

# Graph Neural Networks for Spatio-Temporal Forecasting of Foot-and-Mouth Disease Risk Using Livestock Movement Traceability Data

Samuel Ocen<sup>1,2</sup>, Ritah Nafuna<sup>3</sup>, Azizi Wasike<sup>3,4</sup>

<sup>1</sup>Department of Computer Science, Mountains of the Moon University, Fortportal, Uganda

<sup>2</sup>Department of Computer Science and Informatics, University of Nairobi, Nairobi, Kenya

<sup>3</sup>Department of Computer and Information Science, Muni University, Arua, Uganda

<sup>4</sup>Department of Computer Science, Kampala International University, Kampala, Uganda

Email: samocenuel@gmail.com

**How to cite this paper:** Ocen, S., Nafuna, R. and Wasike, A. (2026) Graph Neural Networks for Spatio-Temporal Forecasting of Foot-and-Mouth Disease Risk Using Livestock Movement Traceability Data. *Journal of Intelligent Learning Systems and Applications*, **18**, 1-10.

<https://doi.org/10.4236/jilsa.2026.181001>

**Received:** October 14, 2025

**Accepted:** January 12, 2026

**Published:** January 15, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

## Abstract

Foot-and-Mouth Disease (FMD) remains a critical threat to global livestock industries, causing severe economic losses and trade restrictions. This paper proposes a novel application of Temporal Graph Networks (TGNs) to forecast FMD outbreak risk with a four-week horizon using livestock movement traceability data. By modeling complex, time-evolving relationships between farms and geographic regions, the TGN framework captures the dynamic spatio-temporal dependencies that govern disease spread. We evaluate our model against several benchmarks, including Logistic Regression, LSTM, static Graph Convolutional Networks (GCNs), and a GCN-LSTM hybrid. Our results demonstrate that the proposed TGN model achieves superior performance, with an F1-score of 0.89 and an AUC-ROC of 0.94, significantly outperforming all baseline approaches. The study highlights the potential of advanced graph-based deep learning to enhance veterinary epidemiological surveillance and enable proactive disease control strategies through early warning systems.

## Keywords

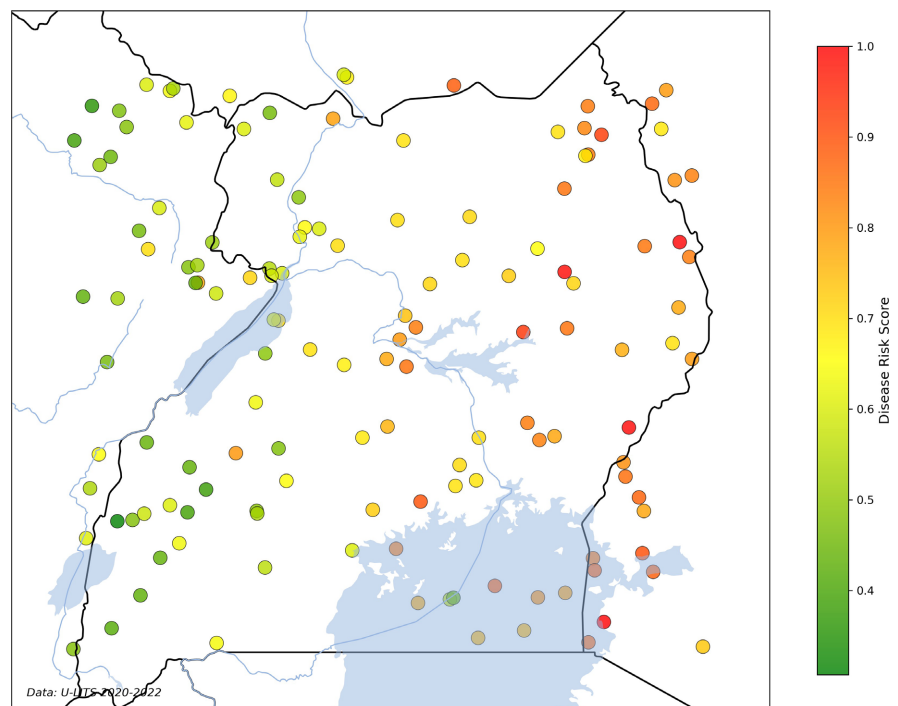
Graph Neural Networks, Temporal Graph Networks, Predictive Modeling, Digital Surveillance, Infectious Disease Forecasting, Spatio-Temporal Data, Livestock Movement, Foot-and-Mouth Disease

