Forecasting of Acute Upper Respiratory Tract Infection in Ukraine during the War Conditions using LSTM Model*

Mykola Butkevych[1,*], Kateryna Vasiuk[1], Iryna Trofymova[1] and Tetyana Chumachenko[2]

[1] *National Aerospace University "Kharkiv Aviation Institute", Vadym Manko str, 17, 61070 Kharkiv, Ukraine*
[2] *Kharkiv National Medical University, Nauky ave., 4, 61001 Kharkiv, Ukraine*

Abstract
Acute upper respiratory tract infections remain a leading cause of outpatient visits worldwide. In Ukraine, the Russian full-scale war has disrupted care and infrastructure, increasing the need for reliable near-term forecasts to support winter preparedness and resource planning. To develop and evaluate a deep learning approach for short-term URTI forecasting under conflict conditions. We built a three-layer LSTM (300–1200–600 units) with a 16-step input window, Adam optimization, MSE loss, and a recursive multi-step strategy. Inputs were scaled to [0,1]. We engineered seasonality using trigonometric month features. Training instabilities linked to gradient explosion were mitigated by normalization and learning-rate tuning. Performance was assessed with MAPE, MSE, and RMSE on held-out data. Cumulative target forecasting improved accuracy versus the raw series: MAPE fell from 6.05% to 3.23%, and RMSE from 111,741.67 to 55,184.62. Seasonal features reduced error and improved fit, while training stabilized after preprocessing and optimizer tuning, despite residual MAPE oscillations during learning. An LSTM with explicit seasonality encoding and careful stabilization can provide actionable week-ahead URTI forecasts for Ukraine's conflict-affected health system.

Keywords
LSTM, machine learning, deep learning, epidemic model

1. Introduction

Acute upper respiratory tract infections (URTIs), including the common cold, pharyngitis, and sinusitis, are among the most frequent acute illnesses worldwide and remain a major driver of outpatient visits and missed productivity [1]. Recent global assessments estimate roughly 12.8 billion new URTI episodes in 2021, underscoring their high incidence and health-system impact across all age groups [2]. Beyond their clinical burden, URTIs contribute substantially to healthcare utilization and costs, and inappropriate antibiotic use remains common despite predominantly viral etiology, reinforcing antimicrobial resistance concerns [3]. Seasonal dynamics further complicate planning: URTIs and other respiratory viruses display strong spatiotemporal patterns, with sentinel systems documenting waves that vary by timing and intensity across regions [4].

In Ukraine, the urgency of effective URTI forecasting has been heightened by the Russian full-scale invasion that began on 24 February 2022, which has disrupted health services, damaged critical infrastructure, and increased exposure to cold conditions that elevate respiratory disease risk [5]. WHO verified over 1,000 attacks on healthcare within the first 15 months of the war, with continued documentation through 2023, contributing to reduced access to care and strained service delivery [6]. Winter risk assessments for Ukraine warn that inadequate heating, damaged housing, and power outages can raise respiratory morbidity, particularly among older adults and people with chronic conditions. Humanitarian health planning for the 2025-2026 winter explicitly prioritizes managing acute respiratory infections, reflecting their expected seasonal surge under conflict-affected living conditions [7]. Despite these pressures, influenza/acute respiratory infection

ProfIT AI'25: 5th International Workshop of IT-professionals on Artificial Intelligence, October 15–17, 2025, Liverpool, UK
1* Corresponding author.
m.v.butkevych@khai.edu (M. Butkevych); k.s.vasiuk@gmail.com (K. Vasiuk); irina.trofymova@gmail.com (I. Trofymova); tatalchum@gmail.com (T. Chumachenko)
0000-0001-8189-631X (M. Butkevych); 0009-0003-8121-6945 (K. Vasiuk); 0000-0002-1537-5601 (I. Trofymova); 0000-0002-4175-2941 (T. Chumachenko)

(ARI) sentinel paradigms continue to provide a foundation for situational awareness in the broader WHO European Region, emphasizing the need for timely, locally adapted analytics that can operate amid data interruptions.

Artificial intelligence (AI) and machine learning methods, particularly sequence models, offer promising tools for forecasting and simulating infectious diseases when traditional assumptions are strained by conflict, displacement, and infrastructure damage [8]. Reviews of AI for infectious disease modeling highlight how deep learning can complement mechanistic frameworks, integrate heterogeneous data streams, and improve short-term predictive accuracy for respiratory pathogens [9-10]. Among deep learning approaches, long short-term memory (LSTM) networks are well-suited to capture nonlinear temporal dependencies in surveillance time series and have shown competitive or superior performance in forecasting influenza and other respiratory virus activity across diverse settings [11]. Emerging scoping and systematic reviews also document growing use of AI-enabled early warning systems that fuse epidemiological indicators with contextual signals (e.g., mobility, weather, or health-service data) to anticipate near-term incidence [12].

