FEATURE INTERPRETATION OF CHATGPT FOR FINANCIAL METRICS PREDICTION FROM TEXTUAL DATA

Vitali Chaiko
University of Library Studies and Information Technologies

Abstract: This study evaluates ChatGPT’s ability to predict Amazon’s quarterly revenue using Twitter data, with a focus on analyzing the features that influence forecasting accuracy. The performance and feature utilization of ChatGPT, particularly the use of financial nouns and verbs, are compared to those of a custom-built Long Short-Term Memory (LSTM) model. Despite ChatGPT’s intuitive selection of features, it underperforms, achieving a 67% accuracy rate compared to the LSTM’s 87%. The findings suggest that although ChatGPT is capable of identifying relevant features, it is less effective and reliable than traditional machine learning methods for automated financial predictions.

Keywords: ChatGPT, financial prediction, LSTM, sentiment, Twitter.

Introduction

Precise predictions of financial metrics and up-to-date insights play significant roles in investments, risk management strategies, and strategic planning within the financial services sector. Given the rapid changes occurring in the current economic landscape, machine learning algorithms offer potential solutions for predicting various economic metrics. For instance, these algorithms have been successfully applied in risk management and asset pricing predictions [1]. The latest advancement in the machine learning space has been the development of large language models (LLMs), with ChatGPT emerging as one of the most versatile and transformative tools across various applications.

The research question of this study is to investigate if ChatGPT, as a new machine learning algorithm, can predict financial metrics using timely and relevant textual data and subsequently investigate which contents of the data lead to specific prediction outcomes. Therefore, the predictive potential of ChatGPT is explored in forecasting quarterly revenue changes for the company Amazon during 2015 – 2020, using social media textual data from Twitter timely related to Amazon. The elements of this data, referred to as features, which ChatGPT uses to inform its predictions, are examined. The prediction outcomes and features used by ChatGPT are compared with those from Long Short-Term Memory (LSTM) neural network models developed in this study. The structure of this study is as follows: Initially, existing literature and solutions are reviewed. The methodology then progresses with the preprocessing of large Twitter datasets, followed by the application of prompt engineering methods to ChatGPT to accurately predict revenue changes and extract features. Concurrently, LSTM models are trained to forecast revenue changes, and features from these models are extracted using the Layer Integrated Gradient (LIG) feature interpretation algorithm. Finally, the prediction performance and features of both ChatGPT and LSTM models are compared, and conclusions are drawn.

Literature review

LLMs [2] are self-supervised machine learning algorithms that have been developed and trained on enormous textual datasets, often consisting of hundreds of billions of text tokens. Its training involves predicting the next word based on all previous words. Among the leading LLMs is ChatGPT[1] which can closely mirror human language [3]. Recent studies have examined ChatGPT's ability to predict financial metrics through analysis of economic news and headlines [4, 5, 6]. Moreover, ChatGPT has been described as capable of interpreting text from Bank of England committee members to predict future interest rate decisions [7]. In a survey within the financial sector, the ability of LLMs to perform various NLP tasks, such as text classification and stock price movement prediction, was evaluated [8]. LLMs trained on financial datasets, such as BloombergGPT and FinBERT, achieve accuracy values ranging from 58% to 90%. Another form of self-supervised learning algorithms are LSTM neural networks [9, 10]. They are part of the recurrent neural network family, use past input and output information for current calculations, which is particularly valuable in tasks requiring an understanding of the temporal dynamics of input data. Each layer of an LSTM consists of memory cells with input, forget, and output gates that manage the flow of information through neurons and cells. For the extraction of features that lead to a certain prediction, LIG can used deep learning neural networks. LIG is based on the concept of path integration and aims to attribute the importance of individual neurons in a specific layer of the neural network [11].

