Short-Term Forecasting of Weekly Dengue Cases in Singapore Using ARIMA and Fourier-ARIMA Models

Wentao Shan [1, *]

[1] Edwardson School of Industrial Engineering, Purdue University, West Lafayette, Indiana, United States

* Corresponding Author Email: Shan50@purdue.edu

Abstract. Dengue remains a recurring public health issue in Singapore, where sudden weekly increases in reported cases can create pressure on surveillance, healthcare response, and vector-control planning. This study builds a short-term forecasting framework using official weekly dengue surveillance data from 2012-W01 to 2026-W17. The dataset was reorganized into a continuous epidemiological-week time series, and total dengue cases were defined by combining dengue fever and dengue hemorrhagic fever cases. Because the raw series contains outbreak spikes and unstable variance, the log transformation $z_t = \log(1 + y_t)$ was used before model fitting. The analysis then applied ACF/PACF inspection, differencing, ARIMA search, Fourier-ARIMA comparison, residual diagnostics, and chronological train-test validation. The first 673 weeks were used for training, and the final 75 weeks were used for one-week-ahead testing. Results indicate that ARIMA(4,1,0)+Fourier(2) offers the strongest balance between prediction accuracy, residual adequacy, and model simplicity. The model should be understood as a practical surveillance baseline rather than a complete explanation of dengue transmission.

Keywords: Dengue, time series, ARIMA, Fourier-ARIMA, public health surveillance.

1. Introduction

1.1. Research Background and Practical Problem

Dengue remains an important surveillance problem in Singapore. Official bulletins and earlier Singapore-based forecasting studies show that weekly case counts can stay moderate for long periods and then rise quickly during epidemic phases [1-7]. Such changes are operationally important because they affect healthcare preparedness, vector-control planning, public communication, and routine monitoring.

The value of forecasting is therefore practical rather than purely descriptive. A short-horizon forecast can give surveillance teams an additional signal before reported cases move further upward. Even when the forecast is imperfect, it may still help decide whether closer observation or early preparation is needed [6, 7].

Dengue forecasting is difficult because the series does not behave like a smooth seasonal curve. Some years contain large peaks, while others remain more moderate, and the timing of high-risk weeks shifts across years. Previous work also suggests that dengue activity can be affected by rainfall, temperature, vector indices, population behavior, serotype patterns, and intervention measures [8-19]. These factors are useful but are not always available in a consistent weekly format for routine modeling.

For this reason, this study uses a narrower baseline design. Instead of modeling every driver of dengue transmission, it tests whether historical weekly case counts alone can support short-term forecasting. A case-only model is easier to reproduce, easier to update, and more realistic when only official surveillance counts are immediately available.

This baseline also creates a fair point of comparison for future extensions. If a more complex model using weather, vector, or spatial information is later developed, its value can be judged against a simple model that uses only routinely reported case counts.

1.2. Related Work

Classical time series models have often been used in infectious disease surveillance because they link current incidence to previous observations. ARIMA-type models are especially useful when recent case counts contain persistence, which is common in weekly disease data [11, 17, 20].

Earlier dengue studies applied ARIMA, SARIMA, or related statistical models to forecast notifications and incidence [8, 11, 17]. More recent studies have compared these methods with machine learning and deep learning models, including random forests, neural networks, LSTM structures, CNN-RNN hybrids, and ensemble forecasting systems [9-19]. These models can improve accuracy when richer predictors are available, but they usually require more data processing, parameter tuning, and technical support.

The recent literature is also useful for identifying the boundary of this project. Systematic reviews emphasize that validation design and reporting quality are major concerns in dengue prediction research [9]. Singapore-focused studies using deep learning or meteorological variables show the value of richer models, but they also rely on additional feature sets [7, 15]. This study therefore uses a simpler case-count baseline to evaluate what can be achieved with limited but reliable surveillance data.

Weekly data also create a long seasonal cycle of about 52 observations per year. Instead of using a highly rigid seasonal ARIMA form, Fourier terms can describe smoother annual movement in a more compact way [20]. This is appropriate for dengue because annual patterns may exist, but outbreak timing is not fixed from year to year.

Together, these studies suggest that more complex models can be useful, but they also show why a transparent benchmark remains necessary. Without a simple baseline, it is difficult to judge whether a high-dimensional model provides enough improvement to justify its added complexity.

1.3. Research Question

The central question is whether official weekly dengue case counts can provide useful short-term forecasts of high-risk weeks in Singapore. The goal is not to explain all biological or environmental causes of dengue transmission, but to test whether recent case movement carries enough information for near-term surveillance.

