A multivariate approach for correcting reporting delays in infectious disease surveillance

Author: Guilherme dos Santos
Advisors: Kelly C. M. Gonçalves (DME-UFRJ) and Leonardo S. Bastos (Fiocruz)

May 31, 2023

Abstract

Frequently, real-time tracking of epidemics is faced with a concerning issue, the reporting delays of cases and deaths. Delays might occur due to logistical problems, laboratory confirmation, and other reasons. Being able to correct the delay is essential to decision-making with the goal of containing an epidemic. In some cases, the epidemic might be associated with more than one disease, Dengue and Chikungunya are common examples of this phenomenon. We propose a multivariate model to correct reporting delays and accommodate the above-mentioned cases. The model is estimated using the Integrated Nested Laplace Approximation method with the aim of providing faster results.

Keywords: Nowcasting. Dengue. Chikungunya. INLA. Bayesian hierarchical model.

1 Introduction

Disease surveillance systems play a pivotal role in the assessment of the severity of an epidemic and risk management. Timeliness is one of the main features desirable of a good surveillance system. It is essential that the time between the occurrence and the report of an event, such as the beginning of symptoms or death by a disease, is as short as possible. In reality, most surveillance systems deal with reporting delays that might occur due to logistical problems, laboratory confirmation, and other reasons. Therefore, correcting delays is essential to real-time tracking of infectious diseases and decision-making.

From a statistical perspective, the notification delay is a momentary censorship problem, where the observed data will eventually be available. Note that there is a distinction between observable and true data since there are cases that will never be reported (underreporting). The observable data are detected and will be eventually reported, while the true counts are the observed count plus the cases that will be never reported.

Bastos et al. (2019) propose a hierarchical modeling under a Bayesian approach that allows estimation of missing observable counts (nowcasting) and prediction for future times. Rotejanaprasert et al. (2020) present a model that takes into account a moving window for the nowcasting of dengue in Thailand. The notification delay problem happens in other areas as well. For example, there are approaches that perform nowcasting to correct the delay between the occurrence of the event of interest and the insurance claim (Renshaw and Verrall, 1998).

In some cases, the epidemic might be associated with more than one disease. Dengue and Chikungunya are common examples of this phenomenon, since these are associated with the same vector (the Aedes aegypti mosquito). Hence, it is reasonable to assume that epidemics of both diseases might occur concurrently. Figure 1 illustrates that, in fact, there were moments of simultaneously high number of cases of both diseases. The dashed line represents the data

observed without delay, i.e. reported at the same week of beginning of symptoms. Note that it hardly represents the real counts of cases.

We propose a multivariate Bayesian framework that allows the joint correction of delays for events that present some relationship, as in this case. Jointly handling corrections can lead to different results and smaller uncertainty when compared to previous methodologies applied independently for each disease.

Since the posterior distributions are not easily obtainable, it is necessary to use some method to sample from the resulting posterior distribution. Instead of using the usual Markov Chain Monte Carlo (MCMC) (Gamerman and Lopes, 2006) approach, we perform inference using the Integrated Nested Laplace Approximation method (INLA) (Rue et al., 2009). This choice was made to obtain faster results and facilitate the future inclusion of the model in a surveillance system. We aim for a flexible and computationally cheap model without losing the ability to accommodate the complex nature of a problem that involves multivariate correction of delays.

