Multiscale forecasting of dengue alerts from notifications, historical seasonality, and socio-spatial vulnerability

Silmara Barnabé[1], **Maximus Borges da Rosa**[1]
Guilherme Malta[1], **Weverton Luis da Costa Cordeiro**[1]
Carla M.D.S. Freitas[1], **Mariana Recamonde-Mendoza**[1,2]

[1]Institute of Informatics – Federal University of Rio Grande do Sul (UFRGS)
Porto Alegre – RS – Brazil

[2]Bioinformatics Core, Hospital de Clínicas de Porto Alegre (HCPA)
Porto Alegre – RS – Brazil

{silmara.barnabe,maxborgesrosa,gbbomalta}@gmail.com
{weverton.cordeiro,carla,mrmendoza}@inf.ufrgs.br

Abstract. *Suspected dengue notifications precede laboratory confirmation and are an early proxy for healthcare demand. We frame surveillance as alert prioritization: ranking municipality-weeks by the risk of high notification burden over the next four weeks, with the alert set per municipality at the 80th percentile of training burden. Following CRISP-DM, we build a multiscale weekly dataset of notification history, acceleration, seasonal anomaly, climate, and socio-spatial vulnerability using training data from 2010 to 2024. On the 2025 out-of-time test, the best classifier reached an* Average Precision *(AP) of* 0.912 *and precision@top20% of* 0.941*, against* 0.837 *for a strong seasonal baseline; the gain held in 2024–2025 (AP* 0.938 *vs.* 0.841*), while finer UDH/locality granularity did not improve on the municipal model.*

1. Introduction

Dengue is one of the main urban arboviruses in Brazil [Zhang et al. 2025], with seasonal epidemics that strain primary care and hospital services [Wilder-Smith et al. 2017]. For operational planning, confirmed-case series become available late, since they depend on laboratory, clinical, or epidemiological case closure. In contrast, suspected notifications, recorded before confirmation, serve as an earlier signal of demand, and predictive systems can turn them into operational alerts [Salim et al. 2024]. The predictive target is therefore the count of suspected notifications, not confirmed cases or outbreak labels. The central question is not "how many cases will occur", but which municipality-weeks to prioritize based on a higher risk of high notification burden in the following weeks.

Much of the related literature operates at the municipal, state, or regional scale [Islam et al. 2025], which can mask intra-urban heterogeneity in income, sanitation, density, and social vulnerability. Human Development Units (Unidades de Desenvolvimento Humano, UDH) represent part of this heterogeneity; here, UDH/locality information enters as a socio-spatial context aggregated to the municipality and as intra-municipal statistics of the local distribution of notifications.

This study aims to develop a spatio-temporal predictive model for prioritizing municipality-weeks at high risk of elevated dengue notification burden over a four-week

forecasting horizon, using suspected notifications and integrating epidemiological, climatic, socioeconomic, healthcare, and intra-urban contextual information. Beyond a predictive model, this work delivers a reproducible *pipeline* that enables testing epidemiological and operational hypotheses. Its main contributions are: (i) a weekly dataset, at the UDH/locality-week and municipality-week grains, integrating notifications, seasonality, acceleration, climate, epidemiological calendar, establishments from the Brazilian National Registry of Health Establishments (CNES), and socioeconomic indicators; (ii) the formulation as a four-week alert prioritization, with complementary future burden regression; (iii) strict temporal evaluation (training 2010–2024, test 2025) with robustness in 2024–2025; and (iv) an ablation of the history, acceleration, seasonality, socioeconomic, and UDH/locality blocks.

2. Related Work

Dengue forecasting and environmental factors. Dengue forecasting models typically combine surveillance and climate data. In Brazil, there are alert systems based on climate data, spatial risk, and probabilistic forecasts [Lowe et al. 2011, Lowe et al. 2013, Lowe et al. 2016]. Studies provide evidence that temperature, precipitation, urbanization, and connectivity contribute to the expansion and intensity of transmission [Barcellos and Lowe 2014, Lowe et al. 2021, Lee et al. 2021], with temperature potentially reducing the generation interval [Siraj et al. 2017]. Recent work applies deep learning to the Brazilian series: LSTM networks with lagged climate variables and spatial effects between states, interpreted with SHAP [Chen and Moraga 2025], later extended with human mobility for outbreak detection [Chen and Moraga 2026]. Reporting delays have also been explicitly treated in forecasts of dengue hospitalizations supported by digital signals [Palmer et al. 2026, Borges et al. 2026]. Together, these studies motivate temporal windows, lags, and climate, but operate at an aggregated scale, whereas we explore features derived from UDH/locality.

