Mathematical Biosciences and Engineering

https://www.aimspress.com/journal/mbe

MBE, 22(10): 2761–2779.
DOI: 10.3934/mbe.2025101
Received: 08 March 2025
Revised: 30 June 2025
Accepted: 28 August 2025
Published: 04 September 2025

Research article

Forecasting infectious disease outbreak risks from vaccine sentiments on social media: A data-driven dynamical systems approach

Zitao He and Chris T. Bauch*

Department of Applied Mathematics, University of Waterloo, Waterloo, ON, N2L 3G1 Canada

* **Correspondence:** Email: cbauch@uwaterloo.ca; Tel: +1-519-888-4567 ext. 32250.

Abstract: Early warning signals are vital in predicting critical transitions in complex dynamical systems. For behavioral epidemiology systems in particular, this includes shifts in vaccine sentiments that may precede disease outbreaks. Conventional statistical indicators, such as variance and lag-1 autocorrelation, often struggle in noisy environments and may fail in real-world scenarios. In this study, we leveraged universal signals of critical slowing down to train deep learning classifiers, specifically using long short-term memory (LSTM) and residual neural network (ResNet) architectures, for detecting early warning signals in disease-related social media time series. These classifiers were trained on simulated data from a stochastic coupled behavior-disease model with additive Lévy noise, a non-Gaussian noise that better reflects the heavy-tailed nature of real-world fluctuations. Our results show that these classifiers consistently outperform conventional indicators in both sensitivity and specificity on theoretical data while delivering quantitatively clear results that are easier to interpret on empirical data. Integrating deep learning with real-time social media monitoring offers a powerful tool for preventing disease outbreaks through proactive public health interventions.

Keywords: deep learning; early warning signals; behavioral epidemiology; vaccination; social media

1. Introduction

Before the development of vaccines, infectious diseases like measles were widespread and often resulted in significant morbidity and mortality, particularly among children [1]. The introduction of vaccines drastically reduced the incidence of these diseases, saving countless lives and alleviating the burden on healthcare systems [2]. Vaccine coverage is crucial not only for individual protection but also for establishing herd immunity, which prevents the spread of diseases within communities, including to those who cannot be vaccinated due to medical reasons or choose not to get vaccinated because of personal interests [3, 4]. When the vaccination rate drops below the herd immunity threshold, the risk of disease transmission increases, and any trigger event, such as an imported case, can potentially lead to a large outbreak. Between 2000 and 2013, California saw a concerning decline in the rate of measles, mumps, and rubella (MMR) vaccination [5,6]. This drop in vaccination coverage resulted in the loss of herd immunity, culminating in a significant measles outbreak at Disneyland, in California, in late 2014 [7–9]. This event underscored the critical need for maintaining high vaccination rates to prevent the resurgence of preventable diseases.

Monitoring vaccination rates is crucial. In the United States, the Centers for Disease Control and Prevention (CDC) oversees vaccination data collection through the National Immunization Survey (NIS) [10]. This annual survey gathers information via telephone from households with children and adolescents, cross-referencing data with healthcare providers' records. In Canada, the Public Health Agency of Canada (PHAC) collaborates with Statistics Canada to collect national vaccination data, conducting the Childhood National Immunization Coverage Survey (cNICS) every two years since 1994 [11]. In England, Public Health England (PHE) manages childhood vaccination data collection through the Cover of Vaccination Evaluated Rapidly (COVER) program, which operates on an annual basis [12]. However, annual or biennial vaccination data collection has notable limitations concerning bias and timeliness [13]. Response bias can significantly affect the reliability of such data, as those who choose to participate in surveys may not accurately represent the broader population, potentially skewing the results. Timeliness is another critical issue. While these surveys provide periodic updates, they may not capture rapid changes or emerging trends in vaccination rates between survey periods. This gap can delay the identification of vaccination coverage issues and the implementation of timely public health interventions. Consequently, the data might not reflect the current state of vaccination coverage, hindering efforts to address potential outbreaks promptly.

Social media data offers a high degree of timeliness, as it is generated and updated instantaneously, providing immediate insights into public behaviors, sentiments, and discussions. This real-time nature makes social media an attractive tool for infectious disease surveillance [14–16]. Public health authorities can quickly detect emerging health concerns [17] and monitor changes in public attitudes toward vaccinations [18,19] by analyzing social media activities, such as posts, shares, and comments. Recent advancements have also shown significant promise in incorporating social media into infectious disease modeling [20]. Social media plays a crucial role in altering people's behaviors by raising awareness of the disease, which generally helps to slow its spread. For example, when the number of infections increases and is reported on social media, susceptible individuals become more cautious and take protective measures, leading to a decrease in the transmission rate between susceptible and infected individuals, which is often modeled by using an exponentially decreasing function [21–23]. Some models also introduce separate compartments for aware and unaware individuals and model the transitions among them [24].

As researchers explore ways to incorporate social media into mathematical models, the potential for enhancing early warning systems becomes increasingly promising [25]. Critical transitions and early warning signals (EWS) are fundamental concepts in understanding and forecasting sudden changes in complex systems [26,27]. Critical transitions occur at moments when a dynamical system reaches a bifurcation point, leading to sudden and often irreversible shifts. EWS, on the other hand, typically precede these transitions and can provide notice of impending changes, potentially enabling the mitigation of adverse outcomes. Specific EWS, such as increased variance and lag-1 autocorrelation, have demonstrated reliability in indicating the emergence of diseases in epidemic models [28,29].

Deep learning (DL) excels at time series classification due to its ability to learn complex patterns within sequential data [30,31]. Research has demonstrated that deep learning models can predict different types of bifurcations in time series data for both continuous- and discrete-time models in physiology, economics, and ecology [32–34]. These deep learning models can recognize subtle changes in the time series data that precede critical transitions by learning generic features of bifurcations from large simulated training datasets. This capability is particularly useful in the context of disease outbreaks, where early detection of critical transitions can significantly improve response times and outcomes.

Certain EWS (e.g., variance, lag-1 autocorrelation) have been observed in vaccine-related Twitter data and Google trends prior to major measles outbreaks [35], indicating that changes in public vaccine sentiments can precede epidemiological events. However, traditional EWS methods typically assume white noise [36], whereas social media dynamics often follow non-Gaussian noise due to occasional high-impact events, such as viral posts from celebrities or influencers. These rare but large fluctuations create heavy-tailed distributions that can distort standard EWS metrics. To explain the phenomenon of "early discussion" and also better model the extreme events on social media, researchers have developed a stochastic coupled behavior-disease model with additive Lévy noise that links social media activity with disease dynamics [37]. This model suggests that heightened online discussion about vaccines and infectious diseases can serve as a signal for shifts in vaccine uptake and disease transmission. Building on these findings, there is promising potential for using deep learning models to detect EWS for vaccine uptake collapse in social media data. These models can learn to recognize patterns suggesting declining vaccination rates by analyzing large datasets, thereby providing timely alerts and enabling public health interventions to prevent outbreaks.

The growing field of behavioral epidemiology emphasizes the importance of incorporating human behavior into the modeling and control of infectious diseases [38–41]. For example, vaccination as an effective measure of preventing infectious diseases has been extensively studied using evolutionary game theory to account for individual decision-making [42–48]. Over time, this line of research has considered factors such as vaccine scares [49,50], social networks [51,52], social norms [53], and homophily [37], leading to more realistic models of vaccination dynamics. These increasingly complex models not only capture key behavioral mechanisms but also offer an opportunity: they can be used to generate synthetic data for training deep learning models. This idea, previously suggested by [54], is extended in this study to contribute new tools for the preparedness and mitigation of potential outbreaks.

In this article, we present a behavioral epidemiological method for using deep learning classifiers to provide EWS for potential vaccine uptake collapse. We train these classifiers using simulated data from a stochastic coupled behavior-disease model based on [37] and test their performance on real-world Twitter data collected before several major measles outbreaks in North America. We evaluate the classifiers' performance relative to traditional EWS indicators such as variance and lag-1 autocorrelation. In addition, we vary the noise amplitude to assess the robustness of these EWS methods.

2. Methods

2.1. Deep learning model framework

We used two deep learning algorithms. The first was a CNN-LSTM framework similar to [32–34], consisting of a 1D convolutional layer with dropout, a max pooling layer, and two LSTM layers with dropout, followed by a dense layer (Figure 1A).

