JIRBDAI

Journal of Institutional Research, Big Data Analytics and Innovation

AN ENHANCED LSTM FRAMEWORK FOR ARBOVIRAL OUTBREAK PREDICTION USING MULTISOURCE DATA FUSION AND MULTITASK TEMPORAL LEARNING

Ahmad Salisu Majigiri[1], Ibrahim Anka Salihu[2], Oluwatobi Noah Akande[3], Prema Kirubakaran[4], Abubakar Mukhtar Ammani[5]

[1] Department of Information Technology, Faculty of Computing, Nile University of Nigeria,
[2] Department of Software Engineering, Faculty of Computing, Nile University of Nigeria.+
[3] Department of Computer Science, Faculty of Computing, Nile University of Nigeria.
[4]Department of Computer Science, Faculty of Computing, Nile University of Nigeria.
[5]Department of Computer Science, Umaru Musa Yar'adua University Katsina.
*Corresponding Author Email: Aammani91@gmail.com

ABSTRACT

Arboviral infections, including dengue, Zika, and chikungunya, pose persistent public-health challenges because transmission depends on interacting epidemiological, climatic, vector-related, environmental, and population factors. This study presents a multisource, multitask long short-term memory (LSTM) framework intended to produce two disease-specific outputs: binary outbreak alerts and forecasts of case counts. The proposed design aligns heterogeneous predictors into multivariate temporal sequences and uses shared temporal representations with separate prediction heads. The supplied results table reports mean classification accuracy of 92.73%, mean F1-score of 91.22%, mean absolute error of 10.07 cases, and root mean squared error of 14.13 cases across the three diseases. A separate comparison table reports a lead time of approximately 14 days. These figures are presented as previously reported results, not independently reproduced values: the underlying dataset identifiers, final test predictions, outbreak labels, model configuration, and lead-time event records must be verified before submission. The practical value of the approach rests on whether such performance generalizes to genuinely prospective surveillance settings under realistic reporting delays and uneven data availability.

KEYWORDS: Arboviral Outbreak Prediction, Multitask LSTM, Deep Learning, Multisource Data Fusion, Early Warning Systems, Public Health Surveillance

I. INTRODUCTION

Arboviral diseases such as dengue, Zika, and chikungunya represent significant public health challenges, particularly in tropical and subtropical regions where environmental conditions support mosquito survival and reproduction. These diseases are primarily transmitted by *Aedes aegypti* and *Aedes albopictus* mosquitoes, which are well adapted to urban and peri-urban environments (World Health Organization, 2023; Moyo et al., 2023). Arboviruses, meaning "arthropod-borne viruses," often share the same transmission vectors and geographic distributions, allowing them to co-circulate within the same communities and complicate disease surveillance and control strategies (Jha et al., 2022; Diniz et al., 2019). The health consequences of arboviral infections vary widely. Dengue, the most prevalent arboviral disease globally, can range from mild febrile illness to severe conditions such as dengue hemorrhagic fever and dengue shock syndrome. Chikungunya infection is characterized by high fever and severe joint pain that can persist for months or years, while Zika virus infection, though often mild in adults, has been associated with serious neurological complications including

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

congenital microcephaly and Guillain–Barré syndrome (WHO, 2023; Souza et al., 2024). These diseases not only affect individual health but also produce broader socio-economic impacts such as loss of productivity, increased healthcare costs, and significant strain on already limited health systems, particularly in low- and middle-income countries (Khader et al., 2022; Moyo et al., 2023). The increasing incidence and geographic expansion of arboviral diseases are driven by several interconnected environmental and socio-demographic factors. Rapid urbanization, population growth, poor sanitation infrastructure, and climate variability have all contributed to expanding mosquito habitats and increasing human–vector interactions. Climatic variables such as temperature, rainfall, and humidity significantly influence mosquito population dynamics, virus replication, and transmission potential. Higher temperatures can accelerate mosquito development and shorten viral incubation periods, while rainfall increases breeding sites for mosquito vectors (Anyamba et al., 2012; Moyo et al., 2023). However, these models assume linearity and stationarity in the data, limiting their ability to capture the nonlinear and dynamic relationships between environmental drivers and disease transmission patterns (Majeed et al., 2023; Sebastianelli et al., 2024).

