Chapter 6
FinTech Applications

The Financial Stability Board defines financial technology (FinTech) as "technology-enabled innovation in financial services."[1] At the 2015 World Economic Forum, experts proposed a taxonomy for financial services[2] that can be classified into six major categories: payments, deposits and lending, market provisioning, capital raising, insurance, and investment management. Because financial opinion mining can be applied to many listed services, we survey various cases in this chapter and show that financial opinion mining is useful and crucial in many financial application scenarios. In Sect. 6.1, we discuss information provision services in the financial domain. In Sect. 6.2, we discuss work on personalized recommendations, which is the goal of emotional banking. In Sect. 6.3, we discuss applications for improving employee efficiency. In this chapter, we demonstrate the importance of financial opinion mining in the financial industry.

6.1 Information Provision

An analyst is a professional information provider who summarizes current events and produces claims based on these events. That is, analysts not only provide the latest market information, but also offer their view based on all available information. We begin this section with the workflow of a professional analyst.

1. **Information collection**: They collect information from sources listed in Chap. 3 such as insiders and news articles.

[1]https://www.fsb.org/wp-content/uploads/P140219.pdf.

[2]http://www3.weforum.org/docs/WEF_The_future__of_financial_services.pdf.

C.-C. Chen et al., *From Opinion Mining to Financial Argument Mining*,
SpringerBriefs in Computer Science,
https://doi.org/10.1007/978-981-16-2881-8_6

Fig. 6.1 Screenshot of Bloomberg Terminal's sentiment analysis function

2. **Information verification**: They verify the collected information by visiting companies or via discussion with other analysts.
3. **Influence inference**: They infer the potential influence of each piece of information.
4. **Opinion formulation**: They sort out the important parts to produce claims and generate a report.

A professional analyst thus "connects all the dots" to get the full picture. When developing an information provision service, we seek to provide analysts with automated assistance by doing the trivial, tedious work for them.

Information vendors such as Bloomberg and Refinitiv play an important role in the first step of the analyst's workflow. They provide the latest news, quotes, and analysis reports from other organizations, combining all essential data on one platform. They provide not only raw data but also sort out this raw data to produce structured data. Of the sources listed in Chap. 3, information vendors most often neglect the opinions of social media users, despite the many studies [19, 40] that demonstrate the informativeness of such opinions. Hence one challenge is collecting opinions and presenting them in a structured form similar to what information vendors do for the views and opinions of insiders and professionals.

In financial opinion mining, sentiment analysis is the most common topic. As shown in Fig. 6.1,[3] Bloomberg Terminal demonstrates how to visualize the extracted sentiments of social media users with market data. They show counts of positive and negative tweets alongside historical price data. As mentioned in previous chapters, such sentiment comes from coarse-grained investor opinion. However, there are many details in a financial opinion: we here discuss how to collect fine-grained information.

[3] https://www.bloomberg.com/company/press/bloomberg-and-twitter-sign-data-licensing-agreement/.

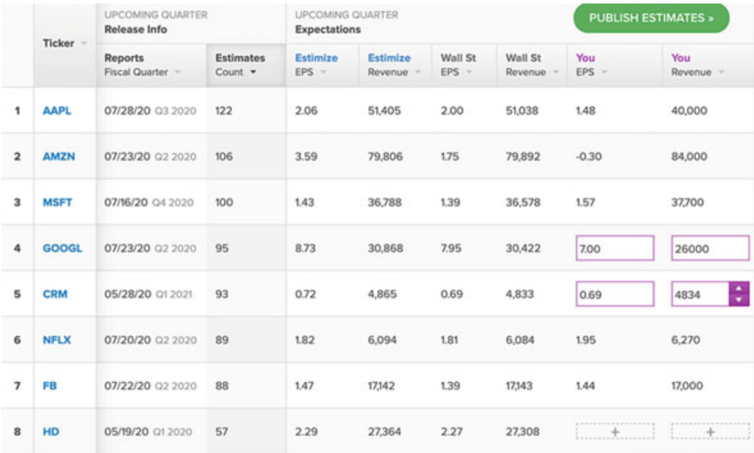

	Ticker	UPCOMING QUARTER Release Info		UPCOMING QUARTER Expectations				PUBLISH ESTIMATES »	
		Reports Fiscal Quarter	Estimates Count	Estimize EPS	Estimize Revenue	Wall St EPS	Wall St Revenue	You EPS	You Revenue
1	AAPL	07/28/20 Q3 2020	122	2.06	51,405	2.00	51,038	1.48	40,000
2	AMZN	07/23/20 Q2 2020	106	3.59	79,806	1.75	79,892	-0.30	84,000
3	MSFT	07/16/20 Q4 2020	100	1.43	36,788	1.39	36,578	1.57	37,700
4	GOOGL	07/23/20 Q2 2020	95	8.73	30,868	7.95	30,422	7.00	26000
5	CRM	05/28/20 Q1 2021	93	0.72	4,865	0.69	4,833	0.69	4834
6	NFLX	07/20/20 Q2 2020	89	1.82	6,094	1.81	6,084	1.95	6,270
7	FB	07/22/20 Q2 2020	88	1.47	17,142	1.39	17,143	1.44	17,000
8	HD	05/19/20 Q1 2020	57	2.29	27,364	2.27	27,308	+	+

