

Deep Learning Models for Predicting Avian Zoonotic Disease Transmission in Poultry-Dominant Mixed Farms

P. NARESH^{1*}, ²MOHAMMED ZAHEER AHMED, M.C. BHANU PRASAD³, P.V. GOPALKRISHNA⁴, B. PRASANTHI⁵, SIVANANDA LAHARI REDDY ELICHERLA⁶

¹Department of Computer Science and Engineering, Dayananda Sagar University, Bengaluru, India; ²Department of Computer Science and Engineering, Vidya Jyothi Institute of Technology, Hyderabad, India; ³Department of CSE, Tadipatri Engineering College, Tadipatri, India; ⁴Department of Mechanical Engineering, Vasavi College of Engineering, Hyderabad, India; ⁵Department of CSE (AI & ML), CVR College of Engineering, Vastunagar, Mangalpalli (V), Ibrahimpatnam (M), Rangareddy District, Telangana, 501510, India; ⁶Department of CSE, CMR University, Bangalore, India.

Abstract | Avian diseases that can spread from birds to humans, especially the highly pathogenic avian influenza (HPAI) subtypes H5N1 and H7N9, continue to threaten poultry farming worldwide and put public health at risk. On farms where chickens are raised alongside pigs, goats, cattle, or other livestock, there's a higher chance of disease spreading because these animals share water, feed, and areas where people work. Traditional ways of tracking these diseases depend on reporting when outbreaks happen and using simple risk scores, which do not capture temporal dynamics adequately and can't show how the disease moves between farms. This paper introduces a new deep learning method that uses both time-based analysis and a layer that considers how animals interact with each other to predict the likelihood of disease spreading between farms over a period of 7 to 14 days. Using a large set of data from 800 mixed farms in a 200 by 200 km area, with daily updates on various factors, the new model called Temp-Contact Net (TCNet) performs better than other methods like XGBoost and standalone LSTM. It has an AUC-ROC of 0.89 and an F1-score of 0.76, which is more than the alternatives. Tests also show that including information about animal interactions improves the model's ability to detect outbreaks, increasing recall by 12% (from 0.52 to 0.82).

Keywords | Avian influenza, Zoonotic transmission, Mixed farms, Deep learning, Graph neural networks, Risk prediction

Received | April 14, 2026; **Accepted** | May 26, 2026; **Published** | June 02, 2026

***Correspondence** | P. Naresh, Department of Computer Science and Engineering, Dayananda Sagar University, Bengaluru, India; **Email:** nareshintell4@gmail.com

Citation | Naresh P, Ahmed MZ, Prasad MCB, Gopalkrishna PV, Prasanthi B, Elischerla SLR (2026). Deep learning models for predicting avian zoonotic disease transmission in poultry-dominant mixed farms. *Adv. Anim. Vet. Sci.*, 14(6):1179-1189.

DOI | <https://dx.doi.org/10.17582/journal.aavs/2026/14.6.1179.1189>

ISSN (Online) | 2307-8316



Copyright: 2026 by the authors. Licensee ResearchersLinks Ltd, England, UK.

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

INTRODUCTION

Zoonotic diseases in birds that can be passed over to humans are a significant concern at the triple intersection of animal health, population health, and economic agriculture. H5N1 and H7N9 subtypes of

avian influenza are highly pathogenic and have resulted in hundreds of cases in humans with case fatality rates over 30% (Webster *et al.*, 1992). Since 2003, the World Health Organization has reported more than 860 human infections with H5N1, approximately 53% of which were fatal. The HPAI outbreak in the United States in 2014-

2015 alone caused the depopulation of around 50 million birds and the estimated losses to the economy were over 3 billion (Roger *et al.*, 2021).

Mixed farms predominantly poultry, such as those found in many parts of Asia, sub-Saharan Africa and some of Latin America, are those that have poultry together with pigs, goats, sheep or cattle (Gilbert *et al.*, 2014). Such integration forms numerous channels of transmission. Pigs are mixing vessels of influenza A viruses and facilitate reassortment that has the potential to produce new influenza strains with pandemic potential. In the synthetic dataset, poultry-pig mixed farms (50% of farms) were modelled with shared housing areas, water sources, and workers between species in 70% of cases (based on typical Southeast Asian mixed-farm practices), explicitly incorporating the 'mixing vessel' effect through elevated reassortment probability in the agent-based transmission dynamics when pigs and poultry shared confined spaces or fomite. Traditional disease surveillance is based on passive clinical reporting, active targeted sampling, laboratory confirmation using rRT-PCR and international notification using WOAHA (World Organisation for Animal Health, 2022). Each stage introduces delays of 24–72 hours. Rule-based risk indices are a blend of weather anomalies and waterfowl patterns using expert-weighted heuristics (Gubbins, 2018). These methods are not able to learn new data or learn nonlinear interactions.

Outbreak prediction with machine learning on XGBoost has been demonstrated to achieve AUC 0.82, but at large spatial scales (district or village level) (Gilbert *et al.*, 2014). Deep learning provides LSTM which depends on time (Hochreiter *et al.*, 1997) and GCNs which depend on relation structure [10]. In this paper, TCNet has been introduced, which fills four gaps: (i) transmission prediction instead of diagnosis over a 7–14 day time period, (ii) daily resolution at the farm level including 12 features, (iii) contact network modelling with distance, market, and veterinary links and (iv) ablation studies that quantitatively measure the contribution of each component. The paper is arranged in the following way. Section II is a review of related work. The problem is defined and the system overview is given in Section III. Section IV explains the generation of synthetic datasets. TCNet architecture is given in section V. Section VI is a report on experimental results. Section VII discusses ablation and interpretability. Section VIII concludes.

RELATED WORK

DEEP LEARNING FOR POULTRY HEALTH MONITORING

Deep learning has become widely adopted of poultry health in the last five years. This growth is because the cost of sensors has gone down and there's more computing power available. A review of 78 studies found that 62%

used CNNs vision-based tasks. Some common uses include recognizing coccidiosis by looking at fecal images, which can be done with 94% accuracy using ResNet50. Also, it can detect respiratory issues by looking at the facial expressions of chickens and classify droppings as healthy or sick. These systems work well in controlled lighting conditions but perform worse on commercial farms (Simonyan *et al.*, 2015). Table 1 Represents research gaps in existing methods.

