Risk-aware Regularization for Opinion-based Portfolio Selection

Ting-Wei Hsu, [1] **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
{twhsu,cjchen}@nlg.csie.ntu.edu.tw, hhhuang@nccu.edu.tw, hhchen@ntu.edu.tw

Abstract

Every investor faces the risk-return tradeoff when making investment decisions. Most of the investors construct a portfolio instead of putting all of their wealth on a certain stock. However, most of the previous works in the NLP community focus on predicting the movement of stocks' prices or volatilities, but do not consider the portfolio selection issue. On the other hand, few works consider unstructured data in the financial community when dealing with this issue. This paper introduces a novel opinion-based portfolio selection task, and proposes new objective functions presenting different risk appetites of investors. The empirical studies of the selecting portfolio are also discussed with both Sharpe ratio and volatility metrics.

Introduction

Portfolio selection is an important issue in the financial domain. Given a pool of assets and a risk-level, we aim to select a subset of assets that constitutes a portfolio, providing the largest expected return. Markowitz (1952) proposes a modern portfolio theory that constructs an efficient frontier to show the trade-off between the risk and return. Figure 1 is an example of the efficient frontier, where the x-axis is an annual return. Note that 0.20 stands for 20%. We use the price data in Yahoo Finance [1] to simulate different kinds of portfolios and visualize the results. We simulate 4,000 portfolios (cloud of points) with weights allocated to the stocks to show the boundary's existence. The portfolios on the frontier are the portfolios that provide the largest expected return under the given risk level. In this paper, we adopt this definition to extend the stock movement prediction task to the portfolio selection task. To explore the proposed task, we provide new labels for the existing dataset, StockNet (Xu and Cohen 2018).

Inspired by previous works, we know that adopting the sentiment and opinions from the textual data such as news articles, social media posts, and formal reports can improve the performance of predicting the price movement (Hu et al. 2018; Liu et al. 2018) and price volatility (Qin and Yang 2019; Yang et al. 2020). Although the trading strategy and portfolio selection are also important, few works consider

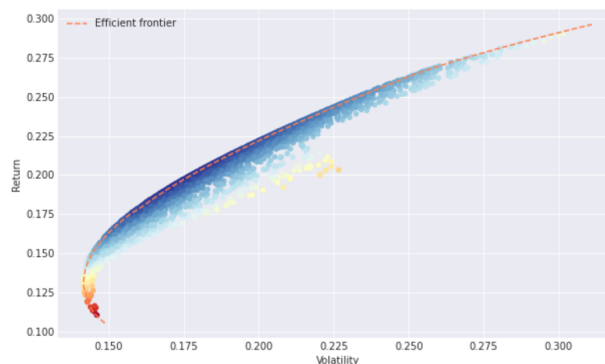[1] https://finance.yahoo.com/

Figure 1: Example of efficient frontier.

such issues with textual information. Recently, Hsu et al. (2020) introduce an opinion-based pair trading strategy by using the textual data from social media platforms and show that hedging strategies perform better than the strategies that only predict the single stock price movement. In this work, we introduce the opinion-based portfolio selection task, and discuss how to leverage the investors' opinions and the market data to select the best portfolio for the investors with different risk appetites.

In order to enable the model to be aware of the risk appetites, we propose new objective functions with novel risk-aware regularization penalties based on market opinions and the trend. Empirical studies confirm the effectiveness of our objective functions. Our contributions are threefold:

1. We introduce a new task and provide additional labels [2] on the publicly-available dataset, StockNet, to increase its value.

2. We demonstrate that NLP models can learn how to select portfolios based on financial social media data.

3. We propose several objective functions that consider both market information and investors' opinions. The empirical studies provide insights for selecting the portfolio for the investors with different risk appetites.

[2] https://github.com/quanthsu/Opinion-based-Portfolio-Selection

Related Work

Portfolio selection (Davis and Norman 1990; Ledoit and Wolf 2017) is a long-term discussed topic in the financial domain. Markowitz (1952) proposes the mean-variance analysis to optimize the portfolio. Black and Litterman (1990) consider investors' views, which are collected from the formatted survey questionnaire when constructing the portfolio. In the AI community, some works (Das, Johnson, and Banerjee 2014; Ding et al. 2018; Zhang et al. 2020) discuss the portfolio selection task. To the best of our knowledge, few previous works take the textual data from social media into consideration when constructing the portfolio. In this work, we explore the portfolio selection task with the investors' opinions on social media platforms.

