

An early warning indicator trained on stochastic disease-spreading models with different noises

Amit K. Chakraborty¹, Shan Gao¹, Reza Miry², Pouria Ramazi², Russell Greiner^{3,4}, Mark A. Lewis⁵, and Hao Wang^{1,*}

¹*Department of Mathematical and Statistical Sciences, University of Alberta, AB, Canada*

²*Department of Mathematics and Statistics, Brock University, ON, Canada*

³*Department of Computing Science, University of Alberta, AB, Canada*

⁴*Alberta Machine Intelligence Institute, AB, Canada*

⁵*Department of Mathematics and Statistics and Department of Biology, University of Victoria, BC, Canada*

**Corresponding author: hao8@ualberta.ca*

Abstract

The timely detection of disease outbreaks through reliable early warning signals (EWSs) is indispensable for effective public health mitigation strategies. Nevertheless, the intricate dynamics of real-world disease spread, often influenced by diverse sources of noise and limited data in the early stages of outbreaks, pose a significant challenge in developing reliable EWSs, as the performance of existing indicators varies with extrinsic and intrinsic noises. Here, we address the challenge of modeling disease when the measurements are corrupted by additive white noise, multiplicative environmental noise, and demographic noise into a standard epidemic mathematical model. To navigate the complexities introduced by these noise sources, we employ a deep learning algorithm that provides EWS in infectious disease outbreak by training on noise-induced disease-spreading models. The indicator's effectiveness is demonstrated through its application to real-world COVID-19 cases in Edmonton and simulated time series derived from diverse disease spread models affected by noise. Notably, the indicator captures an impending transition in a time series of disease outbreaks and outperforms existing indicators. This study contributes to advancing early warning capabilities by addressing the intricate dynamics inherent in real-world disease spread, presenting a promising avenue for enhancing public health preparedness and response efforts.

Keywords— infectious disease, epidemiological model, stochastic differential equations, early warning signals, machine learning, critical transitions

1 Introduction

Infectious diseases have imposed significant burdens on societies, and their impact has intensified in recent decades [1, 2]. While public health measures have made considerable progress in eradicating certain diseases, the continuous emergence of new diseases and the re-emergence of some older ones have resulted in substantial loss of life and economic hardship [3]. Over the past two decades, more than ten severe infectious diseases have (re-)emerged, resulting in approximately 7.3 million global deaths [1, 4–7]. With infectious diseases persistently on the rise at a global level, the necessity for an effective early warning signal (EWS) becomes increasingly imperative [8].

An EWS is a model-independent tool that anticipates an impending critical transition of a disease outbreak before it escalates to an epidemic level, without relying on empirically fitted mathematical models. [3]. This allows for the early identification of outbreaks, which, in turn, can potentially help in reducing the spread of the disease and minimizing its impact through targeted interventions [9–11]. Several authors have explored EWSs for different diseases using time series-based machine learning models. For example, Brett et al. [12] investigated EWSs for dengue and plague outbreaks and mumps re-emergence, Yang et al. [13] focused on influenza, Li et al. [14] studied COVID-19, and Gao et al. [15] predicted mortality risk for COVID-19. Although an effective EWS has the potential to save millions of lives from future diseases, designing an efficient early warning system that can adapt to the complex and uncertain patterns often exhibited by real-world disease dynamics remains a significant challenge [16].

EWSs are diverse and can be generated using a wide range of detection methods, including wastewater-based epidemiology, search engine trends, and human bio-monitoring [17]. While these methods are increasingly popular, their implementation may require additional resources and can differ greatly from one country to another [17]. For instance, wastewater-based epidemiology necessitates the selection of appropriate biomarkers, search engine trends require internet access with precise keyword configurations, and human bio-monitoring demands the careful selection of a control group [17].

Critical slowing down (CSD) serves as the fundamental basis for EWSs to anticipate critical transitions [3]. As a system approaches a critical threshold or undergoes a significant change, such as a phase transition, its ability to recover from perturbations decreases [18]. This slowing down results in an increase in recovery time and a decrease in the recovery rate from small perturbations [19, 20]. Mathematically, CSD occurs due to the diminishing real part of the dominant eigenvalue as a system approaches the bifurcation point, reaching zero precisely at the bifurcation point [18], where bifurcation refers to qualitative changes in the behavior of a dynamical system's steady state as its parameters are varied [21]. Detection of critical slowing down in a system can be achieved by measuring temporal trends in statistics, such as increased autocorrelation (AC), variance, and the magnitude of fluctuations [3, 19].

The reliability of CSD as an indicator of an EWS is not guaranteed when a system involves stochasticity [19, 22]. In highly stochastic scenarios, the dynamics of a system can accelerate rather than slow down near a bifurcation point [23], and a transition often occurs far away from the local bifurcation point [19, 23, 24]. The leading statistical measures, such as variance and AC, are differentially affected due to the external and internal effects of noise, and exhibit different trends for different types of noises [23]. Traditional eigenvalue-based analyses in dynamical systems are also ineffective in predicting state changes in stochastic systems, as unpredictable fluctuations can significantly distort signals generated by linear terms [24]. As stochasticity leads to unforeseen system behavior, the deterministic framework of the underlying mechanisms becomes obscured, potentially resulting in false alarms or delayed responses.

The impact of noise is a critical consideration when generating EWS for disease outbreaks, as the data observed from real-world systems is convoluted with different types of noises [24]. Various forms of randomness can be introduced in a system by both intrinsic and extrinsic noises [25]. Two common forms of noise used when modeling disease are environmental and demographic noise. Environmental noise, also known as extrinsic noise, arises due to the changes in environmental factors outside of population [26, 27]. It plays a crucial role in zoonotic, vector-borne, and waterborne diseases, particularly in Ebola, avian influenza, malaria, and cholera transmission [28, 29]. On the other hand, demographic noise, referred to as intrinsic noise, is associated with temporal variation within the population process, including births, deaths, immigration, emigration, and state transitions [26, 30]. In addition to these two types of noise, additive white noise represents the interplay between deterministic structures and random fluctuations with consistent intensity, possessing the potential to trigger an epidemic in circumstances where there was no initial spread of infection in the absence of this specific noise [31]. Since EWSs arise from the interaction of the dynamics and the noise, the type of noise will affect the nature and intensity of an EWS.

To study the data-driven approach of EWS in infectious diseases, many use generic early warning indicators (EWIs), such as variance and AC [16, 25, 32–39], and machine learning models [12–15, 20, 40, 41]. While generic EWIs offer a computationally efficient method, their pattern can fluctuate, depending on system behavior and noise levels. Moreover, the performance of generic EWIs is sensitive to the data sampling, quantitative and objective measures, length of the time series, and the size of window and bandwidth [40]. This introduces uncertainty regarding their universal applicability or potential influence by specific dynamic features [3, 16, 22]. In addition to these challenges, generic EWIs can not accurately predict the type of bifurcation that is going to happen [20]. Machine learning models, in contrast to generic EWIs, can not only identify the impending change in the system but also categorize the type of bifurcation expected according to the training labels [20]. These models possess the capacity to recognize and learn patterns, allowing them to capture the statistical features inherent in the data [40]. Nevertheless, machine learners, particularly deep learners (DL), demand substantial amounts of training data to produce effective models and are computationally intensive.

Training machine learning models with synthetic data is a fundamental strategy, especially valuable when collecting and labeling extensive real datasets is challenging, time-consuming, resource-intensive, and restricted due to privacy concerns [42]. Synthetic data generation involves the creation of artificial datasets that closely mimic real-world data, which can produce better models as it involves greater data diversity [42]. One of the key advantages of synthetic data lies in the control it provides over data characteristics [42, 43]. This control is invaluable for researchers and practitioners, enabling the exploration of a wide array of scenarios, data distributions, and complexities [43]. Moreover, synthetic data facilitates model training on edge cases and rare events, which may occur in the real world but only rarely [44]. This becomes particularly crucial in applications where the consequences of failure are substantial or the process is highly intricate, such as in medical diagnostics or disease spread [43]. Synthetic scenarios capturing these exceptional cases empower models to make more informed decisions when confronted with unusual situations, ultimately enhancing the model’s robustness and reliability.

