

Article

Multimodal Online Public Opinion Event Extraction and Trend Prediction for Edible Agricultural Products

Yong Han ¹, Zhenqiao Liu ^{2,*}, Hongying Bai ³ and Shaoyi Song ^{2,*}¹ School of Computer Science and Engineering, Beihang University, Beijing 100191, China; hanyong@buaa.edu.cn² School of Computer Science and Artificial Intelligence, Beijing Technology and Business University, Beijing 100048, China³ School of Information Engineering, Ordos College of Applied Technology College, Ordos 017000, China; hongyingbai@oit.edu.cn

* Correspondence: 2431061235@st.btbu.edu.cn (Z.L.); songshaoyi@btbu.edu.cn (S.S.)

Abstract

With the advent of the information age, food safety issues often trigger public panic and even disrupt social order. Traditional public opinion analysis methods struggle to handle massive amounts of data and effectively predict public sentiment trends. To address this problem, this paper proposes a deep reinforcement learning-based method for predicting public opinion trends related to food safety. First, a multimodal event extraction method is designed to extract image information and fuse it with text to generate new textual information. Then, an event detection model is designed based on the HDBSCAN clustering algorithm to more accurately identify emerging issues in safety-sensitive areas. Finally, this paper proposes a deep reinforcement learning-based model to predict public sentiment trends regarding food safety. Experimental results show that, compared with traditional methods, the proposed method has higher accuracy and adaptability in handling sudden or delayed public opinion events.

Keywords: feature analysis; safety of edible agricultural products; reinforcement learning; analysis of online public opinion

Academic Editor: George
Angelos Papadopoulos

Received: 11 November 2025

Revised: 27 November 2025

Accepted: 5 December 2025

Published: 7 December 2025

Citation: Han, Y.; Liu, Z.; Bai, H.; Song, S. Multimodal Online Public Opinion Event Extraction and Trend Prediction for Edible Agricultural Products. *Electronics* **2025**, *14*, 4813. <https://doi.org/10.3390/electronics14244813>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Under the current circumstances, safety of edible agricultural products-related online public opinion is characterized by high public attention, rapid dissemination, and significant social impact. It is more likely to cause public panic and ultimately disrupt social order, posing a threat to social stability. The spread of safety of edible agricultural products public opinion events is a time-evolving process rather than a static prediction task. Reinforcement learning is inherently well-suited for such “sequential decision-making” problems. Public opinion data often exhibit characteristics such as label scarcity (e.g., there is no explicit annotation on whether a certain strategy is effective) and multi-stage dependency (early actions influence later outcomes). In many real-world scenarios, the goal is not only to predict public opinion development but also to guide or control its trajectory, for instance, through public guidance, official responses, or content distribution [1–3]. This resembles a game or multi-stage decision-making system, where reinforcement learning proves to be particularly applicable. Current methods for analyzing and predicting safety of edible agricultural products public opinion events include rule-based keyword detection, text classification using machine learning models, sentiment analysis and trend modeling

with deep learning, social network analysis, and time series forecasting. Each of these approaches offers distinct advantages in opinion detection, dissemination path analysis, and trend prediction—such as fast identification, deep semantic understanding, and clear propagation structure.

However, they also face limitations, including weak semantic recognition capability, high dependency on data quality and quantity, high computational complexity, incomplete platform data, and delayed response to sudden events [4–6]. Overall, there is a trade-off between accuracy, adaptability, and interpretability that remains a core challenge for existing methods [7].

In response to the limitations of existing methods for analyzing and predicting safety of edible agricultural products-related online public opinion, this paper presents a trend prediction approach that combines deep learning and reinforcement learning techniques, leveraging the power of deep reinforcement learning. The deep learning is used to extract sentiment features and keywords from public opinion texts, while reinforcement learning is employed to learn strategies for responding to the development of public opinion events, enabling accurate prediction of public opinion trends. Our contributions are as follows: (1) Construction of a multimodal extraction method for safety of edible agricultural products public opinion. To overcome the problem of information loss associated with traditional event extraction methods that depend only on textual data, this study employs efficient neural network models to extract both visual and semantic features from images related to public opinion. These are then fused with features from the original opinion texts. This not only enhances the semantic representation of the opinion corpus but also mitigates the ambiguity and noise often present in pure text-based approaches. (2) Development of a multimodal event detection model for safety of edible agricultural products public opinion. By aligning and fusing multimodal features, the model enables interpretable feature mapping and reasoning in complex environments. Based on the HDBSCAN clustering algorithm, an event discovery model is constructed to cluster safety of edible agricultural products-related public opinion events. News items describing the same event are grouped into the same cluster, facilitating the identification of thematic safety of edible agricultural products event topics. (3) Construction of the safety of edible agricultural products public opinion trend prediction model. Starting from the characteristics of the opinion events, the carriers of opinion, and the subjects involved, the paper builds the safety of edible agricultural products public opinion risk indicator system using the Analytic Hierarchy Process (AHP) to quantitatively describe opinion risk. Based on this, a trend prediction model leveraging deep reinforcement learning is proposed. This model features strong dynamic decision-making capabilities and can continuously optimize prediction strategies through a “state-action-reward” mechanism. It is well-suited to adapt to the nonlinear and multi-factor evolution of public opinion in complex and rapidly changing environments.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the methods used in the experiments presented in this paper. Section 4 presents experimental results. Section 5 concludes the paper.

2. Related Work

As a key subject in the study of online public opinion, online public sentiment related to the safety of edible agricultural products has attracted widespread attention from researchers across multiple disciplines. The content covered in this paper spans several key areas, including the application of multimodal data, event extraction and detection, the construction of public opinion indicator systems, and time series forecasting. This section offers a summary of the existing research in these key areas.

Early methods of event extraction primarily relied on rules and templates, which are difficult to scale and poorly suited to the rapidly growing volume of online information. Moreover, traditional approaches tend to focus solely on text-based, single-modal data, which are inadequate for handling the widely prevalent image-text fusion content found in today's digital landscape. To tackle the challenges presented by multimodal data, researchers have developed event extraction methods that combine both image and text information to improve extraction accuracy. For example, Chen et al. [8] designed a new method for jointly extracting events from videos and text articles. However, despite the growing maturity of these technologies, model accuracy remains limited in specialized fields such as safety of edible agricultural products public opinion, due to challenges like incoherent text and the frequent use of domain-specific terminology [9–12].

In the area of public opinion event detection, traditional clustering methods focus mainly on clustering candidate events, often requiring multiple rounds of information extraction from online data, which leads to long processing times, low accuracy, and slow convergence. As deep learning has gained prominence, integrating deep neural networks into clustering algorithms has become a popular approach for event detection. Meneghetti et al. [13] combined dimensionality reduction techniques with input-output mapping to achieve effective dimensionality reduction, but there are still some defects. For example, the dimensionality reduction process inevitably leads to the loss of some key features, and some hidden structures or higher-order correlations may not be preserved by the low-dimensional representation, thus affecting the accuracy of the input-output mapping. Therefore, although deep learning-based feature learning methods have shown significant advantages in event detection tasks, most existing methods are applicable to English language events, but not to the field of online public opinion related to the safety of edible agricultural products in Chinese [12–15].

In terms of constructing the indicator system, from the perspective of online public opinion dissemination, Song J. [16] and others integrated the Analytic Hierarchy Process (AHP) with communication dynamics to establish an indicator system focusing on three dimensions: media users, media dissemination, and information dissemination. Their research on integrated media guidance reflects the interrelationships among various indicators of media influence, starting from different aspects of media impact. However, in many studies on constructing online public opinion risk assessment indicator systems, there has been inadequate attention to the components and evolving patterns of online public opinion, which are essential for creating a more efficient and streamlined risk indicator system [17–19].

Regarding online public opinion trend prediction models, researchers have gradually found that Neural network model can effectively address time-series prediction problems [20,21]. For example, Athallah et al. [22] developed a classification system using convolutional neural networks (CNNs) by comparing two text representation techniques, Term Frequency-Inverse Document Frequency (TF-IDF) and Word2Vec, to classify texts according to specific topics or categories. However, TF-IDF in this method is based only on word frequency and cannot capture word order and deep semantics, so it still has obvious limitations.

To overcome the parallelization limitations of RNNs, some researchers have achieved remarkable success in machine translation tasks by applying the Transformer model to the field of text processing [23]. Building upon this, Choudhary [24] used an improved bidirectional Transformer model to uncover hidden features in text style, thereby distinguishing between genuine and fake news.

However, the self-attention mechanism in Transformer utilizes a fully connected architecture, leading to suboptimal model efficiency, limited capability in capturing long-

range semantics, and poor robustness [25–27]. To address these challenges, this paper leverages the strengths of deep learning models and the temporal characteristics of online public opinion events to propose a prediction model based on deep reinforcement learning, which accurately predicts the evolving trends of online public opinion events.