Table 1: Summary of research gaps and contributions.

Gap	Existing limitation	Our contribution
Gap 1	Most DL models are diagnostic, not predictive [4]	7-14 day transmission prediction
Gap 2	Coarse spatial/temporal resolution [3][5][12]	Farm-level daily resolution
Gap 3	Graph methods not tested on avian zoonoses in mixed farms [6][7]	LSTM-GCN for mixed farms
Gap 4	No systematic ablation of contact-network contribution	Comprehensive ablation studies

3: Ferguson *et al.* (2006); 4: Gilbert *et al.* (2014); 5: Gubbins (2018); 6: Hochreiter and Schmidhuber (1997); 7: Kapoor *et al.* (2020); 12: Sun *et al.* (2021)

Using audio has become another way to detect illness. Sick chickens make different sounds that can be seen as patterns in their sound waves. Using recurrent neural networks and 1D CNNs with MFCCs, they can detect diseases like infectious bronchitis and avian influenza with a recall rate over 85%. However, real-world accuracy drops by 15–20% because of background noise. Using data from body postures and feeding it to LSTM classifiers can tell if a chicken is healthy or has symptoms before they start, with AUC values of 0.87 to 0.91. The crucial problem with all these deep learning methods is that they assume the disease is already present and just predict the health status, not forecast future spread or consider how farms interact with each other.

While audio-based detection methods show promising recall (>85%) in controlled settings, their performance drops 15–20% on commercial farms due to background noise. Similarly, TCNet although evaluated on synthetic data with ideal contact structures would require robustness testing against real-world issues such as under-reported biosecurity scores or delayed clinical reporting. Future work will incorporate noisy or missing data augmentation to simulate these challenges.

MACHINE LEARNING FOR ZOOONOTIC DISEASE RISK PREDICTION

When it comes to predicting outbreaks, there has been less research (Gilbert *et al.*, 2014) used XGBoost to predict the risk of H7N9 infections in poultry markets in China and

Southeast Asia. But they focused on whether any infection was present in a market over a long period, not how likely a farm is to spread the disease day by day. The spatial resolution was on a grid that covered large areas with many farms (Vaswani *et al.*, 2017).

In Egypt and Bangladesh, random forests have been used. (Abdelwhab and Hafez, 2011) found that being close to wild birds was the strongest factor for H5N1 transmission in backyard chickens. Using logistic regression with spatial data slightly improved predictions (Vaswani *et al.*, 2017) but still had issues with linear models. Also, the time data was too far apart to help with early warnings. Plus, they didn't consider complex relationships beyond simple distance.

GRAPH NEURAL NETWORKS AND CONTACT-AWARE EPIDEMIC FORECASTING

Graph neural networks have changed how we work with data that has relationships. Kipf and Welling (Kipf *et al.*, 2017) created graph convolutional networks (GCNs) that work on graph-based data. For a graph with an adjacency matrix A and node features H , a single GCN layer calculates $H' = \sigma(A\tilde{W}H)$, where $\tilde{A}$ is a normalized version of the adjacency matrix and W is something the model learns.

In predicting infectious diseases, (Kapoor *et al.*, 2020) used a temporal GCN to forecast COVID-19 cases by looking at how people moved. Compared to temporal GCN approaches for COVID-19 forecasting (Kapoor *et al.*, 2020), TCNet adds farm-level daily resolution, explicit multi-species contact modelling (distance + market + veterinary), a focused 7–14-day zoonotic transmission horizon relevant to avian influenza, and quantitative ablation studies isolating the contribution of the contact-aware layer in a mixed-farming context. They found that considering spatial connections helped more than just distance. In veterinary epidemiology, (Sun *et al.*, 2021) made a GCN model (Szegedy *et al.*, 2015) for foot-and-mouth disease in South Korea, achieving an AUC-ROC of 0.84 and outperforming logistic regression by 11 percentage points.

PROBLEM DEFINITION AND SYSTEM OVERVIEW

PROBLEM FORMULATION

Let $F = \{1, 2, \dots, N\}$ be the set of mixed poultry-dominant farms in a geographical area. In our synthetic dataset, we will define the value of N as 800 to be a medium-density poultry producing area. The farms in F have a farm type: pure poultry (30 percent), poultry-dominant mixed with either pigs (50 percent), or poultry-dominant mixed with ruminants, goats or cattle (20 percent). The proportions are representative of the common distributions found in agricultural census records in Southeast Asia and sub-Saharan Africa.

At each time step of the discrete time, denoted by t where t is one calendar day (the time step here is one calendar day), we have a feature vector, denoted by $x_i(t)$ each with dimension 12 after feature engineering. The variables in the farm i of time t are as follows:

- (1) Poultry density (birds/square meter of housing area), log-transformed to eliminate skew.
- (2) Mixed species proportion, which is non-poultry livestock units per total livestock units, between 0 (pure poultry) and 1 (poultry and other livestock equal).
- (3) Biosecurity score, a composite index between 0 and 100 which is based on farm management practices.
- (4) Proximity to closest wildlife wetland (kilometers), waterfowl-mediated introduction risk.
- (5) Road traffic index (0 to 10), which is a proxy of human and vehicle movement.
- (6) Temperature in degrees Celsius, from daily weather simulation.
- (7) Relative humidity as a percentage, from daily weather simulation.
- (8) 7-day clinical suspicion flag (previously), a binary variable that denotes respiratory distress or mortality.
- (9) Day-of-year sine transformation, seasonal effects.
- (10) Day-of-year cosine transformation combined with sine to represent circular time.
- (11) Proxy human movement, number of market visits the past 7 days (farm worker).
- (12) Neighboring farm historical incidence, a spatial lag variable as a weighted average of clinical suspicion flags on farms within 2 km.

Training set statistics are make use of standardize all continuous features to zero mean and unit variance. Categorical variables are coded into binary indicator variables.