Against this backdrop, we address the problem of near-term URTI forecasting in Ukraine under war conditions using an LSTM-based approach. The study aims to develop an LSTM framework to evaluate the forecast accuracy of URTIs.

The current research is part of a comprehensive information system for assessing the impact of emergencies on the spread of infectious diseases described in [13].

2. Materials and Methods

For the experimental study, we used monthly data on URTI morbidity in Ukraine from 2013 to 2024, collected by the Center for Public Health of the Ministry of Health of Ukraine. The data distribution is presented in Figure 1.

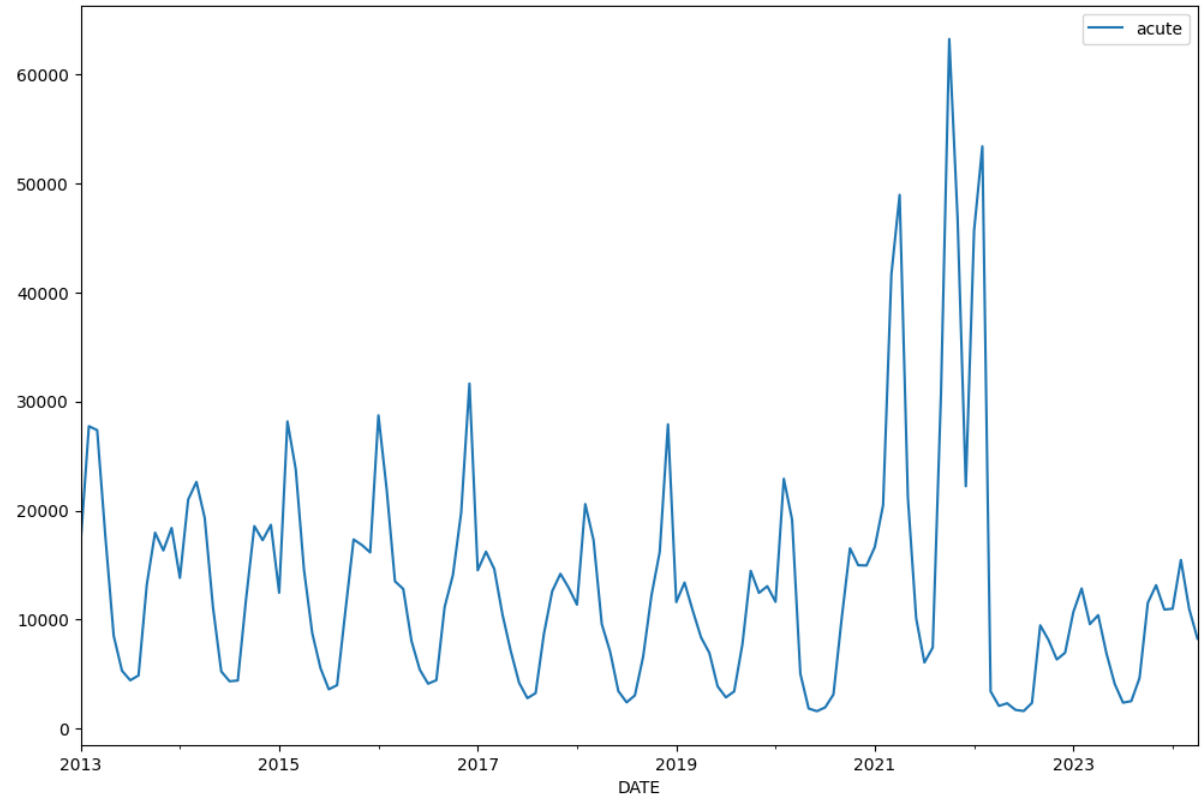

Figure 1: Morbidity by URTI in Ukraine (2013-2024).

We designed a three-layer LSTM network to model non-linear temporal dependencies in the URTI surveillance series. The input to the network is a window of the most recent 16 time steps with $n_{features}$ covariates per step (tensor shape (16, $n_{features}$)). The first LSTM layer has 300 units and returns a sequence (16, 300), followed by dropout to limit overfitting. The second LSTM layer has 1200 units (also returning sequences) to expand capacity for complex patterns, followed by dropout. The third LSTM layer reduces dimensionality to 600 units and returns only the final state (vector of length 600). After a final dropout, a dense layer with one neuron produces the next-step forecast. The model has approximately 11.89 million trainable parameters, concentrated in the LSTM layers, which balances expressive power with training stability (Figure 2).