3. Methods

3.1. Multimodal Event Extraction and Feature Representation of Safety of Edible Agricultural Products Online Public Opinion

Our paper uses web crawlers to obtain multimodal news public opinion corpus information in the area of safety of edible agricultural products and employs splicing technology for alignment and fusion operations to enhance the model's cognitive ability. It acquires multimodal fused information for the extraction of public opinion events and then uses t-SNE to extract event features and verify the reasoning ability of the features.

3.1.1. Multimodal Event Extraction from Online Public Opinion on Safety of Edible Agricultural Products

The processing of multimodal information in many studies can be summarized into five main directions: representation, transformation, alignment, fusion, and collaborative learning. In this paper, the features of images and texts are aligned and fused within the input model, and the combined multimodal text data is subsequently employed for the public opinion event extraction task, greatly improving the model's performance by leveraging diverse information sources. This paper designs a multimodal extraction approach for safety of edible agricultural products online public opinion, and the overall network framework is shown in Figure 1.




Figure 1. Multimodal event extraction method for safety of edible agricultural products internet-based public sentiment.

As shown in Figure 1, the multimodal event extraction method extracts image features from online news images related to the safety of edible agricultural products. Entities are extracted using Faster R-CNN, and complete visual features are obtained through the ViT (Vision Transformer) model. The image is then processed by the BLIP (Bootstrapping Language-Image Pre-training) model to generate a description, which is vectorized using BERT to obtain the meaning-based characteristics of the images. Faster R-CNN is a classic two-stage object detection model. Its core idea is to generate high-quality candidate regions on deep convolutional feature maps and then classify and regress bounding boxes for each region. ViT is a model that directly applies the Transformer structure to image classification tasks. Borrowing from the Transformer, it achieves feature modeling through a global self-attention mechanism. BLIP is a pre-trained model for multimodal tasks in vision and language, with the core goal of unifying the performance of tasks such as image and text

representation. BERT is a pre-trained language model proposed by Google based on the Transformer encoder structure. It utilizes contextual information simultaneously through a bidirectional attention mechanism to obtain deep semantic representations.

In the model, a separate dynamic gating unit is introduced to regulate the fusion of text features generated from image descriptions and the image-based characteristics. The dynamic gating unit assigns a learnable weight to each feature and obtains the final fused vector through normalization. The gating unit continuously optimizes its weight allocation through gradient descent, giving higher weights to features that are more important in specific event scenarios, while suppressing the weights of noisy or irrelevant features. The impact of this dynamic weight adjustment on the final feature vector is mainly reflected in its enhanced discriminativeness of feature representation and its reduction in interference from redundant or noisy information on model prediction. Specifically, the text features derived from the image description are input into the gating unit, where a weight is applied to produce a normalized vector T , as shown in Equation (1):

$$T = H \cdot w \quad (1)$$

Here, w represents the weight of the dynamic gating unit, and H denotes the text features generated from the image description.

Next, the importance of this feature with respect to the visual features is calculated:

$$TKey_{text}, Value_{text} = S(\text{Softmax}(f(T))) \quad (2)$$

In Equation (2), $f(\cdot)$ denotes the Leaky_ReLU activation function, $S(\cdot)$ represents the splitting operation, and Key_{text} and $Value_{text}$, respectively, indicate the importance of the image description relative to the visual features. Finally, the visual features are concatenated with the text features of the image caption,

$$Key = C(Key_{text}, Key_i^{(n)}) \quad (3)$$

$$Value = C(Value_{text}, Value_i^{(n)}) \quad (4)$$

Here, in Equations (3) and (4), $C(\cdot)$ represents a concatenation operation.

Then, during the attention weight calculation process, the text *Query* (Q) interacts with the *Key* and *Value* to integrate the contextual information of the image into the text representation. The formulas for calculating the attention weights are given in Equations (5)–(7):

$$AttentionScore = Query \cdot Key^T \quad (5)$$

$$AttentionWeights = \text{Softmax}(AttentionScore) \quad (6)$$

$$AttentionOutput = AttentionWeight \cdot Value \quad (7)$$

After obtaining the feature fusion vector, it is input into the fully connected layer to perform a multi-classification task to determine the type of safety of edible agricultural products online public opinion event to which the news belongs. After determining the event type to which the news belongs, it is added to the fusion feature vector to strengthen the model's understanding, after which the argument roles associated with the event type are extracted, ultimately achieving multimodal event extraction for safety of edible agricultural products as per online public opinion.

3.1.2. Representation and Reasoning of Event Characteristics

The previous section improved the model's cognitive ability to handle multimodal data through concatenation techniques. This section discovers online public opinion events by outputting multimodal fused feature vectors. In the model design, the input of text information has been replaced. The original news text input is transformed into a simple sentence dataset constructed in this section. This simple sentence dataset is constructed by reorganizing the public opinion event elements (including event type, arguments, and corresponding argument roles) obtained after the multimodal event extraction work. Since it contains effective information from public opinion news, it significantly reduces the impact of noise and makes it easier for the model to identify safety-sensitive domains.

When dealing with increasingly complex and high-dimensional data features, the original high-dimensional dataset should be transformed into a low-dimensional representation while retaining as much of the original meaning of the data as possible. Low-dimensional representations of original data help to overcome the curse of dimensionality and make the data easier to process, analyze, and visualize. Safety of edible agricultural products public opinion corpora form a nonlinear, high-dimensional data space. To enhance data analysis efficiency and accuracy, this section utilizes the t-distributed Stochastic Neighbor Embedding (t-SNE) algorithm, a nonlinear dimensionality reduction technique, to decrease the dimensionality of the data features.

The core idea of t-SNE is to first calculate the pairwise distances between points in the multidimensional space and then map these distances into a two-dimensional or three-dimensional space for visualization. Essentially, t-SNE constructs a similarity model relying on minimal pairwise distances and a dissimilarity model relying on large pairwise distances, ensuring that the characteristics of high-dimensional points remain unchanged after being mapped to a lower-dimensional space. The specific process is as follows.

(1) Calculate similarity probabilities in the high-dimensional space. For each data point pair in the high-dimensional space, determine the conditional probabilities $p_{i|j}$ and $p_{j|i}$, which represent the probability of selecting one point as the neighbor of another, given a specific data point. These probabilities are calculated using a Gaussian (normal) distribution and depend on the Euclidean distance between the two points:

$$p_{i|j} = \frac{\exp(-\|x_j - x_i\|^2 / 2\sigma_j^2)}{\sum_{k \neq j} \exp(-\|x_j - x_k\|^2 / 2\sigma_j^2)} \quad (8)$$

$$p_{j|i} = \frac{\exp(-\|x_i - x_j\|^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|x_i - x_k\|^2 / 2\sigma_i^2)} \quad (9)$$

Here, σ_i represents the standard deviation of the Gaussian distribution centered on the news text x_i . It is automatically determined by the perplexity parameter, which reflects the neighborhood size of each point.

(2) Symmetrize the similarity probabilities. To ensure that, in the high-dimensional space, the total probability of any data point i selecting j as its neighbor is the same, the conditional probabilities are symmetrized to compute the joint probability p_{ij} between high-dimensional vectors.

$$p_{ij} = \frac{p_{i|j} + p_{j|i}}{2N} \quad (10)$$

(3) Calculate similarity probabilities in low-dimensional space. In the low-dimensional space, the similarity probabilities between data points are also computed. A t-distribution is

used to calculate the joint probability q_{ij} between news texts in the low-dimensional space, which helps address the so-called “crowding problem” during dimensionality reduction.

$$q_{ij} = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{k \neq l} (1 + \|y_i - y_l\|^2)^{-1}} \quad (11)$$

Here, y_i represents the point in a space with fewer dimensions corresponding to the news text x_i .

(4) Optimization process. The goal is to make the similarity probabilities PPP in the high-dimensional space and Q in the low-dimensional space as similar as possible. The objective function is defined by the Kullback–Leibler (KL) divergence C , as shown in Equation (12), to minimize the difference between the similarities in the high-dimensional and low-dimensional spaces.

$$C = \sum_i KL(P||Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}} \quad (12)$$

Adjust the position of the point y_i in the low-dimensional space through optimization algorithms such as gradient descent to minimize the KL divergence. This process is typically performed through multiple iterations, and when the objective function reaches a preset accuracy, the algorithm terminates, ultimately obtaining the corresponding data after dimensionality reduction in the high-dimensional space data.

3.2. The Safety of Edible Agricultural Products Network Public Opinion Event Discovery Model Based on HDBSCAN Algorithm

Text clustering algorithms play a crucial role in discovering online public opinion events. By automatically grouping text data with similar content in a corpus, they can reveal underlying themes, thereby identifying key online public opinion events related to food safety. Therefore, this paper clusters a collected dataset of online public opinion events related to food safety, defining the events into five categories: substandard, expired, counterfeit, hygiene issues, and poisoning incidents. This method facilitates model design and subsequent trend prediction research during the event discovery process.