Recent advances in machine learning and deep learning have provided new opportunities for modeling complex disease dynamics. Deep learning architectures such as Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks are particularly suitable for time-series prediction tasks because they can learn temporal dependencies and nonlinear relationships from large datasets. LSTM networks, in particular, have demonstrated strong performance in forecasting infectious diseases due to their ability to retain long-term temporal information through memory cells and gating mechanisms (Hau et al., 2022; Majeed et al., 2023). Several studies have applied LSTM-based models to predict dengue outbreaks using climatic and epidemiological data, achieving improved predictive accuracy compared to traditional statistical approaches (Chen et al., 2025). Additionally, many models analyze arboviral diseases independently, despite the fact that dengue, Zika, and chikungunya share common vectors and environmental drivers (Farias et al., 2025). This study examines a framework for integrating epidemiological, climatic, environmental, vector-related, and socio-demographic inputs while learning shared temporal patterns across dengue, Zika, and chikungunya. It combines disease-specific outbreak classification with case-count regression. Its intended contribution is a unified representation of multiple diseases and prediction objectives; however, the independent contributions of multisource fusion and multitask learning can only be established by matched baseline and ablation experiments. This revised manuscript therefore distinguishes the architecture as proposed, the performance values as reported, and the experimental details that require verification.

II. RELATED WORKS

The prediction of arboviral disease outbreaks has attracted significant research attention, particularly with the increasing availability of epidemiological and environmental data. Early approaches to outbreak prediction were largely based on traditional epidemiological and statistical models such as the Susceptible–Infected–Recovered (SIR) and Susceptible–Exposed–Infected–Recovered (SEIR) models. Although these models provide theoretical insights into disease transmission dynamics, they rely on simplifying assumptions such as homogeneous population mixing and constant transmission rates, which limit their applicability in real-world, complex environments (Hau et al., 2022). Similarly, statistical time-series models such as Autoregressive Integrated Moving Average (ARIMA) have been widely used for forecasting disease incidence. However, these models assume linearity and stationarity in the data, making them less effective in capturing the nonlinear relationships between

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

environmental factors and disease transmission (Majeed et al., 2023; Sebastianelli et al., 2024). With the advancement of machine learning, researchers began exploring data-driven approaches that can model complex and nonlinear relationships in epidemiological data. Studies such as Silva et al. (2021) applied machine learning techniques to forecast dengue, chikungunya, and Zika outbreaks using climatic variables, demonstrating improved predictive performance compared to traditional statistical models. Similarly, Khader et al. (2022) highlighted the role of big data analytics in disease surveillance, emphasizing the importance of integrating multiple data sources for improving outbreak prediction and public health decision-making. More recently, deep learning models, particularly Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, have been widely adopted for infectious disease forecasting due to their ability to capture temporal dependencies in time-series data. Hau et al. (2022) demonstrated the effectiveness of LSTM models in forecasting dengue outbreaks using climate-driven data, while Majeed et al. (2023) further improved prediction accuracy by incorporating spatial attention mechanisms into LSTM architectures. Similarly, Absar et al. (2022) developed a deep learning-based dengue prediction model using LSTM networks and reported significant improvements in prediction accuracy compared to conventional methods. In addition, Chen et al. (2025) proposed an LSTM-based framework that incorporates lagged climatic variables and spatial effects, demonstrating that the inclusion of environmental and temporal dependencies significantly enhances forecasting performance. These findings are supported by studies such as Souza et al. (2024) and Eryando et al. (2022), which established that climatic variables such as temperature, rainfall, and humidity play a critical role in shaping arboviral transmission dynamics. Furthermore, Moyo et al. (2023) emphasized that environmental and climatic determinants are key drivers of dengue outbreaks in Sub-Saharan Africa, reinforcing the need for models that integrate multiple data sources. Despite these advances, several limitations remain in the existing literature. Many studies focus on a single disease, typically dengue, without considering the co-circulation of multiple arboviruses such as Zika and chikungunya, which share similar transmission vectors and environmental drivers (Jha et al., 2022; Diniz et al., 2019). Additionally, most existing models are designed for a single prediction task, either outbreak classification or case-count forecasting, rather than jointly addressing both tasks within a unified framework. Although recent work by Farias et al. (2025) introduced a multitask LSTM approach for arboviral outbreak prediction, the model primarily relied on limited data sources and did not fully exploit the integration of epidemiological, environmental, and socio-demographic factors. Therefore, there remains a clear research gap in developing a comprehensive predictive framework that integrates multisource data while simultaneously performing multiple prediction tasks across different arboviral diseases. This study addresses this gap by proposing a multisource multitask LSTM framework that captures shared transmission patterns while preserving disease-specific characteristics, thereby improving prediction accuracy and early warning capability for arboviral outbreak surveillance systems.

