CRYPTOSENTIMENT: A DATASET AND BASELINE FOR SENTIMENT-AWARE DEEP REINFORCEMENT LEARNING FOR FINANCIAL TRADING

Loukia Avramelou[1], *Paraskevi Nousi*[1], *Nikolaos Passalis*[1], *Stavros Doropoulos*[2], *Anastasios Tefas*[1]

[1]Computational Intelligence and Deep Learning Group, AIIA Lab.
Department of Informatics, Aristotle University of Thessaloniki, Greece
[2]DataScouting, Greece
{*avramell, paranous, passalis*}*@csd.auth.gr, doro@datascouting.com, tefas@csd.auth.gr*

ABSTRACT

Deep Learning (DL) models have been applied in several studies to solve financial trading problems. Most approaches handle these problems as classification or reinforcement learning problems with the objective of developing profitable strategies. Recent works have demonstrated that supplying financial trading agents with sentiment information can lead to improved performance. However, most of these works focus on collecting sentiment in a coarse-grain manner, which is not always appropriate for making fine-grained trading decisions, e.g., on a minute basis. In this paper, we introduce a fine-grained cryptocurrency sentiment dataset, called *CryptoSentiment*, which contains 235,907 sentiment scores for 14 cryptocurrency assets, gathered by various online sources. Moreover, we provide Deep Reinforcement Learning (DRL) baselines using the collected dataset, investigating the impact of multi-modal features on cryptocurrency trading.

Index Terms— Sentiment analysis, Financial Trading, Deep Reinforcement Learning.

1. INTRODUCTION

The purpose of financial trading is to achieve profit by purchasing and selling financial assets. Automated financial trading is performed by algorithms without the involvement of humans. The exponential growth of deep learning (DL) models has resulted in a range of models that focus on such financial trading challenges. The problem can be treated as a supervised learning problem, which means that the models predict either price movements or the action that must be performed, i.e., typically buy, sell or exit. In this case, labeled data must be gathered for either regression or classification problems. On the other hand, automated financial trading is also approached by deep reinforcement learning techniques. Contrary to supervised learning methods, in this case, an agent learns to interact with its environment, decides on the best trading strategy and is rewarded according on the amount of profit that accumulates, without the need of labeled data.

In terms of input features, most financial trading models employ just price-related information [1, 2]. This information can be used for performing trading at various horizons, ranging from minutes and hours [3] to days or longer [4]. Additional information, however, can be used as input features. These kinds of information can be derived from internet sources such as article news and social media platforms that represent public opinion. A number of recent works [5, 6] have used sentiment to enrich price-related features and received positive results. Intuitively, the public sentiment regarding trading assets influences both the price and trading strategies concerning said assets, and vice versa. This means that valuable information can be extracted by taking into consideration public statements relating to the financial market and specifically about the assets in one's portfolio.

However, most of these works focus on collecting sentiment in a coarse-grain manner, e.g., calculate the average sentiment across the collected documents for each day. Even though this information is indeed useful for trading using daily candles [4], it might not contain enough information for making fine-grained trading decisions, e.g., on a minute basis. It is expected that collecting fine-grained sentiment will improve the performance of agents working in finer grained horizons. However, reliably estimating the sentiment in such granularities is not trivial, since it requires collecting an enormous amount of information. For example, assuming collecting just 5 documents per minute requires collecting more than 7,000 documents per day.

The key contribution of this work is to provide a publicly available dataset containing fine-grained sentiment analysis data (minute basis) about cryptocurrency market collected by different online sources. Specifically, we provide the *CryptoSentiment* dataset, which includes 235,907 sentiment scores for 14 different cryptocurrencies gathered from various online sources such as news articles and social media, that exceed the 3,000,000 documents. The collected dataset is pub-

This work has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH - CREATE - INNOVATE (project code: T2EDK-02094).

lished in Zenodo [7] and is publicly available at `https://doi.org/10.5281/zenodo.7684410`. We also provide baselines investigating whether using sentiment information is advantageous when training deep reinforcement learning agents for financial trading. In particular, we examine the impact of features from multiple modals, i.e., price and sentiment, in the training of a DRL agent for financial trading.