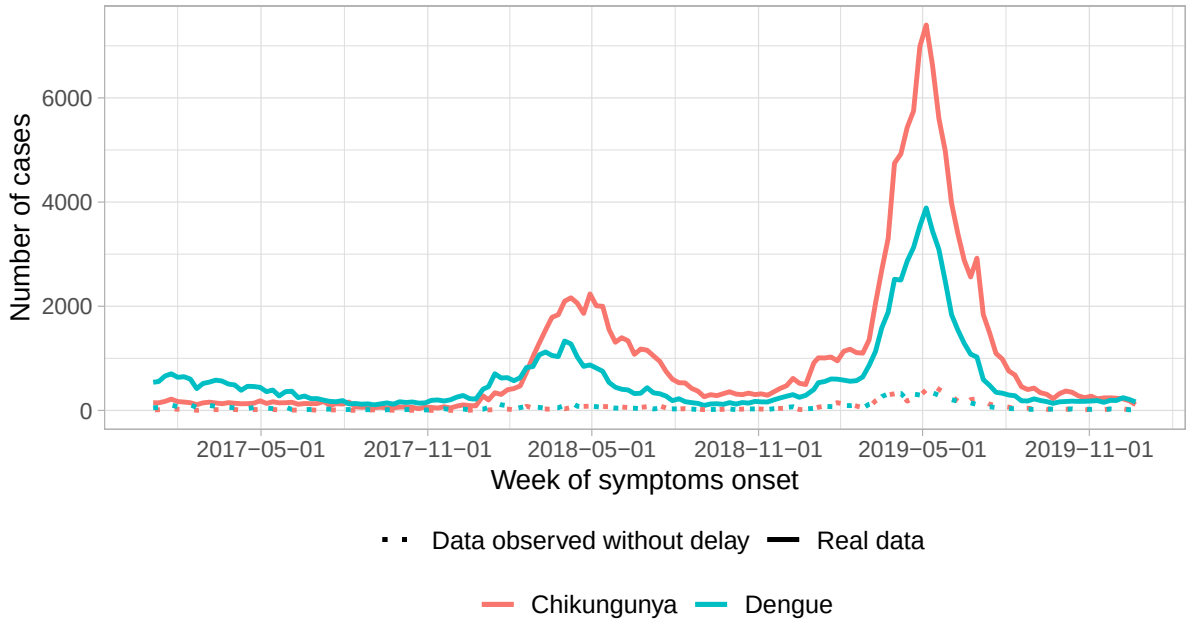

Figure 1: Number of actually reported (solid line) and without delay (dashed line) infections by Dengue (blue) and Chikungunya (red) in the state of Rio de Janeiro by week of symptom onset and disease. 2017-2019.

This document is organized as follows. In section 2, we present the background of a now-casting problem, such as the data structure and notation. The model and inference procedure based on the posterior predictive distribution of the true observed count of events are shown in Section 3. In Section 4 we present an application of the model on the data shown in Figure 1 from Dengue and Chikungunya arboviruses in Rio de Janeiro from 2017 to 2019. A brief discussion and concluding remarks, as well as the future steps of this work, are presented in Section 5.

2 Background

2.1 Run-off triangle

The usual data structure in a reporting delay problem is shown in Table 1, where the rows are indexed by the time steps $t = 1, 2, \ldots, T$ and the columns by the amount of delay $D = 0, 1, \ldots, D$. Here, T represents the current time, and D is the maximum possible delay, measured in the same unit as the time steps.

We adopt the notation used in Höhle and an der Heiden (2014) and Bastos et al. (2019). At any moment t in time, the number of events of interest (e.g., number of cases in a given week) is the total N_t, given by the sum of the row t in the table. The $n_{t,d}$ cell represents the number of events that occurred at time t but reported d units of time later. The light gray cells compose the so-called run-off triangle and are the quantities of interest in a nowcasting problem, the occurred-but-not-yet-reported events. One might also be interested in forecasting H steps ahead, where the focus is on the dark gray cells.

t \ d	0	1	2	$\cdots$	$D-2$	$D-1$	D	N
1	$n_{1,0}$	$n_{1,1}$	$n_{1,2}$	$\cdots$	$n_{1,D-2}$	$n_{1,D-1}$	$n_{1,D}$	N_1
2	$n_{2,0}$	$n_{2,1}$	$n_{2,2}$	$\cdots$	$n_{2,D-2}$	$n_{2,D-1}$	$n_{2,D}$	N_2
3	$n_{3,0}$	$n_{3,1}$	$n_{3,2}$	$\cdots$	$n_{3,D-2}$	$n_{3,D-1}$	$n_{3,D}$	N_3
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\cdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$T-D$	$n_{T-D,0}$	$n_{T-D,1}$	$n_{T-D,2}$	$\cdots$	$n_{T-D,D-2}$	$n_{T-D,D-1}$	$n_{T-D,D}$	N_{T-D}
$T-D+1$	$n_{T-D+1,0}$	$n_{T-D+1,1}$	$n_{T-D+1,2}$	$\cdots$	$n_{T-D+1,D-2}$	$n_{T-D+1,D-1}$	$n_{T-D+1,D}$	N_{T-D+1}
$T-D+2$	$n_{T-D+2,0}$	$n_{T-D+2,1}$	$n_{T-D+2,2}$	$\cdots$	$n_{T-D+2,D-2}$	$n_{T-D+2,D-1}$	$n_{T-D+2,D}$	N_{T-D+2}
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\cdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$T-2$	$n_{T-2,0}$	$n_{T-2,1}$	$n_{T-2,2}$	$\cdots$	$n_{T-2,D-2}$	$n_{T-2,D-1}$	$n_{T-2,D}$	N_{T-2}
$T-1$	$n_{T-1,0}$	$n_{T-1,1}$	$n_{T-1,2}$	$\cdots$	$n_{T-1,D-2}$	$n_{T-1,D-1}$	$n_{T-1,D}$	N_{T-1}
T	$n_{T,0}$	$n_{T,1}$	$n_{T,2}$	$\cdots$	$n_{T,D-2}$	$n_{T,D-1}$	$n_{T,D}$	N_T
$T+1$	$n_{T+1,0}$	$n_{T+1,1}$	$n_{T+1,2}$	$\cdots$	$n_{T+1,D-2}$	$n_{T+1,D-1}$	$n_{T+1,D}$	N_{T+1}
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\cdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$T+H$	$n_{T+H,0}$	$n_{T+H,1}$	$n_{T+H,2}$	$\cdots$	$n_{T+H,D-2}$	$n_{T+H,D-1}$	$n_{T+H,D}$	N_{T+H}