Pre-trained transformer-based text encoders perform well in many NLP tasks (Vaswani et al. 2017; Raffel et al. 2019; Bai et al. 2020). In this paper, we adopt BERT (Devlin et al. 2019) architecture and discuss the performance under the different settings of input features and loss functions. We find that we need to different features and loss functions need to be adopted when constructing the portfolio for investors with different risk appetites.

Methods

Task Formulation

In this paper, we regard the portfolio selection problem as a task of multi-label classification. Given the information at time t and the risk aversion parameter, the model needs to select a set of target stocks that constitute the best-performing portfolio at time $t+1$. For each time t, we can get an efficient frontier (Markowitz 1952) by calculating all kinds of combinations of the candidate stocks under different risk levels. That is, when the risk aversion parameter (γ) is given, we aim to find the optimization variable (w_t) that maximizes the following utility at each time t.

$$\mu_t^T w_t - \gamma w_t^T \Sigma_t w_t \,, \tag{1}$$

where μ_t denotes the returns of each stock at time t. $w_t \in \mathbf{R}_+^n$ is the optimization variable, and $\mathbf{1}^T w_t = 1$. n is the number of stocks in the candidate set, and Σ_t denotes the corresponding covariance matrix of returns. Note that we set $w_t < 0.1$ to enforce the model on selecting at least ten stocks, and we only label the stocks weighted by the above optimization formula for practical purposes. For example, if the weight of stocks ("AAPL", "FB", "GOOG") given by mean-variance is (0.35, 0, 0.65), the labels for this pool is (1, 0, 1). That means models are asked to predict which stocks should be selected to construct the portfolio. The positive real number γ is the risk aversion parameter. The higher risk aversion parameter means that the investor is more conservative. Risk aversion is a non-negative number. For example, Hupman and Abbas (2014) set the risk aversion from 0 to 2. This parameter denotes the risk appetite of an investor. In this paper, we use 0, 1, 2, and 3 as the risk aversion parameters to represent the risk appetites of investors.

Model

As shown in Figure 2, the inputs of the model are the tweets of all the stocks in the candidate asset pool. BERT (Devlin
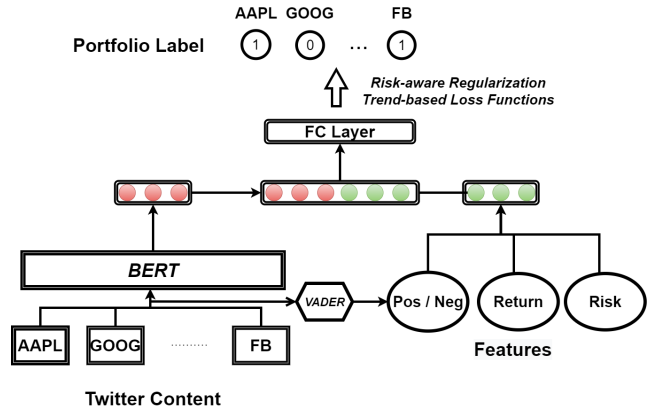

Figure 2: The architecture of the model used in our experiments.

et al. 2019) is adopted to encode the text, and we concatenate the BERT embedding with the proposed features. We further use a fully-connected layer to predict which stock should be selected into the portfolio. In this paper, we focus on discussing the performances of using different features and different loss functions, which are introduced in the following sections.

Features

We use three kinds of features, including the sentiment score of the social media data, averaged returns, and averaged volatility. Valence Aware Dictionary and sEntiment Reasoner (VADER) (Hutto and Gilbert 2014) is adopted to get the positive (Pos) and negative (Neg) sentiment scores of each tweet. We calculate the n-day averaged returns (Return) of each stock and the n-day averaged volatility (Risk) of each stock as the features of market information. The equations are shown below.

$$\text{Return} = \frac{1}{n} \sum_{t=1}^{n} \mu_t \,, \tag{2}$$

$$\mu_t = \frac{P_t - P_{t-1}}{P_{t-1}} \,, \tag{3}$$

$$\text{Risk} = \sqrt{\frac{1}{n-1} \sum_{t=1}^{n} (\mu_t - \text{Return})^2} \,, \tag{4}$$

where P_t is the close price at time t. In this work, n is 5.

