

A Research Agenda for Financial Opinion Mining

Chung-Chi Chen,¹ Hen-Hsen Huang,^{2,3} Hsin-Hsi Chen^{1,3}

¹ Department of Computer Science and Information Engineering, National Taiwan University, Taiwan

² Department of Computer Science, National Chengchi University, Taiwan

³ MOST Joint Research Center for AI Technology and All Vista Healthcare, Taiwan
cjchen@nlg.csie.ntu.edu.tw, hhhuang@nccu.edu.tw, hhchen@ntu.edu.tw

Abstract

Opinion mining is a prevalent research issue in many domains. In the financial domain, however, it is still in the early stages. Most of the researches on this topic only focus on the coarse-grained market sentiment analysis, i.e., 2-way classification for bullish/bearish. Thanks to the recent financial technology (FinTech) development, some interdisciplinary researchers start to involve in the in-depth analysis of investors' opinions. In this position paper, we first define the financial opinions from both coarse-grained and fine-grained points of views, and then provide an overview of the issues already tackled. In addition to listing research issues of the existing topics, we further propose a road map of fine-grained financial opinion mining for future researches, and point out several challenges yet to explore. Moreover, we provide possible directions to deal with the proposed research issues.

Introduction

Dealing with the financial domain data is one of the hot research directions in the artificial intelligence (AI) community. Following the recent trend of financial technology (FinTech), several workshops are held in conjunction with major conferences such as FinNLP¹, ECONLP², and FNP³. These events reflect the increasing interest of AI researchers in financial and economic domains. The special track in IJCAI-2020, AI in FinTech, also evidences this phenomenon.

More and more interdisciplinary research results are published in both finance and computer science communities. Some works (Sedinkina, Breikopf, and Schütze 2019; Qin and Yang 2019) introduce the earning conference call, which is one of the important meetings for announcing the news of a company, to the natural language processing (NLP) community. Some works (Maia et al. 2018; Chen, Huang, and Chen 2019a) pay attention to financial social media data, and propose novel tasks for in-depth investigations. These works indicate the trend of fine-grained opinion mining in the financial domain.

When mentioning the opinion in Finance, bullish/bearish comes into most people's minds. However, market sentiment

is just one component of financial opinion in the financial industry. This paper aims to provide an overview of where we are in fine-grained financial opinion mining and help the community understand where we should be in the future. For understanding the past and the present works, we discuss the components of the financial opinions one-by-one with related works. During the discussion, we will point out some possible research issues. For future research directions, we mainly focus on illustrating the new challenges. We provide a research agenda with the directed graphs toward financial opinions.

Related Work

Pang, Lee et al. (2008) and Liu (2020) provide a general overview of sentiment analysis and opinion mining. The overview and survey papers related to opinion mining in the general domain are updated every year (Abirami and Gayathri 2017; Hussein 2018; Tedmori and Awajan 2019). Most of the works focus on the opinions on social media platforms (Li et al. 2019; Soong et al. 2019). Some works focus on specific topics such as product review (Jebaseeli and Kirubakaran 2012) and reputation evaluation (Chiranjeevi, Santosh, and Vishnuvardhan 2019). Although Kumar and Ravi (2016) provide a survey on text mining in finance, few previous work offers an arrangement of opinion mining in finance. This paper will formulate the financial opinion mining task and illustrate a big picture of this research area.

Although some previous surveys have paid attention to text mining in finance (Das et al. 2014; Fisher, Garnsey, and Hughes 2016), they show less solicitude for opinion mining, and only mention a few coarse-grained financial opinion mining tasks. This paper mainly focuses on provide an in-depth look at the recent trend — fine-grained financial opinion mining and proposes a road map for future research.

Financial Opinion Components

In this paper, we discuss the financial opinions by both coarse-grained and fine-grained viewpoints. The notations used in this paper are shown in Table 1. In this section, we discuss these opinion components one by one.