Layer (type)	Output Shape	Param #
lstm (LSTM)	(None, 16, 300)	362,400
dropout (Dropout)	(None, 16, 300)	0
lstm_1 (LSTM)	(None, 16, 1200)	7,204,800
dropout_1 (Dropout)	(None, 16, 1200)	0
lstm_2 (LSTM)	(None, 600)	4,322,400
dropout_2 (Dropout)	(None, 600)	0
dense (Dense)	(None, 1)	601

Figure 2: The model architecture.

Data are converted into supervised learning pairs using a sliding-window generator: for each time t, the feature tensor aggregates the previous 16 steps, and the target is the value at t+1. We implement this with a standard time series generator that creates batched windows with configurable length, stride, and sampling rate. All continuous inputs are scaled to [0,1] using a transformation fit on the training split. Predictions are inverse-transformed to the original scale for evaluation and visualization. This improves numerical stability during optimization and allows direct comparison of predicted and observed series on plots.

To obtain multi-step horizons, we use a recursive (auto-regressive) strategy. The model is trained for one-step-ahead prediction and then applied iteratively, feeding each forecast back into the input window to predict the next point. Although errors may accumulate with horizon length, this approach is data-efficient and leverages the strong one-step learner. Forecasts are generated for the desired horizon and then compared with ground truth. Predicted and observed series are plotted on the original scale to assess fit to trends and seasonal patterns.

Training uses the Adam optimizer for stochastic gradient-based updates, with mean squared error (MSE) as the loss. We monitor mean absolute percentage error (MAPE) and MSE during training. Regularization includes dropout after each recurrent block, before the output layer, and early stopping on validation loss to prevent overfitting. Hyperparameters (batch size, learning rate, dropout rate) are selected by grid search with a time-ordered validation split.

To encode annual seasonality without discontinuities between December and January, we add cyclical month features using trigonometric transforms $\sin(2*\pi*m/12)$ and $\cos(2*\pi*m/12)$, where m is the month index. These features help the network capture recurring annual structures that simple categorical month encodings may fragment. A before-and-after comparison shows lower RMSE and MAE when these seasonal terms are included, indicating improved short-term accuracy for URTI dynamics (Figure 3).

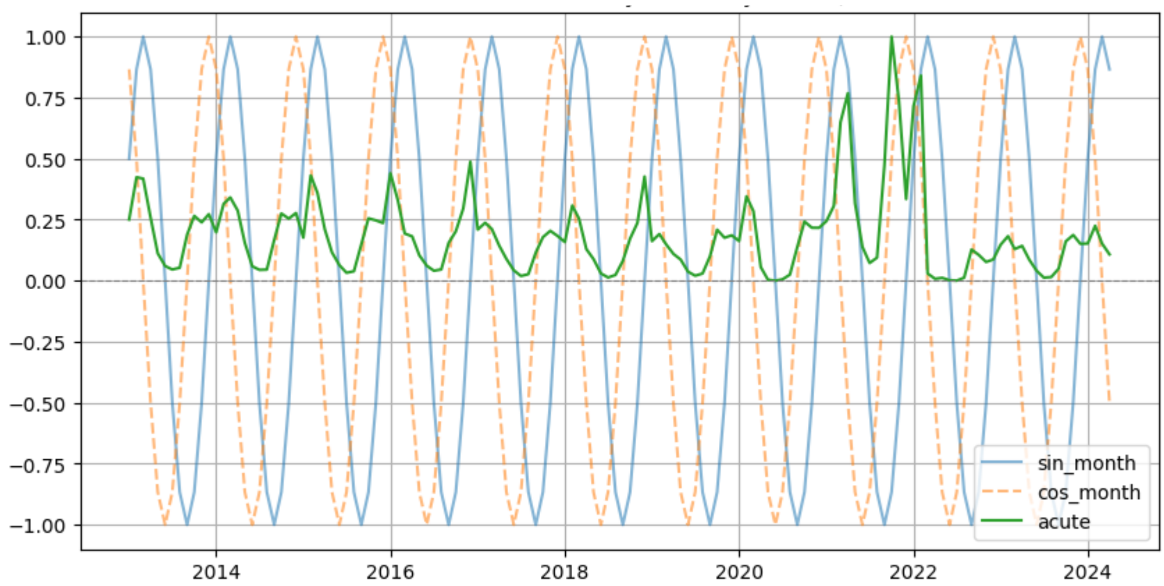

Figure 3: The comparison between seasonality and the input data.