This paper employs loading multimodal features into the HDBSCAN clustering model to achieve event discovery for popular safety of edible agricultural products-related online public opinion incidents. Compared to the traditional DBSCAN, HDBSCAN has significant advantages in processing complex, high-dimensional, and noisy data. In food safety online public opinion analysis, public opinion text data such as news reports, social media comments, and image descriptions may contain repetitive content, off-topic text, or information with significantly different expressions. HDBSCAN can automatically determine the number of clusters without requiring manual parameter specification, while DBSCAN relies on manual parameter tuning and is prone to missed detections or misclassification. Therefore, HDBSCAN, with its density adaptability, hierarchical clustering ability, and strong noise handling capabilities, is more suitable for processing the complex structure and noise in public opinion text data, providing a reliable foundation for multimodal event extraction and public opinion trend prediction.

A fundamental concept in the HDBSCAN algorithm is the mutual reachability distance. This distance between two points is determined by taking the larger value between the direct distance between the points and the distance from each point to its k -th nearest

neighbor, also referred to as the core distance. For two points A and B , their mutual reachability distance is defined as shown in Formulas (13) and (14):

$$\text{core}_k(x) = d(x, N^k(x)) \quad (13)$$

$$d_{\text{mreach}-k}(A, B) = \max\{\text{core}_k(A), \text{core}_k(B), d(A, B)\} \quad (14)$$

Here, $d_{\text{mreach}-k}(A, B)$ refers to the distance of mutual reachability between points A and B , and the $\text{core}_k(x)$ is the distance of a data point to its k -th nearest neighbor; and $d(A, B)$ represents the Euclidean distance between A and B .

Furthermore, HDBSCAN uses a minimum spanning tree to construct a hierarchical tree model between points and compress the clustering hierarchy. This allows the model to automatically determine the optimal clustering result by simply specifying the minimum number of samples needed for a cluster, avoiding complex parameter tuning and significantly improving the model's accuracy and applicability.

3.3. A Predictive Model for the Risk Trend Concerning Online Discussions About Safety of Edible Agricultural Products

This section starts by utilizing the Analytic Hierarchy Process (AHP) to create a risk indicator system for online public opinion. It then proceeds to design a hybrid model that combines LSTM with the reinforcement learning approach PPO to predict the trend of online public opinion related to safety of edible agricultural products.

3.3.1. Construction of the Safety of Edible Agricultural Products Online Public Opinion Risk Indicator System

The paper employs the Analytic Hierarchy Process (AHP) to construct a risk indicator system for online public opinion regarding food safety, aiming to quantify the relative importance of various indicators in public opinion risk. The expert scoring process begins with a review panel composed of multiple experts specializing in edible agricultural product safety and public opinion analysis, who conduct pairwise comparisons of each indicator. Experts independently score the importance of each indicator based on their experience and understanding of the event's impact. The scores are then averaged to form the final judgment matrix, thus reducing the influence of single expert bias on the weight calculation. The internationally accepted 1–9 scale is used in the scoring process, where 1 indicates that two indicators are equally important, 3, 5, 7, and 9 represent increasing preferences from slightly important to extremely important, and 2, 4, 6, and 8 are used for fine-tuning the degree of preference. This evaluation index fully expresses the experts' detailed judgments while ensuring the operability and consistency of the scoring. To verify the consistency of expert scores, this paper calculates the consistency ratio (CR) of the judgment matrix, using the formula $\text{CR} = \text{CI}/\text{RI}$, where CI is the consistency index and RI is the random consistency index. CI is calculated using the largest eigenvalue of the judgment matrix. When the CR is less than 0.1, the consistency of the judgment matrix is considered acceptable; if the CR exceeds 0.1, the scoring needs to be adjusted or the judgment needs to be re-examined. Through the above method, the weights of the food safety online public opinion risk index obtained in this paper reflect both the comprehensive judgment of expert opinions and ensure mathematical consistency, providing a scientific and reliable quantitative basis for subsequent public opinion trend prediction based on LSTM-PPO. This paper selects the safety of edible agricultural products online public opinion risk indicator system from the perspectives of events, netizens, and online media. The specific predictive indicators are shown in Table 1.

Table 1. Description of the safety of edible agricultural products online public opinion risk indicator system.

Objective Layer	Criterion Layer	Indicator Layer	Data Description
The risk of online public opinion regarding the safety of edible agricultural products	Event perspective	Event duration	Event duration
	Netizen perspective	Event attention	Comment volume
	Media perspective	Netizen attention	Baidu index
		News report quantity	Event report quantity

For the risk indicators of safety of edible agricultural products online public opinion, the expert scoring method is adopted to construct a consistency judgment matrix, as shown in Table 2. The 1–9 scale is used to compare and evaluate the relative importance or preference degree among different factors. It utilizes the numbers 1–9 and their reciprocals as the values of a scale. This scale is also commonly referred to as the nine-point scale or nine-level scale, as shown in Table 3.

Table 2. Judgment matrix of event force.

Event Perspective	Duration of Event	Degree of Event Attention
Duration of the event	1	1/2
Attention received by the event	2	1

Table 3. The 1–9 scale judgment matrix and its meanings.

Scale	Definition
1	The former holds equal importance to the latter.
3	The former is somewhat more important than the latter.
5	The former is considerably more important than the latter.
7	The former holds much greater importance than the latter.
9	The former is far more important than the latter.
2, 4, 6, 8	The intermediate values in the aforementioned adjacent judgment matrices.
Count down	If the importance ratio of factor i to j is a_{ij} , and then the importance ratio of factor j to i is $1/a_{ij}$.

Finally, the eigenvectors were calculated, and consistency checks were conducted. After calculation, it was determined that all judgment matrices passed the consistency test. The calculated weights of the indicators are shown in Table 4.

Table 4. Weights of indicators for trend prediction of safety of edible agricultural products online public opinion.

Goal Layer	Criterion Layer	Weight	Indicator Layer	Weight
The risk of online public opinion regarding the safety of edible agricultural products	Event perspective	0.4934	Duration of the event	0.3333
			Number of comments	0.6667
	From the perspective of online media	0.1958	Number of news reports	0.1958
	From the perspective of netizens	0.3108	Public attention on the Internet	0.3108

At this point, the AHP method has been used to compare and evaluate the constructed indicators, and the importance of each indicator has been determined. Thus, the safety of edible agricultural products online public opinion risk indicator system has been established, providing a scientific basis for subsequent trend prediction work.

3.3.2. The Deep Reinforcement Learning Model for Predicting Safety of Edible Agricultural Products Network Public Opinion Trends

The process of predicting online public opinion risk trends can be described as a sequential discrete Markov decision process (MDP) using reinforcement learning. In online public opinion risk trend prediction, the environment is complex, feedback is lagging, and data noise is significant. The Proximal Policy Optimization (PPO) method has clear advantages in this area. PPO ensures the stability and efficiency of the policy optimization process by limiting the policy update magnitude, avoiding overfitting or policy collapse. It is particularly suitable for handling high-dimensional, dynamic, and nonlinear time-series prediction tasks such as public opinion trends. Meanwhile, the Long Short-Term Memory (LSTM) model is a variant of recurrent neural networks in deep learning specifically designed for processing time-series data. This paper combines the advantages of both, constructing an LSTM-PPO model. LSTM in the model encodes time-series risk indicators, transforming states into hidden feature representations and providing time-dependent information for PPO to generate predictive actions. PPO utilizes reward signal optimization strategies to achieve dynamic and high-precision trend prediction of online public opinion related to food safety.

State S : The current public opinion state features, such as topic popularity, sentiment distribution, keyword frequency, and propagation path graphs.

Action A : Predicting the category of public opinion trend at the next time step (rise/decline/stable), hot topics, sentiment fluctuation trends, etc.

Reward $\mathcal{R}$: Positive or negative rewards are given based on prediction accuracy, e.g., +1 for correct trend prediction and −1 for incorrect prediction. Additionally, this paper introduces a penalty term for incorrect predictions, assigning different degrees of penalty according to the magnitude of the actual trend change, thereby constructing a composite reward function.

Policy $\pi_\theta(a_t|s_t)$: The prediction policy learned by the LSTM decision function. The stochastic policy network, parameterized by θ , serves as a downstream risk trend prediction model that predicts the next label based on previously generated labels and source data. Thus, $\log \pi_\theta(\hat{y}_t | \hat{y} < t, x) : S \rightarrow \Delta(A)$, where $\Delta(A)$ denotes the probability distribution over all actions (e.g., target vocabulary). The next action $\hat{y}_t$ will be determined using top-k sampling based on this probability distribution.

