

SETN: Stock Embedding Enhanced with Textual and Network Information

Takehiro Takayanagi
The University of Tokyo
School of Engineering
Tokyo, Japan
m2021ttakayanagi@socsim.org

Hiroki Sakaji
The University of Tokyo
School of Engineering
Tokyo, Japan

Kiyoshi Izumi
The University of Tokyo
School of Engineering
Tokyo, Japan

Abstract—Stock embedding is a method for vector representation of stocks. There is a growing demand for vector representations of stock, i.e., stock embedding, in wealth management sectors, and the method has been applied to various tasks such as stock price prediction, portfolio optimization, and similar fund identifications. Stock embeddings have the advantage of enabling the quantification of relative relationships between stocks, and they can extract useful information from unstructured data such as text and network data. In this study, we propose stock embedding enhanced with textual and network information (SETN) using a domain-adaptive pre-trained transformer-based model to embed textual information and a graph neural network model to grasp network information. We evaluate the performance of our proposed model on related company information extraction tasks. We also demonstrate that stock embeddings obtained from the proposed model perform better in creating thematic funds than those obtained from baseline methods, providing a promising pathway for various applications in the wealth management industry.

Index Terms—Stock embedding, Stock analysis, Natural language processing, Graph embedding, Knowledge graph, Graph neural network

I. INTRODUCTION

Financial experts, such as fund managers, financial advisors, and traders, utilize the relative relationships between stocks, as well as individual stock characteristics, for their professional activities. For example, traders choose stock pairs with high positive correlations based on financial time series in pairs trading, and fund managers select stocks with similar themes to create thematic mutual funds. Therefore, it is important to quantify the relative relationships between stocks for quantitative stock analysis.

One method of capturing relative relationships between stocks is stock embedding. Stock embedding is a powerful tool to extract relative relationships between stocks from various data by expressing stocks with vector representations. Stock embedding has two advantages over existing stock analysis methods. First, it enables the quantification of relative relationships by calculating cosine similarities. Second, unstructured data, such as textual and network data, can be utilized through various embedding methodologies [1]–[5].

Financial text data such as annual reports, integrated reports, and earnings summaries are often utilized to obtain stock embeddings [6], [7]. Textual data are especially useful in

capturing individual stock information, as they contain a large amount of this information, including prevailing business situations, management discussions, and current business risks. Several text embedding methods have been employed to obtain stock embeddings from textual data [8]–[10]. Studies have shown that pretrained transformer-based models effectively learn stock embeddings from textual data [6].

Similarly, network data such as supplier-customer relationships, transaction networks, and fund-asset relationships are used to gain knowledge of stock embeddings [11]–[18]. Network data are useful in capturing cross-stock information, as graphs help represent the complex nature of financial systems with many components and sophisticated relations. Several graph embedding methods have been employed to learn stock embeddings from network data [1]–[3]. Studies have found that graph neural network (GNN) models are effective in learning stock embeddings from network data [15]–[18].

Ideally, stock embeddings should represent stocks from various perspectives. However, many studies do not consider it necessary to train models with both textual and network data, although both types of data have been shown to be effective for stock embedding [6], [7], [11]–[19]. Therefore, it is worth exploring a model that can jointly learn textual and network data.

In this study, we propose SETN, which jointly trains a transformer-based and GNN model. First, SETN further pre-trains a pretrained transformer-based model on a financial corpus to adapt the model to the financial domain [20]–[22]. Second, it learns textual information from the annual reports of a target company using a further pre-trained transformer-based model. It then creates a subgraph of the target stock from the network data and learns the network information using a GNN. Finally, we added a single-layer feedforward neural network to predict the sector and industry of a target stock.

We conducted experiments on the Japanese stock market to evaluate the performance of our proposed model. Our SETN model outperforms the baseline models in related company information extraction tasks, in which stocks in similar business sectors and industries are extracted. The results suggest the benefit of jointly learning both textual and network information. We also analyze the effect of graph type on performance, as there has been limited analysis of what type

of graph better captures network information. Furthermore, we investigate the effectiveness of jointly training a transformer-based and GNN models by separately comparing them with a transformer-based and GNN model. The results suggest that using a directed graph leads to higher performance, and that joint training contributes to higher performance. Furthermore, we introduce thematic fund creation tasks, in which we extract stocks with similar themes. We show that embeddings obtained using our proposed model extract themed stock better than embeddings from the baseline models.

The main contributions of our study are as follows:

- 1) We propose SETN. The results of experiments show our method outperforms baseline models on related company information extraction tasks.
- 2) We conduct a comparative study of graph types and learning architectures in stock embeddings.
- 3) We show that the stock embeddings obtained from SETN can be applied to the creation of thematic funds by introducing thematic fund creation tasks.

II. RELATED WORK

Stock embedding is a powerful tool used for vector representation of various aspects of stocks. It has been applied to multiple tasks, such as stock price prediction, portfolio optimization, and similar fund identification [6], [7], [11]–[19].

Stock embedding methods can be classified into two approaches based on the properties of the data used for embeddings. The first approach uses individual stock information [6], [7], [19]. Traditionally, numerical data, such as financial time series and technical indicators, have been used to obtain latent representations of the stock market. Chen *et al.* [19] extracted stock market embeddings from a financial time series using an attention-based long short-term memory (LSTM) model to predict the daily return ratio. Recently, textual data, such as annual reports and earnings summaries, have been utilized to obtain stock embeddings as textual data contain rich individual stock information. Transformer-based models effectively learn individual stock information from textual data [6], [7]. Ito *et al.* [6] extracted company stock embeddings from a company’s annual reports using bidirectional encoder representations from transformers (BERT) to obtain fine-grained industry information. Du *et al.* [7] extracted stock embeddings from news articles using BERT and achieved portfolio optimization.

