LEARNING DENGUE DYNAMICS THROUGH HYBRID EQUATION-GUIDED AND DATA-DRIVEN MODELS

Americo Cunha Jr
National Laboratory of Scientific Computing
Petrópolis, Brazil &
Rio de Janeiro State University
Rio de Janeiro, Brazil
`americo@lncc.br`

Emanuelle Arantes Paixão
National Laboratory of Scientific Computing
Petrópolis, Brazil
`earantes@lncc.br`

Paulo Esquef
National Laboratory of Scientific Computing
Petrópolis, Brazil
`pesquef@lncc.br`

Marcelo Rubens do Amaral
Escola Nacional de Ciências Estatísticas
Rio de Janeiro, Brazil &
Rio de Janeiro State University
Rio de Janeiro, Brazil
`mrubens@ime.uerj.br`

ABSTRACT

Learning epidemic dynamics from data is a challenging problem due to nonlinear growth, feedback mechanisms, climate-driven nonstationarity, and noisy, aggregated observations. This work investigates how different modeling strategies influence predictive behavior by combining equation-guided mechanistic models, differential-equation-inspired reduced-order representations, and purely data-driven statistical learners within a unified probabilistic forecasting framework. The resulting models span a spectrum of inductive biases, from stochastic climate-modulated differential equations to long-memory autoregressive and exogenously conditioned time-series models. Using state-level dengue incidence data from Brazil, we analyze how these approaches learn and forecast dengue dynamics across heterogeneous regions and outbreak regimes. The results show that no single model is uniformly optimal; instead, predictive skill depends on the alignment between model structure and regional dynamical behavior. We interpret this variability as a manifestation of epistemic uncertainty rather than methodological failure, and argue that model pluralism and ensemble forecasting emerge naturally as principled responses to nonstationary epidemic systems. Beyond dengue, the findings illustrate broader lessons for integrating equation-guided structure and data-driven learning in scientific machine learning for complex, climate-modulated dynamics.

1 INTRODUCTION

Learning epidemic dynamics from data is a prototypical problem of scientific machine learning (SciML) under uncertainty. At the population level, epidemic time series arise from nonlinear growth and saturation mechanisms, feedback effects mediated by immunity and behavior, and exogenous forcing from environmental and climatic variability. These dynamics are only partially observed through surveillance systems that aggregate cases across space and time, introduce reporting delays, and are subject to noise and underreporting. As a result, epidemic forecasting presents several challenges that are central to SciML research: partial observability, strong nonstationarity, structural uncertainty, and the need to balance physical or mechanistic structure with data-driven adaptability. Dengue, in particular, exemplifies these difficulties due to its strong climate sensitivity and pronounced spatial heterogeneity Araujo et al. (2026); Johansson et al. (2019); Bhatt et al. (2013).

In such settings, it is unrealistic to expect a single modeling paradigm to perform robustly across all regions, seasons, and outbreak regimes. Models endowed with strong structural assumptions may extrapolate poorly when the system enters previously unobserved regimes, whereas highly flexible statistical learners may lack interpretability or degrade under distribution shifts. This tension mirrors a broader theme in SciML: predictive performance and scientific insight depend critically on the choice of inductive bias. Differential equations, reduced-order representations, and purely data-driven models encode complementary assumptions about dynamics, memory, and forcing. Understanding how these assumptions perform across regimes is, therefore, as important as optimizing any single architecture.

From a SciML perspective, the above modeling classes can be viewed through a unified lens. Differential-equation models act as equation-guided priors over admissible dynamics, embedding causal structure such as growth, saturation, and thresholds Brauer & Castillo-Chávez (2012); Dantas et al. (2018). Reduced-order or template-based models approximate the system's behavior on a low-dimensional solution manifold, retaining interpretable structure while relaxing full state evolution. Statistical time-series models, in turn, can be interpreted as implicit operator learners that act directly in the observation space, capturing memory effects, delayed feedback, and seasonality without explicit mechanistic commitments. Under nonstationarity and limited observability, none of these inductive biases is uniformly reliable. Consequently, uncertainty-aware forecasting Soize (2017); Cunha Jr et al. (2023) and model pluralism naturally arise as principled responses to epistemic uncertainty, rather than as ad hoc model averaging strategies Johansson et al. (2019); Yamana et al. (2016).

