aims to handle the interactions between enterprises and CEOs. Specifically, given an enterprise entity $u \in U$ and an executive entity $e \in E$, TDGL-net first projects their embeddings into a common space since they are heterogeneous types of entities via a linear layer:

$$\begin{aligned} q'_u &= W_u q_u \\ q'_e &= W_e q_e \end{aligned} \quad (3)$$

where q'_u and q_u represent the original and projection features of u , respectively. Here, the initial feature of q_u can be obtained through m_i^t and the initial feature of q_e can be obtained by averaging the sequential embedding of enterprise where they serve.

For enterprise entity u , different types of inter-class relationships provide different semantics for its embedding, and similarly, different entities with the same relationship also do so. Consequently, TDGL-net employs entity-level and relationship-level attention mechanisms to hierarchi-

cally aggregate signals from neighbors of other types to the target u .

Given the relationship $r_i \in R_{inter}$ between u and e , where R_{inter} denotes the set of inter-class relationships, the importance of executive e to company u is denoted as $w_{ue}^{r_i}$:

$$w_{ue}^{r_i} = \text{att}_{entity}(q'_u, q'_e; r_i) = \text{LeakyReLU}(\mathbf{a}_{r_i}^\top \cdot [q'_u \| q'_e]), \quad (4)$$

where \mathbf{a}_{r_i} is a learnable weight vector. $\|$ denotes the concatenation operation. LeakyReLU is the non-linear activation function. Then, TDGL-net fuses neighbors of different types with a specific relationship r_i based on the importance $w_{ue}^{r_i}$:

$$q_u^{r_i} = \text{Relu}(\sum_{e \in \mathcal{N}_{inter}^{r_i}(u)} w_{ue}^{r_i} \cdot q'_e) \quad (5)$$

where $\mathcal{N}_{inter}^{r_i}(u)$ denotes the sets of executives under the relationship r_i .

Once all the relationship embeddings $q_u^{r_i}$ are obtained, TDGL-net uses the relationship-level attention to fuse them together to get the cross-class relationship embedding z_u of entity u . First, the importance e^{r_i} of each relationship for the final prediction is calculated as follows:

$$e^{r_i} = \frac{\exp\left(\frac{1}{|U|} \sum_{u \in U} W_q \cdot q_u^{r_i}\right)}{\sum_{r_j \in R_{inter}} \exp\left(\frac{1}{|U|} \sum_{u \in U} W_q \cdot q_u^{r_j}\right)}. \quad (6)$$

where W_q is the learnable parameter. By fusing all the relationship embeddings, we obtain the cross-class relationship embedding z_u :

$$z_u = \sum_{r_i \in R_{inter}} e^{r_i} \cdot q_u^{r_i} \quad (7)$$

Intra-class correlation learning aims to learn the interactions between entities of the same type. Specifically, given a relationship $r_k \in R_{intra}$ starting from the enterprise u , where R_{intra} denotes the set of all intra-class relationships of u , we can obtain the neighbor set $\mathcal{N}_{intra}^{r_k}(u)$ based on intra-class relationships. Each intra-class relationship represents a semantic interaction, which can be encoded using relationship-specific attention. We first calculates the entity \bar{u} relative to entity u under the relationship r_k :

$$\alpha_{u\bar{u}}^{r_k} = \frac{\exp(\text{LeakyRelu}(\mathbf{a}_{r_k}^\top \cdot [W_k z_u \| W_k z_{\bar{u}}]))}{\sum_{u' \in \mathcal{N}_{intra}^{r_k}(u)} \exp(\text{LeakyRelu}(\mathbf{a}_{r_k}^\top \cdot [W_k z_u \| W_k z_{u'}]))} \quad (8)$$

where z_u and $z_{\bar{u}}$ represent the output representations of the inter-class attention network. W_k represents the learnable weight matrix, shared by each node of the same type. \mathbf{a}_{r_k} is the node-level attention weight vector under the relationship r_k . $\mathcal{N}_{intra}^{r_k}(u)$ represents the intra-class neighbors of node u under the relationship r_k . Then, given the relationship r_k , the embedding $h_u^{r_k}$ of entity u can be calculated as

follows:

$$h_u^{r_k} = \text{Relu}\left(\sum_{\bar{u} \in \mathcal{N}_{intra}^{r_k}(u)} \alpha_{u\bar{u}}^{r_k} \cdot W_k z_{\bar{u}}\right) \quad (9)$$

Next, we can perform relationship-level attention mechanisms to fuse entity embeddings to obtain the relationship embedding h_u :

$$h_u = \sum_{r_k \in R_{intra}} \left(\frac{\exp\left(\frac{1}{|U|} \sum_{u \in U} W_h \cdot h_u^{r_k}\right)}{\sum_{r_l \in R_{intra}} \exp\left(\frac{1}{|U|} \sum_{u \in U} W_h \cdot h_u^{r_l}\right)} \right) \cdot h_u^{r_k} \quad (10)$$

where W_h represents the learnable parameters.

