

An Explainable Agentic Artificial Intelligence Framework for Early Prediction and Decision Support in Infectious Disease Outbreaks: A Case Study of Lassa Fever in Ebonyi State, Nigeria

Jeremiah, Chukwu¹, Igwe, Sylvester Agbo¹, Nworu, Ursula Amaka¹, Ugwu, Gabriel Evo¹

¹Department of Computer Science, Ebonyi State University, Abakaliki - Nigeria
Corresponding Author: chukwu.jeremiah@ebsu.edu.ng

Abstract

Lassa fever remains a significant public health challenge in Nigeria and other West African countries, with recurrent outbreaks requiring timely surveillance, early detection, and rapid intervention. Conventional surveillance approaches are often dependent on passive reporting and laboratory confirmation, which may limit the ability of health authorities to identify emerging outbreaks at an early stage. Recent advances in artificial intelligence (AI), machine learning (ML), and deep learning have created opportunities for infectious disease prediction through the analysis of epidemiological, environmental, climatic, demographic, and geospatial data. However, many existing AI-based approaches focus primarily on predictive accuracy and provide limited explanation of the factors responsible for model predictions. Furthermore, existing predictive systems generally operate as standalone models rather than autonomous, coordinated systems capable of integrating heterogeneous information, reasoning over evidence, explaining predictions, and supporting public health decision-making. This paper proposes an Explainable Agentic Artificial Intelligence (EAAI) framework for early prediction and decision support in Lassa fever outbreaks in Ebonyi State, Nigeria. The proposed framework integrates specialized AI agents for epidemiological surveillance, environmental analysis, climate assessment, geospatial intelligence, predictive modeling, evidence retrieval, explainability, and decision support. The framework is designed to combine heterogeneous spatio-temporal data and generate localized outbreak-risk predictions accompanied by interpretable and evidence-grounded explanations. The proposed research is expected to contribute a novel architecture for trustworthy and human-centered AI-driven infectious disease surveillance. The study argues that the integration of agentic AI and explainable AI may provide a more transparent, adaptive, and actionable approach to infectious disease early warning and public health decision support.

Keywords: Agentic AI, Explainable Artificial Intelligence, Lassa Fever, Machine Learning, Infectious Disease Prediction, Decision Support, Public Health, Ebonyi State, Nigeria.

Date of Submission: 17-09-2026

Date of Acceptance: 01-10-2026

I. Introduction

Infectious disease outbreaks continue to pose major threats to global health, particularly in low- and middle-income countries where health surveillance infrastructure, diagnostic capacity, and response resources may be constrained. The increasing frequency and complexity of emerging and re-emerging infectious diseases have intensified the need for effective early warning systems capable of identifying potential outbreaks before they escalate into major public health emergencies. Artificial intelligence (AI) and machine learning (ML) have emerged as important technologies for infectious disease surveillance because of their ability to process large volumes of heterogeneous data and identify complex patterns that may be difficult to detect using conventional statistical methods (Melchane *et al.*, 2024).

Lassa fever is an acute viral hemorrhagic disease endemic to several countries in West Africa, including Nigeria. The disease is primarily associated with the rodent reservoir *Mastomys natalensis*, although human-to-human transmission may also occur, particularly in healthcare and household settings. The epidemiology of Lassa fever is influenced by multiple interacting factors, including environmental conditions, climatic variables, human settlement patterns, population characteristics, and ecological relationships. Consequently, effective early prediction requires the analysis of multiple data dimensions rather than relying exclusively on historical case counts.

Recent research has demonstrated the potential of AI and ML techniques for Lassa fever prediction. For example, Osaseri *et al.* (2025) investigated an AI-driven Lassa fever prediction approach based on evolutionary feature selection and Random Forest techniques, demonstrating the increasing application of machine learning to Lassa fever surveillance. Similarly, recent research using Self-Organizing Maps has explored the integration of environmental and settlement variables to improve Lassa fever risk mapping across West Africa (Acta Tropica, 2026). These studies demonstrate that computational intelligence can contribute to improved understanding and prediction of Lassa fever risk.

