FinSenticNet: A Concept-Level Lexicon for Financial Sentiment Analysis

Kelvin Du[♠], Frank Xing[♣], Rui Mao[♠] and Erik Cambria[♠]

[♠]*School of Computer Science and Engineering, Nanyang Technological University, Singapore*
[♣]*Department of Information Systems and Analytics, National University of Singapore, Singapore*
zidong001@e.ntu.edu.sg; xing@nus.edu.sg; rui.mao@ntu.edu.sg; cambria@ntu.edu.sg

Abstract—Sentiment lexicons are important tools for research involving opinion mining and sentiment analysis. They are highly inter-operable, and address critical limitations of learning-based or large language model-based sentiment analysis, providing better reproducibility and explainability. Existing financial sentiment lexicons, manually crafted or automatically constructed, primarily comprise single-word entries despite the fact that jargon, terminologies, and collocations in finance are often multi-word expressions. To address this gap, we present FinSenticNet, a concept-level domain-specific lexicon specifically designed for financial sentiment analysis, where over 65% entries are multi-word expressions. Our construction approach is semi-supervised: the framework consists of a concept parser, a sentiment seeds generation module, and a semantic graph construction module. Each concept (graph node) is subsequently classified in terms of its polarity using the Label Propagation Algorithm and Graph Convolutional Network. Compared to other financial sentiment lexicons, FinSenticNet *captures domain-specific language features and has a broader coverage*. We demonstrate this with superior evaluation results, i.e., sentiment analysis accuracy and F-scores, on multiple well-received benchmark datasets.

Index Terms—financial sentiment analysis, graph convolutional network, sentiment lexicon, label propagation, opinion mining

I. INTRODUCTION

The importance of Financial Sentiment Analysis (FSA) has increased over the past decade [1], with FSA proving to be a powerful tool for understanding investor sentiment and forecasting financial markets [2], [3]. Sentiment analysis is domain-specific and a greater degree of domain-dependence is evident in the finance domain because of the concentration of topics, the use of highly professional language [4], and the distinctive cognitive patterns between different market environments [5]. In the last decade, Web 2.0 has dramatically increased the number and variety of information resources, so we face the challenge of transforming this vast amount of information into computationally tractable information.

Finance has always been a domain where information is needed regarding the opinions of financial institutions, such as central banks, and market evaluations, such as the Michigan Consumer Sentiment Index [6]. During the past few years, text mining techniques have been increasingly used for Financial Sentiment Analysis, which is regarded as a more dynamic and robust way to analyze online texts than traditional surveying methods, and is being increasingly used as an alternative to traditional surveying techniques.

FSA research has made remarkable progress with the advancement of deep learning in recent years. Specifically, FinBERT [7]–[9], a finance domain-specific BERT, is trained by utilizing Reuters Corpora, Yahoo Finance, Raddit Finance, corporate reports, earnings call transcripts and analyst reports in different studies and pushed the boundaries of FSA research. Research into the creation of financial lexicons continues to attract researchers' attention, even though lexicon methods are often used alongside learning-based methods in financial services [10]–[12]. High-quality sentiment lexical resources are essential to achieve good performance in sentiment analysis which is particularly true in the finance domain [6]. Lexicons also help increase the explainability of financial models [13], [14].

A lexicon-based method detects the semantic orientation of the text based on its words and phrases in the text. As most lexicon-based methods are unsupervised [15], [16], their greatest advantage is that they can perform FSA without any annotated dataset, which reduces the need for manual annotations. Meanwhile, lexicons are useful for creating features in supervised learning tasks and integrating lexical knowledge into pre-trained language models to improve model performance [10]. Lexicon-based methods are reproducible, trustworthy, interpretable, and explainable [17]. However, the first challenge with lexicon-based approaches is that it is time-consuming to build lexicons [18]. Secondly, sentiment analysis is sensitive to domain-specific knowledge. Generic domain-independent lexicons are hard to generalize and often ineffective in FSA [19], because financial expressions are different from other text types.

A sentiment lexicon created for general purposes may encounter challenges when dealing with financial text due to the potential misclassification of common financial terms. This issue poses a difficulty in achieving consistent sentiment analysis across different domains, highlighting the necessity for specialized sentiment analysis within the finance industry [20]. To illustrate, words such as "liability" and "debt" are often classified as negative in sentiment analysis, but within financial contexts, they are frequently used and carry a neutral connotation. Further, the construction of lexicons in the financial domain is scarce as compared to general-purpose lexicons [19]. Lexicons such as SenticNet [17], SentiWordNet [21] and Valence Aware Dictionary and Sentiment Reasoner (VADER) [22] are general-purpose lexicons which do not generally perform as expected in finance domain.

