

Predicting Pathogen Emergence Using a Spatiotemporal Graph Neural Network and Explainable AI Approach

Oreofe Jolaolu

jason.jolaolu@gmail.com

Independent Researcher <https://orcid.org/0009-0004-2039-2686>

Dr. Okoh J.C

Independent Researcher <https://orcid.org/0000-0001-7872-5229>

Research Article

Keywords: One Health, Spatiotemporal Analysis, Pathogen Emergence, Graph Neural Networks, Explainable AI, Disease Forecasting, Outbreak Prediction, Public Health Surveillance , Machine Learning, xAI

Posted Date: August 18th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10704919/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Spatiotemporal Data Mining and Forecasting of Pathogen Emergence

(Using Spatiotemporal Graph Neural Network (STGNN) + xAI)

Contributing Authors

Jolaolu O. J.¹, Dr. Okoh J.C.² & Amusan K. A³

¹*ORCID ID: 0009-0004-2039-2686*

²*ORCID ID: 0000-0001-7872-5229*

³*ORCID ID: 0009-0001-9575-7596*

Abstract

Spatiotemporal dynamics of emerging infectious diseases continue to present significant challenges for traditional epidemiological forecasting. Driven by climate change and global mobility, outbreak patterns require advanced, transparent modeling to guide proactive public health interventions. This study presents a framework for integrating Spatiotemporal Graph Neural Networks (STGNN) and Explainable Artificial Intelligence (xAI) to forecast pathogen emergence risk. Utilising a Human Epidemic Database (1,044 events across distinct territories from 2015 to 2020), we construct a robust graph where geographic regions act as nodes connected by spatial and travel edges.

The STGNN architecture captures complex transmission pathways, while subsequent xAI techniques, including SHAP, uncover the key drivers behind these predictions, such as historical outbreak counts and pathogen diversity. Across rolling temporal test sets, the STGNN model outperformed all baseline architectures, achieving an R^2 score of 0.6959 and an MAE of 0.4694.

Uncertainty decomposition on validation instances indicated that aleatoric uncertainty accounted for 98.85% of total predictive variance, while epistemic uncertainty contributed 1.15%. The results demonstrate improved forecasting accuracy while explicitly quantifying the relative contributions of observational noise and model uncertainty.

The proposed framework will empower global health authorities to transition from passive surveillance to proactive, data driven resource allocation. Lastly, the framework presents an opportunity to help strengthen health security against future pathogenic threats, laying a vital foundation for future integration with broader climate data variables.

Keywords: One Health, Spatiotemporal Analysis, Pathogen Emergence, Graph Neural Networks, Explainable AI, Disease Forecasting, Outbreak Prediction, Public Health Surveillance, Machine Learning, xAI

Introduction

The continuous recurrence of infectious diseases remain a critical threat to global public health and societal stability (Fauci & Morens, 2012; Jones et al., 2008). Recent global health crises, notably the COVID-19 pandemic, alongside persistent threats like tuberculosis and dengue, highlight that infectious diseases are deeply intertwined with the welfare of societies (Baker et al., 2022).

Research shows that over sixty percent (60%) of human infectious diseases are zoonotic, originating from pathogens that jump between animals and humans due to human activities like habitat destruction and global mobility (Jones et al., 2008; Karesh et al., 2012). This interconnected reality forms the foundation of the One Health framework, an approach endorsed by the World Health Organization that views human, animal, and environmental health as intrinsically linked (World Health Organization, 2024; Zinsstag et al., 2011; Mackenzie & Jeggo, 2019). In disease forecasting, One Health demands a proactive, ecosystem-based approach that integrates diverse data streams to monitor and predict interrelated risks (Destoumieux-Garzón et al., 2018; Gibbs, 2014).

Developing accurate epidemic forecasting models has become a critical priority for public health management (Viboud et al., 2018; Chretien et al., 2014). Historically, approaches have fallen into three distinct categories, each with its own strengths and critical limitations:

Model Type	Strengths	Limitations
Compartmental Models	Mechanistic clarity and inherent interpretability	Assumes homogeneous mixing, struggles to integrate high-dimensional data like climate or mobility (Keeling & Rohani, 2008; Viboud et al., 2018)
Statistical Models	Good for baseline time-series analysis and regression	Weak out-of-sample generalisability, poorly suited for complex, multimodal digital surveillance data (Brownstein et al., 2009)
Machine Learning (ML)	Discovers complex, nonlinear patterns; highly capable of integrating diverse data streams	Operates as a "black box," prioritising predictive accuracy over transparency (Rudin, 2019)

Table 1.0 Forecasting Model Categories

While Machine Learning offers superior predictive power, its lack of transparency is a major barrier in One Health contexts. Professionals from varying disciplines must collaboratively respond to threats and require clear, actionable rationales for a model's predictions (Desvars-Larrive et al., 2024).