In addition to the feature vectors at the farm level, we also define a contact matrix, represented by $C(t)$, where each entry $C_{ij}(t)$ is a value between 0 and 1. This value shows how much epidemiological contact exists between farm j and farm i at time t . The contact matrix is treated as a fixed element because the main factors that determine contact like geographic distance, shared market participation, or using the same veterinarian change very slowly over time. The strength of contact is calculated as follows:

$$C_{ij} = w_{\text{dist}} \exp(-d_{ij}/\sigma) + w_{\text{market}} M_{ij} + w_{\text{vet}} V_{ij}$$

and where d_{ij} is Euclidean distance in kilometers, $\sigma = 5$ km, M_{ij} is shared live bird market, V_{ij} is shared veterinary provider, $w_{\text{dist}} = 0.5$, $w_{\text{market}} = 0.3$, $w_{\text{vet}} = 0.2$. Self-edges are set to zero: $C_{ii} = 0$ for all i .

Veterinarian visits and binary V_{ij} : The binary veterinary indicator $V_{ij} = 1$ if farms share the same provider represents a conservative proxy for shared risk. In reality, transmission risk varies with visit frequency, timing, and

hygiene practices (e.g., boot washing). This simplification was chosen for computational tractability and because detailed visit logs are rarely available in real surveillance data. A frequency- or time-decayed weight could be explored in extensions with richer data.

The exponential decay parameter $\sigma = 5$ km was chosen based on empirical observations from HPAI outbreaks, where most inter-farm transmission occurs within 5–10 km via local movements, fomites, or short-distance airborne spread. Sensitivity analyses with $\sigma = 3$ km and 10 km yielded similar overall performance but slightly lower recall at longer horizons.

Feature 12 (neighbouring farm historical incidence within 2 km) captures local spatial lag at the feature level, while the graph convolution aggregates learned representations across the broader contact network (including market and veterinary links). Ablation studies (Section 7) confirmed that removing either component reduces performance, indicating complementary rather than redundant information.

The sparsification of the dense contact matrix is done to achieve computational efficiency. We set threshold $\theta = 0.1$ and define binary adjacency matrix $A_{ij} = 1$ if $C_{ij} > \theta$, otherwise $A_{ij} = 0$. The neighborhood of each farm, denoted by the number of strongest connections with other farms, is restricted to the top $k = 5$ connections, thus the computational cost is linear in the number of farms, N .

The prediction target is defined as follows. Suppose that, given a farm i on a day t the ground-truth label that there is at least one new confirmed avian zoonotic case on farm i in the prediction window $(t + 1, t + \tau) = y_i(t)$. There are two horizons that are considered: $\tau = 7$ days (primary) and $\tau = 14$ days (sensitivity analysis). A case is said to be confirmed when the farm has reported clinical signs and a lab test has been positive. The agent-based model simulation is known exactly as ground truth in the synthetic dataset.

The prediction problem is: with known historical feature observations of all farms up to time t , denoted by the observations of all feature vectors of all farms, denoted as: $\{x_i(s) | i \in F, s \leq t\}$, and known contact adjacency matrix, denoted by: A , predict probability: $\hat{y}_i(t)$ in $[0, 1]$, predicting $\hat{y}_i(t) = 1$. The output of the model is a probability that can be used by the downstream users to make decisions depending on risk tolerance.

SYSTEM OVERVIEW

Figure 1 contains the block diagram of the proposed TCNet system that has four consecutive modules: Data Acquisition, Preprocessing, TCNet Model and Output Visualization. Data Acquisition module takes in daily farm records,

weather data, satellite wetlands maps and road network data, whereas Preprocessing module imputes, normalizes Data Acquisition module, forms adjacency matrix A and forms 7-day sliding windows. Primary architecture The TCNet Model uses LSTM temporal encoding, graph convolution to contact-aggregate data, and a feedforward classifier to predict risk, and then the Output Visualization module to produce farm-level risk scores, spatial heatmaps, and probability matrices of transmission. With this pipeline, it is possible to predict the probabilities of transmission, every day, within a 7–14 horizon, and across all 800 farms in the network. Traditional rule-based risk indices, such as the multi-criteria decision analysis (MCDA) frameworks used in Cambodia for H5N1 risk mapping (combining expert-elicited weights for factors like proximity to water bodies, poultry density, and human movement), have shown moderate alignment with observed cases but often overestimate or underestimate risk due to static weighting and limited temporal resolution.




Figure 1: Above diagram of the proposed TCNet system. The system is constituted of four consecutive modules, which are Data Acquisition, Preprocessing and Feature Engineering, TCNet Model and Visualization of the outcome. Arrows show flow of data left to right.

Deep Learning Prediction is the third module, which uses the TCNet architecture (Section [sec:tcnet]). For each farm i , the LSTM encoder processes the 7-day feature sequence to produce temporal hidden state $h_i^{\text{temp}}(t)$. A graph convolutional layer aggregates neighbors' states via A , yielding $h_i^{\text{graph}}(t)$. The concatenated $[h_i^{\text{temp}}(t); h_i^{\text{graph}}(t)]$ feeds a dropout-regularized feedforward classifier for risk probability $\hat{y}_i(t)$. Training uses weighted binary cross-entropy for imbalance, with optional spatial smoothness penalty for coherent nearby predictions.

The fourth module, Output Visualization, produces a farm-level risk score (ranked list or color-coded map) for prioritizing high-risk sites, plus a spatial risk heatmap for regional surveillance. Optionally, it outputs farm-to-farm transmission matrix $\hat{P}_{ij}(t)$, estimating infection spread from j to i (requires transmission pair labels, available in synthetic data). All daily outputs integrate into veterinary dashboards or apps.

DATASET AND EXPERIMENTAL SETUP

SYNTHETIC DATA GENERATION FRAMEWORK

To enable rigorous controlled evaluation of TCNet, we constructed a synthetic dataset using an agent-based modeling (ABM) approach calibrated to HPAI H5N1 and H7N9 outbreak parameters. Five high-density clusters (each containing ~100 farms) account for 500 farms, with the remaining 300 distributed as background farms to replicate this pattern. This spatial scale is able to capture meaningful heterogeneity in the farm density, patterns of contact, and environmental conditions and is computationally tractable to repeatedly simulate.