Coarse-grained Financial Opinion

As the opinion mining task in financial domain, the coarse-grained financial opinions can be separated

¹<http://finnlp.nlpfin.com/>

²<https://sites.google.com/view/econlp-2019>

³<http://wp.lancs.ac.uk/cfie/fnp2021/>

Notation	Denotation	Example in Figure 1
e	Target entity, i.e., mentioed financial instrument	\$AAPL
s	Market sentiment	Bullish
h	Opinion holder	Lisa
t^p	Publishing time	1/3/20 11:44PM
t^v	Validity period of an opinion	1/6/20-1/10/20 (this week)
a	Analysis aspect	technical analysis
d	Degree of market sentiment	[1.91%,3.26%]
C	A set of investor’s claims	Price Target:[303,307]
P	A set of premises	price chart
$M_{t^p}^e$	Market information set of e before t^p	Close price:297.32
q	Opinion quality	Low
ip	Influence Power	Low

Table 1: Notations used in this paper.

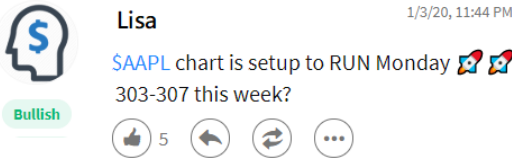


Figure 1: Example of investor’s opinion.

into two (bullish/bearish or positive/negative) or three (bullish/bearish/neutral or positive/negative/neutral) classes, and each opinion is related to one target entity. In most cases, the opinion holder and the publishing time are given. All of the above information can be used to construct a 4-tuple to represent a coarse-grained opinion:

$$(e, s, h, t^p),$$

where e denotes the target entity, s denotes the sentiment, h denotes the opinion holder, and t^p denotes the publishing time. Figure 1 shows an example on one famous financial social media platform, StockTwit⁴. The 4-tuple of this tweet is

$$(\$AAPL, \text{Bullish}, \text{Lisa}, 1/3/20 \ 11:44\text{PM})$$

Because the platform and users provide all essential terms, researchers can easily collect lots of labeled data from the platform. Therefore, many previous works construct a market sentiment lexicon with the data from this platform, and lots of previous works use the labels to test their sentiment analysis models (Oliveira, Cortez, and Areal 2016; Li and Shah 2017; Chen, Huang, and Chen 2018).

Fine-grained Financial Opinion

In this section, we discuss the related components one-by-one to illustrate fine-grained financial opinion. The first one is the aspect of the opinion. Taking the tweet in Figure 1 as an example, the analysis aspect is technical analysis. Maia et al. (2018) and Chen, Huang, and Chen (2019b) provide datasets for extracting the analysis aspect of the investors.

⁴<https://stocktwits.com>

The other important component is the degree of market sentiment (d). Some works in the general domain extend the sentiment into five classes based on the strength (Balikas, Moura, and Amini 2017; Akhtar et al. 2019). In the financial domain, Cortis et al. (2017) label the degree of sentiment into the range between -1 and 1 . In the example in Figure 1, the investor claim that the price of \$AAPL will in the range of 303 to 307 in the coming trading days (one week). The other characteristic of the investor’s opinion is that the market information of the target entity is crucial for understanding the investor’s opinion. For example, the closing price is given every day, and it can provide a base for evaluating the degree of sentiment. Taking 303 and 307 in Figure 1 as instances, these numerals cannot provide any information if we do not compare it with the closing price of \$AAPL. When the closing price 297.32 is given, we can infer d as [1.91%, 3.26%] by a simple calculation. This method is more rational than those labels from -1 to 1 by the intuition of annotators in the previous work (Cortis et al. 2017). Here, we raise the first research question:

(RQ1) How to detect the claims in a financial opinion?

In most opinion mining tasks, the opinions do not have a validity period. Since the financial market changes all the time, the investor’s opinions do have a validity period, even the opinions of professional stock analysts are the same. Most of the analysis reports of professional analysts set the validity period within one year or even shorter. The aforementioned lead to:

(RQ2) How long is the validity period of the investor’s opinions?

