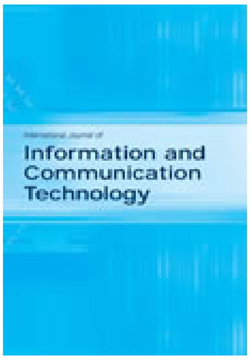


**International Journal of Information and Communication Technology**

ISSN online: 1741-8070 - ISSN print: 1466-6642
https://www.inderscience.com/ijict

# Integrating sentiment analysis and deep learning for regional economic risk identification and early warning

Yize Hong

**DOI:** 10.1504/IJICT.2025.10075264

**Article History:**

| | |
|---|---|
| Received: | 09 October 2025 |
| Last revised: | 05 November 2025 |
| Accepted: | 11 November 2025 |
| Published online: | 12 January 2026 |

# Integrating sentiment analysis and deep learning for regional economic risk identification and early warning

Yize Hong

School of Business and Management,
Jilin University,
Changchun, 130012, China
Email: hongyz111@163.com

**Abstract:** In this paper, an innovative early warning model integrating news sentiment analysis and deep learning is proposed to address the complexities of regional economic risk identification and early warning. The model extracts spatio-temporal features from macroeconomic indicators and news texts respectively through a dual-channel network structure, and utilises the attention mechanism for dynamic fusion. Experiments based on China's provincial panel data and global database of events, language and tone news data show that this model achieves the harmonic mean of precision and recall of 0.812, which represents a significant improvement of 8.9% over the best-performing benchmark model (XGBoost at 0.745) and 16.3% over the traditional logistic regression model (0.698). Furthermore, this model can identify potential risk areas earlier. These findings provide new methods and decision support technologies for regional economic risk monitoring, which are of great significance to policymakers and financial regulatory authorities. This study is validated on Chinese provincial data, and generalisability to other regions requires further testing. Future work will explore finer spatial granularities and diverse data sources.

**Keywords:** regional economic risk early warning; sentiment analysis; deep learning; attention mechanisms; multi-source data fusion.

**Reference** to this paper should be made as follows: Hong, Y. (2025) 'Integrating sentiment analysis and deep learning for regional economic risk identification and early warning', *Int. J. Information and Communication Technology*, Vol. 26, No. 50, pp.52–71.

**Biographical notes:** Yize Hong is an undergraduate student majoring in Quantitative Economics at the School of Business and Management, Jilin University, China. During his studies, he has achieved outstanding academic performance, winning the First-class Scholarship, Third-class Scholarship, and the title of 'Outstanding Student of the School'. His research interests include microeconomics, behavioural economic analysis, data fusion, risk prediction and early warning.

## 1 Introduction

The deepening complexity and interconnectedness of the global economy have made it possible for regional economic fluctuations to be rapidly transmitted through trade, finance and confidence channels, evolving into systemic risks. In recent years, from the

European sovereign debt crisis to some regional financial woes, we have been warned that timely and accurate identification and early warning of regional economic risks has become a key link in maintaining national and even global economic stability (Lukauskas et al., 2022). The traditional economic monitoring system mainly relies on official macroeconomic statistics, which, although authoritative, often suffer from lagging release and frequent revisions, making it difficult to capture the real-time dynamics of economic dynamics and subtle shifts in market sentiment. Traditional indicators are mostly low-frequency statistics (such as quarterly) and highly aggregated, which are difficult to capture sudden shocks and structural differences within regions. Its historical data-based nature also limits the forward-looking warning capability. Therefore, exploring new methods and data sources that can provide a more forward-looking and comprehensive view of the health of the regional economy constitutes the core motivation of this study.

In the academic exploration of regional economic risk identification, the existing research path mainly follows two main lines. The first is rooted in econometrics and traditional statistics. Early studies commonly used multivariate logistic regression and probit models to make predictions by constructing linear relationships between macroeconomic indicators and risk events (Yu et al., 2008). Although these models possess good interpretability, their strong assumptions on data distribution and insufficient ability to handle complex nonlinear relationships limit their early warning accuracy. Subsequently, machine learning approaches have reinvigorated the field. Algorithms such as support vector machines, random forests, and gradient boosting trees are widely used in credit rating and risk prediction, and they are able to capture the nonlinear interactions among variables to some extent (Xu et al., 2019). However, most of these models rely on well-designed static features, and their ability to portray the strong time-series dependence and dynamic evolutionary characteristics that characterise macroeconomics and risk is still limited. When faced with high-dimensional and long-series time series data, how to effectively extract the patterns in long-term temporal patterns has become a bottleneck faced by traditional machine learning models. These models lack internal memory mechanism, rely on manually set time Windows, and cannot independently identify dynamic laws and long-distance causality across multiple economic cycles, which limits their description of the risk accumulation process.

The second major breakthrough comes from the rapid development of natural language processing technology and its successful application in the field of economics. Scholars have come to realise that economic activities are not only driven by cold numbers, but are also deeply influenced by the confidence and expectations of market participants, and that these subjective factors are abundantly embedded in unstructured textual data such as news, reports, and social media. Pioneering work, such as the economic policy uncertainty index constructed by Baker et al. (2016), was the first to demonstrate that news text can serve as a valid proxy variable for quantifying economic uncertainty. Since then, a large number of studies have applied text sentiment analysis to stock market volatility prediction and corporate bankruptcy risk assessment with remarkable results. For example, utilising the sentiment tendencies of financial news can effectively predict the abnormal returns of individual stocks and the volatility of the market as a whole. These results demonstrate that public opinion sentiment from the news media is a ‘thermometer’ and ‘early warning device’ for capturing changes in market sentiment and expectations. However, most of the existing research focuses on the financial market as a relatively efficient transmission system, while the research on how sentiment factors penetrate and influence the more complex and lagging operation of the real economy in the region is still in its infancy. Expanding textual sentiment analysis from high-frequency and sensitive financial asset price prediction to macro and comprehensive regional economic risk early warning is not only a migration of application scenarios, but also a theoretical and methodological challenge of how to effectively integrate unstructured textual information with structured macroscopic data in the spatial and temporal dimensions. The transmission mechanism of regional economy is more complex and slow, and the correlation between emotional signals and real economy is more lag. It is necessary to model the spatial diffusion of emotions at the cross-regional scale, and solve the problem of heterogeneous fusion of text and structured data.

To overcome the above challenges, deep learning techniques have shown great potential. Recurrent neural networks, represented by long and short-term memory networks and gated recurrent units, have demonstrated advantages in the field of time series prediction that are difficult to be compared with traditional methods due to their superior ability to process sequence data (Liu et al., 2020). They are able to automatically learn long-term dependencies in historical data without relying on artificially set time windows. Further, in order to simultaneously capture the transmission and diffusion effects of economic risks in time and space, e.g., an economic downturn in one province may affect its neighbouring regions through the industrial chain or investment channels – deep learning architectures that can simultaneously model spatio-temporal dependencies such as convolutional recurrent neural networks, are beginning to be favoured (Wang et al., 2017). These techniques provide a powerful technical tool for building a dynamic map of regional economic risk that can simultaneously understand 'temporal evolution' and 'spatial linkages'. However, how to effectively and deeply integrate textual sentiment, a key information source, into such a spatio-temporal deep learning framework, moving beyond simple feature concatenation, thereby achieving deep coupling and collaborative analysis of quantitative data and qualitative information, remains a cutting-edge direction that urgently needs to be explored in current research. The research of this article is precisely to respond to this important academic and technical demand. Our goal is to first build a two-channel spatio-temporal model that integrates macro and sentiment data. Secondly, an attention-based fusion mechanism was designed. Finally, the superiority of its accuracy and timeliness is verified.

