CryptoPulse: Short-Term Cryptocurrency Forecasting with Dual-Prediction and Cross-Correlated Market Indicators

Amit Kumar
College of Engineering and Computer Science
Texas A&M University-Corpus Christi
akumar3@islander.tamucc.edu

Taoran Ji
College of Engineering and Computer Science
Texas A&M University-Corpus Christi
taoran.ji@tamucc.edu

2024 IEEE International Conference on Big Data (BigData) | 979-8-3503-6248-0/24/$31.00 ©2024 IEEE | DOI: 10.1109/BIGDATA62323.2024.10982029

***Abstract*—Cryptocurrencies fluctuate in markets with high price volatility, which becomes a great challenge for investors. To aid investors in making informed decisions, systems predicting cryptocurrency market movements have been developed, commonly framed as feature-driven regression problems that focus solely on historical patterns favored by domain experts. However, these methods overlook three critical factors that significantly influence the cryptocurrency market dynamics: 1) the macro investing environment, reflected in major cryptocurrency fluctuations, which can affect investors collaborative behaviors, 2) overall market sentiment, heavily influenced by news, which impacts investors strategies, and 3) technical indicators, which offer insights into overbought or oversold conditions, momentum, and market trends are often ignored despite their relevance in shaping short-term price movements. In this paper, we propose a dual prediction mechanism that enables the model to forecast the next day's closing price by incorporating macroeconomic fluctuations, technical indicators, and individual cryptocurrency price changes. Furthermore, we introduce a novel refinement mechanism that enhances the prediction through market sentiment-based rescaling and fusion. In experiments, the proposed model achieves state-of-the-art performance (SOTA), consistently outperforming ten comparison methods in most cases.**

***Index Terms*—Cryptocurrency Prediction, Large Language Model, Market Sentiment Analysis, Predictive Analytics**

I. Introduction

Cryptocurrencies have recently become a topic of conversation due to their significant impact on the financial world, driven by sudden drops and shocks [1], high return opportunity, and the innovative blockchain technology [2], [3] behind them. Unlike traditional financial markets such as bonds and stocks, the cryptocurrency market is characterized by a comparatively smaller market capitalization and pronounced volatility in short-term fluctuations [4]. On one hand, a large proportion of cryptocurrency investors seek short-term investments to exploit opportunities for rapid and substantial returns [5], thereby intensifying market volatility. On the other hand, given this context, these investors tend to be highly sensitive to market-influencing events reported in news [6], such as regulatory actions and fraud events, with their often exaggerated reactions further fueling market fluctuations. Regardless, cryptocurrency is increasingly recognized as a viable alternative investment avenue by those with higher risk tolerances or an interest in short-term, high-yield opportunities [7]. Therefore, the ability to accurately predict short-term cryptocurrency prices not only holds significant practical importance but also contributes integrally to understanding the dynamics of the financial markets as a whole.

Many studies have employed machine learning techniques like SVM and Random Forests [8] to forecast major cryptocurrency returns based on historical price data. However, these methods often exhibit unstable performance across different timescales and cryptocurrencies [8] due to their inability to capture complex, rapidly changing market dynamics. To address this, recent research has focused on deep learning models such as LSTM, bi-LSTM, GRU [9]–[11], and CNN-LSTM [12]. Yet, these studies are confined to the top cryptos by market capitalization, overlooking those with different behaviors and lower liquidity. Furthermore, they primarily rely on historical price data without incorporating technical indicators and sentiment analysis, potentially missing the influence of overbought or oversold market conditions, market sentiment shifts, and external news events on price volatility.

More recently, researchers have integrated market sentiment by analyzing news data alongside historical price data to predict cryptocurrency prices, specifically focusing on Bitcoin and Ethereum [13], [14]. NLP techniques categorize news sentiment, which is then combined with historical price data and fed into deep learning models like LSTM to predict future prices [15]. However, such studies are rare and typically limited to specific cryptocurrencies because they rely on manually labeling sentiment data, a labor-intensive process that doesn't scale well for real-time predictions across multiple cryptocurrencies [15]. Moreover, using investors' expectations from news alone as a trading strategy has been found inadequate, as concluded by Brown and Cliff [16].

To tackle these challenges, we introduce "CryptoPulse," a novel framework for forecasting next-day closing prices by leveraging three primary factors: 1) broad market sentiment from real-time news, 2) complex price dynamics from historical data and technical indicators, and 3) the macro investing environment indicated by fluctuations in major cryptocurrencies. In particular, the key contributions and highlights of this paper are summarized as follows:

- Formulated a novel framework for next-day cryptocurrency forecasting, leveraging short-term observations of key market indicators including market sentiment, macro investing environment, technical indicators, and inherent pricing dynamics.
- Devised a prompting strategy using few-shot learning and consistency-based calibration for effective LLM-based market sentiment analysis of cryptocurrency news.
- Developed a dual-prediction mechanism that separately forecasts prices based on macro conditions and cryptocurrency dynamics, then fuses them using a market sentiment-driven strategy for enhanced accuracy.
- Validated our model on a large-scale real-world dataset, demonstrating effectiveness against ten comparison methods. This dataset, sourced from Yahoo Finance[1] and Cointelegraph[2], along with the source code, will be publicly available upon acceptance.

II. Problem Formulation

Let $\mathcal{C} = \{\mathbf{c}_i\}_{i=1}^{N}$ denote the historical price data for N cryptocurrencies. For the i-th cryptocurrency, the $\mathbf{c}_i = \{\mathbf{f}_t\}_{t=1}^{T}$ is a sequence of feature vectors $\mathbf{f}_t \in \mathbb{R}^5$, where $\mathbf{f}_t$ includes *opening*, *closing*, *high*, *low* prices, *trading volume* along with the technical indicators such as *stochastic %k*, *stochastic %d*, *momentum*, *williams %r*, *a/d oscillator*, *disparity 7* and *rate of change* of the i-th cryptocurrency on day t. Cryptocurrency news from Cointelegraph is also collected, denoted as $\mathcal{D} = \{\mathbf{d}_t\}_{t=1}^{T}$, where $\mathbf{d}_t = \{d_j\}_{j=1}^{|\mathbf{d}_t|}$ represents the set of articles for day t. Our objective is to predict the next day's closing price of a target cryptocurrency using the past L days of historical market prices, technical indicators and news, formulated as:

$$\hat{p}_{t+1}^{i} = g\left(\mathcal{C}_{t-L+1:t}, \mathcal{D}_{t-L+1:t}\right), \tag{1}$$

where $\hat{p}_{t+1}^{i}$ is the predicted closing price for cryptocurrency i on day $t+1$, and g denotes our proposed model. This prediction is crucial for medium-frequency cryptocurrency trading strategies [17], [18].