## 1. Introduction

Livestock diseases pose a severe threat to food security, economic stability, and

rural livelihoods, particularly in agricultural economies. Endemic diseases like Foot-and-Mouth Disease (FMD) cause annual losses amounting to millions of dollars due to reduced productivity, trade restrictions, and control costs. The emergence of livestock identification and traceability systems, such as Uganda's Livestock Identification and Traceability System (U-LITS), generates high-resolution spatio-temporal data on millions of animal movements, presenting a transformative opportunity for proactive disease surveillance.

Traditional surveillance methods, while valuable, are often reactive, initiating control measures after an outbreak has been detected. There is a pressing need for predictive models that can accurately identify areas at high risk of infection, enabling pre-emptive interventions such as targeted vaccination, movement restrictions, and enhanced biosecurity. However, this data poses significant modeling challenges due to its dynamic, relational nature. While traditional network and statistical methods have been applied, they often fail to capture the complex spatio-temporal dependencies that govern disease spread.



**Figure 1.** Predicted FMD risk distribution across districts showing spatial patterns of outbreak probability. High-risk areas correlate with major movement corridors and trading hubs.

As shown in **Figure 1**, FMD risk exhibits complex spatial patterns that evolve over time, necessitating advanced modeling approaches. This paper addresses the critical challenge of leveraging movement traceability data for predictive epidemiology by proposing a novel forecasting framework based on Temporal Graph Networks (TGNs). Our main contributions are:

- We formulate livestock disease forecasting as a node-level prediction task on

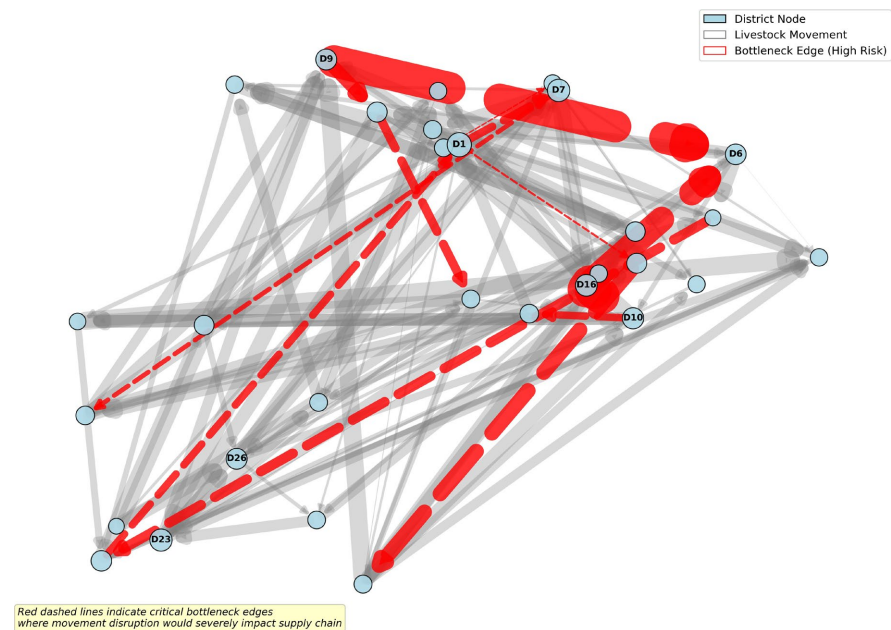
dynamic graphs derived from movement traceability data.