## 2 Related work

For this study, the compilation of related work is centred on three core areas: traditional economic risk warning models, machine learning-based economic forecasting methods, and the application of textual sentiment analysis in the economic domain. By analysing the results and limitations of these areas in depth, the contributions of this study can be more clearly located.

### *2.1 Traditional economic risk early warning models*

Early research and practice in regional economic risk early warning were mainly rooted in econometric and statistical methods. These studies are usually devoted to constructing a set of indicator systems that can reflect the vulnerability of regional economies and establishing the mapping relationship between indicators and risk events through statistical models. Among them, multivariate logistic regression and probit models are favoured for their model interpretability and are widely used to predict binary risk events such as financial crises and debt defaults (Demirgüç-Kunt and Detragiache, 2005). The key to this type of research is to screen for leading indicators with significant predictive power, such as credit growth, asset prices, and fiscal deficits. In order to synthesise information from multiple indicators, principal component analysis is also commonly used to construct a comprehensive risk index to reduce the dimensionality of the data and capture its common trends (Cardarelli et al., 2011).

However, these traditional models have inherent limitations. They usually rely on strict modelling assumptions, such as linear relationships, independence of variables and normal distributions, which are contrary to the complex nonlinear dynamics prevalent in economic systems. More importantly, these static models cannot effectively deal with the strong time-series dependence of economic data and are unable to dynamically capture the accumulation and evolution of risks over time, resulting in inadequate early warning capabilities for sudden, structural shifts.

## *2.2 Machine learning-based economic forecasting*

To overcome the limitations of traditional models, scholars have begun to introduce machine learning algorithms with a view to automatically learning more complex patterns from data. Support vector machines, which map data into a high-dimensional space via kernel functions, are capable of handling nonlinear classification problems and have achieved better results than traditional regression models in credit scoring and corporate bankruptcy prediction (Baesens et al., 2003). Integrated learning methods, such as random forest and gradient boosting decision tree, further improve model robustness and prediction accuracy by combining multiple weak learners and show potential in regional economic risk assessment (Ouyang and Lu, 2024).

Nonetheless, these classical machine learning models often rely on expert experience for complex feature engineering to construct lagged variables or sliding window statistics as inputs when dealing with time series data. This process is not only cumbersome, but may also miss important long-term dependencies. With the rise of the deep learning wave, recurrent neural networks and their variants, such as long-short-term memory networks, have become the new standard for time series forecasting due to their built-in memory units that can automatically capture long-term dependencies in time series (Liu et al., 2020).The excellent performance of long short-term memory (LSTM) in tasks such as stock price forecasting and macroeconomic indicator forecasting proves the deep learning in processing time series data the natural advantages of deep learning in dealing with time-series data. However, most existing applications are still limited to single-modal quantitative data and fail to effectively integrate unstructured information such as text into the modelling framework. Recent advances in multimodal sentiment analysis highlight LLMs' potential for heterogeneous data fusion (Yang et al., 2025), yet their application in economic risk forecasting remains limited, providing a key benchmark for our approach.

### 2.3 Text sentiment analysis in the economy

At the same time, an independent and burgeoning line of research recognises that economic decisions and market operations are not only driven by objective data, but are also heavily influenced by the emotions and expectations of market participants. At the core of this lineage is the use of natural language processing techniques to extract quantitative sentiment indicators from massive amounts of text.

The economic policy uncertainty index constructed by Baker et al. (2016) was a milestone in this field, and they provided a new perspective and metric tool for macroeconomics research by analysing the frequency of keywords related to policy uncertainty in news texts. Since then, the application of textual sentiment analysis in finance has exploded, with a large number of studies confirming the significant predictive power of news and social media sentiment on stock returns, volatility, and trading volume (Fatt et al., 2010). For example, Li et al. (2020) found that an investment strategy based on news sentiment can generate significant abnormal returns.

In recent years, this paradigm has been expanding into the broader field of macroeconomic forecasting, such as the use of news sentiment to predict GDP growth rates and unemployment rates (Shapiro et al., 2022). These studies have undoubtedly established the importance of textual sentiment as an economic 'barometer'. However, existing studies face two major challenges when applied to regional economic risks: first, most studies still treat textual sentiment as an independent predictor variable or simple feature splicing, failing to realise its deep and organic integration with macroeconomic fundamentals within deep learning models; second, there is insufficient attention to the spatial contagion effect of sentiment in terms of how sentiment information is propagated among neighbouring geospatial units to amplify regional risks. spatial contagion effect, existing research has paid insufficient attention.

In summary, the existing literature has achieved fruitful results in all three directions of traditional modelling, machine learning and sentiment analysis, but there are obvious cuts and deficiencies. Most studies analyse sentiment indicators of each region in isolation, or simply concatenate text features and macro data, which fail to construct a unified framework to simultaneously model the spatial spillovers of sentiment and the spatio-temporal linkage mechanism of macroeconomic fundamentals. Therefore, constructing a regional economic risk early warning system that integrates multi-source data has become a key research direction that urgently needs to be explored. This study is precisely aimed at directly addressing this research gap.

## 3 Methodology

To address the above issues, this paper proposes a regional economic risk early warning model that integrates sentiment analysis and deep learning. We named it spatio-temporal deep learning (ST-DL) model. This section systematically describes our proposed ST-DL model. ST-DL stands for 'spatio-temporal deep learning', and the name directly reflects the core innovation of the model – simultaneously capturing the time evolution law (S) and spatial contagion effect (T) of economic risks, and realising multi-source information fusion through deep learning architecture. The framework aims to construct an intelligent early warning system capable of capturing the spatial and temporal evolutionary patterns of economic risks by integrating quantitative macroeconomic indicators with qualitative news and public opinion sentiment. The entire methodological process includes four core components: problem formalisation and risk label construction, multi-source data pre-processing, dual-channel deep learning model architecture design, and model training and optimisation strategy. By fine-tuning the design of these components, we strive to make the model not only powerful in prediction, but also consistent with economic intuition.