III. Methodology

We present our proposed model, CryptoPulse, comprising three main components: 1) macro market environment-based next-day fluctuation prediction, 2) price dynamics-based fluctuation prediction, and 3) market sentiment-based dual-prediction rescaling and fusion. A key preprocessing step calculates technical indicators for each trading day using past price data to capture essential market patterns. An overview of the model is shown in Figure 1.

A. Technical Indicator-Based Preprocessing

The seven technical indicators commonly used in market movement prediction: Stochastic %K, Stochastic %D, Williams %R (W%R), Accumulation/Distribution Oscillator

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

[2]https://cointelegraph.com/

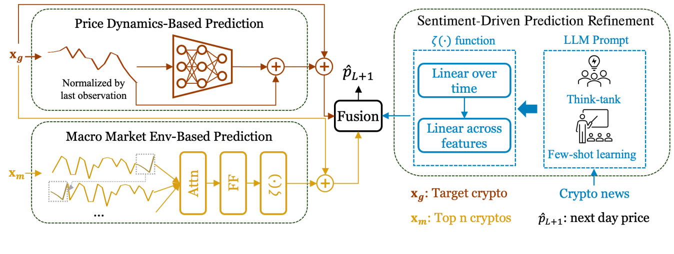

Fig. 1. Overview of CryptoPulse architecture for next-day closing price prediction.

(A/D Osc.), Momentum, Disparity 7 (D7), and Rate of Change (ROC), are implemented as described in [19]–[21] to preprocess the data. These indicators help capture various market dynamics, such as momentum, price trends, and overbought/oversold conditions, providing valuable features for predicting market movements.

B. Macro Market Environment-Based Fluctuation Prediction

The overall macro market environment (e.g., gold and dollar value, policy, public attention) significantly influences cryptocurrency price volatility [22]. However, quantifying this environment is challenging, and most studies [23], [24] focus on specific indicators for individual cryptocurrencies. To address this, we leverage the collective behavior of the top n cryptocurrencies as a proxy for understanding the macro market's influence.

Let $\mathbf{x}_g \in \mathbb{R}^{L\times 5}$ represent a length-L series of observations from the target cryptocurrency, extracted from $\mathbf{c}_i$, using only five direct market features: opening price, closing price, high, low, and trading volume, without technical indicators. Similarly, let $\mathbf{x}_m \in \mathbb{R}^{n\times L\times 5}$ denote a series of the same length L from the top n cryptocurrencies by market capitalization. Both series are processed using a 1D convolutional layer for value and positional embeddings and a sinusoidal positional encoding layer [25]. The embedded observations are represented as $\mathbf{x}_g^{\text{emb}} \in \mathbb{R}^{L\times d_m}$ and $\mathbf{x}_m^{\text{emb}} \in \mathbb{R}^{L\times d_m}$.

Next, we seek to modulate the correlation and interaction between price fluctuation patterns embedded in the target cryptocurrency information $\mathbf{x}_g^{\text{emb}}$ and the macro environment represented by $\mathbf{x}_m^{\text{emb}}$. We formulate this task as *directing the model to learn which sub-series of market behaviors from the top n cryptocurrencies can be aggregated to most effectively approximate the macro investing environment*:

$$\mathbf{h}_m = \sum_{\tau} a_\tau \mathbf{r}_\tau, \mathbf{r}_\tau = \text{roll}\left(\mathbf{x}_m^{\text{emb}}, \tau\right), \tag{2}$$

where $\mathbf{h}_m \in \mathbb{R}^{L\times d_m}$ represents the learned representation of the macro investing environment, and the function $\text{roll}(\cdot, \tau)$ cyclically shifts the input tensor along the temporal dimension by τ steps. The attention weight a_τ for each sub-series is calculated by using the target cryptocurrency $\mathbf{x}_g^{\text{emb}}$ as the query, while all possible shifts of $\mathbf{x}_m^{\text{emb}}$ serve as both keys and values:

$$a_\tau = \text{Softmax}\left(\text{attn}(\mathbf{x}_g^{\text{emb}}, \mathbf{r}_1), \ldots, \text{attn}(\mathbf{x}_g^{\text{emb}}, \mathbf{r}_{L-1}).\right) \tag{3}$$

Technically, the attention function $\text{attn}(\cdot, \cdot)$ can be any time series similarity function. In our experiments, we utilize the period-based similarity calculation method, as introduced in the paper [26].

At last, we use the learned macro investing tensor $\mathbf{h}_m$ to directly predict the next day's closing price fluctuation of the target cryptocurrency $\Delta^{i,1}_{L+1}$. Specifically, $\mathbf{h}_m$ goes through a position-wise feed-forward layer [25], followed by two separate linear layers along the temporal and feature dimensions. Since these linear layers are used multiple times in this paper, we will refer to this process as $\zeta(\cdot)$ in the subsequent sections. The estimated fluctuation is then employed to generate the first prediction for the next-day price: $\hat{p}^{i,1}_{L+1} = p^i_L + \kappa\Delta^{i,1}_{L+1}$, where κ is a scaling factor whose calculation is detailed in Section III-D. In our experiments, we use the top 5 cryptocurrencies to approximate the macro environment.