A recent study conducted by Bury et al. [20] demonstrated that a machine learning model can outperform generic EWIs, providing more accurate EWS by analyzing the pre-transition time series of a system. They employed a DL algorithm trained on a synthetic dataset generated from two-dimensional random dynamical systems comprising a general polynomial of up to third order. The parameter values of the random dynamical system were adjusted to achieve fold, Hopf, and transcritical bifurcations, and the stochastic simulations with additive white noise were then conducted to generate training data. Their study suggested that by training the DL algorithm on a diverse range of datasets, it could capture fundamental features shared by various systems, thus enabling it to provide EWS for systems not explicitly included in the training data. The model was tested on time series from ecology, thermoacoustics, climatology, and epidemiology, which were not part of the training datasets, and successfully anticipated critical transitions in both simulated models and empirical data. Moreover, it accurately identified the specific types of bifurcations, that were occurring within these systems. Nevertheless, when applied to a disease-spreading model, both the DL model and generic EWIs did not perform better than random chance. This motivated us to investigate the performance of EWIs in the context of a disease model.

Our objective is to enhance EWSs for infectious diseases. To achieve this, we evaluate the effectiveness of a DL algorithm, alongside variance and lag-1 AC, in disease-spreading models. To comprehensively analyze the interaction between disease dynamics and noise, we consider a standard epidemic model with additive white noise, multiplicative environmental noise, and demographic noise. Our hypothesis posits that by training the DL algorithm on noise-induced models, the algorithm can outperform other EWIs and demonstrate higher accuracy in providing EWSs for disease outbreaks. We use the time series of the number of infected cases to generate EWSs. We specifically focus on the transcritical form of bifurcation in disease models, identifying the bifurcation point by observing the basic reproduction number (R_0)—a

dimensionless number that denotes the average number of secondary infections arising from a single infected case in a fully susceptible population. A disease model can exhibit other types of bifurcation [45], which are not in our consideration.

2 Mathematical models and the impacts of noises

We consider the susceptible–infectious–recovered (SIR) model [46], which serves as a foundation for most disease-spreading models. The SIR model divides the population into susceptible (S), infectious (I), and recovered (R) compartments and uses a system of ordinary differential equations to describe how individuals move between them over time. The model equations are given by

$$\begin{cases} \frac{dS(t)}{dt} = \Lambda - \beta(t)S(t)I(t) - \mu S(t), \\ \frac{dI(t)}{dt} = \beta(t)S(t)I(t) - \alpha I(t) - \mu I(t), \\ \frac{dR(t)}{dt} = \alpha I(t) - \mu R(t), \end{cases} \quad (2.1)$$

where Λ is the recruitment rate of susceptible populations, $\beta(t)$ is the time-dependent disease transmission rate, μ is the natural death rate, and α is the recovery rate from the disease. When β is treated as a constant parameter, the disease-free equilibrium point of the model (2.1) is $(\Lambda, 0, 0)$, the endemic equilibrium point is $(\frac{\mu+\alpha}{\beta}, \frac{\Lambda}{\mu+\alpha} - \frac{\mu}{\beta}, \frac{\alpha\Lambda}{\mu(\mu+\alpha)} - \frac{\alpha}{\beta})$, and the basic reproduction number, R_0 , is $\frac{\beta\Lambda}{\mu(\alpha+\mu)}$. A disease goes extinct if $R_0 < 1$ and persists if $R_0 > 1$, and a transcritical bifurcation occurs when $R_0 = 1$, where both equilibrium meet and exchange their stability [3, 16, 45, 47]. The bifurcation value, corresponding to $R_0 = 1$, arises from the model with a constant β which is regarded as a bifurcation parameter. The first two equations of (2.1) are independent of R , and the reduced system is given by

$$\begin{cases} \frac{dS(t)}{dt} = \Lambda - \beta(t)S(t)I(t) - \mu S(t), \\ \frac{dI(t)}{dt} = \beta(t)S(t)I(t) - \alpha I(t) - \mu I(t). \end{cases} \quad (2.2)$$

To model the spread of an infectious disease by an exposed individual, the SIR-type model is commonly expanded by including an exposed compartment, leading to the susceptible–exposed–infectious–recovered (SEIR) model [48]. The model equations are given by

$$\begin{cases} \frac{dS(t)}{dt} = \Lambda - \beta(t)S(t)I(t) - dS(t), \\ \frac{dE(t)}{dt} = \beta(t)S(t)I(t) - (d + \kappa)E(t), \\ \frac{dI(t)}{dt} = \kappa E(t) - (d + \gamma)I(t), \\ \frac{dR(t)}{dt} = \gamma I(t) - dR(t), \end{cases} \quad (2.3)$$

where $S(t)$, $E(t)$, $I(t)$, and $R(t)$ denote the susceptible, exposed, infectious, and recovered individuals, respectively, at time t . Moreover, Λ is the recruitment rate of the susceptible popu-

lation, $\beta(t)$ is the disease transmission rate, d is the natural death rate, $1/\gamma$ is the mean infectious period, and $1/\kappa$ is the mean exposed period. The basic reproduction number of the SEIR model is $\frac{\kappa\beta\Lambda}{d(d+\kappa)(d+\gamma)}$, and the bifurcation point is $\beta_c = \frac{d(d+\kappa)(d+\gamma)}{\kappa\Lambda}$.

2.1 Noise-induced models

Deterministic models provide useful insights into disease dynamics under idealized conditions. However, the dynamics of infectious diseases are inherently complex and subject to random fluctuations, thus necessitating the careful consideration of noise in disease models. Noise disrupts the deterministic structure and influences the dynamics of the population compartments. These noises can arise from various sources, including diversity in population behavior, environmental factors, and even the stochastic nature of disease transmission. A noise-induced model is commonly defined through a set of stochastic differential equations.

2.1.1 SIR and SEIR models with additive white noise

The most common noise model is the additive noise model, wherein a perturbed observed compartment is construed as the sum of unaltered deterministic dynamics and an additive noise component [49]. This noise model serves as a means to simulate the inherent stochasticity of natural processes and reproduce their consequential impact on a system [50]. To model additive white noise, we add a white noise term into each equation of the deterministic model, encompassing the continuous intensity of fluctuation associated with the dynamics of individual compartments. A stochastic differential equation with additive white noise is given by

$$dX(t) = f(X(t), t)dt + \sigma dW(t), \quad (2.4)$$

where $X(t)$ is a state variable, σ is noise intensity, and $W(t)$ is a Wiener process. Following this, the SIR model (2.2) with additive white noise has the form

$$\begin{cases} dS(t) = \Lambda dt - \beta(t)S(t)I(t)dt - \mu S(t)dt + \sigma_1 dW_1(t), \\ dI(t) = \beta(t)S(t)I(t)dt - \alpha I(t)dt - \mu I(t)dt + \sigma_2 dW_2(t), \end{cases} \quad (2.5)$$

where σ_i , $i = 1, 2$, are the intensities of white noise and $W_i(t)$, $i = 1, 2$, are independent Wiener processes. Similarly, the SEIR model (2.3) with additive white noise is given by

$$\begin{cases} dS(t) = \Lambda dt - \beta(t)S(t)I(t)dt - dS(t)dt + \sigma_1 dW_1(t), \\ dE(t) = \beta(t)S(t)I(t)dt - (d + \kappa)E(t)dt + \sigma_2 dW_2(t), \\ dI(t) = \kappa E(t)dt - (d + \gamma)I(t)dt + \sigma_3 dW_3(t), \\ dR(t) = \gamma I(t)dt - dR(t)dt + \sigma_4 dW_4(t), \end{cases} \quad (2.6)$$

where σ_i , $i = 1, 2, 3, 4$, are the intensities of white noise and $W_i(t)$, $i = 1, 2, 3, 4$, are independent Wiener processes.