Social vulnerability. Social determinants, such as income, sanitation, and infrastructure, can modulate dengue risk and its detection, but in predictive models, they usually enter at the municipal level, reducing intra-urban contrasts. A recent study showed that the risk of dengue death within the first 15 days after symptoms is higher under unfavorable socioeconomic conditions in Brazil [Cardim et al. 2025]. Extending our preliminary monthly UDH panel [Barnabé et al. 2025], we evaluate whether signals of intra-municipal heterogeneity contribute to alert prioritization in an out-of-time evaluation.

Machine learning for decision support in surveillance. Regarding techniques, machine learning (ML) has been applied to several dengue decision problems: risk and outbreak forecasting [Aleixo et al. 2022], prediction of case acceleration in a subtropical capital [de Oliveira et al. 2025], low-cost clinical triage for dengue prognosis [Rosa et al. 2026], audio-based detection of *Aedes aegypti* [Fernandes et al. 2021], entomological and meteorological models [Ong et al. 2023], and estimation of hospital burden [Wiratmadja et al. 2018]. Public SINAN-derived datasets have also supported patient-level predictive studies [da Silva Neto et al. 2022]. Our work differs by focusing on the operational prioritization of suspected notifications over a four-week horizon, with multiscale evaluation and explicit control of temporal leakage.

3. Methodology

3.1. CRISP-DM-Guided Process

The study follows the *Cross-Industry Standard Process for Data Mining* (CRISP-DM) [Wirth and Hipp 2000]. Business understanding defines the operational proxy — prioritizing municipality-weeks with high future notification burden —. Data, modeling, and evaluation phases are detailed in the following subsections. Deployment with a dashboard and a surveillance routine remains a future step.

3.2. Data Sources

The dataset integrates suspected dengue notifications from the Brazilian Notifiable Diseases Information System (SINAN), CNES units, indicators from the Social Vulnerability Atlas (AVS), territorial division from the Brazilian Institute of Geography and Statistics (IBGE), climate data from the Meteorological Database for Teaching and Research (BDMEP) of the Brazilian National Institute of Meteorology (INMET) and the epidemiological calendar. The panel was reconstructed from 2010 onward, allowing the computation of lags, moving averages, and seasonal statistics before the train/test cut.

The pilot scope considers five capitals from distinct macroregions and epidemiological regimes, which differ in intra-urban granularity: Rio de Janeiro (Southeast; 719 UDHs), Porto Alegre (South; 229), Manaus (North; 157), Recife (Northeast; 145), and Cuiabá (Central-West), which has no UDH subdivision in the Atlas. Considering municipal fallback rows, the panel contains about 1,250 UDHs plus 1,255 localities.

The socioeconomic indicators include the Social Vulnerability Index (IVS), the Municipal Human Development Index (IDHM), and measures of urban infrastructure, income, and vulnerability. As they are structural and static at the weekly scale, they are interpreted as a territorial context, not as short-term causal variables. In the integrated data model, notifications (SINAN) constitute the central fact table, linked by municipality, locality/UDH, and calendar keys to the CNES, AVS, territorial division, and climate (BDMEP/INMET) dimensions.

3.3. Granularity and Feature Engineering

The integrated data are transformed into modeling matrices at two levels, the municipality-week being the main one for training and evaluation.

3.3.1. Municipality-week and locality-week panels

The locality-week matrix aggregates suspected notifications by locality and epidemiological week, and is used to compute local signals (history, active localities, mean local burden, percentiles of the intra-municipal distribution, and territorial concentration), later summarized to the municipality-week level. The municipality-week matrix is the main evaluation grain: municipal notifications are reconciled with the sum of notifications by locality, and the municipal targets are checked against the sum of local targets when there is territorial correspondence. This grain was adopted because it is more actionable for municipal planning and because it reduces the sparsity observed in localities/UDHs.

Regarding missing values, the absence of notifications in the time grid is treated as zero; climate variables use the municipal weekly aggregation, with state fallback when there is no sufficient municipal station; socioeconomic and territorial indicators, when available at finer granularity, are aggregated to the municipality by mean, minimum, maximum, and dispersion. The remaining numeric variables are imputed by the median, and the categorical ones by the mode, with ordinal encoding, within the model pipeline itself (Section 3.4).