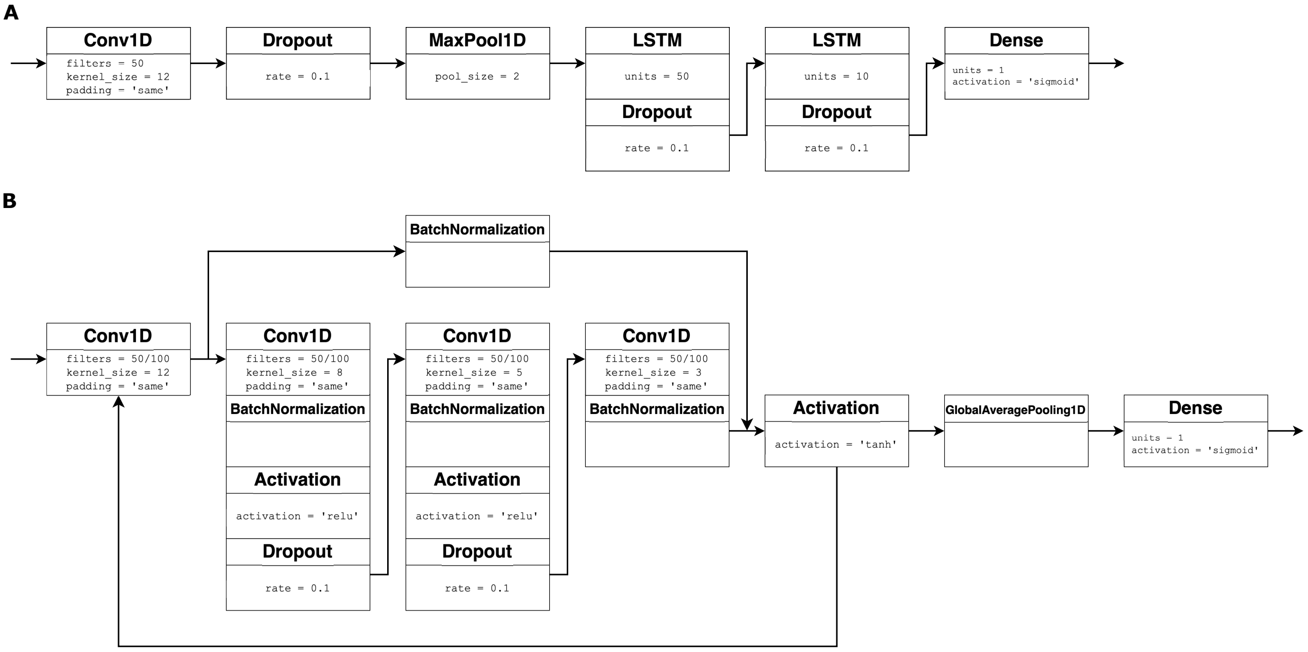

Figure 1. Two deep learning frameworks: (A) CNN-LSTM, (B) ResNet.

The second algorithm applied ResNet [55], as inspired by [30] (Figure 1B). Each residual block consists of three convolutional layers with progressively smaller kernel sizes, enabling the model to capture different levels of abstraction. These layers are followed by batch normalization, ReLU activation, and dropout layers to prevent overfitting. The input to the block is added directly to the output, allowing the model to retain information from earlier layers. The overall architecture includes three such residual blocks. The model concludes with global average pooling and a dense layer, which produces the final output using a sigmoid activation.

Both models were compiled using the Adam optimizer with binary cross-entropy loss and trained for 300 epochs at a learning rate of 0.0005 and a batch size of 1024. All code was written in Python using TensorFlow 2.12.0 and executed in Jupyter Notebook. Notably, the ResNet model contains approximately 20 times more trainable parameters than the CNN-LSTM model and requires a much longer time to train. The CNN-LSTM model achieved precision scores of 86.3% and 99.3% and recall scores of 99.4% and 84.2% for positive and negative cases, respectively, using an equal number of instances for each category. In comparison, the ResNet model achieved precision scores of 96.2% and 80.3% and recall scores of 76.2% and 97.0% for positive and negative cases. The overall accuracy was 91.8% for the CNN-LSTM model and 86.6% for the ResNet model.

2.2. *Theoretical data for training*

We applied a coupled behavior-disease model that incorporates two different social groups (social media users and non-users) [37] with an exponentially decreasing media function [21–23]:

$$\frac{dS}{dt} = \mu(1-x) - \beta e^{-mI} S I - \mu S, \tag{2.1}$$

$$\frac{dI}{dt} = \beta e^{-mI} S I - \gamma I - \mu I, \tag{2.2}$$

$$\frac{dR}{dt} = \mu x + \gamma I - \mu R, \tag{2.3}$$

$$\begin{aligned}\frac{dx_1}{dt} &= \kappa x_1(1-x_1)[-\omega + I + \delta(2x_1 - 1 + (1-h)(2x_2-1))] \\ &\quad + (1-h)\kappa\Big\{x_2(1-x_1)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0] \\ &\quad - (1-x_2)x_1\max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\},\end{aligned} \tag{2.4}$$

$$\begin{aligned}\frac{dx_2}{dt} &= \kappa x_2(1-x_2)[-\omega + I + \delta((1-h)(2x_1-1) + 2x_2 - 1)] \\ &\quad + (1-h)\kappa\Big\{x_1(1-x_2)\max[-\omega + I + \delta(2-h)(x_1+x_2-1), 0] \\ &\quad - (1-x_1)x_2\max[\omega - I - \delta(2-h)(x_1+x_2-1), 0]\Big\},\end{aligned} \tag{2.5}$$

where S, I, and R are the proportions of susceptible, infectious, and recovered individuals, respectively, with $S+I+R = 1$. x_1 (respectively, x_2) is the proportion of social media users (respectively, non-social media users) that are pro-vaccine. The overall pro-vaccine percentage in the population is the weighted average $x = (x_1N_1 + x_2N_2)/(N_1 + N_2)$. The total population size was set to $N = N_1 + N_2 = 200{,}000$, and the initial social media user ratio N_1/N was drawn from $\mathcal{U}[0.05, 0.4]$. We assumed that the number of social media users N_1 had an approximate annual growth rate of 5%. We also assumed that individuals sample others at a rate κ and adopt the sampled strategy with a probability proportional to the payoff difference, provided it is positive [42]. The basic reproduction number R_0 of measles was drawn from the uniform distribution $\mathcal{U}[12, 18]$ [3, 56], the recovery rate γ was drawn from $\mathcal{U}[365/22, 365/13]$ year^{-1} [35,53], the natural birth/death rate μ was drawn from $\mathcal{U}[1/70, 1/50]$ year^{-1}, and the homophily level h was drawn from $\mathcal{U}[0.4, 0.8]$ [37]. For the remaining parameters, the transmission rate was calculated as $\beta = R_0(\gamma + \mu)$ [3], the sampling rate was set to $\kappa = 3000$ year^{-1} [35, 37], the social norm strength was set to $\delta = 3 \times 10^{-4}$ [35,37,53], and $m = 0.6$ quantified the impact of the media coverage, indicating a relatively weak behavioral response to increases in infection prevalence [21, 57]. This coupled model captures a population of social media users and non-users, where individuals imitate others' vaccination strategies both within and across social groups, but tend to interact more within their own group due to homophily [58].

The deterministic model was converted to a stochastic model driven by an α-stable Lévy noise,

$$d\mathbf{X} = \mathbf{f}(\mathbf{X}, t)dt + \boldsymbol{\sigma} dL, \tag{2.6}$$

where $\mathbf{X} = [S, I, x_1, x_2]$. $\mathbf{f}(\mathbf{X}, t)$ is a column vector of the right-hand sides of Eqs 2.1, 2.2, 2.4, and 2.5. $\boldsymbol{\sigma} = [\sigma_S, \sigma_I, \sigma_{x_1}, \sigma_{x_2}] = [10^{-3}, 2\times10^{-6}, 4\times10^{-3}, 10^{-3}]$ is a column vector of scaling factors controlling the magnitude of noise for S, I, x_1, and x_2. dL is the additive α-stable Lévy process. To preserve the total population in the stochastic model, we impose the constraint $\sigma_R = -\sigma_S - \sigma_I$, which ensures that the Lévy noise terms do not alter the overall population size. The number of pro-vaccine social media posts was estimated by

$$T_p = \lambda[cx_1(1-x_1) + \tau I] + \xi. \tag{2.7}$$

Here, the scaling factor $\lambda = 0.2$, while c and τ were drawn from $\mathcal{U}[10, 200]$ and $\mathcal{U}[50{,}000, 100{,}000]$, respectively. ξ is an additive noise that follows a skewed fat-tailed stable distribution, with an expected value of $(1-\lambda)[cx_1(1-x_1) + \tau I]$. The stochastic model was simulated using a self-developed Monte-Carlo algorithm. The additive Lévy noise term dL was generated using the *levy_stable* function from the Python module *scipy.stats* with parameters $\alpha = 1.5$ and $\beta = 0$; extremely large jumps with magnitude larger than 200 were disregarded. The stochastic term ξ was also generated using *levy_stable* with parameters dynamically adjusted based on the baseline level $cx_1(1-x_1)+\tau I$. To reflect the variability in user engagement, T_p was further perturbed using a probabilistic filter that randomly reduced its value based on posting likelihood. Finally, all solutions to Eq 2.6 were clipped to lie within the physical bounds $[0, 1]$ [37].