2.1.2 SIR model with multiplicative environmental noise

Substantial environmental noise can trigger abrupt transitions in disease dynamics even before reaching a critical threshold; nevertheless, in certain instances, environmental fluctuations can alter the stability landscape and promote resilience [51]. Environmental variability introduces multiplicative noise into the modeling process that depends on the system's current state [52]. To model the multiplicative environmental noise, we consider that the fluctuations exhibit a direct correlation with the size of the population compartment. Consequently, we incorporate a multiplication by the population compartment in the noise-induced terms within the dynamics of the respective compartments. A stochastic differential equation with multiplicative environmental noise is given by

$$dX(t) = f(X(t), t)dt + \sigma X(t)dW(t). \quad (2.7)$$

Therefore, following (2.2), the SIR model with multiplicative environmental noise is given by

$$\begin{cases} dS(t) = \Lambda dt - \beta(t)S(t)I(t)dt - \mu S(t)dt + \sigma_1 S(t)dW_1(t), \\ dI(t) = \beta(t)S(t)I(t)dt - \alpha I(t)dt - \mu I(t)dt + \sigma_2 I(t)dW_2(t). \end{cases} \quad (2.8)$$

2.1.3 SIR model with demographic noise

Demographic noise can either mitigate abrupt transitions or promote alternative stable states and abrupt transitions [51]. Moreover, demographic noise has some other epidemiological outcomes such as early decrease of infection, epidemic outbreaks with the collapse of population, and either the extinction of infection or an endemic situation [53]. To characterize demographic stochasticity, we use a system of stochastic differential equations that is based on a diffusion process where both the time and the state variables are continuous. Following the derivation procedure in [26, 28], we incorporate demographic noise into the model (2.2). Examining the interactions within the SIR model (2.2), we identify the list of five possible changes along with their respective probabilities, as presented in Table 1. To incorporate demographic stochas-

Table 1: Transition probabilities associated with changes in model 2.2

i	Description	Change $(\Delta X)_i$	Probability, p_i
1	Recruitment of S	$(1, 0)^{tr}$	$\Lambda \Delta t$
2	Transmission from S to I	$(-1, 1)^{tr}$	$\beta SI \Delta t$
3	Death of S	$(-1, 0)^{tr}$	$\mu S \Delta t$
4	Death of I	$(0, -1)^{tr}$	$\mu I \Delta t$
5	Recovery of I	$(0, -1)^{tr}$	$\alpha I \Delta t$
6	No change	$(0, 0)^{tr}$	$1 - (\Lambda + \beta SI + \mu S + \mu I + \alpha I) \Delta t$

ticity, we need to calculate the expectation $E(\Delta X)$ and the covariance $\Sigma(\Delta X)$ for the time interval Δt . The expectation is given by

$$E(\Delta X) = \sum_{i=1}^6 p_i(\Delta X)_i = \begin{pmatrix} \Lambda - \beta SI - \mu S \\ \beta SI - \alpha I - \mu I \end{pmatrix} \Delta t = \mu \Delta t.$$

The covariance is given by

$$\begin{aligned} \Sigma(\Delta X) &= E((\Delta X)(\Delta X)^{tr}) = \sum_{i=1}^6 p_i(\Delta X)_i(\Delta X)_i^{tr} = \begin{pmatrix} \Lambda + \beta SI + \mu S & -\beta SI \\ -\beta SI & \beta SI + \alpha I + \mu I \end{pmatrix} \Delta t \\ &= \begin{pmatrix} v_{11} & v_{12} \\ v_{21} & v_{22} \end{pmatrix} \Delta t = V \Delta t. \end{aligned}$$

Letting $\Delta t \rightarrow 0$, then the stochastic SIR model with demographic noise has the form

$$dX = \mu dt + G dW(t),$$

where $X = (S, I)^{tr}$, and G is the square root of V defined by

$$G = \sqrt{V} = \frac{1}{e} \begin{pmatrix} v_{11} + d & v_{12} \\ v_{21} & v_{22} + d \end{pmatrix},$$

where $v_{12} = v_{21}$, $d = \sqrt{v_{11}v_{22} - v_{12}^2}$, and $e = \sqrt{v_{11} + v_{22} + 2d}$. Choosing, $v_{11} = a$, $v_{12} = v_{21} = b$, $v_{22} = c$, and substituting d and e in the G matrix, we get the following stochastic SIR model with demographic noise

$$\begin{cases} dS(t) = \Lambda dt - \beta(t)S(t)I(t)dt - \mu S(t)dt + \frac{a(t) + d(t)}{e(t)}dW_1(t) + \frac{b(t)}{e(t)}dW_2(t), \\ dI(t) = \beta(t)S(t)I(t)dt - \alpha I(t)dt - \mu I(t)dt + \frac{b(t)}{e(t)}dW_1(t) + \frac{c(t) + d(t)}{e(t)}dW_2(t), \end{cases} \quad (2.9)$$

where $a(t) = \Lambda + \beta(t)S(t)I(t) + \mu S(t)$, $b(t) = -\beta(t)S(t)I(t)$, $c(t) = \beta(t)S(t)I(t) + \alpha I(t) + \mu I(t)$, $d(t) = \sqrt{a(t)c(t) - b^2(t)}$, $e(t) = \sqrt{a(t) + c(t) + 2d(t)}$, and $W_i(t)$, $i = 1, 2$, is the Wiener process.

3 Methods

3.1 Simulated training data

To train the DL algorithm, we generated the training data from the noise-induced SIR models, i.e., from equations (2.5), (2.8), and (2.9). The training data comprises two types of simulations: *transcritical* and *null*. In transcritical simulations, bifurcation occurs within a specified

time length, while in null simulations there is no bifurcation within the same time frame. The maximum time series length is set to 1500, and we assume that the transition time for the noise-induced models aligns with the deterministic model (2.2). Specifically, the transition in the noise-induced models is the time t when the value of R_0 becomes equal to 1. Therefore, when we use data, we obtain the critical value of the disease transmission rate, β_c , at the bifurcation point by setting $R_0 = 1$, resulting in $\beta_c = \frac{\mu(\alpha + \mu)}{\Lambda}$.

We specify the parameters $\Lambda = 100$, $\alpha = 1$, and $\mu = 1$, and randomly choose other parameters to incorporate variability into the training data, ensuring that the DL algorithm is trained on a wide range of scenarios. The disease transmission rate, β , is assumed to be time-dependent, following a linear increasing pattern: $\beta(t) = \beta_0 + \beta_1 t$, where β_0 and β_1 are chosen from triangular distributions. For β_0 , the lower and upper limits are set to 0 and $\frac{\beta_c}{2}$ respectively. The choice of β_1 determines whether a simulation is transcritical or null. For null simulations, β_1 ranges from 0 (lower limit) to $\frac{\beta_c - \beta_0}{1500}$ (upper limit), whereas for transcritical simulations, β_1 ranges from $\frac{\beta_c - \beta_0}{1500}$ (lower limit) to $\frac{2 \times \beta_c - \beta_0}{1500}$ (upper limit). The modes of these distributions are set at half of their respective upper limits. These parameters result in a random transition between time 0 and 1500 for transcritical simulations and no transition for null simulations (Figure 1). Random transition times are chosen for transcritical simulations to account for the uncertain nature of transition times in real-world scenarios. These parameter choices collectively contribute to the model's robustness by exposing it to a wide range of conditions, fostering adaptability, and enhancing its potential to generalize effectively across different scenarios and uncertainties.

The Wiener process is sampled from a normal distribution, $dW_i \sim \mathcal{N}(0, \sqrt{dt})$, $i = 1, 2$, and the noise intensity, σ_i , $i = 1, 2$, is sampled from a triangular distribution with a lower limit of 0, a mode of 0.5, and an upper limit of 1, and is randomized for each simulation. The random choice of noise intensity introduces different levels of stochasticity into the simulations, making them more adaptable to real-world scenarios where accurately measuring all parameters and noises is often impractical. These variations provide a broader spectrum of possible phenomena.