The second approach utilizes cross-stock information [11]–[18]. This approach utilizes various network data to embed the topological information of the stocks. The network data used for stock embeddings include shareholding graphs [15], [16], news co-occurrence graphs [11], fund-asset relation graphs [12]–[14], and supplier-customer relation graphs [17]. Several graph embedding methods have been proposed for stock embedding. Wu *et al.* [11] obtained stock embeddings using matrix decomposition of news co-occurrence graphs for stock price predictions. Several researchers, such as [12]–[14] obtained stock embeddings by using the node2vec algorithm [3] on fund-asset graphs to identify similar funds and obtain

TABLE I: TOPIX17&TOPIX33, Sector and Industry Classification in Japanese Stock Market

TOPIX17	TOPIX33
FOODS	Fishery, Agriculture & Forestry
	Foods
ENERGY RESOURCES	Mining
	Oil and Coal Products
CONSTRUCTION&MATERIALS	Construction
	Metal Products
	Glass and Ceramics Products
RAW MATERIALS&CHEMICALS	Textiles and Apparels
	Pulp and Paper
	Chemicals
PHAMACEUTICAL	Pharmaceutical
AUTOMOBILES&TRANSPORTATION EQUIPMENT	Rubber Products
	Transportation Equipment
STEEL&NONFERROUS METALS	Iron and Steel
	Nonferrous Metals
MACHINERY	Machinery
ELECTRIC APPLIANCES&PRECISION INSTRUMENTS	Electric Appliances
	Precision Instruments
	Other Products
IT&SERVICES, OTHERS	Information & Communication
	Services
ELECTRIC POWER&GAS	Electric Power and Gas
	Land Transportation
	Marine Transportation
TRANSPORTATION&LOGISTICS	Air Transportation
	Warehousing and Harbor Transportation
COMMERCIAL&WHOLESALE TRADE	Wholesale Trade
RETAIL TRADE	Retail Trade
BANKS	Banks
	Securities and Commodities Futures
FINANCIAL(EXCEPT BANKS)	Insurance
	Other Financing Business
REAL ESTATE	Real Estate

new technical indicators. Recently, the GNN method has evolved and has been shown to be effective in extracting stock embeddings from cross-stock relationships [15]–[18].

Several studies have been conducted on the extraction of stock embeddings using either textual or network information, which have been shown to be useful for extracting stock embeddings [6], [7], [11]–[19]. However, notably, the number of studies that focus on using both textual and network information is quite limited. Some studies such as [23], [24] used both types of information by encoding textual information with a transformer-based model and using them as input to a GNN model. These studies differ from ours in two ways. First, they did not jointly train text embedding and graph embedding models. Specifically, they do not update the parameters for text embedding using existing pre-trained transformer-based models [10], whereas our approach fine-tunes a transformer-based model. Second, their study aimed to predict economic indicators and stock prices, while our study aimed to extract stock embeddings. Therefore, the effectiveness of learning stock embeddings from both textual and network data merits further investigation.

III. METHODOLOGY

In this section, we provide an overview of SETN, a stock embedding enhanced with textual information and network information that jointly trains a transformer-based and GNN model, followed by a detailed explanation of each component. As shown in Figure 1, the model approach consists of the following four steps. Step 1 extracts a subgraph centered on a

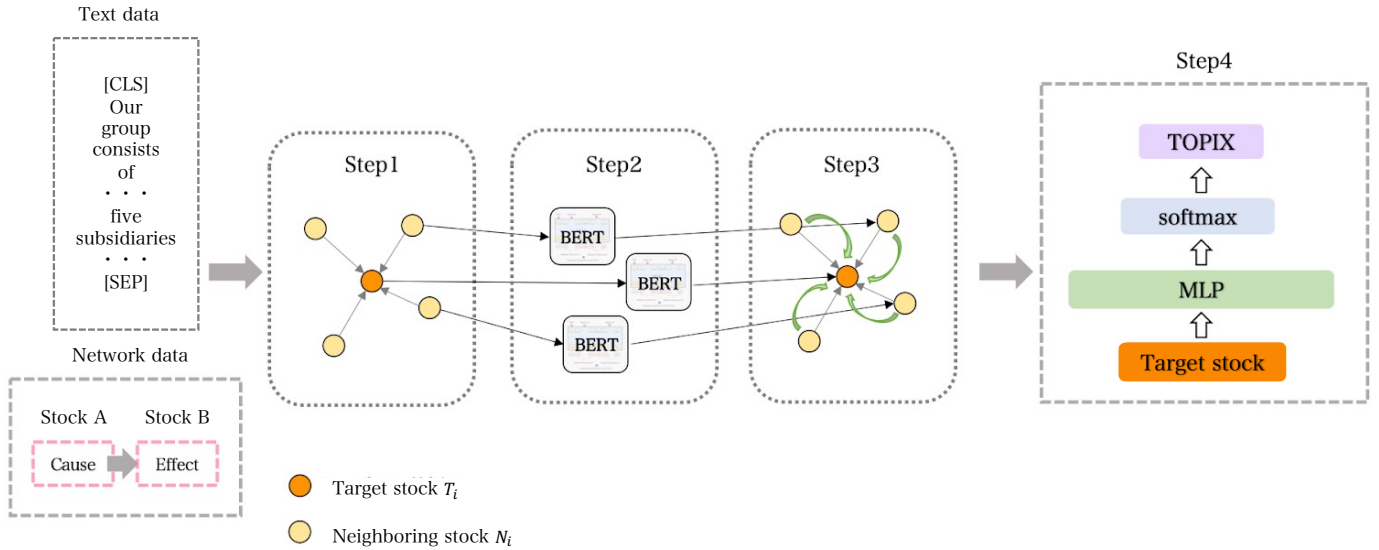


Fig. 1: Outline of our proposed model SETN

target node using network data. Step 2 obtains embeddings for each node from textual information using the domain-adaptive pretrained transformer-based model. Step 3 captures the network information using the GNN model. Step 4 feeds the embeddings of the target stock into the classifier for the sector and industry classifications listed in Table I.

