An Overview of Financial Technology Innovation

Chung-Chi Chen
Artificial Intelligence Research Center, National Institute of Advanced Industrial Science and Technology, Japan
c.c.chen@acm.org

Hiroya Takamura
Artificial Intelligence Research Center, National Institute of Advanced Industrial Science and Technology, Japan
takamura.hiroya@aist.go.jp

Hen-Hsen Huang
Institute of Information Science, Academia Sinica, Taiwan
hhhuang@iis.sinica.edu.tw

Hsin-Hsi Chen
Department of Computer Science and Information Engineering, National Taiwan University, Taiwan
hhchen@ntu.edu.tw

ABSTRACT

In this paper, we provide an overview of financial technology (FinTech) innovation based on our experience of organizing multiple FinTech-related events since 2018, including the FinNLP workshop series, FinWeb workshop series, and FinNum shared task series. These event series aim to provide a forum for blending the research of FinTech and artificial intelligence (AI), and further accelerating the development in the FinTech domain. We hope that with the researchers' sharing in these events, the challenging problems will be identified, and the future research direction will be shaped. Both the development of the technology and the trend of the focused topics in the past four years are discussed in this paper. We also propose a research agenda with the plan of our FinTech-related events.

CCS CONCEPTS

• **Applied computing**;

KEYWORDS

Financial technology, natural language processing, web intelligence

1 FINTECH TREND

Since 2015, financial technology (FinTech) has become a red-hot topic in the financial industry. At the beginning of the FinTech trend, improving online services such as internet banking and mobile banking was the main topic of many financial institutions. During this period, many startups provide new systems with creative ideas to share the common customers in traditional finance. In addition to providing more convenient services to customers, improving the staffs' working efficiency also gets many attentions.

For example, the tool for collecting information for anti money laundering [9] and the model for financial fraud detection [18] are developed. In academia, some long-term discussed topics such as market information prediction [22, 37] and financial report understanding [10, 28] are explored with the latest neural network architectures. More fine-grained analyses of financial narratives are also explored, e.g., numeral understanding [8, 24, 25] and rationale evaluation [5, 38]. In the introduction stage of the FinTech lifecycle, various innovative ideas are probed in both industry and academia.

Based on our observation in organizing FinTech-related events, we find that there are two emerging topics recently, including regulatory technology (RegTech) [27, 31] and decentralized finance (DeFi). Along with FinTech getting mature, the regulation toward new ideas and new services becomes an important issue because the financial industry is still highly regulated. DeFi is an idea to put financial services to Web3, and it is also one of the targets of RegTech developers. Exploring the topics in real-world finance is the new black in the trend of DeFi. For example, stock valuation is a common topic in traditional finance, and the non-fungible token (NFT) valuation [21] is one of the hot topics among DeFi researchers. Among all topics in FinTech, all eyes are on these two directions currently.

Although FinTech has flourished for about one decade, we think that the explored issues are just the tip of the iceberg. Especially, those related to the Metaverse, i.e., finance in the virtual world, are still in the early stage. Discovering and understanding the potential of DeFi would be one of the essential topics in the latest few years. By contrast, embedded finance, i.e., providing financial services everywhere in the real world, will still be the goal of FinTech developers. The customer service bots [15] and personalized services [34] would be two important application scenarios. With the development of new technology, RegTech, of course, should keep up the pace to protect both users and companies.

2 DEVELOPMENT OF THE TECHNOLOGY

This section provides a quick overview of methodologies for addressing FinTech problems. Based on our observation, we can divide the technologies of the FinTech research into the following three categories.

(1) **Adopting State-of-the-Art Models to Financial Applications**: Since the task settings of machine learning models are similar, it is intuitive to adopt best-performing models for dealing with financial issues. For example, BERT

or capsule network is shown to perform well in general classification tasks. These models can also be employed in domain-specific classification tasks such as stock movement prediction [23, 37].

(2) **Adding Domain-Specific Features to Machine Learning Models**: The data in different domains have their own domain-specific features. For example, containing plenty of numbers is one of the characteristics of financial narratives [6]. Capturing this feature can improve the performance of financial document understanding [4]. The empirical analysis results in financial studies could also be taken into account in model construction. For example, the non-GAAP measures are considered as a kind of the critical information in financial documents [3]. We can either add this feature to models for enhancing downstream tasks such as volatility forecasting [7] or analyze the impact of this feature on the performance of downstream tasks like sentiment analysis [36].

