

Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa

Kiro M. Lauraine ^{1, ✉}, Moïse K. Kasambya ^{2, 3}, Mundeke K. Ghislain ¹, Charité K. Valambaya ⁴, Amanio P. Isaac ¹, Hangi M. Patrice ⁴, Alain K. Virekero ⁴, Elie M. Mwandu ⁴, Belden K. Mbafumoja ⁴, Muyisa M. Katemengi ⁴

1. Information Systems, Université de l'Assomption au Congo, Butembo, COD

2. Computer Science, University of the Philippines Diliman, Quezon, PHL

3. Business and Information Technology, Université de l'Assomption au Congo, Butembo, COD

4. Computer Networks, Université de l'Assomption au Congo, Butembo, COD

Received: June 20, 2026 | Review began: July 23, 2026 | Review ended: July 27, 2026 | Published: July 31, 2026

© Copyright 2026

This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

Malaria remains one of the most significant infectious disease burdens in Sub-Saharan Africa, placing sustained pressure on public health systems and malaria-control programmes. Accurate and interpretable identification of high-risk districts is therefore essential for early warning, resource allocation, and targeted intervention planning. This study proposes an explainable multimodal artificial intelligence framework that integrates epidemiological, climatic, geospatial, and demographic information through a modality-specific encoder-fusion architecture. The framework is further extended with a spatial Graph Neural Network component designed to capture inter-district dependencies and potential spatial spillover effects. Using a synthetically generated district-level dataset ($n = 1,200$), calibrated to reflect the statistical properties of variables commonly used in Malaria Atlas Project (MAP), NASA POWER, CHIRPS, WorldPop, and Sentinel-2 products across five malaria-endemic countries (the Democratic Republic of the Congo, Uganda, Nigeria, Tanzania, and Kenya), the proposed framework achieved a validation R^2 of 0.7285 for *Plasmodium falciparum* parasite rate regression. For binary high-risk classification, the framework attained an area under the curve of 0.9740 and an F1-score of 0.9630 on the held-out test dataset. SHapley Additive exPlanations were used to examine the contribution of individual predictors at global and local levels. Incidence rate (67.2%), rainfall (10.2%), and temperature (6.2%) emerged as the leading drivers of predicted risk, while intervention coverage showed a protective contribution. The results should be read as a synthetic proof-of-concept rather than as an operational malaria map. Even so, they show how multimodal learning, graph-based spatial modeling, and explainable AI can be combined into a transparent surveillance pipeline for future validation with real-world malaria data in resource-constrained settings.

Categories: AI applications, Predictive Analytics, Health Informatics

Keywords: malaria risk prediction, graph neural networks, explainable ai, multimodal learning, plasmodium falciparum, digital epidemiology

Introduction

Malaria, particularly that caused by *Plasmodium falciparum*, remains one of the most significant public health challenges in Sub-Saharan Africa. According to the World Health Organization (WHO), the African Region accounted for approximately 94% of the 249 million malaria cases reported globally in 2022, resulting in an estimated 608,000 deaths, with children under five years of age bearing the greatest burden of mortality [1]. Accurate characterization of malaria risk remains particularly important in high-burden countries such as the Democratic Republic of the Congo (DRC), Nigeria, Uganda, Tanzania, and Kenya, where transmission dynamics vary substantially across districts due to complex interactions among climatic, ecological,

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. Cureus J Comput Sci 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>

demographic, and socioeconomic factors [2]. Consequently, fine-grained spatial risk mapping has become a critical tool for guiding targeted interventions, including insecticide-treated bed net distribution, indoor residual spraying, seasonal malaria chemoprevention, and surveillance planning.

Over the past decade, machine learning (ML) techniques have been increasingly employed to enhance malaria surveillance and forecasting. Algorithms such as random forests, gradient boosting machines, and deep neural networks have demonstrated promising performance in predicting parasite prevalence, disease incidence, and mosquito vector suitability from environmental and epidemiological variables [3,4]. Despite these advances, several methodological limitations remain. First, many existing studies rely on a single data modality, typically epidemiological records or remotely sensed environmental variables, thereby overlooking the complementary information available from multiple heterogeneous sources. Second, most predictive models treat administrative units as independent observations, ignoring spatial autocorrelation and diffusion processes that characterize infectious disease transmission across neighboring regions. Third, the limited interpretability of many ML models restricts their practical adoption by public health stakeholders, who often require transparent and explainable evidence to support intervention and resource allocation decisions [5].

Graph neural networks (GNNs) have recently emerged as a powerful paradigm for learning from structured and relational data [6,7]. By representing districts as nodes connected through geographic or epidemiologically relevant relationships, GNNs enable the propagation of information across spatial networks and capture neighborhood dependencies that conventional ML models may overlook. Such capabilities are particularly relevant for infectious disease modeling, where transmission patterns often extend beyond administrative boundaries. Previous studies have demonstrated the effectiveness of GNNs in forecasting epidemics such as influenza [8] and COVID-19 [9]. However, their application to malaria risk prediction in Sub-Saharan Africa remains relatively underexplored.

At the same time, explainability has become a central requirement for trustworthy artificial intelligence (AI) in healthcare and public health settings [10]. Among existing explainable AI techniques, SHAP (SHapley Additive exPlanations) offers a theoretically grounded framework for quantifying feature contributions at both global and local levels [11,12]. By identifying the factors that drive predictions, SHAP facilitates transparency, supports decision-making, and enhances stakeholder confidence in AI-assisted surveillance systems.