More specifically, this study compares a standard ARIMA baseline with Fourier-ARIMA models. The ARIMA component is used to capture short-term dependence, while the Fourier component is used to represent possible annual-scale movement. Models are compared using information criteria, forecast errors, and residual diagnostics [20].

1.4. Research Objective and Application Value

The objective is to build a transparent forecasting workflow for weekly dengue cases in Singapore. The procedure begins with official surveillance records, converts them into a continuous time series, applies transformation and differencing, and then compares ARIMA and Fourier-ARIMA candidates.

The intended application is routine monitoring. The model is not a replacement for epidemiological judgment. Instead, it provides a reproducible statistical signal that can help nontechnical users interpret recent case trends and support early warning discussions [6, 13].

2. Data and Methodology

2.1. Data Source

The dataset contains official weekly dengue surveillance records for Singapore. The unit of analysis is the epidemiological week. The cleaned dataset covers 748 weekly observations from 2012-W01 to 2026-W17. Since 2026 is incomplete, it is mainly used for the short-term forecast window rather than for a full annual comparison [1-5].

Records from 2012 to 2022 come from the data.gov.sg datastore [1]. Records for 2023 are extracted from CDA weekly bulletin PDFs, while records from 2024 to 2026 are taken from CDA weekly bulletin Excel files [2-5]. Because the source files used different formats, all records were standardized into the same weekly structure before modeling. The composition of the data sources is summarized in Table 1.

Table 1. Data sources used in the weekly dengue dataset

Period	Source	Notes
2012-2022	data.gov.sg datastore	574 weekly rows
2023	CDA weekly bulletin PDFs	52 weekly rows
2024	CDA weekly bulletin Excel files	52 weekly rows
2025	CDA weekly bulletin Excel files	53 weekly rows
2026	CDA weekly bulletin Excel files	17 weekly rows to W17

The main fields include year, epidemiological week, week start date, week end date, dengue fever cases, dengue hemorrhagic fever cases, total dengue cases, and source label. Total weekly dengue cases are used as the target variable because this study focuses on overall dengue burden rather than separate clinical categories.

Keeping the data at the weekly level also matches the public reporting structure. This avoids artificial daily interpolation and keeps the modeling unit consistent with the way the surveillance information is released and used.

2.2. Data Cleaning

The raw records required several cleaning steps before time series modeling. Dengue-related fields were retained, and dengue fever and dengue hemorrhagic fever counts were combined into one total weekly count. A single target variable helps keep the model focused on the overall weekly burden of dengue.

The time index was then standardized. Each row was linked to an epidemiological week and sorted in chronological order. Duplicate year-week pairs, missing weeks, non-numeric entries, and negative case counts were checked. Total cases were also compared with the sum of the two dengue categories to confirm that the target variable had been computed consistently.

After cleaning, the data formed a continuous weekly series suitable for transformation, differencing, model fitting, and forecast validation.

These checks are important because small data structure problems can create large modeling problems. For example, one missing week can change the meaning of lagged values, and duplicate weeks can distort the estimated autocorrelation pattern.

2.3. Dataset Partition

The dataset was split by time rather than by random sampling. This is necessary because a forecasting model should be trained on past observations and evaluated on later unseen observations [20].

The first 673 weeks were used as the training set, and the final 75 weeks were used as the test set for one-week-ahead validation. This design is close to a 90/10 split. The training set was used to estimate model parameters and compare candidate structures, while the test set measured future prediction performance.

2.4. Data Transformation

The raw series contains several sharp outbreak peaks. Most weeks have moderate counts, but epidemic weeks can be much higher. This produces right skewness and unstable variance, so direct modeling on raw case counts may give too much influence to extreme weeks.

Raw, square-root, and log-transformed versions were considered. The square-root transformation reduced part of the skewness, but the log scale compressed extreme peaks more effectively. The final modeling target is shown in Equation (1).

$$z_t = \log(1 + y_t) \tag{1}$$

where y_t is the total dengue case count in week t, and z_t is the transformed value. The value 1 is added because some weekly case counts may be zero, and log(0) is undefined.

The transformation therefore should be understood as a preparation step rather than a complete solution. It improves the scale of the data, but it does not remove every temporal pattern or guarantee that the resulting series is ready for ARIMA modeling.

2.5. Stationarity Test: ACF, PACF, and Differencing

After transformation, the series was examined for stationarity. A stationary series should have a relatively stable mean and dependence structure over time. ARIMA models are usually more appropriate after a nonstationary series has been transformed or differenced [20].

ACF was used to examine correlation across lags, while PACF helped identify direct lag relationships. These plots guided model selection, but they were not used alone. Differencing, forecast validation, and residual diagnostics were also considered.