2.3. Trustworthy decision-making mechanism

TDGL-net designs the trustworthy decision-making mechanism to aggregate information between deep embedding representations and graph embedding representations, which dynamically fuses decision information from different views according to the specificity of enterprises.

In generally, traditional classifiers usually use the softmax function to convert continuous outputs into class probabilities, i.e., $p = [p_1, \dots, p_K]$, where $\sum_{k=1}^K p_k = 1$. Mathematically, p can be considered as the parameters of a multinomial distribution describing K mutually exclusive events. Although the softmax function is effective, it often leads to overconfidence. To alleviate this issue, TDGL-net introduces the evidence method to learn the subjective logic $P = \{\{p_1, \dots, p_K\}, \mu\}$ of classification, where $\sum_{k=1}^K p_k \geq 0$ represents the classification probability, μ represents the overall credibility, and $\sum_{k=1}^K p_k + \mu = 1$. Then, the Dirichlet distribution (which can be considered as the conjugate prior of the multinomial distribution) is used to infer the subjective logic, achieving uncertainty estimation for each view.

Specifically, the probability density function of the Dirichlet distribution is defined as follows:

$$\text{Dirichlet}(p | \theta) = \begin{cases} \frac{1}{\beta(\theta)} \prod_{i=1}^K p_i^{\theta_i-1} & \text{if } p \in \Lambda_K \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

where $\theta = [\theta_1, \dots, \theta_K]$ represents the Dirichlet distribution parameters.

$$\Lambda_K = \left\{ p \mid \sum_{i=1}^K p_i = 1, 0 \leq p_1, \dots, p_K \leq 1 \right\} \quad (12)$$

where Λ_K represents the $(K-1)$ -dimensional simplex. $\beta(\theta)$ represents the K -dimensional beta function. Then, based on the Dirichlet distribution parameters, the sample's classification probability is:

$$p_k = \frac{\theta_k - 1}{\sum_{k=1}^K \theta_k} \quad (13)$$

To obtain the Dirichlet distribution, the softplus activation function is used to construct the classifier, so that

the positive outputs are considered as the evidence vector $v = \{v_1, \dots, v_K\}$:

$$\theta_k = v_k + 1 \quad (14)$$

Based on the variational theory, the classification loss function is designed to optimize the classifier:

$$L_c^d = L_1^d + \lambda L_2^d \quad (15)$$

where λ is the balance hyperparameter. The first term is defined as follows:

$$L_1^d = \sum_{k=1}^K y_k (D(\theta_0^d) - D(\theta_k^d)) \quad (16)$$

where $\theta_0^d = \sum_{k=1}^K \theta_k$ represents the Dirichlet strength. y_k is the k -th element of the label y , represented as a one-hot vector. $D(\cdot)$ is the digamma function. At the same time, a reasonable Dirichlet distribution should be concentrated on the corresponding vertex on the simplex, which means that the parameters of the Dirichlet distribution should be as close to 1 as possible except for the parameter corresponding to the correct label. Therefore, the prior regularization term L_2^d is introduced as:

$$L_2^d = D_{KL}(\text{Dirichlet}(p^d | \tilde{\theta}^d) \| \text{Dirichlet}(\bar{p}^d | [1, \dots, 1])) \quad (17)$$

where $\tilde{\theta}^d = y + (1 - y) \odot \theta^d$. $\text{Dirichlet}(\bar{p}^d | [1, \dots, 1])$ is the uniform Dirichlet distribution. By minimizing the classification loss L_c^d , the Dirichlet distribution $\text{Dirichlet}(p^d | \tilde{\theta}^d)$ can be obtained. Then, the subjective logic parameters for deep embedding representations are $L = \{\{p_1^d, \dots, p_K^d\}, \mu^d\}$.