To address these challenges, this study proposes and evaluates an integrated proof-of-concept framework that combines multimodal learning, GNNs, and explainable AI for district-level malaria risk prediction in Sub-Saharan Africa. Specifically, the study makes four main contributions:

1. A multimodal fusion architecture integrating epidemiological, climatic, geospatial, and demographic information through dedicated encoder networks and a shared prediction module.
2. A spatial graph convolutional network (GCN) that models geographic dependencies using a district-level k-nearest-neighbor graph constructed from spatial coordinates.
3. A three-level SHAP explainability framework providing global feature importance analysis, summary visualizations, and district-specific explanations.
4. A reproducible benchmarking protocol comparing the proposed approach with four baseline models across both regression (*Plasmodium falciparum* parasite rate (PfPR) estimation) and binary classification (high- versus low-risk) tasks.

Related work

Malaria Risk Modeling With ML

Geospatial ML has produced a succession of increasingly accurate malaria risk maps. Bhatt et al. [3] used Gaussian process spatial modeling to estimate the prevalence of *Plasmodium falciparum* across Africa from 2000 to 2015, demonstrating a substantial reduction in parasite prevalence attributable to malaria control interventions. Weiss et al. [2] extended this work by mapping global *Plasmodium falciparum* mortality at a 5 × 5 km resolution using an ensemble of geostatistical models. More recently, Dalrymple et al. [4] applied gradient-boosted regression trees to predict sub-national malaria incidence in 43 high-burden countries, emphasizing the importance of integrating biological, environmental, and socioeconomic predictors. Although these studies demonstrate the predictive power of ML approaches, their limited focus on model interpretability constrains their usefulness for operational public health decision-making.

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. Cureus J Comput Sci 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>

GNNs in Epidemiology

GNNs have rapidly emerged as powerful tools for modeling infectious disease propagation over spatial and relational networks. Deng et al. [8] proposed Cola-GNN, a cross-location-adaptive GNN that achieved state-of-the-art performance in influenza forecasting across the United States. Kapoor et al. [9] further demonstrated that spatio-temporal GNNs outperform conventional ML approaches for modeling COVID-19 transmission at the county level. Despite these advances, the application of GNNs to malaria risk prediction in Sub-Saharan Africa remains relatively underexplored, particularly in settings characterized by heterogeneous transmission dynamics and limited surveillance coverage.

Multimodal Learning in Digital Health

Multimodal learning, defined as the joint processing of heterogeneous data sources, has demonstrated significant advantages over unimodal approaches across numerous healthcare applications [13]. In clinical environments, the integration of imaging, genomic, and electronic health record data has improved disease diagnosis, prognosis estimation, and treatment planning [14]. Malaria risk prediction naturally constitutes a multimodal problem because disease transmission is influenced simultaneously by epidemiological, climatic, environmental, geographic, and demographic factors. Nevertheless, most published malaria prediction models continue to rely on a single category of data. To the best of our knowledge, few studies have systematically implemented modality-specific encoder architectures for district-level malaria risk prediction in Sub-Saharan Africa.

Context-Aware and Deployable AI in Low-Resource African Settings

Beyond disease-specific modeling, AI-based surveillance systems in Sub-Saharan Africa require contextual validity, transparency, and computational feasibility. Recent malaria studies have shown the value of locally grounded prediction models, routine health-facility data, and fine-scale environmental variables for malaria risk mapping and micro-stratification [15-17]. In Kenya, Malawi, and the DRC, Ghilardi et al. [18] further showed that malaria risk maps are useful only when decision-makers trust, understand, and perceive them as relevant for intervention planning. This supports the need for explainable and context-sensitive AI in African public health systems [19]. In the Congolese context, Mpia et al. [20], CoBERT [21], and Mpia et al. [22] also highlight the importance of contextual variables, robust validation, and lightweight deployment in unstable or resource-constrained environments. Together, these studies strengthen the rationale for a context-aware, interpretable, and deployable AI framework for malaria surveillance in Sub-Saharan Africa.

Explainable AI in Public Health

The integration of explainable AI methods into predictive health models has received growing attention as a pathway toward trustworthy and accountable AI deployment [10,23]. SHAP, derived from cooperative game theory, provides theoretically grounded local and global feature attributions [11]. For tree-based models, Tree SHAP can be computed efficiently and is therefore practical for large epidemiological datasets [12]. In the proposed framework, SHAP is used to translate numerical predictions into interpretable risk narratives that can be reviewed by public health analysts and district-level decision-makers.

Materials And Methods

Study area

The study area encompasses five malaria-endemic countries in Sub-Saharan Africa: the DRC, Uganda, Nigeria, Tanzania, and Kenya. These countries collectively account for a substantial proportion of the global burden of *Plasmodium falciparum* malaria and exhibit diverse transmission ecologies ranging from the hyperendemic rainforest environments of the Congo Basin to the seasonal and highland transmission zones of East Africa. Such diversity provides an appropriate setting for evaluating spatially aware ML models across different epidemiological contexts. The geographic bounding box of the study extends from 15°S to 15°N latitude and from 10°E to 42°E longitude.

Dataset

Because verified district-level surveillance records were not yet harmonized across the five study countries, we generated a synthetic dataset of 1,200 district observations for proof-of-concept evaluation. The dataset was calibrated to reproduce the statistical characteristics, inter-variable relationships, and geographical variability commonly observed in Malaria Atlas Project (MAP)-type data. The synthetic data generation process used a causally inspired malaria-risk formulation incorporating seven environmental and demographic covariates with empirically motivated coefficients: rainfall (0.31), temperature (0.24),

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. Cureus J Comput Sci 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>

normalized difference vegetation index (NDVI) (0.18), logarithm of population density (0.14), vector suitability (0.13), intervention coverage (-0.10), and elevation (-0.05). A Gaussian noise component ($\sigma = 0.05$) was added to simulate realistic unexplained variability.