(3) **Constructing Tailor-Made Methods for Financial Applications**: In order to encode the domain-specific data to models, some studies investigate tailor-made methods. For instance, asset prices are the special time-series data in the financial domain. The representations of companies can be learned [16]. In addition, co-training with asset price prediction tasks can be applied to enhance the report summarization results [1]. Additionally, encoding transaction networks among companies [30] and contact histories among companies [33] are also shown useful in financial application scenarios.

3 RESEARCH DIRECTIONS

To be complementary to previous survey papers in FinTech from several aspects [12–14, 35], we foresee some research directions and share our thoughts in this section. We will provide a research agenda and point out the future direction of the series of workshops and shared tasks that we are planning to organize.

3.1 Investing for Long-Term Value

When mentioning "Finance", everything about monetary comes into most people's minds. In our community, plenty of research related to market information prediction has been published every year. Most previous studies focus on predicting next-day, 3-day, or 30-day market information, but few attempt to discuss a longer-term valuation issue. We argue that considering the long-term assessment is an important issue in the investment decision-making process. Taking the definition of "Overweight" rating (the stock will outperform the average return of the stocks that have been analyzed by this analyst or this team in the next six to twelve months) in J.P. Morgan's professional analyst's report as an example, professionals focus on more longer-term evaluation instead of next-day forecast. That shows the importance of long-term influence assessment, and the essential to take the time frame into consideration.

Since the long-term influence assessment is a complex issue and may need a sequence of inferences, we plan to organize a series of shared tasks to encourage our community to think about it. Instead of discussing the monetary factors, we pay attention to the concept of ESG (environmental, social, and governance) in the following years. Since 2005, the concept of ESG has been proposed by UN Global Compact. Then, many researchers have started to analyze the impact of these non-monetary factors and find a way to assess these impacts [2, 29]. With the effort of researchers, the idea of ESG has become more mature, and it has gotten lots of attention since 2020. Inspired by these studies, our goal is to borrow the cogitations and findings in the financial domain to provide possible directions to the researchers in our community, and further lead them to think about investing for long-term value and assessing non-monetary factors like environmental and social impact.

The first step is to probe the fundamental NLP issue: Learning Semantic Similarities. The FinSim of shared task series [11, 20, 26] have been held in FinNLP and FinWeb three times. In FinSim-2022, annual reports and sustainability reports released by companies are adopted as the resource. The organizers plan to propose (1) a task related to classifying the given words into ESG-related taxonomy, and (2) a task related to classifying the given sentences into sustainable or unsustainable descriptions. These semantic learning task settings are expected to explore the models' ability to understand ESG-related terms.

With the exploration in FinSim-2022, we plan to expand FinSim-2023 to a more fine-grained setting: adding the time frame and impact into consideration. That is, when given an event, the task is aimed at estimating the impact (high/medium/low) and the expected time horizon (short-term/medium-term/long-term) of the event toward the company. We hope that these tasks could induce the models to learn to be aware of the time frame and the longer-term assessment of non-monetary events. We also believe that it would be helpful in downstream tasks such as asset valuation.

In sum, investing for long-term value is a very complex task in the financial market. In this section, we share our plan in this direction by proposing new ESG-related tasks. Although we start with the non-monetary factors, similar ideas and approaches could be used in monetary factors in the future. Additionally, integrating both monetary and non-monetary factors would be the final step for making a long-term value assessment.

3.2 Causality and Rationale Inference

Causality and rationales play important roles in financial analysis [17]. Every investor can make a claim (prediction) on the price movement, but not all investors can provide a rational analysis (rationale). Previous studies provide pilot experiments to show that the investors' rationales can be classified into several levels based on the quality [38], and we can also sort out the profitable analysis based on the characteristics of rationales [5]. Since the financial market changes all the time, a rational analysis may not lead to a profitable outcome, and vice versa. That raises two research questions: (1) What is the relationship between the market return and the rationality of investors' opinions? (2) How should the rationality in financial narratives be evaluated?

In order to answer the first research question, we propose a shared task (ERAI: Evaluating the Rationales of Amateur Investors) in FinNLP-2022 at EMNLP-2022. This shared task is aimed at sorting out useful investors' opinions from social media platforms. Different from previous studies that put all available textual data into

models for taking advantage of the wisdom of crowds, we argue that only a few of them really matter. The dataset is available, and the pilot exploration is already done [5]. We expect that with the brainstorming of participants, the relationship between the rationality and the profitability of investors' opinions will reveal.