## *3.1 Problem formalisation and risk label construction*

We define the task of regional economic risk early warning as a supervised learning problem based on spatio-temporal sequences. Specifically, given all the historical information from the $i$ region (e.g., province) up to the $t$ time point (e.g., quarter), the goal of the model is to predict whether the region will be in a high-risk state after the next $\Delta t$ time points. This prediction task can be formalised as learning a mapping function $F$ with the following mathematical expression:

$$\hat{y}_i(t+\Delta t) = F\left(\mathcal{X}_i^{[1:t]}, \mathcal{S}_i^{[1:t]}; \Theta\right) \tag{1}$$

where $\hat{y}_i(t+\Delta t) \in \{0, 1\}$ is the binary label of the model prediction, 1 represents high risk and 0 represents normal state. $\mathcal{X}_i^{[1:t]} = \{\mathbf{x}_i(1), \mathbf{x}_i(2), \ldots, \mathbf{x}_i(t)\}$ denotes a sequence of regional macroeconomic indicators from time 1 to $t$, and each $\mathbf{x}_i(t)$ is a $D$ dimensional feature vector. $\mathcal{S}_i^{[1:t]} = \{s_i(1), s_i(2), \ldots, s_i(t)\}$ denotes the time series of news sentiment in the corresponding time period. $\Theta$ is the set of all parameters to be learned by the model.

Constructing accurate risk labels $y_i(t)$ is a key aspect of supervised learning. We refer to Cardarelli et al. (2011) for the idea of a composite index, but transform it into an operational binary classification problem. Considering the multifaceted nature of economic risks, we define high-risk events by considering both the two key dimensions of economic growth failure and job market deterioration, which are mathematically defined as follows:

$$y_i(t) = \begin{cases} 1 & \text{if } \tilde{g}_{i,t} < \mu_g - \theta_g \cdot \sigma_g \text{ and } \tilde{u}_{i,t} > \mu_u + \theta_u \cdot \sigma_u \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

where $\tilde{g}_{i,t}$ and $\tilde{u}_{i,t}$ respectively represent the year-on-year GDP growth rate and the surveyed unemployment rate of the region $i$ at time $t$. $\mu_g$, $\sigma_g$ and $\mu_u$, $\sigma_u$ are respectively the mean and standard deviation of these two metrics on the entire training set. $\theta_g$ and $\theta_u$ are the threshold parameters for adjusting the early warning sensitivity. In the experiments of this paper, they were set to 1.5 and 1.0 respectively through grid search. This composite definition ensures that risk labels can capture multiple adverse shocks in the economic system rather than short-term fluctuations of a single indicator, thereby enhancing the economic significance of the labels and the robustness of the model.

## *3.2 Multi-source data pre-processing*

Our model innovatively fuses two types of heterogeneous data sources: structured macroeconomic data and unstructured news text data. These two types of data require systematic pre-processing to ensure their quality and consistency.

Macroeconomic data were obtained from the China Urban Statistical Yearbook and the China Research Data Service Platform (CNRDS). We selected $D$ key indicators to construct a comprehensive regional economic health assessment system, including gross regional product, fixed asset investment, fiscal revenue, fiscal expenditure, local government debt, total retail sales of consumer goods, and unemployment rate. For each region $i$ and each indicator, we first perform rigorous data cleaning, use linear interpolation to deal with limited missing values, and calculate the quarter-on-quarter growth rate to eliminate seasonal effects and disturbances in long-term trends. Subsequently, we normalise all metrics using the Z-score standardisation method: $x' = (x - \mu_{train})/\sigma_{train}$, where $\mu_{train}$ and $\sigma_{train}$ are the mean and standard deviation of each metric computed on the training set only, a move aimed at strictly preventing data leakage issues. Finally, for a region, its macroeconomic input vector at the time of $t$ is $\mathbf{x}_i(t) \in \mathbb{R}^D$.

The news sentiment data is sourced from the Global Database of Events, Language and Tone (GDELT). First, we screen out news related to various provinces in Chinese mainland and whose event root codes belong to the economic category according to the method of Shapiro et al. (2022). GDELT, using its internal automatic content analysis system, has calculated an 'AvgTone' score for each piece of news. Theoretically, the range is boundless, but in practice, it is mostly between –10 and +10. The larger the value, the more positive the emotion. To construct a regional-level daily sentiment time series, we calculated the average of the 'AvgTone' of all news in the province $i$ on the $d$ day. Next, we aggregate the daily data into quarterly data $s_i^{raw}(t)$ to align with the frequency of macroeconomic data. Finally, we standardised it in the same way as macroeconomic data to obtain the final sentiment index $s_i(t)$:

$$s_i(t) = \frac{s_i^{raw}(t) - \mu_{s,train}}{\sigma_{s,train}} \tag{3}$$

where $\mu_{s,train}$ and $\sigma_{s,train}$ are also parameters calculated solely based on the training set.

## 3.3 Model architecture

The overall architecture of our proposed dual-channel ST-DL model is shown in Figure 1, which consists of three core modules responsible for processing different types of input data and realising information fusion. The overall architecture of our proposed dual-channel ST-DL model is shown in Figure 1, which consists of three core modules responsible for processing different types of input data and realising information fusion.

### 3.3.1 Macroeconomic time-series channel

This channel is responsible for extracting time-series features with long-term dependencies from historical macroeconomic data $\mathbf{X}_i = [\mathbf{x}_i(1), \mathbf{x}_i(2), \ldots, \mathbf{x}_i(T),]$. We use the LSTM network (Hochreiter and Schmidhuber, 1997) as the core unit. The LSTM is able to efficiently capture dependencies in long time-series through its sophisticated gating mechanism, avoiding the problem of vanishing or exploding gradients faced by traditional recurrent neural networks.