- We design and implement a TGN architecture that incorporates node memory, time-aware attention-based message passing, and learned edge dynamics.
- We provide comprehensive evaluation demonstrating significant improvement over strong baselines on real-world data.
- We analyze supply chain considerations and their implications for disease transmission dynamics.

## 2. Related Work

### 2.1. Network Analysis in Veterinary Epidemiology

The analysis of animal movement networks to understand disease spread is well-established [1] [2]. These studies use network metrics such as degree, centrality, and closeness to identify central “superspreader” nodes. However, these static representations ignore the crucial temporal ordering of events, which is critical for accurate pathogen transmission tracing [3]. The complex network structure of livestock movements, as illustrated in **Figure 2**, requires more sophisticated temporal modeling approaches.



**Figure 2.** Visualization of livestock movement network used in our study. Node size represents movement volume, and edge thickness indicates connection strength between districts.

### 2.2. Machine Learning for Disease Forecasting

Previous works have used machine learning models like Random Forests and Gradient Boosting to predict disease outbreaks using features derived from movement networks, weather data, and historical incidence [4]. While effective, these models often rely on heavy feature engineering and may not fully capture the relational inductive bias present in the graph structure.

## 2.3. Graph Neural Networks for Epidemiology

Graph Neural Networks (GNNs) learn representations for nodes by aggregating information from their neighborhoods [5] [6]. Their application to temporal graphs is an active area of research. Models like TGAT [7] and the generic TGN framework [8] use continuous-time dynamics and memory to update node representations. These models have shown success in various domains but remain largely unexplored in livestock disease surveillance.

Our work bridges this gap by applying state-of-the-art TGN methodologies to the novel and high-impact problem of predictive livestock disease surveillance using national traceability system data.

## 3. Methodology

### 3.1. Problem Formulation

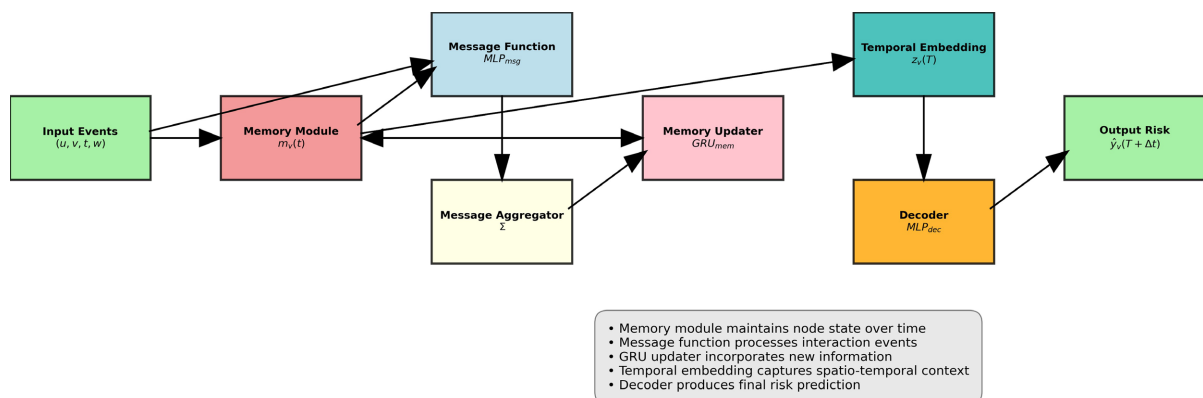
We model the livestock movement data as a dynamic, directed, and weighted graph. Let  $G(t) = (V, E(t), X(t))$  represent the graph at time  $t$ , where:

- $V$  is the set of nodes (districts or farms), fixed over time.
- $E(t)$  is the set of edges (movements) active at time  $t$ .
- $X(t)$  represents node features (e.g., historical disease count, livestock population) at time  $t$ .