The importance of environmental and climatic information has also been demonstrated in recent research involving high-risk Nigerian states. Fuwape *et al.* (2026) analyzed weekly Lassa fever cases from 2018 to 2022 alongside ERA5 climate data and found significant relationships between Lassa fever incidence and several meteorological and environmental variables. Their findings suggest that climate-informed models may provide valuable information for developing early warning systems for Lassa fever in Nigeria.

Despite these advances, important limitations remain. First, many AI-based infectious disease prediction systems are designed around individual predictive algorithms and do not provide an integrated mechanism for coordinating multiple data sources and analytical processes. Second, the emphasis on predictive performance often exceeds attention given to model interpretability and transparency. This is problematic in public health because decision-makers may need to understand why an AI system has classified a particular geographical area as being at high risk. Explainable AI (XAI) has therefore become increasingly relevant to healthcare and clinical decision-support applications (Frasca *et al.*, 2024; Freyer *et al.*, 2024).

Recent literature has also highlighted the importance of evaluating XAI systems from a human-centered perspective. Donoso-Guzmán *et al.* (2025) reported that the practical value of explanations in healthcare remains insufficiently validated in real-world settings, emphasizing the importance of evaluating explanations in terms of their usefulness, trustworthiness, and usability for intended users. Thus, simply attaching an explanation method such as SHAP or LIME to a predictive model may not be sufficient to establish a genuinely trustworthy AI-based decision-support system.

The emergence of agentic AI presents a potentially important direction for addressing these limitations. Unlike conventional predictive systems that produce an output from a predefined pipeline, agentic AI systems can be designed around multiple specialized agents that perform distinct tasks, communicate with one another, retrieve relevant evidence, reason over information, and support complex workflows. When combined with explainability and human oversight, such systems may provide a more transparent and actionable approach to public health intelligence.

This study therefore proposes an Explainable Agentic Artificial Intelligence Framework for Early Prediction and Decision Support in Infectious Disease Outbreaks, using Lassa fever in Ebonyi State, Nigeria, as a case study. The proposed framework seeks to integrate epidemiological, environmental, climatic, demographic, and geospatial information through a coordinated multi-agent architecture capable of producing early risk predictions, explaining the basis of those predictions, and supporting evidence-informed public health decisions.

II. Review of Related Literature

2.1 Artificial Intelligence for Infectious Disease Prediction

AI has increasingly been applied to infectious disease prediction, surveillance, diagnosis, and prevention. Machine learning and deep learning methods are particularly useful because they can identify nonlinear relationships and complex interactions among multiple variables. A recent comprehensive review categorized AI-based infectious disease prediction into approaches based on public health data, patient medical data, and combinations of public and patient data (Melchane *et al.*, 2024).

However, the effectiveness of AI-based infectious disease prediction depends substantially on the quality and characteristics of available data. Infectious disease outbreaks are influenced by temporal, geographical, environmental, social, and demographic factors, which creates a need for models capable of integrating heterogeneous data sources. This requirement is particularly relevant to Lassa fever, where ecological and environmental conditions may influence disease occurrence and transmission.

2.2 Artificial Intelligence and Lassa Fever Prediction

Recent research has begun to investigate the application of machine learning specifically to Lassa fever prediction. Osaseri *et al.* (2025) proposed an AI-driven approach that combined evolutionary feature selection with Random Forest to improve Lassa fever prediction and surveillance. Such work demonstrates the potential of machine learning to identify relevant variables and improve outbreak prediction.

However, conventional machine learning models generally provide predictions without an autonomous mechanism for coordinating multiple analytical tasks. A model may identify a high-risk location, for example, but it may not independently retrieve supporting epidemiological evidence, analyze environmental conditions, explain the prediction to a decision-maker, or formulate a structured recommendation.

Recent research has also expanded the application of AI to Lassa fever risk mapping. A 2026 study using Self-Organizing Maps integrated environmental and human settlement information to classify Lassa fever risk, demonstrating the importance of spatial and environmental factors in disease risk assessment. These findings reinforce the need for approaches capable of integrating spatial, environmental, and epidemiological intelligence into early warning systems (Acta Tropica, 2026).

