





# The Mini-SiTian Array: Real-bogus Classification Using Deep Learning

Jing-Hang Shi<sup>1,2</sup>, Hong-Rui Gu<sup>1,2</sup>, Yang Huang<sup>2,1,3</sup> , Yan-Xia Zhang<sup>1,3</sup>, and Peng-Liang Du<sup>1</sup> 

<sup>1</sup>National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100101, China; [zyx@bao.ac.cn](mailto:zyx@bao.ac.cn)

<sup>2</sup>School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing 100049, China; [huangyang@ucas.ac.cn](mailto:huangyang@ucas.ac.cn)

Received 2024 December 17; revised 2025 February 18; accepted 2025 March 3; published 2025 May 6

## Abstract

The Mini-SiTian (MST) project is a pathfinder for China’s next-generation large-scale time-domain survey, SiTian, aimed at discovering variable stars, transients, and explosive events. MST generates hundreds of thousands of transient alerts every night, approximately 99% of which are false alarms, posing a significant challenge to its scientific goals. To mitigate the impact of false positives, we propose a deep learning–based solution and systematically evaluate 13 convolutional neural networks. The results show that ResNet achieves exceptional specificity (99.70%), EfficientNet achieves the highest recall rate (98.68%), and DenseNet provides balanced performance with a recall rate of 94.55% and specificity of 98.66%. Leveraging these complementary strengths, we developed a bagging-based ensemble classifier that integrates ResNet18, DenseNet121, and EfficientNet\_B0 using a soft voting strategy. This classifier achieved the best AUC value (0.9961) among all models, with a recall rate of 95.37% and specificity of 99.25%. It has now been successfully deployed in the MST real-time data processing pipeline. Validation using 5000 practically processed samples with a classification threshold of 0.798 showed that the classifier achieved 88.31% accuracy, 91.89% recall rate, and 99.82% specificity, confirming its effectiveness and robustness under real application conditions.

*Key words:* techniques: image processing – methods: data analysis – surveys

## 1. Introduction

The universe is a dynamic system where the properties of celestial objects, such as brightness and color, change due to various astrophysical processes. These variations, collectively referred to as transient events, encompass explosive phenomena (e.g., supernovae, kilonovae, gamma-ray bursts) and periodic stellar variability (e.g., flares, pulsating stars). The study of transient sources reveals the nature of extreme energy release and physical processes that are essential to understanding the formation and evolution of the universe. As a precursor to China’s next-generation time-domain survey, the Mini-SiTian (MST) project aims to detect and classify these transient events (Huang et al. 2025; He et al. 2025).

The MST project uses a three-stage variable source detection process (Gu et al. 2025): first, the HOTPANTS image subtraction algorithm (Becker 2015) is used to remove static background sources, generating a difference image. Next, candidate transient sources are extracted from the residual image using SExtractor (Bertin & Arnouts 1996). Finally, a manual screening process is employed to identify reliable samples. However, due to multiple interfering factors, such as optical artifacts, bright star spikes, mismatches during image subtraction, and instrumental effects, the image subtraction algorithm cannot perfectly eliminate static background sources.

This leads to the generation of a large number of candidate sources during the source extraction phase. Observational data show that MST produces hundreds of thousands of transient alerts every night, of which approximately 99% are false positives. This high false positive rate and inefficient manual screening process pose significant challenges to the scientific goals of the project. Therefore, it is essential to develop an efficient, accurate, and automated classifier to distinguish true astrophysical events (real) from false positives (bogus).

In recent years, automated and intelligent deep learning methods, particularly convolutional neural networks (CNNs), have become one of the primary techniques for addressing the classification of transient sources as real or bogus (Goldstein et al. 2015; Cabrera-Vives et al. 2017; Duev et al. 2019; Hosenie et al. 2021; Killestein et al. 2021; Chen et al. 2023). These methods analyze the image features of candidate sources to assess the authenticity of transient sources, ultimately assigning a score to each source on a scale from 0.0 (bogus) to 1.0 (real).

In this paper, we propose a CNN-based ensemble classifier for the MST project, aiming to address the challenge of high false positive rates and the need for efficient real-bogus classification. A high-quality data set has been constructed, consisting of 177,696 bogus samples and 3000 real samples, each represented as a  $64 \times 64$  pixel residual image. We then conducted a rigorous evaluation of 13 CNN architectures,

<sup>3</sup> Corresponding author.

spanning classical architectures (e.g., VGGNet Simonyan & Zisserman 2014; ResNet He et al. 2016; DenseNet Huang et al. 2017), lightweight designs (e.g., EfficientNet Tan & Le 2019; ShuffleNetV2 Ma et al. 2018), and emerging paradigms (e.g., RepVGG Ding et al. 2021; UniRepLKNNet Ding et al. 2024). Experimental analysis revealed that ResNet exhibited outstanding specificity (99.70%), EfficientNet achieved the highest recall (98.68%), and DenseNet demonstrated a balanced performance with a recall of 94.71% and specificity of 98.66%.