Fig. 6.2 Screenshot of Estimize, a service that compiles earnings estimations of its users

Estimize[4] is a FinTech company which compiles earnings estimations of its users. Figure 6.2 shows a screenshot. Users fill out forms, which Estimize uses to calculate the average of all users' estimations. With this information, they compare EPS and revenue estimations from both professional investors and social media users. Jame et al. [21] find that the forecasts provided by Estimize's users improve price discovery. Da and Xing [12] analyze Estimize forecasts from a herding perspective to show that the more public information the user accesses, the less the user shares his/her own private opinion. These works also confirm the accuracy of forecasts from crowdsourcing platforms.

In addition to claims about EPS and earnings, investors also produce forecasts such as price targets. Since many financial opinions are expressed in natural language instead of in a tabular form like that in Fig. 6.2, understanding opinions in unstructured form is another research focus. In Sect. 5.3, we show that price targets from social media users are good predictors of stock movement. In previous work [8], we demonstrate how to visualize this information for investors: Fig. 6.3 shows a screenshot of CrowdPT, the resultant system, which makes it easy for investors compare stock prices with price targets. In addition to price targets, some of the categories in Table 5.3 contain informative opinions. Almost all financial opinions can be converted into an index and shown in charts such as those in Figs. 6.1 and 6.3. In previous work [5], we also show that the distribution of returns based on *buy/sell price* and *support or resistance price* signals from social media users is significantly different from that of randomly selected trading days. These systems and studies all inform methods for automatic information collection, and are also examples of ways to visualize such financial information. These studies support the importance of collecting more fine-grained information as opposed to capturing sentiment only.

[4] https://www.estimize.com/.

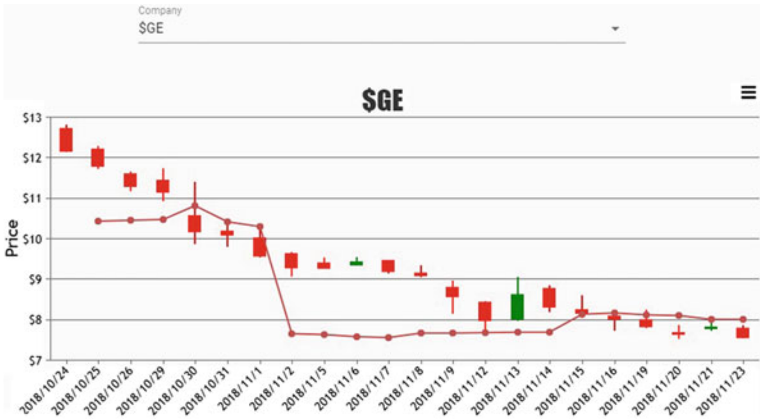

Fig. 6.3 Screenshot of CrowdPT, a system that enables investors to compare stock prices with price targets [8]

Once the information is collected, verification is necessary. Automatic fact-checking is a related research topic. Most objective descriptions of facts in talks or documents released by insiders, professionals, and journalists are correct, reliable information. However, their subjective opinions must be verified. For example, it is important to be able to judge whether a manager's claims in an earnings conference call are rational. It is difficult to design and collect the data needed to train the corresponding models for rationality-checking. In previous work [7], we use market comments to simulate this scenario. According to Chap. 5, numerals are important in financial narratives; managers and investors all focus on numeral information and make claims that include estimations expressed as numerals. In one corresponding verification task, we judge whether a given numeral in a market comment is exaggerated. Take for example (E6.1) and (E6.2): the words in these sentences are the same but the price targets are different. Given a stock which closes at 850, (E6.2)'s price target of 300 is likely an exaggeration. Our experimental results show that models perform well in very irrational cases, but perform worse in instances in which the correct numeral is replaced with a similar value.

(E6.1) We reiterate our buy recommendation and maintain the price target of 900.

(E6.2) We reiterate our buy recommendation and maintain the price target of 300.