2.3 Climate and Environmental Factors in Lassa Fever Prediction

The relationship between environmental conditions and Lassa fever incidence provides an important foundation for predictive modeling. Fuwape *et al.* (2026) investigated six high-risk Nigerian states, including Ebonyi State, using weekly Lassa fever data from 2018 to 2022 and ERA5 environmental and climatic variables. The study identified significant relationships between Lassa fever incidence and variables including precipitation, vegetation, dew point temperature, and soil temperature.

These findings suggest that an effective Lassa fever prediction framework should not rely solely on historical case data. Instead, integrating epidemiological data with environmental and climatic information may improve the capacity of AI systems to detect emerging risk patterns. This supports the development of a multimodal AI architecture that can analyze multiple information streams simultaneously.

2.4 Explainable Artificial Intelligence in Healthcare

The increasing use of AI in healthcare has created concerns regarding transparency and accountability. Complex machine learning and deep learning models may achieve high predictive performance while offering limited insight into the reasoning underlying their outputs. This "black-box" problem can reduce user trust and limit adoption, particularly in high-stakes healthcare and public health contexts. Frasca *et al.* (2024) identified explainability and interpretability as important areas of research in medical AI, reflecting the growing need for AI systems whose predictions can be understood by healthcare professionals. Similarly, Freyer *et al.* (2024) emphasized the ethical importance of explainability in AI-based decision-support systems, particularly because opaque AI systems may make it difficult for users to evaluate the reliability and appropriateness of automated recommendations.

In the context of infectious disease prediction, explainability could allow public health officials to understand the factors contributing to a predicted increase in outbreak risk. For example, an AI system might identify a combination of recent case incidence, climatic conditions, environmental changes, population density, and geographical proximity to previously affected areas as important contributors to an elevated risk classification. However, it is important to distinguish model explanation from causal explanation. A feature attribution method may identify variables that contributed to a model's prediction without establishing that those variables causally produced the disease outbreak. Recent scholarship has emphasized this distinction and cautioned against interpreting predictive explanations as causal mechanisms (Carriero *et al.*, 2025). Therefore, explanations generated by the proposed framework should be presented as model-based evidence and risk factors, rather than definitive causal claims.

2.5 Agentic AI and Public Health Decision Support

The development of agentic artificial intelligence (AI) introduces the possibility of AI systems capable of coordinating multiple specialized computational agents to perform complex, sequential tasks. In public health and epidemic intelligence, such systems can integrate heterogeneous information sources, conduct risk assessment, generate early warnings and support intervention decisions. Recent work has proposed AI-agent architectures in which specialized agents collaboratively perform surveillance, risk evaluation, forecasting, early warning and decision-support functions (Yang *et al.*, 2026). Similarly, emerging work on agentic AI in healthcare highlights the potential of AI systems with greater autonomy to support complex clinical and health-system workflows while emphasizing the need for appropriate ethical and governance safeguards (Hinostroza Fuentes *et al.*, 2025).

In the context of infectious disease surveillance, specialized agents can therefore be assigned complementary responsibilities, including monitoring epidemiological data, analyzing environmental conditions, performing spatial analysis, generating outbreak-risk predictions, retrieving supporting evidence and explaining model outputs. Such an architecture is particularly relevant to epidemic intelligence because outbreak detection requires the integration and interpretation of multiple heterogeneous data streams rather than reliance on a single predictive variable or algorithm (Yang *et al.*, 2026).

The proposed approach differs from conventional AI prediction systems because it conceptualizes outbreak intelligence as a multi-stage reasoning and decision-support process rather than as a single prediction task. For example, a Surveillance Agent could identify an unusual increase in reported cases, while an Environmental Agent could analyse recent climatic and environmental conditions. A Geospatial Agent could identify geographically proximate or interconnected high-risk locations, and a Prediction Agent could estimate

the probability of future outbreak escalation. An Evidence Retrieval Agent could retrieve relevant epidemiological evidence, while an Explainability Agent could provide an interpretable account of the factors contributing to the predicted risk. Such a surveillance-to-decision architecture is consistent with recent proposals for AI-agent-supported epidemic intelligence, which extend traditional surveillance and early-warning systems toward integrated decision support and intervention optimization (Yang *et al.*, 2026).