We modeled the bifurcation parameter, the relative vaccine risk ω, as a linearly increasing function $\omega = (2 + 0.5t) \times 10^{-4}$. A critical transition was defined heuristically when the deviation of x from the equilibrium was greater than 20 times the noise amplitude $\sigma_x = (\sigma_{x_1} N_1 + \sigma_{x_2} N_2)/(N_1 + N_2)$, indicating a likely escape from the basin of attraction [33]. Such a transition should occur before x drops below 95%, the approximate herd immunity threshold for measles [4]. Upon detecting a critical transition, the 500 data points of T_p immediately preceding the transition were classified as a "positive" series; if no critical transition was observed, the time series was disregarded. A "negative" (or "neutral") series was generated by setting ω just below the estimated bifurcation point $\delta(2 - h)$ [37] and collecting the final 500 data points. Each series was then smoothed using the Python package *ewstools* [59] with a Gaussian kernel at a bandwidth of 100 data points. The residual and the smoothed time series together form a two-dimensional sample. A balanced set of 10,000 samples of each type was used for training.

2.3. *Empirical data for testing*

We collected English Tweets containing the keywords "measles" or "MMR" from January 1, 2011, to December 31, 2019, using the Twitter API. After cleaning the text (e.g., removing usernames, URLs, non-English characters, and converting emojis to text), we used the Python package *geopy* and a modified version of the *geodict* library [60] to geocode location information. Deep learning has been widely used for sentiment analysis of vaccine-related content on social media [61,62]. Here, the cleaned and geocoded Tweets were then classified as pro-vaccine, anti-vaccine, or neutral using COVID-Twitter-BERT [63]. Our analysis focused on four major measles outbreaks in the United States and Canada (California 2014–15, New York 2018–19, Washington 2018–19, and British Columbia 2014) [7,64–66] as well as four countries with mandatory measles vaccination policies and low incidence rates between 2014 and 2018 (Singapore, Mexico, Argentina, and Brazil) [67–70]. Pro-vaccine Tweets from the selected regions and time frames were smoothed using *ewstools* with a Gaussian kernel at a bandwidth of 100 data points. The variance and lag-1 autocorrelation of the residual time series were then computed using a rolling window of 200 data points.

3. Results

3.1. *Performance of EWS and DL on theoretical data*

We simulate sample positive and negative (or neutral) time series with different noise amplitudes. Time series with small amplitudes exhibit minor oscillations in the pro-vaccine percentage (Figures 2A and C), while those with large amplitudes display more pronounced oscillations (Figures 2B and D). We monitor two conventional EWS, variance and lag-1 autocorrelation, of the time series data on pro-vaccine percentage (Figures 2E–L) and the number of pro-vaccine social media posts (Figures 3E–L). EWS indicators are expected to show increasing trends as systems approach critical transitions. Con-

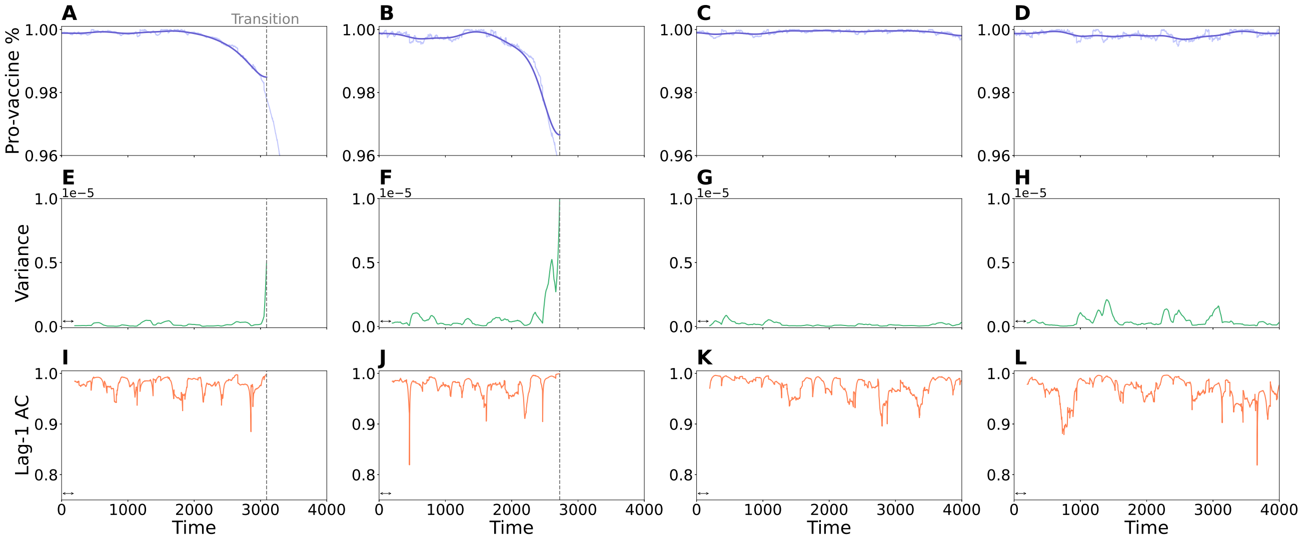

Figure 2. Trends in two conventional EWS (variance and lag-1 autocorrelation) applied to simulated time-series data. (A–D) Pro-vaccine percentage x in the entire population (light purple) and its smoothed trend (dark purple) for different cases: (A) a positive sample with small noise amplitudes, (B) a positive sample with large noise amplitudes, (C) a neutral sample with small noise amplitudes, and (D) a neutral sample with large noise amplitudes. (E–H) Variance of x. (I–L) Lag-1 autocorrelation of x. In (A, B), grey dashed lines mark the transition time.

sistent with this, we observe strong increasing trends in variance for both the pro-vaccine percentage (Figures 2E and F) and the number of pro-vaccine social media posts in positive samples (Figures 3E and F). However, lag-1 autocorrelation does not exhibit a similar upward trend (Figures 2I and J, Figures 3I and J).

The outputs of two deep learning classifiers applied to the number of pro-vaccine social media posts are also shown (Figures 3M–P). The probabilities assigned by the classifiers are smoothed using a Gaussian kernel to reveal an overall trend. For positive cases, both frameworks predict probabilities near zero for a long time before sharply increasing, effectively signaling the onset of a transition. In particular, LSTM transitions earlier than ResNet, indicating greater sensitivity to early signals, and a faster response. For negative cases, ResNet consistently predicts probabilities close to zero throughout the time, whereas LSTM produces several spikes and oscillations following sudden increases in the number of pro-vaccine posts.

The performance of these indicators is assessed using receiver operating characteristic (ROC) curves (Figure 4). The upward trends in variance and lag-1 autocorrelation are quantified by calculating the Kendall tau coefficient, while the final data point in the smoothed probabilities assigned by the deep learning classifiers represents their prediction outcome. We simulate a balanced set of 500 positive samples and 500 neutral samples across five different noise amplitudes. The statistics are generated over the entire pre-transition period, as well as on partial pre-transition periods up to 100, 200, and 400 time points before the transition. For neutral samples, the "pre-transition" period is defined as the entire time series, as no transition will occur. Using the area under the ROC curve (AUC) score, we

