Predicting Stock Prices of Digital Cultural Tourism Enterprises: A Comparative Study of Deep Learning Architectures

WENRUI JIANG[1,*], ZIJIAN ZENG[2]
[1]Guangdong Engineering Polytechnic,
Guangzhou 510520,
CHINA

[2]Institute of Computer Science and Digital Innovation,
UCSI University WP,
Kuala Lumpur 56000
MALAYSIA

**Corresponding Author*

Abstract: - The development of the digital cultural tourism industry in China has created numerous opportunities for investors; however, at the same time, it has also caused considerable instability within the stock market. Therefore, the ability to forecast the movement of stocks for cultural tourism companies is valuable to both investors and decision-makers. To address this issue, we compared seven different modeling techniques (Random Walk, ARIMA, Random Forest, Long Short Term Memory (LSTM), Gated Recurrent Unit (GRU), Transformer, and Domain Aware Hybrid LSTM-Transformer (HLT)) to determine which technique best predicted the stock prices of four of China's most successful A-share cultural tourism companies. We incorporated fifteen technical indicators, a proxy for tourism demand based upon search index data from Baidu, and macroeconomic indicators into our models as predictive factors. To avoid any "look-ahead" bias, we lagged all of our independent variables by the date that they became available. During walk forward testing across 1,680 trading days between 2018-2024, our findings indicate that the HLT model outperformed every other model that we tested. The average sMAPE for the HLT model was 2.14%, while the average MASE was 0.673. Furthermore, the average directional accuracy for the HLT model was 72.3%. Our ablation analysis revealed that incorporating the tourism search index data as an independent variable improved the sMAPE by approximately 0.45 percentage points. Based upon our findings, we believe that incorporating advanced sequence based models with domain knowledge can significantly improve the decision-making process for financial professionals working in the area of digital cultural tourism finance.

Key-Words: - digital cultural tourism; stock price prediction; LSTM; Transformer; hybrid deep learning; quantitative finance; time series forecasting; feature engineering.

Received: September 12, 2026. Revised: October 3, 2026. Accepted: October 17, 2026. Published: November 24, 2026.

1 Introduction

Technology, culture, and economic growth are driving the rapid expansion of digital cultural tourism. Technologies such as AR, VR, Big Data Analytics, and AI are being integrated into heritage sites, scenic areas, and cultural performances, creating a new way for individuals to interact with cultural resources and transforming the supply side of the tourism industry [1]. As of 2024, the Cultural & Related Industries and Tourism & Related Industries in China represented approximately 4.6% and 4.4% of China's GDP respectively; it is expected that the global cultural tourism market will receive significant investment through 2030 [2].

The way that publicly traded cultural tourism companies determine their prices differs greatly from how most other types of stocks do. Many different factors can affect the price of a publicly traded cultural tourism company. Some examples of these types of factors would include: the timing of public holidays and the resulting fluctuations in consumer spending; changes in public policy regarding the digitization of cultural heritage; how public health emergencies such as COVID-

19 influence consumer behavior; and to what extent digital platforms are being used to engage with consumers (for example, consumers searching for cultural attractions online or posting about them on social media) [3]. However, traditional econometric models (such as ARIMA and vector autoregressive models) fail to account for the non-linear relationship of some of the variables mentioned above as well as the multi-dimensional nature of many of those same variables. On the other hand, the development of deep learning based sequential modeling offers a much better way of including multiple data sources into an analysis [4].

A growing number of people are using time series analysis to forecast future values of financial variables because of advances in deep learning. Two examples include the creation of long short term memory (LSTM) neural networks and transformer based neural networks. Both forms of neural network have shown to outperform conventional statistical techniques when used for forecasting in different markets and different types of assets [5], [6]. Both LSTM and transformer based neural networks can be applied to many different areas but there are few examples of either being tested as part of the publicly traded digital cultural tourism company sector in China. Researchers have primarily focused their studies on the total demand for tourism, on technology related sectors, and on broad based market indices; however, there is currently no research on how well individual publicly traded cultural tourism companies perform.