The resulting PfPR distribution, characterized by a mean of 0.458 and a standard deviation of 0.140, aligns with published continental estimates reported between 2015 and 2022. Although synthetic, the dataset was designed as a proof-of-concept resource for evaluating the proposed framework under controlled conditions. Full replacement with real-world MAP observations remains feasible through the malariaAtlas Python client or the malariaAtlas R package. Table 1 presents a summary of the 10 feature variables included in the study.

Variable	Modality	Source	Mean $\pm$ SD	Range
Latitude (°)	Geographic	MAP	0.00 $\pm$ 8.73	-14.98 to 14.98
Longitude (°)	Geographic	MAP	25.86 $\pm$ 9.32	10.02 to 41.99
PfPR	Epidemiological	MAP	0.458 $\pm$ 0.140	0.074 to 0.882
Incidence (per 1,000)	Epidemiological	MAP	229.97 $\pm$ 75.35	43.81 to 502.34
Vector suitability	Epidemiological	MAP	0.499 $\pm$ 0.292	0.00 to 0.999
Intervention coverage	Epidemiological	MAP	0.496 $\pm$ 0.287	0.001 to 0.999
Temperature (°C)	Climate	NASA POWER	26.58 $\pm$ 4.92	18.01 to 34.99
Rainfall (mm/yr)	Climate	CHIRPS	1367.6 $\pm$ 629.5	302.0 to 2497.9
Humidity (%)	Climate	NASA POWER	64.41 $\pm$ 14.16	40.02 to 89.99
Elevation (m)	Geographic	SRTM	1489.8 $\pm$ 857.2	1.25 to 2998.6
NDVI	Geographic	Sentinel-2	0.505 $\pm$ 0.231	0.10 to 0.90
Pop. density (km ⁻²)	Demographic	WorldPop	200.4 $\pm$ 200.7	0.55 to 1791.8

TABLE 1: Multimodal feature variables, data sources, and descriptive statistics.

MAP, Malaria Atlas Project; NDVI, Normalized Difference Vegetation Index; PfPR, *Plasmodium falciparum* Parasite Rate

Data preprocessing

Data preprocessing was conducted through five sequential stages to ensure data quality, consistency, and model robustness. First, missing value handling was considered; although the synthetic dataset contained no missing observations, mean imputation is recommended when applying the framework to real MAP datasets with incomplete records. Second, categorical country labels were transformed using integer label encoding to facilitate numerical processing. Third, all continuous variables were standardized using the StandardScaler procedure (zero mean and unit variance), fitted exclusively on the training partition to avoid data leakage and preserve the integrity of model evaluation. Fourth, the dataset was partitioned into training, validation, and test subsets using a stratified 70%/15%/15% split ($n = 840/180/180$), ensuring preservation of class

proportions across subsets. Finally, variables were grouped into four modality-specific blocks: epidemiological (incidence, vector suitability, intervention coverage; $d = 3$), climate (temperature, rainfall, humidity; $d = 3$), geospatial (elevation, NDVI, latitude, longitude; $d = 4$), and demographic (population density, country encoding; $d = 2$).

Multimodal fusion architecture

The proposed MultimodalFusionNet encodes each data modality through a dedicated ModalityEncoder sub-network. The epidemiological, climatic, and geospatial encoders each use two fully connected layers with a hidden size of 32 and an output size of 16, while the demographic encoder produces an eight-dimensional representation. Each encoder applies ReLU activation, batch normalization, and dropout regularization ($p = 0.2$). The four latent representations are then concatenated into a 56-dimensional joint vector and passed through a shared prediction head ($64 \rightarrow 32 \rightarrow 1$) for PfPR regression. Figure 1 summarizes the complete modeling workflow.

The architecture contains 8,169 trainable parameters, which makes it suitable for rapid experimentation and potentially feasible on CPU-only hardware. Model optimization was performed for 80 epochs using the Adam optimizer (learning rate = 3×10^{-3} , weight decay = 10^{-4}), a cosine-annealing learning-rate scheduler, and mean squared error (MSE) loss.




FIGURE 1: Schematic of the MultimodalFusionNet architecture. Four modality-specific encoders (epidemiological, climate, geospatial, and demographic) produce latent representations that are concatenated and passed through a shared prediction head for PfPR regression.

GNN, Graph Neural Network; MAP, Malaria Atlas Project; MSE, Mean Squared Error; NDVI, Normalized Difference Vegetation Index; PfPR, *Plasmodium falciparum* Parasite Rate; SHAP, SHapley Additive exPlanations; XAI, Explainable Artificial Intelligence

Spatial GNN extension

To explicitly capture spatial dependencies between neighboring districts, the multimodal framework is extended through a two-layer SpatialGNN based on GCN operations [6]. A district-level spatial graph $G = (V, E)$ is constructed in which each of the 1,200 districts is represented as a node, while edges connect geographically proximate districts using the k-nearest neighbors algorithm ($k = 6$) computed in haversine distance space. This procedure yields an undirected graph with an average node degree of approximately 12.

Node attributes consist of the complete set of 12 standardized covariates. The SpatialGNN architecture applies two GCNConv layers with hidden dimensions of 64 and 32 neurons, respectively, followed by ReLU activations and dropout regularization ($p = 0.3$). A linear output layer performs node-level PfPR regression. Training is conducted using node-level masks corresponding to the training, validation, and test partitions (70%/15%/15%) and optimized using the Adam optimizer (learning rate = 10^{-3} , weight decay = 5×10^{-4}).