C. Price Dynamics-Based Fluctuation Prediction

Predicting the next day's closing price $\mathbf{x}_g$ using historical observations and technical indicators is a multivariate-to-univariate time series forecasting task. However, we observed that allowing the model to directly predict the next day's price results in poor projections due to the extreme volatility of cryptocurrencies, leading to overly drastic predictions. To address this, we first predict the next days fluctuation and reconstruct the price using the previous day's closing price:

$$\hat{p}^{i,2}_{L+1} = p^i_L + \kappa\Delta^{i,2}_{L+1}, \Delta^{i,2}_{L+1} = f(\mathbf{x}_g), \tag{4}$$

where $\hat{p}^{i,2}_{L+1}$ is the predicted price, $\Delta^{i,2}_{L+1}$ is the predicted fluctuation, p^i_L is the last observed price, κ is a scaling factor (Section III-D), and f is the prediction model. We observed that Transformer and linear layers yield comparable results, consistent with [27]. For computational efficiency, we adapted the NLinear structure [27] to forecast $\Delta^{i,2}_{L+1}$ by applying a linear layer along the timeline on $\mathbf{x}_g$, normalized by the last-day closing price.

D. Market Sentiment-Guided Dual-Prediction Rescaling and Fusion

As mentioned in Section I, news media significantly influences fluctuations in cryptocurrency markets [6], [28], [29]. However, incorporating this factor into prediction models is challenging because traditional sentiment analysis models [15], [30] often rely on datasets manually annotated for specific scenarios, which are not scalable for real-time analysis in the dynamically changing cryptocurrency market. The recent advancements in LLMs offer an alternative approach for sentiment analysis without requiring extensive fine-tuning on annotated datasets. Nevertheless, designing an effective prompting strategy is crucial for analyzing cryptocurrency news, as recent studies [31] have found that prompt patterns significantly influence the responses of LLMs across various tasks.

In this paper, we combined a "think-tank discussion"-like prompt pattern with the few-shot learning technique to simulate a situation where a group of cryptocurrency traders collaboratively determines the market's reaction to specific news. Recent work suggests that few-shot learning can enhance accuracy and reliability [32], [33]. However, we found that few-shot learning alone is insufficient. Firstly, the LLM's responses are unstable and sometimes yield different outcomes even with the same prompt. Secondly, the model's performance is vulnerable to noisy contexts, which are common in cryptocurrency news. For example, sentences like "the movie is good," if injected into the news, could increase misclassification. As a result, we incorporated a "think-tank discussion"-like prompt pattern into the few-shot learning technique by repeating the following block multiple times with k examples for three different sentiment labels (i.e., 3-way-k-shot learning):

> [m] different cryptocurrency traders are reading this news. Each trader will assign a sentiment label from ["negative", "positive", "neutral"]. Then, each trader will share their label with the group. The majority label will be accepted. Return the majority label without any other text. The news is [news content]
> Label: [True sentiment label]

This approach aligns with consistency-based calibration methods [34], [35], which use agreement scores among LLM "voters" to determine confidence. However, our method is more efficient and cost-effective, as it avoids running the LLM multiple times with the same prompt. In our experiments, we set m to 3 and used GPT-3.5-Turbo [36] for sentiment analysis.

Incorporating cryptocurrency market sentiment is challenging due to its volatility, which may introduce noise. However, we found that sentiment can effectively regularize fluctuation predictions. First, the sentiment vector is embedded during the observation window using the previously introduced structure, producing a tensor $\mathbf{s}^{\text{emb}}$. This tensor serves two purposes: (1) It passes through the $\zeta(\cdot)$ structure from Section III-B, followed by a Tanh activation, to generate $\kappa \in (-1, 1)$ for regularizing price change predictions. (2) It combines with $\mathbf{x}^{\text{emb}}_g$ to determine how to fuse the two predictions. This fusion is critical as *market environment-based predictions are less volatile, while price dynamics predictions are more volatile*, enhancing generality across cryptocurrencies:

$$\hat{p}^i_{L+1} = \gamma * \hat{p}^{i,1}_{L+1} + (1-\gamma) * \hat{p}^{i,2}_{L+1}, \gamma = \zeta([\mathbf{x}^{\text{emb}}_g; \mathbf{s}^{\text{emb}}]). \tag{5}$$

IV. Experiment

A. Experiment Setup

Dataset: The proposed method is evaluated using real-world data, including cryptocurrency prices from Yahoo Finance [37], a leading source of financial information in the U.S., and news articles from Cointelegraph [38], a prominent news outlet for blockchain and cryptocurrency analysis. The price dataset includes historical prices for cryptocurrencies with market valuations above \$8 billion, spanning January 1, 2021, to April 1, 2024, and representing 92.18% of the total market cap. From this price data, we calculated and incorporated seven widely-used technical indicators [19], [21], [39], which are traditionally employed by market analysts to derive market trend insights and enhance predictive capabilities

TABLE I

FORECASTING RESULTS FOR THE TOP 5 INDIVIDUAL CRYPTOCURRENCIES, AS WELL AS AVERAGES FOR THE TOP 10, 15, AND 20. LOWER MAE AND MSE VALUES, AND CORR VALUES CLOSER TO 1, INDICATE BETTER PERFORMANCE. THE BEST-PERFORMING MODEL IS HIGHLIGHTED IN BOLD, WITH THE SECOND-BEST UNDERSCORED. † USES PRICE AND TECHNICAL INDICATORS, AND ‡ USES PRICE, TECHNICAL INDICATORS, AND NEWS SENTIMENT.