Risk-aware Regularization

Different from the price returns that are independent and identically distributed, the volatilities are not. Because of the volatility clustering phenomenon (Mandelbrot 1997), the latest volatility is evidenced useful for revealing the current risk-level. Therefore, adding the risk feature, i.e., volatility, into regularization penalty is expected useful for capturing the risk information. Since the sentiment of investors has been shown related to the stock volatility (Lee, Jiang, and

Indro 2002; Wang et al. 2013), we also add sentiment features into the regularizations. That is to say, to enable the model to be aware of the risk, we add the above features into L1-norm and L2-norm as follows.

$$L1(r) = -\sum_d \sum_i \|r_{i,d} - r_{i,d-1}\| \tag{5}$$

$$L2(r) = -\sum_d \sum_i \|r_{i,d} - r_{i,d-1}\|^2 \tag{6}$$

Here r can be pos, neg and vol denoting the positive sentiment score, negative sentiment score, and volatility of stock i, respectively.

Trend-based Loss Functions

We propose two objective functions, called follow-the-winner (FTW) and follow-the-loser (FTL) (Li and Hoi 2018), for taking the price trend information into account. In FTW, we add punishment to the cases in which the return of the predicted portfolio is far away from the return of the ground truth. In FTL, we believe that the price of the under-performed stock will rise in the future. Therefore, we aim at choosing underperformed stocks. In the following equations, $ret_{i,d}$ denotes the return of stock i at day d.

$$FTW(y, ret) = -\sum_d \sum_i \hat{y}_{i,d} ret_{i,d} - y_{i,d} ret_{i,d} \tag{7}$$

$$FTL(y, ret) = -\sum_d \sum_i \max(0, \hat{y}_{i,d} ret_{i,d} - y_{i,d} ret_{i,d}) \tag{8}$$

Experiments

Dataset

In this paper, we experiment on StockNet (Xu and Cohen 2018), which is collected from Twitter from January 1, 2014 to January 1, 2016. We remove three stocks that do not have enough data from StockNet (88 stocks). Thus, there are 85 stocks in our experiments. The training set contains the data before April 21, 2015. The validation set contains the data from April 21, 2015 to August 10, 2015. The remaining data are included in the test set. The textual data can be downloaded from the site of StockNet[3] (Xu and Cohen 2018). We will release the proposed labels for reproducing the experimental results and supporting future researches.

Experimental Results

Because we formulate the portfolio selection task as a multi-label classification task, we adopt hamming loss as the evaluation metric. Note that lower hamming loss means better performance. Table 1 and Table 2 show the experimental results of different features and different objective functions, respectively. We find that the model with positive sentiment scores achieves the best result. That means the sentiment of social media textual data are useful for the opinion-based portfolio selection task. The model with FTW loss function outperforms the models with other loss functions. This result indicates that the FTW loss function is more suitable for the portfolio selection task than traditional cross-entropy based loss functions.

[3] https://github.com/yumoxu/stocknet-code

	Hamming Loss
Return	0.1968
Risk	0.1955
Pos	**0.1879**
Neg	0.2003

Table 1: Experimental results of different features.

	Hamming Loss
$CE(y)$	0.2182
$CE(y) + L1(vol)$	0.2042
$CE(y) + L1(pos)$	0.1965
$CE(y) + L1(neg)$	0.2063
$CE(y) + L2(vol)$	0.2038
$CE(y) + L2(pos)$	0.1965
$CE(y) + L2(neg)$	0.2115
$FTW(y, ret)$	**0.1918**
$FTL(y, ret)$	0.1943

Table 2: Experimental results of different loss functions.

Backtesting

In this section, we report the backtest results of the data in the test set. We compare the selected portfolios of different models under different risk appetites based on the Sharpe ratio, which simultaneously considers the return and risk. The Sharpe ratio (Sharpe 1994) is calculated by $\frac{R_p - R_f}{\sigma_p}$, where R_p is the return of portfolio, R_f is risk-free rate and σ_p denotes the standard deviation of the returns of the portfolio. Note that we set the risk-free rate as 0 and do not consider the extra cost (e.g., taxes and fees) for backtesting.