In this work, we study the problem of learning dengue dynamics under climate modulation through a structured plural modeling approach. We consider four complementary model classes spanning the spectrum described above: (i) an equation-guided, differential-equation model with stochastic climate-modulated growth; (ii) a differential-equation-inspired reduced-order surge template; (iii) a purely endogenous statistical autoregressive model with long memory; and (iv) a statistical time-series model with explicit exogenous climate conditioning. All models produce probabilistic forecasts and are trained independently under identical data and evaluation constraints. Rather than advocating a single "best" model, our goal is to analyze how different inductive biases align with different dynamical regimes and forecasting targets, and how their comparative behavior informs uncertainty-aware prediction Soize (2017).

We ground this analysis in the context of dengue forecasting in Brazil, which provides a demanding real-world testbed due to strong climate modulation, marked regional heterogeneity, and recent extreme outbreaks. State-level dengue incidence data and associated climate covariates are used to evaluate learning and generalization under realistic operational constraints Araujo et al. (2026). Model evaluation follows a rolling-origin, expanding-window validation protocol across three successive dengue seasons, each defined from epidemiological week (EW) 41 of year Y to EW 40 of year $Y+1$, with training data spanning EW 01 2010 up to EW 25 of the corresponding cutoff year. This setup explicitly probes out-of-sample performance under progressively evolving and potentially nonstationary conditions.

The contributions of this paper are threefold. First, we present a concise SciML formulation of epidemic forecasting as a problem of learning nonstationary dynamics under uncertainty, emphasizing the role of inductive bias. Second, through a unified probabilistic evaluation, we illustrate how equation-guided, reduced-order, and data-driven models exhibit complementary strengths and failure modes across regimes. Third, we distill transferable lessons for SciML research on hybrid modeling and uncertainty quantification in climate-modulated dynamical systems. The remainder of the paper summarizes the forecasting framework, reports comparative results, and discusses implications for scientific machine learning, while detailed model descriptions are deferred to Appendix A.

2 FORECASTING FRAMEWORK AND DATA CONTEXT

2.1 DATA SOURCE

All epidemiological and climate data used in this study were obtained from the Mosqlimate platform, assembled for national-scale dengue forecasting in Brazil. We rely on the publicly released dataset accompanying the Infodengue–Mosqlimate initiative, which provides harmonized surveillance and climate records suitable for retrospective forecasting studies Coelho et al. (2024). The full data-processing pipeline, trained models, and generated forecasts are available in a public D-FENSE repository Cunha Jr et al. (2025) to ensure transparency and reproducibility.

The raw inputs consist of three comma-separated value (CSV) files: `dengue.csv`, containing weekly probable dengue cases at the municipality level; `climate.csv`, containing meteorological variables; and `map_regional_health.csv`, providing the mapping between municipalities and Brazilian states (UFs).

2.2 DATA PROCESSING PIPELINE

To obtain consistent UF-level weekly time series, we implemented a reproducible two-step MATLAB processing pipeline applied uniformly across all states and time periods. *Step 1 — Aggregation:* Raw municipality-level records are cleaned and restricted to epidemiological weeks 201001–202452 (YYYYWW, assuming 52 weeks per year). Data are aggregated from municipality to UF by epiweek, summing reported dengue cases and averaging climate variables, while retaining minimum, median, and maximum values for each climate covariate. The number of rainy days is aggregated using the weekly maximum. This step produces UF-level weekly CSV files and diagnostic plots for quality control. *Step 2 — Filtering and smoothing:* Each UF-level time series is denoised using a combination of singular value decomposition (SVD), Savitzky–Golay filtering, and light spline smoothing, followed by resampling on a weekly grid. Integer-valued fields (dengue cases and rainy days) are rounded and clipped to enforce nonnegativity. The resulting processed series are exported for modeling and visualization. This pipeline is deterministic and applied identically to all states, ensuring consistency across training and evaluation windows.

The final processed dataset is organized as a table indexed by UF and epidemiological week, with the following columns: `epiweek`, `cases`, `temp_min/med/max`, `precip_min/med/max`, `pressure_min/med/max`, `rel_humid_min/med/max`, `thermal_range`, and `rainy_days`.