Research methodology

Dataset

The dataset used for prediction tasks is a subset of a collection of historical tweets from 2015 to 2020, focusing on top NASDAQ companies. This dataset, produced for a stock market research project [12], is available in two parts. The first part contains tweet data, while the second includes market data for the companies, accessible to other researchers on the public data science and coding platform, Kaggle[2]. The dataset comprises over three million tweets, with each record including unique tweet identification, author, creation date, textual content, related NASDAQ company, and metrics of social engagement. An excerpt of the dataset is shown in Table 1. (This work does not provide the exact tweet contents from the dataset, adhering to Twitter's distribution policy. Instead, the tweet content here is paraphrased to maintain semantic similarity with the original tweets.) From this dataset, tweet records directly related to Amazon were selected for further training and evaluation of the prediction models. This choice was made because this company represents a substantial portion of the data, and its quarterly financial metrics are precisely documented online at Statista[3]. Amazon's quarterly revenue figures are presented in Table 2. All subsequent prediction methods in this study are based on the assumption that the textual content of tweets, when aggregated by specific sizes into groups of tweets, correlates with the increase or decrease in Amazon's revenue by the end of the quarter in which the tweets were created.

Table 1. Excerpt of Twitter dataset

tweet_id	writer	post_date	body
701341609	N2Trends	05/01/2015 11:13:12	$AMZN's refusal to use #SecureWeb on all pages endangers acquisitions & online activity, sparking #privacyconcerns. @Amazon
701341610	FinNews	05/01/2015 12:45:15	Risk to Retailers' Operating Cash Flow: Walmart ($WMT), Amazon ($AMZN),Target ($TGT). For a full overview, check out #MarketTrend.
701341614	SarahPalo	06/01/2015 11:23:01	From Equity Trading Forum – Investor Discussions: What's the market direction for tomorrow? $AMZN $MSFT $LYFT #InvestorInsights
701341615	USTrade	06/01/2015 13:03:01	Not a fan of @MarketMaverick's usual pessimism on @CNBCMarketWatch. Glad I ignored your advice, Mark, and bought $AMZN shares under $350. #shiftfrombearishbullish!
701341618	StockBear	06/01/2015 17:11:34	$AMZN shifts focus to expanding cloud strategy. Interested in becoming an AWS Consulting Partner? $VMW $AMZ #CloudStrategy

Table 2. Excerpt of Amazon revenue in billion US dollar

from_date	to_date	value
10/01/2014 00:00:00	31/12/2014 23:59:59	20.58
01/01/2015 00:00:00	31/03/2015 23:59:59	29.33
04/01/2015 00:00:00	30/06/2015 23:59:59	22.72
07/01/2015 00:00:00	30/09/2015 23:59:59	23.18
10/01/2015 00:00:00	31/12/2015 23:59:59	25.36
01/01/2016 00:00:00	31/03/2016 23:59:59	35.75

ChatGPT-based predictions

This section describes how prediction prompts for ChatGPT were developed to forecast changes in revenue and extract relevant features using the content of group of tweets of size *N*. Predictions were made for *N* of 5, 10, and 20. Features were extracted externally via an instruction prompt, which initially defined the rules of the prediction task. No additional information was provided to ChatGPT before this prompt. The interactions with ChatGPT 4 as of 12.14.2023, for the first two tweet groups of size 5 are documented in Table 3. To determine the optimal tweet group size for further feature evaluation, predictions were evaluated with 100 randomly selected tweet groups from the dataset using accuracy [13] and Matthews Correlation Coefficient (MCC) [14], by calculating this metrics across both prediction labels *increase (1)* and *decrease (0)*. The best performance was achieved with a group size of 10, showing an accuracy of 0.67 and an MCC of 0.29.