Cumulative reward function $\hat{A}_\pi^t$: Based on the definition of the generalized cumulative reward function, the cumulative reward at time t is:

$$\hat{A}_\pi^t = \delta_t + \gamma \delta_{t+1} + \dots + \gamma^{T-t+1} \delta_{T-1}, \quad \delta_t = r_t - V_\pi(\hat{y} < t, x) + \gamma V_\pi(\hat{y} < t+1, x) \quad (15)$$

Here, γ is the discount factor, r_t is the reward at time step t , and $V_\pi(s_t)$ is the state value function at time step t . Figure 2 below shows the framework of safety of edible agricultural products online public opinion trend prediction based on deep reinforcement learning.

Objective Function: The goal of LSTM-PPO is to find a policy that maximizes the expected reward when predicting risk trends based on that policy.

$$\max_{\theta} \mathbb{E}_{x \sim \chi, \hat{y} \sim \pi_\theta(\cdot|x)} [\mathcal{R}(\hat{y}, x, y)] \quad (16)$$

Here, χ represents the training set of the source data, $\pi_\theta(\cdot)$ is the policy function, and $\mathcal{R}(\cdot)$ is the reward function. In this paper, the objective function is expressed as

maximizing the advantage function rather than directly maximizing the reward, as shown in the formula below. This is performed to reduce the variance of the prediction results.

$$\max_{\theta} \mathbb{E}_{x \sim \chi, \hat{y} \sim \pi_{\theta}(\cdot|x)} \left[\sum_{t=0}^T \hat{A}_{\pi}^t((\hat{y} < t, x), \hat{y}_t) \right] \quad (17)$$




Figure 2. Framework diagram of safety of edible agricultural products online public opinion trend prediction based on deep reinforcement learning.

The policy gradient method is used to estimate the gradient of the reward-based non-differentiable objective function in the above formula. Therefore, for a given source data x , the update of the policy parameters can be expressed as:

$$\max_{\theta} \mathcal{L}_{\theta}^{PG} = \max_{\theta} \mathbb{E}_{\hat{y} \sim \pi_{\theta}} \left[\sum_{t=0}^T (\log \pi_{\theta}(\hat{y}_t | \hat{y} < t, x) \hat{A}_{\pi}^t) \right] \quad (18)$$

where $\nabla_{\theta} \mathcal{L}_{\theta}^{PG} = \mathbb{E}_{\hat{y} \sim \pi_{\theta}} \left[\sum_{t=1}^T (\nabla_{\theta} \log \pi_{\theta}(\hat{y}_t | \hat{y} < t, x) \hat{A}_{\pi}^t) \right]$, $\nabla_{\theta} \mathcal{L}_{\theta}^{PG}$ represents the estimated gradient of the objective function under the policy parameterized by θ . To further reduce variance and avoid drastic updates to the policy at each iteration, the objective function in the above formula is reformulated into the form shown below, referred to as conservative policy iteration:

$$\begin{aligned} \mathcal{L}_{\theta}^{CPI} &= \mathbb{E}_{\hat{y} \sim \pi_{\theta}} \left[\sum_{t=0}^T \left(\frac{\log \pi_{\theta}(\hat{y}_t | \hat{y} < t, x)}{\log \pi_{\theta_{old}}(\hat{y}_t | \hat{y} < t, x)} \right) \hat{A}_{\pi}^t \right] \\ &= \mathbb{E}_{\hat{y} \sim \pi_{\theta}} \left[\sum_{t=0}^T (c_{\pi}^t(\theta)) \hat{A}_{\pi}^t \right] \end{aligned} \quad (19)$$

where θ_{old} is the policy parameter before the update, $c_{\pi}^t(\theta)$ is the ratio of the logarithmic probabilities between the new policy and the old policy.

Reward Function: A positive reward is given when the model correctly predicts the trend direction; otherwise, a negative reward is assigned:

$$\mathcal{R}_t = \begin{cases} +1, & \text{if } \hat{y}_t = y_t \\ -1, & \text{if } \hat{y}_t \neq y_t \end{cases} \quad (20)$$

This is the most basic “0–1 reward,” suitable for classification accuracy-oriented objectives. For incorrect predictions, a penalty term is introduced, assigning different levels of punishment based on the actual magnitude of trend changes:

$$\mathcal{R}_t = \begin{cases} +1, & \text{if } \hat{y}_t = y_t \\ -\lambda|\Delta_t - \hat{\Delta}_t|, & \text{if } \hat{y}_t \neq y_t \end{cases} \quad (21)$$

Here, Δ_t is the actual magnitude of change in public opinion risk (such as sentiment index or keyword frequency), $\hat{\Delta}_t$ is the predicted magnitude of change by the model, and λ is the penalty coefficient, which needs to be tuned during training. Therefore, a composite reward function can be designed as:

$$\mathcal{R}_t = \mathbb{I}[\hat{y}_t \neq y_t] - \lambda|\Delta_t - \hat{\Delta}_t| \quad (22)$$

Here, $\mathbb{I}[\cdot]$ is an indicator function, and equals 1 if the trend prediction is correct, and 0 if it is incorrect.

KL divergence constraint: The paper introduces the negative KL divergence penalty term $KL(\pi||\rho)$ in the reward to prevent the current policy π from deviating too far from the pre-trained language model ρ . The KL penalty term at time step t can be approximately expressed as:

$$R_{kl}(x, \hat{y} < t) = KL(\pi||\rho) \approx \log \frac{\pi(\cdot|x, \hat{y} < t)}{\rho(\cdot|x, \hat{y} < t)} = \log(\pi(\cdot|x, \hat{y} < t)) - \log(\rho(\cdot|x, \hat{y} < t)) \quad (23)$$

Here, $\log \pi(\cdot|x, \hat{y} < t)$ and $\log \rho(\cdot|x, \hat{y} < t)$ are the log probabilities given by the current policy π and the pre-trained model ρ at time step t , based on the source data x and the previously predicted labels. This reward term can be used to constrain actions and serves as an entropy reward to balance exploration and exploitation in the policy.

Loss Function: This paper adopts Proximal Policy Optimization (PPO) and defines the loss function of LSTM-PPO as follows:

$$\mathcal{L}_\theta = -\mathcal{L}_\theta^{CPI} + \alpha \mathcal{L}_\theta^{VF} \quad (24)$$

$$\mathcal{L}_\theta^{CPI} = \mathbb{E}_{\hat{y} \sim \pi_\theta} \left[\sum_{t=0}^T \min(c_\pi^t(\theta) \hat{y}_\pi^t, \text{clip}(c_\pi^t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_\pi^t) \right] \quad (25)$$

$$\mathcal{L}_\theta^{VF} = \mathbb{E}_{\hat{y} \sim \pi_\theta} \left[\sum_{t=0}^T (V_\pi(\hat{y} < t, x) - (\hat{A}_\pi^t + V_{\pi_{old}}(\hat{y} < t, x)))^2 \right] \quad (26)$$

Here, the loss function $\mathcal{L}_\theta$ is a combination of the policy objective function $\mathcal{L}_\theta^{CPI}$ and the squared discrepancy in the value function $\mathcal{L}_\theta^{VF}$. Therefore, minimizing the loss function $\mathcal{L}_\theta$ can both maximize the surrogate advantage policy objective (i.e., actor optimization) and minimize the value estimation error (i.e., critic optimization). In other words, the actor is guided to maximize the advantage policy objective, which relates to maximizing the expected reward, while the critic is required to minimize the per-token value estimation error defined by the difference between the new policy value function $V_\pi(\hat{y} < t, x)$ and the dense return estimate $\hat{A}_\pi^t + V_{\pi_{old}}(\hat{y} < t, x)$ under the old policy. In the formula above, ϵ represents the clipping range for the proximal policy ratio, and α is the weighting coefficient used in the linear combination between the actor and critic loss terms.

4. Experiments and Result Analysis

4.1. Collecting Data on Online Public Opinion and Tagging Events

The dataset used in this paper was obtained from relevant websites in the field of agricultural product safety through web crawling, constructing a multimodal news and public opinion corpus. The crawled data was preprocessed to meet the basic requirements of the online public opinion analysis model for agricultural product safety. The sources of the public opinion corpus include agricultural product-related information portals, such as Partner.com, Toutiao, China Quality News Network, and Baidu News. A total of 3847 news reports on online public opinion risks related to agricultural product safety were collected. Considering the characteristics of the corpus, we developed an event pattern specifically for online public opinion in the field of agricultural product safety. This pattern uses the same event description to match text and images. The dataset contains 3847 image-text pairs, each labeled with 5 event types and 28 AR-parameter roles. Table 5 shows the crawled and processed multimodal corpus. Table 6 details five categories of online public opinion events related to agricultural product safety, specifically including non-compliance events, hygiene events, poisoning events, counterfeit and infringement events, and expired agricultural product events.