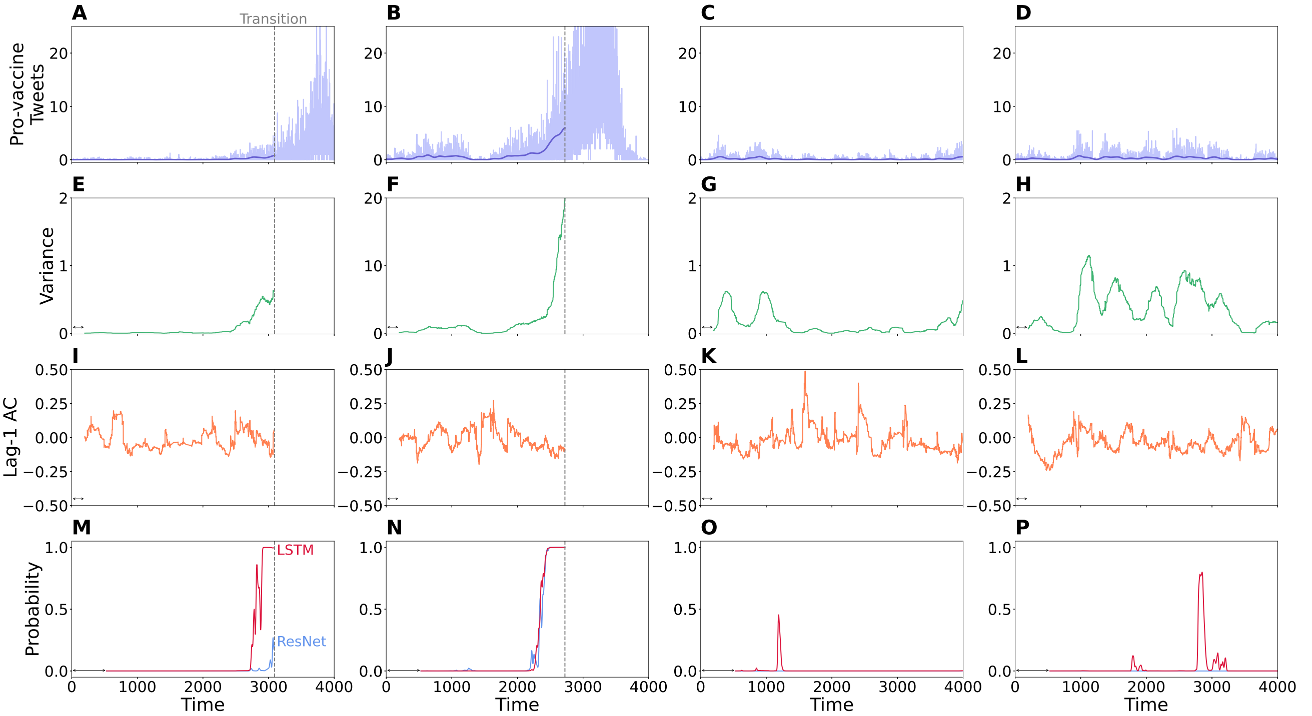

Figure 3. Trends in two conventional EWS (variance and lag-1 autocorrelation) and two deep learning classifiers (ResNet and LSTM) applied to simulated time-series data. (A–D) Corresponding number of pro-vaccine Tweets T_p (light purple) and its smoothed trend (dark purple) of Figures 2A–D. (E–H) Variance of T_p. (I–L) Lag-1 autocorrelation of T_p. (M–P) Probabilities assigned by two deep learning classifiers. In (A, B), grey dashed lines mark the transition time.

find that variance on T_p is the best-performing conventional EWS, while lag-1 autocorrelation on T_p performs the worst (Figures 4E–H). Both indicators on x exhibit weak predictive ability (Figures 4A–D). Deep learning classifiers consistently outperform both variance and lag-1 autocorrelation on x and T_p at all cut-off times in the pre-transition period (Figures 4I–L). When evaluated over the entire pre-transition period, both classifiers achieve near-perfect performance, with ROC curves approaching the top-left corner and AUC scores close to 1, indicating high sensitivity and specificity. However, their performance gradually declines for both classifiers as the cut-off time moves further from the transition, with LSTM generally outperforming ResNet. This trend reverses at the 400-time-point cut-off, where LSTM's performance drops significantly, particularly for negative samples.

3.2. *Performance of EWS and DL on empirical data*

To test our deep learning classifiers on empirical data, we select four measles outbreaks in North America between 2013 and 2020 as positive samples: the 2014–15 outbreak in California, the 2018–19 outbreak in New York, the 2018–19 outbreak in Washington, and the 2014 outbreak in British Columbia, Canada (Figures 5A–D). For negative (or neutral) samples, we choose four countries (Singapore, Mexico, Argentina, and Brazil) that have mandatory measles vaccination policies and reported very few measles cases from 2014 to 2018 (Figures 6A–D).

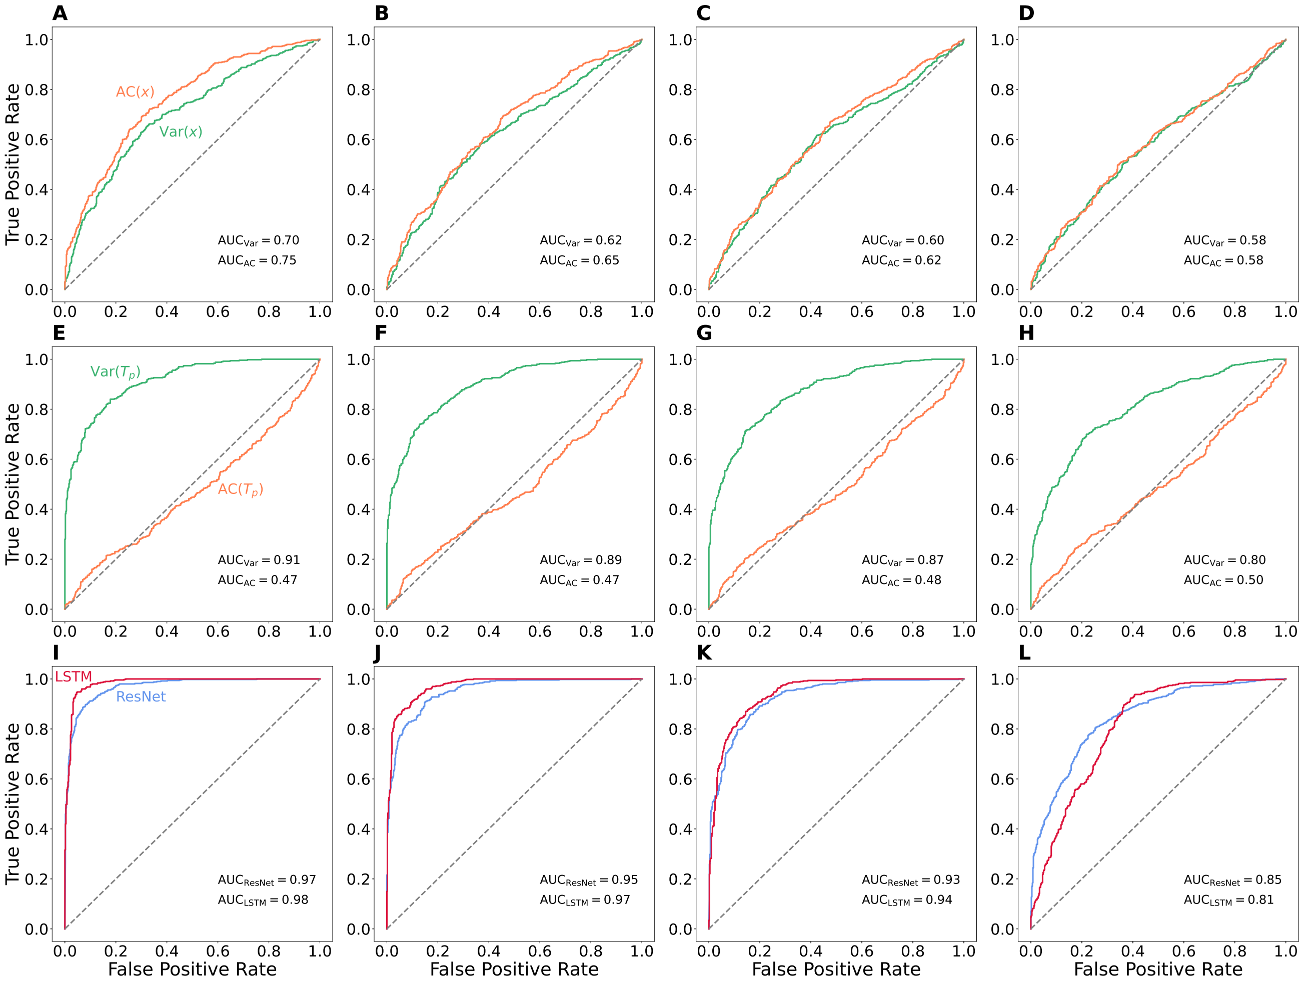

Figure 4. ROC curves for predictions of an upcoming critical transition in theoretical time-series data. (A–D) ROC curves based on Kendall tau values for variance and lag-1 autocorrelation of x. (E–H) ROC curves based on Kendall tau values for variance and lag-1 autocorrelation of T_p. (I–L) ROC curves for probabilities assigned by deep learning classifiers. Performance is evaluated on 500 positive and 500 neutral simulations. Panels correspond to different pre-transition periods used for prediction: (A, E, I) the entire pre-transition period, (B, F, J) up to 100 time points before the transition, (C, G, K) up to 200 time points before the transition, and (D, H, L) up to 400 time points before the transition. Grey dashed lines represent the chance level.