After capturing the Dirichlet distributions from deep embedding representations and graph embedding representations, i.e., $\text{Dirichlet}(p^d | \tilde{\theta}^d)$ and $\text{Dirichlet}(p^s | \tilde{\theta}^s)$, multiple Dirichlet distributions are adaptively integrated based on the uncertainty of the Dirichlet distribution using the evidence theory. Specifically, the subjective logic of two views $L^d = \{\{p_1^d, \dots, p_K^d\}, \mu^d\}$ and $L^s = \{\{p_1^s, \dots, p_K^s\}, \mu^s\}$ are integrated into a fused subjective logic $L = \{\{p_1, \dots, p_K\}, \mu\}$, which is defined by the evidence synthesis formula $L = L^d \oplus L^s$ as follows:

$$p_k = \frac{1}{1 - \sum_{i \neq k} p_i^d p_i^s} (p_k^d p_k^s + p_k^d u^s + p_k^s u^d) \quad (18)$$

$$u = \frac{1}{1 - \sum_{i \neq j} p_i^s p_j^s} u^d u^s$$

The corresponding fused Dirichlet distribution parameters $\text{Dirichlet}(p | \theta)$ are obtained through $\theta_k = p_k \times \sum_{i=1}^K \theta_i + 1$. To obtain a reliable fusion Dirichlet distribution, the fusion

classification loss is designed as:

$$L_c = L_1 + \lambda L_2$$

$$= \sum_{k=1}^K y_k (D(\theta_0) - D(\theta_k)) \quad (19)$$

$$+ \lambda D_{KL}(\text{Dirichlet}(p | \tilde{\theta}) \| \text{Dirichlet}(\bar{p} | [1, \dots, 1]))$$

The trustworthy decision-making mechanism, which combines the evidence neural classification network parameterized by Dirichlet parameters and the theory of evidence fusion, can dynamically identify the risky views for decision-making and utilize the informative views in the final decision.

Thus, TDGL-net utilizes the following loss $Loss$ to optimize the overall architecture with the multi-view feature encoding, the heterogeneous graph information aggregation, the trustworthy decision-making mechanism:

$$Loss = L_c + \alpha L_c^d + \beta L_c^s \quad (20)$$

where α and β are trade-off parameters. After the optimization is completed, a multi-view voting mechanism is designed to achieve the final financial risk prediction \hat{y}_i .

$$\hat{y}_i^t = \arg \max(p_i^t + p_i^{d;t} + p_i^{s;t}), i = 1, 2, \dots, N \quad (21)$$

where \hat{y}_i^t denotes prediction of the i -th enterprise on the t -th data

3. Results and discussion

3.1. Dataset and Setup

Dataset and metric: We have collected publicly available financial data from companies listed on the China Securities Index, which we have named the "SMP" dataset. The SMP dataset encompasses historical price data from November 21, 2017, to December 31, 2019, 50 online financial news articles published across four major financial websites (Sina, Hexun, Sohu, and Eastmoney), as well as 158 companies and 205 CEOs. Following [11], four metrics are employed to evaluate the performance of TDGL-net, including Precision (Pre), Area Under the Curve (AUC), Recall (Rec), and the F1 Score (F1). Pre measures the proportion of true positive predictions out of all positive predictions made by the model. AUC evaluates the model's ability to distinguish between classes by measuring the area under the receiver operating characteristic curve. Rec assesses the proportion of true positive predictions out of all actual positive cases. F1 is the harmonic mean of precision and recall, providing a balanced measure of both metrics.

Implementation Details: TDGL-net is implemented using the PyTorch and PyTorch Geometric libraries. Training

is conducted on a GTX 1080 GPU, with each session averaging 1.5 hours. In the experiments, the lookback window T is set to 20, the learning rate is set to 0.0003, and the dimensions of both the deep embedding representations and the infection graph embedding representations are set to 128. The number of training epochs is configured to 500. To prevent overfitting, an early stopping strategy is employed based on the AUC (Area Under the Curve) from the validation set.

3.2. Performance Comparison

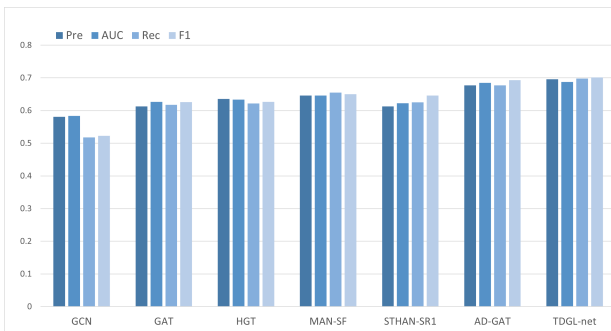


Fig. 2. Comparison of TDGL-net to top-performing methods on SMP in terms of four metrics.