After exploring a single analysis' rationality, we plan to estimate the rationality based on the discussions of a group of investors in 2023. In most cases, an asset can get a deal because two opposite opinions, i.e., bullish/bearish, exist simultaneously. That means investors' opinions may differ at the same time due to various reasons such as different information they observe and different thoughts on the same event. Based on the logic of "the truth becomes clearer through debate", we think that taking the discussions into consideration is a possible direction for guiding us to understand an opinion's rationality. In other words, an investor can either agree or disagree with the other investor's analysis, and further provide support or attack to the causality and rationale of the analysis. Capturing this kind of relationship could help models better understand investors' discussions.

The second research question, evaluating rationality, is still an open question. For example, how to verify the logic in an investor's analysis is challenging. This question is highly relevant to Explainable AI (XAI). In addition to automatically evaluation, which is widely-adopted in the early stage [19], human evaluation is the most common approach to evaluation, especially for the results of generation tasks. However, human evaluation is time-consuming and also causes some problems [32]. To address this issue in financial narrative understanding and generating, our 2024's plan is to take advantage of the results in 2022 and 2023 to propose an approach for verifying the logic of causality and rationale in financial narratives.

In a word, causality and rationale inference is challenging in both task design and evaluation. We believe that the answers will become much clearer with the series of explorations proposed in this section.

3.3 Professional-Like Generation

In addition to the inference of financial narratives, expressing the results, including analysis and prediction, is essential to constructing a more robust AI. We separate the financial opinion into premises (rationale) and claims (prediction) and introduce the components of financial opinions in our previous work [6]. When generating financial documents and narratives, we can still separate the target into premises and claims.

Based on our observation [4], premise generation would be more difficult than claim generation, because investors, especially professionals, always write a claim with some patterns. For example, "We estimate the [Monetary Factor] will be [Direction] for [Estimation] in the following [Time Frame]." is a common pattern when making claims, where [Monetary Factor] could be EPS, earnings, cost, sales, and any monetary factor that previous works attempt to predict; [Direction] reveals the results of the prediction under binary classification task, i.e., increase/decrease; [Estimation] is a numeral that provides more details of the forecast; [Time Frame] is the expected validity period of this claim. Thus, claim generation could be a kind of the cloze test. When given the available data and [Monetary

Factor], models can learn to predict the [Direction], and estimate the [Estimation] and [Time Frame]. With the given pattern and the forecasts, fluent claims can be generated. Thus, in our opinion, the major issue in claim generation is still the accuracy of the forecasts. Among the aforementioned forecasts, [Time Frame] is the topic with fewer discussions. Thus, in addition to the ESG-related time frame dataset proposed in Section 3.1, we plan to release a dataset for financial opinions' validity period identification and inference. After addressing the time frame issue, the performance of claim generation in financial claims would become better.

Regarding premise generation, we think that there are three steps, including (1) fact abstraction, (2) related [Monetary Factor] inference, and (3) importance ordering. Because of several existing facts, such as news, financial reports, earnings conference calls, and crowd discussions at the time point of generating financial analysis, it is important to extract important facts from plenty of information. We call this step — fact abstraction. Firstly, it is still challenging for models to digest long documents or lots of multimodal data simultaneously. Thus, abstracting factors is essential. Secondly, when generating reports, professionals will not say that "Based on the information in the financial report, we estimate ...". For readers, that is too coarse to understand the rationales. Therefore, sorting out the key points of facts, which would influence the decision, is necessary. This is the reason why we need the second step — related [Monetary Factor] inference. Because the facts may influence the change of different [Monetary Factor], identifying which [Monetary Factor] will be influenced by the given factor become the next step of factor abstraction. Till this step, we can generate a sentence-level financial analysis: "Because [Factor Abstract], we estimate the [Monetary Factor] will be [Direction] for [Estimation] in the following [Time Frame]", which contains both premises and claims and the [Monetary Factor] is not a given condition but an inferred result.

However, investors' will not use such cause-effect sentences to form financial documents and narratives. Additionally, we could generate several sentence-level financial analyses based on each factor. Therefore, selecting the most important parts and merging them into a paragraph or article becomes the next issue — importance ordering. The challenge in this step is not to order the sentence-level analysis, but to consider all factors simultaneously. There would have three choices for each factor: keeping, merging, and removing. That is, we need to keep the important fact, merge some similar narratives, and remove those piddling sentences. Based on the results of these operations, we need to re-estimate by considering cross-factor relationships. Finally, we can make inferences based on the generated claims sequentially. For example, investors may make forecasts for the next four quarters. When we generate the analyses for 2022 Q3 at 2022 Q1, we should not make the prediction based on the data of 2022 Q1 only, but also consider the generated analyses for 2022 Q2. In this step, the explorations in Section 3.2 will provide some clues.