Unlike information collected from formal, trustworthy sources, almost all Web information—especially that from social media platforms—must be verified before use. Relevant studies include those on fake news verification [33], fact-checking [16], and even spam detection [22, 34]. The quality evaluation task discussed in Sect. 4.2 is also a related issue.

For the third step—influence inference, in which we estimate the influence of each piece of information—we list some studies that use various kinds of information to predict the impact on future price movement.

- **Financial statements**: Holthausen and Larcker [18] predict excess returns using a logit model with data from financial statements. Their results show that the proposed model earns significant, abnormal returns over the period from 1978 to 1988.
- **Market data**: Liu et al. [28] encode market data at multiple time scales using both an RCNN architecture and a discrete wavelet transform [24] for stock trend prediction. They experiment on two datasets with intra-day price data (the FI-2010 dataset [31] and their CSI-2016 dataset). Their results support the usefulness of considering multi-scale market data for stock trend prediction. Ding et al. [13] experiment on inter-day market data, and also demonstrate the helpfulness of multi-scale representations.
- **Information from insiders**: As mentioned in Sect. 3.1, formal reports and insider talks are all informative and predictive for future returns and risks. For example, Loughran and McDonald [30] show the usefulness of sentiment information for prediction tasks in 10-K reports, and Qin and Yang [32] use earnings conference calls to predict volatility.
- **Investor opinions**: In Sects. 3.2 and 3.3, we discuss the informativeness of professionals' opinions as well as those from users of social media platforms. In Chap. 4, we show how to analyze these opinions. Based on the proposed argumentation structures and the concept of influence power estimation, we can infer the impact of a given opinion on a certain target entity.
- **News**: Event-based market movement prediction has been widely discussed in the NLP community. Many studies use news articles as data sources. For example, Hu et al. [20] propose a hybrid attention network and show that trading based on their model's predictions yields better profits than other baselines. Cheng et al. [10] extract events into tuples, which they then use to construct an event knowledge graph. They also show that their framework is profitable in the stock market.

The above-mentioned studies estimate the probability of future events given financial information. These probabilities can be used in the final step of analysis summarization.

Figure 6.4 illustrates the workflow with the concepts proposed in Chap. 2. We have completed step 3 in the figure. That is, at step 1 we collect information ($I_{i,t}^{e_1}$), and at step 2 we verify this information. That which is identified as fake or exaggerated is removed, and information which is correct becomes the premises ($P_{j,t}^{e_1}$); market data is also a premise. In step 3, we produce inferences based on this verified information. Different models may yield different claims (C^U or C^D). A given model's claims may also vary with the input data. The final step is to summarize the premises and the claims to author a report, which is considered an opinion (O).

The NLP community has proposed datasets and models for use in exploring summarization. For example, Li et al. [26] propose a system that extracts events, then links them, and finally generates a summary based on feature weights. Fabbri et al. [14] publish the Multi-News dataset, which contains more than 50,000 instances, and propose an end-to-end model that merges the pointer-generator network [35] and maximal marginal relevance [3]. For short, text-like tweets, Shapira et al. [36]

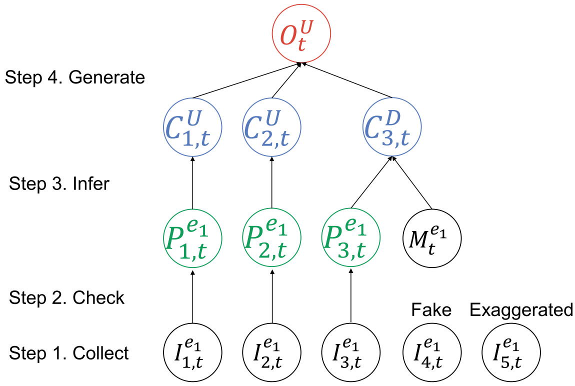

Fig. 6.4 Workflow with the concepts that are introduced in Chap. 2

propose a system based on open knowledge representation [39]. As these works are similar to summarizing premises, future works could borrow their approaches.

Claim generation, however, may be different from premise summarization, because it takes stance into account. Although some studies on argument mining [1, 15, 17] explore claim generation, few generate claims for financial opinions. Many studies in financial opinion mining stop at step 3 in Fig. 6.4. This may be because templates can be used to generate claims. For example, if the model predicts that the price of $AAPL will rise to 200 in the next three months, we can use template (E6.3) to generate the claim (The price target of "$AAPL" is "200".).

(E6.3) The price target of "target entity" is "model's prediction".

Although there are indeed templates that could be used to generate the claims based on the results of step 3, this is still a worthwhile research direction. These claims should be context-aware sentences, and would need to be generated based on the premises. For example, although the price targets are the same in (E6.4) and (E6.5), the meanings of these instances are different: this shows the necessity of exploring context-aware claim generation.