Spatiotemporal modeling has gained particular prominence in infectious disease research, as outbreaks exhibit pronounced spatial and temporal dependencies that are critical for accurate forecasting (Chen et al., 2026; Han et al., 2025). Spatiotemporal Graph Neural Networks (STGNNs) represent a contemporary approach that combines the representational power of graph neural networks with temporal processing to capture complex spatial interactions and temporal dynamics simultaneously (Wu et al., 2021; Han et al., 2025).

Despite their advanced capabilities, current STGNN applications face some challenges including:

- **Oversimplified Spatial Assumptions:** They often assume locations with similar features will share similar future trajectories, ignoring underlying mechanistic heterogeneity like local interventions, medical access, and virus mutations (Han et al., 2025; Liu et al., 2024).

- **Data Dependency:** They rely heavily on fine grained mobility or detailed epidemiological data, which is rarely accessible in real world scenarios (Cao et al., 2026).
- **Lack of Epidemiological Mechanisms:** They largely remain "black boxes" divorced from underlying domain theory, limiting their interpretability and practical utility for decision making, forcing trade off between predictive performance and transparency (Han et al., 2025; Rudin, 2019).

To address this gap, the study introduces Explainable Artificial Intelligence (xAI), comprising tools designed to make machine learning models transparent to human users (Arrieta et al., 2020). Recent successes in forecasting diseases like the Dengue and West Nile virus have proven that xAI can untangle complex drivers of disease spread that traditional models miss.

To build and validate this new framework globally, the research utilises the Human Epidemic Database (HED), constructed from official, open-source surveillance reports (Badker et al., 2025).

- **Scope:** Contains 1,044 epidemic events across various countries and territories.
- **Timeframe:** Disease onset between 2015 and 2020.
- **Pathogens:** Covers over 120 distinct pathogens.

This dataset allows for spatiotemporal forecasting at a global scale, overcoming the limitations of previous studies that focused only on regional outbreaks or single diseases (Badker et al., 2025; Viboud et al., 2018; Chretien et al., 2014).

Research Objectives

Existing forecasting approaches fail to adequately address the need for both spatiotemporal predictive accuracy and mechanistic interpretability (Rudin, 2019; Han et al., 2025). There is a pressing need for a framework that integrates epidemiology-informed modeling, spatiotemporal graphs, and xAI to enable accurate, globally scalable forecasting with actionable insights. The key objectives include:

- **Data Pipeline Development:** Transform raw HED data into a structured dataset by engineering features like historical counts and pathogen diversity. It also included features such as geographic spread metrics.
- **Model Architecture:** Design and validate an STGNN that captures both observed spatiotemporal patterns and hidden mechanistic heterogeneity in transmission.
- **Implementation of xAI:** Apply xAI techniques to extract global feature importance rankings and local, prediction-specific explanations.
- **Evaluation and Mapping:** Compare the framework against baseline models and generate spatiotemporal risk maps that visualise evolving high risk regions.

The remainder of this research work is organised as follows: The literature review examines existing approaches to epidemic forecasting and spatiotemporal modeling. The methodology section describes the research design, data preprocessing, graph construction, model architecture, training procedures, and xAI implementation. The results and discussion section presents the experimental findings and interprets the results. Finally, the conclusion summarises the research contributions and outlines directions for future research.

Literature Review

One Health Framework and Infectious Disease Surveillance

Infectious diseases remain one of the most pressing threats to global health and societal well-being (Baker et al., 2022; Fauci & Morens, 2012). In response, the One Health framework has been proposed as an integrated approach to balance and optimise the health of people, animals, and the environment (World Health Organization, 2024). The central concept is that all three domains are interconnected, with each domain directly or indirectly influencing the others (Zinsstag et al., 2011; Mackenzie & Jeggo, 2019). Practice oriented efforts to operationalise One Health for infectious disease surveillance and forecasting are being adopted. The activities include integrated surveillance platforms that link human, animal, and environmental data streams and big-data technologies for multisource data fusion and situational awareness (Desvars-Larrive et al., 2024).

Machine Learning and Deep Learning for Epidemic Forecasting

In response to the growing volume and variety of data on infectious disease drivers (alongside stringent timeliness requirements), machine learning methods are increasingly being employed (Johansson et al., 2019; Viboud et al., 2018; Chretien et al., 2014). Common approaches include supervised learning (e.g., random forests and deep neural networks), semi-supervised learning, unsupervised learning (e.g., clustering and association rule mining), and reinforcement learning (Scornet et al., 2015; Nsoesie et al., 2013; Reich et al., 2019).

Deep learning models including Recurrent Neural Networks (RNNs), Long Short-Term Memory networks (LSTMs), Gated Recurrent Units (GRUs) and attention-based architectures such as the Transformer (Vaswani et al., 2017) and Temporal Fusion Transformer (Lim et al., 2021) capture complex nonlinear patterns. They have shown superior accuracy over traditional univariate time series models in forecasting healthcare metrics (Souza et al., 2023). Deep neural networks excel in identifying intricate, non-linear relationships within large datasets, thereby enhancing the precision and predictive capabilities of epidemiological analyses (Goodfellow et al., 2016; LeCun et al., 2015).