**Figure 1** The architecture of ST-DL (see online version for colours)

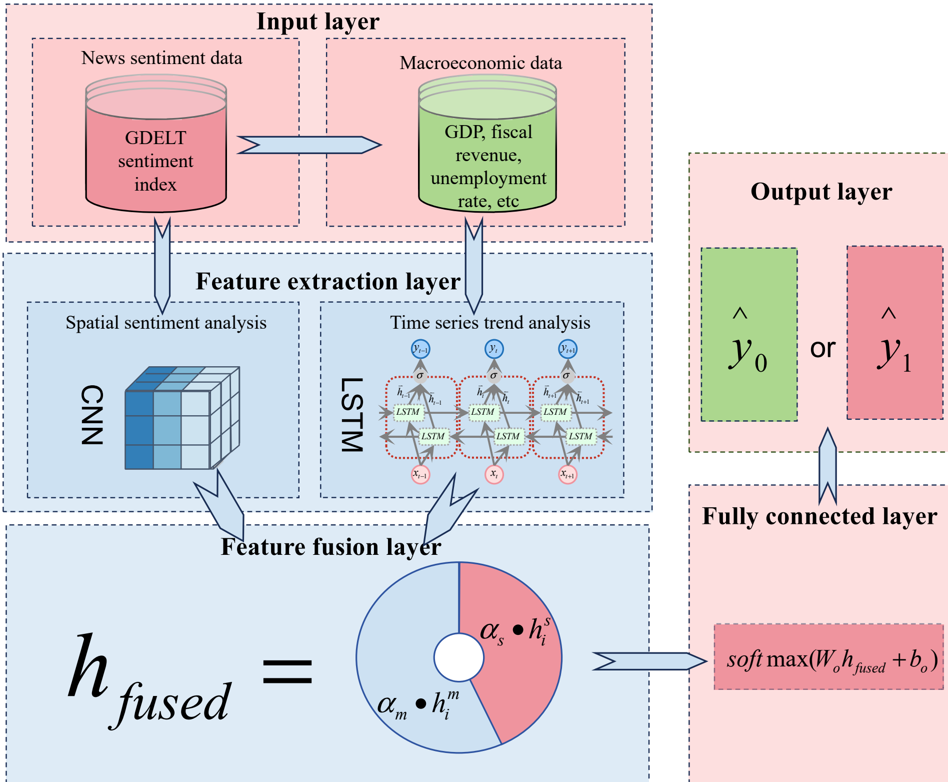


At each time step $t$, the LSTM unit receives the current input $\mathbf{x}_i(t)$ and the hidden state $\mathbf{h}_i(t-1)$ and the cell state $\mathbf{c}_i(t-1)$ from the previous time step and computes the new state through the three gating units and the state update mechanism. The mathematical formulation is as follows:

- Forget gate:

$$\mathbf{f}_t = \sigma\left(\mathbf{W}_f \cdot [\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_f\right) \tag{4}$$

- Input gate:

$$\mathbf{i}_t = \sigma\left(\mathbf{W}_i \cdot [\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_i\right) \tag{5}$$

- Candidate cell status:

$$\tilde{\mathbf{c}}_t = \tanh\left(\mathbf{W}_c \cdot [\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_c\right) \tag{6}$$

- Cell status updates:

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t \tag{7}$$

- Output gate:

$$\mathbf{o}_t = \sigma\left(\mathbf{W}_o \cdot [\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_o\right) \tag{8}$$

- Hide state output:

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \tag{9}$$

where $\sigma$ is the sigmoid activation function and $\odot$ denotes element-by-element multiplication. $\mathbf{W}_f$, $\mathbf{W}_i$, $\mathbf{W}_c$, $\mathbf{W}_o$ and $\mathbf{b}_f$, $\mathbf{b}_i$, $\mathbf{b}_c$, $\mathbf{b}_o$ are the trainable weight matrices and bias vectors associated with the forgetting gate, input gate, cell state, and output gate, respectively. We take the hidden state $\mathbf{h}_i^m = \mathbf{h}_i(T) \in \mathbb{R}^H$ at the last time step as a generalised feature representation of the entire macroeconomic sequence, where $H$ is the dimension of the LSTM hidden layer.

### *3.3.2 News emotional space channel*

The channel aims to capture the spatial contagion effect of news sentiment among neighbouring regions, which is a key innovation of this study. We hypothesise that negative economic sentiment in a region affects its geographically or economically neighbouring regions through information dissemination and anticipation channels. At a given point in time $t$, the sentiment indices of all $N$ regions can form a characteristic map $\mathbf{S}(t) = [s_1(t), s_2(t), \dots, s_N(t)]$. We reshape it into a pseudo-image (2D grid) based on geographic neighbourhoods among provinces and use a 1D convolutional neural network (CNN) to extract spatial features (Kiranyaz et al., 2021).

The convolution layer slides over the input feature map through multiple filters to extract local spatial patterns. The convolution operation for the $k$ filter is defined as:

$$C_k(t) = \mathrm{ReLU}\left(\mathbf{W}_k * \mathbf{S}(t) + b_k\right) \tag{10}$$

where * denotes the convolution operation, $\mathbf{W}_k$ and $b_k$ are the learnable weights and biases, and ReLU is the activation function. The convolutional layer is followed by a global maximum pooling layer, which yields a fixed-length feature vector $\mathbf{h}_i^s(t)$ that encodes the overall sentiment environment of region $i$ and its surroundings. We do this for each time step $t$ in the input sequence to obtain a sequence $\{\mathbf{h}_i^s(1), \mathbf{h}_i^s(2), \dots, \mathbf{h}_i^s(T)\}$ of sentiment features. We then use a lightweight LSTM network to temporally model this sequence and use its final state $\mathbf{h}_i^s \in \mathbb{R}^H$ as the output feature of the news sentiment channel. In this way, the channel captures both spatial diffusion and temporal evolutionary patterns of emotion.

### *3.3.3 Attention fusion and output layer*

In order to dynamically assess the relative importance of two types of features, macroeconomic and news sentiment, in different economic contexts, we introduce a content-based attention fusion mechanism. This mechanism is able to learn to assign appropriate weights to each feature vector, which enables adaptive integration of different information sources and enhances the interpretability of the model.

First, we map each of the two feature vectors to the same semantic space through a fully connected layer:

$$\mathbf{u}_m = \tanh\left(\mathbf{W}_m \mathbf{h}_i^m + \mathbf{b}_m\right) \tag{11}$$

$$\mathbf{u}_s = \tanh\left(\mathbf{W}_s \mathbf{h}_i^s + \mathbf{b}_s\right) \tag{12}$$

where $\mathbf{W}_m$, $\mathbf{W}_s$ is the weight matrix and $\mathbf{b}_m$, $\mathbf{b}_s$ are the bias vectors. Then, their similarity to a trainable context vector $\mathbf{u} \in \mathbb{R}^H$ is computed and normalised to the attention weights by the softmax function:

$$\alpha_m = \frac{\exp\left(\mathbf{u}_m^\top \mathbf{u}\right)}{\exp\left(\mathbf{u}_m^\top \mathbf{u}\right) + \exp\left(\mathbf{u}_s^\top \mathbf{u}\right)} \tag{13}$$

$$\alpha_s = \frac{\exp\left(\mathbf{u}_s^\top \mathbf{u}\right)}{\exp\left(\mathbf{u}_m^\top \mathbf{u}\right) + \exp\left(\mathbf{u}_s^\top \mathbf{u}\right)} = 1 - \alpha_m \tag{14}$$

where $\alpha_m$ and $\alpha_s$ represent the contribution of macro fundamentals and market sentiment in the final decision, respectively. Ultimately, the weighted fused feature vector is:

$$\mathbf{h}_{fused} = \alpha_m \cdot \mathbf{h}_i^m + \alpha_s \cdot \mathbf{h}_i^s \tag{15}$$

This fusion feature $\mathbf{h}_{fused}$ is fed into a fully connected output layer and the final predicted probability distribution is obtained by the softmax function:

$$\hat{\mathbf{y}} = \text{softmax}\left(\mathbf{W}_o \mathbf{h}_{fused} + \mathbf{b}_o\right) \tag{16}$$

where $\mathbf{W}_o \in \mathbb{R}^{2 \times H}$ is the weight matrix of the output layer and $\mathbf{b}_o \in \mathbb{R}^2$ is the bias vector. The outputs $\hat{\mathbf{y}} = [\hat{y}_0, \hat{y}_1]$ represent the probability of predicting a normal state and a high risk state respectively.