Table 1: Data structure in a reporting delay problem with observed number of events (white cells), occurred-but-not-yet-reported number of events (light gray cells) and future number of events we may be interested to forecast (dark gray cells).

2.2 Modelling approaches

A possible approach is to consider the distribution of counts conditioned to the total N_t. Thus, it is reasonable to assume that N_t has a Poisson or Negative Binomial distribution and $(n_{t,0} \ldots n_{t,D}) \mid N_t$ has a multinomial distribution. This approach is used in Höhle and an der Heiden (2014).

An alternative is based on the chain-ladder technique, which directly handles triangle cell counts and is developed as a distribution-free method (Mack, 1993). Renshaw and Verrall (1998) show that the chain-ladder technique can be seen as a generalized linear model, where the mean of $n_{t,d}$ is predicted by a temporal effect and an effect for the notification delay. Bastos et al. (2019) extend the chain-ladder technique to accommodate covariates and spatial effects

that explain the counts. There have been other extensions of this technique in literature, see for instance England and Verrall (2002) and Barbosa and Struchiner (2002).

Stoner and Economou (2020) propose a multivariate approach for the correction of reporting delays, but it is not suitable for approximation with INLA. Salmon et al. (2015) propose a bayesian model for detection of disease outbreaks under reporting delays, the model is implemented in INLA and is available in R though the `surveillance` package. On a more recent note, Günther et al. (2021) use a two step process to nowcast COVID-19 cases in Bavaria. McGough et al. (2020) propose a nowcasting by bayesian smooth approach that presents better performance with varying delays over time, the approach is available through the R package `NobBS`.

3 Model specification and inference procedure

3.1 Model

Let $n_{i,t,d}$ be the $n_{t,d}$ cell of the run-off triangle for the $i-$th ($i = 1, \ldots, p$) disease count, that is, the number of events occurred at time step t but reported d units of times later, for $t = 1, \ldots, T$ and $d = 1, \ldots, D$. We assume that $n_{i,t,d}$ follows a Negative Binomial distribution and the linear predictor combines joint and marginal effects, i.e.,

$$\begin{aligned} n_{i,t,d} &\sim NegBin(\lambda_{i,t,d}, \phi_i), \quad \lambda_{i,t,d} > 0, \quad \phi_i > 0, && (1) \\ \log(\lambda_{i,t,d}) &= \alpha_i + \beta_{i,t} + \gamma_{i,d} + \delta_t + \psi_d + \nu_{i,t} x_t, && (2) \end{aligned}$$

for $i = 1, \ldots, p$; $t = 1, \ldots, T$; and $d = 1, \ldots, D$.

In the parametrization used in Equation 1 for the Negative Binomial distribution, the expected value and variance are given, respectively, by $E(n_{i,t,d}) = \lambda_{i,t,d}$ and $V(n_{i,t,d}) = \lambda_{i,t,d}\,(1 + \lambda_{i,t,d}/\phi_i)$. Note that the Poisson distribution is obtained if we take the limit as ϕ_i tends to infinity. Hence, the Negative Binomial distribution can be seen as a more flexible extension of the Poisson distribution that can accommodate overdispersion.