Six methods are used to validate the performance of TDGL-net, including GCN [11], GAT [12], HGT [13], MAN-SF [14], STHAN-SR1 [15], and AD-GAT [16]. As shown in Fig. 2, experimental results demonstrate that TDGL-net comprehensively outperforms mainstream models in enterprise financial risk prediction tasks. On core evaluation metrics, its precision (0.6958) surpasses baseline models such as GCN (0.5812) and GAT (0.6125) by more than 8.3%-19.7%, while its AUC value (0.6875) and recall rate (0.6977) both achieve the highest levels among comparable methods, validating the model's balanced recognition capability for positive and negative samples. Notably, TDGL-net's F1 score (0.7012) exceeds the 0.7 threshold, representing a 3.5% improvement over the best baseline, AD-GAT (0.6773), which fully demonstrates the model's synergistic optimization of precision and recall. This series of significant gaps indicates that TDGL-net not only enhances single-metric performance but also achieves comprehensive performance breakthroughs across multiple dimensions.

TDGL-net demonstrates superior performance through three architecturally innovative components that form a comprehensive technical framework. The system's advancements manifest in three key dimensions: 1) Multi-perspective Feature Encoding Architecture - By synergizing bilinear neural tensor layers with Transformer networks,

the model achieves multi-dimensional spatiotemporal interaction between financial indicators and media sentiment analytics, effectively overcoming traditional models' deficiencies in capturing dynamic temporal patterns. 2) Heterogeneous Graph Fusion Mechanism - This pioneering framework concurrently models cross-category entity relationships (e.g., enterprise-CEO affiliations) and intra-category market dynamics (e.g., inter-firm momentum spillover effects) through dual-channel attention networks, significantly enhancing feature discriminability while preserving complex semantic relationships. 3) Trustworthy Decision Framework - Incorporating Dirichlet distribution-based uncertainty quantification, this module adaptively fuses deep semantic embeddings with graph structural representations, substantially reducing model overfitting induced by financial data noise through probabilistic regularization.

3.3. Ablation Study

Four ablation experiments are conducted to verify three component effectiveness in TDGL-net. Specifically, (1) TDGL-net utilizes the multi-view feature encoding (MFE) to learn representations of enterprises, and then performs the vanilla classification loss to achieve financial risk prediction. (2) TDGL-net utilizes the heterogeneous graph information aggregation (HIA) to learn representations of enterprises, and then performs the vanilla classification loss to achieve financial risk prediction. (3) TDGL-net utilizes the multi-view feature encoding to learn representations of enterprises, and then performs the trustworthy decision-making mechanism (TDM) to achieve financial risk prediction. (4) TDGL-net utilizes the multi-view feature encoding and the heterogeneous graph information aggregation to learn representations of enterprises, and then performs the vanilla classification loss to achieve financial risk prediction.

Table 1 ablation analysis reveals synergistic mechanisms of TDGL-net's three modules: (1) MFE vs HIA: Multi-view Feature Fusion (MFE) demonstrates superior representation over graph aggregation (HIA) (F1 0.5189 vs 0.4613), yet their integration achieves a 28.7% F1 boost (0.6675), highlighting feature-topology complementarity. (2) TDM Efficacy: The Trustworthy Decision Module (TDM) delivers 23% performance gain through uncertainty modeling, with an additional 5.05% improvement over the MFE+HIA baseline, confirming its confidence-calibration value. (3) System Synergy: The components form an organic system—MFE integrates multi-dimensional attributes, HIA mines topological relationships, and TDM ensures robust risk decisions. Progressive performance enhancements validate the "feature fusion → relation mining → trustworthy

Table 1. Ablation results on SMP dataset in terms of four metrics.

MFE	HIA	TDM	Pre	AUC	Rec	F1
✓			0.5267	0.5153	0.5164	0.5189
	✓		0.4607	0.4536	0.4645	0.4893
✓		✓	0.6260	0.6494	0.6370	0.6381
✓	✓		0.6524	0.6691	0.6543	0.6675
✓	✓	✓	0.6958	0.6875	0.6977	0.7012

Table 2. Sensitivity verification on SMP dataset in terms of four metrics.