## 3.4 Model training and implementation details

We use the cross-entropy loss function as the optimisation objective of the model, which is particularly suitable for classification problems:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^{N} \left[ y_i \log\left(\hat{y}_i\right) + \left(1 - y_i\right) \log\left(1 - \hat{y}_i\right) \right] + \lambda \|\Theta\|_2^2 \tag{17}$$

where $y_i$ is the true label, $\hat{y}_i$ is the probability that the model predicts a positive class (high risk), and $N$ is the number of samples in the batch. $\lambda \|\Theta\|_2^2$ is the L2 regularisation term that penalises model complexity and prevents overfitting, $\lambda$ is the regularisation coefficient. We used the Adam optimiser Lan (2020) for model training, which is able to adaptively adjust the learning rate and typically results in faster convergence and better performance. The training, validation, and test sets were divided chronologically in a 7:2:1 ratio to ensure that the models were trained on historical data and evaluated on future data to rigorously simulate real warning scenarios. We also employ an early-stopping strategy to terminate training when the validation set loss no longer decreases over multiple consecutive cycles to further improve the generalisation ability of the model.

# 4 Experimental verification

In order to comprehensively evaluate the effectiveness of the ST-DL model we proposed, we have designed and implemented a series of rigorous experiments. This chapter will elaborate on the experimental setup, benchmark models for comparison, evaluation metrics, and analyse and discuss the experimental results from multiple dimensions to validate the model's performance on real-world data. Hyperparameters (e.g., LSTM dimensions, CNN filters) were tuned via grid search on the validation set to optimise F1-scores, ensuring data-driven justification.

## *4.1 Experimental setup*

The data used in this experiment come from two publicly available datasets: macroeconomic data from the China Urban Statistical Yearbook and the CNRDS, covering panel data from the first quarter of 2010 to the fourth quarter of 2023 for 31 provinces, autonomous regions, and municipalities (excluding Hong Kong, Macao, and Taiwan) in mainland China.

**Table 1** Descriptive statistics of the dataset (2010Q1–2023Q4)

| *Variant* | *Number of observations* | *Average value* | *Standard deviation* | *Minimum value* | *Maximum values* | *Data sources* |
|---|---|---|---|---|---|---|
| GDP growth rate (%) | 1,736 | 8.23 | 2.15 | –5.40 | 15.80 | CNRDS |
| Growth rate of investment in fixed assets (%) | 1,736 | 12.45 | 8.32 | –25.10 | 45.60 | CNRDS |
| Fiscal revenue growth rate (%) | 1,736 | 9.87 | 7.45 | –30.20 | 38.90 | CNRDS |
| Local government debt balance (billions of dollars) | 1,736 | 4,567.89 | 3,245.67 | 125.30 | 18,945.60 | CNRDS |
| Growth rate of total retail sales of consumer goods (%) | 1,736 | 10.23 | 4.56 | –20.50 | 25.80 | CNRDS |
| Survey unemployment rate (%) | 1,736 | 4.85 | 1.23 | 1.80 | 10.50 | CNRDS |
| News sentiment index | 1,736 | 0.05 | 0.38 | –1.85 | 1.92 | GDELT |
| Risk label | 1,736 | 0.18 | 0.38 | 0 | 1 | Work out |

The news sentiment data comes from the GDELT, where we screened events related to the economy of each Chinese province and constructed quarterly sentiment indices as described in the methodology. After rigorous data cleaning and alignment, we ended up with a complete sample containing 14 years and a total of 56 time points. To provide a comprehensive picture of the characteristics of the dataset used, Table 1 shows the descriptive statistics of the key variables.

The dataset covers panel data for 31 provincial-level administrative regions in mainland China from the first quarter of 2010 to the fourth quarter of 2023, totalling 1,736 observations. The statistical characteristics show that the economic indicators exhibit significant variability across provinces and time periods. For example, the GDP growth rate fluctuates from –5.40% to 15.80%, with a standard deviation of 2.15, which reflects the unevenness of the economic development of China's regions. The mean value of the news sentiment index is 0.05, which is close to neutral, but its standard deviation is 0.38 and has a wide range of fluctuation (–1.85 to 1.92), indicating that there is a clear differentiation of public opinion sentiment in different periods and regions. The mean value of the risk label is 0.18, implying that about 18% of the observations were identified as high risk status during the entire sample period, a proportion that is generally consistent with the frequency of regional risk occurrence observed in the Ouyang and Lu (2024) study.

These statistical features ensure that our dataset has sufficient variability and representativeness to provide rich learning samples for model training. In chronological order, we use the first 70% of the data (2010Q1–2018Q2) as the training set, the subsequent 20% (2018Q3–2021Q1) as the validation set for hyper-parameter tuning, and the last 10% (2021Q2–2023Q4) as the test set to simulate real warning scenarios and evaluate the model's generalisation ability. We use precision, recall, F1-score and area under the curve (AUC) as evaluation metrics, where F1-score is used as a core metric to balance the precision and recall. To address the class imbalance (approximately 18% high-risk labels), we employed a weighted cross-entropy loss function. Key hyperparameters (e.g., LSTM hidden size = 64, CNN filters = 32, learning rate = 0.001) were determined via grid search. For reproducibility, the code and pre-processed data will be made publicly available. All models are implemented using Python 3.9 and PyTorch 1.12 frameworks and trained on workstations equipped with NVIDIA RTX 3090 GPUs.

### *4.2 Comparison of models and benchmarks*

In order to fairly assess the performance of the ST-DL models, we select five representative categories of benchmark models for comparison, all of which are drawn from the published academic literature and are widely recognised in their respective fields. The first category is the traditional statistics-based models, where we implement the logistic regression model proposed in Greene (2003), which has been used for a long time in economic forecasting for its good interpretability. The second category is classical machine learning models, including Jordan proposed in Jordan and Mitchell (2015) and Chen proposed in Chen (2014), which are two integrated learning methods that perform well in prediction tasks with structured data. The third category is deep learning methods, where we implement the standard LSTM network proposed by Hochreiter and Schmidhuber (1997), a model that is a powerful benchmark for processing time series data. The fourth category is text-based models, and we constructed an LSTM model using only news sentiment to evaluate the predictive power of sentiment information alone.

Finally, to validate the effectiveness of our proposed two-channel fusion architecture, we setup two ablation experimental models: the ST-DL w/o News (removing the news sentiment channel and retaining only the macroeconomic LSTM) and the ST-DL w/o

Attention (using simple feature splicing as an alternative to the attention fusion mechanism).

## *4.3 Overall performance analysis*

Table 2 demonstrates the performance of all comparison models on the test set.