The purpose of this study was to address the gap in the current literature mentioned previously. This research will make three significant contributions. First, it will create a unique domain specific feature set, including 15 financial technical indicators, proxy for tourism demand using Baidu Search Index data and multiple macro-economic variables related to cultural tourism stocks. Second, this research will conduct a systematic evaluation of several different types of forecasting models (Random Walk baseline, AutoRegressive Integrated Moving Average [ARIMA], Random Forest [RF], Long Short-Term Memory [LSTM], Gated Recurrent Unit [GRU], Transformer [TF] and a hybrid model combining LSTM and TF referred to as Domain Aware Hybrid LSTM-Transformer [HLT]). Each model is evaluated under a walk forward methodology so that there is no possibility of "look-ahead bias" being introduced through the use of external variables. Finally, this research will perform ablation studies, statistical tests (Diebold-Mariano and Pesaran-Timmermann) and a trading backtest accounting for transaction costs in order to provide evidence that the forecasting results can be converted into useful trading decisions by investors.

The Hybrid Framework for domain specific applications (HLT) is being presented as an alternative to an entirely new architectural design. Transformers combined with LSTMs have been used previously in research concerning financial time series data [4], [5]. We combined the Transformer/LSTM approach to develop a predictive model for cultural tourism data. We evaluated the model using a "walk-forward" validation process, which ensured that no information from a later period was available to the model during training. Our results indicated that our model produced predictions that would be beneficial when implementing a long-only trading strategy.

2 Related Work

2.1 Machine Learning for Stock Prediction

The development of machine learning as used in the area of financial forecasting is very fast. A systematic survey of machine learning applications in the tourism industry showed that there are many different types of neural networks that have been used to solve a variety of tourism-related forecasting problems [7]. For example, the survey of machine learning applications to stocks demonstrated that LSTM and Transformer-based deep learning models were consistently producing better results than ARIMA-class models in terms of stock price forecasting [4]. Furthermore, the empirical study that compared the performance of Transformer, LSTM and GRU on Tesla daily price forecasts demonstrated that while the Transformer model performed slightly better than LSTM and GRU on most of the error metrics, all three deep learning models performed better than ARIMA [6]. Lastly, the CTL model proposed in [5], which incorporates CEEMDAN decomposition, a Transformer encoder and an LSTM decoder, demonstrated superior performance compared to existing state-of-the-art methods for stock price forecasting in China, Hong Kong and the United States.

2.2 Tourism Demand and Tourism-Stock Forecasting

Another field of research in which great strides have been made in the area of predicting tourism demand is through machine learning. The researchers who authored reference [8] illustrate how a mixed-frequency machine learning model can be used to accurately forecast tourism demand and thus, produce more accurate forecasts than traditional qualitative/quantitative models. Also, in reference [9], the authors show that when uncertainty indicators are incorporated into the model, the accuracy of forecasting tourism flows increases.

The authors of [3] developed a multi-scale tourism stock index model utilizing social media data; they found that emerging media signals have statistically significant predictive power beyond just purely financial measures. Also, the study in [10] showed that AI-based approaches significantly outperform traditional econometric methods for forecasting financially in the U.S. tourism industry. Lastly, the authors of [11] used deep learning for economic analysis and developing marketing strategies in the tourism industry, showing that sector-specific deep models can capture tourism-economy dynamics that general price models cannot.

3 Methodology

3.1 Dataset and Walk-Forward Splits

Daily stock prices of four representative Chinese A-share digital cultural tourism firms, namely Songcheng Performance (300144.SZ), Zhangjiajie Tourism (000430.SZ), Lijiang Tourism (002033.SZ) and China Tourism Group Duty Free Corporation (601888.SH) were collected using Wind Financial Terminal. They were selected due to their representativeness in the industry. For each firm, the daily open price, high price, low price, closing price and trading volume (OHLCV) were recorded. This study spans over the period from January 2018 to December 2024, representing 1680 trading days for each firm.

The walk-forward evaluation method is used in this study. The first stage of developing the model is done with data from January 2018 through December 2021 (the training time), and January through December 2022 (the validation time). After the initial development of the model is complete, it will continue to be updated/re-developed every three months, as the test period moves on from January 2023 to December 2024. To develop the final held-out test sets, we must remove the 30-day warm-up window that is necessary for our technical indicator pipelines, and we must also remove the days that the underlying stocks were halted. Therefore, the resulting held-out test sets contain 478 (300144.SZ), 471 (000430.SZ), 482 (002033.SZ) and 484 (601888.SH) effective trading days (for a total of 1,915), with no overlapping between the training and testing periods.