The first difference is written as:

$$\Delta z_t = z_t - z_{t-1} \tag{2}$$

where Δz_t represents the week-to-week change in the transformed series.

Because the data are weekly, a 52-week seasonal difference was also considered:

$$\Delta_{52} z_t = z_t - z_{t-52} \tag{3}$$

where 52 represents the approximate number of weeks in one year.

The combined differencing form is:

$$\Delta\Delta_{52} z_t = (1 - B)(1 - B^{52}) z_t \tag{4}$$

where B is the backshift operator. After differencing, the series became more stable, and the remaining autocorrelation was reduced. This supported the use of ARIMA-type models.

2.6. Model Construction

Standard ARIMA models were first tested. An ARIMA model is written as:

$$\text{ARIMA}(p, d, q) \tag{5}$$

where p is the autoregressive order, d is the differencing order, and q is the moving average order. The autoregressive part uses previous values, the differencing part handles non-stationarity, and the moving average part uses previous forecast errors [20].

The ARIMA search started from ACF/PACF patterns but did not stop there. Several candidate models were fitted and compared. The best standard ARIMA baseline was ARIMA(4,1,0).

After that, Fourier terms were added to ARIMA(4,1,0). Fourier terms were used because dengue may have weak annual movement, but the peak timing is not fixed enough for a rigid seasonal model. Dynamic harmonic regression and Fourier terms are often used for long seasonal periods such as weekly data with an annual cycle [20].

For weekly data with a 52-week period, the Fourier terms are:

$$\sin(2\pi kt / 52), \quad \cos(2\pi kt / 52) \tag{6}$$

where k is the Fourier order. Larger values of K allow more complex seasonal shapes but also increase model complexity. In this study, K values from 1 to 5 were compared.

2.7. Model Evaluation Metrics

Model performance was evaluated using AICc, BIC, MAE, RMSE, and residual diagnostics. AICc and BIC measure model fit while penalizing complexity. Lower values generally indicate a better balance between fit and simplicity [20].

MAE is defined as:

$$MAE = (\frac{1}{n}) \sum_{t=1}^{n} |y_t - \hat{y}_t| \quad (7)$$

where y_t is the observed value and $\hat{y}_t$ is the forecast value. MAE gives the average absolute forecast error.

RMSE is defined as:

$$RMSE = \sqrt{\left[(\frac{1}{n}) \sum_{t=1}^{n} (y_t - \hat{y}_t)^2\right]} \quad (8)$$

RMSE gives more weight to large errors. This is important for dengue forecasting because large misses during outbreak periods may have more serious public health consequences.

Residual diagnostics were also checked. The Ljung-Box test was used to test whether residual autocorrelation remained after model fitting. A high p-value suggests that the residuals are close to random noise.

3. Results and Discussion

3.1. Overall Pattern of Weekly Dengue Cases

The cleaned weekly dengue series does not show a perfectly repeated annual pattern. Some years have moderate activity, while others contain large epidemic peaks. The high activity around 2020 is one example of how sharply the series can rise. These peaks are epidemiologically meaningful, but they also make raw-scale modeling unstable.

The week-by-year pattern also suggests that high-risk periods are not locked to the same weeks every year. This supports the use of a flexible seasonal representation. After applying Equation (1), extreme peaks became less dominant, while the timing of outbreak movement remained visible. The transformed series was therefore more suitable for ARIMA-based modeling.

3.2. Stationarity and Model Identification Results

The log-transformed series still displayed short-term persistence. This is expected in infectious disease data because high case counts in one week often carry into later weeks [11, 17, 20]. However, the log-level series was still not stationary enough for direct ARIMA modeling.

After differencing, the series became more stable, and the ACF/PACF patterns were easier to interpret. The final model was not chosen from plots alone. Candidate models were compared using forecast validation, information criteria, and residual diagnostics, which makes the procedure more reproducible.

3.3. Model Comparison Results

The analysis first compared standard ARIMA candidates and then tested Fourier-ARIMA extensions. Fourier terms were added to ARIMA(4,1,0) with K values from 1 to 5. The quantitative comparison of the leading candidate models is summarized in Table 2.The quantitative comparison of the leading candidate models is summarized in Table 2.