Most of the previous works (Bollen and Mao 2011; Valencia, Gómez-Espinoza, and Valdés-Aguirre 2019) adopt market movement prediction as a downside task of capturing investor’s sentiment, and coarsely use the averaged sentiment score from all investors. Wang et al. (2015) indicate that the top investors, ranking by their history performances, on social media platforms can achieve 75% accuracy on market movement prediction. For reference, the accuracy of the recent market movement prediction model (Feng et al. 2019) is in the range of [53.05%, 57.20%]. In the financial domain, the exaggerated information (Chen et al. 2019) may

influence the market, and the opinions with exaggerated information may also be doubtful. Therefore, the quality of a financial opinion is also an open issue. The above discussion arises the following research questions:

- (RQ3) How to evaluate the quality of a financial opinion?
(RQ4) What kinds of financial opinions are trustable?

Future Research Directions

Argument Mining in Finance

Argument mining is one of the focused topics in the AI community recently. Cabrio and Villata (2018) and Lawrence and Reed (2019) provide the surveys to the recent development of argument mining. In our view, it can be considered as the next stage of fine-grained financial opinion mining. Previous works and the above sections only focus on extracting the opinions of the investors or customers. In this section, we discuss the importance of mining the premises and evaluating rationales of a financial opinion.

In order to clarify the tasks, we use a passage (E2) selected from professional analysis as an example. The target entity of (E2) is TSMC, and there are one fact (F1), three premises (P1-3), and one claim (C1) in (E2).

Example (E2):

(F1) *The overall revenue of semiconductor industry 10–11/2018 is in line with expectations. As (P1) the company’s leading-edge in high-end process production continues to increase, coupled with (P2) Globalfoundries’ withdrawal from competition and (P3) inconsistencies in Intel’s process conversion, we estimate that (C1) TSMC’s revenue in 4Q18 will approximate to 9.35 billion US dollars.*

The first challenge of in-depth opinion analysis is detecting the opinion and the rationales, i.e., the claim and the premise. In the financial market, the investors debate on different financial instruments every day with different stances, bullish or bearish. It is just like the situations where the debaters discuss different topics on the affirmative and negative sides. The detection task is necessary if we attempt to analyze the claims and the premises of the investors.

Aiming to make the AI models becomes explainable, AI scientists strive to prove and evidence for the predictions of models. However, in the financial opinion mining field, people use all kinds of opinions directly without asking why. For example, should we give the same weight to the tweet “\$TSMC Gooooo!” and (E2) when analyzing the investors’ opinions? In most of the previous works, their weights are the same. It shows that there is still room for future researches.

One of the further research issues after claim and premise detection is relation linking. In a narrative of an opinion, investors or customers may propose more than one claim with several premises. After predicting the relationship, an opinion can be transformed into a graph, as shown in Figure 2 (a). Now, the claim set C of an opinion may contain several premises denoted by set P .

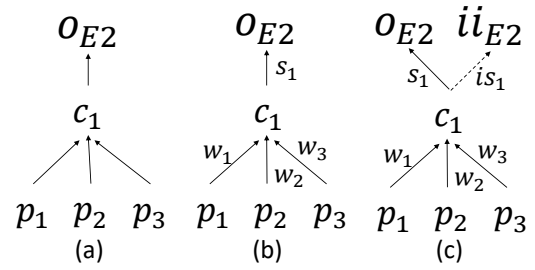


Figure 2: Directed graph of between financial opinion and arguments.

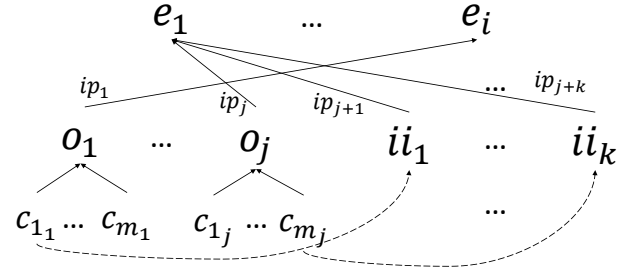


Figure 3: Directed graph between financial opinions and entities.