The integration of these capabilities could transform an AI system from a predominantly predictive tool into an intelligent decision-support framework capable of connecting surveillance signals, risk estimation, explanation and recommended response actions. However, increased AI autonomy also introduces important challenges relating to reliability, data quality, robustness, accountability, transparency and human control. The World Health Organization (WHO) emphasizes that AI systems used in health should protect human autonomy, promote safety, ensure transparency and explainability, and establish clear responsibility and accountability (World Health Organization [WHO], 2021). More recent WHO guidance similarly stresses the importance of appropriate governance, evaluation and oversight for advanced AI systems used in health and public-health applications (WHO, 2025).

Accordingly, the proposed framework should adopt a human-in-the-loop architecture in which AI-generated predictions, explanations and recommendations are treated as decision-support information rather than autonomous public-health decisions. Qualified epidemiologists and public-health professionals should retain authority to validate AI-generated alerts, assess contextual evidence and determine appropriate interventions. This approach is consistent with emerging epidemic-intelligence frameworks that explicitly recommend mandatory human oversight to enhance trust, accountability and safe deployment of AI agents (Yang *et al.*, 2026). It also aligns with WHO's principle that humans should remain in control of health-related decisions and that AI technologies should be subject to appropriate evaluation, transparency and accountability (WHO, 2021).

2.6 Research Gap

The reviewed literature reveals several gaps.

First, existing Lassa fever prediction studies have demonstrated the potential of machine learning, but most focus on individual predictive models rather than coordinated multi-agent AI architectures. Second, although environmental and climatic variables have been shown to contribute to Lassa fever risk prediction, there is limited integration of these heterogeneous data streams into a unified intelligent decision-support framework. Third, existing predictive systems often emphasize model performance while giving comparatively less attention to explaining predictions to public health stakeholders.

Fourth, XAI research has advanced substantially in healthcare, but there remains a need to investigate how explanations can be incorporated into complex, multi-stage, agentic AI workflows. The challenge is no longer only to explain why a predictive model generated an output but also to explain how multiple AI agents reached and justified a decision-support recommendation.

Fifth, recent research on user-centered XAI evaluation suggests that explanations should be evaluated in terms of their practical usefulness, trustworthiness, and usability rather than solely through technical measures (Donoso-Guzmán *et al.*, 2025). This indicates a need for human-centered evaluation involving domain experts.

The central research gap can therefore be expressed as follows:

There is limited research on an integrated, explainable agentic AI framework capable of combining heterogeneous epidemiological, environmental, climatic, demographic, and geospatial data to generate early, localized Lassa fever outbreak predictions and provide transparent, evidence-grounded, human-centered decision support in Ebonyi State, Nigeria.

III. Materials and Methods

3.1 Study Design

This study will employ a retrospective longitudinal spatio-temporal design combined with Design Science Research (DSR) and controlled machine-learning experimentation.

The retrospective component will use historical Lassa fever surveillance records from Ebonyi State. These records will be integrated with weekly climatic, environmental, demographic, and geospatial variables to construct a multimodal dataset.

The DSR component will guide the design and development of the proposed Explainable Agentic Artificial Intelligence (EAAI) framework, while the experimental component will compare the framework against established statistical, machine-learning, and deep-learning models.

The fundamental unit of analysis will be:

Local Government Area (LGA) × Epidemiological Week

Thus, each observation will represent the disease and environmental conditions of one LGA during one epidemiological week.

3.2 Study Area

The study will be conducted in Ebonyi State, southeastern Nigeria. The state will be analyzed at the LGA level because Lassa fever risk is spatially heterogeneous and aggregated state-level data may conceal important local variation.