Method	Bitcoin			Ethereum			Tether			Binance Coin		
	MAE	MSE	CORR	MAE	MSE	CORR	MAE	MSE	CORR	MAE	MSE	CORR
SVM†	0.5530	0.4239	0.0083	0.4420	0.3006	0.2317	0.3884	0.2552	0.2715	0.6887	0.6086	0.6072
RF†	0.5338	0.3778	0.0159	0.4808	0.3398	-0.372	1.2149	3.8364	-0.0804	0.6513	0.5827	-0.1409
GRU‡	0.2299	0.0976	0.9810	0.1427	0.0387	0.9702	0.4731	0.3722	0.5120	0.1249	0.0294	0.9900
LSTM‡	0.3396	0.2458	0.9445	0.1952	0.0888	0.9494	0.5147	0.4988	0.4456	0.2129	0.1036	0.9529
Bi-LSTM‡	0.3235	0.2126	0.9675	0.1947	0.0751	0.9594	0.4464	0.3779	0.4864	0.1933	0.0714	0.9775
CNN-LSTM‡	0.2749	0.1294	0.9403	0.3511	0.2548	0.8420	0.4946	0.4064	0.3361	0.2268	0.0902	0.9649
DLinear‡	0.2975	0.1859	0.9725	0.4009	0.3555	0.4701	0.3963	0.2893	0.3098	0.2213	0.0816	0.9791
Linear‡	0.3625	0.3199	0.9433	0.2600	0.1376	0.8748	0.4474	0.3510	0.2503	0.6565	0.8487	0.3896
NLinear‡	0.1376	0.0306	0.9879	0.1065	0.0202	0.9815	0.3627	0.2283	0.6577	0.0948	0.0212	0.9902
Autoformer‡	0.1604	0.0408	0.9848	0.1594	0.0383	0.9667	0.3929	0.2656	0.6130	0.1627	0.0447	0.9805
CryptoPulse‡	**0.0607**	**0.0095**	**0.9961**	**0.0529**	**0.0065**	**0.9937**	**0.3249**	**0.1891**	**0.6946**	**0.0563**	**0.0103**	**0.9949**
Method	Solana			Top 10			Top 15			Top 20		
	MAE	MSE	CORR	MAE	MSE	CORR	MAE	MSE	CORR	MAE	MSE	CORR
SVM†	0.5375	0.4269	0.1289	0.4305	0.2967	0.3191	0.4377	0.2955	0.3321	0.7813	2.3293	0.2077
RF†	0.6302	0.5464	-0.1028	0.6337	0.8240	-0.1880	0.7689	1.0265	-0.2067	0.9955	2.6771	-0.1491
GRU‡	0.1709	0.0592	0.9822	0.3742	1.9460	0.8295	0.3142	1.6011	0.8839	0.4132	1.7916	0.9091
LSTM‡	0.3246	0.1805	0.9404	0.4811	2.1543	0.7970	0.3520	0.8895	0.8379	0.5340	1.7531	0.8745
Bi-LSTM‡	0.2979	0.1553	0.9210	0.3215	0.6710	0.8076	0.3394	1.1232	0.8468	0.5018	1.7740	0.8770
CNN-LSTM‡	0.2655	0.1565	0.9436	0.4000	1.2353	0.7054	0.3490	0.8692	0.7689	0.5266	1.6552	0.8063
DLinear†	0.4651	0.4820	0.5740	0.3239	0.2873	0.6892	0.3123	0.2499	0.7314	0.3973	0.4790	0.7504
Linear‡	0.2228	0.1013	0.9486	0.3721	0.4585	0.5743	0.3640	0.3925	0.6366	0.4408	0.5934	0.7025
NLinear†	0.1410	0.0297	0.9839	0.1565	0.0517	0.8865	0.1429	0.0430	0.9185	0.1387	0.0421	0.9380
Autoformer‡	0.1474	0.0351	0.9814	0.2037	0.0842	0.8435	0.1919	0.0720	0.8841	0.1939	0.0744	0.9089
CryptoPulse‡	**0.0511**	**0.0064**	**0.9962**	**0.0905**	**0.0301**	**0.9073**	**0.0758**	**0.0224**	**0.9364**	**0.0774**	**0.0225**	**0.9516**

for cryptocurrency forecasting. The news dataset comprises 25,210 articles from the Cointelegraph, spanning the same period. This dataset captures sentiment and market-relevant insights, complementing the price and technical indicator data to support robust and accurate forecasting.

Metrics: Following [27], [40], we use MSE, MAE, and cross-correlation (CORR) to evaluate models performance.

Comparison Methods: Ten SOTA baseline methods are considered for comparison. We adopt the same settings as the original papers. For models capable of directly incorporating technical indicators and sentiment labels alongside price history, we report their performance on the full dataset in the main results to ensure a fair comparison. Superscripts are added in Table I to differentiate model variants based on dataset configurations used for testing. The selected methods include four general time series forecasting models: DLinear [27], NLinear [27], Linear [27], and Autoformer [26]; three RNN-based methods: LSTM [10], [15], GRU [10], and Bi-LSTM [10]; one hybrid RNN model: CNN-LSTM [12]; and two traditional machine learning methods adapted for cryptocurrency forecasting: SVM [8] and RF [8].

We fixed the observation window at 7 days (i.e., $L = 7$) and split the dataset chronologically into training, validation, and test sets using a 7:1:2 ratio. All results are averaged over five experiments

B. Main Results

We evaluate the performance of our proposed model, CryptoPulse, by comparing it with ten SOTA models. Due to space limitations, results for all 75 cryptocurrencies are not included; instead, Table I presents performance for the top five cryptocurrencies by market value and the average for the top 10, 15, and 20 cryptocurrencies. This provides insights into both individual cryptocurrency predictions and broader market trends. All results are averaged over five experiments.

CryptoPulse consistently outperforms all baseline methods across key metrics as shown in Table I. For the top 5 cryptocurrencies, our model improves MAE by 10.4% to 63.8% and MSE by 17.2% to 69.0% compared to the best method. For the top 10, 15, and 20 cryptocurrencies, MAE improves by 42.2% to 46.9% and MSE by 41.8% to 47.9%. These results highlight the effectiveness of CryptoPulses design, which incorporates macroeconomic approximations, technical indicators, and market sentiment analysis to enhance cryptocurrency price forecasting.

In addition to these results, we identified key insights that contribute to improved performance. To understand these factors, we conducted a comprehensive analysis by posing and answering the following questions:

Are traditional machine learning models expressive enough for this task? Deep neural networks often outperform traditional models due to their superior expressive capacity,

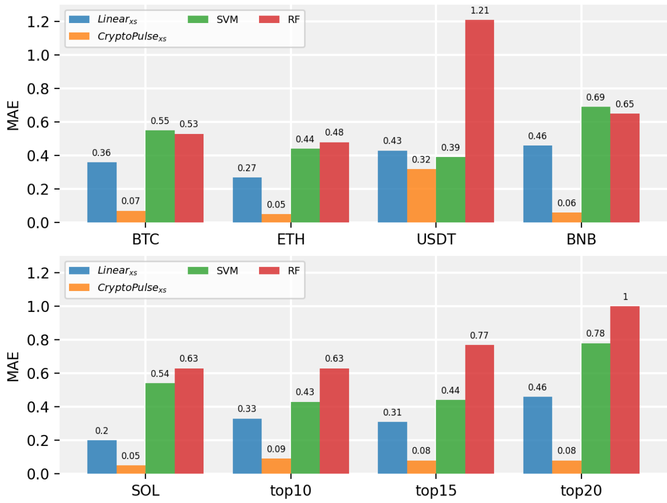

Fig. 2. Deep Learning vs. traditional models on data without sentiment.