Each high-density zone cluster radius will be 10 km or so, that is, the distance between the farms in a cluster is usually 5-10 km. In case of clustered farms, coordinates are generated as a result of a bivariate normal distribution around the centroid of clusters with a standard deviation of 3 km which gives the realistic tapering of the density around the cluster centroid. In the case of background farms, the coordinates are randomly generated all over the region.

Positive samples clarification: The dataset contains 584,000 farm-day samples across 800 farms × 730 days. There are 25,112 positive labels for the 7-day horizon (approximately 4.3% of samples), reflecting the rarity of outbreaks. Recall of 0.82 on the test set therefore corresponds to correctly identifying ~20,590 of these events, which is meaningful for early warning despite class imbalance.

Biosecurity equation typo fix (rewrite the force of infection part): The force of infection $\lambda_i(t)$ is given by: $\lambda_i(t) = \beta_{\text{local}} \times (1 - \text{biosecurity}_i) \times I_{\text{local}}(t) + \beta_{\text{contact}} \times A_{ij} \times I_j(t) + \beta_{\text{env}} \times \text{env}(t)$, where $\text{biosecurity}_i \in [0,1]$ reduces susceptibility (higher score lowers transmission probability). The original expression contained a typographical error; biosecurity acts as a protective multiplier, consistent with reduced transmission on well-managed farms.

Missing values at 2%: Missingness was artificially introduced at 2% to simulate minor real-world gaps while preserving data quality. In practice, real farm records often exhibit 20–30% missingness for variables like market visits; future versions will test higher missingness rates and imputation strategies (e.g., forward-fill or GNN-based).

The baseline values of biosecurity scores are based on a Beta distribution where the shape parameters are 4 and 2, respectively. This distribution has mean $\alpha/(\alpha+\beta) = 4/6 \approx 0.667$, mode $((\alpha-1)/(\alpha+\beta-2)) = 3/4 = 0.75$, and range approximately 0.30 to 0.95. The biosecurity of most farms is moderately good (0.6 to 0.8) with a tail of low-biosecurity farms (below 0.4), which is seasonal changes

where outbreaks are most likely to start. Biosecurity ratings are revised on a monthly basis to capture changes in the season in terms of management practices.




Figure 2: Spatial distribution of 800 simulated farms across a 200*200 km region. Red points indicate pure poultry farms (30%). Blue points indicate poultry-pig mixed farms (50%). Green points indicate poultry-ruminant mixed farms (20%). Five high-density clusters are visible (dense groups of points) with uniform background distribution.

Figure 2 Represents the Spatial Distribution of 800 simulated farm.

Poultry density is assigned based on farm type and cluster membership. Pure poultry farms in clusters have highest densities, drawn from log-normal distribution with mean $\mu=2.5$ on log scale (approximately 12 birds per square meter). Mixed farms have lower densities (means 1.8 for poultry-pig, 2.0 for poultry-ruminant). Background farms have lowest densities (mean 1.2). The mixed species proportion ranges from 0.2 to 0.6 for poultry-pig farms (mean 0.35) and 0.2 to 0.7 for poultry-ruminant farms (mean 0.40). Road traffic indices are assigned based on proximity to 15 major roads. Farms within 0.5 km of major road receive indices 5-10, farms 0.5-2 km receive 2-5, farms beyond 2 km receive 0-2.

Daily weather follows seasonal sine waves with added Gaussian noise. Temperature: $T(t) = 15 + 10 \sin(2\pi t / 365 - \pi/2) + \epsilon_T$, $\epsilon_T \sim N(0,2)$. Humidity: $H(t) = 65 + 15 \sin(2\pi t / 365 + \pi/4) + \epsilon_H$, $\epsilon_H \sim N(0,5)$. These measurements are an approximation of tropical and subtropical climates that have cool dry winters (January–February) and warm humid summers (July–August).

TRANSMISSION DYNAMICS MODEL

The dynamics of disease transmission is the modified Susceptible-Exposed-Infectious-Recovered (SEIR) model on the farm level. Active shedding of the virus takes place in infectious (I) farms and can be transmitted to other

farms. Recovered (R) farms are cured and cannot become infected again in the rest of the simulation horizon.

The force of infection for farm i on day t is given by:

$$\lambda_i(t) = \beta_{\text{local}} (1 - \text{biosecurity}_i) \cdot \text{env}_i(t) + \beta_{\text{contact}} \sum_{j \neq i} A_{ij} \cdot I_j(t)$$

with local transmission rate, contact transmission rate and biosecurity of the farms, respectively, where “biosecurity” _{i} = $-0.12 + 0.08 I_j(t) + 0.06 A_{ij}$. The environmental term [“env”] _{i} (t) is defined as:
[“env”] _{i} (t) = $1 + 0.5 \cdot 1_{(T(t) < 10^\circ \text{C})} + 0.3 \cdot 1_{(H(t) > 80\%)}$

This formula enhances transmission in cold environments (less than 10 °C) by half and in moist environments (more than 80 percent relative humidity) by a third, which is in line with previously known environmental persistence of avian influenza viruses [8]. The contact transmission term reflects the transmission between adjacent infected farms through common markets, veterinary care providers, or proximity.

Upon exposure, each farm will undergo a latent period of 2 days (drawn using Poisson distribution with mean 2) then infectious period of 5 days (drawn using Poisson distribution with mean 5). Once the infectious period has ended, the farm will be recovered and cannot be reinfected within the simulation horizon. This parameterization is based on published estimates for HPAI H5N1 in poultry [8]. The likelihood of an exposed farm on a particular day is a probability, $P(\text{exposed}) = 1 - \exp(-\lambda_i(t))$. This is about the same as the small value of the $\lambda_i(t)$ at the time, i.e. $\lambda_i(t)$. The simulation is deterministic with the exception of the stochasticity of the latent and infectious period. The resulting distribution of the classes is very skewed. Out of all 584,000 samples of farm days (800 farms x 730 days), there are only 4.3 positive labels (transmission within next 7 days). This imbalance is typical of the reality of the rarity of outbreaks and is a challenge to standard classification algorithms.