The rest of the paper has the following structure. In Section 2, we discuss related work and similar approaches. Then, in Section 3, we describe the data collection process and the followed pipeline. After that, in Section 4, we illustrate the results of our experiments and we evaluate the impact of sentiment score used. Finally, in Section 5, we represent our conclusions and we describe future work.

2. RELATED WORK

Financial trading experts and DL researchers have explored the use of DL methods in order to automate financial trading processes with financial assets. This has been done in a number of different research approaches. One common approach is supervised learning [8, 9], with some works focused on regression problems such as forecasting stock prices [10], while others focus on classification problems such as predicting price movements [11]. DRL [12] has also been widely applied in this field, where an agent simulates a financial environment and performs trading transactions. Financial trading agents have produced effective trading strategies that outperform human traders and traditional automated financial trading systems.

Except from price - related data, some research [4] has been done on incorporating additional sources of information in the input space of DRL agents, particularly sentiment analysis of text from news articles or social media, to improve the performance of financial trading models. These works [13, 14, 15] approach the financial trading problem using sentiment information that is extracted from a lexicon-based sentiment analyzers [16] in addition to price data. These surveys show that including sentiment information positively affects agent performance when compared to simply price-related agents.

Since cryptocurrencies have expanded at an exponential growth rate in recent years, cryptocurrency trading has become increasingly popular. This work aims to show how a combination of fine-grained sentiment-related features and price-related data can improve the performance of agents in another financial area, such as cryptocurrency trading. We compared the performance of price-only agents with price and sentiment combined agents and discuss about their performance. Moreover, sentiment data is collected from text about cryptocurrencies using FinBERT [17] - a BERT [18] sentiment extractor focused on financial documents - and provided for further research reasons. Finally, note that we introduce the first - to the best of our knowledge - fine-grained sentiment dataset for financial trading that contains sentiment information at minute level for 14 different cryptocurrencies.

Table 1: Example of cryptocurrencies along with keywords used in the collection of text from online sources

Cryptocurrency	Keywords
XRPUSDT	'Ripple', 'XRP', 'Chris Larsen', 'Jed McCaleb'
BTCUSDT	'Bitcoin', 'BTC', 'XBT', 'Satoshi Nakamoto'
ETHUSDT	'Ethereum', 'ETH', 'Vitalik Buterin'
EOSUSDT	'EOS.IO', 'EOS', 'Dan Larimer'
NEOUSDT	'NEO', 'Da Hongfei', 'Erik Zhang'
XMRUSDT	'Monero', 'XMR', 'Monero Core Team'
XLMUSDT	'Stellar', 'XLM', 'Jed McCaleb'

3. PROPOSED METHOD

3.1. *CryptoSentiment* Dataset

Sentiment data is extracted using text that is collected from various online sources, such as online articles and social media platforms. Text examples have been collected based on keywords related to cryptocurrency, and some of them are shown in the Table 1. For each sample of text we collected, we implemented sentiment analysis using FinBERT [17], a model that is able to predict a 3-label (positive, negative, neutral) sentiment of financial - related content. Sentiment score is determined as standard; in particular, we calculated it as the difference of positive sentiment score minus negative sentiment score. This indicates that positive texts should have an overall score around 1, negative texts should have an overall score close to -1, and neutral texts have a sentiment score around 0.

More specifically, we were able to collect sentiment data from 17-08-2017 to 12-02-2022 aggregated in minute basis. It is worth noting that sentiment data prior to 2020 is quite rare in comparison to sentiment data after 2020, since the popularity of cryptocurrencies has increased dramatically in recent years. In Fig. 1, we can observe the sentiment values from 2020 to 2022. Note the denser sentiment collection period, after approximately March 2021. which leads to more reliable sentiment estimation, compared to the less dense sentiment collection period before March 2021. Note that the collected dataset can be either used as a whole, i.e., including both data collection periods, or by keeping only the more dense and reliable period for training and testing.

3.2. DRL-based Trading Pipeline

First, we gathered two different types of information, namely cryptocurrency price data and sentiment data. The price data was in the form of Open High Low Close (OHLC) price candles, which contain four values: the candle's open price, closing price, high price, and low price. A candle is a type of price

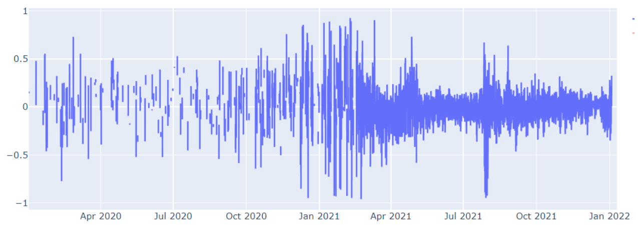

Fig. 1: Mean sentiment value per hour from 2020 to 2022. Note that the average over all 14 cryptocurrencies is reported.