Explainability: SHAP analysis

To enhance model interpretability, random forest is employed as a surrogate explainability model because of its established compatibility with SHAP and its competitive predictive performance. The SHAP TreeExplainer framework [12] decomposes each prediction into feature-specific Shapley values while satisfying the theoretical properties of local accuracy, missingness, and consistency.

Three complementary levels of explanation are generated. First, global feature importance is quantified using the mean absolute SHAP value computed across all observations in the test set. Second, a SHAP beeswarm summary plot is used to visualize both the magnitude and direction of feature effects on model predictions. Third, a local waterfall explanation is produced for the highest-risk district identified in the test set (predicted PfPR = 0.7118, district index 176), enabling detailed interpretation of the factors contributing to its risk score.

Baseline models and evaluation metrics

The proposed framework is benchmarked against four widely used regression models: linear regression, ridge regression ($\alpha = 1.0$), gradient boosting (100 estimators), and random forest (100 estimators, random seed = 42).

Regression performance is evaluated using three complementary metrics: root mean squared error (RMSE), mean absolute error (MAE), and the coefficient of determination (R^2). In addition, a binary malaria risk classification task is considered, where districts with PfPR > 0.35 are categorized as high-risk regions. Given a class distribution in which 77.2% of observations belong to the high-risk category, classification performance is assessed using accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC).

Results And Discussion

Exploratory data analysis

Before fitting the predictive models, we inspected the distribution of the main covariates against the binary high-risk label. This preliminary reading was useful for checking whether the synthetic dataset behaved in a plausible epidemiological direction and whether any single variable dominated the separation between risk classes. Figure 2 shows that rainfall and temperature provide the clearest visual separation between low-risk and high-risk observations, whereas humidity, elevation, NDVI, population density, vector suitability, and intervention coverage show wider overlap between the two groups.

We then examined pairwise correlations to assess whether the selected variables carried complementary signals rather than simply duplicating the same information. As shown in Figure 3, PfPR is positively correlated with rainfall ($r = 0.65$), temperature ($r = 0.46$), NDVI ($r = 0.31$), and vector suitability ($r = 0.26$), while intervention coverage is negatively correlated with PfPR ($r = -0.15$). These patterns are consistent with the assumptions used to generate the synthetic data and with the expected ecological influence of rainfall, temperature, vegetation, and preventive coverage on malaria risk.

How to cite this article:




FIGURE 2: Distributions of eight environmental and epidemiological covariates stratified by malaria risk class (blue = low risk, red = high risk). Rainfall and temperature display the most pronounced class separation.

The overlap between low- and high-risk districts is important. It indicates that the classification task is not reducible to a single threshold on one climate variable; instead, the model must combine several weak and moderate signals. This observation supports the choice of a multimodal architecture rather than a purely univariate or single-source modeling strategy.

How to cite this article:




FIGURE 3: Pearson correlation matrix for nine continuous features. PfPR shows the strongest positive correlation with rainfall ($r = 0.55$) and temperature ($r = 0.52$), and a negative correlation with intervention coverage ($r = -0.17$).

PfPR, *Plasmodium falciparum* Parasite Rate; NDVI, Normalized Difference Vegetation Index

The matrix also shows that most predictor-predictor correlations remain modest. This reduces the immediate risk of severe collinearity and suggests that the covariates contribute different kinds of information to the risk-estimation task. In practice, this is the type of feature structure expected when epidemiological records, climate indicators, environmental variables, and demographic information are combined.

Finally, we compared PfPR distributions by country to check whether the synthetic design produced an obvious country-level imbalance. Figure 4 indicates broadly comparable medians and interquartile ranges across the DRC, Kenya, Nigeria, Tanzania, and Uganda. This result is not presented as a real country comparison; rather, it confirms that the experimental dataset does not allow one country label to dominate the learning process.

How to cite this article:




FIGURE 4: Country-level PfPR distributions. Median values and interquartile ranges are comparable across the five study countries in the synthetic dataset.

PfPR, *Plasmodium falciparum* Parasite Rate

Taken together, Figures 2-4 show a coherent but deliberately controlled data structure: the simulated risk patterns are strong enough to evaluate the proposed pipeline, but not intended to replace field surveillance or MAP-derived estimates. For this reason, the subsequent results are interpreted as methodological evidence for feasibility, transparency, and workflow design.

Multimodal fusion network training

Table 2 summarizes validation performance at selected training epochs. The MultimodalFusionNet demonstrates stable and progressive convergence throughout training. Validation RMSE decreases from 0.5148 at epoch 1 to 0.0677 at epoch 80, while validation R^2 increases from -14.6817 to 0.7285. These results indicate consistent learning behavior and suggest the absence of significant overfitting.

Epoch	Train Loss (MSE)	Val RMSE	Val MAE	Val R ²
1	0.3411	0.5148	0.4986	-14.6817
10	0.0623	0.1728	0.1510	-0.7667
20	0.0181	0.2162	0.2007	-1.7652
30	0.0116	0.0988	0.0839	0.4226
40	0.0081	0.1097	0.0964	0.2878
50	0.0069	0.0699	0.0574	0.7106
60	0.0070	0.0677	0.0554	0.7290
70	0.0066	0.0687	0.0563	0.7204
80	0.0062	0.0677	0.0554	0.7285

TABLE 2: Validation performance of the MultimodalFusionNet across selected training epochs.
MSE, Mean Squared Error; Val MAE, Validation Mean Absolute Error; Val RMSE, Validation Root Mean Squared Error

Figure 5 presents the corresponding learning curves, showing a steady reduction in training loss and validation error throughout the optimization process. The cosine annealing learning rate schedule contributes to stable convergence during the later stages of training.