A. Transformer-based Model for Text Embedding

We extracted textual information by embedding the business description of a company using a transformer-based model to obtain a sequence of token embeddings. We can either use BERT [4] or a robust optimized BERT pre-training approach (RoBERTa) [5]. Before applying pre-trained BERT or RoBERTa to SETN, we implemented domain-adaptive pre-training using financial corpora, as several studies have suggested adapting the BERT architecture to a specific domain, such as medicine and finance, has a significant effect on model performance [20]–[22]. Subsequently, we applied pooling operations to the token embeddings to obtain fixed-size embeddings. Three pooling strategies are possible: (1) taking the output of the special classification (CLS) token, (2) taking the mean of all the output vectors, and (3) applying max-pooling operations over the output vectors. SETN employs the second pooling strategy, taking the mean of all output vectors, because the strategy has been proven to be most effective for semantic textual similarity, natural language inference, and sector and industry classification [6], [25].

$$h_{\text{BERT},i} = \text{BERT}(x_i) \quad (1)$$

where x_i is the business description of target stock i and $h_{\text{BERT},i} \in \mathbb{R}$ is the pooled output of the BERT.

B. GNN Model for Graph Embedding

We extracted network information using the GNN. GNN has been gaining popularity because it enables the capturing

of structural relationships by updating and aggregating node representations in the financial domain [26]. The GNN models that we employed in the experiments were the graph convolutional network (GCN) [1] and graph attention network (GAT) [2]. The GCN is one of the most widely used GNN models that aggregates features of neighborhood nodes and normalizes the aggregated representations by node degree.

$$h_{\text{GCN},i} = \text{ReLU} \left(\sum_{j \in N_i} \hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}} h_j W + b \right) \quad (2)$$

where N_i is the neighboring stock of target stock i , $\hat{A} = A + I$ denotes the adjacency matrix of the graph with inserted self-loops, and $\hat{D}_{i,i} = \sum_{j=0} \hat{A}_{ij}$ is the diagonal degree matrix. h_i is the input for the GCN, W is the weight matrix, and b is the bias.

The GAT assigns neighboring nodes with different weight importance by incorporating an attention mechanism during aggregation.

$$a_{ij} = \frac{\exp(\text{LeakyReLU}(a [Wh_i \| Wh_j]))}{\sum_{k \in N_i} \exp(\text{LeakyReLU}(a [Wh_i \| Wh_k]))} \quad (3)$$

where a is an attention mechanism, and $\|$ represents a concatenation operation. We adopted the leaky ReLU (LReLU or LRel) function with a *negativeslope* = 0.2 for our activation function.

C. Residual Connection

Research shows that adding a residual connection (RC) to a GNN model improves its performance [27]. Thus, we employed a residual connection by adding the output of the

BERT to the output of the GNN. The residual connection is expressed as follows:

$$\begin{aligned} h_{\text{BERT},i} &= \text{BERT}(x_i) \\ h_{\text{GCN},i} &= \text{GCN}(h_{\text{BERT},i}) \\ h_{\text{res},i} &= h_{\text{BERT},i} + h_{\text{GCN},i} \end{aligned} \quad (4)$$

D. Loss Function

In this study, we considered a multitask learning setup for sector and industry classification tasks. For this task, we used the Tokyo Price Index-17 (TOPIX17) labels and the Tokyo Price Index-33 (TOPIX33) labels defined by the Tokyo Stock Exchange (TSE) in Table I. TOPIX33 is an industrial sector index traditionally used in the Japanese market while TOPIX33 reorganizes itself to TOPIX17 for convenience. This loss was due to the addition of TOPIX17 loss and TOPIX33 loss.

$$\text{Loss} = \text{Loss}^{\text{TOPIX17}} + \text{Loss}^{\text{TOPIX33}} \quad (5)$$

The TOPIX17 loss and TOPIX33 loss were computed with a classifier of a single-layer feedforward neural network with dropout and ReLU activation.

$$\begin{aligned} y_i^{\text{TOPIX}} &= \text{Softmax}(W^{\text{TOPIX}}h_i + b^{\text{TOPIX}}) \\ \text{Loss}^{\text{TOPIX}} &= \sum_{i \in \Omega} \text{CE}(d_i^{\text{TOPIX}}, y_i^{\text{TOPIX}}) \end{aligned} \quad (6)$$

where h_i represents the output of the GNN or a residual block, d_i represents the target label of target stock i , Ω represents all stocks, and CE represents the cross-entropy loss function. In Equation 6, we denote TOPIX to show either TOPIX17 or TOPIX33, because we calculate the loss of TOPIX17 and TOPIX33 in a similar manner.

IV. DATASET

We conducted our experiments on the Japanese stock market. For textual information, we used business descriptions from companies’ annual reports in Japanese language datasets on the Japanese stock market. For network information, we used causal chains [28], which represent cross-company economic relationships. Overall, our dataset contained 2,437 Japanese stocks. We split the dataset into training, validation, and test data containing respectively 1,705, 243, and 489 companies.