In Equation 2, α_i is an intercept term, the effects $\beta_{i,t}$ and $\gamma_{i,d}$ capture the temporal evolution and the structure of the delay mechanism, respectively, for the $i-$th disease count. The terms δ_t and ψ_d capture the common structures in time and delay for both counts of interest. These effects are important to take the correlation between diseases into account. Moreover, x_t is a regressor variable and $\nu_{i,t}$ is the coefficient of this external variable varying over time and per disease. A particular case might be adopted when the coefficient ν_i is static, where x_t still varies in time.

In this specification, the effects of time and delay might evolve according to a variety of processes. In particular, a random walk of order 1 for these effects is a possible choice, that is,

$$\begin{aligned} \beta_{i,t} &\sim N(\beta_{i,t-1}, \sigma^2_\beta), && (3) \\ \gamma_{i,d} &\sim N(\gamma_{i,d-1}, \sigma^2_\gamma), && (4) \\ \delta_t &\sim N(\delta_{t-1}, \sigma^2_\delta), && (5) \\ \psi_d &\sim N(\psi_{d-1}, \sigma^2_\psi), && (6) \\ \nu_{i,t} &\sim N(\nu_{i,t-1}, \sigma^2_\nu), && (7) \end{aligned}$$

for $i = 1, \ldots, p$; $t = 2, \ldots, T$; and $d = 2, \ldots, D$.

Since the inference is done under a Bayesian approach, the model must be completed with the prior distributions. We can set, for instance, Gamma or Half-Normal priors for the $\sigma^2_\cdot$

hyperparameters of each random walk. These priors are also suitable for ϕ_i. It is worth noting that time series of infection counts might have longer temporal memory, which can be included in the model through higher-order random walks.

3.2 Inference procedure and nowcasting

We are interested in the posterior distribution for $\boldsymbol{\Theta}$, where

$$\boldsymbol{\Theta} \;=\; \left(\{\alpha_i\}, \{\beta_{i,t}\}, \{\gamma_{i,d}\}, \{\delta_t\}, \{\psi_d\}, \{\nu_{i,t}\}, \{\phi_i\}, \sigma^2_\beta, \sigma^2_\gamma, \sigma^2_\delta, \sigma^2_\psi, \sigma^2_\nu\right), \tag{8}$$

given all the observed data $\boldsymbol{n} = \{n_{t,d}, t + d < T\}$. The posterior distribution of $\boldsymbol{\Theta}$ is given by

$$p(\boldsymbol{\Theta} \mid \boldsymbol{n}) \;\propto\; p(\boldsymbol{\Theta}) \prod_{\substack{t=1 \\ \{t+d<T\}}}^{T} \prod_{d=0}^{D} p\left(n_{t,d} \mid \boldsymbol{\Theta}\right), \tag{9}$$

where $p\left(n_{t,d} \mid \boldsymbol{\Theta}\right)$ is the Negative Binomial probability function and $p(\boldsymbol{\Theta})$ is the prior distribution given by the product of the marginal priors for the effects and hyperparameters.

To avoid the computational expensiveness of MCMC methods, we perform approximate inference through INLA. In short, INLA performs a sequence of Laplace approximations and numerical methods for sparse matrices and is suitable for latent Gaussian models (which is the case for the proposed model). Hence, the model can be promptly implemented in R through R-INLA (`www.r-inla.org`). This procedure is simple and significantly faster than MCMC approaches most of the time (Rue et al., 2017).

Once we learn about the parameters through the posterior distribution, we can access the posterior predictive distribution of the yet unobserved $\{n_{t,d};\; T < t + d < T + D\}$ to perform the nowcast.The posterior predictive is given by

$$\underset{\{T<t+d<T+D\}}{p\left(n_{t,d} \mid \boldsymbol{n}\right)} \;=\; \int_{\boldsymbol{\Theta}} p\left(n_{t,d} \mid \boldsymbol{\Theta}\right) p(\boldsymbol{\Theta} \mid \boldsymbol{n}) d\boldsymbol{\Theta}. \tag{10}$$

The integral in (10) is not solved analytically, although an approximation is obtainable using a Monte Carlo approach. In practice, we take samples from the posterior $p(\boldsymbol{\Theta} \mid \mathbf{n})$ and then, for each sample $\boldsymbol{\Theta}$ we sample from the negative binomial $p(n_{t,d} \mid \boldsymbol{\Theta})$.