To leverage the strengths of these models, we integrated ResNet18, EfficientNet\_B0, and DenseNet121 using a bootstrap aggregating (Bagging) (Breiman 1996) strategy combined with soft voting, resulting in a highly robust classifier with an AUC of 0.9961. This ensemble classifier achieved an optimal balance on the test set, attaining a recall of 95.37% and a specificity of 99.25%. Through dynamic threshold optimization, we identified 0.798 as the optimal threshold, at which both precision and recall reached 92.2%. Validated on 5000 operational samples, the classifier demonstrated strong performance in practical deployment, achieving a precision of 88.31%, a recall of 91.89%, and a specificity of 99.82%. The system has now been stably deployed in the MST real-time processing pipeline, providing crucial technical support for achieving the scientific objectives of the MST project.

The organization of this paper is as follows: Section 2 presents a detailed description of the data set. Section 3 outlines the experimental methods. Section 4 offers an in-depth analysis of the experimental results. Finally, Section 5 provides a concise summary.

## 2. Data


### 2.1. Mini-SiTian

The MST is the pathfinder of the SiTian project (Liu et al. 2021), which aims to provide a technology validation platform and scientific target exploration for the next generation of time-domain surveys. Its core equipment includes three 30 cm Schmidt Complex Achromatic Telescopes (MST1, MST2, and MST3), each equipped with a ZWO ASI6200MM Pro CMOS camera ( $9576 \times 6288$  pixels,  $3.76 \mu\text{m}$  pixel size) and SDSS-like filters ( $i'$ ,  $g'$ ,  $r'$ ) with a pixel scale of  $0''.862 \text{ pixel}^{-1}$ .

From the start of pilot observations in 2022 November–2024 July, the MST has completed a total of 552 operational days, of which 344 were successfully acquired. Its observational targets cover 92 general sky fields, 2 comets, 45 transient sources from the Transient Name Server (TNS), and 33 LIGO gravitational wave alerts (He et al. 2025). These observations provide the basis for the construction of subsequent data sets.

### 2.2. Data Set Creation

In supervised deep learning, model performance heavily relies on accurately labeled, reliable, consistently distributed,



**Figure 1.** Examples of residual images of real candidates during the data set filtering process include a Short Period Variable (SPV), a Long Period Eclipsing Binary (LPEB), a Contact Binary (CB), and three Long Period Variables (LPV). Each example is labeled with its type and its R.A. and decl. coordinates. The top row illustrates problematic images excluded during manual inspection. The bottom row displays high-quality real samples retained after screening.

and representative samples. Based on observational data accumulated during the pilot phase of the MST, we designed a rigorous sample selection protocol that combines automated algorithms with manual review to construct a high-quality data set.

#### 2.2.1. Real Samples

The construction of the real sample set was based on continuous observations of 127 confirmed variable sources from the MST (Gu et al. 2025). To ensure data diversity, we randomly sampled observations from various epochs within each variable source’s observation history. Using their coordinates, we extracted  $64 \times 64$  pixel ( $55''.16$ ) cutout stamp images from the corresponding residual images to form the initial data set.

Despite these stamp images being derived from identified variable sources, there were some problematic samples. For example, some sample centers showed significant spurious features due to suboptimal image subtraction, while others lacked source features due to weak variability signals. To improve data set reliability, we first used SExtractor to detect all stamp images, discarding those without visible central targets. Subsequently, we performed a visual review of the remaining samples and filtered out those with obvious false features.

As a result, a total of 3000 real samples were obtained. Figure 1 shows the problematic examples that were excluded and the high-quality examples that were retained.

### 2.2.2. Bogus Samples

To build the bogus sample data set, we used SExtractor to automatically detect sources from over 200 residual images. To ensure data quality, we excluded detections near known variable sources and performed additional filtering on 190,000 samples.

An initial set of 3000 bogus images was manually selected from the candidate data set. These images, combined with the previously identified real samples, formed an initial training data set. Using this data set and the VGGNet13 network (Simonyan & Zisserman 2014), we obtained an auxiliary screening classifier, achieving an accuracy of 97% on this initial data set. As the training procedure aligns with that described in Section 3, it is not detailed here. The classifier was then applied to all bogus candidates, yielding a probability score for each target to be classified as real.


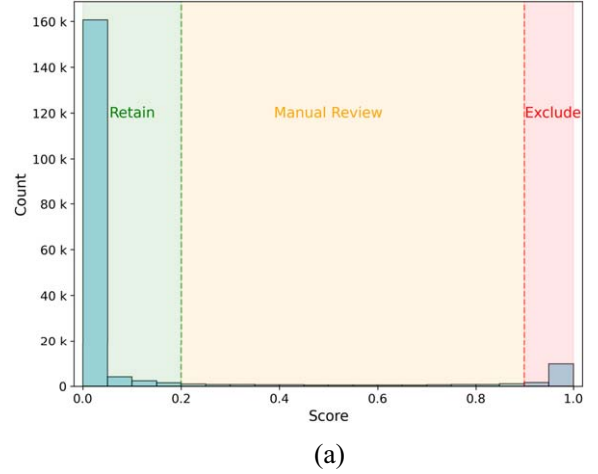
Based on the probability scores provided by the classifier, we screened all candidate samples. Figure 2(a) depicts the distribution of the predicted scores and the corresponding filtering strategies, which are detailed below:

1. *High-score exclusion.* Targets with scores greater than 0.9 were excluded, as the model believed these samples should be classified as real, totaling 11,317. Visual inspection revealed that these targets were mainly residual signals from stationary objects that were not fully subtracted or alignment errors during image subtraction. They typically exhibited features nearly identical to the real samples, as shown in the first row of Figure 2(b).
2. *Intermediate-score manual review.* Targets with scores between 0.1 and 0.8 were manually reviewed. Among the 9994 inspected targets, 9007 were retained. These targets showed intermediate confidence, indicating uncertainty in their classification as real or bogus, so they were carefully examined to ensure accurate labeling. The corresponding samples are shown in the second row of Figure 2(b).
3. *Low-score retention.* Targets with classification scores below 0.2 were directly retained, as the model was highly confident that these sources were not real, and thus they were classified as bogus. A total of 168,689 samples were retained, which are shown in rows 3 and 4 of Figure 2(b).

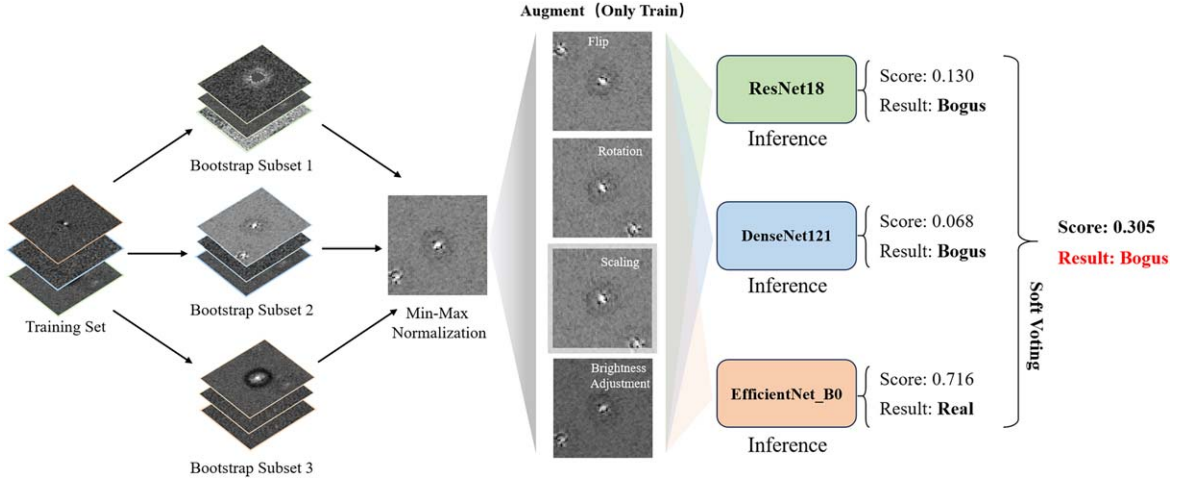
Finally, a total of 177,696 high-quality bogus samples were obtained.

### 2.2.3. Data Set Composition

Using the resulting real and bogus residual images, we constructed a data set. The data set was constructed for model training and validation, consisting of 177,696 bogus samples and 3000 real samples, which were split into training, validation, and test sets with a ratio of 7:1:2.



**Figure 2.** Distribution of classification-based scores and representative examples of bogus samples. (a) Distribution of classification scores for all bogus candidates, as predicted by the VGGNet13 classifier used for auxiliary selection. The  $x$ -axis represents the classification score, where higher values indicate a greater likelihood of the target being classified as real. (b) Examples of residual images of bogus candidates. The value in the upper left corner indicates the real-bogus score obtained using the auxiliary filtering classifier. The top row illustrates high-score samples ( $>0.9$ ), which were excluded. The second row shows intermediate-score samples ( $0.2-0.9$ ), which were manually reviewed. The bottom two rows show low-score samples ( $<0.2$ ), which were directly retained.



**Figure 3.** Overview of the ensemble classifier building process. The training set is divided into bootstrap subsets, and each subset is used to train ResNet18, DenseNet121, and EfficientNet\_B0 with data augmentation. During inference, predictions are combined using a soft voting strategy to achieve final classification results.

### 3. Methods

#### 3.1. CNNs and Ensemble Methods

CNNs are one of the most effective deep learning architectures for image-related tasks, including classification, object detection, and segmentation. By leveraging convolutional operations for feature extraction, pooling layers for dimensionality reduction, and fully connected layers for decision-making, CNNs enable high-precision image recognition. In recent years, various advanced CNN architectures such as VGGNet (Simonyan & Zisserman 2014), ResNet (He et al. 2016), DenseNet (Huang et al. 2017), and EfficientNet (Tan & Le 2019) have been proposed. These models achieve diverse balances between the number of parameters, computational efficiency, and performance, providing a wide range of choices for different application tasks.

To achieve accurate classification of real transient sources and bogus detections in the MST, we trained models with various architectures and parameters using the data set introduced in Section 2.2. All models followed the training strategy described in Section 3.2. The performance metrics for each model are detailed in Section 4.2.1. Based on experimental results and analysis, we selected ResNet18, DenseNet121, and EfficientNet\_B0 as base learners and constructed a more powerful ensemble classifier using the bagging ensemble learning method.