CONTACT MATRIX CONSTRUCTION

Contact matrix is very important part of our modeling scheme as it captures the epidemiological relationship between farms which facilitates transmission of diseases beyond local environmental transmission. In practical mixed agriculture, interfarm transmission can take place in several ways: through common people (farm workers, veterinarians), common equipment and vehicles, animal movements (live bird markets, slaughterhouses), and airborne or fomite transmission by proximity. The parameterization of these various pathways in one contact matrix has to be done with careful consideration depending on published epidemiological work.

We use a 3-component weighted method to construct our contact matrix. Shared veterinary provider is 20% of the weight and represents on the transmission by veterinarians or animal health workers who visit multiple farms, and who may carry pathogens on boots, clothing or equipment. The distance weight is computed by an exponential fall: $\exp(-d_{ij}/5)$ i.e. the Euclidean distance among the farms in kilometers. This operation assigns a big weight to the farms within 5 km ($\exp(-1) = 0.368$) and insignificant weight to those further than 15 km ($= 0.050$). The scale of characteristic 5 km is in line with the research of the HPAI spread in the area, where most of the transmission is usually observed near farms with the infection within 5-10 km. The farms are allotted to the closest provider. The veterinary indicator $V_{ij} = 1$ in case there is the same provider in two farms. This can be attributed to the transmission through veterinarians who visit farms on a number of occasions.

The weight of the raw contact is determined as:

$$C_{ij} = 0.5 \cdot \exp(-d_{ij}/5) + 0.3 \cdot M_{ij} + 0.2 \cdot V_{ij}$$




Figure 3: A subset of 50 farms in contact network visualization. Nodes indicate farms based on color red (pure poultry), blue (poultry-pig mixed), and green (poultry-ruminant mixed). Contact strength is reflected by edge thickness. There is greater density in the network within spatial clusters and less density in cross-cluster links.

Figure 3 Represents the visualization graph of subset of 50 farms

The adjacency matrix A_{ij} is then binarized with a threshold of 0.1, that is, any pair whose sum of weights to be above 0.1 is counted as connected. This cut-off maintains only associations with epidemiological significance. A contact weight of less than 0.1 could be due to, say, two farms at a distance of 12 km (distance weight of about 0.09) but do not have any market or veterinary provider; such a relationship is probably too small to be useful within the 7-day prediction horizon. Lastly, a neighborhood of each farm is restricted to the strongest $k=5$ connections to keep the graph sparse. This is done to make the computation

time bound, a dense adjacency matrix would take $O(N^2)$ operations ($N=800$, $N^2=640,000$) whereas a sparse matrix with average degree 4.2 takes $O(E)$ operations. The resulting graph contains 1,678 edges, mean degree 4.2 and a higher density in the spatial clusters and lower density in the cross-cluster connections.

FEATURE ENGINEERING AND PREPROCESSING

From the raw simulation outputs, we engineer the 12 feature variables listed in Section III-A. Temperature and humidity are used as continuous variables. The previous 7-day clinical suspicion flag is computed by checking whether any clinical signs occurred on the farm in days $t-6$ through t . Day-of-year sine and cosine transformations are computed as $\sin(2\pi \cdot \text{"day"} / 365)$ and $\cos(2\pi \cdot \text{"day"} / 365)$ to encode seasonality as a circular variable.

The human movement proxy is simulated as the number of market visits by farm workers in the past 7 days. For pure poultry farms, mean visits is 2 per week; for poultry-pig mixed farms, mean is 3 per week; for poultry-ruminant mixed farms, mean is 2.5 per week. Missing values are rare in synthetic data (less than 0.1%), but we artificially introduce missingness at random for 2% of feature values to mimic real-world conditions. All continuous features are standardized using training set mean and standard deviation. Standardization parameters are saved and applied to validation and test sets to prevent data leakage.

DATA SPLITTING STRATEGY

The time period is divided into parts in order to avoid look-ahead bias. The first 511 days, which make up 70. These temporal split respects the chronological order to avoid look-ahead bias in which predictions are made. For each day t , starting from day 7 up to day 723, we create a sample that includes an input sequence from days $t-L+1$ to t and a label indicating if transmission happens between $t+1$ and $t+\tau$. With 800 farms, 730 days in total, a sequence length of 7, and a time horizon of 7, this results in $800 \times 717 = 573,600$ samples.

EVALUATION METRICS

the performance of the model is measured using six different metrics. Accuracy is the percentage of correct predictions, but it can be misleading when the data is not balanced. Precision shows the percentage of positive predictions that are actually correct: True positives divided by (true positives plus false positives). Recall, also known as sensitivity, shows the percentage of the actual positive cases that are correctly identified: true positives divided by (true positives plus false negatives). The F1-score is the average of precision and recall, calculated as two times (precision multiplied by recall) divided by (precision plus recall).

The area under the receiver operating characteristic curve (AUC-ROC) is used to evaluate how well the model can distinguish between different outcomes across all possible thresholds. To check if differences in AUC-ROC are statistically significant, DeLong's test is used for correlated ROC curves, and a p-value less than 0.05 is considered significant.

MODEL ARCHITECTURE AND METHODOLOGY

OVERALL ARCHITECTURE

TCNet comprises three main components: The first component is a two-layer stacked LSTM that processes a 7-day sequence of feature vectors for each farm independently. This encoder produces a temporal hidden state $h_i^{\text{temp}}(t)$ that captures recent disease-relevant dynamics on the farm.

The second component is a graph convolutional layer that operates on temporal hidden states of all farms simultaneously. Using the precomputed binary adjacency matrix A , this layer aggregates information from neighboring farms, allowing a farm's risk prediction to be influenced by recent status of epidemiologically connected farms.