2.3 FORECASTING TARGETS AND VALIDATION PROTOCOL

The forecasting target is weekly dengue incidence per UF over a single dengue season. A season is defined operationally as epidemiological week (EW) 41 of year Y through EW 40 of year $Y+1$, aligning with standard surveillance practice in Brazil. All models are evaluated under a rolling-origin, expanding-window protocol designed to probe generalization under nonstationarity.

Three validation tests are considered. Validation 1 trains on data from EW 01 2010 to EW 25 2022 and forecasts the 2022–2023 season (EW 41 2022 to EW 40 2023). Validation 2 trains from EW 01 2010 to EW 25 2023 and forecasts the 2023–2024 season. Validation 3 trains from EW 01 2010 to EW 25 2024 and forecasts the 2024–2025 season. This protocol exposes models to progressively longer histories while evaluating them on successive, potentially regime-shifting seasons, including periods of extreme incidence.

2.4 FAIRNESS AND LEAKAGE CONTROL

All four model classes considered in this work are trained independently but under identical data cutoffs, targets, and seasonal definitions. Each model produces probabilistic forecasts in the form of predictive quantiles or prediction intervals. Any preprocessing steps that could induce information leakage—such as normalization of climate variables—are computed exclusively on the training portion of each validation split and then applied unchanged to the corresponding test season. This ensures a fair comparison across heterogeneous modeling paradigms and isolates differences in performance to inductive bias rather than experimental artifacts.

3 RESULTS AND DISCUSSION

Figure 1 reports the probabilistic forecasts obtained in **Validation Test 3**, corresponding to the 2024–2025 dengue season, for a representative set of Brazilian states. Each panel displays the observed epidemic trajectory for that season (black dots), which was not available during training, together with the predictive envelopes produced by each model. The envelopes summarize central and tail behavior through nested prediction intervals, while the dots provide a direct visual reference for how the unseen epidemic unfolded relative to the forecast distributions. No single curve or statistic is intended to be read as a point prediction; rather, the figure emphasizes distributional forecasts over an entire season.

A first and central observation from Fig. 1 is that the actual trajectories frequently deviate from the median or mean forecast path, sometimes substantially so, particularly around the timing and magnitude of the epidemic peak. This behavior is not exceptional but rather intrinsic to climate-modulated, nonstationary epidemic systems. From an operational perspective, this illustrates why central tendency alone is an inadequate summary for preparedness-oriented forecasting. Public health planning is less concerned with the most likely trajectory than with plausible adverse scenarios, such as early peaks or unusually large outbreaks, which place stress on surveillance, hospital capacity, and vector-control resources. In this context, the width, shape, and asymmetry of predictive envelopes are more informative than point accuracy.

Interpreting Fig. 1 through the lens of inductive bias helps clarify the qualitative differences observed across models. The CLiDENGO forecasts (first row) exhibit smooth and often relatively wide envelopes, reflecting the constraints imposed by an equation-guided growth structure combined with stochastic parameterization. This behavior is consistent with a conservative representation of uncertainty when extrapolating under possible regime shifts, where mechanistic assumptions may only be approximately valid. The SURGE model (second row) produces more compact and symmetric envelopes centered around a canonical outbreak shape. When the timing and morphology of the realized epidemic align with historical seasonal patterns, these envelopes track the unseen trajectory closely; deviations in timing or amplitude, however, are less easily accommodated due to the reduced-order nature of the representation. The ARp forecasts (third row) typically yield narrow envelopes that closely follow historical temporal regularity. This reflects the strong reliance on long-memory endogenous structure and can be effective in stable regimes, but it also highlights sensitivity when the epidemic trajectory departs from previously observed dynamics. Finally, SARIMAX (fourth row) occupies an intermediate position, with envelopes that adapt conditionally through exogenous climate covariates. The resulting behavior varies across states, underscoring that explicit climate conditioning can add flexibility in some settings while remaining limited by statistical assumptions in others.

Across all four model classes, Fig. 1 reinforces that apparent “misses” of the median trajectory do not imply a lack of usefulness for preparedness. In many cases, the realized epidemic remains within or near the forecast envelopes for substantial portions of the season, even when peak timing or amplitude differs from the central path. For decision-makers, this type of information supports anticipatory planning based on upper-tail risk rather than retrospective accuracy. From this perspective, probabilistic forecasting is not an accessory to deterministic prediction but a fundamentally different mode of reasoning aligned with the needs of public health response.