In sum, the professional-like generation task integrates the concepts of long-term value assessment and the notions of causality and rationale inference, and attempts to deliver the models' output in the natural language form. Because our goal is to construct an intelligent machine for the financial domain, in our opinion, this is one of the critical topics in FinTech with NLP.

4 CONCLUSION

This paper briefly summarizes the FinTech trend in the past four years based on the authors' observations during organizing FinTech-related events. Some future directions point out for the researchers interested in this field. The directions of developing the technology for solving financial problems are also present. We hope this overview could help newcomers quickly understand where we are now and where we are going to be. With the sharing of our ongoing plans, we look forward to facilitating the cooperation among all researchers interested in these directions for accelerating the development of the listed crucial issues.

REFERENCES

[1] Yash Agrawal, Vivek Anand, S Arunachalam, and Vasudeva Varma. 2021. Hierarchical model for goal guided summarization of annual financial reports. In *Companion Proceedings of the Web Conference 2021.* 247–254.

[2] Amir Amel-Zadeh and George Serafeim. 2018. Why and how investors use ESG information: Evidence from a global survey. *Financial Analysts Journal* 74, 3 (2018), 87–103.

[3] Dirk E Black, Theodore E Christensen, Jack T Ciesielski, and Benjamin C Whipple. 2018. Non-GAAP reporting: Evidence from academia and current practice. *Journal of Business Finance & Accounting* 45, 3-4 (2018), 259–294.

[4] Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2020. NumClaim: Investor's Fine-grained Claim Detection. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management.* 1973–1976.

[5] Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2021. Evaluating the rationales of amateur investors. In *Proceedings of the Web Conference 2021.* 3987–3998.

[6] Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2021. *From Opinion Mining to Financial Argument Mining.* Springer Nature.

[7] Chung-Chi Chen, Hen-Hsen Huang, Yu-Lieh Huang, and Hsin-Hsi Chen. 2021. Distilling Numeral Information for Volatility Forecasting. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management.* 2920–2924.

[8] Chung-Chi Chen, Hen-Hsen Huang, Yow-Ting Shiue, and Hsin-Hsi Chen. 2018. Numeral understanding in financial tweets for fine-grained crowd-based forecasting. In *2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI).* IEEE, 136–143.

[9] Hao-Yuan Chen, Shang-Xuan Zou, and Cheng-Lung Sung. 2019. Pluto: A deep learning based watchdog for anti money laundering. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing.* 93–95.

[10] Tobias Daudert and Sina Ahmadi. 2019. CoFiF: A corpus of financial reports in French language. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing.* 21–26.

[11] Ismail El Maarouf, Youness Mansar, Virginie Mouilleron, and Dialekti Valsamou-Stanislawski. 2020. The finsim 2020 shared task: Learning semantic representations for the financial domain. In *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing.* 81–86.

[12] Keke Gai, Meikang Qiu, and Xiaotong Sun. 2018. A survey on FinTech. *Journal of Network and Computer Applications* 103 (2018), 262–273.

[13] Peter Gomber, Jascha-Alexander Koch, and Michael Siering. 2017. Digital Finance and FinTech: current research and future research directions. *Journal of Business Economics* 87, 5 (2017), 537–580.

[14] Terrence Hendershott, Xiaoquan Zhang, J Leon Zhao, and Zhiqiang Zheng. 2021. FinTech as a game changer: Overview of research frontiers. , 17 pages.

[15] Jiantao Huang, Yi-Ru Liou, and Hsin-Hsi Chen. 2021. Enhancing Intent Detection in Customer Service with Social Media Data. In *Companion Proceedings of the Web Conference 2021.* 274–275.

[16] Tomoki Ito, Jose Camacho-Collados, Hiroki Sakaji, and Steven Schockaert. 2020. Learning company embeddings from annual reports for fine-grained industry characterization. In *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing.* 27–33.

[17] Kiyoshi Izumi and Hiroki Sakaji. 2019. Economic causal-chain search using text mining technology. In *International Joint Conference on Artificial Intelligence.* Springer, 23–35.

[18] Junxiang Jiang, Boyi Ni, and Chunping Wang. 2021. Financial Fraud Detection on Micro-credit Loan Scenario via Fuller Location Information Embedding. In *Companion Proceedings of the Web Conference 2021.* 238–246.