3.3.2. Temporal, seasonal, and socio-spatial features

For each weekly origin t, the features use only information available up to that week. Recent statistics are computed over windows of 2, 4, 8, 12, 26, and 52 weeks — from very short-term dynamics to the annual cycle — and include lags, moving sums and averages, maxima, standard deviations, and growth ratios. The *acceleration* measures the recent variation of the growth rate (difference between successive growths over 2–4, 4–8, and 8–12 week windows), distinguishing an outbreak that is intensifying from one that is decelerating. The *seasonal anomaly* is the deviation between the level observed at t and the typical level of the same epidemiological week in previous years (in absolute and relative terms). The seasonal statistics use only years prior to the origin, within the history since 2010, avoiding temporal leakage; at the locality/UDH grain a reduced version is adopted, with up to four previous years of the same epidemiological week.

The socio-spatial features summarize, at the municipal level, the IVS, the IDHM, and indicators of urban infrastructure, income, and vulnerability. As they exist at fine granularity (locality/UDH), they are aggregated by mean, minimum, maximum, and dispersion; the *minimum intra-municipal IVS*, for example, captures the least vulnerable locality, while the dispersion reflects internal inequality. Structural variables are also included, such as the *number of localities* with records (a measure of territorial fragmentation) and the number of CNES establishments. Table 1 summarizes the main variable groups. Raw target-date variables, future targets, and `t_plus` components are removed from the feature sets.

Table 1. Feature groups considered as model input.

Group	Examples	Aggregation/scale
Notification history	lags, moving sums, moving averages, maxima, and standard deviations	municipality-week and locality-week
Epidemic dynamics	growth ratios, recent slope, acceleration, and deceleration	computed up to the origin week t
Historical seasonality	mean, maximum, and anomaly relative to the same epidemiological week in previous years	same epid. week in previous years (since 2010)
Intra-municipal heterogeneity	active localities, mean local burden, percentiles, and territorial concentration	locality/UDH aggregated to municipality
Socio-spatial context	IVS, IDHM, sanitation, income, urban infrastructure, and vulnerability	structural indicators aggregated to municipality
Health infrastructure	number and distribution of CNES establishments	municipal aggregation
Calendar	epidemiological week, month, and seasonal encodings	municipality-week

3.4. Predictive Tasks

The main task is alert classification. Let $y_{u,t}$ be the count of suspected notifications in municipality u and epidemiological week t. The four-week future burden is defined as:

$$B_{u,t} = \sum_{k=1}^{4} y_{u,t+k}. \tag{1}$$

The alert label is defined per municipality, using the 80th percentile of the future burden in the training set:

$$a_{u,t} = \mathbf{1}\left[B_{u,t} \geq \tau_u\right], \qquad \tau_u = Q_{0.80}\left(\{B_{u,t} : t \in \mathcal{T}_{train}\}\right). \tag{2}$$

As a complementary task, we fit a regressor to directly predict $B_{u,t}$. Three scale strategies are compared: (i) a direct municipal model; (ii) a municipal model with features derived from locality/UDH; and (iii) a model trained on locality/UDH with the local future-burden predictions summed to obtain the municipal prediction (*post-aggregation*).

3.4.1. Multiscale aggregation

Notifications and future-burden targets are aggregated by summation; the locality/UDH features are summarized by descriptive statistics (sum, mean, minimum, maximum, standard deviation, percentiles, and concentration across localities) and the socioeconomic indicators by mean, minimum, maximum, and dispersion.

3.4.2. Models and hyperparameters

For both tasks (i.e., alert classification and four-week burden regression), we adopt *gradient boosting* [Friedman 2001] with histogram-based trees, efficient on medium-sized tabular datasets [Ke et al. 2017]: the `HistGradientBoostingClassifier` and `HistGradientBoostingRegressor` implementations from the `scikit-learn` library [Pedregosa et al. 2011]. The choice is due to its ability to model nonlinear relationships and interactions among temporal, climatic, and socio-spatial variables in heterogeneous tabular data.

No systematic hyperparameter search was performed; a fixed configuration was adopted, defined in preliminary experiments, and kept constant across scenarios to favor comparability. The classifier used `learning_rate=0.05`, `max_iter=300`, `max_leaf_nodes=15`, `l2_regularization=0.10`, and `random_state=42`; the regressor used the same configuration, with `max_iter=350` and `loss=squared_error`. These values should not be interpreted as global optima, and a systematic optimization (for example, grid or Bayesian search under temporal validation) remains as future work.