Specifically, during training, bootstrap resampling was applied to the training data to generate multiple training subsets with the same sample size as the original training set. This resampling method allows certain samples to appear multiple times in the new subsets, while other samples may be excluded, thus introducing diversity into the training data.

Given the class imbalance in the data set, resampling is performed separately for each class. Each base model is independently trained on its corresponding resampled subset to ensure diversity and variability among the models (Breiman1996).

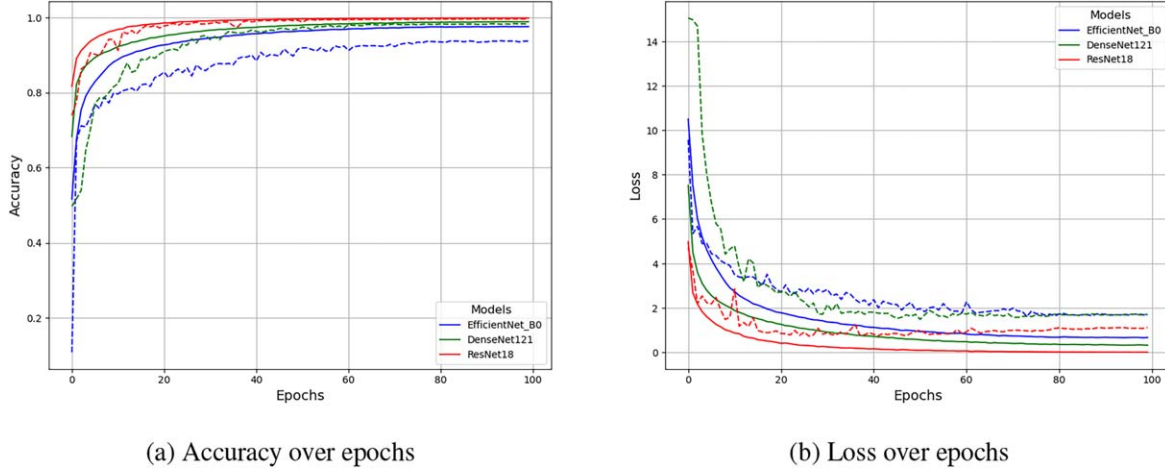
During the inference phase, we employed a soft voting strategy by averaging the predicted probabilities from all models to generate the final prediction. This approach effectively combines the strengths of multiple models while mitigating the weaknesses of individual models, significantly improving the classifier’s performance. Figure 3 provides an intuitive illustration of this method, including bootstrap resampling of training data, individual model training, and the soft voting strategy during the inference phase.

#### 3.2. Training Strategy

In the training of CNNs, the selection of hyperparameters and the design of training strategies are critical to the performance and stability of the models (Mishkin et al. 2017). To ensure consistency and fairness in the training processes of the multiple base models used in this paper, we adopted a uniform hyperparameter configuration for them, as detailed in Table 1. Furthermore, to mitigate the issue of model overfitting, a dropout rate of 0.25 was applied to the convolutional layers and 0.5 to the fully connected layers across all models.

##### 3.2.1. Data Augmentation

To enhance model generalization and mitigate overfitting, we applied online data augmentation techniques during training. These techniques generate new samples by applying random transformations to existing ones, thereby increasing the



**Figure 4.** Training and validation curves for ResNet18, DenseNet121, and EfficientNet\_B0 over 100 epochs of bagging ensemble training. Solid lines represent training metrics, and dashed lines represent validation metrics.

**Table 1**  
Hyperparametric Configurations during the Training Phase

Parameter	Value
Total Epochs	100
Batch Size	256
Optimizer	AdamW (Loshchilov & Hutter 2017)
Weight Decay	0.001
Beta Values ( $\beta_1, \beta_2$ )	(0.9, 0.999)
Scheduler	CosineAnnealingLR
Initial Learning Rate ( $lr$ )	$1 \times 10^{-4}$
Minimum Learning Rate	$1 \times 10^{-7}$

diversity of the training data set. In this paper, we selected the following augmentation methods based on the characteristics of astronomical images:

- (a) *Flip*. Images were flipped along vertical or horizontal axes to simulate different observational perspectives.
- (b) *Rotation*. Images were randomly rotated by multiples of  $90^\circ$  to account for orientation variations.
- (c) *Scaling*. Images were resized randomly between 80% and 120% of their original dimensions to simulate varying observational distances or resolutions.
- (d) *Brightness Adjustment*. The brightness of images was adjusted within a range of 80%–120% to simulate varying exposure conditions.

Each augmentation method was applied with a 50% probability.

### 3.2.2. Class Balancing

Class imbalance, where one class significantly outnumbers another, poses a substantial challenge as it biases the model

toward the majority class. In our data set, the ratio of negative (bogus) samples to positive (real) samples was 59:1, an extreme imbalance. To address this issue, we first increased the representativeness of the minority class (real samples) by duplicating instances in the training set, expanding their presence to five times the original amount. This ensured that the model received adequate exposure to minority-class samples during training. To counter the risk of overfitting associated with direct repetition, we integrated data augmentation techniques to enhance the diversity of the oversampled samples.