[19] Adam J Johs, Denise E Agosto, and Rosina O Weber. 2022. Explainable artificial intelligence and social science: Further insights for qualitative investigation. (2022).

[20] Juyeon Kang, Ismail El Maarouf, Sandra Bellato, and Mei Gan. 2021. FinSim-3: The 3rd Shared Task on Learning Semantic Similarities for the Financial Domain. In *Proceedings of the Third Workshop on Financial Technology and Natural Language Processing.* 31–35.

[21] Arnav Kapoor, Dipanwita Guhathakurta, Mehul Mathur, Rupanshu Yadav, Manish Gupta, and Ponnurangam Kumaraguru. 2022. TweetBoost: Influence of Social Media on NFT Valuation. In *Proceedings of the Web Conference 2022.*

[22] Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Zihan Liu, Yan Xu, Cong Gao, and Pascale Fung. 2019. Learning to learn sales prediction with social media sentiment. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing.* 47–53.

[23] Jintao Liu, Hongfei Lin, Xikai Liu, Bo Xu, Yuqi Ren, Yufeng Diao, and Liang Yang. 2019. Transformer-based capsule network for stock movement prediction. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing.* 66–73.

[24] Maofu Liu, Xinxin Xia, and Wei Wang. 2021. A Chinese Dataset for Exploring Financial Numeral Attributes. In *Companion Proceedings of the Web Conference 2021.* 255–259.

[25] Yuxuan Liu, Maofu Liu, and Mengjie Wu. 2022. A Chinese Dataset for Exploring Financial Numeral Attributes. In *Companion Proceedings of the Web Conference 2022.*

[26] Ismail El Maarouf, Youness Mansar, Virginie Mouilleron, and Dialekti Valsamou-Stanislawski. 2020. The FinSim 2020 Shared Task: Learning Semantic Representations for the Financial Domain. In *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing.* 81–86.

[27] Shivam Mangale, Buggana Sathvik, Deepti Saravanan, Shravya Kanchi, Ujwal Narayan, Lini T. Thomas, Kamalakar Karlapalem, and Natraj Raman. 2022. SEBI Regulation Biography. In *Proceedings of the Web Conference 2022.*

[28] Corentin Masson and Syrielle Montariol. 2020. Detecting Omissions of Risk Factors in Company Annual Reports. In *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing.* 15–21.

[29] Pedro Matos. 2020. ESG and responsible institutional investing around the world: A critical review. (2020).

[30] Naoto Minakawa, Kiyoshi Izumi, Hiroki Sakaji, and Hitomi Sano. 2022. Graph Representation Learning of Banking Transaction Network with Edge Weight-Enhanced Attention and Textual Information. In *Companion Proceedings of the Web Conference 2022.*

[31] Ujwal Narayan, Pulkit Parikh, Kamalakar Karlapalem, and Natraj Raman. 2022. Detecting Regulation Violations for an Indian Regulatory body through multi-label classification. In *Proceedings of the Web Conference 2022.*

[32] Avi Rosenfeld. 2021. Better metrics for evaluating explainable artificial intelligence. In *Proceedings of the 20th international conference on autonomous agents and multiagent systems.* 45–50.

[33] Hiroki Sakaji, Ryota Kuramoto, Hiroyasu Matsushima, Kiyoshi Izumi, Takashi Shimada, and Keita Sunakawa. 2019. Financial Text Data Analytics Framework for Business Confidence Indices and Inter-Industry Relations. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing.* 40–46.

[34] Soumajyoti Sarkar. 2021. Centralized Borrower and Lender Matching under Uncertainty for P2P Lending. In *Companion Proceedings of the Web Conference 2021.* 269–273.

[35] Patrick Schueffel. 2016. Taming the beast: A scientific definition of fintech. *Journal of Innovation Management* 4, 4 (2016), 32–54.

[36] Stacey Taylor and Vlado Keselj. 2020. Using Extractive Lexicon-based Sentiment Analysis to Enhance Understanding of the Impact of Non-GAAP Measures in Financial Reporting. In *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing.* 40–46.

[37] Linyi Yang, Yang Xu, Tin Lok James Ng, and Ruihai Dong. 2019. Leveraging BERT to improve the FEARS index for stock forecasting. In *The First Workshop on Financial Technology and Natural Language Processing, Macao, China, 12 August 2019.*

[38] Annie Ying and Pablo Duboue. 2019. Rationale classification for educational trading platforms. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing.* 14–20.