Table 3. Interactions with ChatGPT

Input Instruction prompt:
Given a group of tweets about the company Amazon in my next input, separated by '<SEP>', for each group, output the following: (X;w1,w2,…,wn), where X is 1 if the text of this tweet group would indicate an increase in the next Amazon's quarterly revenue, or 0 if it would indicate a decrease and w1 to wn are n contents or words of tweet group which led to your output of X. Provide no further explanation.
ChatGPT output:
Sure, I'm ready. Please provide the group of tweets.
Input prediction prompt:
$AMZN acquisitions & #onlineactivity could be compromised – @Amazon refuses to implement #SecureWeb on all sites #privacyconcerns; $OIL Company Profile Updated Sunday, April 5, 2015 7:15:51 PM $PFE $YUM $AMZN $GLUU; Cash Flow Analysis for Retailers 1. $WMT 2. $AMZN 3. $TGT Complete Details: #MarketTrends; $AMZN Recent News Sunday, April 5, 2015 7:08:11 PM $GILD $BMY $IYE $SLX; $MTD Current Updates Sunday, April 5, 2015 7:02:20 PM $GSK $SPY $AMZN $JNK<SEP> $TRIP Forum Updates Available: Check Them Out: $T $VWO $PFE $AMZN; $QCOM Price Movement Summary for Thursday, April 2, 2015 10:07:16 PM $MSFT $CVX $DWT $AMZN; Not endorsing @MarketMaverick from @CNBCMarketWatch, always negative. Glad I ignored your suggestions, Mark, and invested in $amzn when it was below $350 #shiftfrombearishbullish!; 4/2/15 – US: $DIA down 0.33%, $XLK down 0.50%, $VTI up 0.05% RED: $MSFT $GOOGL $AMZN $PYPL $DIS $NVDA Community Stock Update Thursday, April 2, 2015 11:07:16 PM $INTC $DELL $AMZN <SEP>
ChatGPT prediction output:
(0;#privacyconcerns, compromised, refuses, #SecureWeb) (1;dismissed, invested,below, $350, #shiftfrombearishbullish)

LSTM-based predictions

Custom LSTM models were trained using an Amazon tweets dataset. The text content was preprocessed by removing stop words, hyperlinks, and numbers to enhance clarity. Custom word embeddings [15] were then trained on the dataset using the Gensim Python library [16]. Each LSTM input was represented by a list of word embeddings from tweet groups of size *N*. The output target class was either *decrease (0)* or *increase (1)* in revenue compared to the previous value. The LSTM model was built with PyTorch [17] and trained using ten-fold cross-validation. Training parameters are shown in Table 4, with tweet group sizes of 5, 10, and 20. After training, predictions were made on 100 tweet groups, with features extracted via LIG using the Captum library [18]. Table 5 presents predictions and extracted features for the first two tweet groups of size 5. The models were evaluated using accuracy and Matthews Correlation Coefficient (MCC), with the best results (0.87 accuracy, 0.77 MCC) achieved for a tweet group size of 10.

Table 4. LSTM training parameters

Parameter	**Value**
Embedding layer	torch.nn.Embedding.from_pretrained(word_embeddings, freeze=False)
LSTM Layer	torch.nn.LSTM(emb_size=300, hidden_size=512, num_layers=2, batch_first=True, dropout=0.2)
Linear layer 1	torch.nn.Linear(512, 512)
Linear layer 2	torch.nn.Linear(512, 256)
Linear layer 3	torch.nn.Linear(256, 2)
Activation function	torch.nn.functional.relu
Loss function	torch.nn.functional.cross_entropy
Optimizer	torch.optim.Adam
Batch size	100
Epochs	10

Table 5. LSTM based prediction

LSTM model prediction input:
$AMZN acquisitions & #onlineactivity could be compromised – @Amazon refuses to implement #SecureWeb on all sites #privacyconcerns; $OIL Company Profile Updated Sunday, April 5, 2015 7:15:51 PM $PFE $YUM $AMZN $GLUU; Cash Flow Analysis for Retailers 1. $WMT 2. $AMZN 3. $TGT Complete Details: #MarketTrends; $AMZN Recent News Sunday, April 5, 2015 7:08:11 PM $GILD $BMY $IYE $SLX; $MTD Current Updates Sunday, April 5, 2015 7:02:20 PM $GSK $SPY $AMZN $JNK<SEP> $TRIP Forum Updates Available: Check Them Out: $T $VWO $PFE $AMZN; $QCOM Price Movement Summary for Thursday, April 2, 2015 10:07:16 PM $MSFT $CVX $DWT $AMZN; Not endorsing @MarketMaverick from @CNBCMarketWatch, always negative. Glad I ignored your suggestions, Mark, and invested in $amzn when it was below $350 #shiftfrombearishbullish!; 4/2/15 – US: $DIA down 0.33%, $XLK down 0.50%, $VTI up 0.05% RED: $MSFT $GOOGL $AMZN $PYPL $DIS $NVDA Community Stock Update Thursday, April 2, 2015 11:07:16 PM $INTC $DELL $AMZN $CGC<SEP>
LSTM prediction output with most important LIG features:
(0; $JNK, compromised, acquisitions, refuses, $VMW) (1; $DELL, $CGC, movement, invested, endorsed)