Quality Evaluation

After extracting the claims and the premises of an opinion, we can evaluate the opinion quality based on the extracted results. Taking Figure 2 (b) for illustration, we first evaluate the rationality of the premises, and get the rationality scores (w_{1-3}) of the premises. We can further add up the rationality scores as the strength score (s_1) of the claim. In summary, this section provides a possible direction for (RQ3). That is, we can evaluate the opinions based on their claims (Chen, Huang, and Chen 2020) and premises (Chen, Huang, and Chen 2021).

Inferring Implicit Influence

Unlike other argument mining tasks, with the nature in the financial domain, we can infer the implicit influence from an opinion. That is, the bullish opinion of e could be bearish information of the other financial instrument. We illustrate the implicit influence of (E2) in Figure 2 (c). The claim (C1) in the example (E2) may also influence the other company in the semiconductor industry. Therefore, we can infer the implicit influence (ii_{E2}) based on the influence score (is_1). To sum up, this section points out a research issue as follows:

- (RQ5) How to infer the implicit influence embedded in an opinion?

Retrieval and Summarization

Now, we complete the fine-grained opinion mining task on individual opinions. The next stage is to compare the opinions and provide a global view. We provide a directed graph in Figure 3. Here, we use the case in financial instruments as

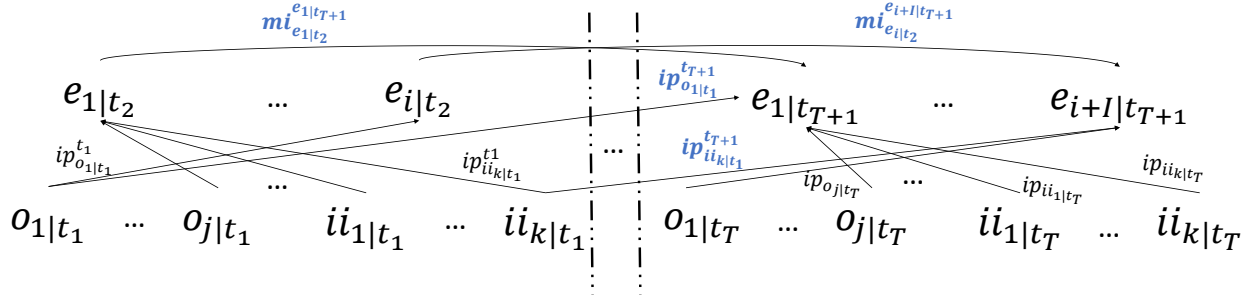


Figure 4: Directed graph in time series. The bold (blue) terms are the new terms that we discuss in Section .

an example. Image that we are in the world with i financial instruments (e), j investors' opinions (O) with k implicit influences (ii), and each investor's opinion has m_j claims. Each opinion node and implicit influence node have their influence power (ip) toward specific financial instruments. Based on the thought, some research questions emerge:

- (RQ6) How to evaluate the influence power of an opinion?
- (RQ7) What is the relationship between the influence power and other components?

Finally, a financial opinion can be represented as an 12-tuple as follows:

$$(e, s, h, t^p, t^v, a, d, C, P, M_{ip}^e, q, ip)$$

Tracing in Time Series

When analyzing financial data, time is an essential component that should be considered. Till now, we only discuss the opinion at a certain time, i.e., t^p . However, the opinion at time t will not influence the status of the target entity at time t . As shown in Figure 4, the opinion at t_1 ($o_{1|t_1}$) may influence e_1 at time $T + 1$ ($e_{1|t_{T+1}}$).

With the concept of time series, three kinds of opinions may exist in time $t + 1$: (1) the new opinion in time $t + 1$, (2) the opinion in time t continuing to exist in time $t + 1$, i.e., t^v has not passed yet, and (3) the opinion in time t changing at time $t + 1$. The opinion may change due to the other opinions in time t . That is, there exists an interaction between the opinions, and here arises the other research question:

- (RQ8) How to evaluate or capture the interaction between the opinions?