Furthermore, to further address the issue of class imbalance, we employed the focal loss function (Lin et al. 2017), an improvement over the standard cross-entropy loss, which dynamically adjusts the model’s focus on different samples. The focal loss is defined as:

$$FL(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t) \quad (1)$$

where  $p_t$  represents the predicted probability of the true class.  $\alpha_t$  is the class weighting factor, set to [0.1, 0.9] in this paper to emphasize the minority class.  $\gamma$  is a focusing parameter, set to 2.0, which amplifies the loss contribution of hard-to-classify samples. By adjusting  $\alpha_t$  and  $\gamma$ , focal loss reduces the contribution of well-classified samples, improving performance on the minority class.

The combination of oversampling and focal loss significantly improved the model’s ability to learn from imbalanced data sets.

### 3.3. Experimental Details

The experimental setup consisted of a system with two Intel Xeon Gold 6230 processors, an NVIDIA A100-PCIE-40 GB GPU running CUDA 12.2, and 251 GB of memory. All experiments were implemented in the Python programming language.

Figure 4 illustrates the evolution of loss and accuracy values for the three base models during the construction process of the ensemble classifier. As the epochs progress, a consistent decline in both training and validation loss can be observed, accompanied by a steady rise in accuracy until it converges. These results suggest that the models were effectively trained, exhibiting stable convergence behavior without signs of overfitting or underfitting throughout the process.

## 4. Results

### 4.1. Evaluation Metrics

To evaluate the model's ability to distinguish between real and bogus samples, the following performance metrics were employed:

*Specificity.* This metric assesses the model's ability to correctly identify bogus samples. A high specificity ensures that the model effectively reduces the misclassification of bogus as real, thereby avoiding unnecessary follow-up observations and resource consumption. It is mathematically expressed as:

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \quad (2)$$

where TN (True Negative) is the number of correctly classified bogus samples, and FP (False Positive) is the number of bogus misclassified as real. Similarly, TP (True Positive) refers to the number of correctly classified real and FN (False Negative) refers to the number of real misclassified as bogus.

*Recall.* This metric measures the model's ability to correctly identify positive samples. A high recall indicates that the model effectively minimizes the misclassification of true transients as bogus. This metric is particularly critical, as missing real samples may result in the loss of subsequent scientific targets. Its formula is as follows:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (3)$$

*Precision.* This metric quantifies the reliability of positive predictions. High precision ensures efficient allocation of observational resources. Its formula is as follows:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (4)$$

*F1-score.* The metric is the harmonic mean of precision and recall, evaluating the trade-off between minimizing missed transients and reducing bogus triggers. Its formula is as follows:

$$\text{F1-score} = \frac{2 \times \text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}} \quad (5)$$

*Area Under Curve (AUC).* The AUC serves as a quantitative metric of the model's overall classification performance. It is obtained by calculating the area under the ROC curve, which is

plotted with the true positive rate (TPR) on the vertical axis and the false positive rate (FPR) on the horizontal axis at various classification thresholds. A value closer to 1 indicates better model performance. The formulas for TPR, FPR, and AUC are as follows:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (6)$$

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (7)$$

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR}) d\text{FPR} \quad (8)$$

### 4.2. Classifier Performance

#### 4.2.1. Individual CNN Classifier

Table 2 summarizes 13 CNN performance metrics on the test set under a threshold of 0.5, sorted by AUC values from largest to smallest. From the results, the ResNet series emerged as top performers in comprehensive metrics, achieving leading AUC values (0.9946–0.9948) and F1-scores (87.22%–87.68%). Notably, their specificity reaches 99.70%, demonstrating exceptional capability in bogus sample identification. However, their recall (91.40%–91.74%) still needs to be improved for MST real-bogus classification scenarios. The EfficientNet series attained the highest recall rates (98.18%–98.68%) across all models. Nevertheless, extreme predictive bias toward real classes severely constrained their precision (21.82%–23.87%) and F1-scores (35.74%–38.29%). In contrast, the DenseNet series achieved a relatively balanced trade-off between recall (94.55%–94.71%) and specificity (98.60%–98.66%), but its overall performance still leaves room for improvement. The performance of the remaining models was generally inferior to the aforementioned three models. Particularly, ShuffleNetV2 and RepVGG exhibited suboptimal performance across multiple metrics, which suggests that their architectures may not be well-suited for the current task.

Furthermore, the experimental results indicate that there is no significant correlation between model performance and parameter count. Despite possessing 138M parameters, making it the model with the largest parameter set, VGGNet11 underperformed across multiple metrics compared to models with fewer parameters, such as ResNet. Moreover, different parameter versions within the same model family (e.g., ResNet18 and ResNet34, EfficientNet\_B0 and EfficientNet\_B2) exhibited comparable performance on key metrics, further highlighting the redundancy of model parameters. This suggests that, in the context of this task, selecting models with varying parameter scales within the same family has a limited impact on performance. Conversely, lower-parameter models offer advantages in terms of deployment and operational