3.4.3. Baselines

As a reference, we compare the models against deterministic *baselines*, without training, computed from the history of each municipality-week using only information available up to the origin. In alert prioritization, the risk score is given by one of these rules: level

at the origin, sum of the last four weeks, four-week moving average, recent acceleration, and seasonal anomaly. In regression, the future burden is estimated by $4\times$ the level at the origin, by the sum of the last four weeks, or by $4\times$ the four-week moving average. These are simple but strong references — especially the seasonal anomaly — so that the models' gains are measured against temporal/seasonal rules. We report the best baseline for each metric.

3.5. Evaluation Protocol

The main scenario uses training between 2010 and 2024, and an out-of-time test in 2025, with 255 municipality-weeks (the five capitals over the weeks of 2025), of which 77 are positive alerts according to the municipal threshold from training. Since the alert is the 80th percentile of the future burden in training, the imbalance is controlled by construction (about 20% positives in training), even though the observed test rate varies with epidemic intensity. Neither resampling nor class reweighting was applied. As temporal robustness, a second scenario trains up to 2023 and tests jointly on 2024–2025. To quantify the variability of the metrics, we also adopt a rolling-origin *temporal cross-validation*: in each fold, the model is trained on the previous target years and validated on the following year (2021–2025), recomputing the alert threshold only on the training set of each fold (Section 4.2).

For classification, we report *Average Precision* (AP), area under the receiver operating characteristic curve (ROC-AUC), Brier score, *precision@top20%* (P@20), and *recall@top20%* (R@20). These metrics evaluate the *ranking* and the ability to identify alert weeks (the minority class), not accuracy under a fixed threshold. AP is prioritized because it summarizes the precision–recall curve and is more informative than ROC-AUC under class imbalance; the top-20 metrics simulate the selection of the 20% highest-risk municipality-weeks, measuring the precision and the coverage of this prioritized list; and the Brier score evaluates the calibration of the probabilities. For regression, we report the *Weighted Absolute Percentage Error* (WAPE), as it is more stable than MAPE in weekly series with low or null counts.

4. Results

The training set (2010–2024) contains 3,895 municipality-weeks, and the rate of positive alerts rises from the 20% imposed by construction in training to 30% in the 2025 test, given the higher epidemic intensity. The threshold τ_u (80th percentile of the four-week future burden in training) is specific to each municipality and varies strongly (from about 132 notifications in Porto Alegre to 2,625 in Rio de Janeiro), which reaffirms the alert as a threshold relative to local history.

4.1. Alert prioritization

Table 2 compares the candidate models and the baselines on the 2025 test set. The best model is a *histogram gradient boosting* classifier at the municipality-week grain, with recent history, acceleration, seasonal anomaly, and socioeconomic indicators (set `mun_history_accel_seasonal_socio_geo`). It achieves the highest AP (0.912) and the highest P@20 (0.941), outperforming the best baseline (seasonal anomaly, AP 0.837). The direct municipal variants are close to each other and above the baselines, while the *post-aggregation* by UDH/locality falls below both.

Table 2. Alert prioritization on the 2025 test set. The selected final model is highlighted in bold.

Model / feature set	AP	ROC-AUC	P@20	R@20	Brier
Direct municipal model (gradient boosting)					
Hist.+accel.+seasonal+socioecon.	**0.912**	**0.952**	**0.941**	**0.623**	**0.094**
Hist.+accel.+seasonal	0.908	0.948	0.922	0.610	0.091
Full multiscale	0.903	0.953	0.922	0.610	0.088
Hist.+accel.+seasonal+UDH	0.903	0.955	0.922	0.610	0.087
Hist.+accel.	0.895	0.943	0.941	0.623	0.102
Hist. (only)	0.884	0.934	0.922	0.610	0.104
Post-aggregation (locality/UDH → municipality)					
Hist.+accel.+seasonal (local)	0.753	0.842	0.725	0.481	0.190
Deterministic baselines					
Seasonal anomaly (best)	0.837	0.911	0.843	0.558	0.161
Level at origin	0.720	0.814	0.706	0.468	0.202
Acceleration	0.422	0.444	0.392	0.260	0.358

Figure 1 shows that the model dominates the baseline over much of the precision–recall curve; Figure 2 is analysed in Section 4.4. Selecting $\lfloor 0.20 \times 255 \rfloor = 51$ municipality-weeks among the 77 real alerts implies a recall ceiling of 0.662; the best model reaches 0.623 (about 94% of it) with a P@20 of 0.941.