We set the initial values of susceptible and infected to 500 and 7 respectively, with a burn-in period of 100 time units. The Euler–Maruyama method is employed for numerical simulations with $dt = 0.01$. For the transcritical simulations, simulation lengths extend up to the transition points, whereas for the null simulations, the length extends up to time 1500. We generated 30,000 simulations from each of the equations (2.5), (2.8) and (2.9), and an additional 10,000 simulations from equation (2.5) by doubling the noise intensity. Thus, a total of 0.1 million simulations were used to train the DL algorithm.

The training data was generated in multiple batches using the Compute Canada server. Time series residuals were computed using the `ewstools` Python package [54], utilizing Lowess (Locally Weighted Scatterplot Smoothing) smoothing [55] with a span of 0.2. Additionally, lag-1 AC and variance of the simulations were computed with a rolling window of 0.25 using




Figure 1: (a) Linear increase samples of $\beta(t)$ for null (green) and transcritical (purple) simulations. A transition occurs when β crosses the critical value β_c (brown horizontal line), with the transition time marked by the intersection of the two lines (red dotted vertical lines). In transcritical simulations, β crosses β_c randomly between time 0 and 1500, while in null simulations, β does not cross β_c before time 1500. (b) If a transition occurs between time 0 and 1500, the number of infected of the preceding 500 (100) points leading up to the bifurcation point are utilized as training data for transcritical simulations. (c) In the absence of a transition during the period from time 0 to 1500, the number of infected of the most recent 500 (100) time points are selected as training data for null simulations.

the same package.

3.2 Deep learning algorithm and training

We assume that the occurrence or absence of an upcoming transition can be reliably determined by examining recent time points within the time series and we define two subsets of observations of each time series: one is 100 time points, and another is 500 time points. Therefore, for transcritical simulations, a subset of observations consists of the preceding 500 (100) time points before a bifurcation point, and for null simulations, it consists of the last 500 (100) time points (Figure 1). If a transition occurs before the 500th (100th) time point, the values beyond the transition point are padded with zeros. This ensures that the transition point is always included in the dataset for transcritical simulations.

We followed the same DL algorithm, training process, and hyperparameters as outlined in [20]. The neural network architecture combined Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) layers [20, 56, 57]. Since we have a two-class classification problem, we modified the output (dense) layer of the neural network to 2 units with a softmax activation function.

The DL algorithm was trained on two subsets of observations, i.e., lengths of 500 and 100 points. The residuals of the time series were used to train the algorithm, normalized by dividing each individual value of a time point by the mean absolute value of the entire residuals. The algorithm underwent 1500 epochs of training, and the learning rate was set at 0.0005. Hyperparameters underwent tuning through a series of grid searches, and the training, test, and validation split was 95%, 1%, and 4%, respectively. The F1 score, precision, and recall for an ensemble of two models for the 500-length trained were all 97.6%, and for the 100-length trained were all 91.0 %. The average prediction at each data point was taken from two trained models of each length (either 500 or 100) and an ensemble was formed.

3.3 Testing

To assess the performance of the DL model, we conducted tests using four noise-induced disease-spreading mathematical models: three corresponding to the models used during training, and the fourth model featuring higher dimensionality compared to our trained models. Additionally, real-world empirical data on COVID-19 cases in Edmonton was included in the test.

3.3.1 Mathematical models

To test the SIR model with additive white noise (2.5), multiplicative environmental noise (2.8), and demographic noise (2.9), we simulated time series using the same parameters and configurations employed during the training phase. Due to the inherent randomness of certain parameters, these time series may exhibit variations compared to the training simulations. Our objective is to determine whether the DL model can accurately classify the transition regardless of these random variations.

To evaluate the predictive performance of the DL model for a higher dimensional model than the trained model, we test the SEIR model with additive white noise (2.6). We use parameter values $\Lambda = 100$, $d = 0.75$, $\kappa = 2$, $\gamma = 1$, and initial conditions $S_0 = 500$, $E_0 = 1$, $I_0 = 2$, $R_0 = 2$, to simulate the model. The disease transmission rate and other parameters remain the same as those chosen for generating the training data discussed in section 3.1.

We generated 20 time series from each test model, comprising 10 transcritical simulations and 10 null simulations. In each transcritical simulation, the time series extended up to the bifurcation point between time 0 and 1500, while each null simulation extended up to time 1500. These simulated time series served as input for trained models to identify the presence of a bifurcation, thereby facilitating the evaluation of the DL model's performance.

3.3.2 Empirical data

We used the COVID-19 dataset of Edmonton sourced from [58], covering daily cases from March 2020 to July 2023. Employing the `EpiEstim` package [59, 60] in the R programming language, we computed the effective reproduction number (R_e) from the reported incidence cases in biweekly sliding windows. This package requires an estimated serial interval distribution of the disease and we considered the mean of the serial interval of COVID-19 to be 6.3, with a standard deviation of 4.2 [60].

We extracted consecutive daily case data from the entire time series based on periods where the mean R_e was less than 1, up to the point where R_e reached 1 (Supplementary Figure S1). This process resulted in obtaining sub-series of varying lengths. To ensure consistency, we selected sub-series with a minimum length of eight weeks. Following this criterion, we derived a total of 7 sliced time series immediately preceding the bifurcation point, which we labeled as transcritical time series. For null time series, we adjusted the length of each transcritical time

series by shortening it by four weeks. This adjustment makes the null time series data up to four weeks before the bifurcation point. Our objective is to assess the capability of the trained DL model to detect the impending transition based on these pre-transition time series.

3.4 Prediction and ROC curve

We employed the DL algorithm trained on two distinct types of data to generate EWSs for both simulated noise-induced time series datasets and empirical data of COVID-19. First, we utilized the 500-classifier DL algorithm from Bury et al. [20], which was trained on data generated from random dynamical systems with polynomial terms, which we will call the PODATR model. Secondly, we employed the DL algorithm trained on data from noise-induced SIR models of time lengths 100 and 500 which we will call the SIDATR-100 and SIDATR-500 models respectively. Of course, we want an accurate classification of both the occurrence and non-occurrence of a transition. Consequently, the trained models assign probabilities for each type of bifurcation in the training set. Thus, based on inputs of the number of infected from a subset of observations, the SIDATR-100 and SIDATR-500 models assign probabilities for both transcritical and null bifurcations, while the PODATR model assigns probabilities for each of fold, Hopf, transcritical, and null bifurcations. A type of bifurcation with the highest probability assigned at the end of the input observation period is considered as a predictive signal for the occurrence of that bifurcation type.

Our prediction process follows the methodology outlined in Bury et al. [20]. We used an expanding window of ten time points for predictions, resulting in a maximum of 50 (10) predictions from a time series spanning up to 500 (100) points. We extracted a set of the last 5 predictions from each of the time series, allowing for a comprehensive assessment of model performance when it reaches the bifurcation point. Thus, for each type of simulation from the mathematical models, we generated a maximum of 50 (10) predictions to ensure a robust and well-balanced analysis.

We compared the performance of the DL models to the lag-1 AC and variance, using the area under the receiver operating characteristic (ROC) curve. The ROC curve illustrates the relationship between sensitivity (true positive rate) and $1 - \text{specificity}$ (false positive rate) across varying discrimination thresholds, offering a comprehensive view of the classifier’s overall performance. An ideal warning system would detect all thresholds without any false alarms, resulting in an area under the ROC curve equal to one. Conversely, an ineffective system would struggle to differentiate between false and true signals.

To assess generic EWIs, variance and lag-1 AC, we use the Kendall τ value using `scipy.stats` package in Python. This statistic serves as an indicator of either an increasing or decreasing trend, where higher and positive values indicate a more strongly increasing trend [20]. Consequently, we employed these indicators to make predictions for a specific outcome when the Kendall τ statistic surpassed a designated discrimination threshold [20].