Results

Evaluation of features

The features responsible for the predictive performance of ChatGPT and LSTM models were evaluated using multiple methods. Initially, for the 100 tweet groups used for predictions, sentiment scores were calculated using the Vader [19]. Sentiment analysis [20] gauges the emotional undertone of the data, categorizing it into positive, negative, and neutral classes. **Error! Reference source not found.** shows sentiment scores corresponding to each prediction label, revealing a modest trend of more positive sentiment in the *increase* label (27 instances) and negative sentiment in the *decrease* label (21 instances). Using Vader and the SpaCy [21], features extracted from the tweet groups were automatically counted and categorized based on sentiment and syntactic part-of-speech (POS) tags, such as nouns or verbs. The frequency of POS tags across both *increase* and *decrease* prediction labels was analyzed. Sentiment classes for each feature within the respective POS tags were quantified, represented by colors: green for positive, red for negative, and gray for neutral. Results are presented in Figures Figure and Figure . Additionally, the total number of POS tags and their associated sentiments within the tweet groups were calculated and are shown in **Error! Reference source not found.**. Both methods predominantly identify proper nouns and nouns as the most informative features for making predictions, with LSTM using significantly more proper nouns than ChatGPT. Proper nouns generally exhibit neutrality, whereas nouns show a positive trend for the *increase* label and vice versa for the *decrease* label. The higher incidence of proper nouns and nouns suggests mentions of specific companies or products are crucial indicators of financial performance. Verbs are the second most important features for predictions, with sentiment trends for each label even more present in features identified by ChatGPT. Adjectives, though the least represented words in the tweet groups, are considered the third most important features. Notably, adverbs, despite being more numerous than adjectives, are not significant for either method. The analysis of extracted POS tags and their sentiments shows that ChatGPT's selection of prediction-relevant information behaves similarly to LSTM. However, a comparison of the exact matching content of features within each POS tag between ChatGPT and LSTM, as illustrated in Figure , showed ChatGPT often selects different features for each label. The highest match, at 47%, is observed in verbs for the *decrease* label, with significantly lower matches for nouns and adjectives. Example features and their sentiment of all POS tags for each prediction label are detailed in Table 6. These example features were selected based on their Captum score for LSTM and their frequency for ChatGPT. ChatGPT's features often align more closely with human intuition, typically exhibiting clear positive or negative sentiments, such as 'winner' or 'risk.' Conversely, LSTM's features tend to be less intuitively human, often incorporating Twitter company tags like '$GPRO' and more features with a neutral sentiment such as 'million' or 'social'.