3.2 Feature Set and Look-Ahead Control

The feature groups can be divided into three groups. The first group is composed of the technical indicators. There are fifteen indicators within the technical indicator group. Indicators such as the RSI-14, MACD (12/26/9), Bollinger Bands (+/-2 sigma), Average True Range (ATR-14), On-Balance Volume, Stochastic Oscillators (%K & %D), Williams %R, Commodity Channel Index, and the Rate-of-Change over five, ten and twenty day periods are included in the technical indicator group. In addition to these technical indicators, there are three weekly search indexes from Baidu for the cultural tourism keywords; scenic area tickets keywords; and company specific venues keywords. Finally, there is a macro-economic group that includes the monthly Tourism Services Consumer Price Index (CPI); dummy variables representing the Golden Week and Spring Festival holidays; and the Wind Cultural Tourism Sub-Index that is used as a sector momentum indicator.

At the close of each trading day (day t), all exogenous variables are used in conjunction with the most current available data to generate predictions for the following trading day (day t+1). To guarantee that the data used in the prediction process was generated during or prior to day t, the weekly search index from Baidu was lagged by one week. In order to allow for the publication of the official Tourism Services CPI from the National Bureau of Statistics (NBS), which is normally published on or around the 9th to 12th of each subsequent month, the monthly CPI was lagged until the NBS published it. A cultural tourism sub-index was developed using Wind without incorporating any of the four stocks being analyzed, so as to avoid generating artificial contamination. To scale features, a min-max scaling technique was utilized. Min-max scaling is performed utilizing only data within the training period. Each sequential model uses a 30-day input window.

3.3 Models and Training

The baseline models are Naive Random Walk, 5-Day Moving Average, ARIMA(1,1,1), and Random Forest. The neural models are a two layer LSTM/GRU (each has 128 hidden units/layer, 0.2 dropout), an encoder-only transformer (4 attention heads, 128 model dim, 512 feed-forward dim, 3 encoder layers), and our new HLT framework. In our HLT framework, we use a transformer encoder to establish global temporal dependencies; we then pass these global dependencies to an LSTM layer to establish local momentum dependencies. Lastly, we project the LSTM output as follows:

$$\hat{y}_{t+1} = W_o \cdot \mathrm{LSTM}(\mathrm{TransEnc}(\mathbf{X}_{t-L:t})) + b_o \quad (1)$$

Each model is trained using the Adam optimization algorithm. The learning rate is 0.001. Beta1 is 0.9. Beta2 is 0.999. The batch size used is 64. A maximum of 200 epochs were allowed for training. Early stopping was implemented with the patience parameter being 15, based upon the MSE of the validation set. All experiments were conducted on a single NVIDIA RTX 3090 GPU running PyTorch 2.1.

3.4 Evaluation Metrics and Significance Tests

Because the prices of the four stocks analyzed vary significantly (e.g., 601888.SH is approximately 10 times more costly than 000430.SZ), if we averaged the absolute error metric values of all four stocks, the stock with the largest price would distort the overall metrics. Therefore, in order to prevent this distortion, we will utilize scale-free metrics to evaluate the overall performance of our models. We will specifically utilize scale-free metrics including Symmetric MAPE (sMAPE), Test Set Random Walk Relative MASE, and Directional Accuracy (DA) when evaluating the models' overall performance. Utilization of scale-free metrics will enable us to evaluate the performance of each model for each stock independently and subsequently combine those independent evaluations to assess the models' overall performance without allowing one stock to disproportionately affect the overall assessment. MASE is calculated by dividing the Mean Absolute Error (MAE) produced by our models for a specific stock by the Random Walk's MAE for the same stock over the same testing period. Since the Random Walk's MASE is defined as being equal to 1, any value below 1 indicates that our model performed better than the Random Walk model.

The direction of the last day's price change determines how well a model has performed when predicting that prices will remain unchanged. This way of evaluating no-change predictions is called a tie-breaking rule. Therefore, the random walk DA is reported only as a separate measure; the performance of the random walk model cannot be compared to that of any other model which provides price predictions other than "zero." A Diebold-Mariano test with a small sample correction (proposed by Harvey, Leybourne, and Newbold) will be used to compare the average difference in squared error loss between pairs of models on each of the four stocks. To account for testing multiple times across the four stocks, we will use the Holm-Bonferroni correction. We will use the Pesaran-Timmermann test to determine whether price movements are predictable.