More importantly, the existing financial lexicons, which include Henry's Financial Dictionary (HFD) [23], Loughran and McDonald (LM) [20] and Stock Market Sentiment Lexicon (SMSL) [19], face generalization problem which fails to perform well on multiple benchmark datasets. Further, the current financial lexicons primarily consist of single-word entries, yet employing a concept-level sentiment analysis approach can augment word-level opinion mining through the incorporation of multi-word expressions and linguistic constructs composed of more than one word. Lastly, earlier approach to automatic lexicon construction encompasses two main methods: supervised learning [19], which necessitates an annotated dataset, and semi-supervised learning through Label Propagation Algorithm (LPA) [24], which fails to adequately consider the significance of node features. To mitigate the aforementioned concerns, we have introduced FinSenticNet, a domain-specific lexicon at the concept level, designed specifically for the FSA task. In particular, our contributions can be summarized from three perspectives:

1) We proposed an automatic approach to construct financial lexicons by building a semantic similarity graph and formulated the lexicon construction into a semi-supervised node classification problem. We implemented a two-stage framework and demonstrated its effectiveness in constructing financial lexicons.

2) We make publicly available FinSenticNet[1], an extensive and all-encompassing concept-level lexicon crafted for the purpose of financial sentiment analysis. FinSenticNet is ever-expanding as new concept entries are discovered from sentences and included based on the semantic graph, thus "a growing organism", using Ranganathan's metaphor.

3) We conducted extensive experiments and achieved competitive performance in comparison to strong benchmark models across four financial sentiment analysis datasets.

II. RELATED WORKS

The development of sentiment lexicons holds substantial significance in the field of Natural Language Processing (NLP) as it constitutes the fundamental basis for conducting research on opinion mining and sentiment analysis [25]. Lexicon construction can be conducted manually, semi-automatically, or automatically [19]. The manual approach requires intensive efforts from creators with expert knowledge, which is slow, but the accuracy is generally higher. On the other hand, the automatic approach is fast and scalable, but often results in sacrifice in the accuracy to some extent. Huang et al. [24] has categorised that the sentiment lexicon can be constructed based on a semantic thesaurus or corpus statistics. In the first category [21], [26]–[28], sentiment polarities are determined using semantic relations and glosses in the existing thesaurus, such as WordNet, HowNet and CiLin etc [29].

The corpus-based method has gained considerable attention and been extensively explored in the field [24], [30], [31]. This approach is founded on the underlying assumption that polar terms, which express similar polarities, tend to co-occur within a domain-specific corpus. To assign polarity, contextual evidence is commonly employed, leveraging the surrounding context of the terms in question [24]. The availability of lexicons specifically designed for the financial domain is relatively limited compared to more general-purpose lexicons [19].

Within the context of FSA, the popular finance domain-specific lexicons include HFD, LM, SMSL and SentiEcon. HFD, the first dictionary tailored to the financial domain, comprises 85 negative words and 104 positive words, is manually constructed from earnings press releases. Its primary application is to gauge the tone of such releases, which serves as a critical aspect of the firm-investor communication process [23]. HFD has found widespread use in financial sentiment analysis. Nevertheless, its limited word coverage poses a weakness. LM lexicon represents a significant contribution to the development of financial lexicons [20]. The authors manually examined the quality of the General Inquirer (GI) lexicon and proposed a revised lexicon specifically designed for financial texts. The LM sentiment word list, derived from annual reports released by firms, stands as the most commonly utilized lexicon in the financial domain.

However, both HFD and LM lack contextual information at the word level [19]. Oliveira et al. [19] proposed a novel and efficient approach for constructing the Stock Market Sentiment Lexicon (SMSL), which was built using labeled data from StockTwits, a microblog specialized in the stock market. This lexicon, consisting of 20,550 words and phrases, demonstrates competitive performance in measuring investor sentiments. Lastly, designed specifically for sentiment analysis applications, SentiEcon represents a comprehensive and domain-specific computational lexicon for the field of Economy and Finance. It comprises 6,470 entries, encompassing both single and multi-word expressions, each annotated with tags indicating their semantic orientation and intensity [6]. It is worth mentioning that SentiEcon is intended to be used in conjunction with a general domain lexicon, as it includes entries where the domain-specific polarity differs from or is not recorded in the general domain lexicon.