Table 5. Examples of multimodal network public opinion corpus.

Number	Corpus
1	On 13 September 2023, edible wheat starch labeled as produced by Chuzhou Xinfu Food Co., Ltd. and sold by Anhui Lu'an Yinshi Trading Co., Ltd. was found by Anhui Public Inspection Research Institute Co., Ltd. to have a total bacterial count that does not meet safety of edible agricultural products standards.
2	On 8 September 2023, natural mineral water labeled as produced by Ningbo. On 14 July 2023, the Market Supervision and Administration Bureau of Shanxi Province issued a notice stating that cucumbers sold by the Dongjie Branch of Tianli Restaurant in Tunliu County were found to contain avermectin in violation of the national safety of edible agricultural products standards by BIVACNO (Shandong) Testing Technology Co., Ltd.
...	...

Table 6. Types of online public opinion events related to the safety of edible agricultural products.

Number	Types of Online Public Opinion Events
1	Substandard edible agricultural products
2	Safety of edible agricultural products incidents involving agricultural products
3	Poisoning incidents involving agricultural products
4	Counterfeit and Infringing Agricultural Products Incidents
5	Expired products incidents

The public opinion data texts were preprocessed, including HTML tag removal, handling of missing values, Chinese word segmentation, and stop-word removal. First, tools such as Beautiful Soup were used to remove HTML tags and clean redundant text. Next, missing values were handled: rows were deleted if the text was empty, and empty URLs were replaced with empty strings. Then, Chinese text was segmented using Jieba's precise mode, dividing the text into words to prepare for subsequent vectorization and word embedding model training. Finally, stop words were removed using the Harbin Institute of Technology stop-word list to clean words without practical meaning, while retaining professional terms related to safety of edible agricultural products.

After preprocessing the collected multimodal public opinion corpus, the paper performed data annotation. First, the safety of edible agricultural products as per online public

opinion events were categorized into five types: non-compliance events, hygiene events, poisoning events, counterfeiting and infringement events, and expiration events. Then, the event texts were annotated, where the event type contained in the corpus was labeled as “1,” and event types not contained were labeled as “0.” A sample of text annotations is shown in Table 7.

Table 7. Event type annotation.

Text Corpus	Event Type	Label
On 30 June 2022, the Market Supervision Bureau of Qingyuan County, Zhejiang Province, announced the 4th batch of sampling inspection results. The announcement showed that two batches of pure milk labeled as produced by Maqu’er Group Co., Ltd. were found to be non-compliant, with propylene glycol identified as the non-compliant item.	Substandard	1
	Hygiene	0
	Poisoning	0
	Counterfeit and Infringement	0
	Expired	0

The argument role extraction task involves extracting the argument roles associated with events from the event corpus based on the event types it contains. Considering the integration with subsequent research, this paper adopts the BIO tagging method, which is easier to implement for text, to accomplish the argument role annotation task. Three letters are used as labels for different argument roles within the same event type. Due to the specificity of the domain corpus, argument roles with identical values are annotated with the same BIO labels. For example, in cases where the “Involved Amount” in the “Counterfeiting and Infringement” event type and the “Fine Amount” in the “Expiration” event type have the same value in the same sentence, both are labeled with the same “B-MOY” and “I-MOY” tags. The detailed annotation of argument roles is shown in Table 8.

Table 8. Argument role annotation.

Event Type	Argument Role	BIO Label
Substandard	Time	B-TIM, I-TIM
	Location	B-LCO, I-LCO
	Non-compliant pollutants	B-POL, I-POL
	Name	B-FOD, I-FOD
	Involved Institutions	B-ORG, I-ORG
	Production Enterprise	B-FAC, I-FAC
	Sales Enterprise	B-SHP, I-SHP
Hygiene	Time	B-TIM, I-TIM
	Location	B-LCO, I-LCO
	Name	B-FOD, I-FOD
	Involved Enterprise	B-COM, I-COM
	Sanitary Condition	B-HYG, I-HYG
Poisoning	Time	B-TIM, I-TIM
	Location	B-LCO, I-LCO
	Name	B-FOD, I-FOD
	Poisoning Substance	B-POL, I-POL
	Poisoning Reaction	B-REA, I-REA
Counterfeiting and infringement	Time	B-TIM, I-TIM
	Location	B-LCO, I-LCO
	Name	B-FOD, I-FOD
	Involved Brand	B-BRD, I-BRD
	Involved Amount	B-MOY, I-MOY
	Involved Quantity	B-QTY, I-QTY

Table 8. *Cont.*

Event Type	Argument Role	BIO Label
Expired	Time	B-TIM, I-TIM
	Location	B-LCO, I-LCO
	Name	B-FOD, I-FOD
	Involved Enterprise	B-COM, I-COM
	Fine Amount	B-MOY, I-MOY

4.2. Analysis of Experimental Results of Multi-Modal Event Extraction Method for Safety of Edible Agricultural Products Online Public Opinion

In this study, this paper allocates 80% of the dataset for multimodal event extraction and uses the remaining data for the test set, as shown in Table 9.

Table 9. Division of datasets.

Dataset	Training Set	Testing Set
Multimodal dataset of online public opinion on the safety of edible agricultural products	3078	769

To rigorously assess the performance of the developed model, several cutting-edge models were chosen for comparison experiments and evaluated against the model presented in this paper. Specifically, they are as follows:

DMCNN: A two-phase event extraction approach utilizing a dynamic multi-pooling convolutional neural network model.

Joint3EE: A jointly trained model based on shared bidirectional GRU hidden layer representations, which can simultaneously complete the prediction tasks of entity mentions, event triggers, and arguments.

BERT-CRF: An event extraction method that acquires trigger word features based on the BERT pre-training model and classifies through Conditional Random Field (CRF).

BERT-BLSTM-CRF: An event extraction method that, after the BERT pre-training model, uses a bidirectional LSTM network in the feature layer to extract text context features and then classifies through Conditional Random Field (CRF).

For the recognition task, the model introduced in this paper is evaluated against the four event extraction methods mentioned earlier in the experiments. The results of these experiments are presented in Table 10 below.

Table 10. The results of event recognition comparative experiments.

Models	Precision (%)	Recall (%)	F1 (%)
DMCNN	75.32	74.38	74.85
Joint3EE	73.64	76.27	74.93
BLSTM-CRF	69.46	70.32	69.89
BERT-BLSTM-CRF	72.55	71.76	72.15
Our Model	81.03	82.70	81.86

The experimental results show that for the safety of edible agricultural products online public opinion dataset constructed in this paper, the proposed multimodal event extraction method delivers the highest performance among all the models compared in the event recognition task. Specifically, when compared with four existing event extraction methods—DMCNN, Joint3EE, BLSTM-CRF, and BERT-BLSTM-CRF—the multimodal event extraction approach shows notable enhancements in precision, recall, and F1 score. The precision

increased from 75.32% with the DMCNN model to 81.03%, the F1 score improved from 76.27% with the Joint3EE model to 82.70%, and the recall rose from 74.93% with the Joint3EE model to 81.86%.

These results fully confirm the outstanding performance of the proposed model in event recognition. The performance improvement stems from the model's effective integration and utilization of multimodal features, as well as its deep understanding of event semantics. On one hand, the multimodal event extraction method integrates information from different modalities, comprehensively capturing the context and intrinsic relationships of events, thereby improving the accuracy and robustness of event recognition. On the other hand, it optimizes the event classification process, demonstrating more efficient information processing capabilities. Experiments prove that introducing image features in addition to traditional event recognition tasks not only supplements missing semantic information for certain public opinion news texts but also helps to eliminate ambiguities, ultimately enhancing the model's event classification ability. For the argument extraction task, the proposed model is also compared with the neural network models discussed earlier, and the experimental results are presented in Table 11.

Table 11. Results of argument extraction comparison experiments.

Models	Precision (%)	Recall (%)	F1 (%)
DMCNN	74.18	72.04	73.09
Joint3EE	72.55	75.76	74.12
BLSTM-CRF	71.26	69.13	70.18
BERT-BLSTM-CRF	73.31	74.12	73.71
Our Model	78.15	79.47	78.80

The experimental results show that in the comparison experiments for the argument extraction task, the proposed model also demonstrates excellent performance. Compared with existing advanced models such as DMCNN, Joint3EE, BLSTM-CRF, and BERT-BLSTM-CRF, the multimodal event extraction method achieves the highest scores in precision, recall, and F1 score. In particular, the highest recall demonstrates that integrating entity information from images along with textual features enhances the model's ability to accurately identify event-related argument roles. Moreover, the precision increased from 73.31% with the BERT-BLSTM-CRF model to 78.15%, the recall rose from 75.76% with the Joint3EE model to 79.47%, and the F1 score improved from 73.71% with BERT-BLSTM-CRF to 78.80%. The experimental results indicate that the proposed model exhibits better generalization ability in the argument extraction task, enabling it to capture entity information that other models fail to annotate. Additionally, compared to joint learning models such as Joint3EE, the proposed model exhibits greater flexibility and accuracy when handling argument extraction tasks. This demonstrates that the model proposed in this paper not only performs well in single tasks but can also effectively handle complex, multi-task learning scenarios.