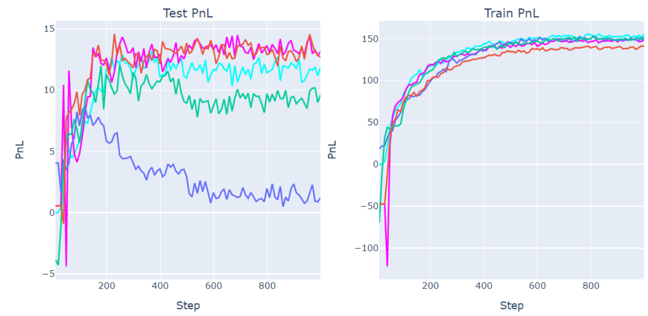

Fig. 2: Test and train PnL during DRL agent training. Features that are used are only price-related. Every different curve represent a different run.

chart that displays the high, low, open, and closing prices of a certain time. The candlestick shows investors whether the closing price was higher or lower than the opening price. Sentiment data is collected as discussed in section 3.1. Thus, we have price and sentiment data at a minute rate. We train DRL agents for financial trading using the previously described price and sentiment data. Price and sentiment data were resampled on a hourly rate. For each hourly time interval, the resampling approach included taking the first value of the open price, the minimum value of the low price, the maximum value of the high price, the final value of the closing price, and the mean of the sentiment values. From this data, we extracted features in order to train the model, such as percentage differences of high, close and low prices, volatility, features that encode time information but also the last-hour sentiment feature.

A DRL agent with Proximal Policy Optimization (PPO) [19] is employed. In further detail, the model used in the studies is a Long Short Term Memory (LSTM) model [20]. Also, dropout is used in order to avoid overfitting. Furthermore, the agent's reward is based on Profit and Loss (PnL) metric, that estimates the expected profit and/or loss of a trading agent over a specific time period. More specifically, PnL is calcualted as:

$$PnL = \sum_{t=1}^{N} \delta_t p_t - |\delta_t - \delta_{t-1}|c, \qquad (1)$$

where N denotes the total duration of the back-testing period (number of time-steps), p_t is the return at time step t, c is the commission paid for realizing profits/losses and δ_t is an index variable used to indicate the current position, which is defined as:

$$\delta_t = \begin{cases} -1, \text{if agent holds a short position at time-step } t \\ 1, \text{if agent holds a long position at time-step } t \\ 0, \text{if the agent is not in the market at time-step } t \end{cases} . \qquad (2)$$

4. EXPERIMENTAL EVALUATION

We train a DRL agent used a LSTM model of 32 neurons followed with a 3 neuron actor and critic. The dropout probability is set to 0.2. Learning rate is initialized as $5 * 10^{-5}$ and commission punishment is set to $2 * 10^{-5}$. The optimizer used was RAdam Optimizer [21]. We used OHLC price data from 14 cryptocurrency-USDT pairs to train the DRL agent. In more detail, *XRPUSDT, BTCUSDT, ETHUSDT, EOSUSDT, ADAUSDT, NEOUSDT, TRXUSDT, XMRUSDT, XLMUSDT, WAVESUSDT, ETCUSDT, VETUSDT, BTCBUSDT, ATOMUSDT* were utilized. We use data before 25-07-2021 for training and data from 25-07-2021 to 12-02-2022 for testing. We evaluated the agent's performance using Profit and Loss (PnL) metric, that estimates the expected profit and/or loss of a trading agent over a specific time period, as introduced in (1).

4.1. Trading using Price

At first, we investigate the agent's behavior in price-only related features to more effectively investigate the effect of sentiment in DRL agent training later. DRL agents can be particularly difficult to train since the training process is frequently unstable. This issue gets worse when training DRL agents for financial trading, where the input time series is particularly noisy. So, we performed 5 randomly initialized experiments with this setup in order to observe agent's behavior. We found that agents learn profitable strategies, although agent behavior is not consistent and test PnL throughout training varies significantly, despite the fact that Train PnL is quite stable, as shown in Fig. 2. The mean test PnL achieved from these experiments is 9.68 ± 4.93 and the mean train PnL achieved is 149.08 ± 5.37.