Before the three outbreaks across the US (CA14, NY18, and WA19), both the ResNet and LSTM classifiers show a clear increase in probability predictions, signaling a potential critical transition in the pro-vaccine percentage among the population (Figures 5M–O). Interestingly, while LSTM consistently responds faster than ResNet in the theoretical data, we find that ResNet is more sensitive to early signals in the empirical data. LSTM, however, proves effective in filtering out "false signals", such as sudden boosts in the number of Tweets caused by the reposting of news reports over a short period. In both CA14 and NY18 data, a small drop in the LSTM prediction occurs after it reaches near 100%

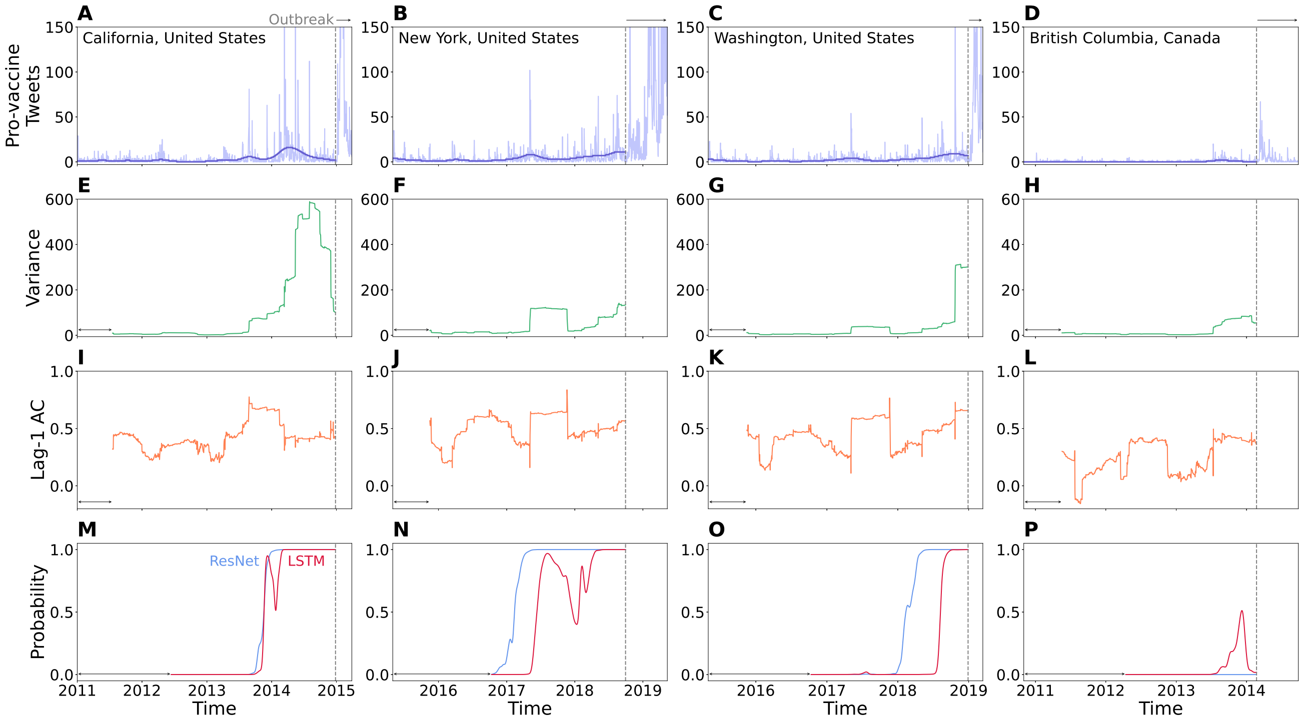

Figure 5. Trends in two conventional EWS and two deep learning classifiers applied to positive empirical Twitter data. (A–D) Number of pro-vaccine Tweets (light purple) and its smoothed trend (dark purple) before and during four selected outbreaks: (A) 2014–15 outbreak in California, (B) 2018–19 outbreak in New York, (C) 2018–19 outbreak in Washington, and (D) 2014 outbreak in British Columbia. Corresponding (E–H) variance, (I–L) lag-1 autocorrelation, and (M–P) probabilities assigned by two deep learning classifiers. Grey dashed lines indicate the start of each outbreak, with outbreak periods marked by arrows at the top of (A–D).

(Figures 5M and N), as these short-lived bursts of discussion did not lead to a sustained discussion over time. Consequently, no significant drop in the current pro-vaccine percentage is expected. The number of Tweets collected in British Columbia is much smaller, making it challenging for deep learning classifiers to capture patterns in the time series data. The ResNet classifier predicts a zero probability for a critical transition to occur, while LSTM shows an increasing probability to around 55% during a period of moderate online discussion before dropping back to zero (Figure 5P).

In the four selected countries where measles vaccination is mandatory (SG, MX, AR, and BR), the number of pro-vaccine Tweets we collected is relatively small. Both ResNet and LSTM predict low probabilities of critical transitions throughout the time range (Figures 6M–P). In early 2015, all four countries experienced a relatively large-scale discussion following the outbreak at Disneyland, California. LSTM captures this slight increase in MX, AR, and BR data. However, the probabilities drop quickly as the discussion does not last long. Notably, no major measles outbreaks have occurred in these countries since then.

Our results reveal a marked increase in variance prior to each selected outbreak (Figures 5E–H). Lag-1 autocorrelation has an overall upward trend, with rises that start earlier but have a smaller scale

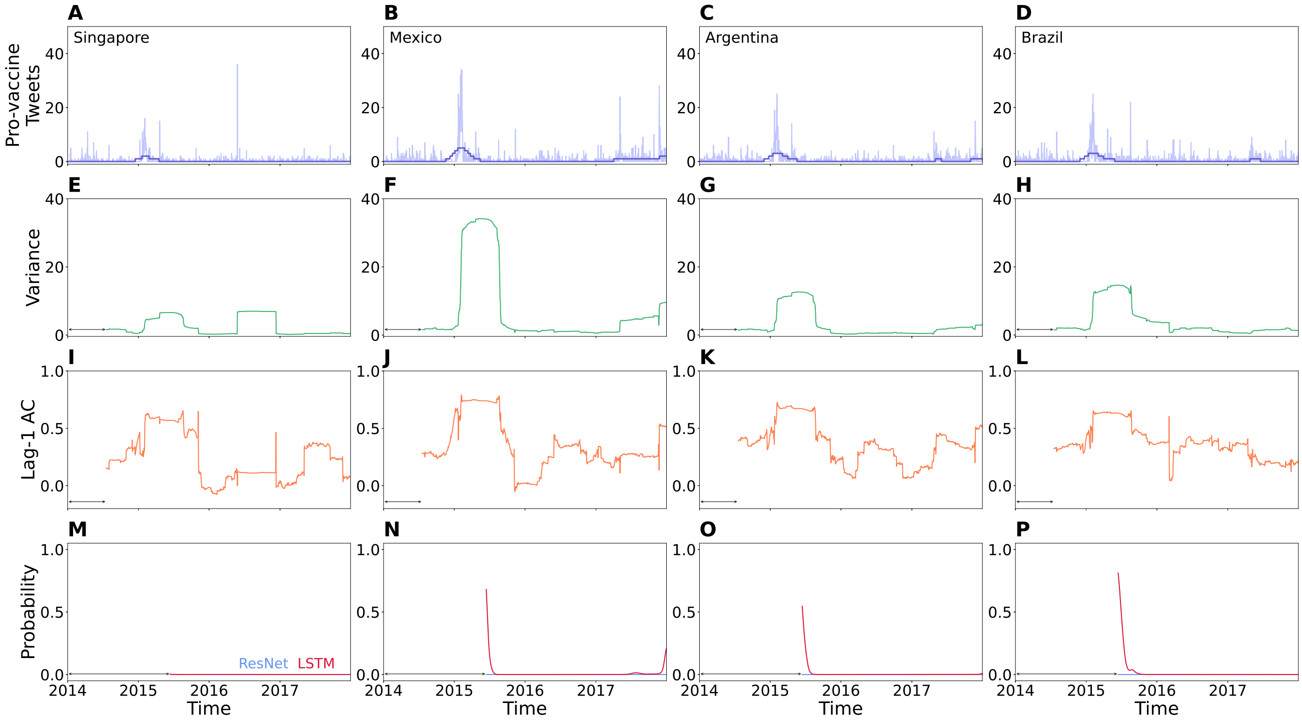

Figure 6. Trends in two conventional EWS and two deep learning classifiers applied to neutral empirical Twitter data. (A–D) Number of pro-vaccine Tweets (light purple) and its smoothed trend (dark purple) between 2014 and 2018 in four selected countries: (A) Singapore, (B) Mexico, (C) Argentina, and (D) Brazil. Corresponding (E–H) variance, (I–L) lag-1 autocorrelation, and (M–P) probabilities assigned by two deep learning classifiers.