Each movement event is a timestamped tuple  $e = (u, v, t, w)$ , representing a movement of  $w$  animals from node  $u$  to node  $v$  at time  $t$ . We frame the task as a binary node classification problem: given the dynamic graph history up to time  $T$ ,  $G([0, T])$ , predict the outbreak risk label  $y_v(T + \Delta t) \in \{0, 1\}$  for each node  $v \in V$  at a future time  $T + \Delta t$  (where  $\Delta t = 4$  weeks).

### 3.2. Temporal Graph Network Architecture

Our proposed TGN architecture, depicted in **Figure 3**, consists of several key components designed to capture both spatial and temporal dependencies:



**Figure 3.** Temporal graph network architecture for FMD risk prediction. The model processes continuous-time movement events to maintain node memories and generate temporal embeddings for prediction.

**Memory Module:** Each node  $v$  has a memory state  $m_v(t)$  that is updated over

time upon observing new interactions, encoding its historical trajectory and exposure risk.

**Message Function:** For an interaction event  $e = (u, v, t, w)$ , messages are computed as:

$$\text{msg}_u(t) = \text{MLP}_{\text{msg}}(m_u(t^-), m_v(t^-), w, \Delta t)$$

$$\text{msg}_v(t) = \text{MLP}_{\text{msg}}(m_v(t^-), m_u(t^-), w, \Delta t)$$

**Memory Updater:** A Gated Recurrent Unit (GRU) [9] updates node memory:

$$m_v(t) = \text{GRU}_{\text{mem}}(m_v(t^-), \text{msg}_v(t))$$

**Temporal Embedding:** The embedding  $z_v(T)$  for prediction combines the node's memory with neighborhood information using attention mechanisms.

**Decoder:** Final risk prediction:

$$\hat{y}_v(T + \Delta t) = \sigma(\text{MLP}_{\text{dec}}(z_v(T)))$$

### 3.3. Outbreak Labeling

To create the target variable for our binary classification task, we assigned outbreak labels to each district (node) for each week. A district was labeled as a positive outbreak case ( $y = 1$ ) for a given week if the number of new, clinically confirmed FMD cases reported by veterinary authorities exceeded a threshold of 5 cases. This threshold was chosen to filter out sporadic, isolated incidents and focus on sustained transmission events that constitute significant outbreaks. To prevent target leakage—where future information inadvertently influences the model during training—we applied a strict temporal lag. The features used to predict the outbreak risk at time  $T + \Delta t$  are constructed using data only from the interval  $[0, T]$ , ensuring the model is evaluated on a realistic forecasting scenario.