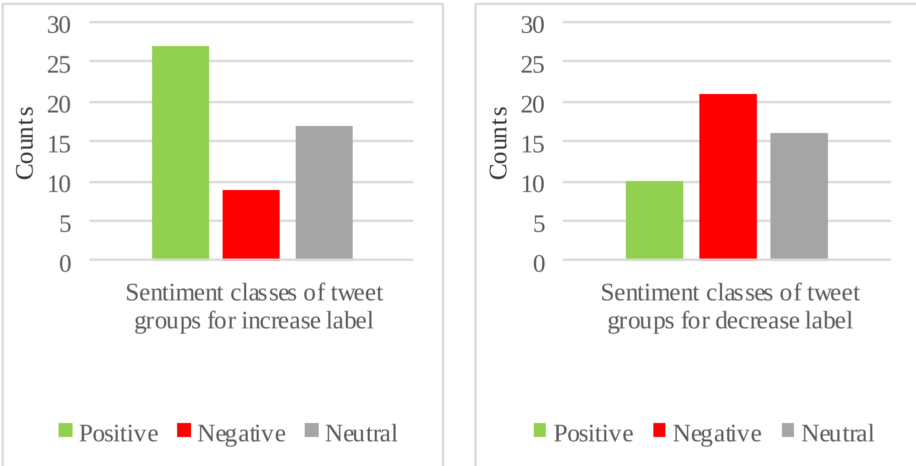

Figure 1. Sentiment of tweet groups

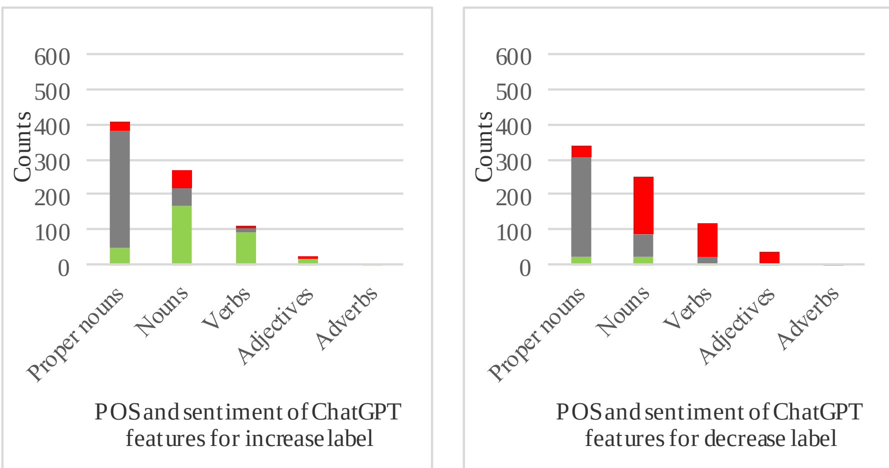

Figure 2. POS and sentiment of ChatGPT features

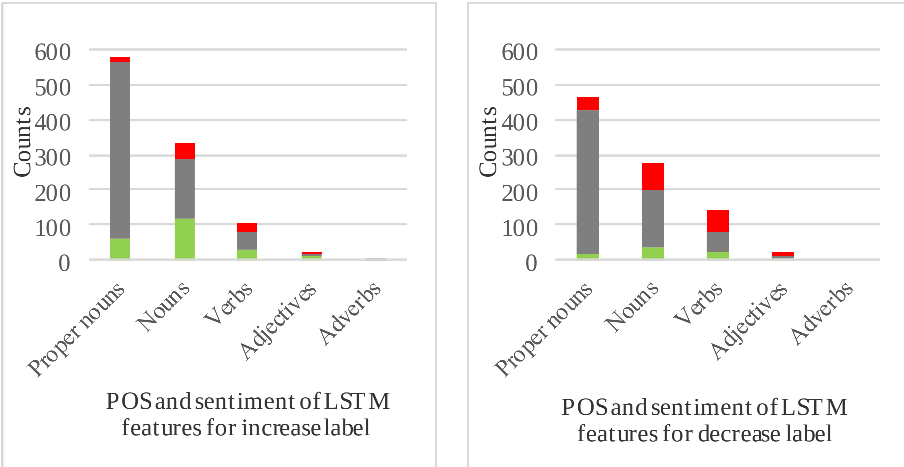

Figure 3. POS and sentiment of LSTM features

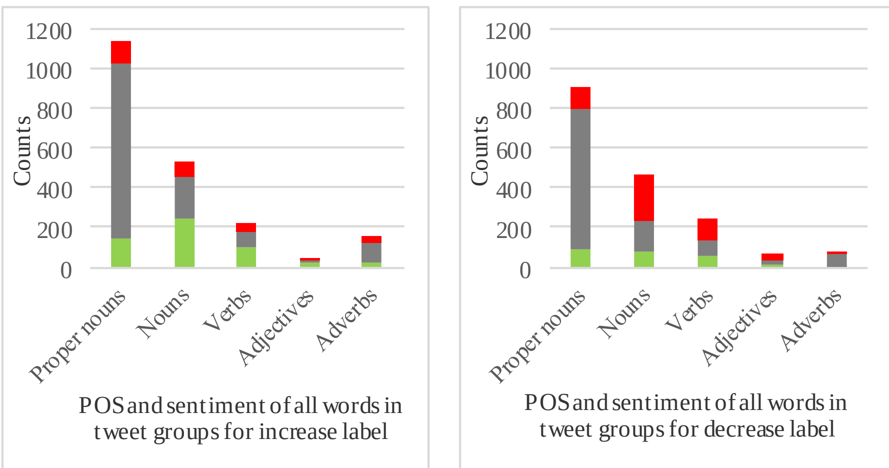

Figure 4. POS and sentiment of all words in tweet groups

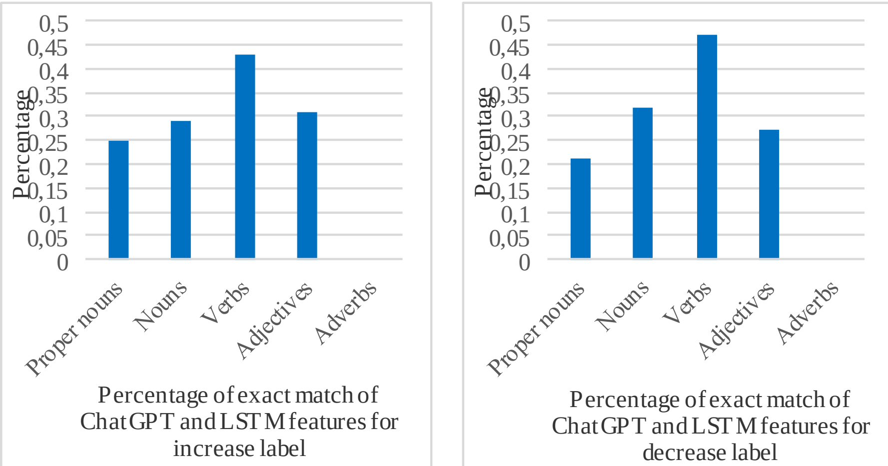

Figure 5. Percentage of exact match of features

Table 6. Example features for each label

	LSTM		**ChatGPT**	
POS tag	**Increase Label**	**Decrease Label**	**Increase Label**	**Decrease Label**
Proper Noun	Kindle, Bezos, $GCI, $GPRO	FireTV, $YHOO, Prime, Linux	Third-party, Bezos, Golden globe, Amazon cloud	DashButton, instant video, IoT, AWS
Noun	leader, strength, battle, million	drone, price, placement, asset	winner, success, marketplace, record	seller, risk, squeeze, correction
Verb	invested, buy, solve, lead	resell, miss, shift, plumber	win, rise, improve, hold	failed, sell, suing, forced
Adverb	wow	neither	again, certainly	badly, never
Adjectiv	big, correct, recent, free	crude, social, greedy, heavy	rising, positive, bullish, huge	poor, sour, fake, worst