(E6.4) Revenue is expected to decline due to COVID-19, so we lowered our target price to 200.

(E6.5) We adjust our target price to 200, because we believe that the economy rebounds in the second half of the year.

Additionally, the structure and strategy of the resulting report may yield different influences on different readers. For example, Yang et al. [41] analyze the persuasion strategies of crowdfunding posts. This research direction may also be worth exploring when generating professional reports.

In this section, we use the workflow of professional analysts as an example. Every function (information collection, information verification, influence inference, and opinion formulation) could be a service that we provide to customers. For example, we could provide verified information to customers, or we could provide them with model predictions. These functions can be explored based on the concepts discussed here about financial opinion mining. We also illustrate the workflow in Fig. 6.4 based on the ideas proposed in Chap. 2. We suggest that future work follow the proposed steps and rationales to produce innovations in the information provision field.

6.2 Personalized Recommendation

Personalized recommendations are an important function in the next generation of banking, i.e., Bank 4.0 [23]. Neural network models and other advanced architectures yield significant improvements in recommendation. In particular, on platforms like e-commerce platforms that have access to a considerable amount of user data, performance has improved significantly. However, as e-commerce products are different from financial products, we face particular challenges when designing personalized recommendation systems for the financial domain. For example, whereas product prices on e-commerce platforms generally do not change constantly, those for financial instruments in financial markets typically do. Indeed, in the financial market, the prices of stocks, bonds, and options change every day; they can even change repeatedly in the space of a second. Also, on e-commerce platforms, product specifications generally stay the same; for instance, the iPhone 12 Pro uses the Apple A14 Bionic, and will not change to use the Apple A13 Bionic (in most cases). However, in the financial market, a company's operations may change every quarter. As companies are the underlying asset for financial instruments such as stocks and bonds, opinions about the iPhone 12 Pro may still be valuable after a year, whereas opinions about $AAPL are worthless after that same year.

Although some methods can be used for both e-commerce platforms and financial markets we must still account for the characteristics of the financial domain to improve performance. For example, just because someone mentions $AAPL does not mean that should we recommend $AAPL-related tweets to them. Instead, we should first seek to understand why they have mentioned $AAPL. For example, maybe they are interested in stocks that have attained a new 52-week high. In this case, instead of recommending $AAPL-related tweets to them, we should recommend other stocks that have made a new high. In previous work [6], we propose a task called next cashtag prediction, in which we attempt to predict cashtag(s)—that is, stock(s)—that the user will mention in the next five days. We present a tailor-made personalized recommendation method for financial social media platforms. As illustrated in Fig. 6.5, the proposed model uses three kinds of latent vectors:

- **User interest vectors**: The interests of the given user, captured from the tweets posted by the user.
- **Analysis vectors**: Background information on candidate cashtags derived from discussions (tweets) from other users.
- **Chart vectors**: Price data in the form of historical prices and volumes.

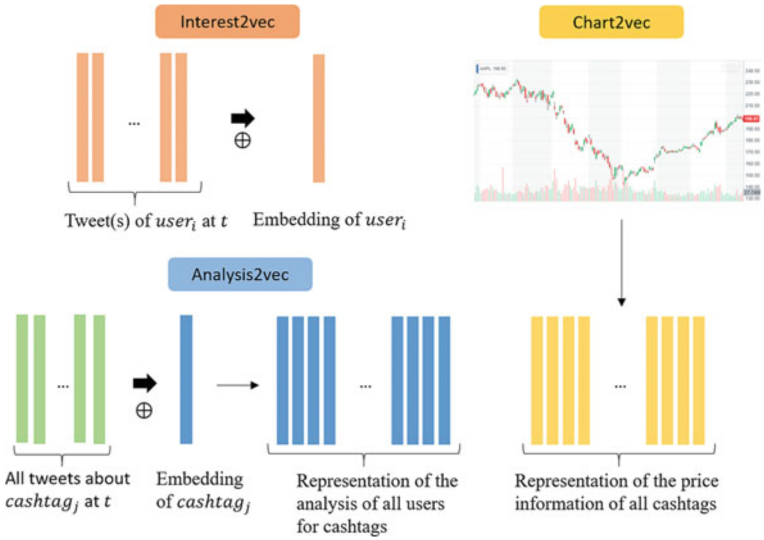

Fig. 6.5 Three kinds of latent vectors, including user interest vectors, analysis vectors, and char vectors

The proposed method achieves a 69.03% $hit@2$ when there are 30 candidate cashtags. This work shows a a direction for personalized investment suggestion.