(Figures 5I–L). For negative samples, both values exhibit only significant fluctuations when intense discussions are observed over a short period (early 2015); they remain relatively stable for the rest of the time (Figures 6E–L). In the case of variance, when we compare specific numerical values, these larger fluctuations are still much smaller compared to the values seen in positive cases.

4. Discussion

The coarse-scale nature of vaccination rate data poses challenges for public health departments in monitoring timely changes, thereby limiting their capacity to provide early warnings of potential declines. In contrast, social media provides high-frequency data on various topics, including vaccine-related sentiments. Our findings indicate that deep learning classifiers are adept at analyzing social media time series data to detect early warning signals of critical transitions in pro-vaccine sentiment. These classifiers not only exhibit superior sensitivity and specificity compared to conventional EWS, like variance and lag-1 autocorrelation in theoretical data, but also demonstrate robust performance when applied to empirical data, underscoring their practical utility in real-world scenarios.

Researchers have evaluated the performance of variance and lag-1 autocorrelation as EWS on a simple SEIRx (Susceptible-Exposed-Infectious-Recovered-Vaccinator) epidemic model [32]. In our study, we observe much lower AUC scores for both indicators on the pro-vaccine percentage, which we attribute to two key factors. First, unlike a model that broadly divides the entire population into pro-vaccine and anti-vaccine individuals, our model considers interactions within and between two distinct social groups. Additionally, the performance of two indicators is assessed not directly on a state variable but on a weighted average of two state variables with time-varying weights. These modifications introduce additional complexity to the system and can potentially distort the intrinsic statistical properties that conventional EWS rely on. Second, instead of Gaussian white noise, we apply fat-tailed α-stable Lévy noise that has an infinite variance when the stability parameter α is less than 2. This type of noise can produce large jumps that obscure the subtle and gradual changes typically expected near critical transitions, reducing the predictive power of variance. Furthermore, these large jumps can also disrupt the temporal relationship between a system's current state and its previous state, making lag-1 autocorrelation a less reliable indicator. The time series representing the number of pro-vaccine social media posts is more complex due to its derivation from transformed state variables through multiple functions, combined with additional Lévy noise. This complexity leads to poor performance in lag-1 autocorrelation, which performs even worse than random guessing. On the other hand, the high performance of variance is expected, as the scale parameter of the additive Lévy noise is assumed to increase as the pro-vaccine percentage decreases [37], and such a decrease is a necessary condition for the system to approach a tipping point. Despite the complexity of the data, our deep learning classifiers still serve as highly effective indicators. These findings highlight the advantage of using deep learning classifiers in capturing patterns in transformed and noisy data. This capability can be particularly valuable in real-world applications, as observational data often undergo multiple transformations and are influenced by non-Gaussian noise.

The performance of any deep learning model depends heavily on the quality of the training data [71]. For detecting early signals of pro-vaccine sentiment collapse on social media, the ideal training set should contain time series data from real-world social networks, paired with detailed vaccination data for the corresponding regions to determine whether a decline in vaccine coverage is occurring. However, collecting sufficient data from social media platforms is challenging, as most posts lack location information, and identifying vaccination rate declines from such coarse-scale data becomes even more difficult. Mathematical models provide an efficient way to simulate large amounts of training data; they also allow us to precisely control key conditions, such as whether a critical transition occurs and when it takes place. This approach has also been adopted in recent work, where theoretical epidemic models were used to generate synthetic data for training deep learning models to detect superspreading events [72]. Our approach uses a coupled behavior-disease model with two social groups as the backbone. We also introduce an exponential media function and randomize most of the parameters in order to make the data more realistic and diverse. However, it is important to acknowledge that our model relies on simplifications that may not fully capture the complexity of real-world social media behaviors. As a result, deep learning classifiers trained on those theoretical data could overfit to the synthetic patterns and struggle to generalize to the unpredictability of empirical data. To address this sim-to-real gap, a potential direction is to pre-train the model on simulated data, followed by fine-tuning on limited real-world data (a specific form of transfer learning) [73]. This procedure usually allows models to retain structural insights from simulation while adapting to empirical noise.

The complex dynamics of social media data require sufficient temporal context to model effectively. As a result, our approach is particularly well-suited for regions that have not experienced outbreaks for extended periods, during which vaccination rates are in decline. Ideally, testing samples should come from regions with outbreaks in the early 2010s, when the resurgence of measles began [74, 75]. However, since social media did not gain widespread popularity until the 2010s, few posts were collected prior to these outbreaks, making the time series data less informative. For outbreaks occurring in the mid-to-late 2010s, the preceding social media data was often influenced by outbreaks in other regions. To address this, we reduce the length of the time series to isolate a cleaner segment of data, minimizing the impact of sudden boosts caused by outbreaks elsewhere. Our choice of 500 data points, corresponding to approximately 1.5 years of daily data, strikes a balance between capturing annual trends and seasonal effects while maintaining computational feasibility. However, a longer time series would be preferable for detecting EWS [76]. To achieve a more accurate reflection of vaccination rate fluctuations, more detailed data preprocessing methods will be required to filter out the effects of outbreaks on social media data.

In the four countries with mandatory measles vaccination policies included in our study, the number of Tweets related to the measles vaccine was low for two main reasons. First, the policy has made vaccination a socially accepted norm, resulting in little public debate. Second, a large portion of the discussions occurred in languages other than English, since English is not an official or commonly spoken language in Mexico, Argentina, and Brazil, although it is one of Singapore's four official languages. Therefore, the English Tweets we collected represent only a small fraction of the total discussions in these countries, which introduces a selection bias and leads to underrepresentation. Future research could adopt multilingual approaches using language-specific sentiment models. This would also help expand the real-world dataset, which could support more effective transfer learning through pre-training and fine-tuning.

We proposed two widely used deep learning frameworks for time series classification. While each has its strengths, we do not consider one approach strictly better than the other. The LSTM model is more computationally efficient due to a much smaller number of parameters. It generally outperforms ResNet but tends to overreact to short-term fluctuations in some cases. The ResNet model is more conservative and predicts a high probability of critical transitions only when enough evidence is observed. We believe that combining insights from both models offers a more balanced and comprehensive perspective, leveraging the sensitivity of LSTM and the stability of ResNet for more robust analysis of time series data.

We hope this study provides new insights into predicting critical transitions in dynamical systems, and contributes to the growing field of behavioral epidemiology, particularly in the context of parental vaccination choices. Human behaviors introduce significant complexity and noise into these systems, often limiting the effectiveness of traditional statistical early warning indicators. Additionally, real-world data shortages can also make theoretically feasible models difficult to implement in practice. Our findings demonstrate that it is possible to bridge the gap between dynamical systems and data-rich sources like social media and let deep learning models extract meaningful patterns from available data and uncover early signals of critical transitions, even in noisy and complex environments. Integrating these models with real-time social media monitoring could serve as a powerful tool for proactive public health interventions, ultimately helping to reduce the risks of disease outbreaks.

Data availability

All code used in this study is available at `https://github.com/z255he/dl-ews-tweet`. The aggregate Twitter data used to generate the plots is provided. Other data can be generated using the provided Python scripts.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

This research was funded by an NSERC Discovery Grant to Chris T. Bauch.

Conflict of interest

The authors declare there is no conflict of interest.

References

1. Centers for Disease Control and Prevention, *History of measles*, 2024. Available from: `https://www.cdc.gov/measles/about/history.html`.
2. Centers for Disease Control and Prevention, *Measles, mumps, and rubella (MMR) vaccination: What everyone should know*, 2024. Available from: `https://www.cdc.gov/vaccines/vpd/mmr/public/index.html`.
3. R. M. Anderson, R. M. May, *Infectious diseases of humans: Dynamics and control*, Oxford University Press, 1991. https://doi.org/10.1093/oso/9780198545996.001.0001
4. A. N. Desai, M. S. Majumder, What is herd immunity? *JAMA*, **324** (2020), 2113. https://doi.org/10.1001/jama.2020.20895
5. California Department of Public Health, *School immunizations in kindergarten by academic year*, 2020. Available from: `https://data.chhs.ca.gov/dataset/school-immunizations-in-kindergarten-by-academic-year`.
6. P. Krishnakumar, S. Karlamangla, *How stricter vaccine laws spared California from a major measles outbreak*, 2019. Available from: `https://www.latimes.com/projects/la-me-measles-us-california-outbreak-vaccine-new-york-disneyland/`.
7. J. Zipprich, K. Winter, J. Hacker, D. Xia, J. Watt, K. Harriman, Measles outbreak — California, December 2014–February 2015, *MMWR. Morb. Mortal. Wkly. Rep.*, **64** (2015), 153–154. `http://www.ncbi.nlm.nih.gov/pmc/articles/pmc4584705/`
8. L. Worden, S. F. Ackley, J. Zipprich, K. Harriman, W. T. A. Enanoria, R. Wannier, et al., Measles transmission during a large outbreak in California, *Epidemics*, **30** (2020), 100375. https://doi.org/10.1016/j.epidem.2019.100375

9. M. K. Doll, J. W. Correira, Revisiting the 2014-15 Disneyland measles outbreak and its influence on pediatric vaccinations, *Hum. Vaccin. Immunother.*, **17** (2021), 4210–4215. https://doi.org/10.1080/21645515.2021.1972707

10. Centers for Disease Control and Prevention, *About the National Immunization Surveys (NIS)*, 2024. Available from: `https://www.cdc.gov/vaccines/imz-managers/nis/about.html`.