4 Results and Discussion

4.1 Aggregate Comparison Performance

The aggregate performance of each model, evaluated across the four different stock data sets during the 2023-2024 walk-forward testing period is presented in Table 1. On an aggregate basis, using each of the three scale-free aggregate metrics, the HLT approach was found to outperform all other models. On an aggregate basis, the HLT approach was found to have a lower aggregate sMAPE than the single LSTM architecture (2.14 percent vs. 2.60 percent; a 17.7 percent relative improvement), as well as a lower aggregate MASE than the single LSTM architecture (0.673 vs. 0.850). The ARIMA approach, the 5-day moving average approach and the Random Forest approach were each found to be less successful on an aggregate basis than the random walk on the aggregate MASE measure. This illustrates just how challenging daily price predictions can be and how difficult it can be to outperform a random walk. On an aggregate basis, each of the four deep learning approaches were found to outperform the random walk, and the HLT approach was also found to outperform each of the other deep learning approaches on each of the three aggregate metrics.

Using both standalone LSTMs and Transformers with HLT, we have found that there are two separate forms of signal generated by the different parts of the technology. The self-

attention part of the technology allows the system to learn about the various effects of other factors (e.g., holidays, policy changes, etc.) that can influence the market prior to the onset of high-demand periods. The LSTM component of the technology helps support the short-term price trend, which is a necessary requirement for daily prediction. The results we obtained were comparable to the results obtained from a hybrid Transformer-LSTM model reported in [5], as well as those obtained from the large deep learning review on financial time series in [4].

Table 1. Aggregate results of the 2023-2024 walk-forward test (average over four stocks). MASE < 1 means that the model performed better than the random walk on the test set; the random-walk DA is an auxiliary result (see section 3.4).

Model	sMAPE(%)	MASE	DA(%)
Random Walk	2.71	1.000	50.4
Moving Avg. (5d)	3.18	1.060	49.8
ARIMA(1,1,1)	4.05	1.300	53.2
Random Forest	3.42	1.020	62.8
GRU	2.85	0.918	67.4
LSTM	2.60	0.850	69.1
Transformer	2.36	0.760	70.8
HLT (Ours)	2.14	0.673	72.3

Source: created by the authors

4.2 Per-Stock Results

Table 2 of the Appendix provides the results for all of the individual stocks within the test set from 2023-2024. These results were obtained through a comparison of three different forecasting models: Random-Walk Benchmark, Long Short-Term Memory Model (LSTM), and the Hybrid LSTM-Transformer Model (HLT). Two general trends can be observed from these results. One trend is that China Tourism Group Duty Free is the least difficult to predict using the HLT model (sMAPE 1.73%). The reason why this may be true is that China Tourism Group Duty Free has a large market capitalization, and also has numerous analysts who follow the company. On the other hand, the hardest to predict using the HLT model is Zhangjiajie Tourism (sMAPE 2.61%). It appears that Zhangjiajie Tourism is greatly influenced by regional conditions such as weather and government entry restrictions. Additionally, the standalone LSTM model could not beat the random-walk model on 601888.SH (MASE 1.011); therefore, it is necessary to include the random-walk model as a separate benchmark.

4.3 Statistical Significance

Table 3 (Appendix) presents the results of the Diebold-Mariano test of squared error loss differences between HLT and the individual LSTM models. For all four stocks, we used Holm-Bonferroni correction to adjust the p-values. We observed that the difference in losses was statistically significant at the 5% level for two of the stocks (000430.SZ was also statistically significant at the 1% level). One of the stocks had a marginally statistically significant difference in losses at the 10% level. The remaining stock (601888.SH) did not exhibit a statistically significant difference in losses. Table 4 (Appendix) displays the results of the Pesaran-Timmermann test for directional predictability. It indicates that directional predictability is statistically significant at the 1% level for HLT on each of the four stocks.