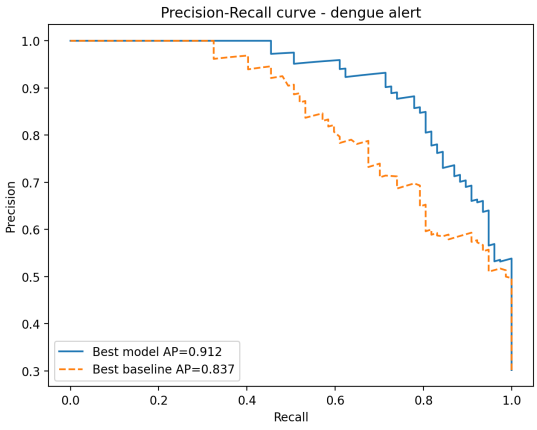


Figure 1. Precision–recall curve on the 2025 test set.

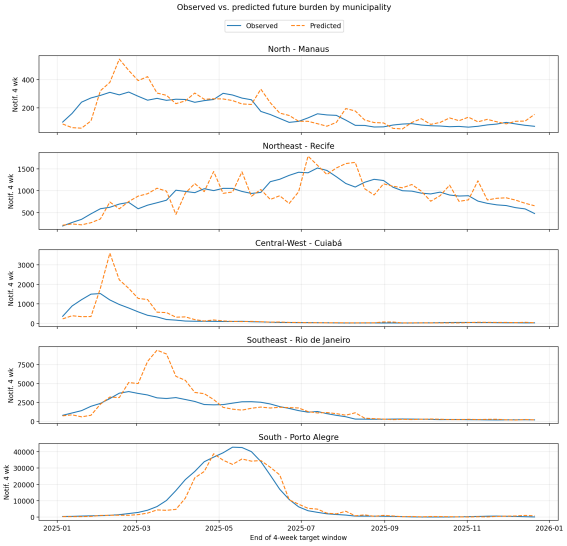


Figure 2. Observed versus predicted future burden, by municipality and macroregion.

4.2. Temporal cross-validation

Table 3 summarizes the rolling-origin temporal cross-validation (five annual folds, 2021–2025). The best model maintains an AP of 0.905 ± 0.039, close to the 2025 value (0.912), and outperforms the best baseline (seasonal anomaly, AP 0.838 ± 0.044) in all folds; the set with UDH features has a lower and more variable AP (0.882 ± 0.052), consistent with the ablation. The higher standard deviation in P@20/R@20 reflects the small size of each fold, while the global ranking (AP and ROC-AUC) remains stable.

4.3. Temporal robustness in 2024–2025

To verify that performance did not depend only on the 2025 hold-out, we repeated the evaluation with training up to 2023 and a joint test in 2024–2025, expanding the test to

Table 3. Temporal CV (mean ± std, five annual folds); best in bold.

Model / set	AP	ROC-AUC	P@20	R@20	Brier
Hist.+accel.+seasonal+socioecon.	**0.905 ± 0.039**	**0.965 ± 0.013**	**0.770 ± 0.252**	**0.728 ± 0.183**	**0.068 ± 0.025**
Hist.+accel.+seasonal	0.892 ± 0.041	0.959 ± 0.012	0.758 ± 0.245	0.719 ± 0.184	0.070 ± 0.025
Hist.+accel.+seasonal+UDH	0.882 ± 0.052	0.955 ± 0.017	0.758 ± 0.245	0.719 ± 0.184	0.078 ± 0.031
Seasonal anomaly (baseline)	0.838 ± 0.044	0.933 ± 0.036	0.693 ± 0.202	0.680 ± 0.237	0.193 ± 0.048

515 municipality-weeks and the positive rate to 0.365. Keeping the same best model and the same best baseline (seasonal anomaly), the model preserved its advantage in AP (Table 4); the P@20 also remained higher (0.961 vs. 0.893) and R@20 rose from 0.489 to 0.527. The gain in alert prioritization, therefore, holds in a broader temporal window.

Table 4. Temporal robustness (best model vs. best baseline).

Scenario	n test	Pos. rate	Model AP	Baseline AP	Δ AP
2025	255	0.302	0.912	0.837	+0.075
2024–2025	515	0.365	0.938	0.841	+0.097

4.4. Future burden regression

The future burden regression is a complementary analysis. On the 2025 test, the model follows the main wave (rise and fall), although it partially anticipates the peak, with a WAPE of 0.285 against 0.444 for the baseline; the comparison between observed and predicted burden by municipality appears in Figure 2. In the 2024–2025 robustness scenario, this gain did not hold: the model's WAPE (0.490) was close to the baseline's (0.481). We thus interpret the regression as complementary and less stable.