The primary spatial resolution will therefore be the LGA, while the temporal resolution will be the epidemiological week.

3.3 Study Period

The study will primarily target 2018–2025, subject to availability and completeness of surveillance data. The 2018–2022 period is particularly valuable because a peer-reviewed study has already documented Lassa fever surveillance in Ebonyi State during this period, reporting 1,578 cases, including 300 laboratory-confirmed cases. However, those published aggregate statistics will be used only as a validation/reference baseline. The actual predictive model will use the underlying weekly observations obtained from authorized surveillance sources.

3.3.1 Dataset Size

The study has **13 LGAs**, approximately 52 epidemiological weeks per year, and 8 years of data:

$$13 \times 52 \times 8 = 5,408$$

the target dataset is therefore approximately: 5,408 LGA-week observations with potentially 20–40+ predictor variables, depending on the availability of environmental and geospatial variables.

3.4 Dataset Development

The study will create a multimodal Lassa Fever Spatio-Temporal Dataset (MLF-STD).

The proposed dataset will combine five major data domains:

1. Epidemiological surveillance;
2. Climate;
3. Environmental/satellite;
4. Demographic;
5. Geospatial.

The resulting dataset will have the structure:

$$D = \{X_{i,t}, Y_{i,t}\}$$

where:

- i = LGA;
- t = epidemiological week;
- $X_{i,t}$ = predictor variables;
- $Y_{i,t}$ = Lassa fever outcome.

3.5 Data Preprocessing

The following sequence will be followed.

Step 1: Data Harmonization

All datasets will be converted to a common: **LGA × epidemiological week** format.

Step 2: Duplicate Detection

Duplicate epidemiological records will be identified and resolved.

Step 3: Missing-Data Analysis

The percentage and pattern of missing observations will be assessed. Missingness will be reported rather than automatically hidden.

Step 4: Imputation

For climate variables, appropriate temporal interpolation or model-based imputation may be used where justified. Epidemiological case counts will not be artificially imputed without epidemiological justification.

Step 5: Outlier Detection

Outliers will be examined using:

- IQR;
- robust z-scores;
- temporal consistency checks.

Extreme epidemiological observations will not automatically be deleted because they may represent genuine outbreaks.

Step 6: Feature Engineering

The following will be generated:

- Lag variables;
- 4-week moving average;
- 8-week moving average;

- cumulative cases;
- case growth rate;
- rainfall anomaly;
- temperature anomaly;
- spatial lag.

IV. Results

The Lassa fever surveillance contained 1,910 records and 311 variables. After restricting the records to cases for which Ebonyi was the responsible region, 1,884 records remained. Among these, 359 records were classified as confirmed Lassa fever cases.

4.1 Temporal distribution of confirmed Lassa fever cases

Table 1: The confirmed cases varied substantially across the six-year study period.

Year	Confirmed cases
2018	86
2019	60
2020	88
2021	19
2022	51
2023	55
Total	359

Table 1 above showed that highest number of confirmed cases occurred in 2020 (88 cases), closely followed by 2018 (86 cases). The lowest number occurred in 2021 (19 cases). The reduction in 2021 followed the relatively high case counts observed in 2018–2020, but confirmed cases subsequently increased to 51 in 2022 and 55 in 2023. This indicates that Lassa fever transmission did not follow a simple monotonic trend but exhibited substantial interannual variability.

This temporal variability supports the rationale for developing a predictive system rather than relying solely on historical annual averages.

The graph presents the number of laboratory-confirmed Lassa fever cases in Ebonyi State for each year from 2018 to 2023.




Figure 1: Annual Trend of Confirmed Lassa Fever Cases, 2018–2023

Figure 1 shows substantial interannual variability in confirmed Lassa fever occurrence in Ebonyi State between 2018 and 2023. Confirmed cases increased from 60 in 2019 to 88 in 2020 before declining sharply to 19 cases in 2021. Cases subsequently increased to 51 and 55 in 2022 and 2023, respectively. This non-linear temporal pattern indicates that Lassa fever occurrence is dynamic and supports the development of predictive approaches capable of modelling temporal dependencies and potential resurgence.