The following work also provides insights.

- Insurance is also a financial product. Bi et al. [2] present a system for recommending insurance products to cold-start users. They employ user latent features from other domains for the insurance domain, showing the possibility of cross-domain features for financial applications.
- An ideal recommendation system proposes the best product to users based on the user's interests or budget, and also explains its decision. Chen et al. [9] discuss a similar scenario with data from an e-commerce platform. In their system they consider both personalized recommendation and explanation, which are also important in the financial domain. For example, the salesperson not only recommends a fund to the customer, but also explains why the recommended fund is suitable for the customer. In this case, the reason may simply be the salesperson's opinion.

Thus, the consensus of e-commerce-based studies is that customer opinions are essential elements to consider when producing personalized recommendations. This also applies to financial applications. In this section, we have laid out a rough outline of an application for financial opinion in recommendation systems. Previous work also shows that latent features of a given domain can be transferred to the financial domain. Additionally, we show why explaining decisions and recommendations are key functions for future work.

Table 6.1 Customer opinion and implicit relation to stock of credit-card-issuing bank in (E6.6)

Meaning	Example in (E6.6)	Implicit relation to market
Target entity	FlyGo	2887.TW
Market sentiment	–	Bearish
Sentiment	Negative	–
Opinion holder	Lisa	Lisa
Publishing time	2020/1/11	2020/1/11
Validity period of an opinion	–	–
Market information set	Cashback: 1%	Close price: 13.3
Analysis aspect	Cashback	Credit card services
Degree of sentiment	−0.8	−0.3
Set of claims	–	–
Set of premises	Cashback canceled	–
Opinion quality	Low	–
Influence power	Low	Low

6.3 Improving Employee Efficiency

In this section, we discuss how to apply techniques for financial opinion mining to improve employee efficiency in related industries. In previous chapters, we discussed financial opinions about investment and trading; these can be considered investor opinion. In the financial industry, services are important immaterial products, and the opinions on financial services are similar to those in the general domain. We take (E6.6) as an example, where FlyGo is a credit card.

(E6.6) Because the cashback of FlyGo was canceled, I cut it directly.

As shown in Table 6.1, the components defined in Chap. 2 can be used to analyze this opinion. Here, note that the customer's opinion may not provide claims for trading and investment. Thus, we can use positive/negative as the sentiment analysis in the general domain. This kind of opinion may also lack a validity period, because the cashback can change every year. In this case, we must also note the market information, i.e., the FlyGo contract. If the cashback changes in the following year, this negative opinion should not be considered for other users interested in FlyGo. To evaluate the quality of this opinion, we analyze the aspect and degree of sentiment and extract the argumentative units. Influence power in this case may be defined differently from that for an investor's opinion. If (E6.6) were posted by an opinion leader on social media platfroms, the market share of FlyGo could drop. This phenomenon also exists on e-commerce platforms, as we discuss in Sect. 4.3. Related work reviewed in Sect. 4.3 shows that customer opinions influence product sales. Table 6.1 lists information that may be related to the stock of the credit-card-issuing bank implied in (E6.6). This shows that customer opinions are also important in the financial domain.

Because customer service staff face customer opinions daily, customer service in the financial industry is another relevant topic. After extracting customer opinion, we attempt to detect their intent. For example, once we have determined that the target entity is related to credit card services, we could put the call through to the credit card coordinator. That is, we leverage the information we have extracted to detect the customer's intent. Moreover, if we discern that the customer is complaining about the low cashback, we could reduce customer churn by suggesting a better plan for the customer. We can also infer the reasoning behind for customer questions. For example, perhaps the customer asking about the cashback rate ratio of foreign spending is planning to travel overseas. In this case, we could encourage him/her to purchase travel insurance. These scenarios are common cases for financial institutions. Although few studies use data in the financial domain, experience from other domains could be adopted in the future. Below we mention some related work.

- Intent detection is domain-specific. The dataset and the taxonomy of intents should be tailor-made for different scenarios. Casanueva et al. [4] present a dataset containing 13,083 instances over 77 intents in the banking domain. They pre-train the sentence encoder on a conversation response selection task, and show that the proposed model is useful for intent detection. They also experiment with cross-domain intent detection datasets such as CLINCI50 [25] and HWU64 [29] to show the robustness of the proposed method.
- Identity fraud can be viewed as a kind of implicit intent. Wang et al. [37] propose an identity fraud detection framework. Their system asks questions drawn from the original personal knowledge graph, and further detects whether the responder is the correct user based on dialogue interactions. They conduct experiments with a simulated dataset and demonstrate promising pilot results.
- Selecting a proper response to the customer is also important. Wang et al. [38] experiment with debt collection. They select policies based on the dialogue state, and further choose the current state script. Their proposed two-state method outperforms a flow-based method for both single- and multi-round dialogues.