A. Text Corpora

We used CoARiJ¹, The Corpus of Annual Reports in Japan, which is organized by Japanese financial reports. The corpus covers 3,016 companies and includes annual reports, integrated reports, corporate social responsibility (CSR) reports, and earnings summaries, ranging from 2014 to 2018. We used all corpora for domain-adaptive pre-training. In conducting the experiments, to embed textual information, use a company’s business description, which contains an overview of the activities of the company.

¹<https://github.com/chakki-works/CoARiJ>

B. Network Data

In this study, we employed a causal chain as company network data. The causal chain concept was proposed by Izumi *et al.* [28] and is composed of causality in textual data. For example, the sentence “U.S. stocks give up earlier gains as oil recovers” includes causality that is composed of a cause “oil recovers” and an effect “U.S. stocks give up earlier gains.” They detect causality using clue expressions such as “as.” The meanings of the cause, effect, and clue expressions are defined in [28]. They then collect cause and effect pairs from financial reports using this technique.

Next, they connect the cause to the effect based on cosine similarities from word2vec. For example, the sentence “World oil prices hit a three- and half-month high on Tuesday after militants said they would broaden attacks on Nigeria’s oil industry” includes a cause “militants said they would broaden attacks on Nigeria’s oil industry,” and an effect “World oil prices hit a three and half-month high on Tuesday.” The effect can be connected to the previous cause “oil recovers”; then, we can obtain the causal chain “militants said they would broaden attacks on Nigeria’s oil industry” - ζ “World oil prices hit a three and half-month high on Tuesday” and “oil recovers” - ζ “U.S. stocks give up earlier gains as oil recovers.” Izumi *et al.* constructed causal chains using Japanese financial reports and accessed the causal chain search system on their website². Moreover, their causal chains are linked to each company because they use the financial reports of each company. Therefore, causal chains can be used as network data for companies. Furthermore, several studies have suggested the effectiveness of representing cross-company economic relationships using causal chains [29], [30]. In our study, a company is represented as a node, and a company’s relationship based on the causal chain is represented as an edge. The network has 2,437 nodes and 12,755 edges.

C. Thematic Fund

For the thematic fund creation tasks, we utilized themed stock listed on a website called “minkabu,”³ a Japanese financial information services website. We extracted 27 themes using the criterion that a theme should include more than 15 stocks in our test data. For the thematic fund creation task, we selected six themes: logistics, IT, biotechnology, semiconductor, 5G and Post-COVID, and assigned them to one of three separate groups. Logistics and IT belong to the first group. The characteristics of the first group are that the theme is related to the sector and industry, and it has a similar label in TOPIX17 and TOPIX33. Biotechnology and semiconductors are in the second group. The characteristic of the second group is that the theme is related to sector and industry, but it does not have similar labels in TOPIX17 or TOPIX33. The third group included the 5G and post-COVID. For the third group, the theme is not about either sector or industry, but other concepts such as social issues and new technology.

²<https://socsim.t.u-tokyo.ac.jp/ccs/>

³<https://minkabu.jp/screening/theme>

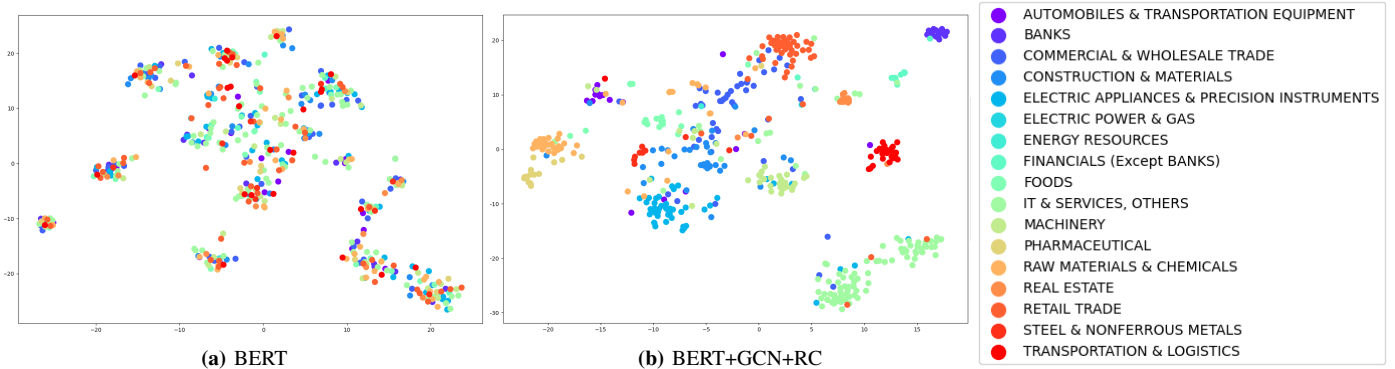


Fig. 2: Visualization of the embeddings with TOPIX17 sector labels. Subfigure (a) shows the embeddings from the BERT model. Subfigure (b) shows the embeddings from the BERT+GCN+RC model