The heterogeneous behavior observed across models and states should therefore be interpreted as a manifestation of epistemic uncertainty rather than as evidence for a single dominant approach. Different inductive biases emphasize different aspects of the underlying dynamics, and their strengths and limitations become apparent under distinct regional and seasonal conditions. This qualitative picture is consistent with broader findings from the Dengue Sprint forecasting challenge Coelho et al. (2025), in which no single model class dominates across targets and regions. In the SciML context, these results support the view that model pluralism and ensemble reasoning arise naturally from the coexistence of complementary representations of complex, nonstationary systems, and that preparedness-oriented forecasting benefits from embracing, rather than suppressing, this diversity.

4 CONCLUSIONS

This work examined dengue forecasting as a problem of learning complex, climate-modulated epidemic dynamics under nonstationarity and epistemic uncertainty, using Brazil as a large-scale and operationally relevant testbed. Rather than advancing a single forecasting architecture, we analyzed a plural modeling strategy that spans equation-guided differential-equation formulations, reduced-order DE-inspired representations, and purely data-driven statistical learners, all embedded within a probabilistic forecasting framework.

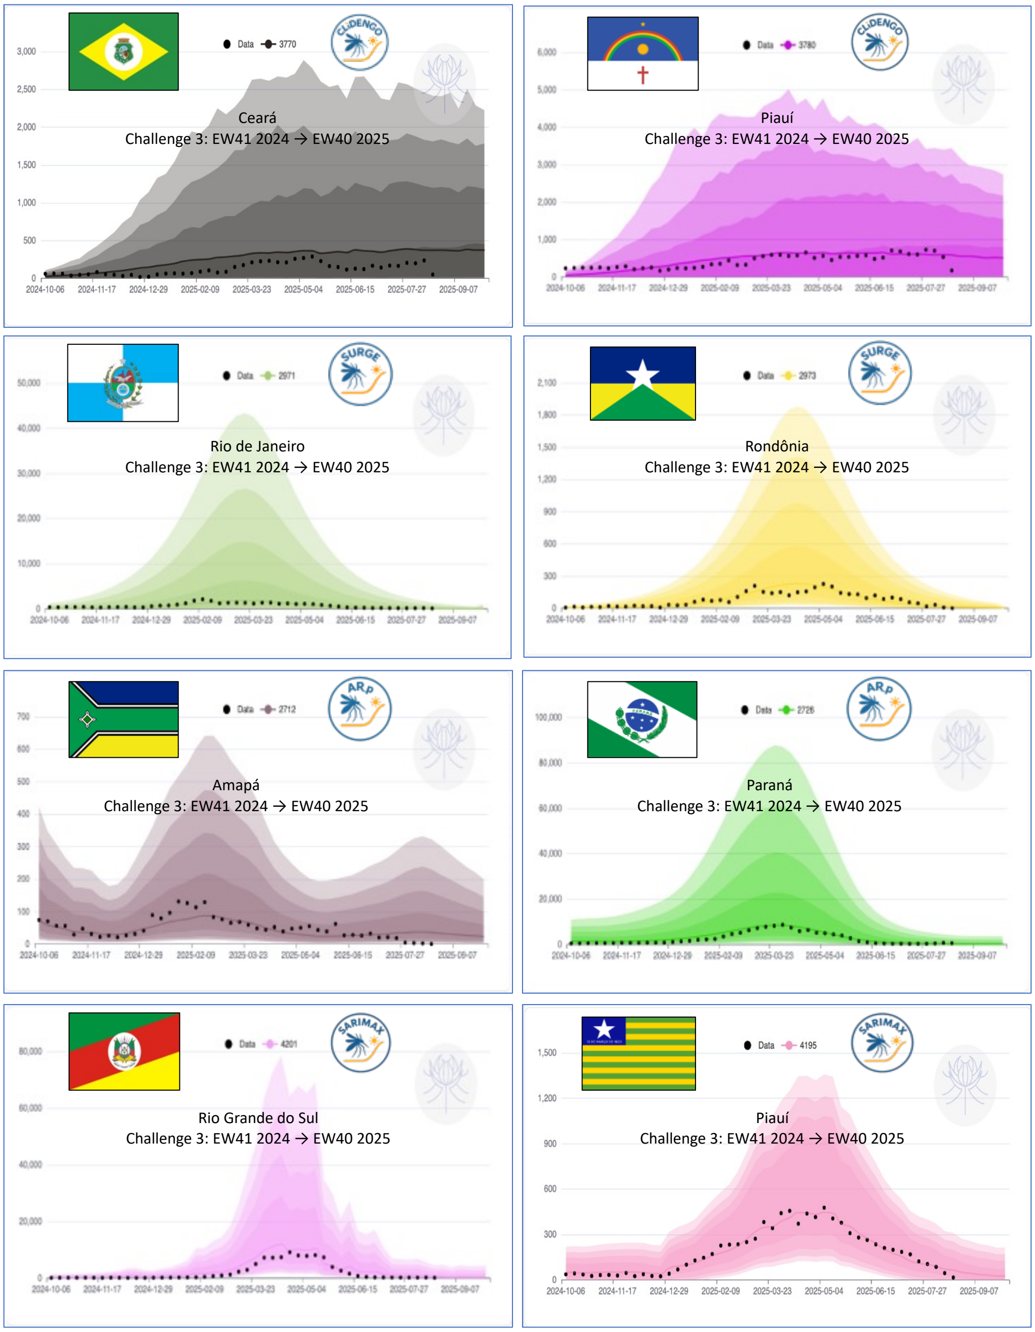

Figure 1: Representative probabilistic season-ahead dengue forecasts for four modeling approaches under the validation setting. Colored bands denote central prediction intervals (50%, 80%, 90%, and 95%), while black dots indicate the actual weekly incidence for the validation season, which was not observed during training. From top to bottom: (row 1) CLiDENGO, a stochastic climate-modulated differential-equation model; (row 2) SURGE, a differential-equation–inspired reduced-order surge template; (row 3) ARp, a high-order autoregressive statistical model; and (row 4) SARIMAX, a seasonal autoregressive model with exogenous climate inputs. The figure highlights how different inductive biases express uncertainty and capture unseen epidemic trajectories under nonstationarity.