4.3. Analysis of Experimental Results of Model Based on HDBSCAN Algorithm

In the experiments on the safety of edible agricultural products online public opinion event discovery model, this paper obtains the event types and argument roles of safety of edible agricultural products online public opinion news through the multimodal event extraction method. Taking the "non-compliance" event type as an example of safety of edible agricultural products online public opinion events, the arguments and their corresponding roles extracted through the multimodal event extraction are shown in Table 12.

Table 12. Example of the role results in online public opinion discourse on the safety of edible agricultural products.

The Role of Online Public Opinion on the Safety of Edible Agricultural Products	
Time	28 June 2022
Location	Zhejiang Province
Production enterprise	Maituer
Sales enterprise	-
Involved institutions	Qingyuan County Market Supervision and Administration Bureau
Name	Pure Milk
Nonconformity items	Propylene Glycol

To improve the accuracy of the event discovery task, argument roles are represented as simple sentences to construct the text dataset. An example is as follows: on 28 June 2022, the Market Supervision Administration of Qingyuan County, Zhejiang Province, reported that pure milk produced by Maqu'er was found to contain the non-compliant item propylene glycol.

Similarly, all news articles are processed using this method to obtain a simple-sentence text dataset. Combined with the image dataset for each news article, image-text pairs are created, and a multimodal event dataset for the event discovery task is ultimately constructed.

To evaluate the differences between the HDBSCAN event discovery clustering algorithm used in this paper and general text clustering methods, three algorithms—K-Means, DBSCAN, and HDBSCAN—were selected for three sets of comparative experiments. Furthermore, to eliminate the influence of variables other than the clustering algorithm, all comparative experiments use the same multimodal feature representation method. This experiment adopts three evaluation metrics—Silhouette Coefficient (SC), Normalized Mutual Information (NMI), and Adjusted Rand Index (ARI)—to quantify clustering performance. The detailed results of the event discovery clustering comparison experiments are shown in Table 13.

Table 13. Comparison experiment results of event discovery by different clustering algorithms.

Clustering Algorithm	SC	NMI	ARI
K-Means	0.8223	0.8326	0.8163
DBSCAN	0.8621	0.8542	0.8603
HDBSCAN	0.8692	0.8754	0.8721

For different clustering algorithms, the larger the values of the clustering evaluation metrics within their respective ranges, the better the quality of the clustering results. Based on the experimental results in the table above, the HDBSCAN clustering algorithm demonstrates a significant performance advantage compared with traditional K-Means and DBSCAN algorithms, achieving the highest SC, NMI, and ARI values, which are 0.8692, 0.8751, and 0.8721, respectively. The three evaluation metric values of DBSCAN are slightly inferior, which may be due to the fact that the parameter *eps* in DBSCAN is a global variable, whereas HDBSCAN is a density-adaptive algorithm. The clustering performance of the K-Means algorithm shows greater variation, likely due to its inherent algorithmic limitations. Given that the text data in this paper contains a certain level of noise and the data distribution may not be uniform, the robustness of the K-Means clustering algorithm performs notably worse compared to the other two algorithms. Therefore, based on a comprehensive comparison of the experimental results, adopting the HDBSCAN clustering

algorithm can efficiently accomplish the clustering-based event discovery for the corpus in this study.

After the event discovery clustering task, each multimodal data point is assigned a cluster label representing its corresponding event topic. By analyzing these cluster labels, all multimodal information belonging to the same event can be filtered, thereby obtaining the complete event clustering results. In this experiment, a total of 432 different event topics were identified, with some of the corpus samples shown in Table 14.

Table 14. Partial text display of some special topics on public opinion events.

Event Topics	Types of Events	Partial Corpus Display of Each Event Topic
1	Substandard edible agricultural products	On 28 June 2022, the Market Supervision Administration of Qingyuan County, Zhejiang Province, announced the results of sampling inspections: two batches of pure milk produced by the dairy company Maqu'er were found to be non-compliant, with the non-compliant item being propylene glycol, which is prohibited according to the standard. On 28 June 2022, pure milk produced by Maqu'er was detected to have excessive levels of propylene glycol in the market of Qingyuan County, Zhejiang Province, with two samples containing 0.318 g/kg and 0.321 g/kg, respectively. It was also labeled as prohibited for use.
2	Poisoning	On 5 October 2019, Wang and nine of his relatives in Jixi City, Heilongjiang Province, died after jointly consuming homemade sour soup at home. The investigation revealed that the ingredient had been frozen in the refrigerator for one year, and preliminary diagnosis pointed to aflatoxin poisoning.
3	Hygiene	On 15 March 2022, the 3.15 Consumer Rights Day Gala exposed the unsanitary practices in pit-fermented pickled cabbage production, specifically naming companies such as Hunan Chaqi Vegetable Industry. At the production sites, workers were seen either barefoot or wearing slippers, stepping directly on the pickled cabbage, and some even threw cigarette butts onto the cabbage. On 15 March 2022, a CCTV evening program exposed the production of pit-fermented pickled cabbage. At a factory in Hunan, workers were seen wearing slippers or barefoot while handling the pickled cabbage in the pits, with phlegm, saliva, and cigarette butts carelessly discarded. The involved companies included Hunan Chaqi Vegetable Industry and others.
4	Counterfeiting and infringement	In June 2021, a video about “fake eggs” went viral on social media. The video showed people producing fake eggs that closely resembled real eggs and selling them to consumers at low prices. This video quickly drew widespread public attention, especially in the food sector, where the production and distribution of fake eggs exposed serious safety of edible agricultural products issues. In June 2021, a video about “fake eggs” circulated widely on social media. In the video, some people produced fake eggs that were almost identical to real eggs and sold them to consumers at low prices. This video quickly drew widespread public attention, especially in the food sector, where the manufacture and distribution of fake eggs exposed the seriousness of safety of edible agricultural products problems.
5	Expired	On 5 August 2020, the Beijing Municipal Administration for Market Regulation, in conjunction with the Shunyi police, successfully cracked a case involving the sale of expired agricultural products. During the investigation, police discovered that the gang illegally purchased agricultural products nearing their expiration date, repackaged them, and sold them at low prices to small supermarkets and farmers' markets, seriously impacting safety of edible agricultural products and consumer rights. On 5 August 2020, the Beijing Municipal Administration for Market Regulation and the Shunyi police jointly cracked a case involving the sale of expired agricultural products, seizing over 5000 kg of expired fruits and vegetables, as well as over 1000 pieces of counterfeit packaging materials. The case involved approximately 500,000 yuan, seriously threatening safety of edible agricultural products and harming consumer interests.

4.4. Analysis of Experimental Results of the Online Public Opinion Risk Trend Prediction Model for Safety of Edible Agricultural Products

The paper aims to construct an intelligent model capable of positively or negatively predicting the future trends of online public opinion events related to food safety, addressing the shortcomings of traditional public opinion analysis methods in understanding temporal evolution and responding to sudden risks. Given the rapid spread, wide impact, and complex evolution of food safety public opinion, its development is manifested not only in changes in emotional intensity but also in the rise (negative evolution) or fall (positive evolution) of risk levels over time. Therefore, this paper constructs a deep learning prediction model to extract multimodal sentiment and semantic features, achieving accurate prediction of future positive or negative trends.

To better illustrate that the LSTM-PPO model designed in this paper outperforms single-model LSTM and PPO, an ablation experiment was designed. Then, based on the dataset collected in this paper, the risk trends of online public opinion on food safety were predicted using multiple models, and the comparative experimental results were analyzed.

(1) Ablation experiment

In the ablation experiments, this paper used four model evaluation metrics: success rate (task completion rate); time per experiment (average time per experiment); convergence steps (number of steps required to train to a stable policy); and stability index (measures the degree of policy fluctuation during training; the lower the index, the more stable the policy). Regarding the experimental environment settings, the LSTM in LSTM and LSTM-PPO is consistent, and the settings for PPO and LSTM_PPO are consistent.