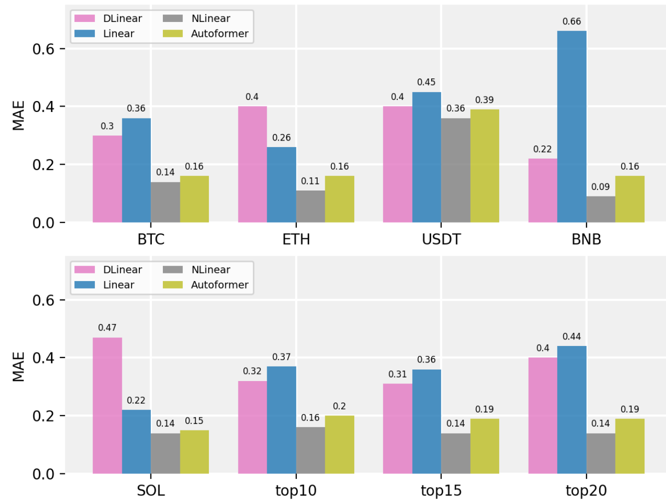

Fig. 4. MAE comparison between linear and transformer-based models.

Fig. 3. Comparison of RNN-based models.

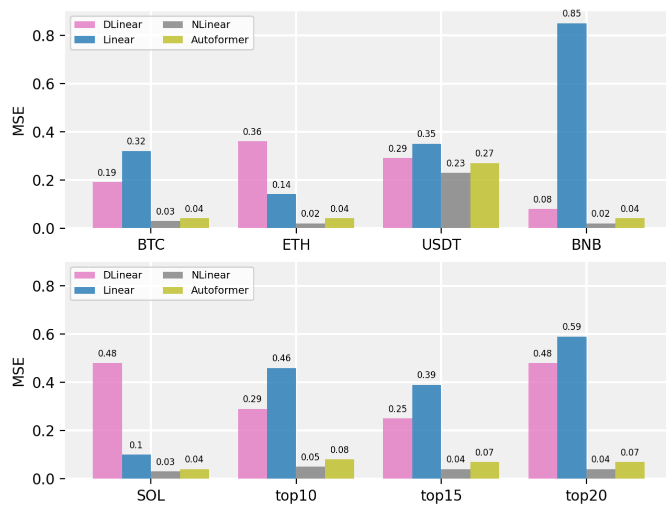

Fig. 5. MSE comparison between linear and transformer-based models.

though traditional models can perform comparably in simpler tasks. As shown in Table I, traditional models (SVM and RF) perform significantly worse than deep learning models. To rule out sentiment data as the cause, we conducted an ablation study on our model and the Linear model (the smallest deep learning model), denoted as $CryptoPulse_{xs}$ and $Linear_{xs}$, by removing sentiment data. Figure 2 shows that traditional models still underperform, except in USDT prediction, where SVM slightly outperforms $Linear_{xs}$. These results suggest that the weak performance of traditional models stems from their limited expressive capability.

Are RNN-based models outdated? RNN-based models can still achieve comparable performance in some cases. Among the four RNN-based models, the best one outperforms the Linear model in MAE or MSE in 12 out of 16 cases, DLinear in 9, and Autoformer in 3. While no single RNN-based model consistently dominates, GRU generally performs better, as shown in Figure 3, likely due to its simple recurrent architecture, which is less prone to overfitting cryptocurrency's dynamic patterns. Due to space constraints, the figure for MSE performance is omitted, but Table I shows similar patterns as MAE. Another key observation is that RNN-based predictions are more stably correlated to the ground truths than those of DLinear and Linear models. Thus, RNN-based models remain important benchmarks in our scenarios. However, our model not only outperforms RNN-based models across all cases but is also more computationally efficient.

Are linear models always better than Transformer-based models? We investigate whether Linear models consistently outperform Transformers, as observed by [27] in other tasks. Our findings show this is not necessarily the case. As shown in Figure 4 and 5, DLinear and Linear perform worse than Autoformer, while NLinear consistently outperforms Autoformer, with comparable results. Linear models, which do not explicitly capture correlations across time series, benefit from price-related data, technical indicators, and sentiment

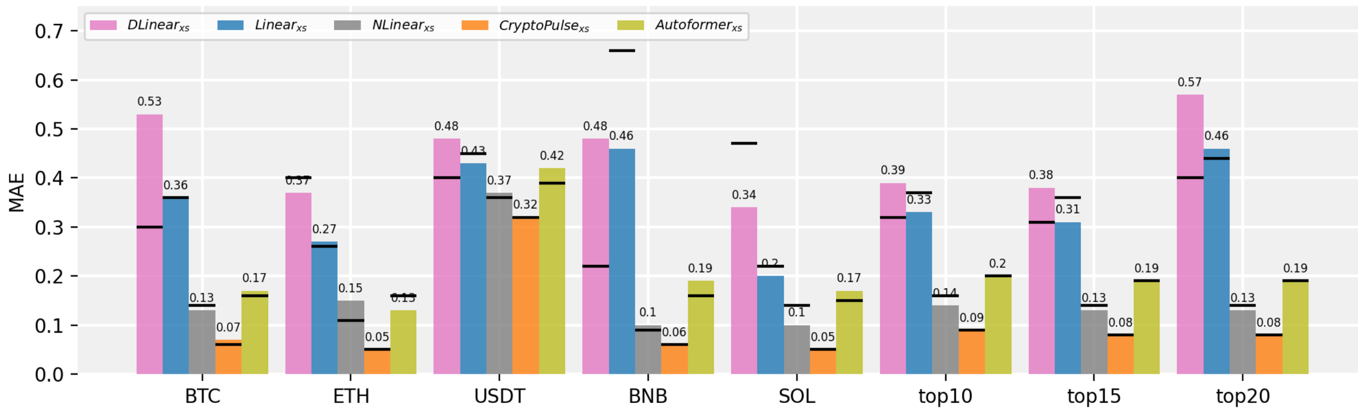

Fig. 6. MAE comparison of models with sentiment data removed.