The results highlight that no single inductive bias is uniformly reliable across regions, seasons, and outbreak regimes. Mechanistic structures provide interpretability and conservative uncertainty quantification; reduced-order models offer parsimony and robustness when seasonal patterns persist; and statistical time-series learners excel when historical temporal regularity remains informative. The observed variability in performance across models and states should therefore be understood as reflecting genuine epistemic uncertainty in the underlying system, rather than as a failure of any individual modeling approach.

A key practical implication is that central tendency is an insufficient summary for preparedness-oriented epidemic forecasting. For decision support, the ability to characterize plausible adverse scenarios through predictive envelopes is more informative than point accuracy, particularly under regime shifts and extreme outbreaks. In this sense, probabilistic forecasting and model pluralism are not ad hoc additions, but principled responses to the structural and observational uncertainties inherent in epidemic systems.

From a SciML perspective, this case study illustrates how differential equations can be used not only as explicit simulators but also as sources of inductive bias that shape reduced-order and data-driven learning. The continuum explored here—from equation-guided dynamics to implicit operator learning in observation space—offers a concrete example of how hybrid modeling can enhance robustness and interpretability without sacrificing flexibility. Although motivated by dengue, the lessons extend to other climate-driven, partially observed systems in which learning from data requires balancing structure, uncertainty, and adaptability.

ACKNOWLEDGMENTS

The authors acknowledge the financial support provided by the Coordination for the Improvement of Higher Education Personnel (CAPES) – Finance Code 001; the National Council for Scientific and Technological Development (CNPq), grant 305476/2022-0; and the Carlos Chagas Filho Research Foundation of Rio de Janeiro State (FAPERJ), grant 204.477/2024. Furthermore, the authors would like to express their gratitude to Christian Soize (Université Gustave-Eiffel), Marcelo Provenza (UERJ), Vinicius Layter (UERJ), and Marcio Borges (LNCC) for their insightful discussions and valuable contributions to the development of this work.