4 Results

4.1 Performance on mathematical models

The area under the ROC curve (AUC) for lag-1 AC exceeded 0.85 in the simulations of the SIR model with white noise, the demographic noise, and the SEIR model, outperforming variance (Figure 2). However, a contrasting trend emerged in the simulations of SIR model with environmental noise, where variance achieved an AUC of 0.82, exceeding the predictive accuracy of lag-1 AC (Figure 2).

The PODATR model demonstrated strong performance in certain scenarios, achieving an AUC exceeding 0.95 in the simulations of the SIR models with white noise and environmental noise, surpassing both variance and lag-1 AC (Figure 2). However, its performance decreased in SEIR model, where it achieved an AUC of approximately 0.80 (Figure 2). The model outperformed only variance in the demographic noise model, achieving an AUC of 0.73 (Figure 2).

The AUC for the SIDATR-500 model consistently reached around one across all test model simulations, illustrating its exceptional predictive capability and robustness. This performance surpassed not only that of the PODATR model but also measures of variance and lag-1 AC, indicating its superiority in capturing the underlying dynamics of the systems under study (Figure 2). On the other hand, the SIDATR-100 model outperformed the variance and PODATR model in the SIR model with demographic noise (Figure 2). Additionally, it outperformed both variance and lag-1 AC in the SIR model with environmental noise (Figure 2).




Figure 2: Area under the ROC curve of the generic EWIs—variance, and lag-1 AC—in addition to the PODATR, SIDATR-500 and SIDATR-100 DL model. Performance was assessed on the last five predictions of transcritical and null simulations of the SIR model with white noise (yellow), SIR model with multiplicative environmental noise (green), SIR model with demographic noise (cyan), SEIR model with white noise (brown), and COVID-19 dataset of Edmonton (blue).

In transcritical simulations of the SIR model with white noise, both the SIDATR-100 and SIDATR-500 models successfully identified 48 out of 50 frequencies (Figure 3). However, dur-

ing null simulations, the SIDATR-100 model detected only 29 frequencies, while the SIDATR-500 model captured all frequencies (Figure 3).

In both types of simulations of SIR model with environmental noise, the SIDATR-500 model successfully identified all frequencies, whereas the SIDATR-100 model identified all frequencies in transcritical simulations and 48 out of 50 frequencies in null simulations (Figure 3).

Interestingly, in transcritical simulations of both the SIR model with demographic noise and SEIR models, the SIDATR-100 models identified slightly more frequencies compared to the SIDATR-500 models, whereas in null simulations, the SIDATR-500 model exhibited better frequency detection than the SIDATR-100 model (Figure 3).

The PODATR model consistently identified impending transitions in transcritical simulations across all test models, demonstrating its effectiveness in recognizing critical changes in the system. However, it encountered challenges in accurately classifying most of the null simulations within the test models. Despite its success in identifying impending transitions in transcritical simulations, the PODATR model occasionally misclassified bifurcation types among frequencies (Figure 3).

Overall, the SIDATR-500 model consistently exhibited superior frequency performance in the last five predictions across the simulated test models compared to both the PODATR and SIDATR-100 models (Figure 3). However, in most simulations, particularly during shorter time spans and farther from the bifurcation point, the SIDATR-500 model displayed variations in probability assignments (Supplementary Figures S2, S4, S6, and S8). Despite this, as the simulations progressed toward the bifurcation point, the model’s predictions significantly improved, showcasing precise probability assignments for the correct bifurcation type.

4.2 Performance on empirical data

We employed seven transcritical and null time series from the COVID-19 dataset of Edmonton as inputs for the generic EWIs and the trained DL models and investigated the EWSs of the outbreak. The AUC values for variance, lag-1 AC, the PODATR model, and the SIDATR-500 model were all found below 0.50 (Figure 2). Notably, the SIDATR-500 model showed an improvement in AUC compared to other indicators, yet overall performance remained below random chance. Intriguingly, the SIDATR-100 model demonstrated superior performance within the COVID-19 dataset, achieving an AUC of 0.71 and outperforming other indicators (Figure 2).

Each set of transcritical and null simulations yielded 35 frequencies derived from the seven time series. The PODATR model successfully predicted transitions in over half of the frequencies within transcritical simulations (Figure 3). However, in null simulations, most frequencies were inaccurately identified (Figure 3). Among the accurately identified frequencies in transcritical simulations, the majority corresponded to the Hopf bifurcation type (Figure 3).




Figure 3: Frequency distribution of the favored probability among transcritical (blue) and null (orange) simulations generated by the (a) SIDATR-100, (b) SIDATR-500, and (c) PODATR model. In each input simulation, SIDATR-100 and SIDATR-500 models assign probabilities for transcritical and null bifurcations, while PODATR assigns probabilities for fold, Hopf, transcritical, and null bifurcations. Based on the last five favored probabilities in each simulation, a total of 50 frequencies from the transcritical simulations and another 50 frequencies from the null simulations were extracted from each mathematical model. Additionally, for the COVID-19 data, a total of 35 frequencies of each type were extracted.

In contrast, the SIDATR-500 model successfully predicted all frequencies within transcritical simulations but struggled to accurately identify most frequencies in null simulations (Figure 3). The SIDATR-100 model exhibited a superior frequency distribution compared to the other models, accurately capturing all frequencies within transcritical simulations and approximately half of the frequencies within null simulations (Figure 3). These findings suggest that the DL algorithm, trained by the noise-induced model, provides enhanced EWSs for impending transitions of a disease outbreak.

5 Discussion

We investigated EWSs in the context of noise-induced disease-spreading models, incorporating additive white noise, multiplicative environmental noise, and demographic noise. We analyzed

these signals by employing a DL algorithm alongside two standard EWIs: variance and lag-1 AC. To effectively train the DL algorithm, we introduced a training dataset originating from noise-induced SIR models. We tested both the trained DL algorithm and generic EWIs over the noise-induced disease models, scrutinizing their EWSs in scenarios with and without critical transitions. We also assessed a pre-trained DL algorithm from Bury et al. [20], trained on simulated data from randomly generated dynamical systems characterized by general polynomial terms. For a comprehensive evaluation, we compared the area under the ROC curve of all the indicators. In assessing the DL algorithm’s performance on real-world disease outbreaks, we employed an empirical COVID-19 cases dataset. Remarkably, the DL algorithm demonstrated significant improvement across all test models and empirical data when trained on noise-induced disease model data, outperforming other EWIs.

The performance of the classifiers, except for SIDATR-100, fell below random chance when evaluated on the empirical COVID-19 dataset from Edmonton. While the SIDATR-500 model demonstrated superior performance in transcritical simulations of empirical data, it performed poorly in null simulations. This inconsistency is associated with the short length of the empirical dataset, particularly in null simulations. The PODATR and SIDATR-500 models were trained on 500 time points, whereas the null simulations of COVID-19 consisted of less than 100 time points. Conversely, the performance notably improved with the SIDATR-100 model, which was trained on a shorter time span of 100 points. Consequently, enhancing the model’s performance on empirical data can be achieved by training it on shorter time lengths, as demonstrated by the efficacy of the SIDATR-100 model.

The initial probabilities assigned by the SIDATR-500 model were found to be inaccurate for a very short time series. However, the model’s predictive performance significantly improved as the observed time point neared the bifurcation point. This inconsistency can be explained to the DL model’s enhanced ability to provide more accurate probabilities when it is close to the transition point [40], whereas the initial predictions may include data points that are considerably distant from this transition point. This phenomenon aligns with the concept of critical slowing down in a dynamical system. Critical slowing down typically manifests when a system approaches a bifurcation point, leading to the emergence of distinctive patterns in the time series [19, 22]. However, when analyzing a lengthy time series that does not sufficiently approach the bifurcation point, these characteristic patterns of critical slowing down may not be readily apparent.