The last interaction that we should consider is the one between e_i . The status of e_i at time t may influence the status of itself at time $t + 1$ and the status of other entities at time $t + 1$. The influence between entities is denoted as mi (market influence). It can be linked to the research question in the microeconomics field. Now, the overall picture from a financial opinion to the target entity is complete.

Possible Solutions

Relations of Components

Some components could be inferred based on other components. We list some examples as follows. The sentiment (s)

and the degree of sentiment (d) could be deduced from the comparison between market information (M_{tp}) and the fine-grained claims (c) such as price target. The opinion quality (q) could be analyzed based on the claims' strength (C). The influence power (ip) could be a function of the opinion holder (h) and the opinion quality (q). For example, the tweet of the president of the United States may have a higher ip than that of a common person. The implicit influence (ii) is also an interesting and complex research issue.

Entity Status Evaluation and Prediction

Analyzing the status of e_i is the final purpose of analyzing opinions. For example, the status of e_i could be the stock price. In Figure 3, we show that the current status of the entity can be formulated by the opinions related to it. This information could be used for evaluating the current reputation of the entity. As shown in Figure 4, it could also be the cue for predicting the entity's future status.

Based on the discussions in the paper, we suggest the researchers interested in this field pay more attention to completing the financial opinion tuple and the proposed graph of financial opinions. Future works can be an in-depth analysis of the relations between the target entities and the components. In this way, the decision-making process will become explainable and more rational.

Conclusion

This position paper provides an overview of fine-grained financial opinion mining and proposes the comprehensive directed graphs for real-world interaction between financial opinions and entities. We indicate 8 research questions for future works and provide feasible research directions for them. Besides, we also point out several important but untackled challenges in fine-grained financial opinion mining. Our intent is to depict a big picture for researchers who involve in expediting the development of this topic.

Acknowledgments

This research was partially supported by Ministry of Science and Technology, Taiwan, under grants MOST 109-2218-E-009-014, MOST 109-2634-F-002-040, MOST 109-2634-F-002-034, and MOST 110-2634-F-002-028.