TABLE II: Results for Related Company Information Extraction

	TOPIX17			TOPIX33		
	MAP@5	MAP@10	MAP@50	MAP@5	MAP@10	MAP@50
BERT	0.669	0.628	0.513	0.601	0.572	0.508
RoBERTa	0.669	0.631	0.532	0.581	0.543	0.488
BERT+GCN	0.635	0.594	0.484	0.551	0.516	0.455
BERT+GAT	0.635	0.588	0.467	0.555	0.517	0.436
RoBERTa+GCN	0.632	0.598	0.5	0.548	0.516	0.451
RoBERTa+GAT	0.616	0.576	0.474	0.527	0.496	0.44
BERT+GCN+RC	0.682	0.651	0.547	0.606	0.587	0.528
BERT+GAT+RC	0.679	0.657	0.561	0.598	0.585	0.524
RoBERTa+GCN+RC	0.68	0.651	0.553	0.599	0.579	0.528
RoBERTa+GAT+RC	0.666	0.635	0.528	0.581	0.555	0.495

V. EXPERIMENT

A. Training Details

We leveraged the pre-trained Japanese BERT and RoBERTa models for our proposed model. We used the following pre-trained models from Hugging Face: bert-base-Japanese⁴, and roberta-base-Japanese⁵. The models were trained on Japanese text with 12 layers, 768 hidden units, and 12 attention heads. The text was tokenized by MeCab⁶ and WordPiece [31] for bert-base-japanese, The text was first segmented into words by Juman++ [32] in advance, and each word was tokenized by SentencePiece [33] for roberta-base-Japanese.

For domain-adaptive pre-training, we used all text data from the CoARiJ corpus (approximately 200,000 text files and 5.0 GB). The pre-training was based on the masked language model and next-sentence prediction for the BERT model and masked language model for RoBERTa. We trained all 12 layers in 200,000 steps in both models.

For our proposed model, we utilized the first 512 tokens of business descriptions in annual reports as textual information. In training SETN, we trained only the last layer in BERT or RoBERTa. For extracting a subgraph, we performed neighborhood sampling by one hop in this study. We backpropagated only through the loss of the target node in the subgraph to

avoid leakage. We trained our models on the loss function using the Adam optimizer [34] for 20 epochs, with a learning rate of 0.001 and dropout rate of 0.2. For all models, PyTorch and PyTorch-geometric were used for implementation. All experiments were conducted on an NVIDIA A6000 and Intel i9-10980XE.

B. Evaluation Tasks

First, we evaluated the performance of our proposed method on a related company information extraction task. We then conducted two types of ablation studies. Finally, we evaluated the performance of our proposed model on a thematic fund creation task.

Related company information extraction task

The purpose of the related company information extraction task is to evaluate the extent to which stock embeddings capture sector and industry information. Following the work in [35], we evaluated the embeddings using the Mean Average Precision at K (MAP@K) metric. We set K in MAP@K to $K = 5, 10, 50$.

To verify the usefulness of our proposed model further, we conducted two types of ablation studies. The first ablation study analyzed the effect of the graph type on performance. Most studies use undirected graphs for graph construction [11]–[13], [15], [18] although the relationships between stocks are often described in a directed form, such as supply chain

⁴<https://huggingface.co/cl-tohoku/bert-base-japanese>

⁵<https://huggingface.co/nlp-waseda/roberta-base-japanese>

⁶<http://mecab.sourceforge.jp>

and transaction relationships. Therefore, we compared the performance of the models with a directed graph and evaluated the model’s performance with an undirected graph. The second ablation study verified the improvement in model performance by jointly training a text embedding and graph embedding model, compared to several studies in which text embedding and graph embedding models were trained separately [23], [24]. Specifically, we compared our proposed model with a model that does not update the parameters of a transformer-based model. In other words, the text features are encoded via a transformer-based model and fed as an input to a graph embedding model.

Thematic fund creation tasks

In this task, we evaluated the model’s capability to find relevant companies with respect to a given theme. We experimented with thematic fund creation tasks based on six different themes. To this end, we obtained $|\mathbf{t}|$ most similar companies for each stock i in a given theme in terms of cosine similarities. Then, we evaluated how many stocks out of $|\mathbf{t}|$ stocks shared the same theme as stock i using the following evaluation metric.

$$\text{theme metric} = \sum_{\mathbf{t} \in \mathcal{T}} \sum_{i,j=1}^{|\mathbf{t}|} \frac{\mathcal{H}_{i,j,\mathbf{t}}}{|\mathcal{T}||\mathbf{t}|^2} \quad (7)$$

where \mathcal{T} is a set of all themes and \mathbf{t} represents a set of stocks in a given theme, such as biotechnology and post-COVID. $\mathcal{H}_{i,j,\mathbf{t}}$ are calculated as follows:

$$\mathcal{H}_{i,j,\mathbf{t}} = \begin{cases} 1, & \text{if } \text{pred}_{i,j} \text{ is in the same theme as stock } i \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

where $\text{pred}_{i,j}$ is the stock closest to stock i with respect to cosine similarity.

C. Baseline Models

We compared the performance of the proposed model with the following baselines:

BERT_{Japanese}: Pre-trained BERT with single-layer feedforward neural network. The pre-trained BERT was pre-trained on Japanese Wikipedia and further pre-trained on the CoARiJ dataset.

RoBERTa_{Japanese}: RoBERTa was pre-trained with a single-layer feedforward neural network. The pre-trained RoBERTa was pre-trained on Japanese Wikipedia and the Japanese portion of CC-100 and further pre-trained on the CoARiJ dataset.

D. Result

The results for the related company information extraction task are reported in Table II. Tables III and IV show the nearest neighbors for the selected stock using the BERT and BERT+GCN+RC models from the test data.

Figure 2 shows a visualization of the embeddings with TOPIX17 labels. T-distributed stochastic neighbor embedding

(t-sne) was applied to reduce the dimensions. The reduced embeddings were plotted and labeled with color using TOPIX17 labels. Figure 2a shows the embeddings from the BERT model, whereas Figure 2b shows the embeddings from the BERT+GCN+RC model. Table V shows the results of the two types of ablation studies on the graph type and learning architecture. Table VI shows the results of thematic fund creation for the six selected themes.