Based on the results in Table 2, FTW and FTL perform better than other loss functions. Therefore, in Table 3, we adopt FTW and FTL with different kinds of features. The results in Table 3 show that the models with risk features and FTL loss functions get the highest Sharpe ratio for the investors with lower risk aversion parameters (0 and 1). The results also indicate that the models with FTW loss function get the best Sharpe ratio for the investors with higher risk aversion parameters (2 and 3).

To answer whether the models successfully capture the investors' risk appetites, we compare the volatility of the selected portfolios under different settings. Table 4 shows the results of the settings that get the best performances for different kinds of investors in Table 3. We find that most of the results follow the same trend of the ground truth. That is, this model recommends the portfolio with higher volatility to the investors who have higher risk appetites (lower risk aversion parameter). Based on these results, we infer that the models can capture the risk appetites of investors.

The $1/N$ portfolio strategy (DeMiguel, Garlappi, and Uppal 2009) is the naive strategy, which puts equal weight on all stocks in the candidate pool. It is the usual baseline of the portfolio selection problem. In our experiment, the Sharpe ratio of $1/N$ portfolio strategy is 0.1299, and the volatility of $1/N$ portfolio strategy is 0.0085. That means this strategy takes a lower risk, which leads to a lower return. In contrast,

	0	1	2	3
FTW	0.5172	0.5420	0.5699	0.5427
FTL	0.6693	0.5790	0.5040	0.5115
Return	0.5724	0.6779	0.5504	0.5710
Return + FTW	0.5613	0.5768	0.5771	**0.6547**
Return + FTL	0.6224	0.5782	0.4461	0.5442
Risk	0.6939	0.7265	0.4512	0.5472
Risk + FTW	0.5233	0.5048	0.5219	0.5253
Risk + FTL	**0.7883**	**0.7331**	0.6245	0.5263
Pos	0.5837	0.5503	0.5881	0.4975
Pos + FTW	0.5093	0.4862	**0.6673**	0.5774
Pos + FTL	0.6637	0.5796	0.5073	0.4984
Neg	0.4176	0.5358	0.6159	0.5550
Neg + FTW	0.5383	0.6541	0.5384	0.5868
Neg + FTL	0.6210	0.5946	0.5730	0.5803
GT (Upper Bound)	1.2301	1.2500	1.2061	1.2324

Table 3: Sharpe ratio of different settings under different risk aversion parameter. GT denotes the results of using the ground truth, which can be considered as the upper bound of the performance.

	0	1	2	3
FTW	0.9654	0.9495	0.9173	0.9143
FTL	0.9575	**0.9028**	0.9169	0.9171
Return	0.9444	0.9949	0.9166	0.9157
Return + FTW	1.0040	0.9604	**0.8988**	**0.9135**
Return + FTL	0.9944	0.9431	0.9346	0.9155
Risk	0.9536	0.9327	0.9110	0.9156
Risk + FTW	0.9837	0.9479	0.9206	0.9181
Risk + FTL	1.0146	0.9673	0.9591	0.9148
Pos	0.9494	0.9467	0.9193	0.9169
Pos + FTW	0.9841	0.9094	0.9557	0.9308
Pos + FTL	0.9482	0.9471	0.9286	0.9189
Neg	0.9987	0.9484	0.9256	0.9157
Neg + FTW	**0.9402**	0.9433	0.9240	0.9151
Neg + FTL	0.9938	0.9313	0.9412	0.9169
GT	0.9710	0.9562	0.9313	0.9183

Table 4: Volatility of different settings under different risk aversion parameter. The models are expected to select a portfolio with a lower risk to the investor with a higher risk aversion parameter.

the proposed methods take reasonable risks to earn more, which is more practical in the investment scenario.

In sum, this work introduces a new research direction that leverages social media data for portfolio selection, and have the following findings.

1. Sentiment, especially the positive sentiment, is useful for portfolio selection.

2. FTW and FTL loss functions perform better than cross-entropy loss functions in portfolio selection.

3. The models with the proposed features and loss functions can capture the investors' risk appetites and suggest the portfolio based on the risk appetites.

Conclusion

In this paper, we introduce the portfolio selection task to our community and provide additional labels on the publicly-available dataset. We further show that the opinions from social media platforms are useful for the proposed task. We discuss the results of BERT architecture with different kinds of features and loss functions. Our results indicate that the tailor-made loss function, FTW and FTL perform better than the vanilla cross-entropy loss function in the proposed task. We also find that the models can capture the investors' risk appetites.