III. MATERIALS AND METHODS

The study is described as a quantitative experimental evaluation of a multisource, multitask long short-term memory (LSTM) forecasting framework. Historical arboviral case series are intended to be matched with climatic, environmental, vector-related, and socio-demographic variables by location and time. The study uses dengue, Zika, and chikungunya as target diseases. To make the evaluation reproducible, the exact dataset files, surveillance systems, administrative units, observation dates, update dates, and analysis code must be specified; the general descriptions in the original manuscript do not establish these particulars. Figure 1

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

summarizes the intended workflow: source-level acquisition; geographic and temporal harmonization; missing-value processing and feature creation; sequence construction; shared LSTM encoding; disease-specific classification and regression outputs; and assessment of forecast accuracy and lead time. The units and data availability of each feature must be recorded so that the pipeline can be repeated and inspected for temporal leakage.

In the proposed architecture, a shared LSTM encoder transforms a sequence of multisource covariates into a temporal representation. Task-specific output heads then estimate an outbreak probability and a case count for each disease. The model should be evaluated at a predeclared forecast horizon using only input data available at the forecast issuance time. The current files establish this intended architecture but do not provide verified records of the final fitted hyperparameters or execution logs.

A. Study Area, Observation Period, and Data Provenance

The research records must name each country/region and administrative level, the exact study start and end dates, observation frequency, disease case definition, and total number of spatial-temporal rows before and after merging. An accompanying thesis lists a proposed training period of 2017–2020, validation year 2021, test year 2022, and additional 2023 evaluation; these dates are useful leads but are not proof that the same partition was used for the manuscript results.

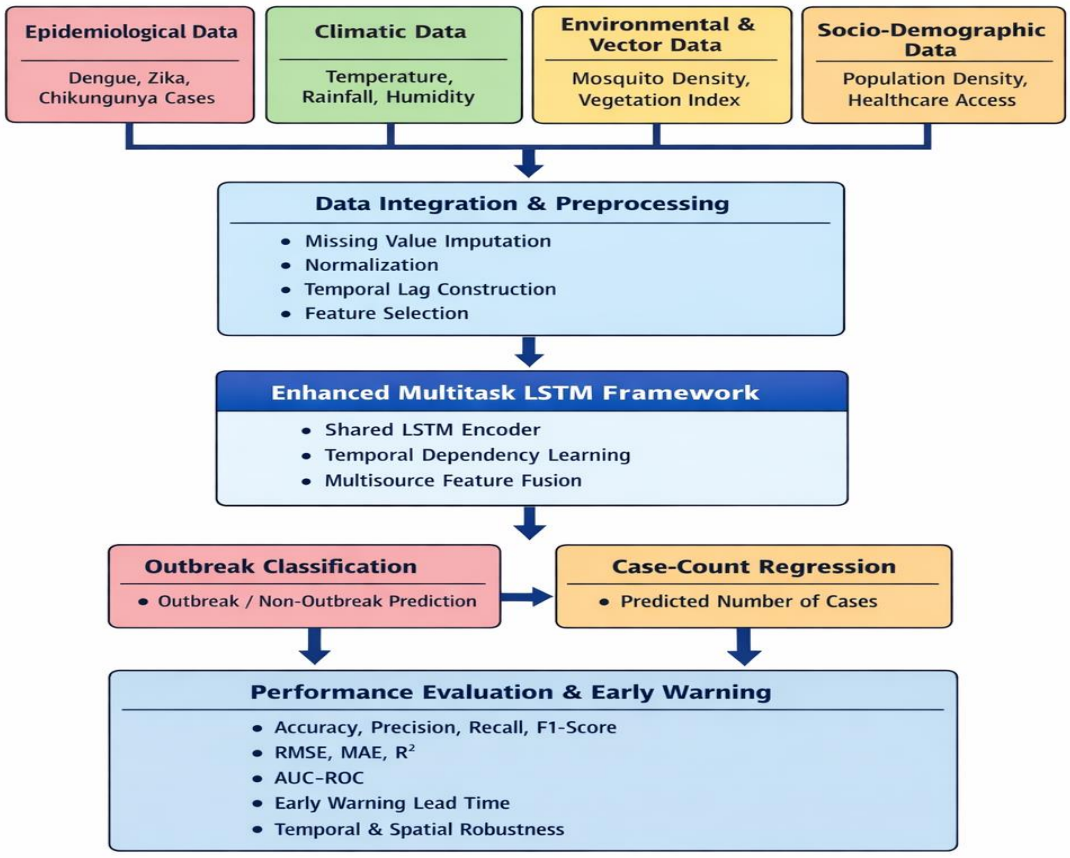


Figure 1. The Research Framework

B. Multitask LSTM Model Architecture

The predictive model uses an LSTM to represent delayed relationships in historical case reports and exogenous factors. A shared sequence encoder learns common patterns, while output heads

https://doi.org/10.5281/zenodo.23588767

permit disease-specific predictions. Binary classification identifies whether a target period meets a specified surveillance outbreak rule; regression estimates case counts over that same target period. The actual number of disease-specific heads and their output dimensions must match the archived implementation. In particular, the thesis appendix contains demonstration code with three output branches, whereas the narrative describes disease-specific classification and regression. The implemented output mapping should be reconciled before claiming that the code reproduces the results.