3. Results

During testing, we observed episodes in which the loss became undefined (NaN). Diagnostic checks indicated a gradient explosion as the underlying cause, resulting from insufficient input normalization, an unsuitable loss choice, or a poorly tuned learning rate. To address this, we first normalized all inputs to the [0,1] range before training, which reduced numerical instability. We then tuned the learning rate of the Adam optimizer to stabilize weight updates without abrupt jumps. After these changes, training proceeded without NaNs, though the error trajectory sometimes remained sensitive. As shown in Figure 4, the MAPE curve exhibits marked oscillations across epochs, indicating forecast instability and periods with large deviations during learning.

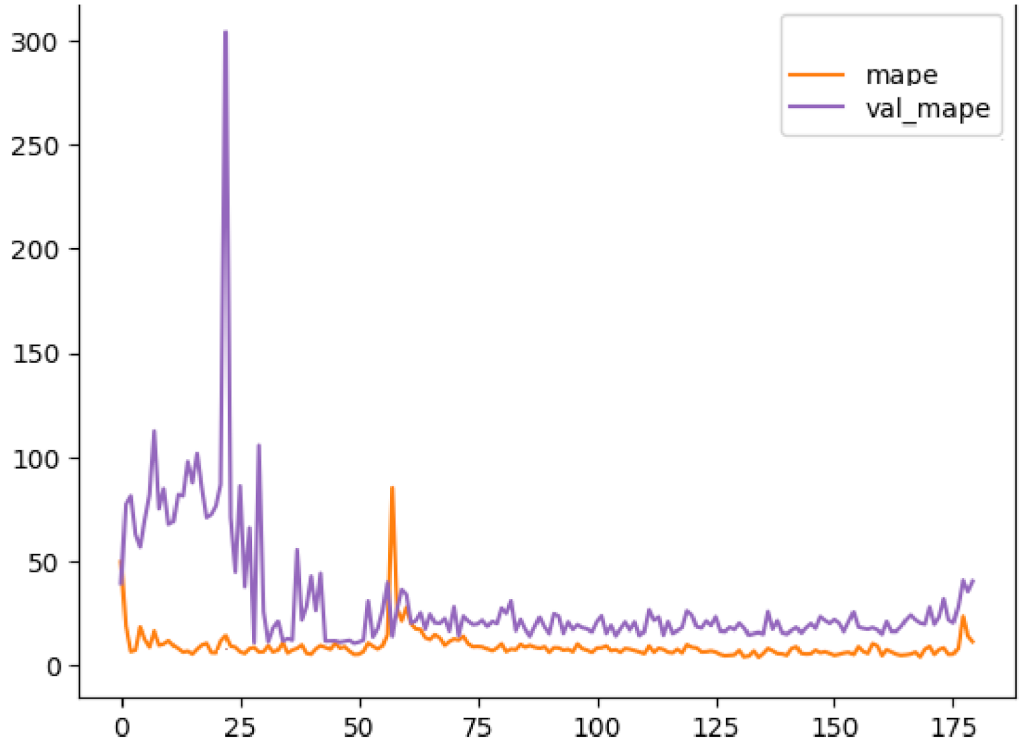

Figure 4: Metrics optimization.

Next, we analyzed forecasts generated in the cumulative form of the series. Unlike point-by-point prediction of the raw series, the cumulative forecast aggregates values over time. This aggregation smooths random fluctuations and noise, revealing long-term tendencies more clearly. In settings where the cumulative effect is of interest, such as financial forecasts (accumulated revenue or costs), product-demand analysis, or assessment of the overall impact of a process, this representation can improve interpretability and short-term robustness. At the same time, cumulative values may blur local changes and mask critical breaks in trend because the accumulated trajectory has inherent inertia. Sharp jumps or structural shifts can be less visible than in the raw series. Visual inspection is consistent with these properties: the pre-invasion cumulative forecast tracks the general trajectory with damped variability (Figure 5), while the full-series cumulative forecast preserves the broad trend despite conflict-related shocks (Figure 6).