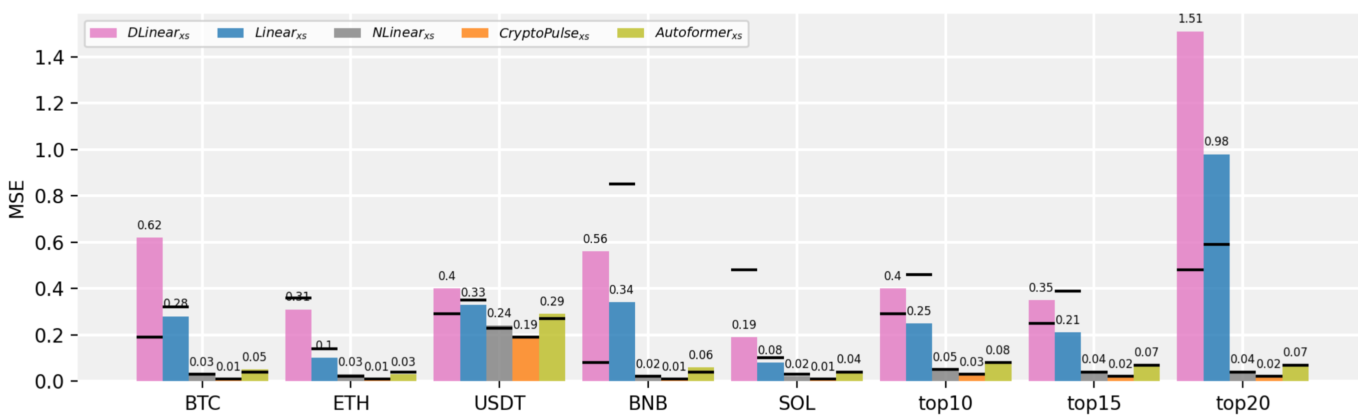

Fig. 7. MSE comparison of models with sentiment data removed.

when combined in forecasting tasks. Transformer-based models better handle these complex correlations due to their larger size. However, DLinear and Linear exhibit instability in high-volatility scenarios, particularly with MSE as the metric, as shown in Figure 5.

Can trend analysis benefit performance? Series decomposition is a common method in time-series forecasting, and we examined its role in our task. DLinear, our model, and Autoformer explicitly account for trend patterns, while RNN-based models capture changes across time points implicitly. Trend analysis proves to be a double-edged sword. On one hand, as shown in Figure 4 and 5, improperly modulated moving-average-based trends can destabilize models like DLinear due to cryptocurrency volatility disrupting the moving average. On the other hand, Autoformer balances seasonal and trend-cyclical components effectively, producing more stable forecasts. The short observation window likely amplifies this effect but is necessary, as long-term patterns are rarely present in cryptocurrency markets for next-day predictions.

C. Ablation Study

In this subsection, we analyzed the impact of each group of financial features on forecasting performance. First, we examined the effect of sentiment data on cryptocurrency prediction by removing news sentiment from the dataset for all Linear-based and Transformer-based models, denoted with an xs subscript. Figures 6 and 7 visualize the results, where bar heights indicate performance, and black horizontal lines represent full-feature performance, including price history, technical indicators, and market sentiments. Sentiment data generated using our proposed LLM-based approach improves forecasting performance overall, as shown by the black line within the bars. However, NLinear outperformed its full-feature version in 5 of 8 cases, likely due to its reliance on time series continuity, where sentiment labels introduce noise from missing values on days without cryptocurrency-related news. DLinear and Linear models showed instability, with significant performance differences when sentiment data was excluded.

Second, we removed technical indicators from the feature set and conducted ablation studies on all linear-based and Transformer-based models, as well as our own. These models are denoted with an xi subscript. Figures 8 and 9 present the results, with black lines marking full-feature performance. Including technical indicators improved performance overall, with DLinear and Autoformer benefiting the most, while our model showed slight improvement. These indicators, derived from financial domain knowledge, are designed to be less sensitive to short-term fluctuations, providing trend analysis insights that enhance automated models.

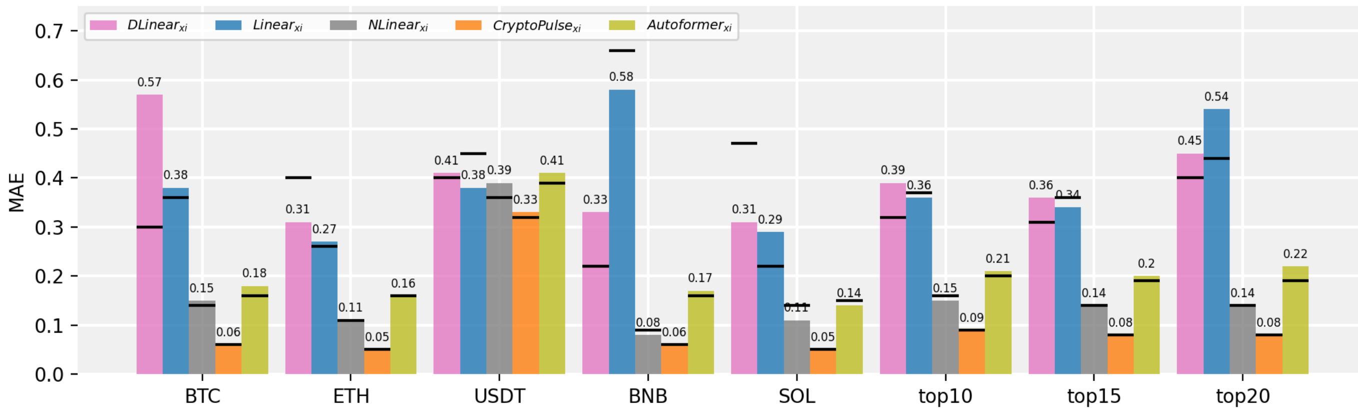

Fig. 8. MAE comparison of models with technical indicators data removed.