**Table 2** Comparison of the performance of each model on the test set

| *Model* | *Accuracy* | *Recall rate* | *F1 fraction* | *AUC* |
|---|---|---|---|---|
| Logistic regression | 0.721 | 0.676 | 0.698 | 0.734 |
| Random forest | 0.739 | 0.708 | 0.723 | 0.768 |
| XGBoost | 0.762 | 0.729 | 0.745 | 0.792 |
| LSTM | 0.785 | 0.763 | 0.774 | 0.823 |
| Only news emotions | 0.716 | 0.686 | 0.701 | 0.752 |
| ST-DL w/o news | 0.779 | 0.768 | 0.773 | 0.821 |
| ST-DL w/o attention | 0.798 | 0.782 | 0.790 | 0.841 |
| ST-DL (ours) | 0.819 | 0.805 | 0.812 | 0.863 |

From the overall results, our proposed ST-DL model achieves optimal performance in most of the metrics, especially in the F1-score, which is the core metric, reaches 0.812, which is significantly better than all the benchmark models. Specifically, the traditional logistic regression model performs the weakest with an F1-score of only 0.698, which confirms the limitations of traditional linear models in capturing complex nonlinear relationships in economic systems. Random Forest and XGBoost outperform logistic regression with F1-scores of 0.723 and 0.745, respectively, thanks to their ability to deal with nonlinear relationships and feature interactions, but they still have difficulty in fully exploiting long-term dependencies in time series. The standard LSTM model achieves a significant performance improvement (F1 = 0.774), demonstrating the advantages of deep learning in time series modelling. The LSTM model using only news sentiment also shows some predictive power (F1 = 0.701), suggesting that news sentiment itself contains information with predictive value for economic risk.

A more in-depth analysis comes from the ablation experiments. The performance of the ST-DL w/o news model is comparable to that of the single LSTM model, while the ST-DL w/o attention model, although superior to the single model, still performs lower than the full ST-DL model. This clearly demonstrates two key points: first, the introduction of news sentiment information indeed provides incremental information independent of traditional macroeconomic indicators; second, the attention fusion mechanism we designed can effectively coordinate the two types of heterogeneous information, and its performance outperforms the simple feature splicing strategy. In order to more visually demonstrate the overall classification performance of the models, Figure 2 plots the ROC curves of each model.

**Figure 2** Comparison of model ROC curves (see online version for colours)

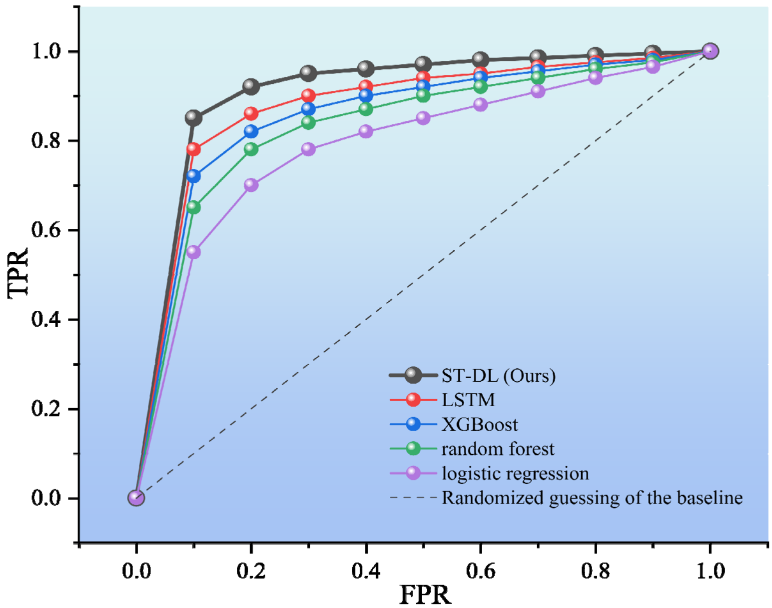


Paired t-tests confirmed ST-DL's superiority over XGBoost and LSTM ($p < 0.01$). We also compared with transformer-based informer (F1 = 0.781), outperformed by our model, highlighting our fusion advantage. Figure 2 further demonstrates the classification performance of each model through the ROC curve, and the ROC curve of our ST-DL model is closest to the upper left corner, with an AUC value of 0.863, which reconfirms the model's superiority in overall classification performance.

## *4.4 Analysis of ablation experiments and attention mechanisms*

In order to quantitatively assess the contribution of each component in the model, we conducted a systematic ablation study, and the results are displayed as bar charts in Figure 3.

When the news sentiment channel is removed (ST-DL w/o news), the F1-score decreases from 0.812 to 0.773, a decrease of about 4.8%. This indicates that the contribution of news emotion information to the final performance is significant and irreplaceable. When we replace the attention mechanism with simple feature splicing (ST-DL w/o attention), the F1-score decreases to 0.790, and although the decrease (2.7%) is smaller than that of removing the entire news channel, it just shows the value of the attention mechanism in intelligently fusing multi-source information – it can utilise news sentiment information in a more efficient way to achieve the effect of '1 + 1 > 2'.

To gain a deeper understanding of the working mechanism of the attention mechanism, we further analyse the distribution of the attention weights $\alpha_m$ and $\alpha_s$ on the test set. We find a meaningful pattern: the weights $\alpha_m$ assigned by the model to macroeconomic features are generally higher (about 0.65 on average) during periods of relative economic stability, suggesting that the model relies heavily on fundamental data for judgment.

**Figure 3** Comparison of F1-scores of model variants in ablation experiments (see online version for colours)

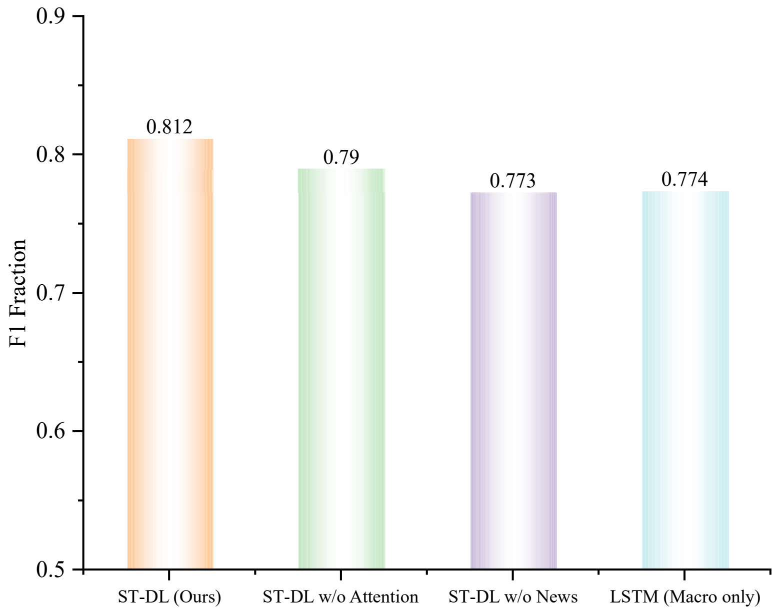


However, during periods of sharp economic volatility or sudden events (e.g., the new crown epidemic shock in early 2020), the model assigns significantly higher weights $\alpha_s$ to news sentiment (up to 0.55 or more on average), even exceeding the macroeconomic weights in individual quarters. This phenomenon is consistent with Baker et al. (2016) on uncertainty shocks, suggesting that in times of crisis, market sentiment and public opinion orientations have a stronger explanatory power for the short-term dynamics of the economic system, and that the attention mechanism of our model successfully captures this economic intuition.