VI. DISCUSSION

A. Related Company Information Extraction

We compared the performance of the SETN with that of the baseline models in the task of related company information extraction. Table II shows that the proposed modeling approach outperformed the baseline models. This indicates that jointly learning both textual and network information improves the performance of a model with regard to related company information extraction tasks. In addition, adding a residual connection layer to the model improves its performance significantly. The residual connection layer enhances the model’s performance because it helps to equalize the weights of textual and network information by adding the output of BERT or RoBERTa, either of which encodes textual information, to the output of GCN or GAT, either of which encodes network information.

Figure 2 shows a visualization of the embeddings from the BERT and BERT+GCN+RC models with TOPIX17 labels. We can observe that the stocks in the same TOPIX17 sector tend to be grouped together more in the BERT+GCN+RC model than in the BERT model, which is consistent with the MAP@K evaluation in related company information extraction tasks.

Tables III and IV show the three nearest neighbors for the selected stocks in the test data using the BERT and BERT+GCN+RC models. The tables show that the three neighbors in the BERT+GCN+RC model generally have the same TOPIX17 and TOPIX33 labels as those in the BERT model. The case of Kintetsu Group Holdings Co., Ltd provides an example where the TOPIX17 labels of the three nearest neighbors are correct in the BERT+GCN+RC model but incorrect in the BERT model. In the BERT model, the Kintetsu Department Store Co., Ltd. was chosen as one of the closest neighbors. Inspecting the textual data reveals that the business descriptions of both stocks are quite similar, which might be because Kintetsu Department Store Co., Ltd. is a subsidiary of Kintetsu Group Holdings Co., Ltd. We consider the BERT model closely represents the two stocks because it is difficult to tell the difference between each stock’s sector and industry in this example, as the model only learns from textual information. Conversely, the BERT+GCN+RC model learns from textual information and cross-company economic relationship information. Therefore, the embeddings become more informative for the stock sector and industry when the model is trained with both types of information. This example suggests that the proposed model can utilize both textual and network information.

TABLE III: Three nearest neighbors for selected stocks in the test set for the BERT model

Company	TOPIX17	TOPIX33
Kintetsu Group Holdings Co., Ltd.	Transportation & Logistics	Land Transportation
Tobu Railway Co., Ltd. Takashimaya Co., Ltd. Kintetsu Department Store Co., Ltd.	Transportation & Logistics Retail Trade Retail Trade	Land Transportation Retail Trade Retail Trade
Nippon Express Co., Ltd.	Transportation & Logistics	Land Transportation
Kintetsu World Express, Inc Niigata Kotsu Co., Ltd. Kyushu Railway Company	Transportation & Logistics Transportation & Logistics Transportation & Logistics	Warehousing and Harbor Transportation Land Transportation Land Transportation
Fujitsu General Limited	Electric appliances & Precision instruments	Electric Appliances
Komori Machinery (Nantong) Co., Ltd Towa Corporation Shibaura Mechatron Corporation	Machinery Machinery Electric appliances & Precision instruments	Machinery Machinery Electric Appliances
SBI Holdings, Inc	Financial (except banks)	Securities and Commodities Futures
Yodobashi Holdings INC Hotman Co., Ltd. Square Enix Holdings Co., Ltd.	Retail Trade Retail Trade IT & Services, Others	Retail Trade Retail Trade Other Products

TABLE IV: Three nearest neighbors for selected stocks in the test set for the BERT+GCN+RC model

Company	TOPIX17	TOPIX33
Kintetsu Group Holdings Co., Ltd.	Transportation & Logistics	Land Transportation
Tobu Railway Co., Ltd. Kyushu Railway Company Kobe Electric Railway Co., Ltd.	Transportation & Logistics Transportation & Logistics Transportation & Logistics	Land Transportation Land Transportation Land Transportation
Nippon Express Co., Ltd	Transportation & Logistics	Land Transportation
Inui Global Logistics Co., Ltd. Kintetsu World Express, Inc Nippon Yusen Kabushiki Kaisha	Transportation & Logistics Transportation & Logistics Transportation & Logistics	Marine Transportation Warehousing and Harbor Transportation Marine Transportation
Fujitsu General Limited	Electric appliances & Precision instruments	Electric Appliances
Daihen Corporation Hokuetsu Industries Co., Ltd. Nissan Electric Co., Ltd.	Electric appliances & Precision instruments Machinery Electric appliances & Precision instruments	Electric Appliances Machinery Electric Appliances
SBI Holdings, Inc.	Financial (except banks)	Securities and Commodities Futures
Mizuho Leasing Company, Limited Money Partners Group Co., Ltd. Toyo Securities Co., Ltd.	Financial (except banks) Financial (except banks) Financial (except banks)	Other Financing Business Securities and Commodities Futures Securities and Commodities Futures