α	β	λ	Pre	AUC	Rec	F1
0.1	0.1	0.1	0.6267	0.6153	0.6151	0.6189
0.1	1	0.1	0.6607	0.6536	0.6664	0.67489
0.1	1	1	0.6556	0.6449	0.6437	0.6481
1	0.1	0.1	0.6924	0.6857	0.6965	0.6980
1	0.1	1	0.6812	0.6832	0.6888	0.6907
1	1	0.1	0.6958	0.6875	0.6977	0.7012

Table 3. Impact of representation dimension on TDGL-net performance on the SMP dataset.

z	m	h	Pre	AUC	Rec	F1
64	64	64	0.6725	0.6643	0.6789	0.6752
128	128	128	0.6958	0.6875	0.6977	0.7012
256	256	256	0.6843	0.6765	0.6856	0.6867
512	512	512	0.6634	0.6547	0.6689	0.6656

decision" architecture. Functional coupling drives collaborative optimization, advancing financial risk assessment to new heights.

3.4. Sensitivity verification

Table 2 presents the results of experiments conducted to evaluate the effects of different parameter, i.e., α , β , and λ , values on TDGL-net's performance. As shown in the table, increasing the value of α generally enhances the model's performance. When α is increased from 0.1 to 1, the precision increases from 0.6267 to 0.6958, the AUC from 0.6153 to 0.6875, the recall from 0.6151 to 0.6977, and the F1 score from 0.6189 to 0.7012. This indicates that a higher α value improves the integration of deep and graph embeddings. Similarly, increasing β also positively impacts performance. For instance, setting β to 1 (with α at 0.1 and λ at 0.1) results in improved precision, AUC, recall, and F1 score compared to β at 0.1. This suggests that a higher β value better captures the heterogeneous graph information. The parameter λ significantly affects performance. A lower λ value (0.1) generally yields better results than a higher value (1). For example, with α and β set to 1, a λ of 0.1 results in a pre-

cision of 0.6958, AUC of 0.6875, recall of 0.6977, and F1 score of 0.7012, which are the highest values achieved. In conclusion, the optimal parameter settings are $\alpha = 1$, $\beta = 1$, and $\lambda = 0.1$, providing the best balance and highest prediction accuracy for TDGL-net in enterprise financial risk prediction.

3.5. Representation analysis

Table 3 presents the impact of different representation dimensions on TDGL-net's performance. The dimensions of the multi-view fusion representation z , sequential representation m , and relationship representation h were varied together as 64, 128, 256, and 512. Results show that the model's performance peaks when dimensions are set to 128, achieving the highest F1 score of 0.7012. At 64 dimensions, while performance is reasonable, there is a noticeable drop in all metrics compared to 128 dimensions, indicating that 64 may not be sufficient to capture complex feature interactions. Increasing dimensions beyond 128 leads to diminishing returns, with performance declining at 256 and 512 dimensions. This suggests that higher dimensions may introduce noise or overfitting, harming predictive accuracy. Thus, choosing an appropriate representation dimension is crucial for balancing model complexity and performance.

4. Conclusion

This paper introduces TDGL-net, a novel deep graph learning framework for financial risk prediction, featuring three innovations: multi-view feature encoding with neural tensors and Transformers, heterogeneous graph aggregation via dual-attention mechanisms, and a subjective logic-based decision mechanism. These advancements enhance enterprise relationship modeling and prediction reliability. Future research directions include exploring high-order graph architectures, integrating multi-modal data, expanding applications to stock prediction and credit scoring, refining uncertainty quantification, developing explainable AI, and enhancing real-time prediction capabilities.