The third component is a feedforward classification head that takes concatenated representation $[h_i^{\text{temp}}(t); h_i^{\text{graph}}(t)]$ and outputs probability $\hat{y}_i(t)$. Dropout is applied before output layer to reduce overfitting.




Figure 4: TCNet architecture diagram. (A) Input: 7-day sequence of feature vectors (12 features per day). (B) Temporal encoding: Two-layer LSTM with layer normalization and dropout produces temporal hidden state $h_i^{\text{temp}}(t) \in \mathbb{R}^{64}$. (C) Contact-aware aggregation: Graph convolutional layer with normalized adjacency matrix $\tilde{A}$ aggregates neighbor states to produce $h_i^{\text{graph}}(t) \in \mathbb{R}^{32}$. (D) Concatenation and classification: Combined representation passes through MLP with dropout and sigmoid to produce risk probability $\hat{y}_i(t)$.

Figure 4 Represents the TCNet Architecture.

INPUT REPRESENTATION

For each farm i at time t , the input is a 7-day sequence:

$$X_i(t) = [x_i(t-6), x_i(t-5), \dots, x_i(t)]$$

Where; each $x_i(s) \in \mathbb{R}^{12}$

TEMPORAL ENCODING WITH LSTM

The LSTM recurrence at time step s within the lookback window is:

$$\begin{aligned} f_{s\&} &= \sigma(W_f \cdot [h_{s-1}, x_s] + b_f) @ i_{s\&} = \sigma(W_i \cdot [h_{s-1}, x_s] + b_i) @ c_{s\&} \\ &= \tanh(W_c \cdot [h_{s-1}, x_s] + b_c) @ c_{s\&} \\ &= f_s \odot c_{s-1} + i_s \odot c_{s\&} @ o_{s\&} \\ &= \sigma(W_o \cdot [h_{s-1}, x_s] + b_o) @ h_{s\&} = o_s \odot \tanh(c_{s\&}) \end{aligned}$$

Where; σ is sigmoid, $\odot$ is element-wise multiplication. Hidden dimension $h_{\text{lstm}} = 64$ for both layers. After processing all 7 steps, final hidden state $h_i^{\text{temp}}(t) \in \mathbb{R}^{64}$.

CONTACT-AWARE AGGREGATION

All farms' temporal states are stacked into $H^{\text{temp}}(t) \in \mathbb{R}^{(800 \times 64)}$. Normalized adjacency matrix: $A \sim D^{(-1/2)} A D^{(-1/2)}$. Graph convolution:

$$H^{\text{graph}}(t) = \text{"ReLU"}(A^{\text{H}^{\text{temp}}}(t) W_g + b_g) \text{ where } W_g \in \mathbb{R}^{(64 \times 32)}, H^{\text{graph}}(t) \in \mathbb{R}^{(800 \times 32)}$$

CLASSIFICATION HEAD

Concatenated representation: $z_i(t) = [h_i^{\text{temp}}(t); h_i^{\text{graph}}(t)] \in \mathbb{R}^{96}$.

First MLP layer: $z_i^{(1)}(t) = \text{"ReLU"}(W_1 z_i(t) + b_1)$, $W_1 \in \mathbb{R}^{(64 \times 96)}$. Dropout 0.3 applied.

Second MLP layer (sigmoid output): $y_i(t) = \sigma(w_2^T z_i^{(1)}(t) + b_2)$, where $\sigma(z) = 1/(1 + e^{-z})$.

LOSS FUNCTION

Weighted binary cross-entropy for each sample:

$$L_i(t) = -[w \cdot y_i(t) \cdot \log(y_i(t)) + (1 - y_i(t)) \cdot \log(1 - y_i(t))]$$

Where; $w = \text{"#negative"} / \text{"#positive"} \approx 22.3$

Spatial smoothness penalty:

$$L_{\text{smooth}} = \lambda_{\text{smooth}} \sum_{((i,j): d_{ij} < 2 \text{ "km"})} (y_i - y_j)^2, \quad \lambda_{\text{smooth}} = 0.01$$

$$\text{Total loss: } L_{\text{total}} = L + L_{\text{smooth}}$$

Table 2 represents TCNet hyper-parameters of various Layers

BASLINE MODELS

Four baseline models are compared against TCNet. LSTM-only uses the same two-layer LSTM as TCNet but no graph convolution. GCN-only uses graph convolution on single-day features with no temporal modeling.

Table 2: TCNet hyperparameters.

Parameter	Value
LSTM layers	2
LSTM hidden dimension	64
LSTM dropout	0.2
Graph output dimension	32
MLP hidden dimension	64
MLP dropout	0.3
Lookback window	7 days
Prediction horizon	7 days (primary), 14 days (sensitivity)
Optimizer	Adam
Initial learning rate	0.001
Learning rate decay	0.95 every 10 epochs
Batch size	256
Max epochs	80
Early stopping patience	10 epochs
Positive class weight	22.3
Smoothness coefficient	0.01

Table 3 Represents research gaps in existing methods.

Table 3: Baseline models.

Model	Temporal	Contact	Description
Logistic regression	No	No	Single-day features only
XGBoost	Rolling means	No	7-day rolling averages
LSTM-only	Yes (sequence)	No	LSTM without graph
GCN-only	No	Yes	Graph on single day
TCNet	Yes	Yes	LSTM + GCN hybrid

LSTM variants: We evaluated GRUs and bidirectional LSTMs during preliminary experiments. The unidirectional two-layer LSTM with 128 hidden units provided the best trade-off between performance and computational cost on the temporal sequences. Bidirectional variants showed marginal gains but increased training time without improving generalization on the forward-prediction task.

Spatial smoothness penalty: The spatial smoothness penalty is a Laplacian regularization term: $L_{\text{smooth}} = \lambda \sum_i \sum_{j \in N(i)} \|h_i - h_j\|^2$, where h denotes the temporal hidden states and $N(i)$ are neighbors in A . It was included with $\lambda = 0.01$ in all reported experiments to encourage coherent predictions among connected farms.

Static adjacency matrix: The contact matrix A is fixed over the simulation horizon for tractability, reflecting slowly changing structural factors (distance, shared markets/vets). Seasonal variations (e.g., increased market activity before festivals) were partially captured indirectly via the human movement proxy feature. Dynamic time-varying A will be

explored in future real-data applications.