The capability of the SIDATR-500 and SIDATR-100 models to identify bifurcations in a time series beyond those on which they were specifically trained was assessed using an SEIR model with additive white noise and real-world COVID-19 cases. The SEIR model, characterized by a higher dimensionality compared to the training models, serves as a versatile framework for describing the dynamics of various infectious diseases [48], and the COVID-19 pandemic stands as a recent and significant public health challenge. The performance of generating EWSs of these DL models can be ascribed to the diverse training data deliberately

generated with a wide array of random parameters, encompassing all conceivable scenarios that may manifest in a disease outbreak. Additionally, by incorporating different types of noise with random intensity, our training data replicated the varying levels of fluctuation, mirroring the uncertainty inherent in real-world outbreaks [61]. Notably, our training data accounted for the linear increase in disease transmission rate, reflecting the common trend of initially low infection levels that intensify over time in infectious diseases [62]. Moreover, we accommodated cases where the rate increased but did not reach the threshold for a full-blown outbreak. By incorporating all these potential phenomena in the training data, the DL model gained the necessary adaptability to provide EWSs in models that closely resemble real-world scenarios.

One of the hypotheses examined in [20] was whether the DL algorithm could provide more effective EWSs compared to generic EWIs. Within the context of our study, the performance of the PODATR model both supported and contradicted this hypothesis in different test models. The training data for the PODATR model was derived from a two-dimensional random dynamical system featuring a general polynomial up to the third order and incorporating additive white noise in the simulations. The results revealed that the performance of the PODATR model was less effective than lag-1 AC, SIDATR-100, and SIDATR-500 models for the simulations of SIR with demographic noise. Notably, infectious diseases often involve a range of noise sources in their dynamics [63], and these noises can significantly impact the robustness of early warning systems [64]. Beside these noises, mathematical models are generally built on a set of underlying assumptions, and the dynamics of population compartments are contingent on these assumptions [48, 65]. As a result, certain models may exhibit distinct dynamical behavior, potentially leading to challenges for the trained model in accurately capturing trends in those cases. As evidenced by the results, the PODATR model exhibited the lowest accuracy compared to some other indicators in the SEIR model, even when considering only additive white noise.

The performance of Lag-1 AC and variance exhibited variations in the SIR model with environmental and demographic noise. Specifically, variance demonstrated superior performance in one model, while Lag-1 AC outperformed it in the other. This variability in performance can be attributed to the differing levels of noise intensities present in the simulations and different types of noises in the system [24]. The simulations used for performance assessment were generated by random noise intensity, and it is worth noting that the generic EWIs are sensitive to noise intensity [66] as well as the type of noise [25], underscoring the need to explore alternative statistical measures for noise-induced models in future research.

The SIR model is a generalized framework for characterizing a wide range of infectious diseases [48], with additive white noise, multiplicative environmental noise, and demographic noise representing common types of noise encountered in disease outbreaks [26]. Consequently, the noise-induced SIR models employed in this study serve as fundamental representations capable of capturing the dynamics of various infectious disease outbreaks. By amalgamating data from a hundred thousand simulations across all three models, we introduced a

substantial level of variability into the training dataset. This diversity equips the DL model with the ability to adapt to a spectrum of scenarios. Since DL models typically excel when applied to data that closely resembles their training dataset [40], the SIDATR-100 and SIDATR-500 models, having been trained on data closely aligned with these noise-induced models, possess the adaptability to exhibit heightened sensitivity and specificity in providing EWSs for emerging and re-emerging infectious disease outbreaks that exhibit similarities with these noise-induced models.

The training data focused on the transcritical bifurcation from a basic SIR model, where $R_0 < 1$ serves as the critical threshold for disease elimination. To incorporate other types of bifurcations, future research should consider more advanced models that account for different bifurcation scenarios. For example, backward bifurcation occurs when $R_0 < 1$ is insufficient for disease elimination [67], and Hopf bifurcation is associated with the appearance or disappearance of a limit cycle around the equilibrium point [68].

Our study has several limitations. First, despite having a maximum simulation length of 1500 time points, the DL algorithm was exclusively trained on a subset of observations consisting of 500 and 100 time points. This limited exposure to both shorter and longer time points, potentially restricting its ability to generalize across a broader range of temporal dynamics. Second, the trained DL model's sensitivity to data peaks could lead to misclassification when applied to null simulations, as these simulations do not exhibit peaks over time. Third, our training process primarily focused on linear variations in the disease transmission rate, overlooking other forms of rate variations and potential fluctuations in other parameters within the SIR model, which may be crucial in real-world scenarios [69, 70]. Fourth, our training data pertained to three specific types of noise-induced SIR models, which might limit the model's suitability for capturing transitions in models with different behaviors. Although we aimed to closely simulate real-world scenarios, it is essential to note that our model was not specifically trained on real-world data. Real-world datasets frequently face challenges such as low quality, variability, and reporting delays [71], which can introduce various types of noise and bias. Additionally, the intricate dynamics of real-world diseases, shaped by diverse environmental, social, and biological factors, may not be entirely reflected in our synthetic training data, potentially undermining the accuracy and reliability of the model. As observed, the performance on the COVID-19 dataset was comparatively lower. Future directions should explore training on empirical datasets to assess models' performance in real-world scenarios. Fifth, we explored only one DL algorithm and did not investigate the capabilities of other DL architectures. This will be a focus of future research. Sixth, it is important to acknowledge that none of the models can predict the exact timing of bifurcations, highlighting the need for further investigation in this area.

For a real-world outbreak, this method can be applied by detrending the time series of infected cases and inputting the residuals into the trained classifier to generate early warning signals. A higher probability assigned by the classifier to a specific type of bifurcation (tran-

scritical or null) serves as the signal for that bifurcation type. Therefore, when the classifier assigns a higher probability to a transcritical bifurcation than to a null bifurcation, it suggests an imminent outbreak of the disease.

This research has a significant impact on the generation of reliable EWSs for impending transitions in infectious disease outbreaks. While acknowledging the model's limitations and offering suggestions, it opens the door for future research aimed at developing even more effective early warning systems.

Ethics

This work did not require ethical approval from a human subject or animal welfare committee.

Data accessibility

Training data and codes are available in a Zenodo repository (<https://doi.org/10.5281/zenodo.10841969>) [72] under version v1 and v2, respectively. The COVID-19 data are available at the city of Edmonton's open data portal (<https://data.edmonton.ca/Community-Services/COVID-19-in-Edmonton-Daily-Active-Cases/qkyj-dqjp>) [58].

Declaration of AI use

We have not used AI-assisted technologies in creating this article.

Authors' contributions

A.K.C.: conceptualization, formal analysis, investigation, methodology, software, validation, visualization, writing-original draft, and writing-review; S.G.: conceptualization, methodology, and editing; R.M.: conceptualization, methodology, and software; P.R.: conceptualization, funding acquisition, methodology, and editing; R.G.: conceptualization, funding acquisition, methodology, software, and editing; M.A.L.: conceptualization, funding acquisition, methodology, project administration, supervision, writing-original draft, and editing; H.W.: conceptualization, funding acquisition, methodology, project administration, supervision, writing-original draft, and editing.

Conflict of interest declaration

We declare we have no competing interests.

Funding

This project (Amit K. Chakraborty, Shan Gao, Reza Miry) was primarily supported by One Health Modelling Network for Emerging Infections (OMNI), Amii (Alberta Machine Intelligence Institute) matching fund, and the Department of Mathematical and Statistical Sciences (IUSEP funding) at the University of Alberta. Hao Wang’s research was partially supported by the Natural Sciences and Engineering Research Council of Canada (Individual Discovery Grant RGPIN-2020-03911 and Discovery Accelerator Supplement Award RGPAS-2020-00090) and the Canada Research Chairs program (Tier 1 Canada Research Chair Award). Mark A. Lewis gratefully acknowledges support from an NSERC Discovery Grant and the Gilbert and Betty Kennedy Chair in Mathematical Biology. Pouria Ramazi also acknowledges the support from an NSERC Discovery Grant. Russell Greiner was also partially funded by CIFAR and NSERC.