Table 2. Top model comparison results

Rank	Model	Family	AICc	BIC	Ljung-Box p	MAE	RMSE
1	ARIMA(4,1,0)+Fourier(2)	Fourier-ARIMA	-345.711	-304.458	0.99993	37.866	65.743
2	ARIMA(4,1,0)	ARIMA	-330.746	-307.773	0.99995	38.475	69.671
3	ARIMA(4,1,0)+Fourier(1)	Fourier-ARIMA	-333.899	-301.775	0.99990	37.585	67.991
4	ARIMA(0,1,3)	ARIMA	-330.712	-312.323	0.99646	38.580	69.692
5	ARIMA(1,1,3)	ARIMA	-328.993	-306.020	0.99810	38.622	69.800

As shown in Table 2, the results show that ARIMA(4,1,0)+Fourier(2) had the best overall balance. Its RMSE was 65.743, lower than the ARIMA(4,1,0) baseline RMSE of 69.671. This decrease is meaningful in a public health forecasting context. RMSE penalizes large forecast errors more heavily, so a lower RMSE suggests fewer large misses during weeks when accurate monitoring may matter most.

Fourier(1) had a slightly lower MAE than Fourier(2), but its RMSE was higher. Because large forecast errors are more costly in outbreak monitoring, Fourier(2) is a stronger final choice. The high Ljung-Box p-value also suggests that the final model did not leave strong residual autocorrelation.

3.4. Visual Model Comparison

Fig. 1 summarizes the MAE and RMSE values of the top five candidate models. The figure is included to make the model comparison easier to read. Since lower forecast error values indicate better accuracy, the final model should appear among the lowest bars, especially for RMSE.

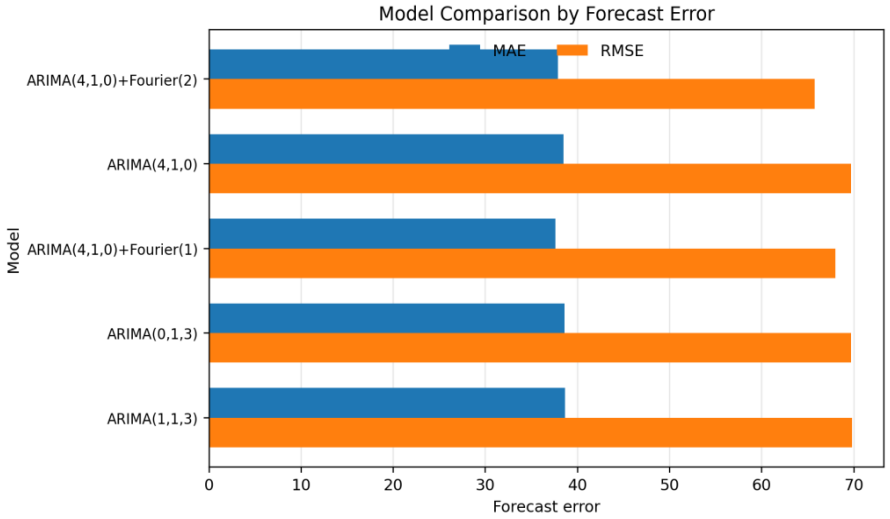


Fig. 1 Model comparison by forecast error

As shown in Fig. 1, ARIMA(4,1,0)+Fourier(2) has the lowest RMSE among the compared models. This supports its selection as the final forecasting model. The chart also shows that the improvement is not only a numerical detail. In disease surveillance, reducing larger forecast errors is useful because large misses can affect risk interpretation and response preparation.

3.5. Final Forecasting Performance

The final model was evaluated on the final 75 weeks using one-week-ahead forecasts. In this test period, it achieved an MAE of about 11.6 cases and an RMSE of about 16.0 cases. The forecast errors were centered close to zero, suggesting that the model did not consistently overestimate or underestimate dengue cases.

The 8-week future forecast also stayed around the low-to-mid 20s. The highest local forecast was 26.5 cases in 2026-W25, slightly below the outbreak threshold of about 27.0 cases under the Mean + 1.5SD rule. Therefore, the model did not indicate a clear outbreak signal in the next 8 weeks. This result should still be interpreted cautiously because a case-only model cannot fully capture sudden external changes.

In practice, the forecast is best used as one layer of evidence. If field surveillance, weather conditions, or vector indices suggest rising risk, public health interpretation should not rely only on the statistical forecast from historical cases.

4. Conclusion

This study developed a short-term forecasting framework for weekly dengue cases in Singapore. The raw surveillance data were cleaned into a continuous weekly series and transformed using Equation (1) to reduce the influence of extreme outbreak peaks. ACF/PACF analysis, differencing, ARIMA search, Fourier-ARIMA comparison, and residual diagnostics were then used to select the final model.