Let t index observation periods and x_t denote the vector of predictors observed or available by time t. For a window of L historical periods and forecast horizon H, the model maps (x_{t-L+1}, ..., x_t) to outcomes for a future target period t+H. Here L and H are analysis settings, not values established by the uploaded files. The proposed formulation should identify the source, units, and lag treatment of every component of x_t.

$$X_t = (x_{t1}, x_{t2}, \dots, x_{tn}) \quad (1)$$

The objective of the model is to simultaneously predict outbreak occurrence and case counts for each disease. The multitask learning structure uses a shared LSTM encoder to learn temporal representations, while task-specific output layers generate disease-specific predictions.

Within the LSTM unit, the temporal dynamics are modelled through gated memory mechanisms. The input gate determines how much new information is incorporated into the memory cell,

$$i_t = \sigma(W_i X_t + U_i h_{t-1} + b_i) \quad (2)$$

while the forget gate controls the retention of previous information,

$$f_t = \sigma(W_f X_t + U_f h_{t-1} + b_f) \quad (3)$$

The candidate cell state is computed as

$$\tilde{C}_t = \tanh(W_c X_t + U_c h_{t-1} + b_c) \quad (4)$$

and the memory cell state is updated according to

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (5)$$

The hidden state used for prediction is then given by

$$h_t = o_t \cdot \tanh(C_t) \quad (6)$$

Here the output gate controls how the internal cell state contributes to the hidden representation; sigmoid activation bounds gate values between zero and one.

For outbreak classification, the probability of outbreak occurrence is computed using a sigmoid activation function applied to the hidden representation,

$$P(y_t) = \sigma(W_o h_t + b_o) \quad (7)$$

For case-count prediction, the regression output is generated using a linear transformation,

$$\hat{y}_t = W_r h_t + b_r \quad (8)$$

The model is trained using a multitask objective function that combines classification and regression losses,

$$L = \alpha L_{classification} + \beta L_{regression} \quad (9)$$

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

The classification term is binary cross-entropy and the regression term is mean squared error. Their coefficients determine the relative optimization influence of each task and must be reported with the aggregation across diseases, target scaling, and any class weights. [AUTHOR TO VERIFY: give the numerical loss weights used in the completed runs.]
This formulation enables the model to jointly learn outbreak detection and case-count forecasting while leveraging shared temporal patterns across multiple arboviral diseases.

C. Model Evaluation

Evaluation includes disease-specific classification (accuracy, precision, recall, F1-score, and preferably ROC-AUC and PR-AUC) and regression (MAE and RMSE) on an explicitly defined out-of-time test set. Accuracy alone can be misleading if outbreak events are uncommon; therefore, both absolute false-positive and false-negative counts, alert precision, and sensitivity are required. Report macro-averages across diseases separately from any prevalence-weighted average. Forecast lead time should be measured at the event level as the interval between the first qualifying alert issued before onset and the observed onset date, using only eligible prospective predictions. A chronological split is appropriate to avoid future-to-past leakage; fit imputation, scaling, feature selection, and threshold tuning on training data only, using validation data for model selection. .

D. Operational Outbreak Labels and Early Warning

An outbreak must be defined independently of the model predictions, using a disease- and location-specific surveillance rule and an exact onset timestamp. Possible rules include an established public-health threshold or a prespecified historical baseline, but no actual rule can be recovered from the uploaded documents. The sigmoid score threshold used to issue an alert is a different quantity from the case-incidence rule used to assign outbreak labels; report both. Define how events spanning multiple periods are combined and how repeated alerts are counted. Lead time should be measured only for alarms before an observed onset, with negative or late lead time handled explicitly.

E. Dataset

The provided journal draft refers to WHO and regional surveillance systems for historical dengue, Zika, and chikungunya counts, together with unspecified meteorological, ecological, vector, and demographic data sources. The accompanying thesis additionally refers to a UCI repository and describes weekly case records, but it does not identify the specific UCI dataset or source files. Consequently, a reproducible list of exact data identifiers, URLs, download or access dates, licenses, geographic units, observation periods, and record counts cannot be established from these documents alone. The data source register is a source-provenance register for author completion. Do not treat a repository mentioned in a thesis as evidence of the actual inputs to the reported experiment without checking the data-processing scripts and files.