**Table 2**  
Performance Comparison of 13 Different CNNs on MST Real-bogus Classification

Model	Parameter (M)	Specificity (%)	Recall (%)	Precision (%)	F1-score (%)	AUC
ResNet50	25.6	99.70	91.74	83.96	87.68	0.9948
ResNet18	11.7	99.70	91.40	83.92	87.50	0.9948
ResNet34	21.7	99.69	91.40	83.41	87.22	0.9946
EfficientNet_B0	5.3	94.09	98.68	22.14	36.16	0.9932
EfficientNet_B2	9.2	94.64	98.18	23.78	38.29	0.9929
DenseNet121	8.0	98.66	94.55	54.58	69.21	0.9928
EfficientNet_B1	7.8	93.98	98.68	21.82	35.74	0.9926
DenseNet169	14.1	98.60	94.71	53.55	68.42	0.9926
VGGNet11	138.0	97.46	92.07	38.18	53.97	0.9903
UniRepLKNet_A	4.4	98.36	90.25	48.40	63.01	0.9883
ViT_Tiny	5.7	91.45	94.05	15.77	27.02	0.9736
ShuffleNetV2_x0_5	1.4	99.64	12.89	38.05	19.26	0.7597
RepVGG_A0	8.3	94.05	26.12	6.95	10.98	0.5616

**Note.** Models are sorted in descending order of AUC to highlight high-performing architectures.

**Table 3**  
Performance Evaluation of Bagging Ensembles for MST Real-bogus Classification, Covering Individual Bagged Models and Combinations of ResNet18, DenseNet121, and EfficientNet\_B0

Model	Specificity (%)	Recall (%)	Precision (%)	F1-score (%)	AUC
ResNet18	99.85	91.90	91.00	91.45	0.9951
EfficientNet_B0	94.00	98.35	21.82	35.71	0.9928
DenseNet121	98.76	94.38	56.48	70.67	0.9926
DenseNet121, EfficientNet_B0	96.99	96.03	35.21	51.53	0.9938
EfficientNet_B0, ResNet18	98.33	95.04	49.23	64.86	0.9953
DenseNet121, ResNet18	99.30	94.54	69.76	80.28	0.9955
DenseNet121, EfficientNet_B0, ResNet18	99.25	95.37	68.45	79.70	0.9961

efficiency, presenting a more optimal choice for practical applications.

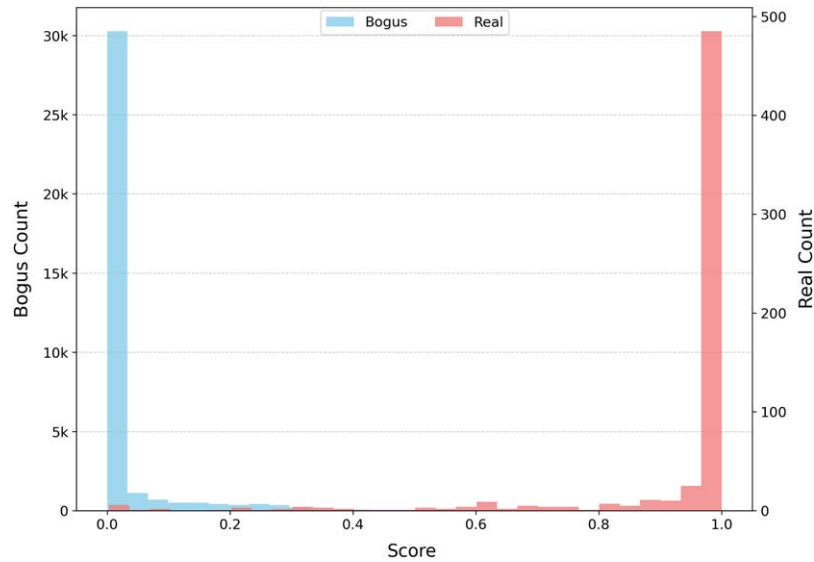
#### 4.2.2. Ensemble Classifier

To further develop a more robust classifier that meets the requirements of the MST project, we conducted bagging training with three models—ResNet18, EfficientNet\_B0, and DenseNet121—based on the experimental analysis presented in Section 4.2.1. Subsequently, we implemented ensemble learning through soft voting strategies with different model combinations and evaluated their performance metrics. Table 3 comprehensively documents the performance of individual retrained models and the ensemble results of various model combinations.


The experimental results demonstrate that the retrained models maintain strong consistency with previous training outcomes (record in Table 2), thereby validating the effectiveness of ensemble training. Notably, the tri-model ensemble configuration comprising DenseNet121, EfficientNet\_B0, and ResNet18 achieved the highest AUC value of 0.9961. This

finding indicates that this particular ensemble exhibits superior overall classification performance compared to all tested model combinations. Although the ensemble model shows slight decreases in F1-score (79.70%) and precision (68.45%) metrics relative to some other combinations, the significant improvement in AUC, coupled with maintained excellence in recall (95.37%) and specificity (99.25%), substantiates its enhanced discriminative power between real and bogus samples. Given that integrating heterogeneous architectures effectively enhance model generalization capabilities and robustness while ultimately achieving optimal overall performance, we ultimately selected the ensemble configuration of DenseNet121, EfficientNet\_B0, and ResNet18 as our preferred solution.