Once the simulated values from the posterior predictive distribution are obtained, we can calculate the total number of notifications at time t, N_t, by summing over the rows of Table 1. Some values will be known, and others will be taken from the samples of the posterior predictive.

In short, these are the steps taken to obtain the posterior predictive using the Monte Carlo approach after the approximation has been found.

1. Sample $(\{\alpha_i\}, \{\beta_{i,t}\}, \{\gamma_{i,d}\}, \{\delta_t\}, \{\psi_d\}, \{\nu_{i,t}\})$ from the joint posterior;
2. Sample $n_{t,d}$ from the likelihood evaluated at the sampled parameters;
3. Compute $N_t = \sum_{d=0}^{D} n_{t,d}$, some values are known, others have been sampled in the previous step.

4 Application to Dengue and Chikungunya in Rio de Janeiro

In this section, we present an application of the model to the Dengue and Chikungunya data in the state of Rio de Janeiro from 2017-2019, as shown in Figure 1. In Brazil, dengue and chikungunya diseases of mandatory notification. The data consists of suspected cases of Dengue and Chikungunya in the city of Rio de Janeiro reported in the Sistema de Informação de Agravos de Notificação (SINAN). The aggregated data is available ate InfoDengue (`https://info.dengue.mat.br/`).

We aim to nowcast the number of cases in a given week using the proposed model assuming $p = 2$. For that, some model structures assuming different real possibilities for a bivariate nowcasting problem were experimented. The number of tweets containing Dengue-related terms per week was used as a covariate. The model formulations considered are presented in Table 2.

Model description	Linear predictor ($log(\lambda_{i,t,d})$)
M0 - Two univariate models (model without common effects)	$\alpha_i + \beta_{i,t} + \gamma_{i,d}$
M1 - Complete model	$\alpha_i + \beta_{i,t} + \gamma_{i,d} + \delta_t + \psi_d + \nu_{i,t}x_t$
M2 - Complete model with static covariate effect	$\alpha_i + \beta_{i,t} + \gamma_{i,d} + \delta_t + \psi_d + \nu_i x_t$
M3 - Model without the covariate	$\alpha_i + \beta_{i,t} + \gamma_{i,d} + \delta_t + \psi_d$
M4 - Model considering the same delay structure for both time series	$\alpha_i + \beta_{i,t} + \delta_t + \psi_d$
M5 - Model considering only the common terms of time and delay	$\alpha_i + \delta_t + \psi_d$

Table 2: Models considered in the application and respective linear predictor structures.

All models presented in this section considered a random walk of order 2 for the time effects and order 1 for delay effects and the covariate's coefficient. We considered the default INLA priors, i.e. a vague Gamma prior, for the precision hyperparameters $1/\sigma^2$ in equations 3 to 7. We set a vague Log-Gamma prior for the logarithm of ϕ_i, the size of the Negative Binomial distribution, which is equivalent to setting a Gamma prior for this hyperparameter.

We chose to evaluate the model at the second peak shown in Figure 1 (20th epidemiological week of 2019) since there was a simultaneously high number of infections by Dengue and Chikungunya. Hence, it is a moment when correction of reporting delays would be demanding. After truncating the data on the 20th epidemiological week of 2019, we need to consider that every case reported after is yet unobserved. The maximum possible delay considered in this case was $D = 15$ weeks. After the truncation, $T = 124$ time steps remain (rows of Table 1), where the last 15 are under the influence of reporting delays.

After calculating the approximation, a sample of size 1000 was taken from the joint posterior in the inference procedure. The point estimations are given by the median of the posterior predictive distribution, and we took the 95% credible intervals considering the same probability for the tails of the posterior predictive for the interval estimation.

Table 3 presents some selection criteria used to assess the performance of the models. The table contemplates information-based criteria, like the Watanabe-Akaike information criterion (WAIC) (Watanabe, 2013) and the deviance information criterion (DIC) (Spiegelhalter et al., 2002), and measures of error for point estimations through the mean absolute percentage error (MAPE). To evaluate the quality of the interval estimation, we present the range of the last 95% credible interval (since it is the one under most uncertainty) and the mean interval score (MIS) over the nowcasting period. The interval score is a measure that penalizes points outside the interval based on its level as well as its range (Gneiting and Raftery, 2007). Other metrics as the mean absolute and mean squared error for point estimation, mean range and interval score for 95% and 50% intervals were computed, but omitted from the table since the results

were often redundant. For every criterion in the table, smaller values indicate better results.