REFERENCES

E C Araujo, L M Carvalho, F Ganem, L B Vacaro, L S Bastos, L P Freitas, M Bastos, R Alencar, L Bianchi, R Capellán, X Chen, O Cruz, A Cunha Jr, H K Das, C Fletcher, R M Lana, R Lowe, D Lührsen, G Moirano, P Moraga, L M Stolerman, F Valente, C T Codeço, and F C Coelho. Leveraging probabilistic forecasts for dengue preparedness and control: the 2024 dengue forecasting sprint in brazil. *Proceedings of the National Academy of Sciences*, 123(7): , 2026. doi: ~.

S Bhatt, P W Gething, O J Brady, J P Messina, A W Farlow, C L Moyes, J M Drake, J S Brownstein, A G Hoen, O Sankoh, M F Myers, D B George, T Jaenisch, G R W Wint, C P Simmons, T W Scott, J J Farrar, and S I Hay. The global distribution and burden of dengue. *Nature*, 496(7446): 504–507, 2013. doi: 10.1038/nature12060.

F Brauer and C Castillo-Chávez. *Mathematical Models in Population Biology and Epidemiology*. Springer, New York, 2nd edition, 2012. doi: 10.1007/978-1-4614-1686-9.

F C Coelho, C T Codeço, I Almeida, L M Carvalho, E C Araújo, L Bastos, L B Vacaro, and R M Lana. Full dataset for dengue forecasting in brazil for infodengue-mosqlimate sprint 2024, August 2024. URL `https://doi.org/10.5281/zenodo.13328231`.

F C Coelho, C T Codeco, E C Araujo, F S Ganem dos Santos, I Ferreira de Almeida, L Bastos, L M Carvalho, O Cruz, R Martins Lana, S S Oliveira, T Riback, L Vacaro, L Alves, V Godinho, M Bastos, J Souza, B Laiate, L Monteiro Bianchi, L Freitas, D Andreza da Cruz Ferreira, A Gouveia, R Alencar, D Sales Barreira, Zuilho Segundo, and S Rocha. Relatório imdc 2025, November 2025. URL `https://doi.org/10.5281/zenodo.17516518`.

A Cunha Jr, D A W Barton, and T G Ritto. Uncertainty quantification in mechanistic epidemic models via cross-entropy approximate bayesian computation. *Nonlinear Dynamics*, 111:9649–9679, 2023. doi: 10.1007/s11071-023-08327-8.

A Cunha Jr, E A Paixão, M M Provenza, M R S Amaral, M R Borges, P A A Esquef, S Luque, T M Porcino, and V L Xavier. D-FENSE: Dynamics for Epidemic Surveillance and Evaluation. `https://github.com/americocunhajr/D-FENSE`, 2025.

E Dantas, M Tosin, and A Cunha Jr. Calibration of a SEIR-SEI epidemic model to describe Zika virus outbreak in Brazil. *Applied Mathematics and Computation*, 338:249–259, 2018. doi: 10.1016/j.amc.2018.06.024.

M A Johansson, K M Apfeldorf, S Dobson, J Devita, A L Buczak, B Baugher, L J Moniz, T Bagley, S M Babin, E Guven, T K Yamana, J Shaman, T Moschou, N Lothian, A Lane, G Osborne, G Jiang, L C Brooks, D C Farrow, S Hyun, R J Tibshirani, R Rosenfeld, J Lessler, N G Reich, D A T Cummings, S A Lauer, S M Moore, H E Clapham, R Lowe, T C Bailey, M García-Díez, M S Carvalho, X Rodó, T Sardar, R Paul, E L Ray, K Sakrejda, A C Brown, X Meng, O Osoba, R Vardavas, D Manheim, M Moore, D M Rao, T C Porco, S Ackley, F Liu, L Worden, M Convertino, Y Liu, A Reddy, E Ortiz, J Rivero, H Brito, A Juarrero, L R Johnson, R B Gramacy, J M Cohen, E A Mordecai, C C Murdock, J R Rohr, S J Ryan, A M Stewart-Ibarra, D P Weikel, A Jutla, R Khan, M Poultney, R R Colwell, B Rivera-García, C M Barker, J E Bell, M Biggerstaff, D Swerdlow, L Mier-y Teran-Romero, B M Forshey, J Trtanj, J Asher, M Clay, H S Margolis, A M Hebbeler, D George, and J P Chretien. An open challenge to advance probabilistic forecasting for dengue epidemics. *Proceedings of the National Academy of Sciences*, 116(48):24268–24274, 2019. doi: 10.1073/pnas.1909865116.