Figure 5 presents a bimodal histogram comparing the predicted score distributions of the ensemble model for the real (605 samples) and bogus (35,538 samples) classes in the test set. As depicted in the figure, the distributions of the real (red) and bogus (blue) samples exhibit a clear separation. The bogus samples form a sharp peak near 0, with over 95% concentrated within the 0–0.2 score range. Conversely, the real



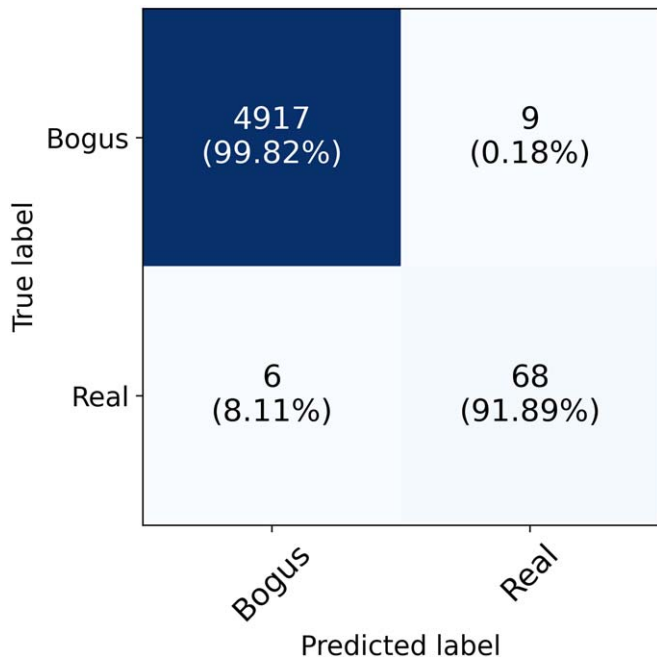
**Figure 5.** Histogram of ensemble model scores for 605 real samples and 35,538 bogus samples from the test set. The horizontal axis represents the prediction score, the left side of the vertical axis represents the number of bogus samples, and the right side represents the number of real samples.



**Figure 6.** Precision and recall curves of the ensemble model as a function of the classification threshold. The vertical dashed red line highlights the optimal threshold (0.798), where precision and recall both reach 0.922, indicating a balanced performance.

samples demonstrate a dominant peak near 1, with the majority of predictions clustering between 0.8 and 1.0. Only a minority of samples reside within the intermediate probability range of 0.2–0.8. This evident separation suggests the classifier is capable of making high-confidence classification decisions in most instances. However, the small number of samples in the intermediate range reflects the model’s uncertainty in handling boundary cases, suggesting potential directions for further optimization.

In binary classification tasks, the classification threshold is a critical parameter that determines how the model assigns samples to the positive or negative class. Figure 6 visualizes the variations in precision and recall of the ensemble model under different thresholds, thereby assisting in the selection of the optimal decision point. As illustrated in the figure, the ensemble model achieves its best balanced performance at a threshold of 0.798, where both precision and recall reach an impressive 0.922. This indicates that by setting the threshold to



**Figure 7.** Confusion matrix of the ensemble model on 5000 randomly selected real-world application samples. The ground truth is determined through manual inspection, while the predicted labels are generated by the classifier.

0.798, the model can effectively identify real samples with high precision while minimizing the risk of missing true positives. This not only further solidifies the robust capability of the ensemble model in accurately distinguishing between real and bogus samples but also significantly enhances its reliability in practical application scenarios. In the actual deployment of the MST project, 0.798 serves as a crucial reference for the initial threshold, which is dynamically adjusted around this value based on specific task requirements.

#### 4.3. Real-world Application

The ensemble classifier proposed in this study has been successfully deployed in the MST real-time data processing pipeline. Based on previous experimental analyses, we set the classification threshold at 0.798 to balance precision and recall in practical applications. To further validate the operational effectiveness of the classifier, we randomly selected 5000 samples from actual operations and conducted manual verification of the classification results. Among these samples, the classifier identified 77 as real and 4923 as bogus. Using the manually reviewed results as the ground truth, we obtained the confusion matrix shown in Figure 7, which reflects the model's classification performance in real deployment environments.

Based on calculations derived from this confusion matrix, the model achieved a precision of 88.31%, a recall of 91.89%, and a specificity of 99.82% during practical use. These results

indicate a slight degradation in performance relative to the optimal values observed during the validation phase. This discrepancy can be attributed to subtle differences between the distributions of the operational and training data, as well as variations in the proportion of genuine to bogus samples encountered in real-world scenarios. Nonetheless, the model maintains a high level of performance, effectively detecting true targets while mitigating the risk of false alarms.

In conclusion, the model's performance in real-world applications is sufficient to meet the operational requirements of the MST project, thereby providing robust technical support for its efficient and stable implementation. As the classifier continues to operate within the MST project, we will progressively accumulate more real-world operational data. These valuable data sets will be systematically incorporated to enhance and optimize our training corpus, thereby driving iterative improvements in model performance through continuous learning cycles.

## 5. Conclusions

In this study, we developed a deep learning-based ensemble classifier specifically designed to distinguish real transient sources from bogus detections in the MST. To support this work, we constructed a data set. The first data set, used for model training and evaluation, comprised 177,696 bogus samples and 3000 real samples, carefully selected through a combination of automated filtering and manual inspection.