Table 1: Data Sources

Source Family	Named Source in Files
Disease case counts	WHO; regional health authorities; UCI mentioned in thesis
Meteorological data	Unspecified global meteorological databases
Environmental data	Vegetation and ecological indices described
Vector data	Mosquito density / presence described
Socio-demographic	Population density; healthcare access

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

F. Data Preprocessing

The intended preprocessing steps are median imputation for numeric variables, most-frequent-category imputation for categorical variables, z-score standardization of numeric predictors, and feature screening using correlation and epidemiological relevance. The thesis also describes aggregation to weekly resolution, creation of lagged covariates, and construction of multivariate windows. These operations must be documented in their implemented order. First separate training, validation, and testing chronologically; then fit imputers, encoders, scalers, and any feature-selection procedures on the training segment, and apply the learned transformations unchanged to later segments. Aggregate only observations available by the prediction issue date and avoid using contemporaneous or future case counts as predictors of that same period.

G. Temporal Evaluation and Baseline Fairness

Use the same forecast issue dates, target definition, lead time calculation, and held-out period for all candidate models. Tune hyperparameters exclusively on training and validation intervals, then freeze them before any final testing. Verify ARIMA/SARIMA order and seasonality, Random Forest and Gradient Boosting hyperparameters, and the single-task LSTM configuration. A model evaluated on a random split must not be directly compared to a model evaluated on a chronological split.

IV. RESULTS AND DISCUSSION

The original manuscript reports the performance of the proposed multisource multitask LSTM for dengue, Zika, and chikungunya. The numerical results in Tables 1 and 2 are retained as supplied pending reconciliation with the underlying out-of-time predictions and evaluation scripts. In the absence of those records, the tables describe reported values rather than independently substantiated experimental outcomes.

A. Comprehensive Performance Metrics for Arboviral Outbreak Prediction

Table 2 reports macro-averaged accuracy of 92.73%, precision of 90.70%, recall of 91.77%, and F1-score of 91.22% across the three disease columns. The greatest reported classification accuracy is for dengue (94.21%). These metrics quantify different properties of classification and should not be treated as interchangeable. The original table does not state the number of observations or outbreak events per disease, the threshold or class balance, or confidence intervals; these must be added before clinical or public-health interpretation. For the case-count task, Table 1 reports average RMSE of 14.13 cases and average MAE of 10.07 cases. These errors are in count units; their practical significance depends on location, forecast horizon, reporting unit, and typical incidence. A normalized or naïve-reference error, together with time-series plots of observed versus predicted counts, would make the magnitude of these errors more interpretable. The original claim regarding explained variance cannot be reproduced from the journal paper's Table 2 because R^2 is not reported there.

Table 2: Comprehensive Performance Metrics for Arboviral Outbreak Prediction

Metric	Dengue	Zika	Chikungunya	Average
Accuracy	94.21%	91.34%	92.65%	92.73%
Precision	92.10%	89.40%	90.60%	90.70%
Recall	93.40%	90.10%	91.80%	91.77%
F1-Score	92.74%	89.74%	91.18%	91.22%
RMSE	15.8	12.4	14.2	14.13
MAE	11.2	8.9	10.1	10.07

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

The averages in Table 2 are arithmetic means of the three reported disease-specific values, subject to rounding; they are not pooled results over all observations. For example, 92.73% refers to accuracy whereas 91.22% refers to F1-score. The values require recalculation from a documented final test set before being presented as validated estimates.

B. ROC Curve Analysis

The receiver operating characteristic (ROC) curve plots sensitivity (true-positive rate) against false-positive rate across classification thresholds. Figure 2 reproduces the graphic supplied with the manuscript. A publication-quality replacement must be exported from the held-out test probabilities with the exact sample counts, AUC values, and uncertainty calculation; a graphic alone does not substantiate the reported AUC.