			Chikungunya correction			Dengue correction		
Model	WAIC	DIC	MAPE	Range of last 95% CI	MIS	MAPE	Range of last 95% CI	MIS
M0	25195.23	25167.95	7.11	11161.30	54.15	8.75	4549.55	57.75
M1	25179.51	25154.80	6.89	8893.00	49.56	8.89	**3759.30**	49.15
M2	25176.96	25153.08	**6.68**	8697.00	**49.36**	8.78	3912.62	48.50
M3	**25173.92**	**25150.11**	6.70	9003.17	50.39	8.93	3879.72	51.00
M4	25378.97	25360.83	11.33	**8282.85**	90.04	6.79	4349.70	**34.01**
M5	26647.28	26635.84	14.05	8935.25	171.75	**6.49**	6142.53	34.09

Table 3: Model selection criteria for the six models used in the application.

Apart from the information-based criteria, there is no clear best model for delays' corrections of both diseases. However, note that M3 is quite close to the best model according to MAPE and MIS criterion for the Chikungunya correction. For the Dengue correction, this model still performs well, being not too far from others when we look at the MAPE and the second best model considering the range of the last interval.

M5 yields the best point estimations for the observed Dengue number of cases, although the opposite happens for Chikungunya estimations with the same model, which might indicate that the model prioritized Dengue corrections. Note that M0 (equivalent to two univariate models) performs worse than at least three of the other models for most metrics presented in the Table, which indicates that there is, in fact, a gain in choosing a multivariate model in this case.

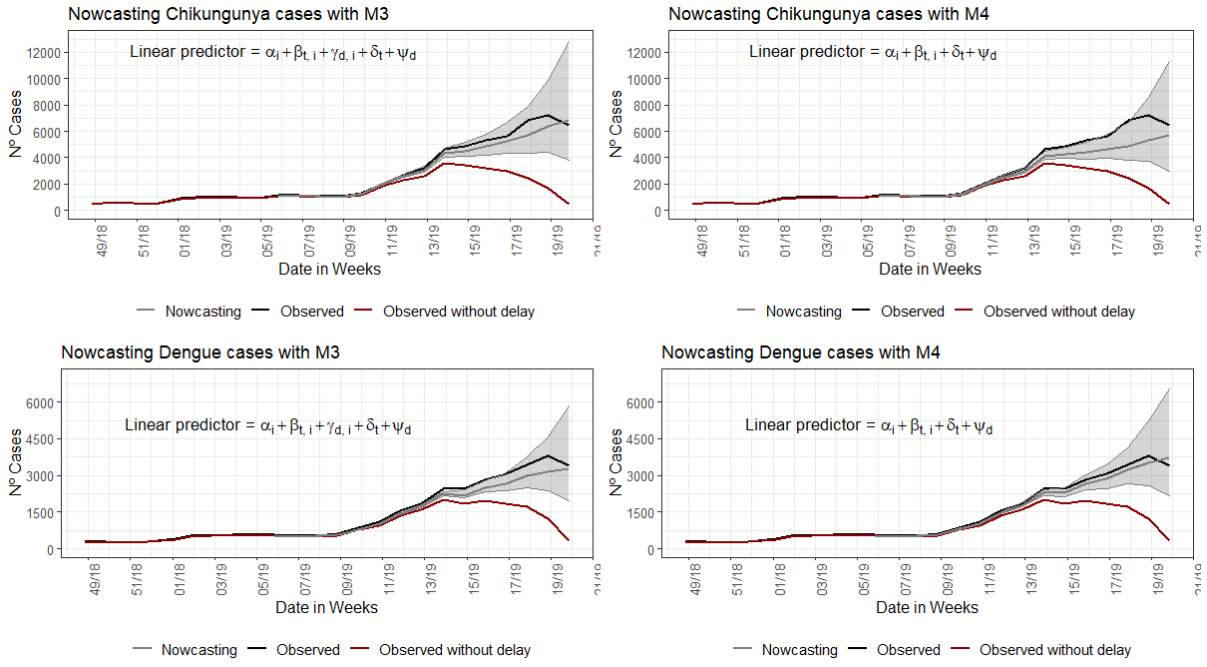

Figure 2: Nowcasting (grey), Observed data (black) and Observed data without delay (red) for M3 and M4.