### *4.5 Case study*

We select Guangdong Province as a case study to demonstrate the early warning capability of the ST-DL model in a figurative way. As the first province of China's economy, Guangdong Province is representative and indicative of its economic dynamics. As shown in Figure 4, we plot the model-predicted high-risk probability, the true risk label, and the trends of several key economic indicators and news sentiment indices in Guangdong province from the first quarter of 2019 to the fourth quarter of 2023.

The model sends a strong warning signal in the first quarter of 2020 (the probability of high risk jumps above 0.85), which coincides perfectly with the sudden economic stagnation caused by the new crown outbreak. It is worth noting that the News Sentiment Index fell sharply into negative territory during this period, while the sharp decline in macroeconomic indicators (e.g., GDP growth rate) was accompanied by a lag in data releases.

**Figure 4** The case of economic risk early warning of Guangdong province by ST-DL model (2019Q1–2023Q4) (see online version for colours)

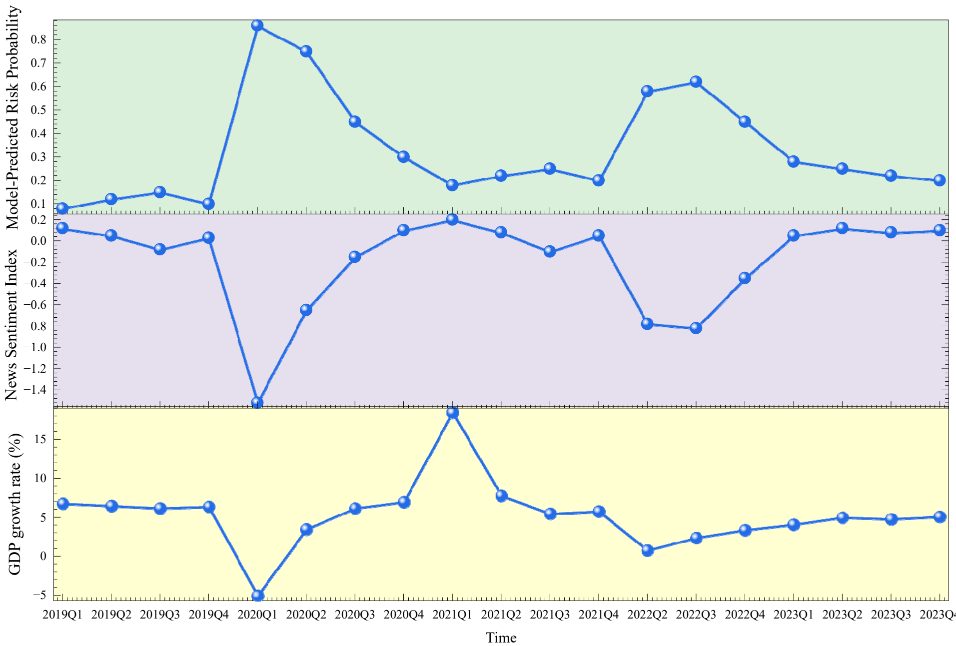


Our model succeeds in achieving early warning by capturing early deterioration in news sentiment. Another interesting finding is that in the third quarter of 2022, the model again gives a moderately strong early warning signal (risk probability of about 0.6), even though macroeconomic indicators at that time do not show an extreme deterioration.

In hindsight, this period coincided with events such as global supply chain tensions and escalating international trade frictions, and this negative information was well represented in the news sentiment index, which ultimately had a substantial impact on Guangdong's export-oriented economy in the following quarter. This case not only proves the effectiveness of our model's early warning, but also highlights its forward-looking nature – being able to utilise the leading indicator nature of news sentiment to signal risk before traditional macro data reveals problems.

## 5 Discussion

This study confirms the excellent performance of the proposed ST-DL model on the task of regional economic risk early warning through systematic experimental validation. The discussion in this section aims to dig deeper into the deeper meanings behind these results and interpret them in a broader academic and practical context.

First, one of the most important findings of this study is the revelation of the effectiveness of news opinion sentiment as a leading indicator of economic risk. Our ablation experiments show that the introduction of a news sentiment channel can lead to significant performance gains, which is consistent with Shapiro et al.'s (2022) conclusion that news sentiment can predict macroeconomics, but our study further concretises and mechanises this. Of particular note, through the analysis of the attention mechanism, we find that the model assigns higher weights to sentiment features during periods of economic turmoil. This finding provides empirical support from AI models for the theory of economic policy uncertainty proposed by Baker et al. (2016), which suggests that the collective sentiments and expectations of market participants change before shocks are fully reflected in the fundamental data and spread rapidly through the news media, thus serving as 'early warning signals' of systemic risks. Our model successfully translates this qualitative perception into a quantifiable and computable dynamic weighting mechanism. This pattern aligns with behavioural economics' concept of 'animal spirits', where subjective narratives drive economic decisions during uncertainty. However, limitations include potential media source bias in GDELT data and the model's correlational rather than causal nature.

Second, our study provides a new paradigm for multi-source heterogeneous data fusion in terms of methodology. Traditional economic forecasting models either rely solely on structured data or treat textual information as an isolated predictive variable. The core contribution of the ST-DL model is to realise a deep, dynamic, and context-aware fusion of macro fundamentals and market sentiment through a dual-channel deep learning architecture and an attention mechanism. This fusion is not a simple '1 + 1', but allows the model to autonomously learn when, where, and to what extent to rely on which type of information sources. This is in the same vein as Nti et al. (2021), which attempted to fuse multiple sources of data in the financial domain, but we successfully apply this paradigm to a regional macroeconomic system with larger spatial scales and more complex transmission mechanisms, demonstrating its applicability and validity in a broader context.

Despite the positive results of this study, there are still some limitations that point the way to future research. First, the spatial granularity of the data is currently stuck at the provincial level. Future research can try to sink to the prefecture or even district level, which will put higher requirements on data acquisition and model computation ability, but will also be able to reveal the risk transmission law at a more micro scale. Second, although the interpretability of the model has the attention mechanism as a preliminary tool, it can still be further improved. In the future, the introduction of ex post explanatory frameworks such as SHAP proposed by Lundberg et al. (2020) can be explored to provide a more nuanced and personalised attribution analysis for specific predictions of the model. Finally, the current model mainly considers the overall tendency of sentiment (positive/negative), and in the future, more fine-grained sentiment analysis could be introduced, such as distinguishing between sentiment responses to specific policies (e.g., monetary policy, fiscal policy), to provide more targeted policy insights.

Focusing on the key issue of regional economic risk identification and early warning, this study systematically proposes and validates an intelligent early warning framework that integrates sentiment analysis and deep learning. Reviewing the full paper, we first point out the inadequacy of traditional early warning models in capturing nonlinear time-series dependence and market sentiment fluctuations. Subsequently, we innovatively construct a two-channel ST-DL model, which utilises an LSTM network to capture the temporal dynamics of macroeconomic fundamentals, a CNN to extract the spatial diffusion characteristics of news sentiment across regions, and ultimately realises the adaptive fusion of the two heterogeneous sources of information through an attention mechanism.