4.2 Seasonal pattern

The monthly analysis showed a pronounced concentration of confirmed cases during the first quarter of the year.

Table 2: Monthly Analysis

Month	Confirmed cases
January	—
February	111
March	101
Other months	Remaining cases

February and March together accounted for **212 of the 359 confirmed cases**, or approximately **59.1%** of all confirmed cases in the downloadable dataset.

This is a strong seasonal signal. Approximately six out of every ten confirmed cases occurred during February and March. This finding is particularly important for the proposed AI framework because it suggests that temporal/seasonal variables should be included as predictors. However, seasonality alone cannot explain the observed spatial and interannual variation. This provides a strong justification for subsequently integrating rainfall, temperature, vegetation and other environmental variables.

This graph displays the number of confirmed Lassa fever cases according to month of report.




Figure 2: Monthly Seasonal Distribution

Figure 2 demonstrates a marked seasonal concentration of confirmed Lassa fever cases. February and March recorded 111 and 101 confirmed cases, respectively, together accounting for 59.1% of the 359 confirmed cases identified in the downloadable surveillance dataset. The concentration of cases during these months indicates substantial temporal seasonality in Lassa fever occurrence. Although the surveillance data establish the seasonal pattern, they do not independently establish the environmental mechanisms underlying this distribution. Consequently, climate and environmental variables should be incorporated in subsequent modelling to determine whether they provide additional predictive information beyond historical disease incidence and seasonality.

4.3 Geographic Distribution

The confirmed cases were highly concentrated geographically.

Table 3: Geographical Concentration

LGA	Confirmed cases	Percentage
Abakaliki	181	50.4%
Izzi	41	11.4%
Ezza North	33	9.2%
Ikwo	25	7.0%
Ebonyi	24	6.7%

Other LGAs	55	15.3%
Total	359	100%

Abakaliki LGA alone accounted for approximately half of all confirmed cases (50.4%). Izzi, Ezza North, Ikwo and Ebonyi LGAs were the next major contributors. The substantial geographic concentration suggests that Lassa fever risk is spatially heterogeneous within Ebonyi State. Consequently, an AI model operating only at state level could conceal important local variation.

For the proposed study, this supports the construction of an:

LGA × epidemiological-week prediction framework

rather than modelling Ebonyi as a single homogeneous geographic unit.

The graphical representation of the distribution of confirmed Lassa fever cases across the Local Government Areas in the surveillance data.

The graph below shows the distribution of confirmed Lassa fever cases across the Local Government Areas represented in the surveillance data.




Figure 3: Confirmed Cases by LGA

Figure 3 revealed considerable spatial heterogeneity in confirmed Lassa fever cases across Ebonyi State. Abakaliki LGA accounted for 181 (50.4%) of the 359 confirmed cases, followed by Izzi (11.4%), Ezza North (9.2%), Ikwo (7.0%) and Ebonyi (6.7%). The concentration of cases in particular LGAs suggests that state-level aggregation may obscure important local variation in disease occurrence. Nevertheless, reported case counts may reflect differences in population size, healthcare access, surveillance intensity and diagnostic capacity in addition to differences in transmission. These findings support the development of an LGA-level spatio-temporal prediction framework.

4.4 AI prediction problem

For the preliminary AI experiment, the individual surveillance records were aggregated into LGA-week observations.

The prediction target was defined as:

Whether at least one confirmed Lassa fever case would occur in an LGA during the following week.

Historical surveillance information was transformed into predictors including:

- cases in the previous week;
- cases two weeks previously;
- cases four weeks previously;

- cases eight weeks previously;
- cases twelve weeks previously;
- four-week rolling mean;
- eight-week rolling mean;
- twelve-week rolling mean;
- seasonal sine component; and
- seasonal cosine component.

This formulation avoids using future information to predict the future and is therefore more appropriate than randomly splitting the surveillance observations.

4.5 Preliminary AI model results

A chronological evaluation was used, with earlier years used for model development and **2023 reserved as the temporal test period**.