4.5. Ablation

Figure 3 shows the incremental effects on AP of feature ablation. Starting from history only (AP 0.884), the seasonal anomaly and the acceleration are the block with the largest positive contribution, and the socioeconomic indicators add a smaller gain that still comprises the best final model: the long history and seasonality already capture a relevant part of the temporal heterogeneity.

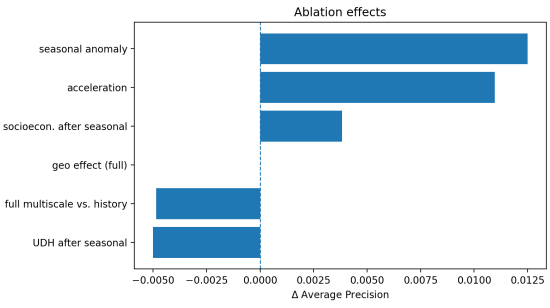


Figure 3. Ablation effects on Average Precision.

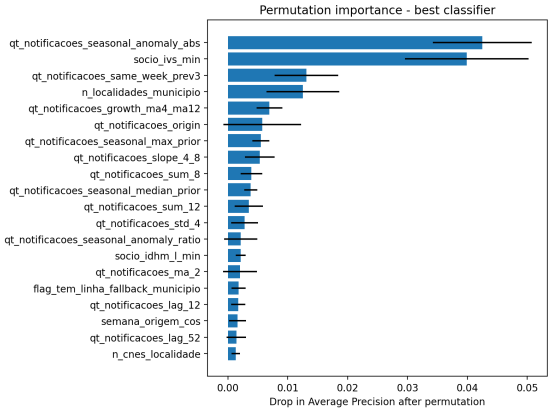


Figure 4. Permutation importance of the best classifier (2025 test set).

Adding UDH-derived features after seasonality, using the full multiscale set, or aggregating local predictions did not improve AP. This does not invalidate the use of UDH: it indicates that, with the data and period evaluated, the territorial granularity still requires refinement — more history, better population weighting, or explicit spatial modeling — to contribute to prioritization, while remaining useful for territorial investigation.

4.6. Permutation importance

The permutation importance of the best classifier (direct municipal, set `mun_history_accel_seasonal_socio_geo` with 136 variables, evaluated on the 2025 test; Figure 4) indicates the absolute seasonal anomaly, the minimum intra-municipal IVS, notifications of the same week in previous years, the number of localities, and measures of recent growth as the most relevant variables. The model, therefore, does not depend solely on the last week: it combines deviations from what is expected for the time of year, local history, and socio-spatial context.

The presence of socioeconomic variables among the most important ones should be interpreted with caution. The result suggests predictive association and utility for prioritization, but does not allow causal inference about social vulnerability and dengue transmission.

Permutation importance also quantifies only the magnitude of each variable's contribution, not the direction of its effect. Complementing it with SHAP values or partial dependence and accumulated local effects plots would show, for instance, whether a higher intra-municipal IVS raises or lowers the alert score, and would expose interactions between the seasonal anomaly and the socio-spatial context; we leave this deeper interpretability analysis as future work.

4.7. Regional heterogeneity and performance by municipality

The stratified evaluation reveals distinct epidemic regimes, and the local metrics should be read together with the rate of positive weeks (Figure 2). In Porto Alegre, 2025 had an intense and prolonged wave (positive rate 0.863); the model captured the critical period (AP 0.971, P@20 1.000), but R@20 stayed at 0.250 because the 20% budget recovers only part of many critical weeks. In Rio de Janeiro and Cuiabá, with more concentrated peaks, the model identified the risk windows well (AP 0.967 and 0.944), although it overestimated the magnitude. In Recife, with diffuse burden, the regression was aligned in the aggregate (WAPE 0.225). However, classification was more difficult (AP 0.413). The aggregate result thus stems from the ability to rank municipality-weeks under distinct regimes.

5. Limitations

The study prioritizes suspected notifications, not confirmed cases; notifications are more timely but subject to underreporting, delays, and reclassification, and the alert is relative to municipal history rather than a universal outbreak threshold. The evaluation in five pilot municipalities does not replace a national validation or prospective external temporal validation, and municipal metrics can be unstable at a very low or very high local positive rate. The socioeconomic indicators are static and census-based; the UDH/locality features do not capture spatial adjacency or diffusion and may reflect the notifying establishment, not the residence; climate data is approximated by station/municipality or state (fallback), not by UDH; and no population weighting was applied, so the need for more explicit spatial and temporal modeling persists.