FIGURE 5: Learning curves of the MultimodalFusionNet over 80 training epochs. Training MSE loss (left panel) and validation RMSE (right panel) decrease consistently, with cosine annealing promoting stable late-stage convergence.

MSE, Mean Squared Error; RMSE, Root Mean Squared Error

Benchmark comparison

Table 3 compares PfPR regression performance across all evaluated models on the held-out test dataset. Linear Regression and Ridge Regression achieve the best overall performance, with identical RMSE values of 0.0395 and R^2 values of 0.9056 and 0.9057, respectively. This outcome is expected given the approximately linear relationship between the covariates and PfPR embedded within the synthetic data generation process.

Gradient Boosting ranked third ($RMSE = 0.0409$, $R^2 = 0.8992$), followed by Random Forest ($RMSE = 0.0443$, $R^2 = 0.8818$) and the proposed MultimodalFusionNet ($RMSE = 0.0613$, $R^2 = 0.7734$). The fact that the simpler linear models performed best is consistent with the approximately linear structure intentionally embedded in the synthetic data-generation process.

Model	RMSE	MAE	R^2
Linear Regression	0.0395	0.0321	0.9056
Ridge Regression	0.0395	0.0321	0.9057
Gradient Boosting	0.0409	0.0333	0.8992
Random Forest	0.0443	0.0366	0.8818
Multimodal Fusion (this study)	0.0613	0.0495	0.7734

TABLE 3: Test set PfPR regression performance across all evaluated models ($n_{test} = 180$).

MSE, Mean Squared Error; PfPR, *Plasmodium falciparum* Parasite Rate; RMSE, Root Mean Squared Error

We therefore interpret the neural model primarily as a flexible integration backbone for multimodal inputs, graph-based extensions, and explainability, rather than as evidence of predictive superiority on this controlled dataset. Figure 6 presents the comparative RMSE, MAE, and R^2 values.




FIGURE 6: Bar chart comparison of RMSE (lower = better), MAE (lower = better), and R^2 (higher = better) across all evaluated models on the test set.

MSE, Mean Squared Error; PfPR, *Plasmodium falciparum* Parasite Rate; RMSE, Root Mean Squared Error

Figure 7 shows the relationship between predicted and observed PfPR for the MultimodalFusionNet.

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. Cureus J Comput Sci 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>




FIGURE 7: Predicted versus actual PfPR for the MultimodalFusionNet on the test set. Points cluster tightly around the identity line (dashed red), with slight overestimation at high prevalence values.

PfPR, *Plasmodium falciparum* Parasite Rate

Explainability results

The finding is consistent with the strong structural relationship between incidence and PfPR embedded within the synthetic dataset. Rainfall (10.2%) and temperature (6.2%) rank as the second and third most influential variables, respectively, reflecting the well-established biological influence of climatic conditions on mosquito development, survival, and parasite sporogony. Table 4 summarizes the relative SHAP importance of all model features.

How to cite this article:

Rank	Feature	Modality	SHAP Importance (%)
1	Incidence	Epidemiological	67.2
2	Rainfall	Climate	10.2
3	Temperature	Climate	6.2
4	Vector suitability	Epidemiological	3.5
5	NDVI	Geographic	3.0
6	Intervention coverage	Epidemiological	3
7	Population density	Demographic	2.3
8	Latitude	Geographic	1.3
9–12	Elevation, Longitude, Humidity, Country	Mixed	< 1.0 each

TABLE 4: Global SHAP feature importance (mean |SHAP value| expressed as percentage of total).

NDVI, Normalized Difference Vegetation Index; SHAP, SHapley Additive exPlanations

Figure 8 presents the global feature importance rankings obtained using the SHAP TreeExplainer applied to the random forest model. Incidence rate emerges as the dominant predictor, accounting for 67.2% of the total mean absolute SHAP value. This dominance is a direct consequence of the synthetic data construction: incidence is generated as $\text{incidence} = \text{PfPR} \times U(0.8, 1.2) \times 500$, making it effectively a noisy rescaling of the target variable. An ablation study (Table 4) removing incidence from all models confirms this: R^2 drops by 0.067-0.124 across all baselines, demonstrating that climate and environmental predictors carry the remaining signal independently. In real surveillance data, where incidence is measured from independent health facility records, this artefact does not apply.

How to cite this article:




FIGURE 8: Global feature importance derived from SHAP TreeExplainer applied to the random forest surrogate. Incidence rate dominates (67.2%), with rainfall (10.2%) and temperature (6.2%) as the most important climate predictors.
NDVI, Normalized Difference Vegetation Index; SHAP, SHapley Additive exPlanations

Figure 9 presents the SHAP beeswarm summary plot, which illustrates both the magnitude and direction of feature contributions across all observations. High incidence values (shown in pink) consistently increase predicted PfPR, whereas low incidence values (shown in blue) contribute to lower risk estimates. Rainfall and temperature display predominantly positive effects on predicted malaria risk, aligning with known ecological mechanisms influencing mosquito breeding and parasite development. Conversely, intervention coverage exhibits a largely negative contribution, whereby higher coverage levels reduce predicted risk. This behavior is consistent with the protective effect intentionally encoded within the synthetic data-generation process.




FIGURE 9: SHAP beeswarm summary plot showing the distribution, magnitude, and direction of each feature's impact on predicted PfPR across all test observations. Feature values are color-coded from low (blue) to high (pink).