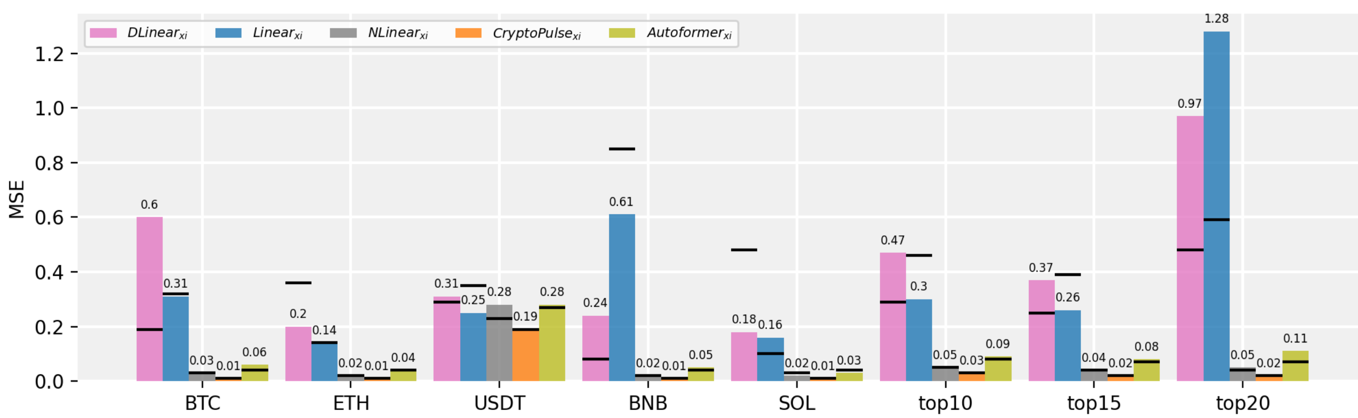

Fig. 9. MSE comparison of models with technical indicators data removed.

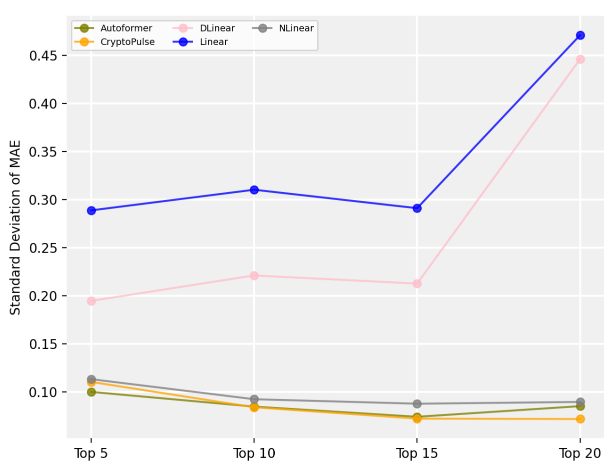

Fig. 10. Standard deviation of MAE across models for top cryptocurrencies.

D. Robustness

The robustness of the models was evaluated by calculating the standard deviation of MAE across five independent experiments, averaged over the top 5, 10, 15, and 20 cryptocurrencies. A lower standard deviation indicates higher consistency and robustness across training runs. To avoid overcrowding in figures, the analysis focuses on Linear-based and Transformer-based methods, along with our proposed model, as these demonstrated superior performance. As shown in Figure 10, our approach achieves the lowest standard deviation for the top 10, 15, and 20 cryptocurrencies and performs comparably to the best for the top 5. Furthermore, it demonstrates notable robustness when handling smaller market-cap cryptocurrencies, which are typically more volatile, a challenge where Linear and DLinear methods struggle.

V. Conclusion

In this paper, we present "CryptoPulse", a new approach to predicting the next-day closing prices of cryptocurrencies. This model integrates three primary factors: fluctuations in the macro environment, changes in individual cryptocurrency prices and technical indicators, and overall cryptocurrency market sentiment. By leveraging a dual prediction mechanism, the model captures both the macro market environment and the specific price and technical indicator dynamics of the target cryptocurrency. Moreover, a fusion component based on the market sentiment information integrates these predictions to improve the results. The experimental evaluation shows that our model achieves higher accuracy in predicting cryptocurrency fluctuations compared to ten different methods, making it suitable for application in the highly unpredictable cryptocurrency market.

VI. Acknowledgment

This work is supported in part by the National Science Foundation via grants NSF CNS-2431176 and NSF ITE-2431845. The US Government is authorized to reproduce and distribute reprints of this work for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF, or the U.S. Government.

References

[1] J. Fry and E. Cheah, "Negative bubbles and shocks in cryptocurrency markets," *International Review of Financial Analysis*, vol. 47, pp. 343–352, 2016.

[2] G. W. Peters and E. Panayi, "Understanding modern banking ledgers through blockchain technologies: future of transaction processing and smart contracts on the internet of money," *SSRN Electronic Journal*, 2015.

[3] M. H. Joo, Y. Nishikawa, and K. Dandapani, "Cryptocurrency, a successful application of blockchain technology," *Managerial Finance*, vol. 46, pp. 715–733, 2019.

[4] E.-T. Cheah and J. Fry, "Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin," *Economics letters*, vol. 130, pp. 32–36, 2015.

[5] F. Fang, C. Ventre, M. Basios, L. Kanthan, D. Martinez-Rego, F. Wu, and L. Li, "Cryptocurrency trading: a comprehensive survey," *Financial Innovation*, vol. 8, no. 1, p. 13, 2022.

[6] M. A. Yamin and M. Chaudhry, "Cryptocurrency market trend and direction prediction using machine learning: a comprehensive survey," 2023.

[7] N. Trabelsi, "Are there any volatility spill-over effects among cryptocurrencies and widely traded asset classes?" *Journal of Risk and Financial Management*, vol. 11, no. 4, p. 66, 2018.

[8] E. Akyildirim, A. Goncu, and A. Sensoy, "Prediction of cryptocurrency returns using machine learning," *Annals of Operations Research*, vol. 297, pp. 3–36, 2021.

[9] W. Zhengyang, L. Xingzhou, R. Jinjin, and K. Jiaqing, "Prediction of cryptocurrency price dynamics with multiple machine learning techniques," in *Proceedings of the 2019 4th International Conference on Machine Learning Technologies*, 2019, pp. 15–19.