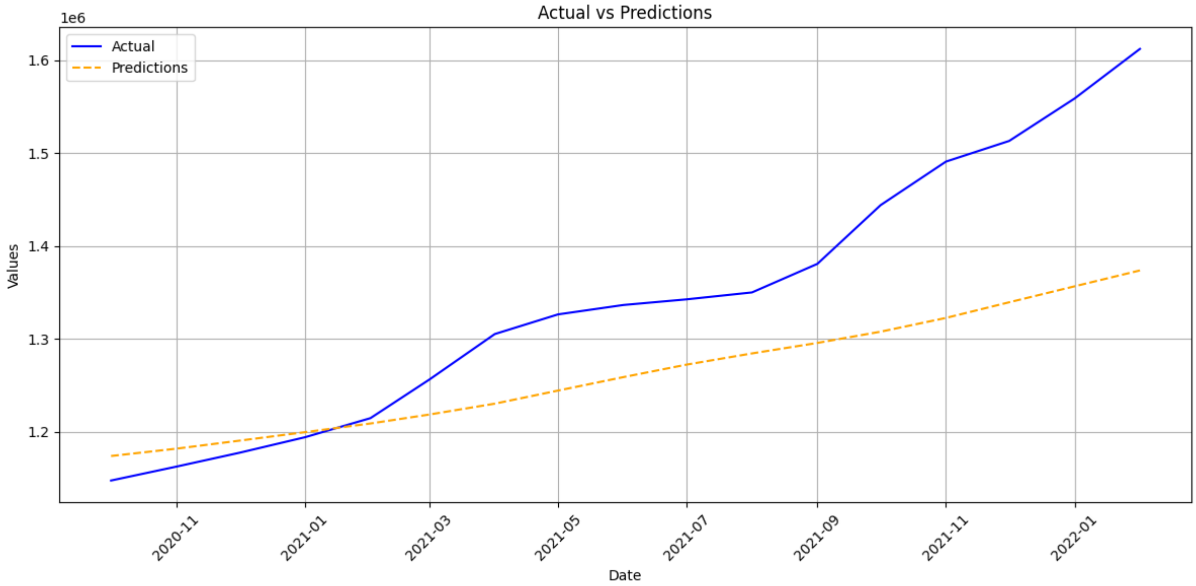

Figure 5: Forecast of cumulative morbidity before the Russian full-scale invasion.

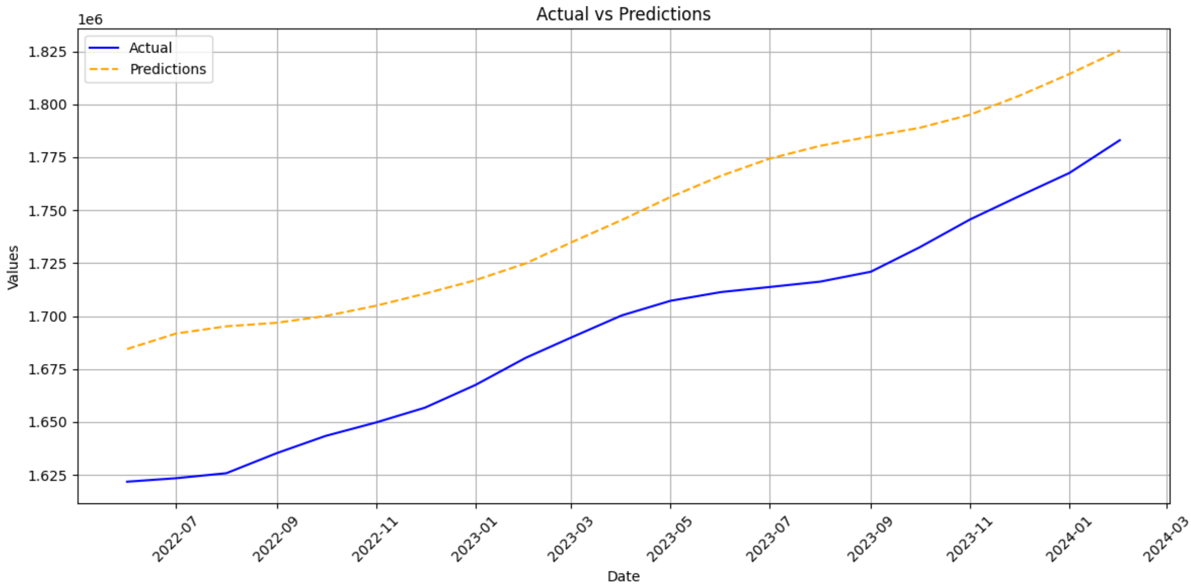

Figure 6: Forecast of cumulative morbidity including the data after the Russian full-scale invasion.

Quantitatively, cumulative modeling yielded lower error when the post-invasion period was included. As summarized in Table 1, MAPE decreased from 6.05% (cumulative data, pre-invasion) to 3.23% (cumulative data, including post-invasion), a relative reduction of ~46.6%. MSE fell from

12,486,199,938 to 3,045,342,366 (−75.6%), and RMSE declined from 111,741.67 to 55,184.62 (−50.6%). These differences align with the visual patterns noted above: aggregation enhances signal-to-noise for medium-term tendencies, though it can attenuate abrupt local shifts. The diagnostics and comparative evaluation indicate that careful normalization and learning-rate tuning mitigate training pathologies linked to gradient explosion, training dynamics remain somewhat volatile as reflected by MAPE oscillations in Figure 4, and cumulative forecasts improve headline accuracy while requiring cautious interpretation of short-lived changes, as illustrated in Figures 5 and 6.