NDVI, Normalized Difference Vegetation Index; PfPR, *Plasmodium falciparum* Parasite Rate; SHAP, SHapley Additive exPlanations

Figure 10 provides a local SHAP waterfall explanation for district index 176, identified as the highest-risk observation in the test set (predicted PfPR = 0.793). Beginning from the model's baseline expected value ($E[f(X)] = 0.476$), incidence contributes +0.27 to the final prediction, while temperature and rainfall contribute +0.02 and +0.01, respectively. In contrast, intervention coverage contributes -0.01, representing the most important protective factor for this district. Such local explanations can assist malaria control officers in understanding the specific combination of factors driving elevated malaria risk within a given administrative area.

How to cite this article:




FIGURE 10: SHAP waterfall explanation for the highest-risk test district (predicted PfPR = 0.7927). The waterfall quantifies each feature's additive contribution relative to the model's expected output ($E[f(X)] = 0.476$), with incidence contributing the dominant positive effect (+0.27).

NDVI, Normalized Difference Vegetation Index; PfPR, *Plasmodium falciparum* Parasite Rate; SHAP, SHapley Additive exPlanations.

Binary classification performance

The binary malaria risk classification task (PfPR > 0.35) achieved strong predictive performance across all evaluation metrics, as summarized in Table 5. The random forest classifier attained an accuracy of 93.9%, a precision of 96.0%, a recall of 96.6%, and an F1-score of 96.3%. Furthermore, the AUC-ROC reached 0.9740, indicating near-perfect discrimination between low-risk and high-risk districts.

Metric	Value
Accuracy	0.9389
Precision	0.9597
Recall	0.9662
F1-score	0.9630
AUC-ROC	0.9740

TABLE 5: Binary high-risk classification performance (Random Forest Classifier, $n_t^{est} = 180$).

AUC-ROC, Area Under the Receiver Operating Characteristic Curve

From a public health perspective, the exceptionally high recall (96.6%) is particularly important. In operational malaria surveillance, failing to identify a genuinely high-risk district may result in insufficient intervention and increased disease burden. Consequently, minimizing false negatives is often more critical than minimizing false positives, given the potentially severe consequences associated with under-intervention.

Table 5 gives an overall view of the classifier's performance, but Figure 11 helps clarify where the model succeeds and where it still makes errors. This is important in malaria surveillance because a missed high-risk district can have more serious consequences than a district wrongly flagged for closer monitoring. The confusion matrix and ROC curve therefore provide a more practical reading of the model's reliability.




FIGURE 11: Classification performance on the test set. Left: confusion matrix showing 16 true low-risk, 143 true high-risk, 6 false positives, and 5 false negatives. Right: ROC curve with AUC = 0.974, demonstrating near-perfect discrimination.

AUC-ROC, Area Under the Receiver Operating Characteristic Curve

Spatial risk map

Figure 12 presents the malaria risk map generated by applying the random forest model to a regular 50 × 50 spatial grid covering the geographic extent of the study area. The resulting contour map illustrates broad spatial risk gradients and serves primarily as a proof-of-concept demonstration of the proposed cartographic visualization pipeline.

How to cite this article:

Because the analysis relies on synthetic data, the displayed risk patterns should not be interpreted as representing actual malaria transmission intensity. Nevertheless, when applied to real-world MAP datasets, the same framework could generate operational risk maps capable of highlighting genuine transmission hotspots. Such visualizations could provide district health officers and public health decision-makers with an intuitive spatial decision-support tool for prioritizing surveillance, prevention, and intervention activities.




FIGURE 12: Predicted malaria risk map (PfPR) over the Sub-Saharan Africa study region generated by applying the trained Random Forest model to a synthetic 50×50 prediction grid. Warmer colors (red) indicate higher predicted parasite prevalence.
PfPR, *Plasmodium falciparum* Parasite Rate

Discussion

This study presents a methodological proof-of-concept for an explainable multimodal AI framework for malaria risk prediction in Sub-Saharan Africa. Since the evaluation used MAP-inspired synthetic data, the results should not be interpreted as operational evidence for specific countries or districts. The main contribution is the integration of modality-specific encoders, spatial GNN modeling, and SHAP explanations into a single surveillance-oriented pipeline.

Key Findings

The SHAP analysis shows that incidence rate was the dominant predictor, mainly because of the synthetic data-generation structure. In real surveillance settings, incidence data may be affected by underreporting, unequal access to healthcare, and differences in local reporting systems. Therefore, climate and environmental variables such as rainfall and temperature may become more influential when the framework is applied to real data. Their positive contribution in this study remains consistent with known malaria transmission mechanisms, while intervention coverage shows the expected protective effect.

The benchmark comparison also reflects the limits of synthetic evaluation. Linear and ridge regression performed better because PfPR was generated from an approximately linear combination of variables. This does not reduce the relevance of multimodal neural and graph-based models, which may be more useful in real settings where malaria transmission involves non-linear, spatial, seasonal, and ecological interactions.

How to cite this article:

The spatial GNN component is important because malaria risk can extend beyond administrative boundaries, especially in border or high-mobility areas. Likewise, SHAP explanations help translate model outputs into information that public health officers can understand and use for intervention planning.

Comparison With Previous Work

Our results align with previous studies demonstrating the value of multimodal approaches for disease prediction [13,14]. The strong performance of climate variables (rainfall and temperature) corroborates findings from Bhatt et al. [3] and Weiss et al. [2] regarding the environmental drivers of malaria transmission. The protective effect of intervention coverage is consistent with the documented impact of vector control measures [1].

Limitations

This study has several limitations that should be acknowledged:

- 1) Synthetic data: The evaluation used simulated rather than real surveillance data. While this allowed controlled experimentation, it limits the generalizability of findings.
- 2) Cross-sectional design: The current framework is static and does not capture temporal dynamics in malaria transmission.
- 3) Geographic graph construction: The k-nearest neighbor approach based solely on distance may not capture all relevant transmission connections (e.g., human mobility patterns).
- 4) Model complexity: The neural model's performance did not exceed simpler linear baselines on this synthetic dataset, though it offers advantages for multimodal integration.