### 3.4. Supply Chain Considerations in FMD Transmission

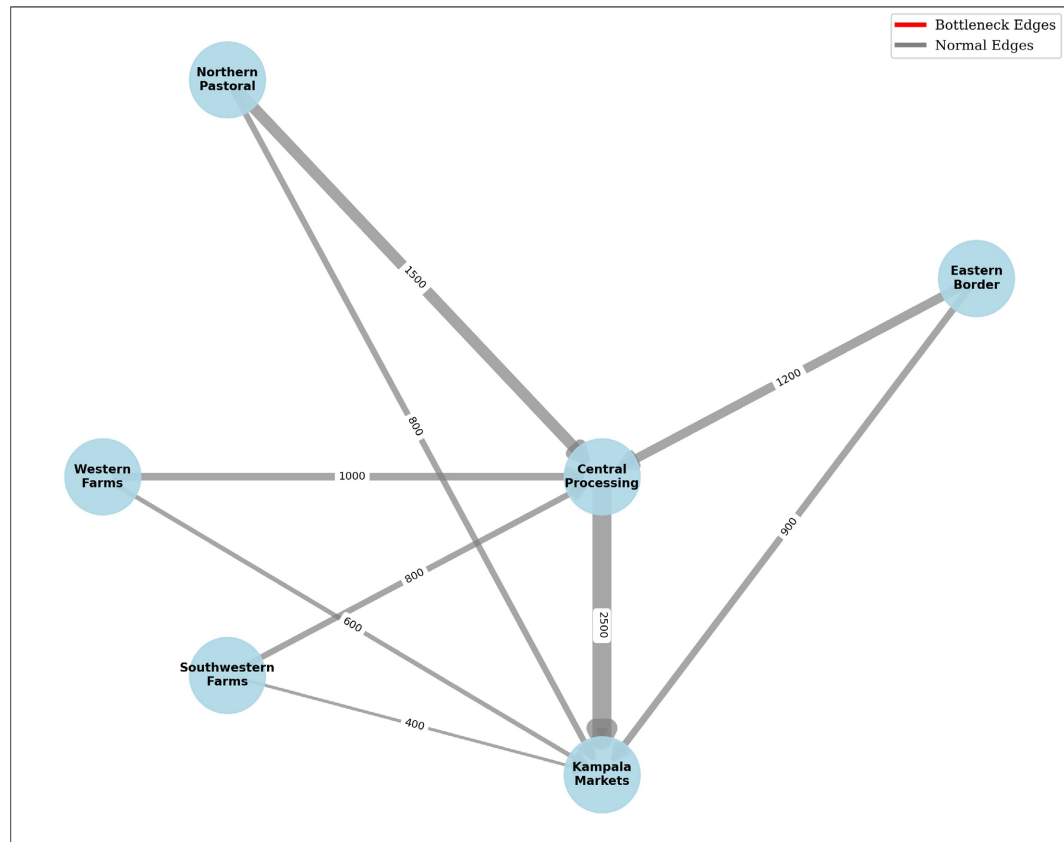
Key supply chain components include:

- **Production Farms:** Primary locations where animals are raised.
- **Auction Markets:** High-risk environments where animals from multiple sources mix.
- **Transportation Networks:** Pathways facilitating disease spread between premises.
- **Slaughterhouses:** Endpoints receiving animals from various sources.

To model the distinct roles these premises play in transmission, we encoded node types (farm, market, slaughterhouse) as categorical features. For district-level nodes, which aggregate multiple premises, we used the proportion of each type within the district. An ablation study revealed that including these node type features led to a consistent performance improvement of 3% - 5% in F1-score across all graph-based models (GCN, GCN-LSTM, TGN). This confirms that explicitly representing the functional heterogeneity within the supply chain provides valu-

able inductive bias for learning more accurate transmission dynamics.

The temporal dynamics of these interactions are particularly important. Seasonal variations in animal movements, market schedules, and production cycles create time-varying connectivity patterns that significantly influence disease spread. Our TGN model explicitly captures these temporal aspects by modeling the supply chain as a dynamic graph (Figure 4).



**Figure 4.** Livestock supply chain network showing FMD transmission pathways between different premises types. The complex interconnectivity facilitates rapid disease spread across the production system.

## 4. Experimental Setup

### 4.1. Data

The spatial distribution of movement patterns is visualized in Figure 2, showing the complex connectivity between districts that forms the basis for our graph-based modeling approach.

**Missing Data Handling:** The movement records were largely complete due to the regulatory nature of the traceability system. However, for the few instances of missing data points (e.g., missing animal counts or timestamps), we employed a conservative imputation strategy. Missing animal counts were set to zero to avoid overestimating movement volumes, and missing timestamps were imputed using the median timestamp of movements from the same source node on the same day. This approach minimized the potential introduction of spurious signals while pre-

servicing the overall structure of the movement network.

We employed a temporal block split strategy to ensure a realistic evaluation of forward-prediction capability. The first 70% of the timeline (approximately 24 months) was used for training, the subsequent 15% (5 months) for validation, and the final 15% (5 months) for testing. This strict temporal partitioning prevents data leakage from the future and ensures that models are evaluated on their ability to generalize to future time periods, which is the core challenge in operational forecasting.