The above studies show that some methods can be used in several domains; intents specific to a certain domain, though, may still necessitate customization. One goal of this research direction is providing automatic customer services. One open issue in the financial domain is how to reply to customers based on their opinions. Given the development of current NLP methods, human-machine cooperation probably remains the most likely method for real-world applications.

Information extracted from various sources can be used to improve the working efficiency of employees. For example, in previous work [27], we proposed FinSense, a system that suggests stocks that are implicitly related to a given news article; a screenshot is shown in Fig. 6.6. When a news article is pasted into the box on the left, FinSense extracts the stocks mentioned in the article and lists them in the middle box. This function is provided for journalists to streamline their job, because they no longer need to provide labels after they complete the article. Nevertheless, they may

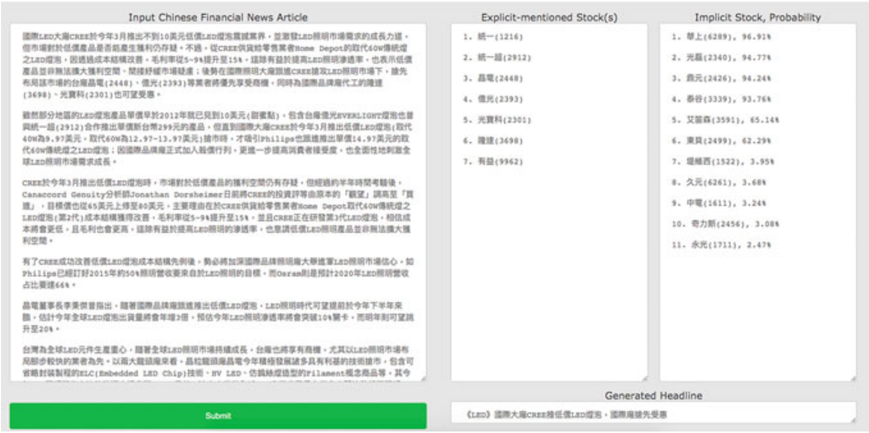

Fig. 6.6 Screenshot of FinSense, a system that suggests stocks that are implicitly related to a news article [27]

need to provide additional labels for implicit stocks, that is, stocks that are related but not explicitly mentioned. FinSense also recommends such stocks, based on the implicit information inference techniques mentioned in Sect. 4.3. Journalists must also compose a headline for the news article. FinSense suggests a headline based on the Transformer model [11]. This figure is thus one example of an application that uses financial opinion mining to improve employee efficiency. Although the recommended tags and headline may need some tweaking, the system does narrow down the journalist's choices.

In the financial domain, we also discuss another type of opinion: customer opinion. We go over scenarios that involve extracting the components of a customer's opinion, and discuss intent detection and dialog generation in customer services as potential applications. As an example, we show how implicit information inference can be used to streamline a journalist's job.

6.4 Summary

In this chapter we describe applications of financial opinion mining. Providing information to the customer is the primary purpose of many financial institutions. We provide a detailed discussion of the workflow of professional analysts, and present selected investment scenarios. In Sect. 6.3, we discuss personalized recommendations and domain-specific features. Various studies show the feasibility of transferring other domains' latent features to the financial field. We show how financial opinion components can be extracted and used to improve employee efficiency. We show relations between customer and investor opinions in the financial domain. Because customer opinions in the financial domain are similar to those in other fields, we believe that they can be leveraged using methods proposed for opinions in other domains. Hence in this book we focus on investor opinions. In the next chapter, we summarize the proposed research directions and show how to apply the results of financial opinion mining research to other domains.