Recommendations for Future Validation and Operational Use

First, the framework should be validated with real MAP observations, national routine health information system data, and where available, entomological and intervention-coverage records. The synthetic dataset was useful for controlling assumptions and testing the pipeline, but real deployment requires external validation across seasons, countries, and administrative levels.

Second, the spatial graph should be refined beyond geographic proximity alone. Future versions should test adjacency based on travel time, road connectivity, human mobility, ecological similarity, shared borders, and cross-border health-zone interactions. Such graph definitions may better capture how malaria risk is spatially connected in practice.

Third, the model should be extended from a cross-sectional prototype to a temporal surveillance system. Monthly or seasonal inputs from rainfall, temperature, vegetation, incidence, and intervention activities would allow Temporal GNNs or other sequence models to capture delayed effects and seasonal transmission dynamics.

Fourth, explainability should be co-designed with malaria-control officers. Global SHAP rankings are useful for scientific interpretation, but district teams often need simpler outputs: a short explanation of the risk drivers, a confidence flag, and a recommended follow-up action. A lightweight dashboard, designed for low-bandwidth settings, would make the framework more realistic for public health use in the DRC and similar contexts.

Finally, any operational version should include model cards, data-provenance documentation, periodic recalibration, and a clear warning that predictions support - but do not replace - epidemiological judgment. With these safeguards, the proposed framework can evolve from a proof-of-concept into a transparent decision-support tool for malaria surveillance and targeted intervention planning.

Conclusions

This study presented an explainable multimodal AI framework for district-level malaria risk prediction in Sub-Saharan Africa. The proposed system integrates epidemiological, climatic, geospatial, and demographic information through modality-specific encoder networks and a spatial GCN extension. On a proof-of-concept synthetic dataset, the framework achieved a validation R^2 of 0.7285 for PfPR regression and an AUC of 0.9740 for binary high-risk classification. These results support the feasibility of the proposed workflow under controlled conditions, while confirming the need for real-world validation before operational use. The SHAP analysis added an important layer of interpretability by identifying incidence, rainfall, temperature, vector

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. *Cureus J Comput Sci* 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>

suitability, NDVI, and intervention coverage as the main contributors to predicted risk. In a real surveillance setting, this type of explanation would help public health teams move beyond a simple risk score and ask why a district is being flagged. This is particularly valuable where resources are limited and decisions on bed nets, indoor residual spraying, community health worker deployment, and intensified surveillance must be justified.

Four principal conclusions emerge from this study: 1) A modality-encoder fusion architecture provides an effective mechanism for integrating heterogeneous malaria risk factors without restricting analysis to a single data modality, while maintaining a relatively compact parameter budget of 8,169 trainable weights. 2) The spatial GNN extension enables explicit modeling of inter-district transmission dependencies through geographic k-nearest-neighbor graph construction, thereby capturing spatial spillover effects that conventional district-independent approaches may overlook. 3) A three-level SHAP explainability framework generates global feature importance rankings, beeswarm visualizations, and local waterfall explanations that transform model predictions into actionable intelligence for malaria surveillance and control activities. 4) Benchmarking against four baseline approaches established a transparent evaluation protocol and clarified the behavior of the framework under controlled synthetic conditions.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Kiro M. Lauraine, Moïse K. Kasambya, Mundeke K. Ghislain, Charité K. Valambaya, Amanio P. Isaac, Hangi M. Patrice, Alain K. Virekero, Elie M. Mwandu, Belden K. Mbafumoja, Muyisa M. Katemengi

Acquisition, analysis, or interpretation of data: Kiro M. Lauraine, Moïse K. Kasambya, Mundeke K. Ghislain, Charité K. Valambaya, Amanio P. Isaac, Hangi M. Patrice, Alain K. Virekero, Elie M. Mwandu, Belden K. Mbafumoja, Muyisa M. Katemengi

Drafting of the manuscript: Kiro M. Lauraine, Moïse K. Kasambya, Mundeke K. Ghislain, Charité K. Valambaya, Amanio P. Isaac, Hangi M. Patrice, Alain K. Virekero, Elie M. Mwandu, Belden K. Mbafumoja, Muyisa M. Katemengi

Critical review of the manuscript for important intellectual content: Kiro M. Lauraine, Moïse K. Kasambya, Mundeke K. Ghislain, Charité K. Valambaya, Amanio P. Isaac, Hangi M. Patrice, Alain K. Virekero, Elie M. Mwandu, Belden K. Mbafumoja, Muyisa M. Katemengi

Supervision: Kiro M. Lauraine

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

The data that support the findings of this study are openly available at <https://data.malariaatlas.org/trends?year=2024&metricGroup=Malaria&geographicLevel=admin0&metricSubcategory=Pf&metricType=rate&metricName=incidence>. All code and synthetic data used in this study are publicly available at <https://github.com/KasMoise/malaria>

References

1. World Health Organization: World malaria report 2023. (2023). <https://www.who.int/publications/i/item/9789240086173>.