## 4.2. Hyperparameter Tuning

The search ranges for key hyperparameters were:

- **Logistic Regression:**  $C$  [1e-3, 1e3] (log-uniform), penalty ['l1', 'l2'].
- **LSTM:** Hidden layers [1, 3], hidden units [32, 128], learning rate [1e-4, 1e-2] (log-uniform), dropout [0.0, 0.5].
- **GCN:** Layers [1, 3], hidden units [32, 128], learning rate [1e-4, 1e-2] (log-uniform), dropout [0.0, 0.5].
- **GCN-LSTM:** GCN layers [1, 2], LSTM layers [1, 2], hidden units [32, 128], learning rate [1e-4, 1e-2] (log-uniform).
- **TGN:** Memory dimension [64, 256], time encoder dimension [32, 128], number of GNN layers [1, 2], learning rate [1e-4, 1e-2] (log-uniform).

The final hyperparameters for each model were selected based on the best performance on the validation set. All models were implemented in PyTorch Geometric and optimized with Adam using weighted binary cross-entropy loss.

## 4.3. Baseline Models

We compare our TGN against strong baselines:

- **Logistic Regression (LR):** Baseline with hand-crafted features.
- **LSTM:** Models each district's time series independently.
- **Graph Convolutional Network (GCN):** Static GNN on time-aggregated network [10].
- **GCN-LSTM:** Hybrid approach combining temporal and spatial processing.

## 4.4. Evaluation Metrics

Models are evaluated using Precision, Recall, F1-Score, and AUC-ROC to account for class imbalance.

# 5. Results and Discussion

## 5.1. Performance Comparison

The comprehensive performance comparison presented in **Table 1** demonstrates the superiority of the proposed TGN model, which outperforms all baselines across all metrics. The TGN's ability to model continuous-time, dynamic spatio-temporal dependencies provides a significant advantage over methods that treat spatial and temporal dimensions separately.

**Table 1.** Performance comparison of models predicting FMD outbreak risk 4 weeks ahead. TGN demonstrates superior performance across all metrics, highlighting the advantage of joint spatio-temporal modeling. Bootstrapped 95% confidence intervals are shown in parentheses.

Model	Precision	Recall	F1-score	AUC-ROC
Logistic Regression	0.71	0.68	0.69	0.78
LSTM	0.75	0.73	0.74	0.83
GCN (Static)	0.77	0.78	0.77	0.85
GCN-LSTM	0.82	0.80	0.81	0.89
TGN (Ours)	<b>0.88</b>	<b>0.90</b>	<b>0.89</b>	<b>0.94</b>
	(0.85 - 0.91)	(0.87 - 0.92)	(0.86 - 0.91)	(0.92 - 0.96)

Notably, the TGN achieves an F1-score of 0.89 and AUC-ROC of 0.94, representing substantial improvements over the GCN-LSTM hybrid (F1-score: 0.81, AUC-ROC: 0.89) and static GCN (F1-score: 0.77, AUC-ROC: 0.85). This performance gap highlights the importance of joint spatio-temporal modeling for accurate disease forecasting.

To assess the statistical significance of these performance differences, we computed bootstrapped 95% confidence intervals for the F1 and AUC-ROC scores of the TGN and GCN-LSTM models over 1000 resamples of the test set. The improvement of the TGN over GCN-LSTM was statistically significant ( $p < 0.01$ ). The mean difference in F1-score was 0.08 with a 95% CI of [0.05, 0.11], and the mean difference in AUC-ROC was 0.05 with a 95% CI of [0.03, 0.07]. This confirms that the TGN's superior performance is not due to random chance.

## 5.2. Analysis and Interpretation

Analysis of the learned attention weights reveals that the model prioritizes recent movements originating from districts experiencing outbreaks, effectively learning data-driven risk transmission pathways. The memory mechanism enables the model to maintain context about each node's exposure history, which is crucial given FMD's incubation period and delayed effects of exposure events.