Conclusions/discussion

The prediction performance of ChatGPT in predicting changes in financial figures from social media data is significantly lower than that of the specialized LSTM model, with 67% accuracy compared to 87%. By evaluating the features responsible for prediction results, it becomes clear that ChatGPT and LSTM models share similarities in terms of syntactical POS tags; primarily, proper nouns, nouns, and verbs are most prevalent in tweets and carry significant predictive information for financial predictions. However, a deeper analysis of the POS tags reveals that ChatGPT utilizes distinctly different features within each tag. The words ChatGPT selects as features appear intuitive from a human perspective, as it employs financial terms with appropriate sentiments to predict corresponding changes in financial figures. Conversely, LSTM features often present ambiguous sentiments, exemplified by terms like 'social' and 'price.' Despite having less intuitive features, LSTM performs significantly better than ChatGPT, suggesting that its features may possess hidden meanings not captured by ChatGPT's current prompts. ChatGPT may be overly human-like in its assumptions, quickly relying on apparent information. To answer the research question of this study, ChatGPT cannot be seen as a reliable tool for financial predictions from social media data at the moment. However, its choice of prediction-relevant features is intuitive to humans, and it shares similarities with the specialized LSTM in the choice of prediction-relevant features regarding sentiment and POS tags, but not in the exact contents of features. Enhancing prompt design or providing a broad set of accurate tweet samples for pre-training could improve ChatGPT's ability to recognize more subtle relationships, potentially enhancing its prediction accuracy. However, it is important to not overly complicate the prompts, as LSTM models already deliver quick and reliable predictions. Experimenting with different textual datasets – including additional social media platforms, news articles, and company-specific forums – and employing a variety of financial metrics such as stock price movements, trading volumes, EPS, and market capitalization might also help validate the hypothesis that ChatGPT lacks the reliable predictive capabilities of specialized machine learning algorithms. This hypothesis is supported by the fluctuating prediction performance of the models discussed in the literature review. It is important to note that features were extracted from ChatGPT using prompt instructions. Access to ChatGPT's internal architecture and the application of techniques like LIG could yield different insights. However, due to the vast size of ChatGPT's model, LIG might be impractical or excessively time-consuming, and the internal architecture is not publicly accessible.

Conclusion

This study predicted changes in Amazon’s quarterly revenue from social media textual data using ChatGPT and specialized LSTM models. While ChatGPT achieved a prediction accuracy of 67%, it was significantly outperformed by the LSTM model, which achieved 87% accuracy. The analysis further explored the features that influenced predictions, with both ChatGPT and LSTM focusing on similar syntactical POS-tags, yet differing significantly in the sentiment and specific contents within these tags. Proper nouns, nouns, and verbs contained the most predictive information. ChatGPT's features were more intuitively human, often reflecting a clear financial sentiment, either positive or negative. Despite its capabilities, ChatGPT cannot be considered a reliable tool for automated financial predictions, as it performs differently and worse than traditional machine learning algorithms. Enhancing ChatGPT's use of features or hidden contexts similar to those used by LSTM by experimenting with more complex prompts could potentially improve its prediction performance. Additionally, further experiments involving more diverse datasets and various financial metrics could also enhance ChatGPT's effectiveness.

Notes

1. OpenAI, "ChatGPT" 2024. [Online]. Available: https://openai.com/chatgpt/. [Accessed 01.06.2024].
2. Kaggle, "Tweets about the Top Companies from 2015 to 2020," 2023. [Online]. Available: https://www.kaggle.com/datasets/omermetinn/tweets-about-the-topcompanies-from-2015-to-2020. [Accessed 20.01.2024].
3. Statista, "Net sales revenue of Amazon from 1st quarter 2007 to 4th quarter 2023," 2024. [Online]. Available: https://www.statista.com/statistics/273963/quarterly-revenue-ofamazoncom/. [Accessed 12.02.2024].