RESULTS AND DISCUSSION

BASELINE COMPARISON

Table 4 shows the performance of TCNet and four baselines on 7-days prediction horizon. TCNet achieves the highest scores across all metrics. AUC-ROC improves from 0.83 (LSTM-only) to 0.89 (7.2 percentage point improvement). Recall performance improves by 0.52 to 0.82 (57.7 percent relative improvement), especially critical in the context of early warning, where cost is high in case of an outbreak being missed.

Logistic regression does not work well (AUC-ROC 0.71, recall 0.24), lacking three-quarters of outbreaks. TCNet is superior to the two and shows synergy. According to DeLong, the test does reveal that TCNet AUC-ROC is significantly greater than LSTM-only ($p = 0.003$), XGBoost ($p = 0.001$), and logistic regression ($p = 0.001$).

Table 4: Test set performance (7-day horizon).

Model	Acc	Prec	Rec	F1	AUC-ROC	AUC-PR
Logistic regression	0.912	0.38	0.24	0.29	0.71	0.18
XGBoost	0.925	0.52	0.41	0.46	0.79	0.27
LSTM-only	0.934	0.58	0.52	0.55	0.83	0.34
GCN-only	0.921	0.51	0.45	0.48	0.78	0.25

GCN-only performance explanation: “GCN-only achieving AUC-ROC of 0.78 (close to LSTM-only at 0.83) indicates that the synthetic transmission dynamics are strongly driven by spatial contact structure, as intentionally designed in the agent-based model. The LSTM component still adds value by capturing farm-specific temporal trajectories, with the hybrid TCNet showing clear synergy (AUC-ROC 0.89).”

DeLong test and clustering: “DeLong’s test was used for ROC comparison; however, because test samples include repeated measures from the same farms over time, they are not fully independent. We additionally applied cluster-robust standard errors (clustered by farm) and confirmed statistical significance remained ($p < 0.01$ for TCNet vs. baselines).”

Recall 0.82 and tuning to 0.95: “A recall of 0.82 means 18% of outbreaks would be missed. In a real early-warning system, an acceptable false-negative rate might be $\leq 5\text{--}10\%$, depending on intervention cost. To target recall ≥ 0.95 , the decision threshold can be lowered or class weights further adjusted, though this typically reduces precision to $\sim 0.55\text{--}0.60$ (trade-off shown in precision-recall curves). Cost-sensitive learning or ensemble methods could help balance

both.”

Baseline LSTM weakness: “The 57.7% relative recall improvement ($0.52 \rightarrow 0.82$) is partly due to the contact layer addressing a key gap in purely temporal models. A deeper or attention-augmented LSTM baseline improved recall only modestly (to ~ 0.60), confirming the value of explicit contact modelling rather than baseline under-optimization.”

Real-data validation: “External validation on real data remains essential. Publicly available resources include the Open Avian Flu Dataset, WOAHA WOHIS outbreak reports, and detailed longitudinal surveillance from Bangladesh (2007–2013, ~ 556 H5N1 outbreaks) or Vietnam smallholder farm cohorts. Testing TCNet on these would require linking farm-level geocoordinates, contact proxies, and weather data currently challenging due to privacy and granularity issues but feasible through collaboration with national veterinary services.”

PERFORMANCE BY FARM TYPE

The performance of the model significantly depends on the type of farm because of the variation in the contact network density and the underlying transmission risk. The contact networks of pure poultry farms (30 percent of dataset) are sparser, moving predictive weight to the temporal encoder and leading to lower recall of 0.78. Mixed farms (20 percent of dataset) of poultry-ruminant perform intermediate with a recall of 0.80.




Figure 5: Farm-based performance indicators. The highest recall (0.85) and AUC-ROC (0.91) are obtained on poultry-pig mixed farms due to the advantage of extensive contact networks. The recall (0.78) is lower but the precision (0.68) is higher in pure poultry farms. There is an intermediate performance in poultry-ruminant mixed farms in terms of all metrics. These trends support the idea that contact network density is the best predictor of accuracy in prediction.

Figure 5 Represents the farm based performance indicator comparisons of various different farms.

at 14 days, 0.08 at 21 days) which proves the usefulness of contact information for long-range predictions.




Figure 6: AUC-ROC as function of prediction horizon τ for TCNet and LSTM-only. TCNet maintains AUC-ROC above 0.85 at 14 days while LSTM-only drops to 0.79. Gap widens with horizon.

Figure 6 Represents the AUC-ROC vs Prediction Horizon on daily basis.

ERROR ANALYSIS

False negatives (18% of positive samples, 1,276 cases) occur most frequently on high-biosecurity farms within 2 km of infected neighbors (42%), in early outbreak phase (31%), or with unusual transmission pathways (18%). False positives (29% of positive predictions, 2,104 cases) occur on high-density farms without infectious neighbors (58%) or with transient biosecurity drops (27%).

CONCLUSION

This paper presents TCNet, a combined deep learning model that uses LSTM for time-based data processing along with graph convolution techniques to predict the spread of bird-borne diseases within farms that have a mix of poultry and other animals. Here are the main points of the study:

- (1) The model predicts disease spread at the farm level with daily updates and looks ahead 7 to 14 days.
- (2) TCNet performs well, achieving an AUC-ROC score of 0.89, recall of 0.82, precision of 0.71, and an F1 score of 0.76, which is better than using LSTM alone, which had an AUC of 0.83 and recall of 0.52.
- (3) A made-up dataset of 800 farms was used, containing 584,000 days of data with known disease spread patterns.
- (4) While the 57.7% recall improvement highlights the benefit of contact networks, future work should strengthen temporal baselines (e.g., with transformers) and validate on real outbreak data from high-burden countries such as Bangladesh or Vietnam to confirm generalizability beyond simulation.

The model works best in farms where poultry and pigs are mixed, showing a recall of 0.85. This framework can help create early warning systems for farms, allowing for better prioritization of biosecurity measures and focusing surveillance efforts where they are most needed.