A notable trend in the development of financial lexicons is the shift from focusing only on single words to including also multi-word expressions, which also encapsulate direction-dependent expressions. This shift is particularly significant in the finance domain, as the sentiment associated with a financial term can vary depending on directional words. For instance, the concept of 'making profit' is considered positive, while 'making loss' is deemed negative. Consequently, it is crucial to construct context-aware and direction-dependent lexicons [32]. An illustrative example is the presence of direction-dependent words such as 'profit_up' and 'loss_down', which convey positive sentiment, whereas 'profit_down' and 'loss_up' indicate negative sentiment.

---

[1] https://github.com/senticnet/finsenticnet

Park et al. [32] improved the performance on the Phrase-Bank dataset by developing a context-aware sentiment lexicon for the finance domain. This approach involved incorporating direction-dependent words and combining them with the LM lexicon [20] to conduct sentiment analysis based on lexicons. Despite the predominance of manual creation as the primary method for constructing financial lexicons, which demands significant efforts from knowledgeable creators and generally yields higher accuracy, there is ongoing research [19] that advocates for the transition from manual to automatic approaches. This shift allows us to overcome the limitations of the slow and limited coverage associated with manual construction, thereby enabling the development of lexicons with enhanced speed and coverage. Currently, manual creation remains the predominant method for constructing financial lexicons, demanding substantial efforts from knowledgeable creators and generally yielding higher accuracy. However, the recent research by Oliveira et al. [19] is driving the shift from manual to automatic approaches in lexicon construction. Toward the same direction, FinSenticNet aims to address the inherent drawbacks of the manual approach, such as the slow process and limited coverage, by building lexicons with improved speed and broader coverage automatically. Meanwhile, the framework emphasizes the creation of concept-level lexicons to enhance the accuracy, reproducibility, trustworthiness, interpretability, and explainability. Moreover, FinSenticNet can serve as a knowledge source for integration into learning-based approaches to enhance financial sentiment analysis.

III. PROPOSED APPROACH

Our proposed approach is illustrated in Fig. 1. Specifically, it consists of Concept Parser, Semantic Graph, Seed Words and Concepts, Label Propagation, and Graph Convolutional Network (GCN).

A. Concept Parser

The concept-level sentiment analysis approach enhances word-level opinion mining by utilizing multi-word expressions and linguistic objects formed by more than one word, that exhibit formal or functional idiosyncratic properties as they relate to free word combinations [33]. The concept parser consists of tokenization, multi-word token expansion, Part-of-Speech (POS) tagging, lemmatization and dependency parsing. The dependency parsing constructs a tree-like structure of words based on the input sentence, reflecting the syntactic dependency relationships among the words. These resulting tree representations, adhering to the Universal Dependencies formalism, hold significance in various downstream applications [34]. In the preprocessing phase, each sentence undergoes several steps including tokenization, multi-word token expansion, POS tagging, lemmatization, and parsing into a structure of universal dependencies, where the head index of each word can be accessed through the property 'head' providing information about the word's syntactic dependency relation with respect to other words, represented by the property 'deprel'. For example, nsubj is
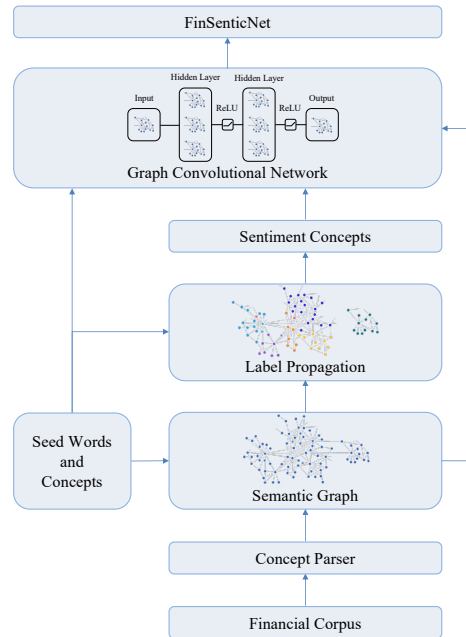

Fig. 1. Proposed Framework for FinSenticNet.

selected and nsubj ('nsubj', 'drop', 'profit') is obtained after parsing the sentence "L&T's net profit for the whole 2010 drop to eur 36 million from eur 45 million for 2009", and profit_drop is parsed as a concept.