An empirical study on Chinese provincial panel data demonstrates that the ST-DL model significantly outperforms multiple benchmark approaches including traditional econometric models, classical machine learning models, and a single deep learning model in both early warning accuracy (F1-score) and forward-lookingness (AUC). Ablation experiments confirm the respective indispensable contributions of news emotional information and attention fusion mechanisms. The case study further demonstrates that the model is able to utilise the prior indicator property of news sentiment to send effective early warning signals of economic shocks before they are fully captured by traditional data.

## 6 Conclusions

The theoretical contributions of this study are mainly reflected in three aspects: first, it successfully expands the application scenario of textual sentiment analysis from high-frequency financial markets to the field of macro-regional economic risks, which enriches the research paradigm of economic forecasting; second, it proposes and validates an in-depth multi-source data fusion architecture, which provides a technological blueprint for dealing with complex and heterogeneous information in the economic system; and third, it provides new empirical evidence for behavioural macroeconomics by means of the interpretable attention weights, it reveals the dynamic change patterns of the relative importance of fundamentals and sentiment at different stages of the economic cycle, providing new empirical evidence for behavioural macroeconomics.

At the practical level, this study has important implications for policy makers and financial regulators. We suggest that a real-time sentiment monitoring subsystem should be integrated into the existing macroeconomic monitoring system based on quantitative indicators. Our model can be used as a core computational engine to help policymakers identify potential risk areas earlier, realising the shift from ‘after-the-fact response’ to ‘before-the-fact warning’. For example, when the model’s probability of risk in a certain region is persistently high and the weight of attention is obviously skewed towards negative news, policymakers can conduct targeted stress tests and prepare plans, thus enhancing the precision and foresight of macroeconomic governance and providing effective scientific and technological support for the maintenance of the country’s economic and financial stability. Future work includes: first, applying ST-DL to city-level data; secondly, GNNs were used to establish the economic network model. The third incorporates SHAP for interpretability.

## Declarations

All authors declare that they have no conflicts of interest.

## References

Baesens, B., Van Gestel, T., Viaene, S., Stepanova, M., Suykens, J. and Vanthienen, J. (2003) 'Benchmarking state-of-the-art classification algorithms for credit scoring', *Journal of the Operational Research Society*, Vol. 54, No. 6, pp.627–635.

Baker, S.R., Bloom, N. and Davis, S.J. (2016) 'Measuring economic policy uncertainty', *The Quarterly Journal of Economics*, Vol. 131, No. 4, pp.1593–1636.

Cardarelli, R., Elekdag, S. and Lall, S. (2011) 'Financial stress and economic contractions', *Journal of Financial Stability*, Vol. 7, No. 2, pp.78–97.

Chen, T. (2014) 'Introduction to boosted trees', *University of Washington Computer Science*, Vol. 22, No. 115, pp.14–40.

Demirgüç-Kunt, A. and Detragiache, E. (2005) 'Cross-country empirical studies of systemic bank distress: a survey', *National Institute Economic Review*, Vol. 192, pp.68–83.

Fatt, C., Ujum, E. and Ratnavelu, K. (2010) 'The structure of collaboration in the *Journal of Finance*', *Scientometrics*, Vol. 85, No. 3, pp.849–860.

Greene, W.H. (2003) *Econometric Analysis*, Vol. 1, pp.3–5, Pretence Hall.

Hochreiter, S. and Schmidhuber, J. (1997) 'Long short-term memory', *Neural Computation*, Vol. 9, No. 8, pp.1735–1780.

Jordan, M.I. and Mitchell, T.M. (2015) 'Machine learning: trends, perspectives, and prospects', *Science*, Vol. 349, No. 6245, pp.255–260.

Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M. and Inman, D.J. (2021) '1D convolutional neural networks and applications: a survey', *Mechanical Systems and Signal Processing*, Vol. 151, p.107398.

Lan, G. (2020) 'First-order and stochastic optimization methods for machine learning', *First-Order and Stochastic Optimization Methods for Machine Learning*, Vol. 1, pp.3–10.

Li, X., Wu, P. and Wang, W. (2020) 'Incorporating stock prices and news sentiments for stock market prediction: a case of Hong Kong', *Information Processing & Management*, Vol. 57, No. 5, p.102212.

Liu, Y., Gong, C., Yang, L. and Chen, Y. (2020) 'DSTP-RNN: a dual-stage two-phase attention-based recurrent neural network for long-term and multivariate time series prediction', *Expert Systems with Applications*, Vol. 143, p.113082.

Lukauskas, M., Pilinkienė, V., Bruneckienė, J., Stundžienė, A., Grybauskas, A. and Ruzgas, T. (2022) 'Economic activity forecasting based on the sentiment analysis of news', *Mathematics*, Vol. 10, No. 19, p.3461.

Lundberg, S.M., Erion, G., Chen, H., DeGrave, A., Prutkin, J.M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N. and Lee, S-I. (2020) 'From local explanations to global understanding with explainable AI for trees', *Nature Machine Intelligence*, Vol. 2, No. 1, pp.56–67.

Nti, I.K., Adekoya, A.F. and Weyori, B.A. (2021) 'A novel multi-source information-fusion predictive framework based on deep neural networks for accuracy enhancement in stock market prediction', *Journal of Big Data*, Vol. 8, No. 1, p.17.

Ouyang, Z. and Lu, M. (2024) 'Systemic financial risk forecasting with decomposition-clustering-ensemble learning approach: evidence from China', *Symmetry*, Vol. 16, No. 4, p.480.

Shapiro, A.H., Sudhof, M. and Wilson, D.J. (2022) 'Measuring news sentiment', *Journal of Econometrics*, Vol. 228, No. 2, pp.221–243.

Wang, Y., Long, M., Wang, J., Gao, Z. and Yu, P.S. (2017) 'Predrnn: recurrent neural networks for predictive learning using spatiotemporal LSTMs', *Advances in Neural Information Processing Systems*, Vol. 30, p.5.

Xu, Y-Z., Zhang, J-L., Hua, Y. and Wang, L-Y. (2019) 'Dynamic credit risk evaluation method for e-commerce sellers based on a hybrid artificial intelligence model', *Sustainability*, Vol. 11, No. 19, p.5521.

Yang, H., Zhao, Y., Wu, Y., Wang, S., Zheng, T., Zhang, H., Ma, Z., Che, W., Wang, S. and Wei, S. (2025) 'Large language models meet text-centric multimodal sentiment analysis: a survey', *Science China Information Sciences*, Vol. 68, No. 10, pp.1–29.

Yu, L., Wang, S. and Lai, K.K. (2008) 'Credit risk assessment with a multistage neural network ensemble learning approach', *Expert Systems with Applications*, Vol. 34, No. 2, pp.1434–1444.