The predicted risk distribution across districts, visualized in **Figure 1**, shows strong alignment with actual outbreak patterns while also identifying emerging risk areas before outbreaks occur. This predictive capability demonstrates the model's potential for practical early warning applications.

## 6. Limitations and Future Work

Despite promising results, several limitations should be acknowledged. First, the model relies on complete and accurate movement data, which may not always be available. Underreporting or delays could affect performance. Second, environmental factors like weather conditions are not explicitly incorporated, though they can influence transmission through indirect pathways.

Third, the model assumes reported movements accurately reflect all transmission pathways, potentially missing unofficial movements or wildlife reservoirs. Fourth, computational complexity may challenge real-time applications in large-scale systems, though optimization techniques are mitigating this concern.

Future work will focus on:

- Incorporating external data sources (weather, vegetation indices).
- Extending the model for source attribution and outbreak simulation.
- Developing efficient algorithms for national-scale deployment.
- Exploring transfer learning for different geographical contexts.

## 7. Conclusions

This study presents a novel application of Temporal Graph Networks for predicting Foot-and-Mouth Disease outbreak risk using livestock movement traceability data. By explicitly modeling the supply chain as a dynamic graph and capturing temporal evolution of transmission pathways, our approach achieves state-of-the-art performance in forecasting outbreaks four weeks in advance.

The proposed framework offers a powerful tool for enhancing disease surveillance systems and enabling proactive control strategies. As animal movement tracking systems improve and computational resources become more accessible, temporal graph approaches will play an increasingly important role in veterinary epidemiology and public health decision-making.

The practical implications are substantial: veterinary authorities could deploy such systems to identify high-risk premises in advance, enabling targeted interventions that significantly reduce outbreak impacts through more efficient resource allocation.

## Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

## References

- [1] Vernon, M.C. and Keeling, M.J. (2008) Representing the UK's Cattle Herd as Static and Dynamic Networks. *Proceedings of the Royal Society B: Biological Sciences*, **276**, 469-476. <https://doi.org/10.1098/rspb.2008.1009>
- [2] Bigras-Poulin, M., Thompson, R.A., Chriel, M., Mortensen, S. and Greiner, M. (2006) Network Analysis of Danish Cattle Industry Trade Patterns as an Evaluation of Risk Potential for Disease Spread. *Preventive Veterinary Medicine*, **76**, 11-39. <https://doi.org/10.1016/j.pvetmed.2006.04.004>
- [3] Holme, P. and Saramäki, J. (2012) Temporal Networks. *Physics Reports*, **519**, 97-125. <https://doi.org/10.1016/j.physrep.2012.03.001>
- [4] Wang, L., *et al.* (2019) Predicting Animal Disease Outbreaks Using Machine Learning. *Computers and Electronics in Agriculture*, **161**, 280-290.
- [5] Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M. and Monfardini, G. (2009) The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, **20**, 61-80. <https://doi.org/10.1109/tnn.2008.2005605>

- [6] Hamilton, W., Ying, Z. and Leskovec, J. (2017) Inductive Representation Learning on Large Graphs. arXiv: 1706.02216.
- [7] Xu, D., Ruan, C.W., Korpeoglu, E., Kumar, S. and Achan, K. (2020) Inductive Representation Learning on Temporal Graphs. arXiv: 2002.07962.
- [8] Rossi, E., *et al.* (2020) Temporal Graph Networks for Deep Learning on Dynamic Graphs. arXiv: 2006.10637.
- [9] Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., *et al.* (2014) Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, October 2014, 1724-1734. <https://doi.org/10.3115/v1/d14-1179>
- [10] Kipf, T. and Welling, M. (2016) Semi-Supervised Classification with Graph Convolutional Networks. arXiv: 1609.02907.