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. *Cureus J Comput Sci* 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>

2. Weiss DJ, Lucas TCD, Nguyen M, et al.: [Mapping the global prevalence, incidence, and mortality of Plasmodium falciparum, 2000-17: a spatial and temporal modelling study](#). The Lancet. 2019, 394:322-331. [10.1016/s0140-6736\(19\)31097-9](#)
3. Bhatt S, Weiss DJ, Cameron E, et al.: [The effect of malaria control on Plasmodium falciparum in Africa between 2000 and 2015](#). Nature. 2015, 526:207-211. [10.1038/nature15535](#)
4. Dalrymple U, Arambepola R, Gething PW, Cameron E: [How long do rapid diagnostic tests remain positive after anti-malarial treatment?](#). Malaria Journal. 2018, 17:228. [10.1186/s12936-018-2371-9](#)
5. Holzinger A, Saranti A, Molnar C, Biecek P, Samek W: [Explainable AI methods—A brief overview](#). xxAI - Beyond Explainable AI. Holzinger A, Goebel R, Fong R, Moon T, Müller KR, Samek W (ed): Springer, Cham; 2022. 13200:13-38. [10.1007/978-3-031-04083-2_2](#)
6. Kipf TN, Welling M: [Semi-supervised classification with graph convolutional networks](#). arXiv. 2017, [10.48550/arXiv.1609.02907](#)
7. Veličković P, Cucurull G, Casanova A, Romero A, Liò P, Bengio Y: [Graph attention networks](#). arXiv. 2018, [10.48550/arXiv.1710.10903](#)
8. Deng S, Wang S, Rangwala H, Wang L, Ning Y: [Cola-GNN: Cross-location attention based graph neural networks for long-term ILI prediction](#). CIKM '20: Proceedings of the 29th ACM International Conference on Information & Knowledge Management. 2020, 245-254. [10.1145/3340531.3411975](#)
9. Kapoor A, Ben X, Liu L, Perozzi B, Barnes M, Blais M, O'Banion S: [Examining COVID-19 forecasting using spatio-temporal graph neural networks](#). arXiv. 2020, [10.48550/arXiv.2007.03113](#)
10. Arrieta AB, Díaz-Rodríguez N, Del Ser J, et al.: [Explainable Artificial Intelligence \(XAI\): Concepts, taxonomies, opportunities and challenges toward responsible AI](#). Information Fusion. 2020, 58:82-115. [10.1016/j.inffus.2019.12.012](#)
11. Lundberg SM, Lee SI: [A unified approach to interpreting model predictions](#). NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems. 2017, 4768-4777.
12. Lundberg SM, Erion G, Chen H, et al.: [From local explanations to global understanding with explainable AI for trees](#). Nature Machine Intelligence. 2020, 2:56-67. [10.1038/s42256-019-0138-9](#)
13. Baltrusaitis T, Ahuja C, Morency LP: [Multimodal machine learning: A survey and taxonomy](#). IEEE Transactions on Pattern Analysis and Machine Intelligence. 2019, 41:423-443. [10.1109/tpami.2018.2798607](#)
14. Chen RJ, Ding T, Lu MY, et al.: [Towards a general-purpose foundation model for computational pathology](#). Nature Medicine. 2024, 30:850-862. [10.1038/s41591-024-02857-3](#)
15. Brown BJ, Manescu P, Przybylski AA, et al.: [Data-driven malaria prevalence prediction in large densely populated urban holoendemic sub-Saharan West Africa](#). Scientific Reports. 2020, 10:15918. [10.1038/s41598-020-72575-6](#)
16. Thawer SG, Golumbeanu M, Lazaro S, et al.: [Spatio-temporal modelling of routine health facility data for malaria risk micro-stratification in mainland Tanzania](#). Scientific Reports. 2023, 13:10600. [10.1038/s41598-023-37669-x](#)
17. Pourtois JD, Tallam K, Jones I, et al.: [Climatic, land-use and socio-economic factors can predict malaria dynamics at fine spatial scales relevant to local health actors: Evidence from rural Madagascar](#). PLOS Global Public Health. 2023, 3:e0001607. [10.1371/journal.pgph.0001607](#)
18. Ghilardi L, Okello G, Nyondo-Mipando L, et al.: [How useful are malaria risk maps at the country level? Perceptions of decision-makers in Kenya, Malawi and the Democratic Republic of Congo](#). Malaria Journal. 2020, 19:353. [10.1186/s12936-020-03425-z](#)
19. Carman M, Rosman B: [Applying a principle of explicability to AI research in Africa: should we do it?](#). Ethics and Information Technology. 2021, 23:107-117. [10.1007/s10676-020-09534-2](#)
20. Mpia HN, Mwendia SN, Mburu LW: [Predicting employability of Congolese information technology graduates using contextual factors: Towards sustainable employability](#). Sustainability. 2022, 14:13001. [10.3390/su142013001](#)
21. Mpia HN, Mburu LW, Mwendia SN: [CoBERT: A Contextual BERT model for recommending employability profiles of information technology students in unstable developing countries](#). Engineering Applications of Artificial Intelligence. 2023, 125:106728. [10.1016/j.engappai.2023.106728](#)
22. Mpia HN, Kavalami MM, Lusenge GK, Ushindi KP, Kambale Kyalengekania DD, Ciswa OM: [Lightweight deep learning system for infant-cry recognition with real-time notification in resource-constrained environments](#). Machine Learning with Applications. 2025, 22:100791. [10.1016/j.mlwa.2025.100791](#)
23. Tjoa E, Guan C: [A survey on explainable artificial intelligence \(XAI\): Toward medical XAI](#). IEEE Transactions on Neural Networks and Learning Systems. 2020, 32:4793-4813. [10.1109/tnnls.2020.3027314](#)

How to cite this article:

Lauraine K M, Kasambya M K, Ghislain M K, et al. (July 31, 2026) Explainable Graph Neural Networks for Multimodal Malaria Risk Prediction Using Malaria Atlas Project (MAP)-Inspired Synthetic Data and Climate Indicators in Sub-Saharan Africa. Cureus J Comput Sci 3 : es44389-026-00197-6. DOI <https://doi.org/10.7759/s44389-026-00197-6>