TABLE V: Ablation Study

BERT Train	Graph Type		TOPIX17			TOPIX33		
			MAP@5	MAP@10	MAP@50	MAP@5	MAP@10	MAP@50
w/ train	Directed	BERT+GCN+RC	0.682	0.651	0.547	0.606	0.587	0.528
		BERT+GAT+RC	0.679	0.657	0.561	0.598	0.585	0.524
		RoBERTa+GCN+RC	0.68	0.651	0.553	0.599	0.579	0.528
		RoBERTa+GAT+RC	0.666	0.635	0.528	0.581	0.555	0.495
w/ train	Undirected	BERT+GCN+RC	0.647	0.616	0.512	0.568	0.536	0.485
		BERT+GAT+RC	0.612	0.569	0.468	0.532	0.488	0.44
		RoBERTa+GCN+RC	0.649	0.615	0.524	0.552	0.522	0.471
		RoBERTa+GAT+RC	0.663	0.627	0.523	0.553	0.514	0.468
w/o train	Directed	BERT+GCN+RC	0.637	0.611	0.51	0.567	0.547	0.481
		BERT+GAT+RC	0.66	0.623	0.501	0.572	0.545	0.482
		RoBERTa+GCN+RC	0.673	0.633	0.539	0.597	0.56	0.507
		RoBERTa+GAT+RC	0.661	0.624	0.509	0.571	0.541	0.484

TABLE VI: Results for the Thematic Fund Creation

	LOGISTICS	IT	BIOTECHNOLOGY	SEMICONDUCTOR	5G	POST-COVID
Random Guess	0.031	0.033	0.033	0.033	0.037	0.035
BERT	0.308	0.227	0.251	0.067	0.062	0.118
RoBERTa	0.375	0.188	0.267	0.063	0.077	0.080
RoBERTa+GCN+RC	0.313	0.294	0.220	0.125	0.065	0.080
RoBERTa+GAT+RC	0.424	0.247	0.239	0.125	0.096	0.094
BERT+GCN+RC	0.388	0.208	0.204	0.114	0.105	0.125
BERT+GAT+RC	0.388	0.314	0.227	0.145	0.096	0.146

Table V presents the results of the two ablation studies. The first ablation study investigated whether a directed graph model performed better than an undirected graph model. For the results, we compared the first and second sections in Table V. The results show that the model with a directed graph outperforms the model with an undirected graph for related company information extraction tasks. We observe that a directed graph model better captures cross-stock relationships because the economic relationship between stocks is often in a directed form, such as supply chain and transaction relationships.

The second ablation study investigated how model performance improves by jointly training both the text embedding model and the graph embedding model. The results can be observed by comparing the first and third sections of Table V. The result suggests that jointly training both text-embedding and graph embedding models outperforms separately encoding textual information with the transformer-based model and using it as an input to the GNN model. We observe that the embeddings learn both textual and network information more effectively by training both the transformer-based and GNN models in an end-to-end framework.

B. Thematic Fund Creation

The results of the thematic fund creation experiment reveal that our proposed models outperform baseline models in most settings. First, comparing the performances of the different models shows that our modeling approach outperformed the baseline and random guess models. This indicates that jointly learning both textual and network information improves the performance of models with regard to thematic fund creation tasks. Second, a comparison of the performances of different theme groups revealed that the performances deteriorated from the first to the last group as the task gradually became different from the related company information extraction task. The first group, logistics and IT, had the highest performance among all groups. An explanation for this is that the group’s task is similar to the related company information extraction task, and it has already been shown that the embeddings capture TOPIX17 and TOPIX33 information well. The tasks of the second group, biotechnology and semiconductor, can be seen as zero-shot learning of related company information extraction because the label is an unseen sector and industry label. The performance is not the highest but is close to that of the first group. We recall that during training, the embeddings

learned the latent variables of the sector and industry associated with the stocks. This leads to competitive results, even for an unseen theme. The third group, 5G and Post-COVID, had the lowest performance, as the theme was completely different from the related company information extraction task, and it was not even about the company stocks associated industry or sector. However, our model outperformed the baseline and random-guess models. Our analysis suggests that the embeddings obtain information about social situations, new technology, and other general concepts by learning from business descriptions and economic relationships captured by causal chains, although the scope of the information that the embeddings learned is not yet fully understood. Therefore, the properties of the embeddings must be analyzed thoroughly in future work.

VII. CONCLUSION

We presented the SETN, a framework that jointly learns from textual and network information. SETN learns textual information with transformer-based models, such as BERT and RoBERTa, and learns network information using GNN models, such as GCN and GAT. We tested the validity of our proposed model on the related company information extraction task using MAP@K, which evaluated the extent to which the candidate models captured the information of the associated industrial sector. Our proposed models, especially the BERT+GAT+RC model, outperformed the baseline models in related company information extraction tasks. Furthermore, we conducted ablation studies to confirm the validity of using a directed graph to learn network information and the advantages of jointly training a transformer-based text embedding model and a GNN graph embedding model over training them separately. Finally, we applied the embeddings obtained from our proposed model to a thematic fund creation task. Our proposed model outperformed the baseline models, and the results showed that the embeddings captured the latent variables of stocks in the associated sector and industry.

We plan to extend this work to incorporate rich information about stocks, such as shareholding and supply chain relationships, in three ways. First, we shall widen the scope of the data to capture different types of network information. Second, we shall expand our model to heterogeneous graphs that can handle different types of nodes and edges. Third, we shall identify new evaluation metrics to quantify the extent to which the model captures diverse network information.

ACKNOWLEDGMENT

This work was supported by the Japan Science and Technology Future Program (JST-Mirai), Grant Number JP-MJMI20B1, Japan. This work was also supported by Daiwa Securities Group Inc.