The corrections for M3 and M4 are presented in Figure 2. Both models can correct the delay and retrieve the original yet-to-be-observed time series to a certain extent. M4 presents better point estimates for Dengue cases, but with larger intervals, and underestimates Chikungunya cases by a considerable amount. Therefore, we will be looking at the results of M3 from now on, since it is the model that presents jointly better results.

The posterior median effects of time ($\beta_{i,t} + \delta_t$) and delay ($\gamma_{i,d} + \psi_d$) for both time series obtained from M3 are presented in Figure 3. When compared to Figure 1, we see that the time effects undergo the same shifts as the original series of cases, with a higher number of Dengue cases at the begging of the period when compared to the Chikungunya cases. The delay effects indicate that Dengue delayed reports are usually more incident than the Chikungunya ones in the first four weeks, and then the behavior changes.

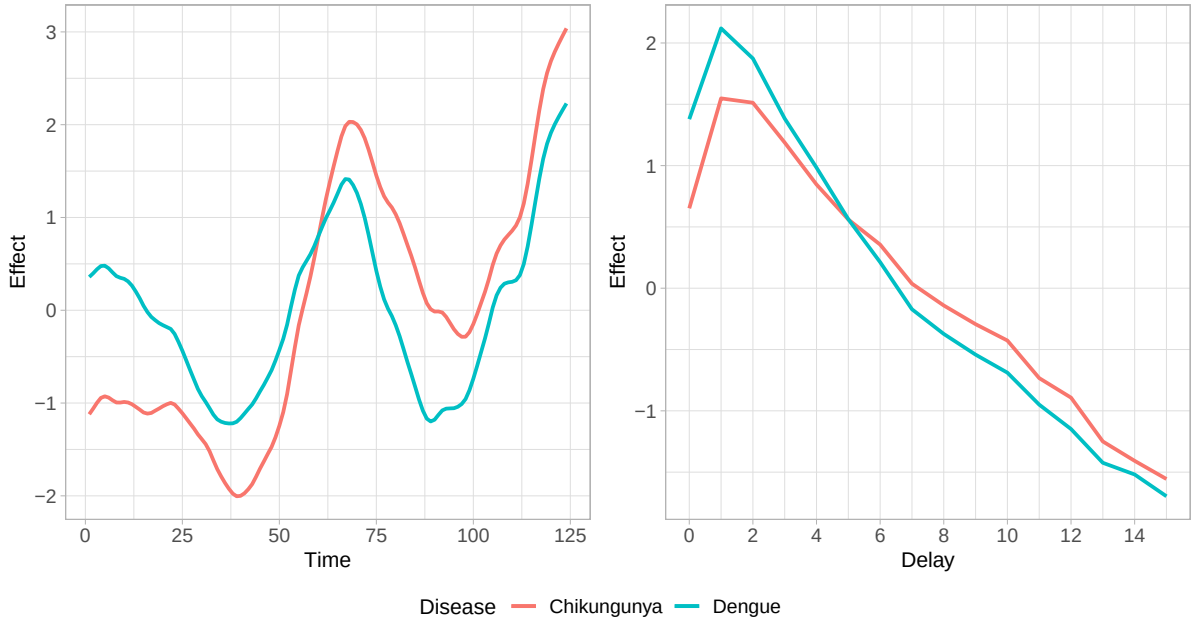

Figure 3: Effects of time and delay for Dengue (blue) and Chikungunya (red).

The posterior median of ϕ_i and precision hyperparameters are presented in Table 4. The low values for ϕ_i indicate that the Poisson distribution is likely, not fit for this application. A Poisson distribution was also considered, but the model did not perform well. The smaller precision for the correlation-inducing parameters, δ_d, and ψ_d, show that these have more importance in the linear predictor and the model is not close to a double univariate approach.

Hyperparameter	Median
ϕ_1	3.06
ϕ_2	3.90
$1/\sigma^2_\beta$	2546.44
$1/\sigma^2_\gamma$	126.86
$1/\sigma^2_\delta$	451.10
$1/\sigma^2_\psi$	10.11

Table 4: Posterior median of hyperparameters for M3.

5 Preliminary conclusions and future steps

We propose a multivariate framework for correcting reporting delays in disease surveillance data and show its application in a real dataset considering two diseases. So far, the model can accommodate different effects of time and delay for both time series and also covariates with effects varying over time. Considering an interaction term between time and delay is a future possibility.