4.2. Trading using Sentiment

Similarly, we can train a DRL agent using sentiment features from the previous hour as well as time features. We discov-

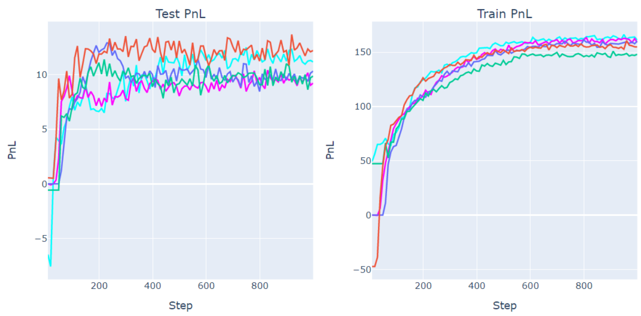

Fig. 3: Test and train Pnl during DRL agent training. Features that are used are price and sentiment features. Every different curve represent a different run.

Fig. 4: Mean test PnL achieved from different experimental setups. The average and standard deviation over 5 different runs are reported.

ered that training a DRL agent with only sentiment information is not as profitable or stable as training a DRL agent using only price features.The mean test PnL reached with these experiments is 6.69 ± 6.43 and mean train PnL is 116.92 ± 20.58. This could be explained since sentiment data are noisy, and training a DRL agent with just this kind of data might lead to the agent learning ineffective strategies.

4.3. Trading using both Price and Sentiment

After that, we investigate the impact of the sentiment features on the agent's performance. We used the previous experiment's price features as well as the previous hour's sentiment. The mean test PnL reached with these experiments is 10.55 ± 1.16 and mean train PnL is 156.68 ± 5.38. The results of these experiments are shown in Fig. 3. We noticed that using sentiment features in addition to price-related features resulted in more profitable agents than price-only agents.

We can observe in Fig 4 and 5, as well as in Table 2 the

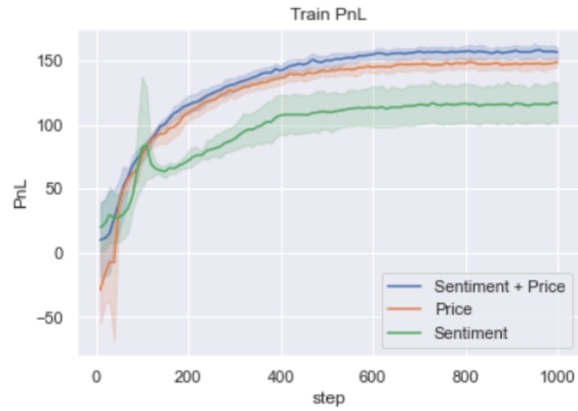

Fig. 5: Mean Train PnL achieved from different experimental setups. The average and standard deviation over 5 different runs are reported.

Table 2: Mean train and test PnL after DRL agent training

	Test PnL	**Train PnL**
Sentiment-only	6.69 ± 6.43	116.92 ± 20.58
Price only	9.68 ± 4.93	149.08 ± 5.37
Price + Sentiment	$\mathbf{10.55 \pm 1.16}$	$\mathbf{156.68 \pm 5.38}$

mean PnL achieved from different DRL agent training setups. We notice that sentiment information in addition to price information has improved the agent's performance. When we use sentiment information in addition to price data, we not only get a better test PnL from the agent, but we also get more stable behavior than price-only related agents.

5. CONCLUSIONS

In this work, we investigated the impact of sentiment in the training of DRL agents for financial trading. We, also, provide a dataset of sentiment values that represent emotion of cryptocurrency - related published text from online sources. We observed that applying sentiment features to a financial trading agent results in more profitable and stable models. The study's findings have inspired our interest in future investigation. First, it is interesting in looking deeply into the significance of sentiment data on financial trading. For example, how additional sentiment features and/or different sentiment values for each cryptocurrency effect agents learning trading strategies. Moreover, it is crucial to investigate the agent's behavior using different types of data, such as forex data with their corresponding sentiment value to determine whether sentiment is beneficial, and other domains other than cryptocurrency trading.