As shown in Table 15, the comprehensive analysis of the results for the three tasks—pure memory tasks, interference/noise tasks, and temporal inference tasks—leads to the conclusion that the LSTM-PPO combined model performs best across all tasks, based on the analysis of four key indicators: success rate, time efficiency, convergence steps, and policy stability. Specifically, LSTM-PPO achieves a higher success rate than either LSTM or PPO alone, reaching up to 92%, and maintaining a 78% success rate even in noisy or interference environments. In terms of time efficiency, each round takes only 0.11–0.12 s, significantly faster than both LSTM and PPO, resulting in the fastest training and inference speeds. It requires the fewest convergence steps, only 70k steps in pure memory tasks, and significantly fewer than single models in other tasks, indicating the highest learning efficiency. Regarding policy stability, the stability index of LSTM-PPO is 0.80–0.85, far higher than LSTM and PPO, indicating small reward fluctuations and strong policy reliability during training. Overall, the LSTM model combining sequence features and the PPO model optimized by reinforcement learning policies can give full play to the advantages of both, achieving high success rate, high efficiency, fast convergence and stable policies in different task environments, demonstrating the comprehensive performance advantages of the combined model in complex reinforcement learning scenarios.

Table 15. Comparison of ablation experiments.

Task Type	Models	Success Rate	Time per Episode (s)	Convergence Steps (k)	Stability Index
Pure memory tasks	LSTM	0.88	0.12	90	0.75
	PPO	0.80	0.15	100	0.70
	PSTM-PPO	0.92	0.11	70	0.85

Table 15. Cont.

Task Type	Models	Success Rate	Time per Episode (s)	Convergence Steps (k)	Stability Index
Interference/noise tasks	LSTM	0.65	0.14	120	0.50
	PPO	0.72	0.13	110	0.65
	PSTM-PPO	0.78	0.12	90	0.80
Temporal inference tasks	LSTM	0.75	0.13	110	0.70
	PPO	0.77	0.12	100	0.72
	PSTM-PPO	0.85	0.11	80	0.85

(2) Comparative Experiment on Public Opinion Risk Trend Prediction

This section the paper selected 432 event topics related to safety of edible agricultural products from public opinion events and chose the “Maqu’er Propylene Glycol Non-compliance” incident as a typical case study. The “Maqu’er Propylene Glycol Non-compliance” public opinion event began on 29 June 2022 and ended on 24 August 2022. Using web crawling technology, indicator data related to this event during its dissemination period were collected, including event duration, number of comments, number of reports, and public attention. The public opinion risk indicator data of the “Maqu’er Propylene Glycol Non-compliance” event were used for subsequent trend prediction experiments. For model training and evaluation, these data were divided into training and testing sets, with a ratio of 80% and 20%, respectively.

Through a comprehensive evaluation of multiple models, this paper selected four typical deep learning prediction models—RNN, LSTM, Transformer, and Autoformer—for comparative experiments, examining their performance differences compared to the LSTM-PPO model proposed in this study for safety of edible agricultural products online public opinion trend prediction tasks. During the experiments, multiple predictions were conducted on the same dataset, and the model’s performance was comprehensively assessed using various metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Root Mean Squared Error (RMSE), and Symmetric Mean Absolute Percentage Error (SMAPE). By comparing the results of different models on these metrics, we can gain a deeper understanding of the advantages and limitations of each prediction model. The experimental results are presented in Table 16.

Table 16. Comparison experiment results of different prediction models.

Models	MAE	MSE	MAPE	RMSE	SMAPE
RNN	0.3814	0.3358	0.3715	0.5429	0.7343
LSTM	0.2643	0.2825	0.2929	0.4558	0.4927
Transformer	0.1123	0.1212	0.1624	0.3842	0.2924
Autoformer	0.0592	0.1124	0.1322	0.3318	0.1548
LSTM-PPO	0.0422	0.1019	0.1122	0.2962	0.1035

The larger the values of the evaluation metrics for the prediction models, the greater the prediction error. The results in the table clearly indicate that the safety of edible agricultural products online public opinion trend prediction model based on LSTM-PPO outperforms the other four deep prediction models (RNN, LSTM, Transformer, Autoformer), achieving the best performance. In particular, the LSTM-PPO model achieved the highest prediction accuracy across all evaluation metrics, with its MAE, MSE, MAPE, RMSE, and SMAPE values being notably lower than those of the other models. Among them, the MAE of the

LSTM-PPO model was the smallest, at only 0.0422, while the MAE values of RNN, LSTM, Transformer, and Autoformer were 0.3814, 0.2643, 0.1123, and 0.0592. This presents that the improved LSTM-PPO model performed best in terms of mean absolute error, with the smallest gap between predicted and actual values, achieving the highest prediction accuracy. The Autoformer model produced error results closest to LSTM-PPO but was still slightly inferior, which proves that the improved PPO model incorporating LSTM can more accurately identify public opinion risk information.

Moreover, the SMAPE values of all prediction models were relatively high, possibly due to the presence of outliers or extreme values in the public opinion prediction data, such as data spikes during short-term public opinion surges. These values may affect the calculation of SMAPE, resulting in larger SMAPE values.

Overall, the excellent predictive performance of LSTM-PPO can be attributed to its model design, which employs sequence-level feature aggregation that better aligns with the continuity of time series, allowing it to capture dependencies within the input sequence more effectively. In addition, the LSTM component enhances the temporal feature representation of public opinion risk data, leading to more accurate predictions. Other models, constrained by their traditional RNN, LSTM, or Transformer structures, fail to fully exploit the characteristics of sequential data. Therefore, an examination of the experimental results reveals that the LSTM-PPO model has higher prediction accuracy and stability compared to other deep prediction models in the safety of edible agricultural products online public opinion trend prediction task, making it a more effective prediction model.

Furthermore, the experiment also visualized the prediction performance of the models by plotting prediction curves to intuitively compare how well the model predictions matched the actual trends. The prediction curve clearly shows the comparison between the predicted and actual values for the LSTM-PPO model, directly reflecting the performance of each model in the safety of edible agricultural products online public opinion trend prediction task, as shown in Figure 3. Through the model designed in this paper, we can more accurately predict the hot development trends of safety of edible agricultural products time in safety-sensitive domains.




Figure 3. Comparison curve of LSTM–PPO network event risk prediction results and actual risk values.

In order to better verify the performance of the LSTM-PPO model, 100 events were extracted from each of the five types of data in the dataset, and the prediction accuracy of the model was calculated. The results are shown in Figure 4. Figure 4 illustrates that the

LSTM-PPO-based food safety online public opinion trend prediction model demonstrates high prediction accuracy across various event types. In three independent experiments, the model's average accuracy was 90%, 89%, and 91%, respectively, showing stable overall performance and good robustness and reliability. In terms of event categories, poisoning events showed the highest prediction accuracy at 92%, 91%, and 93%, indicating that these events exhibit clear public opinion characteristics, and the model can easily capture their development trends. Expired events showed slightly lower accuracy at 88%, 87%, and 89%, respectively. Other event categories, such as substandard products, hygiene conditions, and counterfeit products, showed intermediate accuracy, indicating balanced overall performance. Overall, the LSTM-PPO-based food safety online public opinion trend prediction model not only provides high-accuracy prediction results but also possesses the ability to handle various types of public opinion events, providing reliable data support for the design of public opinion guidance and response strategies.




Figure 4. LSTM-PPO Predictive Analysis of Trends in Five Types of Events.

5. Conclusions

Safety of edible agricultural products online public opinion is marked by intense public interest and swift spread of information, and significant social impact, often triggering public panic and disrupting social order. Traditional public opinion analysis methods face challenges such as insufficient semantic recognition, strong data dependency, and high computational complexity. Therefore, accurately predicting and controlling the evolving patterns of public opinion has become an urgent task. This paper presents a method for predicting public opinion trends using deep reinforcement learning, which improves the model's cognitive capabilities by integrating event feature representation with reasoning techniques. Through multimodal event extraction, analysis of interpretability and reasoning capabilities in feature mapping under complex environments, event discovery, and trend prediction model design, this method effectively addresses the issue of feature information loss in traditional methods and demonstrates stronger predictive capabilities in dynamic decision-making environments, with a focus on safety-sensitive domains. The proposed method continuously optimizes strategies in complex and variable public opinion environments, providing more scientific and precise support for public opinion manage-

ment. The proposed LSTM-PPO prediction model based on deep reinforcement learning exhibits excellent accuracy and can capture the time-series dependencies of public opinion data. However, in some application scenarios, it often suffers from increased computational complexity and prolonged training time. For example, for large-scale, multimodal public opinion data, the training resources are substantial, potentially becoming a bottleneck for rapid response. Secondly, while multimodal data processing (text + image) enhances semantic understanding, it also increases the overhead of data alignment and feature fusion. Since the data used in this paper consists mainly of simple news texts and single images, there will be time delays when dealing with large amounts of complex data. Therefore, although the LSTM-PPO model designed in this paper demonstrates excellent prediction accuracy and stability, it has potential bottlenecks in real-time performance, computational resource consumption, and the ability to handle data sparsity. In practical applications, a comprehensive approach considering real-time performance, computational resource consumption, and the ability to handle data sparsity should be taken to improve model efficiency, thereby ensuring practicality and timely response capabilities in real-world food safety public opinion monitoring.