However, these approaches predominantly emphasise temporal dynamics and often neglect spatial heterogeneity and contact patterns, which shape transmission and limit predictive robustness in realistic settings. Furthermore, deep learning based approaches typically disregard intrinsic transmission dynamics and lack interpretability, operating as black-box learners that sever the link to domain theory (Rudin, 2019; Lipton, 2016).

Spatiotemporal Modeling and Graph Neural Networks

Spatiotemporal modeling has gained particular prominence in infectious disease research, as outbreaks exhibit pronounced spatial and temporal dependencies that are critical for accurate forecasting (Chen et al., 2026; Han et al., 2025).

Spatiotemporal Graph Neural Networks (STGNNs) represent an approach that combines the representational power of graph neural networks with temporal processing to capture complex spatial interactions and temporal dynamics simultaneously (Wu et al., 2021; Han et al., 2025). STGNNs model locations as nodes in a network and their relationships as edges. This enables the learning of both geographic dependencies and time-varying transmission patterns. Recent advances have demonstrated the potential of STGNNs in epidemic forecasting, with

applications to COVID-19, dengue, influenza, and other infectious diseases showing superior predictive performance compared to traditional methods (Han et al., 2025; Cao et al., 2026). STGNNs have also been successfully applied beyond infectious diseases, as demonstrated by Chen et al. (2026) in the context of opioid overdose death prediction using urine drug test data as dynamic features and Social Determinants of Health as static features, highlighting the versatility of STGNN architectures for public health surveillance.

Despite these advances, most STGNN methods bear an over simplified assumption that two locations with similar observed features in previous time steps will develop similar infection numbers in the future (Zheng et al., 2024). In fact, for any epidemic disease, there exists strong heterogeneity of its fundamental evolution mechanisms across geolocation and time (which can eventually lead to diverged infection numbers in two "similar" locations). However, such mechanistic heterogeneity is non trivial to be captured due to the existence of numerous influencing factors like medical resource accessibility, mobility patterns, etc., most of which are spatio-temporal yet unreachable or even unobservable (Zheng et al., 2024; Han et al., 2025). Most methods rely on over simplified assumptions about spatial heterogeneity, treating locations with similar observed features as having similar future infection trajectories while overlooking the underlying mechanistic heterogeneity of epidemic transmission processes (Han et al., 2025).

Explainable Artificial Intelligence (xAI) in One Health and Epidemic Forecasting

To address the limitations of both purely mechanistic and purely data-driven approaches, researchers have proposed epidemiology-informed machine learning (EML) models that incorporate mechanistic laws into a machine learning backbone to enhance generalisability and robustness (Han et al., 2025; Zheng et al., 2024; Liu et al., 2024). This trend also extends to STGNNs, where the time/ location sensitive parameters in epidemic models are predicted to support dynamic epidemic forecasting (Han et al., 2025; Cao et al., 2026).

Explainable Artificial Intelligence (xAI) comprises a growing set of tools designed to make machine learning models transparent and interpretable to human users (Arrieta et al., 2020). Techniques such as SHapley Additive exPlanations (SHAP) (Lundberg & Lee, 2017) and Local Interpretable Model-agnostic Explanations (LIME) (Ribeiro et al., 2016) demonstrate how model outputs can be broken down into human understandable components, identifying the most influential contributors to predictions and explaining the drivers behind individual predictions in specific contexts.

The lack of interpretability in AI models poses serious challenges in One Health, where public health professionals and policymakers must collaboratively respond to emerging threats (Desvars-Larrive et al., 2024). Transparent and explainable models are more likely to foster stakeholder trust, and support timely responses (Desvars-Larrive et al., 2024; Arrieta et al., 2020).

Cao et al. (2026) argued that embedding xAI within a One Health framework offers a new organising principle for infectious disease intelligence and xAI functions as an implementation layer that operationalises One Health by enabling mechanism-informed and transparent, accountable decision-making. Recent use cases have demonstrated the value of xAI integrated models in dengue forecasting, where SHAP and LIME revealed complex, non-linear thresholds and provided actionable targets that linear coefficients missed.

The Human Epidemic Database (HED)

The Human Epidemic Database (HED), constructed from official, open-source surveillance reports, provides a valuable resource for advancing spatiotemporal epidemic research. The dataset includes 1,044 epidemic events with onset between 2015 and 2020, covering more than 230 countries and territories and 120 pathogens (Badker et al., 2025).

The extensive and standardised nature of the dataset, structured through a methodology designed to produce the most reliable spatiotemporal view of outbreaks possible, makes it well-suited for both descriptive epidemiology and exploring outbreak dynamics and disease emergence (Badker et al., 2025). The dataset has also been used in the development of an Epidemic Preparedness Index, a quantitative metric for capturing country-level pandemic preparedness and has informed working papers on drivers of pandemic risk.