C Soize. *Uncertainty Quantification: An Accelerated Course with Advanced Applications in Computational Engineering*. Springer International Publishing, 2017.

T K Yamana, S Kandula, and J Shaman. Superensemble forecasts of dengue outbreaks. *Journal of the Royal Society Interface*, 13(123):20160410, 2016. doi: 10.1098/rsif.2016.0410.

A APPENDIX

This appendix provides detailed descriptions of the four forecasting models used in this study. The models span a spectrum of inductive biases, from equation-guided mechanistic dynamics to purely statistical time-series learning. All models are trained and evaluated under the common data-processing and validation protocol described in Section 2, and all produce probabilistic forecasts in the form of predictive quantiles or prediction intervals.

A.1 CLIDENGO: CLIMATE-MODULATED STOCHASTIC DIFFERENTIAL EQUATION MODEL

CLiDENGO (*CLimate Logistic DENGue Outbreak Simulator*) is a mechanistic, equation-guided model that represents epidemic evolution through a stochastic, climate-modulated growth law. The state variable is the cumulative number of reported dengue cases, denoted by $C(t)$, from which weekly incidence is obtained as its rate of change.

The epidemic dynamics are governed by a β-logistic growth equation,

$$\frac{dC(t)}{dt} = r(t)\, C(t)^q \left[1 - \left(\frac{C(t)}{K}\right)^{\alpha}\right]^p , \tag{1}$$

where $r(t)$ is an effective growth rate, K is the final epidemic size, q controls early-time growth, α governs asymmetry, and p controls late-time saturation. The baseline growth scale r_0 is modulated by climate through

$$r(t) = r_0\, B(T_t, P_t, H_t)\,, \tag{2}$$

where T_t, P_t, and H_t denote temperature, precipitation, and relative humidity, respectively, and $B(\cdot) \in [0,1]$ is a smooth suitability function. Each climatic factor is mapped to $[0,1]$ using bounded, biologically motivated response curves (of Brière type), and the effective modulation may involve temperature alone or multiplicative combinations of temperature, precipitation, and humidity. Integer-valued lags are allowed between climate covariates and epidemiological response to account for delayed effects.

Uncertainty is represented through stochastic parameterization rather than demographic noise. Around calibrated mean values, parameters are randomized per Monte Carlo realization using simple distributions: r_0 follows a Gamma distribution; p follows a shifted Gamma distribution bounded below by one; K, q, and α are sampled from uniform distributions on symmetric intervals around their means; and the initial condition C_0 is sampled empirically from the pool of historical season onsets. For each realization, Eq. 1 is integrated weekly over a 52-week horizon using a standard ODE solver. An ensemble of trajectories yields probabilistic forecasts summarized by medians and central prediction intervals (50%, 80%, 90%, and 95%).

CLiDENGO encodes a strong causal inductive bias through explicit growth and saturation mechanisms and interpretable climate forcing. This structure supports scenario exploration but may be sensitive to structural mismatch under extreme or unprecedented outbreak regimes.

A.2 SURGE: DIFFERENTIAL-EQUATION-INSPIRED REDUCED-ORDER TEMPLATE MODEL

SURGE is a reduced-order forecasting model explicitly grounded in the structure of a logistic-type differential equation, while avoiding full state-space simulation. The starting point is the canonical logistic growth equation for cumulative cases $C(t)$,

$$\frac{dC(t)}{dt} = k\,C(t)\left(1 - \frac{C(t)}{L}\right), \tag{3}$$

where L denotes the final epidemic size and k the intrinsic growth rate. Rather than modeling the full cumulative dynamics, SURGE operates directly on the implied incidence curve, which corresponds to the right-hand side of Eq. 3 evaluated along its solution trajectory.

By substituting the closed-form solution of Eq. 3 into its right-hand side, the weekly incidence profile takes the form

$$S(t) = \frac{dC(t)}{dt} = \frac{Lk\,e^{-k(t-t_0)}}{\left(1 + e^{-k(t-t_0)}\right)^2}, \tag{4}$$

where t_0 denotes the time of peak incidence. Equation 4 defines a unimodal, symmetric surge whose shape is entirely determined by (L, k, t_0). In SURGE, this expression is interpreted as a *template* for epidemic surges, representing a low-dimensional solution manifold inherited from the underlying differential equation.