REFERENCES

- [1] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [2] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," in *Proceedings of the 6th International Conference on Learning Representations*, 2017.
- [3] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "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*, 2019, pp. 4171–4186.
- [5] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A robustly optimized bert pretraining approach," in *Proceedings of the 8th International Conference on Learning Representations*, 2019.
- [6] T. Ito, J. C. Collados, H. Sakaji, and S. Schockaert, "Learning company embeddings from annual reports for fine-grained industry characterization," in *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing*, 2020, pp. 27–33.
- [7] X. Du and K. Tanaka-Ishii, "Stock embeddings acquired from news articles and price history, and an application to portfolio optimization," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 3353–3363.
- [8] J. Pennington, R. Socher, and C. Manning, "GloVe: Global vectors for word representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 2014, pp. 1532–1543.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [10] D. Cer, Y. Yang, S. Kong, N. Hua, N. Limtiaco, R. St. John, N. Constant, M. Guajardo-Cespedes, S. Yuan, C. Tar, B. Strope, and R. Kurzweil, "Universal sentence encoder for English," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 169–174.
- [11] Q. Wu, C. G. Brinton, Z. Zhang, A. Pizzoferrato, Z. Liu, and M. Cucuringu, "Equity2vec: End-to-end deep learning framework for cross-sectional asset pricing," in *Proceedings of the Second ACM International Conference on AI in Finance*, 2021, pp. 1–9.
- [12] V. Satone, D. Desai, and D. Mehta, "Fund2vec: Mutual funds similarity using graph learning," in *Proceedings of the Second ACM International Conference on AI in Finance*, 2021, pp. 1–8.
- [13] S. Li, X. Wang, and X. Jiang, "Mining for the preference of funds based on subgraph embedding of fund-stock networks," in *2020 International Conference on Communications, Computing, Cybersecurity, and Informatics*, 2020, pp. 1–5.
- [14] Z. Li, J. Bian, D. Yang, T. Qin, L. Zhao, and T. Y. Liu, "Individualized indicator for all: Stock-wise technical indicator optimization with stock embedding," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 894–902.
- [15] Y. Chen, Z. Wei, and X. Huang, "Incorporating corporation relationship via graph convolutional neural networks for stock price prediction," in *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 2018, pp. 1655–1658.
- [16] J. Ye, J. Zhao, K. Ye, and C. Xu, "Multi-graph convolutional network for relationship-driven stock movement prediction," in *2020 25th International Conference on Pattern Recognition*, 2021.
- [17] D. Matsunaga, T. Suzumura, and T. Takahashi, "Exploring graph neural networks for stock market predictions with rolling window analysis," in *Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy*, 2019.
- [18] J. Gao, X. Ying, C. Xu, J. Wang, S. Zhang, and Z. Li, "Graph-based stock recommendation by time-aware relational attention network," *ACM Transactions on Knowledge Discovery from Data*, vol. 16, no. 1, pp. 1–21, jul 2021.
- [19] Y. Chen, J. Wu, and H. Bu, "Stock market embedding and prediction: A deep learning method," in *2018 15th International Conference on Service Systems and Service Management*, 2018, pp. 1–6.
- [20] Z. Liu, D. Huang, K. Huang, Z. Li, and J. Zhao, "Finbert: A pre-trained financial language representation model for financial text mining," in *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, 2020, pp. 4513–4519.
- [21] B. Peng, E. Chersoni, Y.-Y. Hsu, and C.-R. Huang, "Is domain adaptation worth your investment? comparing BERT and FinBERT on financial tasks," in *Proceedings of the Third Workshop on Economics and Natural Language Processing*, 2021, pp. 37–44.
- [22] D. Araci, "Finbert: Financial sentiment analysis with pre-trained language models," 8 2019. [Online]. Available: <http://arxiv.org/abs/1908.10063>
- [23] R. Sawhney, S. Agarwal, A. Wadhwa, and R. 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*, 2020, pp. 8415–8426.
- [24] R. Sawhney, P. Khanna, A. Aggarwal, T. Jain, P. Mathur, and R. R. Shah, "VolTAGE: Volatility forecasting via text audio fusion with graph convolution networks for earnings calls," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 8001–8013.
- [25] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," in *Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds., 2019, pp. 3982–3992.
- [26] J. Wang, S. Zhang, Y. Xiao, and R. Song, "A review on graph neural network methods in financial applications," 11 2021. [Online]. Available: <http://arxiv.org/abs/2111.15367>
- [27] X. Liu, J. Ding, W. Jin, H. Xu, Y. Ma, Z. Liu, and J. Tang, "Graph neural networks with adaptive residual," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021, pp. 9720–9733.
- [28] I. Kiyoshi and S. Hiroki, "Economic causal-chain search using text mining technology," in *the First Workshop on Financial Technology and Natural Language Processing In conjunction with the 28th International Joint Conference on Artificial Intelligence*, 2019, pp. 61–65.
- [29] N. Kei, S. Shingo, S. Hiroki, and I. Kiyoshi, "Economic causal chain and predictable stock returns," in *7th International Conference on Smart Computing and Artificial Intelligence*, 2019, pp. 655–660.
- [30] N. Kei, S. Shingo, and S. Hiroki, "Investment strategy via lead lag effect using economic causal chain and ssestm model," in *10th International Congress on Advanced Applied Informatics*, 2022.
- [31] M. Schuster and K. Nakajima, "Japanese and korean voice search," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2012, pp. 5149–5152.
- [32] A. Tolmachev, D. Kawahara, and S. Kurohashi, "Juman++: A morphological analysis toolkit for scriptio continua," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 54–59.
- [33] T. Kudo and J. Richardson, "SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 66–71.
- [34] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations*, Y. Bengio and Y. LeCun, Eds., 2015.
- [35] K. Yu, B. Zhang, H. Zhu, H. Cao, and J. Tian, "Towards personalized context-aware recommendation by mining context logs through topic models," in *Pacific-Asia conference on knowledge discovery and data mining*, 2012, pp. 431–443.