References

1. **Kelly**, B. and D. **Xiu**. Financial Machine Learning. Foundations and Trends, Finance, vol. 13, no. 3-4, 2023, pp. 205 – 363.
2. **Bender**, E. M., T. **Gebru**, A. McMillan-Major and S. Shmitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big" in 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event Canada, 2021.
3. **Bahrini**, A, M. **Khamoshifar**, H. **Abbasimehr**, R. **Riggs**, M. **Esmaeili**, R. **Majdabadkohne** and M. **Pasehvar**. ChatGPT: Applications, Opportunities, and Threats. – In IEEE Systems and Information Engineering Design Symposium, Charlottesville, 2023.
4. **Lopez-Lira**, A and Y. **Tang**, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models", arXiv, no. arXiv:2304.07619, 2023.
5. **Zhang**, H., F. **Hua**, C. **Xu**, J. **Guo**, H. **Kong** and R. & **Zuo**. "Unveiling the Potential of Sentiment: Can Large Language Models Predict Chinese Stock Price Movements?", arXiv, no. arXiv:2306.14222 , 2023.
6. **Ding**, Y., S. **Jia**, T. **Ma**, B. **Mao**, X. **Zhou**, L. **Li** and D. & **Han**. "Integrating Stock Features and Global Information via Large Language Models for Enhanced Stock Return Prediction", no. arXiv:2310.05627, 2023.
7. **Woodhouse**, D. and A. **Charlesworth**, "Can ChatGPT Predict Future Interest Rate Decisions?", SSRN Electronic journal, no. 4572831, 2023.

8. **Lee**, J., N. **Stevens**, S. C. **Han** and M. **Song**, “A Survey of Large Language Models in Finance”, arXiv, no. arXiv:2402.02315, 2024.
9. **Brownlee**, J. Long Short-Term Memory Networks With Python, 1.5 ed., San Juan: Machine Learning Mastery, 2018, pp. 2 – 17.
10. **ochreiter**, S. and J. **Schmidhuber**, “Long Short-Term Memory”., Neural Computation, No. 9 (8), pp. 1735 – 1780, 1997.
11. **Sundararajan**, M., A. Taly and Q. Yan, “Axiomatic Attribution for Deep Networks” in Proceedings of the 34th International Conference on Machine Learning 70, Sydney, 2017.
12. **Doğan**, M., Ö. **Metin**, S. E. **Tek**, Yumuşak and K. Öztoprak, “Speculator and Influencer Evaluation in Stock Market by Using Social Media”, in 020 IEEE International Conference on Big Data (Big Data), Altanta, 2020.
13. **Foody**, G. M. “Challenges in the real world use of classification accuracy metrics: From recall and precision to the Matthews correlation coefficient”, PLoS ONE, vol. 18, no. 10, 2023.
14. **Chicco**, D. and G. **Jurman**. “The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation”, BMC Genomics, no. 21(1), p. 6, 2020.
15. **Mikolov**, T., K. **Chen**, G. Corrado and J. & Dean, “Efficient Estimation of Word Representations in Vector Space”, arXiv, no. arXiv:1301.3781, 2013.
16. **Rehurek**, R. and P. **Sojka**. “Software framework for topic modelling with large corpora”, in Proceedings of the LREC 2010 workshop on new challenges for NLP frameworks, Valletta, 2010.
17. **Paszke**, A., S. **Gross**, F. **Massa**, A. **Lerer**, J. **Bradbury**, G. **Chanan**, T. **Killeen**, Z. **Lin**, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai and S. Chintala, “PyTorch: An Imperative Style, High-Performance Deep Learning Library”, arXiv, no. arXiv:1912.01703, 2019.
18. **Kokhlikyan**, N., V. **Miglani**, M. Martin, E. Wang, B. Alsallakh, J. Reynolds, A. Melnikov, N. Kliushkina, C. Araya, S. Yan and O. Reblitz-Richardson, “Captum: A unified and generic model interpretability library for PyTorch”, arXiv, no. arXiv:2009.07896 , 2020.
19. **Hutto**, C. and E. **Gilbert**, “VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text”, in Proceedings of the International AAAI Conference on Web and Social Media, Ann Arbor, 2014.
20. **Liu**, B. “Sentiment Analysis and Opinion Mining”, Synthesis Lectures on Human Language Technologies 5(1), pp. 1-167, 2012.
21. **Honnibal**, M. and I. **Montani**, “spaCy: Industrial-strength Natural Language Processing in Python”, 2015.

About the Author

Vitali Chaiko is a doctoral student at the University of Library Studies and Information Technologies in Sofia and also works as a software engineer. His professional interests are data science, machine learning and clean code development.

To contact the Author: chaikovitali.cv@gmail.com.