Table 1
Results of the experimental study

Metric	Before the war	Including the war data
MAPE	6.05	3.23
MSE	12486199938	3045342366
RMSE	111741.67	55184.62

4. Discussion

This study shows that an LSTM configured as a deep, three-layer sequence model can produce stable short-term forecasts of acute URTI activity in Ukraine when inputs are carefully normalized and the optimizer is tuned. However, that training remains sensitive, with periods of loss of stability and oscillating MAPE during learning. The reduction in error for the cumulative series, MAPE from 6.05% to 3.23%, MSE from 12.49×10^9 to 3.05×10^9, and RMSE from 111,742 to 55,185, indicates that aggregating the signal improves headline accuracy, consistent with its noise-suppressing effect. At the same time, visual inspection confirms that cumulative trajectories can mask local regime shifts, especially around conflict-related shocks. The observed benefits of cyclical month encodings align with prior evidence that Fourier-type terms capture annual periodicity without artificial breaks between December and January.

These findings fit broader results in infectious disease forecasting. Deep learning can complement classical models by learning non-linear dependencies and combining heterogeneous signals, but it is sensitive to nonstationarity and evaluation choices. Recent reviews highlight that LSTM-based and related deep architectures improve short-term forecasts for respiratory infections when the data pipeline is well specified and exogenous context is available [9]. The empirical gain we see for cumulative modeling is also consistent with the literature's emphasis on smoothing to stabilize near-term predictions, while cautioning that aggregation can hide abrupt shifts requiring separate detection [14].

From a modeling perspective, the architecture's capacity is high relative to the monthly data length, which explains the training sensitivity and NaN episodes before normalization and learning-rate tuning [15]. Beyond the steps we used, practical stability can be improved with gradient-norm or adaptive gradient clipping, which have theoretical and empirical support for controlling exploding updates in deep networks, including recurrent models [16]. Given the depth and parameter count, additional regularization and early stopping remain advisable. A leaner recurrent backbone or dilated temporal convolutions may reduce variance without sacrificing accuracy. Comparative studies and hybrid LSTM-Transformer designs in epidemic forecasting suggest that architectures with attention and mobility cues can enhance 1-4-week horizons [12].

Using a recursive multi-step strategy is data-efficient and straightforward, but it accumulates error as the horizon grows. Alternative strategies, including direct, multi-input–multi-output, or hybrid schemes, can mitigate compounding error and deserve testing in this setting. Recent work unifies these approaches and shows when each is advantageous, suggesting that hybridization may outperform pure recursion under distribution shift [17]. In parallel, evaluation should follow best practices for time-ordered splits and rolling origins and report metrics that are robust to scale and near-zero denominators. While MAPE offers interpretability, it can misbehave when values approach zero. Complementing it with RMSE/MAE and scale-free measures such as MASE or interval scores will give a fuller picture of forecast quality for URTIs [18].

Seasonality engineering appears important for this application. The improvement observed when adding trigonometric month encodings accords with established guidance on Fourier terms for seasonal effects, which avoid discontinuities introduced by categorical month dummies [19]. Future experiments could expand these seasonal bases or allow time-varying seasonality to capture changes in school calendars, heating disruptions, or mobility during wartime winters.

Contextual signals could further strengthen forecasts. Studies show that integrating mobility, meteorology, and digital trace indicators improves early warning and near-term accuracy for respiratory infections by explaining deviations from historical seasonal patterns [20]. For Ukraine, candidate covariates include local temperature, humidity, outages of power or heat, displacement flows, and mobility between oblasts. LSTM/Transformer hybrids leveraging mobility have improved 1-4-week predictions elsewhere, and reviews of early-warning systems point to consistent gains when such context is added to surveillance baselines. A spatiotemporal extension could also capture diffusion between oblasts under displacement.

Uncertainty quantification is essential for decision support. Point forecasts alone can overstate confidence, especially under conflict-driven shocks [21]. Conformal prediction methods offer distribution-free intervals that adapt under nonstationarity and can be wrapped around existing LSTM predictors [22]. This would allow planners to set risk-aware thresholds for surge staffing and supplies. In parallel, ensembles that blend statistical, machine-learning, and deep models have improved influenza forecasts and could stabilize week-ahead URTI predictions in Ukraine.