[10] P. L. Seabe, C. R. B. Moutsinga, and E. Pindza, "Forecasting cryptocurrency prices using LSTM, GRU, and bi-directional LSTM: A deep learning approach," *Fractal and Fractional*, vol. 7, no. 2, p. 203, 2023.

[11] M. J. Hamayel and A. Y. Owda, "A novel cryptocurrency price prediction model using GRU, LSTM and bi-LSTM machine learning algorithms," *Ai*, vol. 2, no. 4, pp. 477–496, 2021.

[12] Y. Li and W. Dai, "Bitcoin price forecasting method based on CNN-LSTM hybrid neural network model," *The journal of engineering*, vol. 2020, no. 13, pp. 344–347, 2020.

[13] C. Lamon, E. Nielsen, and E. Redondo, "Cryptocurrency price prediction using news and social media sentiment," *SMU Data Sci. Rev*, vol. 1, no. 3, pp. 1–22, 2017.

[14] Y. Pang, G. Sundararaj, and J. Ren, "Cryptocurrency price prediction using time series and social sentiment data," in *Proceedings of the 6th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*, 2019, pp. 35–41.

[15] A.-D. Vo, Q.-P. Nguyen, and C.-Y. Ock, "Sentiment analysis of news for effective cryptocurrency price prediction," *International Journal of Knowledge Engineering*, vol. 5, no. 2, pp. 47–52, 2019.

[16] G. W. Brown and M. T. Cliff, "Investor sentiment and the near-term stock market," *Journal of empirical finance*, vol. 11, no. 1, pp. 1–27, 2004.

[17] S. Peerzade, D. Wayal, and G. Kale, "Automated algorithmic trading for cryptocurrencies," *International Journal of Advanced Research in Science, Communication and Technology*, pp. 326–330, 2021.

[18] Delfabbro, P. and King, D. L. and Williams, J. N., "The psychology of cryptocurrency trading: risk and protective factors," *Journal of Behavioral Addictions*, vol. 10, pp. 201–207, 2021.

[19] P. Oncharoen and P. Vateekul, "Deep learning for stock market prediction using event embedding and technical indicators," in *2018 5th international conference on advanced informatics: concept theory and applications (ICAICTA)*. IEEE, 2018, pp. 19–24.

[20] J. Kaur and K. Dharni, "Data mining–based stock price prediction using hybridization of technical and fundamental analysis," *Data Technologies and Applications*, vol. 57, no. 5, pp. 780–800, 2023.

[21] Y. Zhai, A. Hsu, and S. K. Halgamuge, "Combining news and technical indicators in daily stock price trends prediction," in *Advances in Neural Networks–ISNN 2007: 4th International Symposium on Neural Networks, ISNN 2007, Nanjing, China, June 3-7, 2007, Proceedings, Part III 4*. Springer, 2007, pp. 1087–1096.

[22] S. Dastgir, E. Demir, G. Downing, G. Gozgor, and C. K. M. Lau, "The causal relationship between Bitcoin attention and Bitcoin returns: Evidence from the Copula-based Granger causality test," *Finance Research Letters*, vol. 28, pp. 160–164, 2019.

[23] A. H. Dyhrberg, "Bitcoin, gold and the dollar–A GARCH volatility analysis," *Finance research letters*, vol. 16, pp. 85–92, 2016.

[24] N. Gandal and H. Halaburda, "Can we predict the winner in a market with network effects? Competition in cryptocurrency market," *Games*, vol. 7, no. 3, p. 16, 2016.

[25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[26] H. Wu, J. Xu, J. Wang, and M. Long, "Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting," *Advances in neural information processing systems*, vol. 34, pp. 22 419–22 430, 2021.

[27] A. Zeng, M. Chen, L. Zhang, and Q. Xu, "Are transformers effective for time series forecasting?" in *Proceedings of the AAAI conference on artificial intelligence*, vol. 37, no. 9, 2023, pp. 11 121–11 128.

[28] K. A. Coulter, "The impact of news media on bitcoin prices: modelling data driven discourses in the crypto-economy with natural language processing," *Royal Society Open Science*, vol. 9, 2022.

[29] J. J. Schulp, "Crypto Crash: Why the FTX Bubble Burst and the Harm to Consumers," *CATO Institute*, 2022.

[30] H. Xu, B. Liu, L. Shu, and P. S. Yu, "BERT post-training for review reading comprehension and aspect-based sentiment analysis," *arXiv preprint arXiv:1904.02232*, 2019.

[31] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith, and D. C. Schmidt, "A prompt pattern catalog to enhance prompt engineering with chatgpt," *arXiv preprint arXiv:2302.11382*, 2023.

[32] W. Zhang, Y. Deng, B. Liu, S. J. Pan, and L. Bing, "Sentiment analysis in the era of large language models: A reality check," *arXiv preprint arXiv:2305.15005*, 2023.

[33] C. Si, Z. Gan, Z. Yang, S. Wang, J. Wang, J. Boyd-Graber, and L. Wang, "Prompting gpt-3 to be reliable," *arXiv preprint arXiv:2210.09150*, 2022.

[34] X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou, "Self-consistency improves chain of thought reasoning in language models," *arXiv preprint arXiv:2203.11171*, 2022.

[35] M. Zhang, J. He, T. Ji, and C.-T. Lu, "Don't Go To Extremes: Revealing the Excessive Sensitivity and Calibration Limitations of LLMs in Implicit Hate Speech Detection," *arXiv preprint arXiv:2402.11406*, 2024.

[36] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, "Training language models to follow instructions with human feedback," *Advances in neural information processing systems*, vol. 35, pp. 27 730–27 744, 2022.

[37] Lawrence, A. and Ryans, J. P. and Sun, E. and Laptev, N., "Earnings announcement promotions: a yahoo finance field experiment," *Journal of Accounting and Economics*, vol. 66, pp. 399–414, 2018.

[38] L. Phan, S. Li, and K. Mentzer, "Blockchain technology and the current discussion on fraud," 2019.

[39] K. Kim, "Financial time series forecasting using support vector machines," *Neurocomputing*, vol. 55, pp. 307–319, 2003.

[40] T. Derrick and J. Thomas, "Time series analysis: the cross-correlation function," 2004.