Acknowledgement

We would like to thank Pijush Panday for the initial discussion and Tianyu Guan for providing valuable feedback.

References

- [1] Rachel E Baker, Ayesha S Mahmud, Ian F Miller, Malavika Rajeev, Fidisoa Rasambainarivo, Benjamin L Rice, Saki Takahashi, Andrew J Tatem, Caroline E Wagner, Lin-Fa Wang, et al. Infectious disease in an era of global change. *Nature Reviews Microbiology*, 20(4):193–205, 2022.
- [2] Waleed M Sweileh. Global research activity on mathematical modeling of transmission and control of 23 selected infectious disease outbreak. *Globalization and Health*, 18(1):1–14, 2022.
- [3] Emma Southall, Tobias S Brett, Michael J Tildesley, and Louise Dyson. Early warning signals of infectious disease transitions: a review. *Journal of the Royal Society Interface*, 18(182):20210555, 2021.
- [4] David E Bloom and Daniel Cadarette. Infectious disease threats in the twenty-first century: strengthening the global response. *Frontiers in immunology*, 10:549, 2019.
- [5] Nicholas LePan et al. Visualizing the history of pandemics. *Visual Capitalist*, 14:00060–20, 2020.
- [6] World Health Organization (WHO). Coronavirus (covid-19) data. <https://www.who.int/data#dashboards>, October 2023.

- [7] World Health Organization. *Managing epidemics: key facts about major deadly diseases*. World Health Organization, 2018.
- [8] Thushara Kamalrathne, Dilanthi Amaratunga, Richard Haigh, and Lahiru Kodituwakku. Need for effective detection and early warnings for epidemic and pandemic preparedness planning in the context of multi-hazards: Lessons from the covid-19 pandemic. *International Journal of Disaster Risk Reduction*, 92:103724, 2023.
- [9] Weizhong Yang. *Early warning for infectious disease outbreak: theory and practice*. Academic Press, 2017.
- [10] Zuiyuan Guo, Kevin He, and Dan Xiao. Early warning of some notifiable infectious diseases in china by the artificial neural network. *Royal Society open science*, 7(2):191420, 2020.
- [11] Sylvester Languon and Osbourne Quaye. Filovirus disease outbreaks: a chronological overview. *Virology: research and treatment*, 10:1178122X19849927, 2019.
- [12] Tobias S Brett and Pejman Rohani. Dynamical footprints enable detection of disease emergence. *PLoS biology*, 18(5):e3000697, 2020.
- [13] Liuyang Yang, Jiao Yang, Yuan He, Mengjiao Zhang, Xuan Han, Xuancheng Hu, Wei Li, Ting Zhang, and Weizhong Yang. Enhancing infectious diseases early warning: A deep learning approach for influenza surveillance in china. *Preventive Medicine Reports*, page 102761, 2024.
- [14] Mingzhang Li, Shuo Ma, and Zhengrong Liu. A novel method to detect the early warning signal of covid-19 transmission. *BMC Infectious Diseases*, 22(1):626, 2022.
- [15] Yue Gao, Guang-Yao Cai, Wei Fang, Hua-Yi Li, Si-Yuan Wang, Lingxi Chen, Yang Yu, Dan Liu, Sen Xu, Peng-Fei Cui, et al. Machine learning based early warning system enables accurate mortality risk prediction for covid-19. *Nature communications*, 11(1):5033, 2020.
- [16] Daniele Proverbio, Françoise Kemp, Stefano Magni, and Jorge Gonçalves. Performance of early warning signals for disease re-emergence: A case study on covid-19 data. *PLOS Computational Biology*, 18(3):e1009958, 2022.
- [17] Natalie Sims and Barbara Kasprzyk-Hordern. Future perspectives of wastewater-based epidemiology: monitoring infectious disease spread and resistance to the community level. *Environment international*, 139:105689, 2020.
- [18] Tobias S Brett, John M Drake, and Pejman Rohani. Anticipating the emergence of infectious diseases. *Journal of The Royal Society Interface*, 14(132):20170115, 2017.

- [19] Marten Scheffer, Stephen R Carpenter, Timothy M Lenton, Jordi Bascompte, William Brock, Vasilis Dakos, Johan Van de Koppel, Ingrid A Van de Leemput, Simon A Levin, Egbert H Van Nes, et al. Anticipating critical transitions. *science*, 338(6105):344–348, 2012.
- [20] Thomas M Bury, RI Sujith, Induja Pavithran, Marten Scheffer, Timothy M Lenton, Madhur Anand, and Chris T Bauch. Deep learning for early warning signals of tipping points. *Proceedings of the National Academy of Sciences*, 118(39):e2106140118, 2021.
- [21] Steven H Strogatz. *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*. CRC press, 2018.
- [22] Carl Boettiger, Noam Ross, and Alan Hastings. Early warning signals: the charted and uncharted territories. *Theoretical ecology*, 6:255–264, 2013.
- [23] Mathew Titus and James Watson. Critical speeding up as an early warning signal of stochastic regime shifts. *Theoretical Ecology*, 13:449–457, 2020.
- [24] Rui Liu, Pei Chen, Kazuyuki Aihara, and Luonan Chen. Identifying early-warning signals of critical transitions with strong noise by dynamical network markers. *Scientific reports*, 5(1):17501, 2015.
- [25] Suzanne M O’Regan and Danielle L Burton. How stochasticity influences leading indicators of critical transitions. *Bulletin of mathematical biology*, 80(6):1630–1654, 2018.
- [26] Linda JS Allen. *An introduction to stochastic processes with applications to biology*. CRC press, 2010.
- [27] Carl Boettiger. From noise to knowledge: how randomness generates novel phenomena and reveals information. *Ecology letters*, 21(8):1255–1267, 2018.
- [28] Linda JS Allen. A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infectious Disease Modelling*, 2(2):128–142, 2017.
- [29] Philippe Carmona and Sylvain Gandon. Winter is coming: Pathogen emergence in seasonal environments. *PLoS computational biology*, 16(7):e1007954, 2020.
- [30] George WA Constable, Tim Rogers, Alan J McKane, and Corina E Tarnita. Demographic noise can reverse the direction of deterministic selection. *Proceedings of the National Academy of Sciences*, 113(32):E4745–E4754, 2016.
- [31] Henry C Tuckwell, Laurent Toubiana, and Jean-Francois Vibert. Enhancement of epidemic spread by noise and stochastic resonance in spatial network models with viral dynamics. *Physical Review E*, 61(5):5611, 2000.