Several limitations temper interpretation. First, monthly data constrain sample size and can hide rapid shifts in consultation behavior or reporting during attacks on infrastructure. Higher-frequency data would help. Second, the gains observed for cumulative targets do not guarantee better detection of local peaks. Operational monitoring should pair cumulative forecasts with raw series nowcasts and anomaly detection. The present results support the feasibility of LSTM-based forecasting for URTIs in a conflict-affected health system, provided that training stability is secured and seasonality is encoded explicitly. The main practical lessons are: aggressive preprocessing and optimizer tuning are necessary but may be insufficient without gradient clipping; cumulative targets can improve headline accuracy but must be paired with raw-series monitoring to avoid missing abrupt changes; and adding contextual and spatial signals, uncertainty intervals, and ensemble baselines is a promising path to robust, actionable forecasts for winter-season planning in Ukraine.

5. Conclusions

We developed and evaluated a three-layer LSTM model to forecast acute URTIs in Ukraine under war conditions. The pipeline used feature scaling to [0,1], Adam optimization with MSE loss, recursive multi-step prediction, and cyclical month features. Normalization and learning rate tuning mitigated training instabilities linked to gradient explosion. Cumulative target forecasting improved accuracy relative to the raw series, with MAPE falling from 6.05% to 3.23%, MSE from 12,486,199,938 to 3,045,342,366, and RMSE from 111,741.67 to 55,184.62.

The paper contributes a deep LSTM configuration tailored to conflict-affected surveillance streams with explicit safeguards for training stability; an empirical comparison showing that cumulative modeling can materially reduce short-horizon error for URTI forecasting in this setting; and a simple, effective seasonality encoding with trigonometric month features that improves fit while avoiding discontinuities between calendar years.

The approach provides forecasts that can inform staffing, supply planning, and winter preparedness in a health system stressed by infrastructure damage and displacement. The pipeline is lightweight, reproducible, and compatible with routine surveillance. Its outputs (point forecasts and error metrics) are easily integrated into operational dashboards.

Next steps include adding gradient clipping and weight decay to stabilize training further, testing alternative forecasting strategies and model families, incorporating contextual and spatial covariates, and reporting uncertainty via conformal prediction or ensemble methods. Extending from monthly to weekly data and performing rolling-origin evaluations with a richer metric set will strengthen external validity and operational readiness.

Acknowledgements

The study was funded by the National Research Foundation of Ukraine in the framework of the research project 2023.03/0197 on the topic "Multidisciplinary study of the impact of emergency situations on the infectious diseases spreading to support management decision-making in the field of population biosafety".

Declaration on Generative AI

During the preparation of this work, the authors used GPT-5-mini and Grammarly in order to: grammar and spelling check and text polishing. After using these services, the authors reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] Thomas, M.; Bomar, P.A. Upper Respiratory Tract Infection Available online: https://www.ncbi.nlm.nih.gov/books/NBK532961 (accessed on 5 August 2025).

[2] Sirota, S.B.; et al. Global, Regional, and National Burden of Upper Respiratory Infections and Otitis Media, 1990–2021: A Systematic Analysis from the Global Burden of Disease Study 2021. The Lancet Infectious Diseases 2024, 25, 36–51, doi: 10.1016/s1473-3099(24)00430-4.

[3] Cui, C.; Timbrook, T.T.; Polacek, C.; Heins, Z.; Rosenthal, N.A. Disease Burden and High-Risk Populations for Complications in Patients with Acute Respiratory Infections: A Scoping Review. Frontiers in Medicine 2024, 11, 1325236, doi: 10.3389/fmed.2024.1325236.

[4] European Centre for Disease Prevention and Control Seasonal Influenza - Annual Epidemiological Report for 2023/2024 Available online: https://www.ecdc.europa.eu/en/publications-data/seasonal-influenza-annual-epidemiological-report-20232024 (accessed on 5 August 2025).

[5] Haque, U.; et al. A Comparison of Ukrainian Hospital Services and Functions before and during the Russia-Ukraine War. JAMA Health Forum 2024, 5, e240901, doi: 10.1001/jamahealthforum.2024.0901.

[6] World Health Organization WHO Records More than 1000 Attacks on Health Care in Ukraine over the Past 15 Months of Full-Scale War Available online: https://www.who.int/europe/news/item/30-05-2023-who-records-1-000th-attack-on-health-care-in-ukraine-over-the-past-15-months-of-full-scale-war (accessed on 5 August 2025).