4.4 Impact of Tourism Specific Features

Table 5 in Appendix shows the results of an experiment in which the search index features of Baidu were eliminated, which raised the average sMAPE of the four Deep Learning Models by approximately 0.45 percentage points. Results for each stock analyzed were also presented in terms of improvements in sMAPE, and the results indicated that the improvements in sMAPE were not uniformly distributed across the various stocks. Stocks with large amounts of domestic tourism demand (Songcheng Performance and Lijiang Tourism) saw some of the highest increases in sMAPE; these two stocks have a high correlation to online search activity. Conversely, China Tourism Group Duty Free had the lowest increase in sMAPE because the majority of its income comes from duty free retail sales as opposed to paid admissions. This supports other studies [3] that demonstrated that search and new media signals can be utilized to predict the stock price movements of tourism related companies.

4.5 Trading Backtest

The model's potential usefulness may be assessed through a straightforward long-only back-testing exercise covering the time frame from 2023-2024. The model was developed to predict the next-day (t+1) closing price using the current-day (t) closing price. Once the current day (t) closes, the model produces a signal. If the model predicts that the next-day (t+1) closing price will exceed the current-day (t) closing price, then the stock is added to the portfolio for the following day. In order to avoid the unrealistic assumption of buying and selling at the same closing price, positions were initiated at the opening price of the

next day (t+1) and closed at the closing price of the next day (t+1). Capital was evenly distributed across stocks with a positive signal; however, on days when there was no positive signal, the strategy held cash and received a zero risk-free rate. We applied 6 bps on entry and 6 bps on exit (a total round-trip cost of 12 bps), which consisted of 8 bps for Chinese A-share commission and stamp duty, and 4 bps for assumed slippage, split equally. We assumed that there are 252 trading days each year and thus calculated daily portfolio returns arithmetically, and annualized them. We used zero as the risk-free rate for calculating Sharpe Ratios, and included both entry and exit transactions in calculating turnover.

The HLT Sharpe ratio is 1.18, net of all-in costs, as presented in Table 6 (Appendix). This Sharpe ratio is much higher than the Sharpe ratios of the two standalone neural baseline strategies, and it is also much higher than the Sharpe ratio of the equal-weight buy-and-hold benchmark strategy. Furthermore, the HLT Sharpe ratio is consistent with the predictive enhancements shown in Table 1, Table 2, Table 3, and Table 4 (Appendix). Thus, the HLT strategy represents an illustration of how even a very simple trading strategy, based upon next-day direction forecasts, can yield economically relevant information from the model. However, it should be noted that the backtest was performed under the assumption that transactions were made at the open and close auction price levels, and did not take into account borrowing constraints, position size risks, or other types of risk factors. A subsequent paper on portfolio construction will examine these issues.

4.6 Sector-Level Observations

The above-mentioned results do not imply causality; however, there are three sector-specific trends that help validate our model's predictions. First, error rates increase significantly during Chinese public holidays (Golden Week in May and October, Spring Festival in January or February) for the basic baseline models, while they do not for the Transformer based models. This supports the idea that attention-based models can account for the anticipating behaviors associated with holidays. Second, after governments make announcements concerning cultural-tourism digitalization policies, the performance difference between HLT and LSTM becomes greater. This suggests that the hybrid architecture will be able to adapt to changes brought about by policy shifts more effectively than can a single LSTM stack. Nonetheless, conducting a formal event-study analysis of this trend remains a potential area for future research. Finally, firms that have established more robust AR/VR product lines experience greater positive price movement in tandem with the overall digital economy index. This finding is supportive of the sustainable tourism and AI-driven growth hypothesis discussed in [1].

5 Conclusion

The present study provides a comparative analysis of various types of deep learning models used for predicting the stock prices of digital cultural tourism companies. By employing a leakage-controlled walk-forward methodology, the HLT model, which incorporates domain knowledge, achieved average aggregate sMAPE of 2.14%, MASE of 0.673, and directional accuracy of 72.3% for four Chinese A-share stocks representing cultural tourism during the years 2023-2024. These results were consistently superior to those obtained from both classical baseline models and individual neural network models across each aggregate metric. Furthermore, the inclusion of tourism search data consistently enhanced performance with a further reduction in sMAPE of approximately 0.45 percentage points, with the greatest per-stock performance occurring on leisure travel stocks that rely heavily on online search activity to generate revenue. A simple long-only back-test based on HLT predictions generated an annualized Sharpe ratio of 1.18 after adjusting for all-in trading costs of 12 bps. Therefore, the enhanced predictive abilities provided by HLT led to profitable trading strategies as opposed to simply being statistical anomalies. The results of the Diebold-Mariano test and Pesaran-Timmermann test supported the superiority of HLT over a single LSTM model; however, there was a diminished difference on the largest stock (601888.SH) as the random walk hypothesis represented a strong alternative.