ARIMA(4,1,0)+Fourier(2) provided the best overall balance between forecast accuracy, information criteria, residual adequacy, and model simplicity. Compared with the ARIMA baseline, the Fourier(2) model reduced RMSE, which is useful because larger forecast errors can distort outbreak risk interpretation. The final one-week-ahead validation also showed reasonable short-term accuracy.

The main limitation is that the model uses only historical case counts. It does not include rainfall, temperature, mosquito density, serotype information, population movement, immunity, or intervention variables. It should therefore be treated as a transparent monitoring baseline rather than a complete outbreak-causation model.

Future work should add external predictors and compare more advanced time series, machine learning, or hybrid models. Weather data, mosquito surveillance, serotype information, and intervention records may improve outbreak prediction and longer-horizon planning. The current model is best used as a reproducible starting point for short-term dengue monitoring.

References

[1] Shi Yuan, Liu Xu, et al. Three-month real-time dengue forecast models: An early warning system for outbreak alerts and policy decision support in Singapore. Environmental Health Perspectives, 2016, 124(9): 1369-1375.

[2] Tian Na, Zheng Jin-Xin, et al. Precision prediction for dengue fever in Singapore: A machine learning approach incorporating meteorological data. Tropical Medicine and Infectious Disease, 2024, 9(4): 72.

[3] Government of Singapore. Weekly Infectious Disease Bulletin Cases. data.gov.sg, 2026.

[4] Communicable Diseases Agency Singapore. Weekly Infectious Diseases Bulletin 2023. Communicable Diseases Agency Singapore, 2023.

[5] Communicable Diseases Agency Singapore. Weekly Infectious Diseases Bulletin 2024. Communicable Diseases Agency Singapore, 2024.

[6] Communicable Diseases Agency Singapore. Weekly Infectious Diseases Bulletin 2025. Communicable Diseases Agency Singapore, 2025.

[7] Communicable Diseases Agency Singapore. Weekly Infectious Diseases Bulletin 2026. Communicable Diseases Agency Singapore, 2026.

[8] Hii Yien Ling, Zhu Huaiping, et al. Forecast of dengue incidence using temperature and rainfall. PLOS Neglected Tropical Diseases, 2012, 6(11): e1908.

[9] Leung Xing Yu, Islam Rakibul M., et al. A systematic review of dengue outbreak prediction models: Current scenario and future directions. PLOS Neglected Tropical Diseases, 2023, 17(2): e0010631.

[10] Ong Song Quan, Isawasan Pradeep, et al. Predicting dengue transmission rates by comparing different machine learning models with vector indices and meteorological data. Scientific Reports, 2023, 13(1): 19129.

[11] Earnest Arul, Tan Say Beng, et al. Comparing statistical models to predict dengue fever notifications. Computational and Mathematical Methods in Medicine, 2012, 2012-758674.

[12] Zhao Ning, Charland Katelyn, et al. Machine learning and dengue forecasting: Comparing random forests and artificial neural networks for predicting dengue burden at national and sub-national scales in Colombia. PLOS Neglected Tropical Diseases, 2020, 14(9): e0008056.

[13] Colon-Gonzalez Felipe J., Soares Bastos Leonardo, et al. Probabilistic seasonal dengue forecasting in Vietnam: A modelling study using superensembles. PLOS Medicine, 2021, 18(3): e1003542.

[14] Nguyen Van H., Tuyet-Hanh Thi Tuyet, et al. Deep learning models for forecasting dengue fever based on climate data in Vietnam. PLOS Neglected Tropical Diseases, 2022, 16(6): e0010509.

[15] Zhao Xinxing, Li Kainan, et al. A deep learning based hybrid architecture for weekly dengue incidences forecasting. Chaos, Solitons & Fractals, 2023, 168: 113170.

[16] Panja Madhurima, Chakraborty Tanujit, et al. An ensemble neural network approach to forecast dengue outbreak based on climatic condition. Chaos, Solitons & Fractals, 2023, 167: 113124.

[17] Sebastianelli Alessandro, Spiller Dario, et al. A reproducible ensemble machine learning approach to forecast dengue outbreaks. Scientific Reports, 2024, 14: 3807.

[18] Chen Xuan, Moraga Paula. Assessing dengue forecasting methods: A comparative study of statistical models and machine learning techniques in Rio de Janeiro, Brazil. Tropical Medicine and Health, 2025, 53: 52.

[19] Chen Xuan, Moraga Paula. Forecasting dengue across Brazil with LSTM neural networks and SHAP-driven lagged climate and spatial effects. BMC Public Health, 2025, 25: 973.

[20] Hyndman Rob J., Athanasopoulos George. Dynamic harmonic regression. Forecasting: Principles and Practice, 2021.