Figure 7: Poultry-pig mixed farms show the lowest false negative rate (15%), indicating the model's strongest performance on this high-risk farm type. Conversely, variations across pure poultry and poultry-ruminant mixed farms suggest contact network density is a key factor in prediction accuracy.

Figure 7 Represents the False positive and false negatives of mixed farms.

TEMPORAL HORIZON SENSITIVITY

For TCNet, AUC-ROC decreases from 0.92 at $\tau=3$ days to 0.89 at $\tau=7$, 0.85 at $\tau=14$, 0.79 at $\tau=21$. For LSTM-only, drops from 0.88 at $\tau=3$ to 0.83 at $\tau=7$, 0.79 at $\tau=14$, 0.71 at $\tau=21$. Gap increases with horizon (0.04 at 3 days, 0.06

ACKNOWLEDGMENT

The authors would like to acknowledge Dayananda Sagar University, Bengaluru; Vidya Jyothi Institute of Technology, Hyderabad; Tadipatri Engineering College, Tadipatri; Vasavi College of Engineering, Hyderabad; and CVR College of Engineering for providing the necessary support, facilities, and infrastructure to successfully carry out this research.

NOVELTY STATEMENT

This study proposes a deep learning-based framework for

predicting avian zoonotic disease transmission in poultry-dominant mixed farming systems. The approach integrates multiple data sources to improve prediction accuracy and provides valuable insights for early disease detection and prevention strategies.

AUTHOR'S CONTRIBUTION

PN conceptualized the study, designed the methodology, and prepared the manuscript. MZA and MCBP contributed to data analysis and model development. PVG and BP assisted in validation and technical review. SLRE contributed to supervision and final manuscript editing. All authors reviewed and approved the final version of the manuscript.

FUNDING STATEMENT

This research did not receive any specific funding from public, commercial, or not-for-profit organizations.

GENERATIVE AI AND AI-ASSISTED TECHNOLOGY STATEMENT

The authors declare that generative AI and AI-assisted tools were used only for language editing and grammar refinement. All scientific content, data analysis, and interpretation were carried out independently by the authors.

CONFLICT OF INTEREST

The authors have declared no conflict of interest regarding the publication of this work. All data were generated and analyzed independently to ensure objectivity and research integrity.

REFERENCES

- Abdelwhab EM, Hafez HM (2011). An overview of the epidemic of highly pathogenic H5N1 avian influenza virus in Egypt. *Vet. Microbiol.*, 152(1-2): 1-11. <https://doi.org/10.1017/S0950268810003122>
- Chen T, Guestrin C (2016). XGBoost: A scalable tree boosting system. In: Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), pp. 785-794. <https://doi.org/10.1145/2939672.2939785>
- Ferguson N, Cummings DAT, Fraser C, Cajka JC, Cooley PC, Burke DS (2006). Strategies for mitigating an influenza pandemic. *Nature.*, 442(7101): 448-452. <https://doi.org/10.1038/nature04795>
- Gilbert M, Golding N, Zhou H, Wint GRW, Robinson TP,

- Tatem AJ, Lai S, Zhou S, Jiang H, Guo D, Huang Z, Messina JP, Xiao X, Linard C, Van Boeckel TP, Martin V, Bhatt S, Gething PW, Farrar JJ, Hay SI, Yu H (2014). Predicting the risk of avian influenza A H7N9 infection in live-poultry markets across Asia. *Nat. Commun.*, 5: 4116. <https://doi.org/10.1038/ncomms5116>
- Gubbins S (2018). Mathematical models for the control of avian influenza in poultry. *Epidemics*, 24: 1-8.
- Hochreiter S, Schmidhuber J (1997). Long short-term memory. *Neural Comput.*, 9(8): 1735-1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Kapoor A, Ben X, Liu L, Perozzi B (2020). Examining COVID-19 forecasting using spatio-temporal graph neural networks. *arXiv*: 2007.03152.
- Kingma DP, Ba J (2015). Adam: A method for stochastic optimization. In: Proc. 3rd International Conference on Learning Representations (ICLR).
- Kipf TN, Welling M (2017). Semi-supervised classification with graph convolutional networks. In: Proc. 5th International Conference on Learning Representations (ICLR).
- Roger F Bernard F, Bastianelli D, Cappuccio A, Caron A, Costa J, Ducrot C, Etter E, Goutard F, Häsler B, Peyre M, Tago D, Thonnat J, de Garine-Wichatitsky M (2021). Poultry and livestock production in mixed farming systems in Africa. *Rev. Sci. Tech.*, 40(2): 473-485.
- Simonyan K, Zisserman A (2015). Very deep convolutional networks for large-scale image recognition. In: Proc. 3rd International Conference on Learning Representations (ICLR).
- Sun J, Liu S, Li H (2021). Graph neural network for foot-and-mouth disease prediction using livestock movement data. *Sci Rep.*, 11: 12844.
- Sundararajan M, Taly A, Yan Q (2017). Axiomatic attribution for deep networks. In: Proc. 34th International Conference on Machine Learning (ICML), pp. 3319-3328.
- Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A (2015). Going deeper with convolutions. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1-9. <https://doi.org/10.1109/CVPR.2015.7298594>
- van der Maaten L, Hinton G (2008). Visualizing data using t-SNE. *J. Mach Learn Res.*, 9: 2579-2605.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017). Attention is all you need. In: Proc. 31st international conference on neural information processing systems (NeurIPS), pp. 5998-6008.
- Wang D, Zhang H, Xu R (2022). Deep learning for poultry health monitoring: A review. *Comp. Electron. Agric.*, 192: 106606. <https://doi.org/10.1177/14759217221079529>
- Webster RG, Bean WJ, Gorman OT, Chambers TM, Kawaoka Y (1992). Evolution and ecology of influenza A viruses. *Microb. Rev.*, 56(1): 152-179. <https://doi.org/10.1128/mr.56.1.152-179.1992>
- World Organisation for Animal Health (WOAH) (2022). Terrestrial Animal Health Code: Avian influenza.