A specific limitation concerns the alert definition. The threshold τ_u is estimated once, from the 80th percentile of the future burden over the entire training period, and then held fixed during the test. Over fourteen years, however, population growth, changes

in the notification process, and shifts in the circulating serotypes move the baseline level of notifications, so a static threshold may become progressively miscalibrated: it turns conservative when the reporting level rises structurally — as in the 2025 test, whose positive rate reached 30% against the 20% imposed by construction in training — and overly sensitive in periods of low circulation, as in Manaus, which had no positive week in 2025, and therefore uninformative local metrics. Adaptive definitions — rolling-window percentiles, thresholds normalized by population or by recent trend, or recalibration at each epidemiological year — would keep the alert comparable over time and are a natural extension of the protocol. Finally, the model supports, but does not replace, epidemiological surveillance: deployment and feedback from managers are the next steps in the CRISP-DM cycle.

6. Conclusion

This paper presented a multiscale formulation, safe against temporal leakage, for prioritizing dengue alerts from suspected notifications. The best model, at the municipality-week grain and enriched with history, acceleration, seasonal anomaly, and socioeconomic indicators, outperformed a strong temporal/seasonal baseline on the 2025 test. Under an operational constraint, by selecting only the 20% of the highest-risk municipality-weeks, it maintains high precision in the prioritized list, close to the recall ceiling of this budget.

The ablation shows gains from seasonal anomaly and acceleration, with a smaller contribution from the socioeconomic context, while the UDH/locality granularity should be incorporated sparingly, without discarding its value for territorial investigation. The robustness in 2024–2025 confirmed that the classification gain does not depend only on the 2025 hold-out, although the burden-regression task proved less stable. As future work, we intend to expand the evaluation to more municipalities, adopt a systematic hyperparameter optimization (grid or Bayesian search) under rolling-origin temporal validation, test adaptive alert thresholds and prospective validation, and integrate the model into an automated surveillance routine.[1]

Acknowledgments

This study was partially funded by CAPES - Finance Code 001, FAPERGS [Proj. No. 22/2551-0000390-7 and Proj. No. 25/2551-0002826-4] and CNPq [Proj. No. 310655/2026-9 and Proj. No. 308642/2022-8].

References

Aleixo, R., Kon, F., Rocha, R., et al. (2022). Predicting dengue outbreaks with explainable machine learning. In *Workshop on Artificial Intelligence for Health (AI4Health)*.

Barcellos, C. and Lowe, R. (2014). Expansion of the dengue transmission area in Brazil: The role of climate and cities. *Trop. Med. Int. Health*, 19(2):159–168.

Barnabé, S., Rosa, M. B. d., Freitas, C. M. D. S., et al. (2025). Previsão de surtos de dengue em nível de unidade de desenvolvimento humano (UDH) no Brasil: Uma abordagem integrada e explicável. In *Anais da I Escola Regional de Aprendizado de Máquina e Inteligência Artificial da Região Sul*.

[1]The curated dataset and the preparation pipeline are available at `https://github.com/mlab-inf-ufrgs/eniac-dengue-forecasting-idh`.

Borges, M. E., Codeço, C. T., Machado, D., and Almeida, A. (2026). Epidemiological and digital syndromic surveillance data on dengue, chikungunya, and SARI in Brazil. *Scientific Data*, 13:68.

Cardim, L. L., Teixeira, M. d. G., Costa, M. d. C. N., Mitka, W. M., Teixeira, C. S. S., Santos, A. C., Santos, G., Souza, A. P. F. d., Smeeth, L., Barreto, M. L., Brickley, E. B., Paixao, E. S., and Pescarini, J. M. (2025). Socioeconomic markers of dengue mortality in the 100 Million Brazilian Cohort (2007–2018): A nationwide registry-based cohort study. *PLOS Negl. Trop. Dis.*, 19(11):1–16.

Chen, X. and Moraga, P. (2025). Forecasting dengue across brazil with lstm neural networks and shap-driven lagged climate and spatial effects. *BMC Public Health*, 25(1).

Chen, X. and Moraga, P. (2026). Dengue forecasting and outbreak detection in Brazil using LSTM: Integrating human mobility and climate factors. *Infect. Dis. Model.*, 11(1):338–354.

da Silva Neto, S. R., de Oliveira, T. T., Teixeira, I. V., et al. (2022). Arboviral disease record data – dengue and chikungunya, Brazil, 2013–2020. *Scientific Data*, 9:198.

de Oliveira, F. S., Santos, D. E., Bruhn, F. R. P., et al. (2025). Forecasting dengue outbreaks with machine learning: A case study in Porto Alegre. In *Anais da I Escola Regional de Aprendizado de Máquina e Inteligência Artificial da Região Sul*. SBC.