Table 4: Preliminary Results

Model	ROC-AUC	PR-AUC	Sensitivity	Precision	F1
Logistic Regression	0.836	0.371	85.7%	—	0.193
XGBoost	0.823	0.350	82.1%	—	0.210
Random Forest	0.792	0.298	64.3%	—	0.199

The results demonstrated in table 4 indicate that all three models demonstrated useful discriminatory ability, with ROC-AUC values above 0.79.

Logistic Regression achieved the highest ROC-AUC (0.836), indicating the strongest overall ranking/discrimination among the preliminary models. XGBoost produced a slightly lower ROC-AUC (**0.823**) but achieved the highest preliminary F1-score (**0.210**). Random Forest had the lowest ROC-AUC (**0.792**) and sensitivity (**64.3%**) among the three.

However, the relatively low F1 scores indicate that accurately identifying the relatively infrequent positive outbreak weeks remains difficult. Therefore, ROC-AUC alone should not be used to claim that the models are highly effective.

For an outbreak-warning system, PR-AUC, sensitivity, calibration and lead time are particularly important because missing an outbreak may have greater operational consequences than generating a false alert.

The graph illustrates the three comparable machine-learning models used in the analysis:

1. **Logistic Regression**
2. **XGBoost**
3. **Random Forest**




Figure 4: Comparative AI Model Performance

V. Conclusion

This study demonstrated substantial temporal and spatial heterogeneity in confirmed Lassa fever occurrence in Ebonyi State between 2018 and 2023. Analysis of the available surveillance records identified 359 confirmed cases within the downloadable dataset, with considerable variation in annual disease occurrence. Confirmed cases increased from 60 in 2019 to 88 in 2020, declined sharply to 19 in 2021, and subsequently increased to 51 and 55 cases in 2022 and 2023, respectively. This fluctuating pattern indicates that Lassa fever occurrence is dynamic and cannot be adequately represented by static annual estimates.

A marked seasonal pattern was also observed, with February and March accounting for approximately 59.1% of the confirmed cases. This concentration provides empirical evidence of temporal seasonality and supports the incorporation of temporal and environmental predictors into predictive modelling. Spatially, the disease burden was highly heterogeneous, with Abakaliki LGA accounting for approximately half of the confirmed cases, followed by Izzi, Ezza North, Ikwo and Ebonyi LGAs. These findings demonstrate the importance of modelling Lassa fever risk at a geographically disaggregated level rather than relying exclusively on state-level aggregates.

The preliminary machine-learning experiment further demonstrated that historical surveillance information contains predictive signals for subsequent Lassa fever occurrence. Logistic Regression achieved the highest ROC-AUC (0.836), while XGBoost produced the highest F1-score (0.210) among the preliminary models evaluated. The differences in model performance indicate that model selection should not be based on a single evaluation metric, particularly in an epidemiological setting characterised by an imbalanced outcome and potentially costly false-negative predictions.

Overall, the findings provide an empirical basis for developing an explainable, spatio-temporal and decision-oriented artificial intelligence framework for Lassa fever early warning in Ebonyi State. However, the present results represent a surveillance-only baseline and do not, by themselves, establish the superiority of an agentic AI framework or demonstrate the incremental value of climate and environmental predictors. The next phase of the research should therefore integrate surveillance, climatic, environmental, demographic and spatial information and evaluate whether multimodal modelling improves predictive accuracy, calibration, interpretability and outbreak lead time. Such an approach has the potential to transform routine surveillance data into an actionable early-warning and decision-support mechanism for public-health authorities.

The study establishes a reproducible spatio-temporal surveillance baseline and provides the foundation for an explainable multimodal AI framework capable of converting heterogeneous disease and environmental data into interpretable, prospective outbreak-risk estimates and actionable decision support.