- [32] John M Drake, Tobias S Brett, Shiyang Chen, Bogdan I Epureanu, Matthew J Ferrari, Éric Marty, Paige B Miller, Eamon B O’dea, Suzanne M O’regan, Andrew W Park, et al. The statistics of epidemic transitions. *PLoS computational biology*, 15(5):e1006917, 2019.
- [33] Adjani Gama Dessavre, Emma Southall, Michael J Tildesley, and Louise Dyson. The problem of detrending when analysing potential indicators of disease elimination. *Journal of theoretical biology*, 481:183–193, 2019.
- [34] Suzanne M O’Regan, Jonathan W Lillie, and John M Drake. Leading indicators of mosquito-borne disease elimination. *Theoretical ecology*, 9:269–286, 2016.
- [35] Emma Southall, Michael J Tildesley, and Louise Dyson. Prospects for detecting early warning signals in discrete event sequence data: Application to epidemiological incidence data. *PLoS computational biology*, 16(9):e1007836, 2020.
- [36] Eamon B O’Dea and John M Drake. Disentangling reporting and disease transmission. *Theoretical Ecology*, 12:89–98, 2019.
- [37] Tobias Brett, Marco Ajelli, Quan-Hui Liu, Mary G Krauland, John J Grefenstette, Willem G van Panhuis, Alessandro Vespignani, John M Drake, and Pejman Rohani. Detecting critical slowing down in high-dimensional epidemiological systems. *PLoS computational biology*, 16(3):e1007679, 2020.
- [38] Suzanne M O’Regan and John M Drake. Theory of early warning signals of disease emergence and leading indicators of elimination. *Theoretical Ecology*, 6:333–357, 2013.
- [39] Tobias S Brett, Eamon B O’Dea, Éric Marty, Paige B Miller, Andrew W Park, John M Drake, and Pejman Rohani. Anticipating epidemic transitions with imperfect data. *PLoS computational biology*, 14(6):e1006204, 2018.
- [40] Smita Deb, Sahil Sidheekh, Christopher F Clements, Narayanan C Krishnan, and Partha S Dutta. Machine learning methods trained on simple models can predict critical transitions in complex natural systems. *Royal Society Open Science*, 9(2):211475, 2022.
- [41] Daniel Dylewsky, Timothy M Lenton, Marten Scheffer, Thomas M Bury, Christopher G Fletcher, Madhur Anand, and Chris T Bauch. Universal early warning signals of phase transitions in climate systems. *Journal of the Royal Society Interface*, 20(201):20220562, 2023.
- [42] Sergey I Nikolenko. *Synthetic data for deep learning*, volume 174. Springer, 2021.
- [43] Debbie Rankin, Michaela Black, Raymond Bond, Jonathan Wallace, Maurice Mulvenna, Gorka Epelde, et al. Reliability of supervised machine learning using synthetic data in health care: Model to preserve privacy for data sharing. *JMIR medical informatics*, 8(7):e18910, 2020.

- [44] Benjamin N Jacobsen. Machine learning and the politics of synthetic data. *Big Data & Society*, 10(1):20539517221145372, 2023.
- [45] Abba B Gumel. Causes of backward bifurcations in some epidemiological models. *Journal of Mathematical Analysis and Applications*, 395(1):355–365, 2012.
- [46] William Ogilvy Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.
- [47] Md Aminul Haque, Amit Kumar Chakraborty, and Md Anowarul Islam. Mathematical modeling and analysis of the dynamics of chikungunya in bangladesh. *Matematika*, pages 91–108, 2021.
- [48] Fred Brauer, Pauline Van den Driessche, Jianhong Wu, and Linda JS Allen. *Mathematical epidemiology*, volume 1945. Springer, 2008.
- [49] Alan C Bovik. Basic gray level image processing. In *The essential guide to image processing*, pages 43–68. Elsevier, 2009.
- [50] Axel Hutt. Additive noise-induced system evolution (anise). *Frontiers in Applied Mathematics and Statistics*, 8:879866, 2022.
- [51] Sabiha Majumder, Ayan Das, Appilineni Kushal, Sumithra Sankaran, and Vishwesha Guttal. Finite-size effects, demographic noise, and ecosystem dynamics. *The European Physical Journal Special Topics*, 230(16):3389–3401, 2021.
- [52] Qun Liu, Daqing Jiang, Ningzhong Shi, Tasawar Hayat, and Ahmed Alsaedi. Stationary distribution and extinction of a stochastic sirs epidemic model with standard incidence. *Physica A: Statistical Mechanics and its Applications*, 469:510–517, 2017.
- [53] Sébastien Lambert, Pauline Ezanno, Mathieu Garel, and Emmanuelle Gilot-Fromont. Demographic stochasticity drives epidemiological patterns in wildlife with implications for diseases and population management. *Scientific Reports*, 8(1):16846, 2018.
- [54] Thomas M Bury. ewstools: A python package for early warning signals of bifurcations in time series data. *Journal of Open Source Software*, 8(82):5038, 2023.
- [55] William S Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association*, 74(368):829–836, 1979.
- [56] Ronald Mutegeki and Dong Seog Han. A cnn-lstm approach to human activity recognition. In *2020 international conference on artificial intelligence in information and communication (ICAIIIC)*, pages 362–366. IEEE, 2020.

- [57] Andrés Vidal and Werner Kristjanpoller. Gold volatility prediction using a cnn-lstm approach. *Expert Systems with Applications*, 157:113481, 2020.
- [58] City of Edmonton’s Open Data Portal. Covid-19 in edmonton: Daily active cases. <https://data.edmonton.ca/Community-Services/COVID-19-in-Edmonton-Daily-Active-Cases/qkyj-dqjp>, 2024.
- [59] Anne Cori, Simon Cauchemez, Neil M Ferguson, Christophe Fraser, Elisabeth Dahlgvist, P Alex Demarsh, Thibaut Jombart, Zhian N Kamvar, Justin Lessler, Shikun Li, et al. Package ‘epiestim’. *CRAN: Vienna Austria*, 2020.
- [60] Rebecca K Nash, Samir Bhatt, Anne Cori, and Pierre Nouvellet. Estimating the epidemic reproduction number from temporally aggregated incidence data: A statistical modelling approach and software tool. *PLOS Computational Biology*, 19(8):e1011439, 2023.
- [61] Martin Tobias Huber, Hans Albert Braun, and Jürgen Christian Krieg. Effects of noise on different disease states of recurrent affective disorders. *Biological Psychiatry*, 47(7):634–642, 2000.
- [62] Stefan Thurner, Peter Klimek, and Rudolf Hanel. A network-based explanation of why most covid-19 infection curves are linear. *Proceedings of the National Academy of Sciences*, 117(37):22684–22689, 2020.
- [63] Yassine Sabbar, Driss Kiouach, SP Rajasekar, and Salim El Azami El-Idrissi. The influence of quadratic lévy noise on the dynamic of an sic contagious illness model: New framework, critical comparison and an application to covid-19 (sars-cov-2) case. *Chaos, Solitons & Fractals*, 159:112110, 2022.
- [64] Shanshan Qin and Chao Tang. Early-warning signals of critical transition: Effect of extrinsic noise. *Physical Review E*, 97(3):032406, 2018.
- [65] Pouria Ramazi, Arezoo Haratian, Maryam Meghdadi, Arash Mari Oriyad, Mark A Lewis, Zeinab Maleki, Roberto Vega, Hao Wang, David S Wishart, and Russell Greiner. Accurate long-range forecasting of covid-19 mortality in the usa. *Scientific reports*, 11(1):13822, 2021.
- [66] Daniele Proverbio, Alexander Skupin, and Jorge Gonçalves. Systematic analysis and optimization of early warning signals for critical transitions using distribution data. *iScience*, 2023.
- [67] Karl P Haderler and Pauline Van den Driessche. Backward bifurcation in epidemic control. *Mathematical Biosciences*, 146(1):15–35, 1997.

- [68] Brian D Hassard, Nicholas D Kazarinoff, and Yieh-Hei Wan. *Theory and applications of Hopf bifurcation*, volume 41. CUP Archive, 1981.
- [69] Jie Long, AQM Khaliq, and Khaled M Furati. Identification and prediction of time-varying parameters of covid-19 model: a data-driven deep learning approach. *International journal of computer mathematics*, 98(8):1617–1632, 2021.
- [70] Xiunan Wang, Hao Wang, Pouria Ramazi, Kyeongah Nah, and Mark Lewis. A hypothesis-free bridging of disease dynamics and non-pharmaceutical policies. *Bulletin of Mathematical Biology*, 84(5):57, 2022.
- [71] Caroline O Buckee, Maria IE Cardenas, June Corpuz, Arpita Ghosh, Farhana Haque, Jahirul Karim, Ayesha S Mahmud, Richard J Maude, Keitly Mensah, Nkengafac Villyen Motaze, et al. Productive disruption: opportunities and challenges for innovation in infectious disease surveillance. *BMJ global health*, 3(1):e000538, 2018.
- [72] Amit K. Chakraborty. Training data and Codes: An early warning indicator trained on stochastic disease-spreading models with different noises, March 2024. Zenodo, doi: 10.5281/zenodo.10841970, <https://doi.org/10.5281/zenodo.10841969>.