Research Gaps

- Existing STGNN based epidemic forecasting methods largely operate as black box learners that lack embedded epidemiological mechanisms, severing the link to domain theory and limiting their interpretability and actionable utility for public health decision-making (Han et al., 2025; Rudin, 2019)
- Second, most methods overlook mechanistic heterogeneity (the underlying intrinsic transmission mechanisms that determine divergent outbreak trajectories across locations) focusing instead on spatiotemporal heterogeneity in observed features (Zheng et al., 2024; Liu et al., 2024).
- Third, while hybrid epidemiology informed models show promise, few integrate explainable AI techniques to provide transparent, actionable insights for One Health decision-making (Desvars-Larrive et al., 2024).
- Fourth, forecasting frameworks that leverage comprehensive event databases have the potential for development. This study addresses these gaps by developing a novel framework that integrates STGNNs with epidemiology-informed modeling and xAI to enable both accurate forecasting and actionable insights at a global scale using the HED dataset.

Methodology

Graph Formulation and Input Structure

The framework structures epidemiological surveillance records as a dynamic spatial-temporal tensor to capture both localised geographic dependencies and longitudinal trends. The input data is formulated into a multi-dimensional tensor layout with a shape of $[Batch, Nodes, Timesteps, Features]$ (or $[B, S, F]$, depending on flattened configurations), derived from the cleaned dataset.

- **Spatial Nodes (N):** Defined across discrete spatial regions or country level entities using ISO identifiers (*iso2/country*). Spatial relationships between nodes are modeled using a symmetric spatial adjacency matrix A , where edge weights reflect relational connectivity.
- **Temporal Horizon (T):** Incorporates a structured lookback horizon spanning 225 timesteps of historical observation windows.
- **Feature Channels (F):** Comprises primary domain feature channels, specifically mapping to cumulative cases (*cuml_cases*, designated as F_0), cumulative deaths (*cuml_deaths*, F_1), pathogen taxonomy categorical groupings (*pathogen_group*, F_2).

We structure historical observations into lookback windows containing consecutive time steps. The complete input tensor feeds into the model with dimensions matching the batch size and the total number of graph nodes. It also includes the length of the temporal window, and the feature count.

Architecture Design and Spatio-Temporal Operations

The proposed Spatio-Temporal Graph Neural Network (STGNN) architecture is designed to jointly process spatial graph topologies and temporal dynamics for count-based epidemiological modeling:

1. **Spatial Graph Convolution:** A spatial graph convolutional layer aggregates localised structural information, allowing neighbouring nodes to share contextual signals through the adjacency matrix.
2. **Temporal Sequence Layers:** Sequential patterns and longitudinal trends are processed using temporal hidden layers configured with 128 hidden units.
3. **Poisson Rate Head (λ):** To properly model discrete epidemiological event counts, the network terminates in a Poisson Rate Head.
4. **Activation Function:** Utilising a Softplus activation function, the model outputs the Poisson rate parameter (λ), representing the expected mean count per node-snapshot.
5. **Training Configuration:** The model is optimised using the Adam optimiser with an initial learning rate of 0.001, an L2 weight decay of 0.0001, and a batch size of 32. Training is executed for up to 100 epochs employing an early stopping patience mechanism set to 10 epochs. The loss formulation utilises the Poisson Negative Log-Likelihood (NLL) loss function.




Fig 1.0 STGNN Architecture

Baseline Comparison and Benchmarking Protocol

We test the proposed STGNN framework against alternative baseline approaches to isolate the value of joint spatio-temporal message passing:

- **Standard GCN:** Evaluates spatial topological propagation without accounting for dynamic temporal sequence lookbacks.
- **Standard LSTM:** Evaluates performance using longitudinal patterns without explicit spatial graph structures.
- **XGBoost:** Serves as a non deep learning baseline operating on flat feature representations.

Evaluation and comparison are performed using standard statistical metrics including the coefficient of determination (R^2), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE).

Hyperparameter	Selected Value
Optimisation Algorithm	Adam Optimiser
Initial Learning Rate	0.001
Weight Decay (L2 Regularization)	0.0001
Batch Size	32
GCN Hidden Dimensions	64 channels
LSTM Hidden Units	128 units
Monte Carlo Dropout Rate	0.20 (20%)
Monte Carlo Inference Passes	50 Stochastic Passes
Loss Formulation	Poisson Negative Log-Likelihood (NLL)
Training Epochs / Early Stopping	100 Epochs (Patience = 10)

Table 3.0 Model Hyperparameters and Training Setup

Feature Attribution and Dimensional Unflattening

To interpret model decisions and ensure transparency, Explainable AI (xAI) modules integrate Kernel SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations):

- **Memory-Optimised Sampling:** Kernel SHAP is executed using a background baseline sample cap of 30 instances and an evaluation sample cap of 50 instances (with LIME evaluated on 5 representative local instances using 250 local perturbation samples).
- **Dimensional Unflattening:** Because tabular explainers require 2D input matrices, high-dimensional tensors are flattened during explanation passes. A dedicated unflattening pipeline processes the raw explanation vectors to map flat indices back to explicit structural coordinates, yielding a decoded spatial-temporal mapping that details the *Node Index*, *Time Lag*, *Feature Name*, and *Mean Absolute SHAP* impact.