11. Government of Canada, *Vaccination coverage in Canada*, 2024. Available from: `https://www.canada.ca/en/public-health/services/immunization-vaccines/vaccination-coverage.html`.

12. UK Health Security Agency, *Cover of vaccination evaluated rapidly (COVER) programme: annual data*, 2015. Available from: `https://www.gov.uk/government/publications/cover-of-vaccination-evaluated-rapidly-cover-programme-annual-data`.

13. F. T. Cutts, P. Claquin, M. C. Danovaro-Holliday, D. A. Rhoda, Monitoring vaccination coverage: Defining the role of surveys, *Vaccine*, **34** (2016), 4103–4109. https://doi.org/10.1016/j.vaccine.2016.06.053

14. A. E. Aiello, A. Renson, P. N. Zivich, Social media- and Internet-based disease surveillance for public health, *Annu. Rev. Public Health*, **41** (2020), 101–118. https://doi.org/10.1146/annurev-publhealth-040119-094402

15. A. E. Wilson, C. U. Lehmann, S. N. Saleh, J. Hanna, R. J. Medford, Social media: A new tool for outbreak surveillance, *Antimicrob. Steward. Healthc. Epidemiol.*, **1** (2021), e50. https://doi.org/10.1017/ash.2021.225

16. P. Pilipiec, I. Samsten, A. Bota, Surveillance of communicable diseases using social media: A systematic review, *PLoS One*, **18** (2023), e0282101. https://doi.org/10.1371/journal.pone.0282101

17. E. Aramaki, S. Maskawa, M. Morita, Twitter catches the flu: Detecting influenza epidemics using Twitter, in *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing* (eds. R. Barzilay, M. Johnson), Association for Computational Linguistics, (2011), 1568–1576. `https://dl.acm.org/doi/10.5555/2145432.2145600`

18. M. Salathé, S. Khandelwal, Assessing vaccination sentiments with online social media: Implications for infectious disease dynamics and control, *PLoS Comput. Biol.*, **7** (2011), e1002199. https://doi.org/10.1371/journal.pcbi.1002199

19. S. Ajovalasit, V. M. Dorgali, A. Mazza, A. d'Onofrio, P. Manfredi, Evidence of disorientation towards immunization on online social media after contrasting political communication on vaccines. Results from an analysis of Twitter data in Italy, *PLoS One*, **16** (2021), e0253569. https://doi.org/10.1371/journal.pone.0253569

20. J. Sooknanan, D. M. G. Comissiong, Trending on social media: Integrating social media into infectious disease dynamics, *Bull. Math. Biol.*, **82** (2020), 86. https://doi.org/10.1007/s11538-020-00757-4

21. J. Cui, Y. Sun, H. Zhu, The impact of media on the control of infectious diseases, *J. Dyn. Diff. Equat.*, **20** (2007), 31–53. https://doi.org/10.1007/s10884-007-9075-0

22. Y. Xiao, T. Zhao, S. Tang, Dynamics of an infectious diseases with media/psychology induced non-smooth incidence, *Math. Biosci. Eng.*, **10** (2013), 445–461. https://doi.org/10.3934/mbe.2013.10.445

23. L. Mitchell, J. V. Ross, A data-driven model for influenza transmission incorporating media effects, *R. Soc. Open Sci.*, **3** (2016), 160481. https://doi.org/10.1098/rsos.160481

24. G. O. Agaba, Y. N. Kyrychko, K. B. Blyuss, Mathematical model for the impact of awareness on the dynamics of infectious diseases, *Math. Biosci.*, **286** (2017), 22–30. https://doi.org/10.1016/j.mbs.2017.01.009

25. E. Southall, T. S. Brett, M. J. Tildesley, L. Dyson, Early warning signals of infectious disease transitions: A review, *J. R. Soc. Interface*, **18** (2021), 20210555. https://doi.org/10.1098/rsif.2021.0555

26. M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, et al., Early-warning signals for critical transitions, *Nature*, **461** (2009), 53–59. https://doi.org/10.1038/nature08227

27. V. Dakos, S. R. Carpenter, W. A. Brock, A. M. Ellison, V. Guttal, A. R. Ives, et al., Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data, *PLoS One*, **7** (2012), e41010. https://doi.org/10.1371/journal.pone.0041010

28. T. S. Brett, J. M. Drake, P. Rohani, Anticipating the emergence of infectious diseases, *J. R. Soc. Interface*, **14** (2017), 20170115. https://doi.org/10.1098/rsif.2017.0115

29. T. S. Brett, E. B. O'Dea, É. Marty, P. B. Miller, A. W. Park, J. M. Drake, et al., Anticipating epidemic transitions with imperfect data, *PLoS Comput. Biol.*, **14** (2018), e1006204. https://doi.org/10.1371/journal.pcbi.1006204

30. Z. Wang, W. Yan, T. Oates, Time series classification from scratch with deep neural networks: A strong baseline, in *2017 International Joint Conference on Neural Networks (IJCNN)*, (2017), 1578–1585. https://doi.org/10.1109/IJCNN.2017.7966039

31. H. Ismail Fawaz, G. Forestier, J. Weber, L. Idoumghar, P.-A. Muller, Deep learning for time series classification: A review, *Data Min. Knowl. Dis.*, **33** (2019), 917–963. https://doi.org/10.1007/s10618-019-00619-1

32. T. M. Bury, R. I. Sujith, I. Pavithran, M. Scheffer, T. M. Lenton, M. Anand, et al., Deep learning for early warning signals of tipping points, *Proc. Natl. Acad. Sci. U.S.A.*, **118** (2021), e2106140118. https://doi.org/10.1073/pnas.2106140118

33. T. M. Bury, D. Dylewsky, C. T. Bauch, M. Anand, L. Glass, A. Shrier, et al., Predicting discrete-time bifurcations with deep learning, *Nat. Commun.*, **14** (2023), 6331. https://doi.org/10.1038/s41467-023-42020-z

34. D. Dylewsky, M. Anand, C. T. Bauch, Early warning signals for bifurcations embedded in high dimensions, *Sci. Rep.*, **14** (2024), 18277. https://doi.org/10.1038/s41598-024-68177-1

35. A. D. Pananos, T. M. Bury, C. Wang, J. Schonfeld, S. P. Mohanty, B. Nyhan, et al., Critical dynamics in population vaccinating behavior, *Proc. Natl. Acad. Sci. U.S.A.*, **114** (2017), 13762–13767. https://doi.org/10.1073/pnas.1704093114

36. R. Liu, P. Chen, K. Aihara, L. Chen, Identifying early-warning signals of critical transitions with strong noise by dynamical network markers, *Sci. Rep.*, **5** (2015), 17501. https://doi.org/10.1038/srep17501

37. Z. He, C. T. Bauch, Effect of homophily on coupled behavior-disease dynamics near a tipping point, *Math. Biosci.*, **376** (2024), 109264. https://doi.org/10.1016/j.mbs.2024.109264

38. S. Funk, M. Salathé, V. A. A. Jansen, Modelling the influence of human behaviour on the spread of infectious diseases: A review, *J. R. Soc. Interface*, **7** (2010), 1247–1256. https://doi.org/10.1098/rsif.2010.0142

39. P. Manfredi, A. d'Onofrio, *Modeling the interplay between human behavior and the spread of infectious diseases*, Springer, 2013. https://doi.org/10.1007/978-1-4614-5474-8

40. Z. Wang, C. T. Bauch, S. Bhattacharyya, A. d'Onofrio, P. Manfredi, M. Perc, et al., Statistical physics of vaccination, *Phys. Rep.*, **664** (2016), 1–113. https://doi.org/10.1016/j.physrep.2016.10.006

41. F. Verelst, L. Williem, P. Beutels, Behavioural change models for infectious disease transmission: A systematic review (2010–2015), *J. R. Soc. Interface*, **13** (2016), 20160820. https://doi.org/10.1098/rsif.2016.0820

42. C. T. Bauch, Imitation dynamics predict vaccinating behaviour, *Proc. R. Soc. B*, **272** (2005), 1669–1675. https://doi.org/10.1098/rspb.2005.3153