Fernandes, M. S., Cordeiro, W., and Recamonde-Mendoza, M. (2021). Detecting aedes aegypti mosquitoes through audio classification with convolutional neural networks. *Comput. Biol. Med.*, 129:104152.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232.

Islam, J., Frentiu, F. D., Devine, G. J., et al. (2025). A state-of-the-science review of long-term predictions of climate change impacts on dengue transmission risk. *Environ. Health Perspect.*, 133(5):056002.

Ke, G., Meng, Q., Finley, T., et al. (2017). LightGBM: A highly efficient gradient boosting decision tree. In *NeurIPS*, volume 30, pages 3146–3154.

Lee, S. A., Economou, T., Catão, R. d. C., et al. (2021). The impact of climate suitability, urbanisation, and connectivity on the expansion of dengue in 21st century Brazil. *PLOS Negl. Trop. Dis.*, 15(12):e0009773.

Lowe, R., Bailey, T. C., Stephenson, D. B., et al. (2011). Spatio-temporal modelling of climate-sensitive disease risk: Towards an early warning system for dengue in Brazil. *Computers & Geosciences*, 37(3):371–381.

Lowe, R., Bailey, T. C., Stephenson, D. B., et al. (2013). The development of an early warning system for climate-sensitive disease risk with a focus on dengue epidemics in Southeast Brazil. *Statistics in Medicine*, 32(5):864–883.

Lowe, R., Coelho, C. A. S., Barcellos, C., et al. (2016). Evaluating probabilistic dengue risk forecasts from a prototype early warning system for Brazil. *eLife*, 5:e11285.

Lowe, R., Lee, S. A., O'Reilly, K. M., et al. (2021). Combined effects of hydrometeorological hazards and urbanisation on dengue risk in Brazil: A spatiotemporal modelling study. *The Lancet Planetary Health*, 5(4):e209–e219.

Ong, S. Q., Isawasan, P., Ngesom, A. M. M., et al. (2023). Predicting dengue transmission rates by comparing different machine learning models with vector indices and meteorological data. *Scientific Reports*, 13(1):19129.

Palmer, D. Q., Motta, M., Moura, E., et al. (2026). Dengue hospitalizations in Brazil: Forecasting with climatic and physicians' digital search data under real-world reporting delays. *PLOS Digital Health*, 5(5):e0001206.

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830.

Rosa, M. B., Tavares, G. C., and Recamonde-Mendoza, M. (2026). Low-cost dengue triage: Predicting disease severity using machine learning without laboratory biomarkers. In *Anais do XXVI Simpósio Brasileiro de Computação Aplicada à Saúde*. SBC.

Salim, M. F., Satoto, T. B. T., Daniel, D., et al. (2024). Digital health interventions in dengue surveillance to detect and predict outbreak: a scoping review. *Open Public Health J.*, 17(1):e18749445283264.

Siraj, A. S., Oidtman, R. J., Huber, J. H., et al. (2017). Temperature modulates dengue virus epidemic growth rates through its effects on reproduction numbers and generation intervals. *PLOS Negl. Trop. Dis.*, 11(7):e0005797.

Wilder-Smith, A., Gubler, D. J., Weaver, S. C., et al. (2017). Epidemic arboviral diseases: priorities for research and public health. *Lancet Infect. Dis.*, 17(3):e101–e106.

Wiratmadja, I. I., Salamah, S. Y., and Govindaraju, R. (2018). Healthcare data mining: Predicting hospital length of stay of dengue patients. *J. Eng. Technol. Sci.*, 50(1):110–126.

Wirth, R. and Hipp, J. (2000). CRISP-DM: Towards a standard process model for data mining. In *Proc. 4th Int. Conf. Practical Application of Knowledge Discovery and Data Mining*, pages 29–39. Springer.

Zhang, W.-X., Zhao, T.-Y., Wang, C.-C., He, Y., Lu, H.-Z., Zhang, H.-T., Wang, L.-M., Zhang, M., Li, C.-X., and Deng, S.-Q. (2025). Assessing the global dengue burden: Incidence, mortality, and disability trends over three decades. *PLOS Negl. Trop. Dis.*, 19(3):1–18.