Author Contributions: Conceptualization, Y.H. and S.S.; Methodology, Y.H.; Software, H.B.; Validation, Z.L. and S.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Key Technology R&D Program of China (No. 2019YFC1606401), the National Natural Science Foundation of China (Grant No. 62433002), the Project of Construction and Support for high-level Innovative Teams of Beijing Municipal Institutions (Grant No. BPHR20220104), Beijing Scholars Program (Grant No. 099), and Inner Mongolia Autonomous Natural Science Foundation Project (Grant No. 2022QN06003).

Data Availability Statement: The data used in this article can be found at: https://github.com/hanyonggithub/food_event_extraction (accessed on 15 October 2025).

Conflicts of Interest: The research content in this paper is entirely original, with clear sharing among the authors and no conflict of interest.

References

1. Hassan, M.M.; Xu, Y.; Sayada, J.; Zareef, M.; Shoaib, M.; Chen, X.; Li, H.; Chen, Q. Progress of machine learning-based biosensors for the monitoring of safety of edible agricultural products: A review. *Biosens. Bioelectron.* **2025**, *267*, 116782. [CrossRef] [PubMed]
2. Xu, X.; Chen, X.; Hou, J.; Cheng, T.; Yu, Y.; Zhou, L. Should live streaming be adopted for agricultural supply chain considering platform's quality improvement and blockchain support? *Transp. Res. Part E Logist. Transp. Rev.* **2025**, *195*, 103950. [CrossRef]
3. Chon, M.G.; Xu, L.; Kim, J.; Liu, J. Understanding active communicators on the safety of edible agricultural products issue: Conspiratorial thinking, organizational trust, and communicative actions of publics in China. *Am. Behav. Sci.* **2025**, *69*, 168–186. [CrossRef]
4. Liu, L.; Zhang, T.; Wu, Q.; Xie, L.; Zhao, Q.; Zhang, Y.; Cui, Y.; Wang, C.; He, Y. Highly sensitive detection of carbendazim in agricultural products using colorimetric and photothermal lateral flow immunoassay based on plasmonic gold nanostars. *Talanta* **2025**, *281*, 126891. [CrossRef] [PubMed]
5. Meijer, N.; Safitri, R.; Tao, W.; Hil, E.H.-V.D. European Union legislation and regulatory framework for edible insect production—Safety issues. *Animal* **2025**, *19*, 101468. [CrossRef]
6. Rodríguez-Seijo, A.; Santás-Miguel, V.; Arenas-Lago, D.; Arias-Estévez, M.; Pérez-Rodríguez, P. Use of nanotechnology for safe agriculture and food production: Challenges and limitations. *Pedosphere* **2025**, *35*, 20–32. [CrossRef]
7. Kumar, D.; Pawar, P.P.; Ananthan, B.; Rajasekaran, S.; Prabhakaran, T.V. Optimized support vector machine based fused IOT network security management. In Proceedings of the 2024 3rd International Conference on Artificial Intelligence For Internet of Things (AIIoT), Vellore, India, 3–4 May 2024; IEEE: New York, NY, USA, 2024; pp. 1–5.
8. Chen, B.; Lin, X.; Thomas, C.; Li, M.; Yoshida, S.; Chum, L.; Ji, H.; Chang, S.-F. Joint multimedia event extraction from video and article. *arXiv* **2021**, arXiv:2109.12776. [CrossRef]
9. Vignesh, T.; Selvakumar, D.; Jayavel, R. Detecting ferric oxide adulteration in chilli Powder: A Multimodal analytical approach for enhanced safety of edible agricultural products. *Microchem. J.* **2025**, *208*, 112332. [CrossRef]

10. Addula, S.R.; Meesala, M.K.; Ravipati, P.; Sajja, G.S. A Hybrid Autoencoder and Gated Recurrent Unit Model Optimized by Honey Badger Algorithm for Enhanced Cyber Threat Detection in IoT Networks. *Secur. Priv.* **2025**, *8*, e70086. [\[CrossRef\]](#)
11. Mondal, M.; Khayati, M.; Sandlin, H.; Cudré-Mauroux, P. A survey of multimodal event detection based on data fusion. *VLDB J.* **2025**, *34*, 9. [\[CrossRef\]](#)
12. Ashqar, H.I.; Jaber, A.; Alhadidi, T.I.; Elhenawy, M. Advancing object detection in transportation with multimodal large language models (mllms): A comprehensive review and empirical testing. *Computation* **2025**, *13*, 133. [\[CrossRef\]](#)
13. Meneghetti, L.; Demo, N.; Rozza, G. A dimensionality reduction approach for convolutional neural networks. *arXiv* **2021**, arXiv:2110.09163. [\[CrossRef\]](#)
14. Upadhyay, A.; Meena, Y.K.; Chauhan, G.S. SatCoBiLSTM: Self-attention based hybrid deep learning framework for crisis event detection in social media. *Expert Syst. Appl.* **2024**, *249*, 123604. [\[CrossRef\]](#)
15. Balali, A.; Asadpour, M.; Jafari, S.H. COFEE: A comprehensive ontology for event extraction from text. *Comput. Speech Lang.* **2025**, *89*, 101702. [\[CrossRef\]](#)
16. Song, J.; Zhu, X. Research on public opinion guidance of converging media based on AHP and transmission dynamics. *Math. Biosci. Eng.* **2021**, *18*, 6857–6886. [\[CrossRef\]](#) [\[PubMed\]](#)
17. Zhao, X.; Yan, H.; Liu, Y. Hierarchical Multilabel Classification for Fine-Level Event Extraction from Aviation Accident Reports. *Inf. J. Data Sci.* **2025**, *4*, 51–66. [\[CrossRef\]](#)
18. Cavero-Redondo, I.; Saz-Lara, A.; Martínez-García, I.; Otero-Luis, I.; Martínez-Rodrigo, A. Validation of an early vascular aging construct model for comprehensive cardiovascular risk assessment using external risk indicators for improved clinical utility: Data from the EVasCu study. *Cardiovasc. Diabetol.* **2024**, *23*, 33. [\[CrossRef\]](#)
19. Hu, Y.; Ouyang, J.; Xia, Y.; Sheng, Y. The effects of different exercise intensities on body composition and cardiovascular risk indicators in children with metabolic syndrome: A RCT network meta-analysis. *BMC Sports Sci. Med. Rehabil.* **2025**, *17*, 292. [\[CrossRef\]](#)
20. Bhanja, S.; Das, A. A Black Swan event-based hybrid model for Indian stock markets' trends prediction. *Innov. Syst. Softw. Eng.* **2024**, *20*, 121–135. [\[CrossRef\]](#)
21. Li, R.; Wang, Z.; Du, X. Efficient Document-level Event Relation Extraction. In Proceedings of the 10th Workshop on Representation Learning for NLP (RepL4NLP-2025), Albuquerque, NM, USA, 4 May 2025; pp. 92–99.
22. Athallah, M.R.; Lhaksmana, K.M. Hadith Text Classification Based on Topic Using Convolutional Neural Network (CNN) and TF-IDF. *J. Renew. Energy Electr. Comput. Eng.* **2025**, *5*, 30–36. [\[CrossRef\]](#)
23. Zhang, H.; Shafiq, M.O. Survey of transformers and towards ensemble learning using transformers for natural language processing. *J. Big Data* **2024**, *11*, 25. [\[CrossRef\]](#)
24. Choudhary, A.; Arora, A. Assessment of bidirectional transformer encoder model and attention based bidirectional LSTM language models for fake news detection. *J. Retail. Consum. Serv.* **2024**, *76*, 103545. [\[CrossRef\]](#)
25. Xiong, S.; Tian, W.; Batra, V.; Fan, X.; Xi, L.; Liu, H.; Liu, L. Safety of edible agricultural products news events classification via a hierarchical transformer model. *Heliyon* **2023**, *9*, e17806. [\[CrossRef\]](#)
26. Ma, Y.; Yang, L.; Bai, X.; Wang, K. Sensitive detection of organophosphorus pesticides in agricultural food products by a highly luminescent coordination polymer. *Spectrochim. Acta Part A Mol. Biomol. Spectrosc.* **2025**, *341*, 126471. [\[CrossRef\]](#)
27. Wu, Y.; Han, H.; Chen, J.; Zhai, W.; Cao, Y.; Zha, Z. BRAT: Bidirectional Relative Positional Attention Transformer for Event-based Eye tracking. In Proceedings of the Computer Vision and Pattern Recognition Conference, Nashville, TN, USA, 11–15 June 2025; pp. 5136–5144.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.