43. T. C. Reluga, C. T. Bauch, A. P. Galvani, Evolving public perceptions and stability in vaccine uptake, *Math. Biosci.*, **204** (2006), 185–198. https://doi.org/10.1016/j.mbs.2006.08.015

44. A. P. Galvani, T. C. Reluga, G. B. Chapman, Long-standing influenza vaccination policy is in accord with individual self-interest but not with the utilitarian optimum, *Proc. Natl. Acad. Sci. U.S.A.*, **104** (2007), 5692–5697. https://doi.org/10.1073/pnas.0606774104

45. A. d'Onofrio, P. Manfredi, E. Salinelli, Vaccinating behaviour, information, and the dynamics of SIR vaccine preventable diseases, *Theor. Popul. Biol.*, **71** (2007), 301–317. https://doi.org/10.1016/j.tpb.2007.01.001

46. E. Shim, G. B. Chapman, J. P. Townsend, A. P. Galvani, The influence of altruism on influenza vaccination decisions, *J. R. Soc. Interface*, **9** (2012), 2234–2243. https://doi.org/10.1098/rsif.2012.0115

47. X. Chen, F. Fu, Imperfect vaccine and hysteresis, *Proc. R. Soc. B*, **286** (2019), 20182406. https://doi.org/10.1098/rspb.2018.2406

48. S. L. Chang, M. Piraveenan, P. Pattison, M. Prokopenko, Game theoretic modelling of infectious disease dynamics and intervention methods: A review, *J. Biol. Dyn.*, **14** (2020), 57–89. https://doi.org/10.1080/17513758.2020.1720322

49. C. T. Bauch, S. Bhattacharyya, Evolutionary game theory and social learning can determine how vaccine scares unfold, *PLoS Comput. Biol.*, **8** (2012), e1002452. https://doi.org/10.1371/journal.pcbi.1002452

50. E. Shim, J. J. Grefenstette, S. M. Albert, B. E. Cakouros, D. S. Burke, A game dynamic model for vaccine skeptics and vaccine believers: Measles as an example, *J. Theor. Biol.*, **295** (2012), 194–203. https://doi.org/10.1016/j.jtbi.2011.11.005

51. F. Fu, D. I. Rosenbloom, L. Wang, M. A. Nowak, Imitation dynamics of vaccination behaviour on social networks, *Proc. R. Soc. B*, **278** (2010), 42–49. https://doi.org/10.1098/rspb.2010.1107
52. M. L. Ndeffo Mbah, J. Liu, C. T. Bauch, Y. I. Tekel, J. Medlock, L. A. Meyers, et al., The impact of imitation on vaccination behavior in social contact networks, *PLoS Comput. Biol.*, **8** (2012), e1002469. https://doi.org/10.1371/journal.pcbi.1002469
53. T. Oraby, V. Thampi, C. T. Bauch, The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases, *Proc. R. Soc. B*, **281** (2014), 20133172. https://doi.org/10.1098/rspb.2013.3172
54. R. E. Baker, J.-M. Peña, J. Jayamohan, A. Jérusalem, Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biol. Lett.*, **14** (2018), 20170660. https://doi.org/10.1098/rsbl.2017.0660
55. K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2016), 770–778. https://doi.org/10.1109/CVPR.2016.90
56. F. M. Guerra, S. Bolotin, G. Lim, J. Heffernan, S. L. Deeks, Y. Li, et al., The basic reproduction number (R_0) of measles: a systematic review, *Lancet Infect. Dis.*, **17** (2017), e420–e428. https://doi.org/10.1016/s1473-3099(17)30307-9
57. R. Liu, J. Wu, H. Zhu, Media/psychological impact on multiple outbreaks of emerging infectious diseases, *Comput. Math. Methods Med.*, **8** (2007), 612372. https://doi.org/10.1080/17486700701425870
58. M. McPherson, L. Smith-Lovin, J. M. Cook, Birds of a feather: Homophily in social networks, *Annu. Rev. Sociol.*, **27** (2001), 415–555. https://doi.org/10.1146/annurev.soc.27.1.415
59. T. M. Bury, Ewstools: A Python package for early warning signals of bifurcations in time series data, *J. Open Source Softw.*, **8** (2023), 5038. https://doi.org/10.21105/joss.05038
60. P. Warden, *Geodict: A simple Python library/tool for pulling location information from unstructured text*, 2010. Available from: `https://github.com/petewarden/geodict`.
61. İ. Aygün, B. Kaya, M. Kaya, Aspect based Twitter sentiment analysis on vaccination and vaccine types in COVID-19 Pandemic with deep learning, *IEEE J. Biomed. Health Inform.*, **26** (2022), 2360–2369. https://doi.org/10.1109/JBHI.2021.3133103
62. W. Alotaibi, F. Alomary, R. Mokni, COVID-19 vaccine rejection causes based on Twitter people's opinions analysis using deep learning, *Soc. Netw. Anal. Min.*, **13** (2023), 62. https://doi.org/10.1007/s13278-023-01059-y
63. M. Müller, M. Salathé, P. E. Kummervold, COVID-Twitter-BERT: A natural language processing model to analyse COVID-19 content on Twitter, *Front. Artif. Intell.*, **6** (2023), 1023281. https://doi.org/10.3389/frai.2023.1023281
64. R. McDonald, P. S. Ruppert, M. Souto, D. E. Johns, K. McKay, N. Bessette, et al., Notes from the field: Measles outbreaks from imported cases in orthodox Jewish communities — New York and New Jersey, 2018–2019, *MMWR. Morb. Mortal. Wkly. Rep.*, **68** (2019), 444–445. https://doi.org/10.15585/mmwr.mm6819a4

65. A. Carlson, M. Riethman, P. Gastañaduy, A. Lee, J. Leung, M. Holshue, et al., Notes from the field: Community outbreak of measles — Clark County, Washington, 2018–2019, *MMWR. Morb. Mortal. Wkly. Rep.*, **68** (2019), 446–447. https://doi.org/10.15585/mmwr.mm6819a5
66. L. Sherrard, J. Hiebert, S. Squires, Measles surveillance in Canada: Trends for 2014, *Can. Commun. Dis. Rep.*, **41** (2015), 157–168. https://doi.org/10.14745/ccdr.v41i07a01
67. W. See, Y. K. Ng, L. Cui, Y. Kita, S P.-L. Ooi, V. Lee, et al., Singapore's efforts to achieve measles elimination in 2018, *Western Pac. Surveill. Response J.*, **12** (2021). https://doi.org/10.5365/wpsar.2019.10.4.002
68. T. Marks, S. Vanderslott, *Which countries have mandatory childhood vaccination policies?* 2021. Available from: `https://ourworldindata.org/childhood-vaccination-policies`
69. World Health Organization, *Measles vaccination coverage*, 2025. Available from: `https://immunizationdata.who.int/global/wiise-detail-page/measles-vaccination-coverage`.
70. World Health Organization, *Measles reported cases and incidence*, 2025. Available from: `https://immunizationdata.who.int/global/wiise-detail-page/measles-reported-cases-and-incidence`.
71. G. Fenza, M. Gallo, V. Loia, F. Orciuoli, E. Herrera-Viedma, Data set quality in machine learning: Consistency measure based on group decision making, *Appl. Soft Comput.*, **106** (2021), 107366. https://doi.org/10.1016/j.asoc.2021.107366
72. A. Tasciotti, F. Urban, F. de Dea, L. Bortolussi, G. Caravagna, A. D'Onofrio, Deep learning-informed Bayesian model-based analysis to estimate superspreading events in epidemic outbreaks, *IEEE Access*, **12** (2024), 161375–161400. https://doi.org/10.1109/ACCESS.2024.3490374
73. H. Ismail Fawaz, G. Forestier, J. Weber, L. Idoumghar, P.-A. Muller, Transfer learning for time series classification, in *2018 IEEE International Conference on Big Data (Big Data)*, (2018), 1367–1376. https://doi.org/10.1109/BigData.2018.8621990
74. A. P. Fiebelkorn, S. B. Redd, P. A. Gastañaduy, N. Clemmons, P. A. Rota, J. S. Rota, et al., A comparison of postelimination measles epidemiology in the United States, 2009–2014 versus 2001–2008, *J. Pediatric Infect. Dis. Soc.*, **6** (2015), 40–48. https://doi.org/10.1093/jpids/piv080
75. S. Bidari, W. Yang, Global resurgence of measles in the vaccination era and influencing factors, *Int. J. Infect. Dis.*, **147** (2024), 107189. https://doi.org/10.1016/j.ijid.2024.107189
76. N. Boers, Early-warning signals for Dansgaard-Oeschger events in a high-resolution ice core record, *Nat. Commun.*, **9** (2018), 2556. https://doi.org/10.1038/s41467-018-04881-7

AIMS Press