B. Semantic Graph

The semantic graph $G(V, E, W)$ represents the sentiment words and concepts as nodes V and the similarity between them as edge weights W. By constructing the graph, we predict the sentiment polarity of unlabeled candidate sentiment words and concepts using node classification.

The construction of the graph involves the integration of WordNet [35] and word embeddings. Word embeddings, a form of word representation, enable words with similar meanings to possess comparable representations. In this study, FinText [36], a recently developed and highly advanced financial word embedding derived from the Dow Jones Newswires Text News Feed Database, is employed. WordNet, on the other hand, is an extensive lexical database of the English language that organizes nouns, verbs, adjectives, and adverbs into collections of cognitive synonyms known as synsets, each representing a distinct concept [35].

A limitation of non-contextual word embeddings, such as Word2Vec, is their reliance solely on co-occurrence statistics, which can encompass both synonyms and antonyms. For instance, when examining the word 'increase' within FinText, the words with high cosine similarity scores include 'increased', 'increases', 'decrease', 'increasing', 'reduction', 'decline', 'rise', 'improvement', 'reduce', and 'reduced'. It is important to note that among these words, 'decrease', 'reduction', 'decline', 'reduce' and 'reduced' are antonyms of 'increase'. To address this issue, we employ WordNet synsets, which group together words that are synonymous and express the same underlying concept, enabling us to

refine the results obtained from FinText. The selection and labeling of seed nodes are carried out by considering finance-related entries from SenticNet 7, along with manually crafted direction-dependent expressions in finance, and terms with high Pointwise Mutual Information (PMI) scores, extracted from the training dataset of SemEval 2017 Task 5. PMI is an association metric that measures the relationship between two outcomes by comparing the probability of observing two outcomes together with the probability of observing two outcomes independently [37]. There are several studies that have discussed the effectiveness of PMI in capturing associations between words and sentiments [38]. The PMI between the word w and the sentiment polarity s is defined as follows:

$$
\begin{aligned}
PMI(w, s) &= \log_2(P(w|s)/P(w)) \\
&= \log_2(P(w, s)/P(w)P(s))
\end{aligned}
\tag{1}
$$

Specifically, we compute $PMI(w = 0; s = 0)$, $PMI(w = 0; s = 1)$, $PMI(w = 1; s = 0)$ and $PMI(w = 1; s = 1)$. In this context, $w = 1$ denotes the presence of a word or concept, while $w = 0$ indicates its absence. Similarly, $s = 1$ signifies a positive sentiment polarity, whereas $s = 0$ represents a negative sentiment polarity.

C. Label Propagation Algorithm

Label Propagation Algorithm is a semi-supervised learning algorithm introduced by [39], [40]. Semi-supervised learning is based on the prior assumption of consistency which means neighboring nodes tend to have the same label. LPA represents a transductive learning method that propagates known labels to unlabeled nodes. The fundamental concept involves finding a collective labeling for all nodes given a graph and a limited number of labelled nodes.

Given that A is the adjacency matrix which includes self-loops, the entry a_{ij} in A corresponds to the weight of the edge connecting nodes v_i and v_j, the labels of nodes are represented by the matrix L and the number of nodes that have labels is m, and D is a diagonal matrix where the element at position (i, i) is equal to the sum of the elements in the i-th row of A. We iterate the following steps until convergence [41]:

$$
L^{(k+1)} = D^{-1/2} A D^{-1/2} L^{(k)},
\tag{2}
$$

$$
l_i^{(k+1)} = l_i^{(0)}, \forall i \leq m
\tag{3}
$$

D. Graph Convolutional Network

GCN is a multi-layer feed-forward neural network designed to propagate and transform node features within a graph. The GCN model demonstrates linear scalability concerning the quantity of edges in the graph, and acquires hidden layer representations that effectively capture the local structure of the graph and the distinctive features associated with each individual node [41]. The layer-wise propagation rule of GCN is represented by the equation 4:

$$
X^{(k+1)} = \sigma(D^{-1/2} A D^{-1/2} X^{(k)} W^{(k)})
\tag{4}
$$

where σ represents an activation function such as ReLU, $W^{(k)}$ denotes a weight matrix in the k-th layer that is trainable, and $X^{(k)} = [\mathrm{x}_1^{(k)}, \mathrm{x}_2^{(k)}, ..., \mathrm{x}_n^{(k)}]^T$ corresponds to the node representations in the k-th layer, with $X^{(0)}$ being the initial node features [42]. In our particular case, nodes displaying a higher level of confidence in being classified as positive or negative, such as those falling within the 90th percentile and above, as determined by the label propagation model, are added into the seed nodes and subsequently passed to GCN for further training. In addition, FinBERT embeddings [7] are adopted as the initial node features $X^{(0)}$.