Uncertainty Quantification and Variance Decomposition

Uncertainty Quantification (UQ) is integrated to evaluate prediction reliability and isolate error sources using the Monte Carlo Dropout inference:

- **Stochastic Passes:** The model executes 50 stochastic inference passes utilising a Monte Carlo dropout rate of 0.20 (20%) on feature inputs.
- **Variance Decomposition:** Total predictive variance is decomposed into
 - **Aleatoric Uncertainty** (inherent data noise, which constitutes 98.85% of the variance, mathematically consistent with Poisson distributions where the mean equals the variance) and
 - **Epistemic Uncertainty** (model parameter ambiguity, accounting for 1.15%, verifying high network convergence and weight stability).
- **Calibration Assessment:** Model calibration is validated via the correlation between predicted standard deviation and absolute prediction error, achieving a strong positive correlation ($r = 0.6101$) that confirms the model reliably signals when predictions carry higher risk.

Results and Discussion

To evaluate the predictive capacity of the proposed Spatio-Temporal Graph Neural Network (STGNN), we benchmarked its performance against three standard comparative baselines: a non-spatial temporal architecture (Standard LSTM), a pure spatial baseline (Standard GCN), and a gradient-boosted decision tree pipeline (XGBoost).

Comparative Experimental Results

The experimental evaluation metrics across all candidate architectures are summarised in Table 4.0. (Test set results).

Model	MSE	RMSE	MAE	R^2
Proposed STGNN	1.5339	1.2385	0.4694	0.6959
XGBoost	2.1762	1.4752	0.5583	0.5685
Standard GCN	5.2572	2.2928	0.9495	-0.0423
LSTM	4.2434	2.0599	0.8171	0.1587

Table 4.0 Experimental Results

The assessment demonstrates that the Proposed STGNN outperforms all baseline architectures across the evaluation results, validating the architectural necessity of jointly modeling spatial and temporal dimensions for epidemiological forecasting. It also aligns with prior findings that dynamic spatial graph messaging is essential to capture cross-regional transmission pathways that uncoupled sequence models miss (Panagopoulos et al., 2021).

- **Superior Predictive Accuracy and Limitations of Isolated Dimensions:** The Proposed STGNN achieves an R^2 score of 0.6959, an MSE of 1.5339, an RMSE of 1.2385, and an MAE of 0.4694. This represents a substantial performance improvement over non-deep learning tabular baselines and isolated structural models. The Standard GCN fails to capture predictive variance, yielding a negative R^2 score of -0.0423 (RMSE: 2.2928). This confirms that static spatial graph topologies alone cannot handle temporal evolution without sequential memory mechanisms. The LSTM model achieves an R^2 of 0.1587 (RMSE: 2.0599), showing that temporal sequence tracking captures some longitudinal patterns but misses critical cross-node spatial dependencies.
- **Benchmark Comparison:** While XGBoost provides a competitive non-deep learning baseline (R^2 : 0.5685, RMSE: 1.4752), it is basically constrained by flat feature representations, failing to match the relational propagation captured by the STGNN graph layers.




Fig 2.0 Evaluation Summaries

Explainability Metrics Interpretation

Feature attribution insights derived from Kernel SHAP and LIME link directly to the core feature mapping of the dataset, comprising the following tracked features:

- **Cumulative Cases (F_0):** Serves as the primary global driver of the model's output rate parameter (λ). Longitudinal case accumulation dictates immediate transmission trajectories across neighbouring spatial nodes.
- **Cumulative Deaths (F_1):** Functions as a secondary severity and impact weighting factor, anchoring the network's understanding of outbreak intensity.
- **Pathogen Group Taxonomy (F_2):** Acts as a structural categorical control, allowing the network to condition its baseline propagation rates to accommodate differing transmission characteristics across distinct pathogen classes.

While feature attribution analysis identify cumulative cases (F_0) as a primary global driver of the model's output rate parameter (λ), this reflects standard autoregressive behaviour where past realisations of the target variable condition future trajectories. Utilising lagged target values as inputs is a principle in temporal forecasting (Hyndman & Athanasopoulos, 2021), though future iterations of this framework could decouple historical target metrics from extrinsic indicators to further isolate external drivers of pathogen emergence.




Fig 3.0 SHAP Drivers

Uncertainty Quantification and Risk Calibration

To assess model reliability for operational deployment, predictive variance was decomposed into epistemic uncertainty (model weight ignorance) and aleatoric uncertainty (inherent observational noise) using 50 stochastic Monte Carlo Dropout passes under a heteroscedastic loss formulation (Kendall & Gal, 2017).