References

- Abirami, A.; and Gayathri, V. 2017. A survey on sentiment analysis methods and approach. In *ICoAC*.
- Akhtar, S.; Ghosal, D.; Ekbal, A.; Bhattacharyya, P.; and Kurohashi, S. 2019. All-in-One: Emotion, Sentiment and Intensity Prediction using a Multi-task Ensemble Framework. *IEEE Transactions on Affective Computing* .
- Balikas, G.; Moura, S.; and Amini, M.-R. 2017. Multitask learning for fine-grained twitter sentiment analysis. In *SIGIR*.
- Bollen, J.; and Mao, H. 2011. Twitter mood as a stock market predictor. *Computer* 44(10).
- Cabrio, E.; and Villata, S. 2018. Five Years of Argument Mining: a Data-driven Analysis. In *IJCAI*.
- Chen, C.-C.; Huang, H.-H.; and Chen, H.-H. 2018. NTUSD-Fin: a market sentiment dictionary for financial social media data applications. In *FNP*.
- Chen, C.-C.; Huang, H.-H.; and Chen, H.-H. 2019a. Crowd view: converting investors' opinions into indicators. In *IJCAI*.
- Chen, C.-C.; Huang, H.-H.; and Chen, H.-H. 2019b. Numerical attachment with auxiliary tasks. In *SIGIR*.
- Chen, C.-C.; Huang, H.-H.; and Chen, H.-H. 2020. NumClaim: Investor's Fine-grained Claim Detection. In *CIKM*.
- Chen, C.-C.; Huang, H.-H.; and Chen, H.-H. 2021. Evaluating the Rationales of Amateur Investors. In *The World Wide Web Conference*.
- Chen, C.-C.; Huang, H.-H.; Takamura, H.; and Chen, H.-H. 2019. Numeracy-600K: Learning Numeracy for Detecting Exaggerated Information in Market Comments. In *ACL*.
- Chiranjeevi, P.; Santosh, D. T.; and Vishnuvardhan, B. 2019. Survey on Sentiment Analysis Methods for Reputation Evaluation. In *Cognitive Informatics and Soft Computing*.
- Cortis, K.; Freitas, A.; Daudert, T.; Huerlimann, M.; Zarrouk, M.; Handschuh, S.; and Davis, B. 2017. Semeval-2017 task 5: Fine-grained sentiment analysis on financial microblogs and news. In *SemEval*.
- Das, S. R.; et al. 2014. Text and context: Language analytics in finance. *Foundations and Trends® in Finance* .
- Feng, F.; Chen, H.; He, X.; Ding, J.; Sun, M.; and Chua, T.-S. 2019. Enhancing stock movement prediction with adversarial training. In *IJCAI*.
- Fisher, I. E.; Garnsey, M. R.; and Hughes, M. E. 2016. Natural language processing in accounting, auditing and finance: A synthesis of the literature with a roadmap for future research. *Intelligent Systems in Accounting, Finance and Management* 23(3).
- Hussein, D. M. E.-D. M. 2018. A survey on sentiment analysis challenges. *Journal of King Saud University-Engineering Sciences* 30(4).
- Jebaseeli, A. N.; and Kirubakaran, E. 2012. A survey on sentiment analysis of (product) reviews. *International Journal of Computer Applications* 47(11).
- Kumar, B. S.; and Ravi, V. 2016. A survey of the applications of text mining in financial domain. *Knowledge-Based Systems* 114.
- Lawrence, J.; and Reed, C. 2019. Argument Mining: A Survey. *Computational Linguistics* .
- Li, Q.; and Shah, S. 2017. Learning stock market sentiment lexicon and sentiment-oriented word vector from stocktwits. In *CoNLL*.
- Li, Z.; Fan, Y.; Jiang, B.; Lei, T.; and Liu, W. 2019. A survey on sentiment analysis and opinion mining for social multimedia. *Multimedia Tools and Applications* 78(6).
- Liu, B. 2020. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge university press.
- Maia, M.; Handschuh, S.; Freitas, A.; Davis, B.; McDermott, R.; Zarrouk, M.; and Balahur, A. 2018. WWW'18 Open Challenge: Financial Opinion Mining and Question Answering. In *WWW*.
- Oliveira, N.; Cortez, P.; and Areal, N. 2016. Stock market sentiment lexicon acquisition using microblogging data and statistical measures. *Decision Support Systems* 85.
- Pang, B.; Lee, L.; et al. 2008. Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval* 2(1-2).
- Qin, Y.; and Yang, Y. 2019. What You Say and How You Say It Matters: Predicting Stock Volatility Using Verbal and Vocal Cues. In *ACL*.
- Sedinkina, M.; Breitkopf, N.; and Schütze, H. 2019. Automatic Domain Adaptation Outperforms Manual Domain Adaptation for Predicting Financial Outcomes. In *ACL*.
- Soong, H.-C.; Jalil, N. B. A.; Ayyasamy, R. K.; and Akbar, R. 2019. The Essential of Sentiment Analysis and Opinion Mining in Social Media: Introduction and Survey of the Recent Approaches and Techniques. In *ISCAIE*.
- Tedmori, S.; and Awajan, A. 2019. Sentiment Analysis Main Tasks and Applications: A Survey. *Journal of Information Processing Systems* 15(3).
- Valencia, F.; Gómez-Espinosa, A.; and Valdés-Aguirre, B. 2019. Price Movement Prediction of Cryptocurrencies Using Sentiment Analysis and Machine Learning. *Entropy* 21(6).
- Wang, G.; Wang, T.; Wang, B.; Sambasivan, D.; Zhang, Z.; Zheng, H.; and Zhao, B. Y. 2015. Crowds on Wall Street: Extracting Value from Collaborative Investing Platforms. In *CSCW*.