IV. Experimental Setup

A. Evaluation Datasets

To enable fair comparisons of different lexicons, we conduct sentiment analysis experiments on four well-received datasets, i.e., the PhraseBank, SemEval 2017 Task 5, FiQA Task 1 and SEntFiN. The PhraseBank dataset, established by [4] in 2014, represents a significant milestone dataset for FSA. It comprises 4,846 pieces of news that have been categorized into positive, neutral, and negative sentiments by 16 individuals possessing expertise in financial markets from an investor perspective. The dataset includes four reference datasets, each based on the level of agreement among annotators, namely 100%, 75%, 66%, and 50% agreement. In this study, the 100% agreement dataset is adopted as the benchmark dataset. The SemEval 2017 Task 5 dataset was created with a focus on fine-grained sentiment analysis. Each entry in the dataset is linked to specific target entities and accompanied by sentiment scores. The training data comprises 1,142 headlines from financial news and 1,694 posts from microblogs [43]. The test data encompass 491 financial news headlines and 794 posts.

The annotation process for this dataset involved manual annotations conducted by three independent financial experts, following the guidelines defined by [43]. The dataset used for FiQA Task 1 comprises a collection of 498 financial news headlines and 675 posts. Notably, all entries in the dataset have undergone rigorous manual annotation, encompassing the identification of target entities, aspects, and the assignment of corresponding sentiment scores [44]. Furthermore, the SEntFiN dataset is released and made publicly available by [45] to promote FSA research. SEntFiN is a human-annotated dataset that includes 10,753 news headlines with their entities and corresponding sentiment polarities i.e., positive, neutral and negative. The experiment was designed as a multi-class classification problem. When explicit polarity labels are not provided in the dataset, we assigned the polarity label as positive if the sentiment score exceeded 0.25, negative if it was below -0.25, and neutral if it is in the range of -0.25 to 0.25, inclusive.

B. Benchmark Lexicons and Models

The financial lexicons utilized as benchmark resources include HFD, LM, and SMSL. HFD, developed by Henry in 2008, holds the distinction of being the first dictionary specifically designed for the financial domain. It encompasses

85 negative words and 104 positive words and serves the primary purpose of evaluating the tone expressed in earnings press releases, which play a crucial role in the communication process between firms and investors [23]. LM, as the most widely utilized sentiment word list in the finance domain, is constructed based on the analysis of annual reports released by companies. LM comprises 2,355 negative words, 354 positive words, 19 strong modal words, 27 weak modal words, 297 uncertainty-related words, 904 litigious words, and 184 constraining words [20]. SMSL, the Stock Market Sentiment Lexicon, is developed based on labeled data obtained from StockTwits, a microblogging platform specializing in stock market discussions.

With a collection of 20,550 words and phrases, SMSL has demonstrated promising results in measuring investor sentiments [19]. Its creation from real-time social media data makes it particularly relevant for capturing sentiment within the finance domain. We applied the method introduced by [46] to perform FSA on benchmark datasets using SMSL. Furthermore, in addition to domain-specific lexicons, general-purpose lexicons are also employed as benchmark resources. Notably, these include the more recent SenticNet 7 and VADER. SenticNet 7 is a neurosymbolic AI framework that exploits subsymbolic models, including auto-regressive language models and kernel methods, in order to construct symbolic representations that facilitate the transformation of natural language into a form of protolanguage. This conversion process enhances the system's ability to accurately infer polarity from textual data. It functions as a comprehensive knowledge base comprising 361,654 words and concepts, each intricately linked to comprehensive emotional information [17].

On the other hand, VADER is specifically designed for sentiment analysis in social media contexts. It encompasses 7,520 emoticons, emojis, and words, each assigned sentiment scores [22]. Lastly, the most recent FinBERT, which is developed and released by [9] has also been employed as the benchmark model.

V. Result and Analysis

The sentiment analysis performance is evaluated by adopting accuracy and macro-averaged F1-Score as metrics, with the results presented in Table I. It is observed that FinSenticNet consistently achieves top three performance, outperforming both general-purpose and financial lexicons by a significant margin across evaluation metrics.