The uncertainty quantification metrics computed across $N = 944$ validation instances are summarised in Table 4.0.

Evaluation Parameter	Value	Structural Percentage / Interpretation
Total Evaluated Instances	944	Complete validation window
Actual Target Mean	0.8517	Baseline ground-truth mean
Predicted Target Mean	0.8381	Model predicted mean
Mean Absolute Error (MAE)	0.4715	Validation sample absolute deviation
Root Mean Squared Error (RMSE)	1.2516	Outlier-sensitive sample deviation
Mean Aleatoric Variance (σ^2_{alea})	0.8381	98.85% (Data and Observational Noise)
Mean Epistemic Variance (σ^2_{epis})	0.0098	1.15% (Model Parameter Ignorance)
Total Predictive Variance (σ^2_{total})	0.8479	100.00% (Combined Uncertainty Scale)
Mean Total Standard Deviation (σ_{total})	0.6788	Average predicted error bound
Error-Uncertainty Correlation (r)	0.6101	Calibration alignment metric

Table 4.0 Uncertainty Analysis Metrics

- Epistemic Convergence:** Model Ignorance is 1.15%, the epistemic variance across 50 Monte Carlo Dropout iterations is extremely small ($\sigma^2_{\text{epis}} = 0.0098$), confirming that performing stochastic forward passes can capture model weight ambiguity without adding significant computational overhead (Gal & Ghahramani, 2016). This low variance proves that the STGNN architecture has converged efficiently, maintaining

high structural parameter stability without suffering from model ambiguity or underfitting.

- **Aleatoric Dominance:** Data Noise is 98.85%, the majority of total uncertainty ($\sigma^2_{\text{alea}} = 0.8381$, representing 98.85%) exists from inherent data noise. This is mathematically expected in Poisson-distributed epidemiological count data, where event variance equals the mean.
- **Risk Calibration:** A positive correlation ($r = 0.6101$) between predicted standard deviation Higher predicted standard deviations genuinely signal higher operational risk and potential anomaly.

Operational Implications

Combining the empirical benchmarking and xAI attribution with the uncertainty quantification yields the following primary conclusions for spatial-temporal disease modeling:

- **Mandatory Joint Modeling:** Epidemiological phenomena cannot be accurately forecasted using isolated spatial or temporal components; fully integrated spatio-temporal architectures are required to capture complex outbreak dynamics.
- **Robust Convergence:** The minimal epistemic uncertainty (1.15%) confirms that residual prediction errors are driven by real-world stochastic shock events (aleatoric noise) rather than deficiencies in network training or parameter tuning.
- **Actionable Risk Mitigation:** Because predicted uncertainty correlates strongly with actual error ($r = 0.6101$), public health decision-makers can leverage the model's output standard deviation (σ) as a dependable early-warning indicator to flag high-risk spatial-temporal anomalies before severe outbreaks escalate.

Conclusion

The key findings and contributions of this work are listed below:

- **Superior Predictive Performance:** The proposed STGNN conclusively outperformed all baseline architectures, achieving an R^2 score of 0.6959 and an MAE of 0.4694. Comparative results confirmed that neither isolated temporal modeling (LSTM, $R^2 = 0.1587$) nor static spatial graph approaches (Standard GCN, $R^2 = -0.0423$) can independently capture the multi-dimensional nature of epidemiological data.
- **Transparent Feature Attribution:** Explainable AI analyses via Kernel SHAP and LIME validated that the model's predictions are anchored by logical clinical drivers, identifying cumulative cases (F_0), cumulative deaths (F_1), and the pathogen group taxonomy (F_2) as the primary determinants of outbreak trajectories.
- **Reliable Uncertainty Quantification:** The Monte Carlo-based variance decomposition proved that the network has achieved high stability, with epistemic model uncertainty accounting for only 1.15% of total variance. Furthermore, the strong positive correlation ($r = 0.6101$) between predicted standard deviation and absolute error establishes that the model accurately quantifies its own operational risk.

Regarding future research possibilities, currently, the spatial edges rely on static proximity. In reality, human movement changes constantly due factors such as public holidays, and local lockdowns. Introducing dynamic adjacency matrices that adapt to changing spatial connections over time would allow the graph structure to reflect real world human mobility much more accurately. The integration of demographic/socioeconomic indicators such as population density will also add crucial context regarding local vulnerability, helping the model distinguish between densely populated transit hubs and rural regions.

This study demonstrates the potential of Spatiotemporal Graph Neural Networks with xAI to advance One Health surveillance by providing explainable, and uncertainty-aware forecasts of pathogen emergence. While limitations related to data quality and model complexity remain, the framework offers a promising path toward proactive, data-driven public health decision-making. The integration of multi-relational dynamic graphs and explainability tools addresses key challenges in current outbreak forecasting, bringing us closer to a future where emerging health threats can be anticipated rather than merely reacted to.