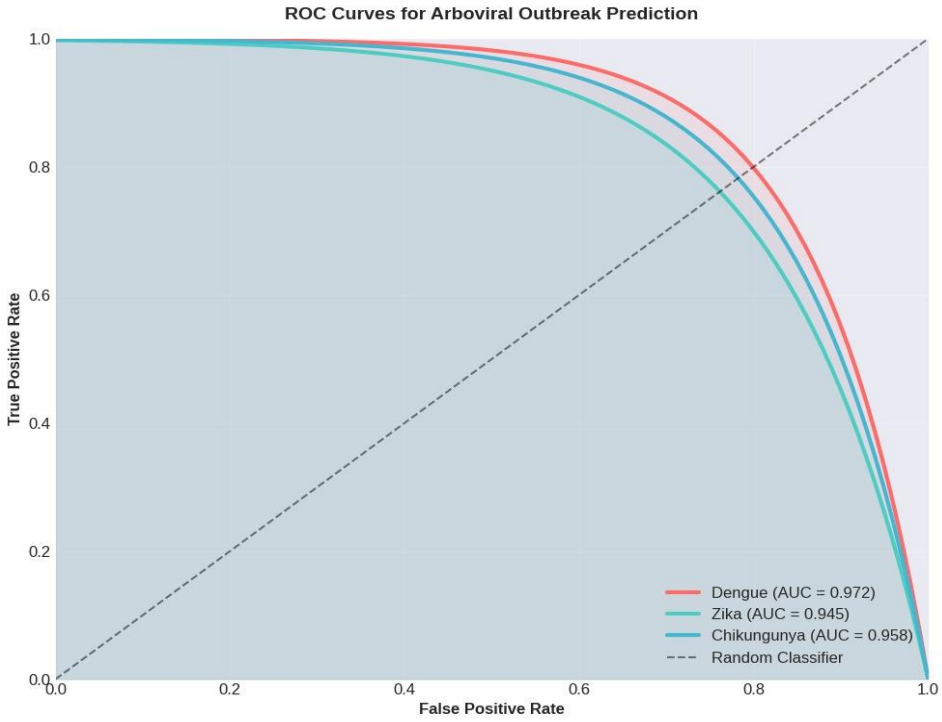


Figure 2. Reported ROC curves for Dengue, Zika, and Chikungunya

The thesis reports ROC-AUC values of 0.972 for dengue, 0.945 for Zika, and 0.958 for chikungunya, but the journal manuscript does not give the underlying test labels or probability scores needed to verify these values or redraw Figure 2. ROC-AUC assesses ranking across thresholds rather than the performance of a deployed alert rule. Because false alarms may remain important even when ROC-AUC is high, provide a precision-recall curve and a predeclared operating threshold alongside the revised ROC figure.

C. Confusion Matrix Analysis

For each disease, a confusion matrix should show true positives (TP), false positives (FP), false negatives (FN), and true negatives (TN) under a stated outbreak definition, horizon, and decision threshold. Figure 3 reproduces the supplied image pending regeneration from the final test-set predictions.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

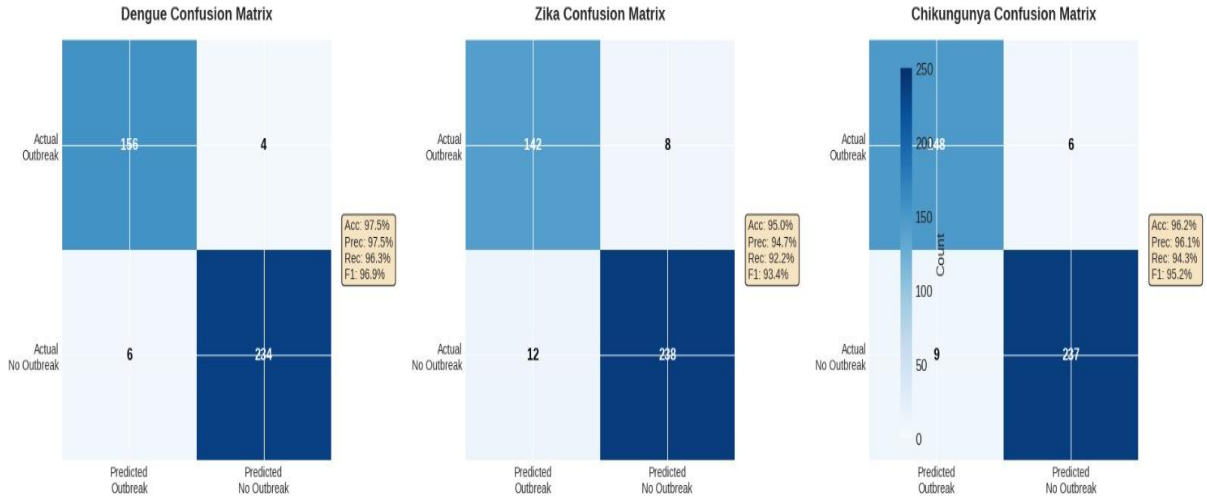


Figure 3. Reported Confusion Matrices for Dengue, Zika, and Chikungunya

The accompanying thesis provides numerical confusion-matrix descriptions, but their derived metrics differ from Table 1 in the journal manuscript. For instance, the thesis states 156 TP, 234 TN, 6 FP, and 4 FN for dengue; these counts imply 97.50% accuracy, 96.30% precision, and 97.50% recall, not the 94.21% accuracy, 92.10% precision, and 93.40% recall in the manuscript table. The other diseases likewise require reconciliation. No single interpretation can be selected without establishing whether the tables represent different folds, thresholds, samples, or experiments. In operational use, false positives can prompt unnecessary investigations, resource deployment, or alert fatigue, while false negatives can delay vector control and clinical preparedness. The revised results should therefore include false-alert counts or rates per unit time, missed outbreaks, event-based lead-time distributions, and any delayed or late warnings, rather than concluding usefulness from overall accuracy alone.