It also surpasses FinBERT on SemEval 2017 Task 5 and FiQA Task 1. Such performance highlights the effectiveness of FinSenticNet in accurately capturing and analyzing sentiment within the financial domain. It is important to note that FinBERT attains the highest score on the PhraseBank dataset; however, this is primarily due to its fine-tuning specifically on that particular dataset. In reality, FinBERT encounters notable challenges in terms of generalization and exhibits suboptimal performance when applied to previously unseen datasets such as SemEval 2017 Task 5 and FiQA Task 1. Lastly, SenticNet demonstrates good performance on SemEval 2017 Task 5

TABLE I
COMPARISON WITH OTHER LEXICONS ON 4 BENCHMARK DATASETS. BOLDFACE INDICATED THE TOP 3 RESULT.

Lexicon	PhraseBank		SemEval 2017 Task 5		FiQA Task 1		SEntFiN	
	Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1
SentiWordNet	0.5525	0.3159	0.3260	0.2734	0.3069	0.2617	0.3795	0.2797
VADER	0.5711	0.4871	0.4459	**0.4441**	0.4552	0.4503	0.5112	0.5089
SenticNet	0.2698	0.2194	**0.4910**	0.3753	**0.5183**	0.3803	0.4075	0.3433
HFD	**0.8105**	**0.7714**	0.4311	0.4259	0.4595	**0.4589**	**0.5983**	**0.5864**
LM	0.6444	0.3688	0.3655	0.2596	0.3768	0.3806	0.4926	0.4566
SMSL	0.2800	0.2890	**0.4972**	0.4092	**0.5396**	0.4402	0.4358	0.4171
FinBERT	**0.9169**	**0.8970**	0.4536	**0.4553**	0.4859	**0.4898**	**0.6656**	**0.6676**
FinSenticNet	**0.7619**	**0.7216**	**0.4972**	**0.5048**	**0.5012**	**0.5092**	**0.5920**	**0.5939**

TABLE II
ABLATION STUDY ON 4 BENCHMARK DATASETS. BOLDFACE INDICATED THE BEST RESULT.

Module	PhraseBank		SemEval 2017 Task 5		FiQA Task 1		SEntFiN	
	Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1
FinSenticNet (w/ LPA)	0.7429	0.6752	0.4754	0.4787	0.4663	0.4711	0.5502	0.5397
FinSenticNet (w/ GCN)	0.7513	0.7127	0.4848	0.4949	0.4901	0.5014	0.5676	0.5671
FinSenticNet (w/ LPA-GCN)	**0.7619**	**0.7216**	**0.4972**	**0.5048**	**0.5012**	**0.5092**	**0.5920**	**0.5939**

and FiQA Task 1, while HFD shows promising results on PhraseBank and SEntiFin datasets.

VI. Ablation Study

To validate the effectiveness of different components within the framework, an ablation analysis is conducted, and the corresponding results are presented in Table II. The result shows that lexicons derived from the two-stage LPA-GCN approach demonstrate superior performance compared to lexicons generated solely from either LPA or GCN. The integration of high-confidence sentiment words and concepts generated by LPA has enhanced the training of GCN. It is worth noting that the lexicons constructed using GCN exhibit better performance than those generated solely from LPA, emphasizing the significance of incorporating node feature information. In summary, the LPA component facilitates the dissemination of node label information across graph edges, while the GCN component promotes the propagation and transformation of node feature information.

VII. Conclusion

An automatic semi-supervised framework consisting of a concept parser, a semantic graph, and a two-stage LPA-GCN, is proposed for the construction of concept-level financial lexicons in this paper. The framework's effectiveness is demonstrated by the accurate financial sentiment analysis results. We achieved competitive results when compared to robust benchmark models on four financial sentiment analysis datasets. FSA represents a challenging yet crucial problem in the field of computational finance. The learning-based methods often encounter the issue of overfitting on the training dataset, leading to challenges in generalization when applied to new data. Conversely, lexicon-based methods present their own limitations, including concerns related to accuracy and coverage, which can impact the reliability and comprehensiveness of the results.

Our future research will focus on expanding and refining the semantic graph by including additional nodes and edges

to improve the representation of semantic relationships within the graph. Additionally, we aim to enhance the coverage of FinSenticNet by incorporating a wider range of financial concepts and terms, and also introduce sentiment intensity analysis into FinSenticNet, enabling a more nuanced understanding of sentiment expressions within financial texts.

REFERENCES

[1] M. Van de Kauter, D. Breesch, and V. Hoste, "Fine-grained analysis of explicit and implicit sentiment in financial news articles," *Expert Systems with Applications*, vol. 42, no. 11, pp. 4999–5010, 2015.