VI. Recommendations

1. Develop an LGA-level early-warning system: Public-health authorities should adopt geographically disaggregated Lassa fever surveillance and prediction to identify LGAs with increasing outbreak risk and support timely allocation of health resources.
2. Integrate multimodal data sources: Future predictive models should combine routine Lassa fever surveillance with climatic, environmental, demographic and geospatial variables to improve the prediction of emerging outbreaks.
3. Prioritise robust and explainable AI evaluation: AI models should be evaluated using sensitivity, specificity, PR-AUC, F1-score, calibration, false-alert rate and outbreak lead time, while explainability techniques such as SHAP should be used to provide interpretable evidence for AI-generated predictions.
4. Strengthen validation and surveillance data quality: The final model should undergo rigorous temporal and, where possible, external validation using high-quality surveillance data. Inconsistencies, missing observations, duplicate records and geographic or temporal coding errors should be addressed before deployment.
5. Implement human-in-the-loop prospective decision support: Following retrospective validation, the AI framework should be prospectively piloted within the public-health surveillance system, where AI-generated risk alerts and explanations support—but do not replace—epidemiologists and public-health authorities in outbreak investigation and response.

References

- [1]. Acta Tropica. (2026). Improving Lassa fever risk mapping using self-organizing maps and spatial determinants. *Acta Tropica*, 277, 108040. <https://doi.org/10.1016/j.actatropica.2026.108040>
- [2]. Carriero, A., de Hond, A., Cappers, B., Paulovich, F., Abeln, S., Moons, K. G. M., & van Smeden, M. (2025). Explainable AI in healthcare: To explain, to predict, or to describe? *[Preprint]*.
- [3]. Donoso-Guzmán, I., Kacafirková, K. S., Szymanski, M., Jacobs, A., Parra, D., & Verbert, K. (2025). A systematic review of user-centred evaluation of explainable AI in healthcare. *[Preprint]*.

- [4]. Frasca, M., La Torre, D., Pravettoni, G., & Cutica, I. (2024). Explainable and interpretable artificial intelligence in medicine: A systematic bibliometric review. *Discover Artificial Intelligence*, 4, 15. <https://doi.org/10.1007/s44163-024-00114-7>
- [5]. Freyer, N., Groß, D., & Lipprandt, M. (2024). The ethical requirement of explainability for AI-DSS in healthcare: A systematic review of reasons. *BMC Medical Ethics*, 25, 104. <https://doi.org/10.1186/s12910-024-01103-2>
- [6]. Fuwape, I., Ogunjo, S., Erinfolami, F., Osunmakinwa, O., Ajanaku, O., Akerele, A., Akinlami, A., & Akinsusi, J. (2026). Quantifying the climatic influence on Lassa fever incidences in six high-risk Nigerian states using multilinear regression and generalized additive models. *Discover Public Health*, 23, 276. <https://doi.org/10.1186/s12982-026-01550-9>
- [7]. Hinostroza Fuentes, V. G., Karim, H. A., Tan, M. J. T., & AlDahoul, N. (2025). AI with agency: A vision for adaptive, efficient, and ethical healthcare. *Frontiers in Digital Health*, 7, 1600216. doi:10.3389/fdgth.2025.1600216
- [8]. Melchane, S., Elmir, Y., Kacimi, F., & Boubchir, L. (2024). Artificial intelligence for infectious disease prediction and prevention: A comprehensive review. *[Preprint]*.
- [9]. Osaseri, R. O., Usiobaifo, A. R., & Ighodaro, U. E. (2025). AI-driven prediction of Lassa fever using evolutionary and Random Forest: A machine learning approach for enhanced surveillance in West Africa. *International Journal of Data Science*, 6(1), 40–53. <https://doi.org/10.18517/ijods.6.1.40-53>
- [10]. World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. Geneva, Switzerland.
- [11]. World Health Organization. (2025). *Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models*. Geneva, Switzerland.
- [12]. Yang, L., Shan, L., Cao, X., Cui, J., Tong, M., Niu, Y., & Zhang, T. (2026). AI agents and epidemic intelligence on respiratory infectious diseases: Toward a conceptual framework integrating decision support. *Journal of Medical Internet Research*, 28, e86936. doi:10.2196/86936