References

1. Atanasova, P., Wright, D., Augenstein, I.: Generating label cohesive and well-formed adversarial claims. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP) (Online, Nov. 2020). Association for Computational Linguistics, pp. 3168–3177 (2020)
2. Bi, Y., Song, L., Yao, M., Wu, Z., Wang, J., Xiao, J.: DCDIR: a deep cross-domain recommendation system for cold start users in insurance domain. In: Proceedings of the Forty-Third International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1661–1664 (2020)
3. Carbonell, J., Goldstein, J.: The use of MMR, diversity-based reranking for reordering documents and producing summaries. In: Proceedings of the Twenty-First Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 335–336 (1998)
4. Casanueva, I., Temčinas, T., Gerz, D., Henderson, M., Vulić, I.: Efficient intent detection with dual sentence encoders. In: Proceedings of the Second Workshop on Natural Language Processing for Conversational AI (Online, July 2020). Association for Computational Linguistics, pp. 38–45 (2020)
5. Chen, C.-C., Huang, H.-H., Chen, H.-H.: Crowd view: Converting investors' opinions into indicators. IJCAI, pp. 6500–6502 (2019)
6. Chen, C.-C., Huang, H.-H., Chen, H.-H.: Next cashtag prediction on social trading platforms with auxiliary tasks. In: 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 525–527. IEEE (2019)
7. Chen, C.-C., Huang, H.-H., Takamura, H., Chen, H.-H.: Numeracy-600K: learning numeracy for detecting exaggerated information in market comments. In: Proceedings of the Fifty-Seventh Annual Meeting of the Association for Computational Linguistics (Florence, Italy, July 2019). Association for Computational Linguistics, pp. 6307–6313 (2019)
8. Chen, C.-C., Huang, H.-H., Tsai, C.-W., Chen, H.-H.: CrowdPT: summarizing crowd opinions as professional analyst. In: The World Wide Web Conference, pp. 3498–3502 (2019)
9. Chen, T., Yin, H., Ye, G., Huang, Z., Wang, Y., Wang, M.: Try this instead: personalized and interpretable substitute recommendation. In: Proceedings of the Forty-Third International ACM SIGIR Conference on Research and Development in Information Retrieval (New York, NY, USA, 2020), SIGIR '20. Association for Computing Machinery, pp. 891-900 (2020)
10. Cheng, D., Yang, F., Wang, X., Zhang, Y., Zhang, L.: Knowledge graph-based event embedding framework for financial quantitative investments. In: Proceedings of the Forty-Third International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2221–2230 (2020)

11. Chu, J., Chen, C.-C., Huang, H.-H., Chen, H.-H.: Learning to generate correct numeric values in news headlines. In: Companion Proceedings of the Web Conference 2020, pp. 17–18 (2020)
12. Da, Z., Huang, X.: Harnessing the wisdom of crowds. Manag. Sci. **66**(5), 1847–1867 (2020)
13. Ding, Q., Wu, S., Sun, H., Guo, J., Guo, J. Hierarchical multi-scale Gaussian transformer for stock movement prediction. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20 (7 2020). International Joint Conferences on Artificial Intelligence Organization, pp. 4640–4646. Special Track on AI in FinTech (2020)
14. Fabbri, A., Li, I., She, T., Li, S., Radev, D.: Multi-news: a large-scale multi-document summarization dataset and abstractive hierarchical model. In: Proceedings of the Fifty-Seventh Annual Meeting of the Association for Computational Linguistics (Florence, Italy, July 2019). Association for Computational Linguistics, pp. 1074–1084 (2019)
15. Gretz, S., Bilu, Y., Cohen-Karlik, E., Slonim, N.: The workweek is the best time to start a family – a study of GPT-2 based claim generation. In: Findings of the Association for Computational Linguistics: EMNLP 2020 (Online, Nov. 2020). Association for Computational Linguistics, pp. 528–544 (2020)
16. Hassan, N., Arslan, F., Li, C., Tremayne, M.: Toward automated fact-checking: detecting check-worthy factual claims by ClaimBuster. In: Proceedings of the Twenty-Third ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1803–1812 (2017)
17. Hidey, C., McKeown, K.: Fixed that for you: generating contrastive claims with semantic edits. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers) (Minneapolis, Minnesota, June 2019). Association for Computational Linguistics, pp. 1756–1767 (2019)
18. Holthausen, R.W., Larcker, D.F.: The prediction of stock returns using financial statement information. J. Account. Econ. **15**(2–3), 373–411 (1992)
19. Hsu, T.-W., Chen, C.-C., Huang, H.-H., Chang Chen, M., Chen, H.-H.: Hedging via opinion-based pair trading strategy. In: Companion Proceedings of the Web Conference 2020, pp. 69–70 (2020)
20. Hu, Z., Liu, W., Bian, J., Liu, X., Liu, T.-Y.: Listening to chaotic whispers: a deep learning framework for news-oriented stock trend prediction. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, pp. 261–269 (2018)
21. Jame, R., Johnston, R., Markov, S., Wolfe, M.C.: The value of crowdsourced earnings forecasts. J. Account. Res. **54**(4), 1077–1110 (2016)
22. Jindal, N., Liu, B.: Review spam detection. In: Proceedings of the Sixteenth International Conference on World Wide Web, pp. 1189–1190 (2007)
23. King, B.: Bank 4.0: Banking Everywhere, Never at a Bank. Wiley, New York (2018)
24. Lahmiri, S.: Wavelet low-and high-frequency components as features for predicting stock prices with backpropagation neural networks. J. King Saud Univ.-Comput. Inf. Sci. **26**(2), 218–227 (2014)
25. Larson, S., Mahendran, A., Peper, J. J., Clarke, C., Lee, A., Hill, P., Kummerfeld, J. K., Leach, K., Laurenzano, M. A., Tang, L., Mars, J.: An evaluation dataset for intent classification and out-of-scope prediction. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the Ninth International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (Hong Kong, China, Nov. 2019). Association for Computational Linguistics, pp. 1311–1316 (2019)
26. Li, W., He, L., Zhuge, H.: Abstractive news summarization based on event semantic link network. In: Proceedings of COLING 2016, the Twenty-Sixth International Conference on Computational Linguistics: Technical Papers (Osaka, Japan, Dec. 2016). The COLING 2016 Organizing Committee, pp. 236–246 (2016)