References

- [1] H. Xia, H. Ao, L. Li, Y. Liu, S. Liu, G. Ye, and H. Chai. "CI-STHPAN: pre-trained attention network for stock selection with channel-independent spatio-temporal hypergraph". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 38. 8. 2024, 9187–9195. DOI: [10.1609/aaai.v38i8.28770](https://doi.org/10.1609/aaai.v38i8.28770).
- [2] C. Xu, H. Huang, X. Ying, J. Gao, Z. Li, P. Zhang, J. Xiao, J. Zhang, and J. Luo, (2022) "HGNN: Hierarchical graph neural network for predicting the classification of price-limit-hitting stocks" *Information Sciences* 607: 783–798. DOI: [10.1016/j.ins.2022.06.010](https://doi.org/10.1016/j.ins.2022.06.010).
- [3] J. Gao, M. Liu, P. Li, J. Zhang, and Z. Chen, (2024) "Deep Multiview Adaptive Clustering With Semantic Invariance" *IEEE Transactions on Neural Networks and Learning Systems* 35(9): 12965–12978. DOI: [10.1109/TNNLS.2023.3265699](https://doi.org/10.1109/TNNLS.2023.3265699).
- [4] M. Leippold, Q. Wang, and W. Zhou, (2022) "Machine learning in the Chinese stock market" *Journal of Financial Economics* 145(2): 64–82. DOI: [10.1016/j.jfineco.2021.08.017](https://doi.org/10.1016/j.jfineco.2021.08.017).
- [5] P. Li, J. Gao, J. Zhang, S. Jin, and Z. Chen, (2022) "Deep Reinforcement Clustering" *IEEE Transactions on Multimedia*: DOI: [10.1109/TMM.2022.3233249](https://doi.org/10.1109/TMM.2022.3233249).
- [6] P. Li, Z. Chen, L. T. Yang, Q. Zhang, and M. J. Deen, (2017) "Deep convolutional computation model for feature learning on big data in internet of things" *IEEE Transactions on Industrial Informatics* 14(2): 790–798. DOI: [10.1109/TII.2017.2739340](https://doi.org/10.1109/TII.2017.2739340).
- [7] P. Li, A. A. Laghari, M. Rashid, J. Gao, T. R. Gadekallu, A. R. Javed, and S. Yin, (2022) "A deep multimodal adversarial cycle-consistent network for smart enterprise system" *IEEE Transactions on Industrial Informatics* 19(1): 693–702. DOI: [10.1109/TII.2022.3197201](https://doi.org/10.1109/TII.2022.3197201).
- [8] W. Bao, Y. Cao, Y. Yang, H. Che, J. Huang, and S. Wen, (2024) "Data-driven stock forecasting models based on neural networks: A review" *Information Fusion*: 102616. DOI: [10.1016/j.inffus.2024.102616](https://doi.org/10.1016/j.inffus.2024.102616).
- [9] R. Cheng and Q. Li. "Modeling the momentum spillover effect for stock prediction via attribute-driven graph attention networks". In: *Proceedings of the AAAI conference on artificial intelligence*. 35. 1. 2021, 55–62. DOI: [10.1609/aaai.v35i1.16077](https://doi.org/10.1609/aaai.v35i1.16077).
- [10] K. Du, R. Mao, F. Xing, and E. Cambria. "A dynamic dual-graph neural network for stock price movement prediction". In: *2024 International Joint Conference on Neural Networks (IJCNN)*. 2024, 1–8. DOI: [10.1109/IJCNN60899.2024.10650440](https://doi.org/10.1109/IJCNN60899.2024.10650440).
- [11] T. N. Kipf and M. Welling, (2016) "Semi-supervised classification with graph convolutional networks" *arXiv preprint arXiv:1609.02907*:
- [12] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, et al., (2017) "Graph attention networks" *stat* 1050(20): 10–48550.
- [13] Z. Hu, Y. Dong, K. Wang, and Y. Sun. "Heterogeneous graph transformer". In: *Proceedings of the web conference 2020*. 2020, 2704–2710. DOI: [10.1145/3366423.338002](https://doi.org/10.1145/3366423.338002).
- [14] R. Sawhney, S. Agarwal, A. Wadhwa, and R. Shah. "Deep attentive learning for stock movement prediction from social media text and company correlations". In: *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*. 2020, 8415–8426. DOI: [10.18653/v1/2020.emnlp-main.676](https://doi.org/10.18653/v1/2020.emnlp-main.676).
- [15] R. Sawhney, S. Agarwal, A. Wadhwa, T. Derr, and R. R. Shah. "Stock selection via spatiotemporal hypergraph attention network: A learning to rank approach". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 35. 1. 2021, 497–504. DOI: [10.1609/aaai.v35i1.16127](https://doi.org/10.1609/aaai.v35i1.16127).
- [16] R. Cheng and Q. Li. "Modeling the momentum spillover effect for stock prediction via attribute-driven graph attention networks". In: *Proceedings of the AAAI conference on artificial intelligence*. 35. 1. 2021, 55–62. DOI: [10.1609/aaai.v35i1.16077](https://doi.org/10.1609/aaai.v35i1.16077).