For each state (UF), historical seasons are aligned by their observed peak weeks and averaged to obtain a typical, centralized surge. The parameters (L, k, t_0) of Eq. 4 are then estimated via nonlinear least squares. Season-to-season variability is modeled through a positive amplitude gain G multiplying the template,

$$S_{\text{season}}(t) = G\,S(t), \tag{5}$$

where G captures interannual variability in outbreak magnitude. The set of gains inferred from historical seasons is assumed to follow a log-normal distribution, $G \sim \log\mathcal{N}(\mu_G, \sigma_G^2)$.

Probabilistic forecasts are generated by sampling gains from this distribution and applying them to the surge template placed at the appropriate seasonal timing. Ensemble quantiles yield median forecasts and central prediction intervals. By explicitly deriving its template from a differential equation, SURGE preserves mechanistic shape information while reducing the forecasting problem to low-dimensional stochastic variability in amplitude. This makes the model robust and interpretable when outbreak timing and shape are stable, while naturally limiting its flexibility under strongly asymmetric or multi-peak epidemic regimes.

A.3 ARP: HIGH-ORDER AUTOREGRESSIVE STATISTICAL MODEL

The ARp model is a purely data-driven approach based on a high-order autoregressive process that exploits long-memory temporal structure in dengue incidence. Let C_t denote weekly case counts. After a variance-stabilizing log transformation, the dynamics are modeled as

$$C_t = \sum_{i=1}^{p} \phi_i C_{t-i} + \varepsilon_t, \tag{6}$$

with autoregressive order $p = 92$, chosen empirically to span multiple seasonal cycles. The coefficients $\{\phi_i\}$ are estimated using covariance-based methods.

Stochasticity enters through the innovation term ε_t, assumed to be Gaussian white noise with variance calibrated from historical residuals organized in yearly blocks. Forecasting proceeds by Monte Carlo simulation with fixed autoregressive coefficients and randomized innovations. Predictions are mapped back to the original count scale, smoothed using singular spectrum analysis, and cropped to the target seasonal window. Prediction intervals are obtained from empirical quantiles of the simulated ensemble.

From a dynamical-systems perspective, ARp can be interpreted as learning a finite-memory approximation of the evolution operator acting on observed incidence. The model is robust under strong temporal regularity but lacks explicit causal or exogenous structure, which may limit extrapolation under regime shifts.

A.4 SARIMAX: SEASONAL STATISTICAL MODEL WITH EXOGENOUS CLIMATE INPUTS

SARIMAX extends classical seasonal autoregressive integrated moving-average models by incorporating exogenous climate covariates. The transformed response $y_t = \log(C_t + 100)$ is modeled as

$$\Phi(B^s)\phi(B)\nabla^d\nabla_s^D y_t = c + \Theta(B^s)\theta(B)\varepsilon_t + \boldsymbol{\beta}^\top \mathbf{X}_t, \tag{7}$$

where B is the backshift operator, $s = 52$ denotes weekly seasonality, and $\mathbf{X}_t$ contains exogenous regressors. The regressors include the weekly median temperature and a 52-week rolling mean of precipitation, capturing accumulated rainfall effects.

Model orders are selected per validation window using automated procedures for guidance, followed by manual adjustment to favor parsimonious specifications suitable for long-horizon forecasting. Future climate inputs are supplied using seasonal-naive extrapolation based on previous years. Predictive uncertainty is quantified on the transformed scale using Gaussian assumptions and propagated through inverse transformation with truncation at zero to enforce nonnegativity.

SARIMAX explicitly separates endogenous temporal dependence from exogenous climate forcing, offering a controlled statistical counterpart to mechanistic climate modulation. Its performance depends on the stability and relevance of learned climate–incidence relationships across seasons.

A.5 SYNTHESIS

Together, CLiDENGO, SURGE, ARp, and SARIMAX span a continuum of inductive biases, from equation-guided mechanistic structure to purely statistical learning with endogenous and exogenous conditioning. All models yield probabilistic forecasts compatible with ensemble evaluation, enabling a controlled analysis of how structural assumptions influence learning and generalization under climate-driven nonstationarity.