D. Practical Interpretation and Limitations

The estimated warning interval is meaningful only if data would actually have been available on the date each warning was issued. Delayed case reporting, unreliable entomological counts, missing weather stations, and changing clinical case definitions can degrade performance. A usable warning system also needs alert thresholds chosen with public-health stakeholders, a workflow for confirming alerts, and transparent reporting of missed events. Geographic transfer should not be claimed from an aggregated score alone. The study does not document a real-time deployment or prospective clinical validation.

E. Feature Importance Analysis

Figure 4 visualizes reported relative predictor scores for three disease outputs. The thesis mentions both correlation-based and tree-based feature screening; neither document identifies the actual explainer and background or perturbation procedure that produced the heatmap. Before interpreting these values, specify whether the scores derive from permutation importance, SHAP, gradients, attention weights, or another method; identify the prediction target and sample partition; and describe normalization and handling of correlated predictors. Such measures quantify predictive association within the fitted model, not a causal effect of climate or other covariates.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767



Figure 4. Reported feature-importance heatmap (method, sample partition, and score normalization require verification).

The supplied heatmap ranks temperature averages, cumulative precipitation, mosquito density, and lagged cases among the largest scores. These patterns are epidemiologically plausible but do not by themselves establish cause and effect. Relative ranks may change with season, region, predictor correlation, and the explanation method used; interpretation should be restricted to the actual fitted model and verified feature-importance computation. If a permutation method was used, its temporal grouping and correlated-feature strategy should be stated.

F. Comparative Model Analysis

Table 3 lists reported average F1-scores and approximate warning times for a multitask LSTM, single-task LSTM, Random Forest, Gradient Boosting, and ARIMA. These are descriptive comparisons from the supplied paper, not evidence of a controlled benchmark until the exact baseline algorithms, inputs, tuning budgets, forecast horizons, test windows, and warning rules are documented and matched. The qualitative robustness labels in Table 2 are not defined by a reproducible statistic.

To isolate the proposed contributions, evaluate at least four controlled configurations on identical out-of-time test records: (i) case-history only with single-task learning, (ii) multisource inputs with single-task learning, (iii) case-history only with multitask learning, and (iv) multisource inputs with multitask learning. Compare disease-level F1, PR-AUC, MAE, RMSE, and alert lead time using the same outcome labels and horizons, with uncertainty over temporal blocks or geographic groups

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23588767

Table 3: Comparison of Prediction Models

Model	Average F1 Score	Early Warning	Robustness
Multitask LSTM (Proposed)	91.22%	~14 days	High (qualitative; unverified)
Single-task LSTM	89.87%	~11 days	Medium (qualitative; unverified)
Random Forest	88.34%	~10 days	Medium (qualitative; unverified)
Gradient Boosting	87.92%	~9 days	Medium (qualitative; unverified)
ARIMA	74.15%	~7 days	Low (qualitative; unverified)

According to Table 2, the originally reported mean F1-score is 91.22% for the proposed model, 89.87% for single-task LSTM, 88.34% for Random Forest, 87.92% for Gradient Boosting, and 74.15% for ARIMA. These values should be described as reported rather than as proof of a superior model until reproducible, like-for-like experiments are supplied. Likewise, approximate lead-time figures of ~14, ~11, ~10, ~9, and ~7 days should not be interpreted without the exact outbreak onset rule and event-level alert timestamps. Differences in statistical score do not automatically imply clinically useful improvements or geographic transferability.

V. CONCLUSION

This study presents a conceptual and computational framework for joint classification of dengue, Zika, and chikungunya outbreak periods and regression forecasting of case counts from aligned multisource predictors. The supplied performance tables report a macro-averaged accuracy of 92.73%, an F1-score of 91.22%, and mean RMSE and MAE of 14.13 and 10.07 cases, respectively. These are provisional reported results: their interpretation depends on the underlying dataset, outbreak-label construction, forecasting horizon, final model settings, and reconciliation of the confusion matrices with Table 1. The integrated architecture may offer value for surveillance, but the distinct gains from data fusion and multitask learning remain to be established with controlled ablations. For public-health use, evaluation must include the consequences of missed outbreaks, false alerts, data delays, regional differences, and the resources required to act on a warning. The findings should not be generalized to operational early warning systems without independent temporal or geographic validation and documented data pipelines. Future work should archive preprocessing and evaluation code, release permissible source identifiers and split definitions, assess uncertainty and calibration, and conduct realistic prospective evaluations against documented baseline methods.