[2] Y. Ma, R. Mao, Q. Lin, P. Wu, and E. Cambria, "Multi-source aggregated classification for stock price movement prediction," *Information Fusion*, vol. 91, pp. 515–528, 2023.

[3] F. Xing, E. Cambria, and R. Welsch, "Intelligent asset allocation via market sentiment views," *IEEE Computational Intelligence Magazine*, vol. 13, no. 4, pp. 25–34, 2018.

[4] P. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala, "Good debt or bad debt: Detecting semantic orientations in economic texts," *Journal of the Association for Information Science and Technology*, vol. 65, no. 4, pp. 782–796, 2014.

[5] R. Mao, K. Du, Y. Ma, L. Zhu, and E. Cambria, "Discovering the cognition behind language: Financial metaphor analysis with MetaPro," in *IEEE ICDM*, 2023.

[6] A. Moreno-Ortiz, J. Fernández-Cruz, and C. P. C. Hernández, "Design and evaluation of senticon: A fine-grained economic/financial sentiment lexicon from a corpus of business news," in *Proceedings of LREC*, 2020, pp. 5065–5072.

[7] D. Araci, "Finbert: Financial sentiment analysis with pre-trained language models," *arXiv preprint arXiv:1908.10063*, 2019.

[8] 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 IJCAI*, 2021, pp. 4513–4519.

[9] A. H. Huang, H. Wang, and Y. Yang, "Finbert: A large language model for extracting information from financial text," *Contemporary Accounting Research*, vol. 40, no. 2, pp. 806–841, 2023.

[10] K. Du, F. Xing, and E. Cambria, "Incorporating multiple knowledge sources for targeted aspect-based financial sentiment analysis," *ACM Trans. on Management Information Systems*, vol. 14, pp. 1–24, 2023.

[11] Z. Wang, Z. Hu, F. Li, S.-B. Ho, and E. Cambria, "Learning-based stock trending prediction by incorporating technical indicators and social media sentiment." *Cognitive Computation*, vol. 15, no. 3, pp. 1092–1102, 2023.

[12] F. Xing, F. Pallucchini, and E. Cambria, "Cognitive-inspired domain adaptation of sentiment lexicons," *Information Processing and Management*, vol. 56, no. 3, pp. 554–564, 2019.

[13] W. Yeo, W. Van Der Heever, R. Mao, E. Cambria, R. Satapathy, and G. Mengaldo, "A comprehensive review on financial explainable AI," *arXiv preprint arXiv:2309.11960*, 2023.

[14] K. Ong, W. Van Der Heever, R. Satapathy, G. Mengaldo, and E. Cambria, "FinXABSA: Explainable finance through aspect-based sentiment analysis," *arXiv preprint arXiv:2303.02563*, 2023.

[15] R. Mao, C. Lin, and F. Guerin, "Word embedding and WordNet based metaphor identification and interpretation," in *Proceedings of ACL*, vol. 1, 2018, pp. 1222–1231.

[16] R. Mao, X. Li, K. He, M. Ge, and E. Cambria, "MetaPro Online: A computational metaphor processing online system," in *Proceedings of ACL*, 2023, pp. 127–135.

[17] E. Cambria, Q. Liu, S. Decherchi, F. Xing, and K. Kwok, "SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis," *Proceedings of LREC*, pp. 3829–3839, 2022.

[18] S. Sohangir, D. Wang, A. Pomeranets, and T. M. Khoshgoftaar, "Big data: Deep learning for financial sentiment analysis," *Journal of Big Data*, vol. 5, no. 1, pp. 1–25, 2018.

[19] N. Oliveira, P. Cortez, and N. Areal, "Stock market sentiment lexicon acquisition using microblogging data and statistical measures," *Decision Support Systems*, vol. 85, pp. 62–73, 2016.

[20] T. Loughran and B. McDonald, "When is a liability not a liability? textual analysis, dictionaries, and 10-ks," *The Journal of Finance*, vol. 66, no. 1, pp. 35–65, 2011.

[21] S. Baccianella, A. Esuli, F. Sebastiani *et al.*, "Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining." in *Proceedings of LREC*, vol. 10, no. 2010, 2010, pp. 2200–2204.

[22] C. J. Hutto and E. Gilbert, "Vader: A parsimonious rule-based model for sentiment analysis of social media text," in *ICWSM*, 2014.