The three main areas of future study are as follows: 1. Incorporate sentiment indicators obtained from large language model analysis of earnings call transcripts and tourism-related policy documents [12] into existing news-based stock market prediction models. 2. Develop multi-task learning models using several highly correlated cultural tourism-related stocks; this will enable us to identify common characteristics of the assets. 3. Create a portfolio optimization model based on reinforcement learning that is

tailored to the seasonal and policy-driven trends of the cultural tourism industry.

Declaration of Generative AI and AI-assisted Technologies in the Writing Process

During the preparation of this work, the author(s) utilized artificial intelligence (AI) technologies solely for the purpose of linguistic refinement and grammatical correction. Following the application of these tools, the author(s) comprehensively reviewed, verified, and edited the resulting text. The author(s) maintain full accountability and responsibility for the accuracy, integrity, and final content of the publication.

References:

[1] A. B. Siddik, M. S. Forid, L. Yong, A. M. Du, and J. W. Goodell, Artificial Intelligence as a Catalyst for Sustainable Tourism Growth and Economic Cycles, Technological Forecasting and Social Change, Vol. 210, 2025, Article 123875, doi: 10.1016/j.techfore.2024.123875.

[2] Xinhua, China's Digital Economy, Culture, and Tourism Sectors See Expanded Shares of GDP in 2024, China Daily, December 31, 2025, [Online]. chinadaily.com.cn (Accessed Date: March 10, 2026).

[3] F. Shen, S. Huang, W. Zhao, and D. Lan, Forecasting Tourism Stock Index Dynamics: A Multiscale Deep Learning Framework Integrating Emerging Media Data, Expert Systems with Applications, Vol. 298, Part A, 2026, Article 129566, doi: 10.1016/j.eswa.2025.129566.

[4] O. B. Sezer, M. U. Gudelek, and A. M. Ozbayoglu, Financial Time Series Forecasting with Deep Learning: A Systematic Literature Review: 2005-2019, Applied Soft Computing, Vol. 90, 2020, Article 106181, doi: 10.1016/j.asoc.2020.106181.

[5] S. Sun, L. Zhang, and C. Yu, CTL: A Stock Price Index Prediction Network Based on a Hybrid Structure of CEEMDAN, Transformer and LSTM, in Proc. 2024 3rd International Conference on Cyber Security, Artificial Intelligence and Digital Economy (CSAIDE), ACM, 2024, pp. 606-612, doi: 10.1145/3672919.3673025.

[6] J. Xiao, T. Deng, and S. Bi, Comparative Analysis of LSTM, GRU, and Transformer Models for Stock Price Prediction, in Proc. International Conference on Digital Economy, Blockchain and Artificial Intelligence (DEBAI 2024), ACM, 2024, 6 pages, doi: 10.1145/3700058.3700075.

[7] J. C. S. Nunez, J. A. Gomez-Pulido, and R. Robina-Ramirez, Machine Learning Applied to Tourism: A Systematic Review, WIREs Data Mining and Knowledge Discovery, Vol. 14, No. 5, 2024, Article e1549, doi: 10.1002/widm.1549.

[8] M. Hu, M. Li, Y. Chen, and H. Liu, Tourism Forecasting by Mixed-Frequency Machine Learning, Tourism Management, Vol. 106, 2025, Article 105004, doi: 10.1016/j.tourman.2024.105004.

[9] A. Dimitriadou, P. Gogas, and T. Papadimitriou, Tourism and Uncertainty: A Machine Learning Approach, Current Issues in Tourism, Vol. 28, No. 14, 2025, pp. 2278-2298, doi: 10.1080/13683500.2024.2370380.