In the future, we plan to extend this work by adopting the opinions from different sources such as professional analysts' reports and the transcriptions of companies' earnings calls. We also plan to explore the performances of FTW and FTL loss functions on trading strategy construction. Probing the portfolio selection task with a larger asset pool is also worth discussing. We will enlarge the asset pool with currency, bond, commodity, and other financial instruments to test whether neural network models can select the portfolio near to the efficient frontier with the larger asset pool.

References

Bai, H.; Shi, P.; Lin, J.; Tan, L.; Xiong, K.; Gao, W.; and Li, M. 2020. SegaBERT: Pre-training of Segment-aware BERT for Language Understanding. *arXiv preprint arXiv:2004.14996* .

Black, F.; and Litterman, R. 1990. Asset allocation: combining investor views with market equilibrium. *Goldman Sachs Fixed Income Research* 115.

Das, P.; Johnson, N.; and Banerjee, A. 2014. Online Portfolio Selection with Group Sparsity. In *AAAI*, 1185–1191.

Davis, M. H.; and Norman, A. R. 1990. Portfolio selection with transaction costs. *Mathematics of operations research* 15(4): 676–713.

DeMiguel, V.; Garlappi, L.; and Uppal, R. 2009. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The review of Financial studies* 22(5): 1915–1953.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.

Ding, Y.; Liu, W.; Bian, J.; Zhang, D.; and Liu, T.-Y. 2018. Investor-imitator: A framework for trading knowledge ex-

traction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1310–1319.

Hsu, T.-W.; Chen, C.-C.; Huang, H.-H.; Chang Chen, M.; and Chen, H.-H. 2020. Hedging via Opinion-based Pair Trading Strategy. In *Companion Proceedings of the Web Conference 2020*, 69–70.

Hu, Z.; Liu, W.; Bian, J.; Liu, X.; and Liu, T.-Y. 2018. 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*, WSDM '18, 261–269. 9781450355810.

Hupman, A. C.; and Abbas, A. E. 2014. Optimizing fixed targets in organizations through simulation. In *Proceedings of the Winter Simulation Conference 2014*, 986–995. IEEE.

Hutto, C. J.; and Gilbert, E. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth international AAAI conference on weblogs and social media*.

Ledoit, O.; and Wolf, M. 2017. Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks. *The Review of Financial Studies* 30(12): 4349–4388.

Lee, W. Y.; Jiang, C. X.; and Indro, D. C. 2002. Stock market volatility, excess returns, and the role of investor sentiment. *Journal of banking & Finance* 26(12): 2277–2299.

Li, B.; and Hoi, S. C. H. 2018. *Online portfolio selection: principles and algorithms*. Crc Press.

Liu, Q.; Cheng, X.; Su, S.; and Zhu, S. 2018. Hierarchical Complementary Attention Network for Predicting Stock Price Movements with News. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM '18, 1603–1606. 9781450360142.

Mandelbrot, B. B. 1997. The variation of certain speculative prices. In *Fractals and scaling in finance*, 371–418. Springer.

Markowitz, H. 1952. Portfolio Selection. *The Journal of Finance* 00221082, 15406261.

Qin, Y.; and Yang, Y. 2019. What You Say and How You Say It Matters: Predicting Stock Volatility Using Verbal and Vocal Cues. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 390–401. Florence, Italy: Association for Computational Linguistics.

Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683* .

Sharpe, W. F. 1994. The sharpe ratio. *Journal of portfolio management* 21(1): 49–58.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.

Wang, C.-J.; Tsai, M.-F.; Liu, T.; and Chang, C.-T. 2013. Financial sentiment analysis for risk prediction. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, 802–808.

Xu, Y.; and Cohen, S. B. 2018. Stock Movement Prediction from Tweets and Historical Prices. In *ACL*.

Yang, L.; Ng, T. L. J.; Smyth, B.; and Dong, R. 2020. HTML: Hierarchical Transformer-Based Multi-Task Learning for Volatility Prediction. In *Proceedings of The Web Conference 2020*, WWW '20, 441–451. 9781450370233.

Zhang, Y.; Zhao, P.; Li, B.; Wu, Q.; Huang, J.; and Tan, M. 2020. Cost-sensitive portfolio selection via deep reinforcement learning. *IEEE Transactions on Knowledge and Data Engineering* .