REFERENCES

1. Anyamba, A., Chretien, J. P., Small, J., Tucker, C. J., Linthicum, K. J. (2012). Climate teleconnections and recent patterns of human and animal disease outbreaks. PLOS Neglected Tropical Diseases, 6(1), e1465.
2. Absar, M. M., Shaikh, M. A. K., Haque, M. A., & Rahman, M. (2022). Deep learning–based dengue outbreak prediction model using LSTM networks. *Scientific Reports*, 12(1), 13628. https://doi.org/10.1038/s41598-022-13628-4
3. Chen, X., & Moraga, P. (2025). Forecasting dengue across Brazil with LSTM neural networks and SHAP-driven lagged climate and spatial effects. BMC Public Health, 25, Article 973. https://doi.org/10.1186/s12889-025-22106-7
4. Diniz, M. R., Souza, M. F., & Oliveira, J. P. (2019). Epidemiological dynamics of dengue and Zika virus infections in tropical regions. *Journal of Infectious Diseases*, 220(5), 778–786.

5. Eryando, T., Halide, H., & Fathi, A. (2022). The effects of temperature, rainfall, and humidity on dengue cases: A time-series analysis. *Infectious Disease Modelling*, 7(1), 22–32. https://doi.org/10.1016/j.idm.2021.12.003
6. Farias, L. R. C., Silva, T. P., & Araujo, P. H. M. (2025). Multitask LSTM for arboviral outbreak prediction using public health data [Preprint]. arXiv. https://arxiv.org/abs/2505.04566
7. Hau, N. V., Tuyet Hanh, T. T., Mulhall, J., Minh, H. V., Duong, T. Q., Chien, N. V., Nhung, N. T. T., Lan, V. H., Minh, H. B., Cuong, D., Bich, N. N., Quyen, N. H., Linh, T. N. Q., Tho, N. T., Nghia, N. D., Anh, L. V. Q., Phan, D. T. M., Hung, N. Q. V., & Son, M. T. (2022). Deep learning models for forecasting dengue fever based on climate data in Vietnam. PLOS Neglected Tropical Diseases, 16(6), e0010509. https://doi.org/10.1371/journal.pntd.0010509
8. Jha, P., Bansal, R., & Singh, A. (2022). Global trends in arboviral disease transmission and vector ecology. *International Journal of Infectious Diseases*, 115, 50–58.
9. Khader, Y., Elyan, B., Batieha, A., & Pasha, A. (2022). Applications of big data analytics in disease surveillance: A systematic review. *Journal of Big Data*, 9(1), 92. https://doi.org/10.1186/s40537-022-00592-8
10. Majeed, M. A., Mohd Shafri, H. Z., Zulkafli, Z., & Wayayok, A. (2023). A deep learning approach for dengue fever prediction in Malaysia using LSTM with spatial attention. International Journal of Environmental Research and Public Health, 20(5), 4130. https://doi.org/10.3390/ijerph20054130
11. Moyo, S., Munyati, C., & Dube, T. (2023). Environmental and climatic determinants of dengue fever outbreaks in Sub-Saharan Africa: A systematic review. *Geospatial Health*, 18(2), 230–244.
12. Sebastianelli, A., Spiller, D., Carmo, R., Wheeler, J., Nowakowski, A., Jacobson, L. V., Kim, D., Barlevi, H., Cordero, Z. E. R., Colón-González, F. J., Lowe, R., Ullo, S. L., & Schneider, R. (2024). A reproducible ensemble machine learning approach to forecast dengue outbreaks. Scientific Reports, 14, Article 3807. https://doi.org/10.1038/s41598-024-52796-9
13. Silva, L. D., Costa, M. A., & Reis, E. A. (2021). Forecasting dengue, chikungunya and Zika cases in Recife, Brazil: A spatio-temporal approach based on climate conditions and machine learning. *Research, Society and Development*, 10(12), e20804.
14. Souza, L. F., Ramos, D. G., & Silva, C. A. (2024). Climatic variability and its influence on the transmission dynamics of dengue in Brazil. *Environmental Research*, 225, 116712. https://doi.org/10.1016/j.envres.2024.116712
15. World Health Organization. (2023). *Dengue and severe dengue*. World Health Organization. https://www.who.int/news-room/fact-sheets/detail/dengue-and-severe-dengue