Furthermore, we conducted a comprehensive evaluation of 13 CNN models and, based on the experimental results, adopted a bagging strategy to organically integrate ResNet18, EfficientNet\_B0, and DenseNet121 into an ensemble classifier. This classifier not only achieves an excellent AUC (0.9961) but also exhibits balanced performance, with a recall of 95.37% and a specificity of 99.25%. The prediction scores show a distinctly bimodal distribution: the vast majority of bogus samples cluster near 0, whereas real samples concentrate near 1, thereby effectively reducing decision uncertainty at the boundaries.

In addition, the classifier's performance in real-world environments further validates its practical applicability. Although there is a slight degradation when applied to actual MST observational data, the model still attains a precision of 88.31%, a recall of 91.89%, and a specificity of 99.82%. This outstanding performance robustly demonstrates the benefits of employing an ensemble of heterogeneous architectures.

Overall, the deep learning-based ensemble classifier that we developed effectively meets the operational requirements of the MST project. Looking ahead, as more data are incorporated and the model undergoes continuous iterative optimization, its performance and adaptability are expected to improve further. These ongoing enhancements are poised to consolidate its critical role in the MST real-time processing pipeline and ultimately advance in-depth scientific research into dynamic astrophysical phenomena.

## Acknowledgments

This work is supported by the National Key Basic R&D Program of China via 2023YFA1608303, the Strategic Priority Research Program of the Chinese Academy of Sciences (XDB0550103), and the National Natural Science Foundation of China under grant Nos. 12273076, 12133001, 12422303 and 12261141690.

The SiTian project is a next-generation, large-scale time-domain survey designed to build an array of over 60 optical telescopes, primarily located at observatory sites in China. This array will enable single-exposure observations of the entire northern hemisphere night sky with a cadence of only 30 minutes, capturing true color (gri) time-series data down to about 21 mag. This project is proposed and led by the National Astronomical Observatories, Chinese Academy of Sciences (NAOC). As a pathfinder for the SiTian project, the Mini-SiTian project utilizes an array of three 30 cm telescopes to simulate a single node of the full SiTian array. The Mini-SiTian has begun its survey since 2022 November. The SiTian and Mini-SiTian have been supported from the Strategic Pioneer Program of the Astronomy Large-Scale Scientific Facility, Chinese Academy of Sciences and the Science and Education Integration Funding of University of Chinese Academy of Sciences.

## ORCID iDs

Yang Huang  <https://orcid.org/0000-0003-3250-2876>  
Peng-Liang Du  <https://orcid.org/0009-0008-0836-6732>

## References

- Becker, A. 2015, HOTPANTS: High Order Transform of PSF ANd Template Subtraction, Astrophysics Source Code Library, ascl:1504.004
- Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393
- Breiman, L. 1996, *Mach. Learn.*, **24**, 123
- Cabrera-Vives, G., Reyes, I., Förster, F., Estévez, P. A., & Maureira, J.-C. 2017, *ApJ*, **836**, 97
- Chen, Z., Zhou, W., Sun, G., et al. 2023, arXiv:2309.09937
- Ding, X., Zhang, X., Ma, N., et al. 2021, in Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) (Washington, DC: IEEE), 13733
- Ding, X., Zhang, Y., Ge, Y., et al. 2024, in Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) (Washington, DC: IEEE), 5513
- Duev, D. A., Mahabal, A., Masci, F. J., et al. 2019, *MNRAS*, **489**, 3582
- Goldstein, D., D'andrea, C., Fischer, J., et al. 2015, *AJ*, **150**, 82
- Gu, H.-R., Huang, Y., Sun, Y.-K., et al. 2025, *RAA*, **25**, 044007
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (Washington, DC: IEEE), 770
- He, M., Wu, H., Ge, L., et al. 2025, *RAA*, **25**, 044005
- Hosenie, Z., Bloemen, S., Groot, P., et al. 2021, *ExA*, **51**, 319
- Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. 2017, in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (Washington, DC: IEEE), 4700
- Huang, Y., Liu, J., Wu, H., et al. 2025, *RAA*, **25**, 044001
- Killestein, T., Lyman, J., Steeghs, D., et al. 2021, *MNRAS*, **503**, 4838
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollar, P. 2017, in Proc. Int. Conf. Computer Vision (ICCV) (Washington, DC: IEEE), 2980
- Liu, J., Soria, R., Wu, X.-F., Wu, H., & Shang, Z. 2021, *AnABC*, **93**, 20200628
- Loshchilov, I., & Hutter, F. 2017, arXiv:1711.05101
- Ma, N., Zhang, X., Zheng, H.-T., & Sun, J. 2018, in Proc. European Conf. Computer Vision (ECCV) (Berlin: Springer), 116
- Mishkin, D., Sergievskiy, N., & Matas, J. 2017, *Computer Vision and Image Understanding*, **161**, 11
- Simonyan, K., & Zisserman, A. 2014, arXiv:1409.1556
- Tan, M., & Le, Q. 2019, in Proc. 36th Int. Conf. Mach. Learn. (ICML) (Long Beach, CA: PMLR), 6105