6. REFERENCES

[1] Ahmet Murat Ozbayoglu, Mehmet Ugur Gudelek, and Omer Berat Sezer, "Deep learning for financial applications: A survey," *Applied Soft Computing*, vol. 93, pp. 106384, 2020.

[2] Avraam Tsantekidis, Nikolaos Passalis, Anastasios Tefas, Juho Kanniainen, Moncef Gabbouj, and Alexandros Iosifidis, "Using deep learning to detect price change indications in financial markets," in *EUSIPCO*, 2017, pp. 2511–2515.

[3] Au Vo and Christopher Yost-Bremm, "A high-frequency algorithmic trading strategy for cryptocurrency," *Journal of Computer Information Systems*, vol. 60, no. 6, pp. 555–568, 2020.

[4] Nikolaos Passalis, Loukia Avramelou, Solon Seficha, Avraam Tsantekidis, Stavros Doropoulos, Giorgos Makris, and Anastasios Tefas, "Multisource financial sentiment analysis for detecting bitcoin price change indications using deep learning," *Neural Computing and Applications*, vol. 34, no. 22, pp. 19441–19452, 2022.

[5] Selin Duz Tan and Oktay Tas, "Social media sentiment in international stock returns and trading activity," *Journal of Behavioral Finance*, vol. 22, no. 2, pp. 221–234, 2021.

[6] Robert P Schumaker, Yulei Zhang, Chun-Neng Huang, and Hsinchun Chen, "Evaluating sentiment in financial news articles," *Decision Support Systems*, vol. 53, no. 3, pp. 458–464, 2012.

[7] Loukia Avramelou, Paraskevi Nousi, Nikolaos Passalis, Stavros Doropoulos, and Anastasios Tefas, "CryptoSentiment: A large scale sentiment dataset for cryptocurrencies," 2023, Available at `https://doi.org/10.5281/zenodo.7684410`.

[8] Tej Bahadur Shahi, Ashish Shrestha, Arjun Neupane, and William Guo, "Stock price forecasting with deep learning: A comparative study," *Mathematics*, vol. 8, no. 9, pp. 1441, 2020.

[9] Jingyi Shen and M Omair Shafiq, "Short-term stock market price trend prediction using a comprehensive deep learning system," *Journal of big Data*, vol. 7, no. 1, pp. 1–33, 2020.

[10] Sanjiban Sekhar Roy, Dishant Mittal, Avik Basu, and Ajith Abraham, "Stock market forecasting using lasso linear regression model," in *First International Afro-European Conference for Industrial Advancement*, 2015, pp. 371–381.

[11] Wen Long, Zhichen Lu, and Lingxiao Cui, "Deep learning-based feature engineering for stock price movement prediction," *Knowledge-Based Systems*, vol. 164, pp. 163–173, 2019.

[12] Zihao Zhang, Stefan Zohren, and Stephen Roberts, "Deep reinforcement learning for trading," *The Journal of Financial Data Science*, vol. 2, no. 2, pp. 25–40, 2020.

[13] Johnson D Kinyua, Charles Mutigwe, Daniel J Cushing, and Michael Poggi, "An analysis of the impact of president trump's tweets on the djia and s&p 500 using machine learning and sentiment analysis," *Journal of Behavioral and Experimental Finance*, vol. 29, pp. 100447, 2021.

[14] Prahlad Koratamaddi, Karan Wadhwani, Mridul Gupta, and Sriram G Sanjeevi, "Market sentiment-aware deep reinforcement learning approach for stock portfolio allocation," *Engineering Science and Technology*, vol. 24, no. 4, pp. 848–859, 2021.

[15] Abhishek Nan, Anandh Perumal, and Osmar R Zaiane, "Sentiment and knowledge based algorithmic trading with deep reinforcement learning," in *Database and Expert Systems Applications*, 2022, pp. 167–180.

[16] Clayton Hutto and Eric Gilbert, "Vader: A parsimonious rule-based model for sentiment analysis of social media text," in *AAAI Conference on Web and Social Media*, 2014, vol. 8, pp. 216–225.

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

[18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[19] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[20] Sepp Hochreiter and Jürgen Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[21] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han, "On the variance of the adaptive learning rate and beyond," *arXiv preprint arXiv:1908.03265*, 2019.