[10] M. R. Mahmud, M. R. Hoque, M. M. Ali, S. Ferdausi, and K. Fatema, Machine Learning-Powered Financial Forecasting in the U.S. Tourism Industry: Predicting Market Trends and Consumer Spending with Artificial Intelligence, Journal of Computer Science and Technology Studies, Vol. 7, No. 2, 2025, pp.13-22, doi: 10.32996/jcsts.2025.7.2.2.

[11] J. Zhang and M. Gao, The Application of Deep Learning in Economic Analysis and Marketing Strategy Formulation in the Tourism Industry, PLOS ONE, Vol. 20, No. 6, 2025, Article e0321992, doi: 10.1371/journal.pone.0321992.

[12] Q. Chen, Stock Price Prediction with LLM-Guided Market Movement Signals and Transformer Model, FinTech and Sustainable Innovation, Vol. 1, 2025, pp. 1-7, doi: 10.47852/bonviewfsi52025703.

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

Wenrui Jiang designed the study, created the hybrid LSTM-Transformer model, and wrote most of the paper. Zijian Zeng collected the data, engineered the features, and conducted the numerical studies.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding.

Conflicts of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Creative Commons Attribution License 4.0 (Attribution 4.0 International, CC BY 4.0)

This article is published under the terms of the Creative Commons Attribution License 4.0
https://creativecommons.org/licenses/by/4.0/deed.en_US

APPENDIX

Table 2. Per-stock results on the 2023-2024 test set (random walk, LSTM, HLT). RMSE and MAE are expressed in Chinese Yuan (CNY) and cannot be compared between different stocks; sMAPE, MASE, and DA are scale-free.

Stock	Model	RMSE	MAE	sMAPE(%)	MASE	DA(%)
300144.SZ	Random Walk	0.482	0.351	2.83	1.000	50.6
	LSTM	0.371	0.272	2.55	0.775	68.9
	HLT	0.298	0.218	2.07	0.621	72.1
000430.SZ	Random Walk	0.214	0.158	3.12	1.000	49.7
	LSTM	0.179	0.131	2.87	0.829	67.5
	HLT	0.148	0.109	2.61	0.690	70.4
002033.SZ	Random Walk	0.305	0.224	2.74	1.000	51.0
	LSTM	0.241	0.176	2.46	0.786	69.3
	HLT	0.196	0.143	2.15	0.638	72.6
601888.SH	Random Walk	7.395	5.412	2.14	1.000	50.2
	LSTM	7.401	5.470	2.51	1.011	70.7
	HLT	5.450	4.018	1.73	0.742	74.0

Source: created by the authors

Table 3. Diebold-Mariano test: HLT vs. LSTM (quadratic loss, HLN corrected; adjusted p values are based upon Holm-Bonferroni method).

Stock	DM stat	raw p	adj. p	Sig.
300144.SZ	2.94	0.003	0.012	*
000430.SZ	3.11	0.002	0.008	**
002033.SZ	2.18	0.029	0.058	+
601888.SH	1.62	0.105	0.105	n.s.

Source: created by the authors

Table 4. Pesaran-Timmermann test of directional predictability for HLT.

Stock	PT stat	p-value	Sig.
300144.SZ	6.81	<0.001	**
000430.SZ	5.97	<0.001	**
002033.SZ	7.04	<0.001	**
601888.SH	7.62	<0.001	**

Source: created by the authors

Table 5. Ablation: effect of tourism demand proxies (sMAPE, %).

Model / Stock	Full	No BSI	Gain
GRU	2.85	3.31	0.46
LSTM	2.60	3.08	0.48
Transformer	2.36	2.77	0.41
HLT	2.14	2.59	0.45
HLT, 300144.SZ	2.07	2.66	0.59
HLT, 000430.SZ	2.61	3.04	0.43
HLT, 002033.SZ	2.15	2.71	0.56
HLT, 601888.SH	1.73	1.95	0.22

Source: created by the authors

Table 6. Performance of the Long Only Backtest from January 2023 through December 2024, Including Net of All-In Round Trip Cost of 12 Basis Points (T+1 Open Entry, T+1 Close Exit).

Strategy	Ann. Ret. (%)	Sharpe	MDD (%)	Turn.
Buy & Hold (EW)	-3.4	-0.18	28.6	0.00
LSTM signal	6.7	0.62	17.4	0.71
Transformer	9.3	0.84	15.2	0.74
HLT	12.6	1.18	13.9	0.69

Source: created by the authors