[7] Humanitarian Action Ukraine Winter Response Plan (October 2025 – March 2026) | Ukraine Winter Response Plan (October 2025 – March 2026) | Humanitarian Action Available online: https://humanitarianaction.info/document/ukraine-winter-response-plan-october-2025-march-2026/article/ukraine-winter-response-plan-october-2025-march-2026 (accessed on 5 August 2025).

[8] Chumachenko, D.; Butkevych, M.; Lode, D.; Frohme, M.; Schmailzl, K.J.G.; Nechyporenko, A. Machine Learning Methods in Predicting Patients with Suspected Myocardial Infarction Based on Short-Time HRV Data. Sensors 2022, 22, 7033, doi: 10.3390/s22187033.

[9] Mashtalir, V.P.; Yakovlev, S.V. Point-Set Methods of Clusterization of Standard Information. Cybernetics and System Analysis, 2001, vol. 37, pp. 295-307, doi: 10.1023/A:1011985908177.

[10] Mohammadi, A.; et al. Comparative Study of Linear Regression and SIR Models of COVID-19 Propagation in Ukraine before Vaccination. Radioelectronic and Computer Systems 2021, 2021, 5–18, doi: 10.32620/reks.2021.3.01.

[11] Chumachenko, D.; Piletskiy, P.; Sukhorukova, M.; Chumachenko, T. Predictive Model of Lyme Disease Epidemic Process Using Machine Learning Approach. Applied Sciences 2022, 12, 4282, doi: 10.3390/app12094282.

[12] Jiao, S.; Wang, Y.; Ye, X.; Nagahara, L.; Sakurai, T. Spatio-Temporal Epidemic Forecasting Using Mobility Data with LSTM Networks and Attention Mechanism. Scientific Reports 2025, 15, 9603, doi: 10.1038/s41598-025-94089-9.

[13] Chumachenko, D.; Bazilevych, K.; Butkevych, M.; Meniailov, I.; Parfeniuk, Y.; Sidenko, I.; Chumachenko, T. Methodology for Assessing the Impact of Emergencies on the Spread of Infectious Diseases. Radioelectronic and Computer Systems 2024, 2024, 6–26, doi: 10.32620/reks.2024.3.01.

[14] Taylor, K.S.; Taylor, J.W. Interval Forecasts of Weekly Incident and Cumulative COVID-19 Mortality in the United States: A Comparison of Combining Methods. PLOS ONE 2022, 17, e0266096, doi: 10.1371/journal.pone.0266096.

[15] Ramaswamy, A. Gradient Clipping in Deep Learning: A Dynamical Systems Perspective. Proceedings of the 12th International Conference on Pattern Recognition Applications and Methods 2023, 107–114, doi: 10.5220/0011678000003411.

[16] Mienye, I.D.; Swart, T.G.; Obaido, G. Recurrent Neural Networks: A Comprehensive Review of Architectures, Variants, and Applications. Information 2024, 15, 517, doi: 10.3390/info15090517.

[17] Green, R.; Stevens, G.; Abdallah, Z.S.; Filho, S. Stratify: Unifying Multi-Step Forecasting Strategies. Data Mining and Knowledge Discovery 2025, 39, 64, doi: 10.1007/s10618-025-01135-1.

[18] Hewamalage, H.; Ackermann, K.; Bergmeir, C. Forecast Evaluation for Data Scientists: Common Pitfalls and Best Practices. Data Mining and Knowledge Discovery 2022, 37, 788–832, doi: 10.1007/s10618-022-00894-5.

[19] Hao, C.; Zhao, Z.; Zhang, P.; Wu, B.; Ren, H.; Wang, X.; Qiao, Y.; Cui, Y.; Qiu, L. Application of Improved Harmonic Poisson Segmented Regression Model in Evaluating the Effectiveness of Kala-Azar Intervention in Yangquan City, China. Frontiers in Public Health 2024, 12, 1326225, doi: 10.3389/fpubh.2024.1326225.

[20] Hu, W.-H.; Sun, H.-M.; Wei, Y.-Y.; Hao, Y.-T. Global Infectious Disease Early Warning Models: An Updated Review and Lessons from the COVID-19 Pandemic. Infectious Disease Modelling 2024, 10, 410–422, doi: 10.1016/j.idm.2024.12.001.

[21] Stankeviciute, K.; Alaa, A.; van der Schaar, M. Conformal Time-Series Forecasting. Advances in Neural Information Processing Systems 2021, 34, 6216–6228.

[22] Tsang, T.K.; Du, Q.; Cowling, B.J.; Viboud, C. An Adaptive Weight Ensemble Approach to Forecast Influenza Activity in an Irregular Seasonality Context. Nature Communications 2024, 15, 8625, doi: 10.1038/s41467-024-52504-1.