27. Liou, Y.-T., Chen, C.-C., Tang, T.-H., Huang, H.-H., Chen, H.-H.: FinSense: an assistant system for financial journalists and investors. In: Proceedings of the 14th International Conference on Web Search and Data Mining (2021)
28. Liu, G., Mao, Y., Sun, Q., Huang, H., Gao, W., Li, X., Shen, J., Li, R., Wang, X.: Multi-scale two-way deep neural network for stock trend prediction. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20 (7 2020). International Joint Conferences on Artificial Intelligence Organization, pp. 4555–4561. Special Track on AI in FinTech (2020)
29. Liu, X., Eshghi, A., Swietojanski, P., Rieser, V.: Benchmarking natural language understanding services for building conversational agents. In: Tenth International Workshop on Spoken Dialogue Systems Technology (2019)
30. Loughran, T., McDonald, B.: When is a liability not a liability? textual analysis, dictionaries, and 10-Ks. J. Financ. **66**(1), 35–65 (2011)
31. Ntakaris, A., Magris, M., Kanniainen, J., Gabbouj, M., Iosifidis, A.: Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods. J. Forecast. **37**(8), 852–866 (2018)
32. Qin, Y., Yang, Y.: What you say and how you say it matters: predicting stock volatility using verbal and vocal cues. In: Proceedings of the Fifty-Seventh Annual Meeting of the Association for Computational Linguistics (Florence, Italy, July 2019). Association for Computational Linguistics, pp. 390–401 (2019)
33. Rashkin, H., Choi, E., Jang, J. Y., Volkova, S., Choi, Y.: Truth of varying shades: analyzing language in fake news and political fact-checking. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 2931–2937 (2017)
34. Rayana, S., Akoglu, L.: Collective opinion spam detection: bridging review networks and metadata. In: Proceedings of the Twenty-First ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 985–994 (2015)
35. See, A., Liu, P. J., Manning, C. D.: Get to the point: summarization with pointer-generator networks. In: Proceedings of the Fifty-Fifth Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1073–1083 (2017)
36. Shapira, O., Ronen, H., Adler, M., Amsterdamer, Y., Bar-Ilan, J., Dagan, I.: Interactive abstractive summarization for event news tweets. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations (Copenhagen, Denmark, Sept. 2017). Association for Computational Linguistics, pp. 109–114 (2017)
37. Wang, W., Zhang, J., Li, Q., Zong, C., Li, Z.: Are you for real? Detecting identity fraud via dialogue interactions. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the Ninth International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 1762–1771 (2019)
38. Wang, Z., Liu, J., Cui, H., Jin, C., Yang, M., Wang, Y., Li, X., Mao, R.: Two-stage behavior cloning for spoken dialogue system in debt collection. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20 (7 2020). International Joint Conferences on Artificial Intelligence Organization, pp. 4633–4639. Special Track on AI in FinTech (2020)
39. Wities, R., Shwartz, V., Stanovsky, G., Adler, M., Shapira, O., Upadhyay, S., Roth, D., Martínez-Cámara, E., Gurevych, I., Dagan, I.: A consolidated open knowledge representation for multiple texts. In: Proceedings of the Second Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics, pp. 12–24 (2017)
40. Xu, Y., Cohen, S. B.: Stock movement prediction from tweets and historical prices. In: Proceedings of the Fifty-Sixth Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (Melbourne, Australia, July 2018). Association for Computational Linguistics, pp. 1970–1979 (2018)

41. Yang, D., Chen, J., Yang, Z., Jurafsky, D., Hovy, E.: Let's make your request more persuasive: modeling persuasive strategies via semi-supervised neural nets on crowdfunding platforms. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers) (Minneapolis, Minnesota, June 2019). Association for Computational Linguistics, pp. 3620–3630 (2019)