References

- Ajao-Olarinoye, M., Palade, V., He, F., Wark, P. A., Mukandavire, Z., & Mousavi, S. (2025). A Hybrid Physics-Informed Neural Network: SEIRD Model for Forecasting COVID-19 Intensive Care Unit Demand in England. CRC Press <https://doi.org/10.1201/9781003570882-15>
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., & Herrera, F. (2020). Explainable artificial intelligence (xAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115.
- Attai, K., Amannejad, Y., Vahdat Pour, M., Obot, O., & Uzoka, F. M. (2022). A Systematic Review of Applications of Machine Learning and Other Soft Computing Techniques for the Diagnosis of Tropical Diseases. *Tropical medicine and infectious disease*, 7(12), 398. doi:10.3390/tropicalmed7120398
- Badker, R., Kipperman, N., Ash, B., Madhav, N.K., Oppenheim, B., Savage, P.T., Stephenson, N., & Pardee, C. (2025). Constructing a global human epidemic database using open-source digital biosurveillance. *Scientific Data*, 12.
- Bai, S., Kolter, J. Z., & Koltun, V. (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.
- Baker, R. E., Mahmud, A. S., Miller, I. F., Rajeev, M., Rasambainarivo, F., Rice, B. L. & Metcalf, C. J. E. (2022). Infectious disease in an era of global change. *Nature Reviews Microbiology*, 20(4), 193-205.
- Begum, M. & Ferdush, J. (2026) Opportunities and Challenges of Explainable AI (xAI) in HealthCare: A Review. *Journal of Computer and Communications*, 14, 71-86. doi: 10.4236/jcc.2026.146006
- Brownstein, J. S., Freifeld, C. C., & Madoff, L. C. (2009). Digital disease detection, harnessing the Web for public health surveillance. *New England Journal of Medicine*, 360(21), 2153-2157. doi: 10.1056/NEJMp0900702
- Cao, Y., Lancaster, E., Lee, J., & Wu, J. (2026). Integrating explainable AI and One Health: A new frontier in combating infectious diseases. *eBioMedicine*, 126, Article 106207. doi: 10.1016/j.ebiom.2026.106207
- Carlson, C. J., Albery, G. F., Merow, C., Trisos, C. H., Zipfel, C. M., Eskew, E. A., & Bansal, S. (2022). Climate change increases cross-species viral transmission risk. *Nature*, 607(7919), 555-562.
- Chen, X., Yin, C., Myers, J. V., Slover, B., Thomas, N., Marks, C., Kim, J., Fernández, S., Whitley, P., Fareed, N., & Zhang, P. (2026). Area-specific autoencoder spatiotemporal graph neural networks for opioid overdose death prediction. *JAMIA Open*, 9(2), Article ooag063. <https://doi.org/10.1093/jamiaopen/ooag063>
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Chretien, J. P., George, D., Shaman, J., Chitale, R. A., & McKenzie, F. E. (2014). Influenza forecasting in human populations: A scoping review. *PLoS ONE*, 9(4), e94130. doi: 10.1371/journal.pone.0094130
- Destoumieux-Garzón, D., Mavingui, P., Boetsch, G., Boissier, J., Darriet, F., Duboz, P., & Voituron, Y. (2018). The One Health concept: 10 years old and a long road ahead. *Frontiers in Veterinary Science*, 5, 14.
- Desvars-Larrive, A., Vogl, A. E., Puspitarani, G. A., Yang, L., Joachim, A., & Käsbohrer, A. (2024). A One Health framework for exploring zoonotic interactions demonstrated through a case study. *Nature Communications*, 15(1), Article 5650
- Fauci, A. S., & Morens, D. M. (2012). The perpetual challenge of infectious diseases. *New England Journal of Medicine*, 366(5), 454-461. doi: 10.1056/NEJMr1108296

- Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian Approximation: Representing Model Uncertainty In Deep Learning. International Conference on Machine Learning (pp. 1050–1059). PMLR.
- Gibbs, E. P. J. (2014). The evolution of One Health: A decade of progress and challenges for the future. *Veterinary Record*, 174(4), 85-91. doi: 10.1136/vr.g143
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*, The MIT Press, ISBN: 9780262035613,
- Han, S., Stelz, L., Sokolowski, T. R., Zhou, K., & Stöcker, H. (2025). Epidemiology-informed spatiotemporal graph neural network for heterogeneity-driven interpretable epidemic forecasting. *Engineering Applications of Artificial Intelligence*, 162(Part E), Article 112764. doi: 10.1016/j.engappai.2025.112764
- Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts.
- Johansson, M. A., Apfeldorf, K. M., Dobson, S. & Chretien, J. (2019). An open challenge to advance probabilistic forecasting for dengue epidemics. *Proceedings of the National Academy of Sciences*, 116(48), 24268-24274. doi: 10.1073/pnas.1909865116
- Jones, K. E., Patel, N. G., Levy, M. A., Storeygard, A., Balk, D., Gittleman, J. L., & Daszak, P. (2008). Global trends in emerging infectious diseases. *Nature*, 451(7181), 990–993. doi: 10.1038/nature06536
- Karesh, W. B., Dobson, A., Lloyd-Smith, J. O., Lubroth, J., Dixon, M. A., Bennett, M., Aldrich, S., Harrington, T., Formenty, P., Loh, E. H., Machalaba, C. C., Thomas, M. J., & Heymann, D. L. (2012). Ecology of zoonoses: Natural and unnatural histories. *The Lancet*, 380(9857), 1936–1945. doi: 10.1016/S0140-6736(12)61678-X
- Keeling, M. J., & Rohani, P. (2008). *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press.
- Kendall, A., & Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 5580–5590.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764 doi: 10.1016/j.ijforecast.2021.03.012
- Lipton, Z. C. (2016). The mythos of model interpretability. *Queue*, 16(3), 31-57.
- Liu, Z., Wan, G., Prakash, B. A., Lau, M. S. Y., & Jin, W. (2024). A review of graph neural networks in epidemic modeling. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, 6577–6587. doi: 10.1145/3637528.3671455
- Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
- Mackenzie, J. S., & Jeggo, M. (2019). The One Health approach. Why is it so important? *Tropical Medicine and Infectious Disease*, 4(2), 88.
- Morse, S. S., Mazet, J. A. K., Woolhouse, M., Parrish, C. R., Carroll, D., Karesh, W. B., & Daszak, P. (2012). Prediction and prevention of the next pandemic zoonosis. *The Lancet*, 380(9857), 1956-1965.
- Nsoesie, E. O., Brownstein, J. S., Ramakrishnan, N., & Marathe, M. V. (2013). A systematic review of studies on forecasting the dynamics of influenza outbreaks. *Influenza and Other Respiratory Viruses*, 8(3), 309-316. doi: 10.1111/irv.12226
- Panagopoulos, G., Nikolentzos, G., & Vazirgiannis, M. (2021). Transfer graph neural networks for pandemic forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6), 4838–4845. <https://doi.org/10.1609/aaai.v35i6.16616>

- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
- Reich, N. G., Brooks, L. C., Fox, S. J., Kandula, S., McGowan, C. J., Moore, E., Osthus, D., Ray, E. L., Tushar, A., Yamana, T. K., Biggerstaff, M., Johansson, M. A., Rosenfeld, R., & Shaman, J. (2019). A collaborative multiyear, multimodel assessment of seasonal influenza forecasting in the United States. *Proceedings of the National Academy of Sciences*, 116(8), 3146–3154, doi: 10.1073/pnas.1812594116
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144.
- Riley, S. (2007). Large-scale spatial-transmission models of infectious disease. *Science*, 316(5829), 1298-1301.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215.
- Scornet, E., Biau, G., & Vert, J. P. (2015). Consistency of random forests. *The Annals of Statistics*, 43(4), 1716-1741.
- Souza, G. N., Jr, Mendes, A. G. B., Costa, J. D. S., Oliveira, M. D. S., Lima, P. V. C., de Moraes, V. N., Silva, D. C. C., Rocha, J. E. C. D., Botelho, M. D. N., Araujo, F. A., Fernandes, R. D. S., Souza, D. L., & Braga, M. B. (2023). Deep learning framework for epidemiological forecasting: A study on COVID-19 cases and deaths in the Amazon state of Pará, Brazil. *PloS one*, 18(11), e0291138. doi: 10.1371/journal.pone.0291138
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (pp. 6000–6010). Curran Associates Inc.
- Viboud, C., Sun, K., Gaffey, R., Ajelli, M., Fumanelli, L., Merler, S., Zhang, Q., Chowell, G., Simonsen, L., & Vespignani, A. (2018). The RAPIDD ebola forecasting challenge: Synthesis and lessons learnt. *Epidemics*, 22, 13–21. doi: 10.1016/j.epidem.2017.08.002
- Wallinga, J., & Teunis, P. (2004). Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology*, 160(6), 509-516. doi: 10.1093/aje/kwh255
- World Health Organization. (2024). One Health: Adaptation guide on the One Health approach for the WHO European Region
- World Health Organization. (2025). COVID-19 Dashboard. WHO Data. <https://data.who.int/dashboards/covid19/summary>
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A Comprehensive Survey on Graph Neural Networks," in *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4-24, Jan. 2021, doi: 10.1109/TNNLS.2020.2978386
- Zheng, Y., Nguyen, Q. V. H., Jiang, W., Zhan, C., Zhou, A., & Chen, T. (2024). Epidemiology-informed graph neural network for heterogeneity-aware epidemic forecasting. *arXiv preprint arXiv:2411.17372*
- Zinsstag, J., Schelling, E., Waltner-Toews, D., & Tanner, M. (2011). From "one medicine" to "one health" and systemic approaches to health and well-being. *Preventive Veterinary Medicine*, 101(3-4), 148-156.