[23] E. Henry, "Are investors influenced by how earnings press releases are written?" *The Journal of Business Communication*, vol. 45, no. 4, pp. 363–407, 2008.

[24] S. Huang, Z. Niu, and C. Shi, "Automatic construction of domain-specific sentiment lexicon based on constrained label propagation," *Knowledge-Based Systems*, vol. 56, pp. 191–200, 2014.

[25] G. Xiong, Y. Fang, and Q. Liu, "Automatic construction of domain-specific sentiment lexicon based on the semantics graph," in *ICSPCC*, 2017, pp. 1–6.

[26] M. Hu and B. Liu, "Mining and summarizing customer reviews," in *Proceedings of KDD*, 2004, pp. 168–177.

[27] F. Sebastiani and A. Esuli, "Sentiwordnet: A publicly available lexical resource for opinion mining," in *Proceedings of LREC*, 2006, pp. 417–422.

[28] A. Esuli and F. Sebastiani, "Pageranking wordnet synsets: An application to opinion mining," in *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, 2007, pp. 424–431.

[29] K. Wang and R. Xia, "A survey on automatic construction methods of sentiment lexicons," *Acta Automatica Sinica*, vol. 42, no. 4, pp. 495–511, 2016.

[30] P. D. Turney, "Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews," in *ACL*, 2002.

[31] R. Y. K. Lau, C. L. Lai, P. B. Bruza, and K. F. Wong, "Leveraging web 2.0 data for scalable semi-supervised learning of domain-specific sentiment lexicons," in *Proceedings of CIKM*, 2011, pp. 2457–2460.

[32] J. Park, H. J. Lee, and S. Cho, "Automatic construction of context-aware sentiment lexicon in the financial domain using direction-dependent words," *arXiv preprint arXiv:2106.05723*, 2021.

[33] E. Cambria, R. Mao, S. Han, and Q. Liu, "Sentic parser: A graph-based approach to concept extraction for sentiment analysis," in *IEEE ICDM Workshops*, 2022, pp. 413–420.

[34] P. Qi, Y. Zhang, Y. Zhang, J. Bolton, and C. D. Manning, "Stanza: A python natural language processing toolkit for many human languages," in *Proceedings of ACL*, 2020, pp. 101–108.

[35] C. Fellbaum, "Wordnet and wordnets," in *Encyclopedia of Language and Linguistics*, A. Barber, Ed. Elsevier, 2005, pp. 2–665.

[36] E. Rahimikia, S. Zohren, and S.-H. Poon, "Realised volatility forecasting: Machine learning via financial word embedding," *arXiv preprint arXiv:2108.00480*, 2021.

[37] Z. Ahanin and M. A. Ismail, "A multi-label emoji classification method using balanced pointwise mutual information-based feature selection," *Computer Speech & Language*, vol. 73, p. 101330, 2022.

[38] M. Darwich, S. A. Mohd, N. Omar, and N. A. Osman, "Corpus-based techniques for sentiment lexicon generation: A review." *J. Digit. Inf. Manag.*, vol. 17, no. 5, p. 296, 2019.

[39] X. Zhu and Z. Ghahramani, "Learning from labeled and unlabeled data with label propagation," Carnegie Mellon University, Tech. Rep., 2002.

[40] X. Zhu, *Semi-supervised learning with graphs*. Carnegie Mellon University, 2005.

[41] H. Wang and J. Leskovec, "Combining graph convolutional neural networks and label propagation," *ACM Trans. Inf. Syst.*, vol. 40, no. 4, pp. 1–27, 2021.

[42] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations (ICLR)*, 2017.

[43] K. Cortis, A. Freitas, T. Daudert, M. Huerlimann, M. Zarrouk, S. Handschuh, and B. Davis, "SemEval-2017 task 5: Fine-grained sentiment analysis on financial microblogs and news," in *Proceedings of SemEval-2017*, 2017.

[44] M. Maia, S. Handschuh, A. Freitas, B. Davis, R. McDermott, M. Zarrouk, and A. Balahur, "WWW'18 open challenge: financial opinion mining and question answering," in *Companion Proceedings of the Web Conference 2018*, 2018, pp. 1941–1942.

[45] A. Sinha, S. Kedas, R. Kumar, and P. Malo, "Sentfin 1.0: Entity-aware sentiment analysis for financial news," *Journal of the Association for Information Science and Technology*, 2022.

[46] H. Tatsat, S. Puri, and B. Lookabaugh, *Machine Learning and Data Science Blueprints for Finance*. O'Reilly Media, 2020.