The proposed model has shown to perform better than two univariate models in the application, providing better point estimates and smaller intervals. The estimated effects seem to capture the dynamic of the actual number of cases over time, which is desirable.

A simulation study fitting only the INLA approximation was performed (but not included). Since we can obtain samples from the posterior considerably fast with INLA, a sensitivity analysis for the prior distributions will be done in future steps.

Other covariates apart from number of tweets will be investigated in applications. The index of google searches containing terms related to Dengue or Chikungunya, obtainable from google trends, is also a possibility that has been shown useful in nowcasting problems (Miller et al., 2022). Environmental variables such as climate associated ones are also a possibility.

Other metrics to evaluate the predictions are currently being investigated, such as the continuous ranked probability score (CRPS). It would be interesting to find some metric that considers the samples from the posterior predictive distribution instead of only the quantiles that provide the intervals.

Some formulations to contemplate other nowcasting problems, e.g., splitting the total number of cases of severe acute respiratory illness (SARI) into Influenza and COVID-19, are a possibility. An application in arboviruses in other places in Brazil are also a possibility. A future intention is to develop an R package and implement the model in a surveillance system.

References

Barbosa, M. T. S. and Struchiner, C. J. (2002). The estimated magnitude of aids in brazil: a delay correction applied to cases with lost dates. *Cadernos de Saúde Pública*, 18:279–285.

Bastos, L. S., Economou, T., Gomes, M. F., Villela, D. A., Coelho, F. C., Cruz, O. G., Stoner, O., Bailey, T., and Codeço, C. T. (2019). A modelling approach for correcting reporting delays in disease surveillance data. *Statistics in medicine*, 38(22):4363–4377.

England, P. D. and Verrall, R. J. (2002). Stochastic claims reserving in general insurance. *British Actuarial Journal*, 8(3):443–518.

Gamerman, D. and Lopes, H. F. (2006). *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. CRC press.

Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378.

Günther, F., Bender, A., Katz, K., Küchenhoff, H., and Höhle, M. (2021). Nowcasting the covid-19 pandemic in bavaria. *Biometrical Journal*, 63(3):490–502.

Höhle, M. and an der Heiden, M. (2014). Bayesian nowcasting during the stec o104: H4 outbreak in germany, 2011. *Biometrics*, 70(4):993–1002.

Mack, T. (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. *ASTIN Bulletin: The Journal of the IAA*, 23(2):213–225.

McGough, S. F., Johansson, M. A., Lipsitch, M., and Menzies, N. A. (2020). Nowcasting by bayesian smoothing: A flexible, generalizable model for real-time epidemic tracking. *PLoS computational biology*, 16(4):e1007735.

Miller, S., Preis, T., Mizzi, G., Bastos, L. S., Gomes, M. F. d. C., Coelho, F. C., Codeço, C. T., and Moat, H. S. (2022). Faster indicators of chikungunya incidence using google searches. *PLOS Neglected Tropical Diseases*, 16:1–16.

Renshaw, A. E. and Verrall, R. J. (1998). A stochastic model underlying the chain-ladder technique. *British Actuarial Journal*, 4(4):903–923.

Rotejanaprasert, C., Ekapirat, N., Areechokchai, D., and Maude, R. J. (2020). Bayesian spatiotemporal modeling with sliding windows to correct reporting delays for real-time dengue surveillance in thailand. *International Journal of Health Geographics*, 19(1):1–13.

Rue, H., Martino, S., and Chopin, N. (2009). Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of the royal statistical society: Series b (statistical methodology)*, 71(2):319–392.

Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P., and Lindgren, F. K. (2017). Bayesian computing with inla: a review. *Annual Review of Statistics and Its Application*, 4:395–421.

Salmon, M., Schumacher, D., Stark, K., and Höhle, M. (2015). Bayesian outbreak detection in the presence of reporting delays. *Biometrical Journal*, 57(6):1051–1067.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the royal statistical society: Series b (statistical methodology)*, 64(4):583–639.

Stoner, O. and Economou, T. (2020). Multivariate hierarchical frameworks for modeling delayed reporting in count data. *Biometrics*, 76(3):789–798.

Watanabe, S. (2013). A widely applicable bayesian information criterion. *Journal of Machine Learning Research*, 14(27):867–897.