Filtering and improved Uncertainty Quantification in the dynamic estimation of effective reproduction numbers

Marcos A. Capistrán[*]   Antonio Capella[†]
J. Andrés Christen[*‡]

03DIC2020

Abstract

The effective reproduction number R_t measures an infectious disease's transmissibility as the number of secondary infections in one reproduction time in a population having both susceptible and non-susceptible hosts. Current approaches do not quantify the uncertainty correctly in estimating R_t, as expected by the observed variability in contagion patterns. We elaborate on the Bayesian estimation of R_t by improving on the Poisson sampling model of Cori, Ferguson, Fraser, and Cauchemez (2013). By adding an autoregressive latent process, we build a Dynamic Linear Model on the log of observed R_ts, resulting in a filtering type Bayesian inference. We use a conjugate analysis, and all calculations are explicit. Results show an improved uncertainty quantification on the estimation of R_t's, with a reliable method that could safely be used by non-experts and within other forecasting systems. We illustrate our approach with recent data from the current COVID19 epidemic in Mexico.

Keywords: Effective reproduction number; Renewal equation; Bayesian inference; Dynamic Linear Models; Uncertainty Quantification.

[*]Centro de Investigación en Matemáticas (CIMAT), Jalisco S/N, Valenciana, Guanajuato, GTO, 36023, MEXICO. *marcos at cimat.mx, jac at cimat.mx*

[†]Instituto de Matemáticas, Universidad Nacional Autónoma de México (UNAM). *capella at im.unam.mx*

[‡]Corresponding author.

1

1   Introduction

The effective reproduction number R_t is the proportion of new cases generated by active cases at calendar time t. Estimating R_t has proved to be an important tool for monitoring ongoing epidemics, to monitor interventions or general changes in contagion patterns (e.g. Ferguson et al., 2006; Fraser, Riley, Anderson, and Ferguson, 2004). This has been the case in the past and even more recently with the current COVID19 epidemics. In Germany, a calculation of R_t is used as public monitor of the evolution of the COVID19 epidemic and has direct consequences on public life and decision making RKI (2020). Regarding the estimation of R_t, in basically all realistic cases, a proxy for incidence data has to be used, and along with the inherent stochastic variability in reported cases, proxies are subject to noise.

There are methods to estimate R_t as a sub product of larger mechanistic compartmental models. The main limitation of these methods is precisely that an underlying mechanistic model must be postulated and calibrated beforehand, thus inheriting possible model drawbacks (see Bettencourt and Ribeiro, 2008; Cintrón-Arias et al., 2009, for example). Recently, more generic methods for estimating R_t, requiring incidence data only, have been proposed Cauchemez, Boëlle, Thomas, and Valleron (2006); Cori et al. (2013); Fraser (2007); Thompson et al. (2019); Wallinga and Teunis (2004). Cauchemez et al. (2006) propose a hierarchical model and Bayesian estimation using Markov Chain Monte Carlo (MCMC) methods to estimate R_t. MCMC is hardly recommended for non-experts and simpler analytic methods are sough.

Let I_t be the incidence of cases, to be used for the estimation of R_t, $t = 1, 2, \ldots$. Using the renewal equation and other assumptions Fraser (2007) explains that the effective reproduction number R_t may be defined and estimated with

$$R_t = \frac{I_t}{\sum_{s=1}^{t} I_{t-s} w_s} = \frac{I_t}{\Lambda_t}. \tag{1}$$

Here w_s is a probability mass function (i.e. $w_s \geq 0$ and $\sum_{s=1}^{\infty} w_s = 1$) that accounts for the infectious disease generation interval. It is assumed that infected individuals have a generation interval w_s, dependent on the number on days since infection s but independent of current time. At time t, $I_{t-s} w_s$ represents a "force of infection" of individuals infected s days ago. Then Λ_t represents the "total infectiousness of infected individuals" Cori et al. (2013). It may also be seen as an estimation of the current total of active cases. R_t is

the ratio of secondary cases I_t produced by the current total of active cases. From here onwards we will assume that time t is measured in days.

Note that I_t may be only a proxy of the total number of infected people, as for example, people arriving at hospitals to seek help, etc. In most common cases, the new infected individuals at time t, I_t^*, including asymptomatic infections, is impossible to measure. Ideally, as a proxy, we may regard instead $I_t = KI_t^*$ with K an unknown constant. However, R_t is a unit less measure of the proportion of secondary infections and is indeed independent of K. Any estimation of R_t should reflect this, providing the same results if a different K is used or if the proxy is given proportional to the total population etc. Moreover, the uncertainty we have on the actual R_t should not be bound to the absolute values of I_t but to the recent variation in the observed ratio I_t/Λ_t.

Moreover, epidemic data has far more caveats that need to be address. Reporting delays , along with right truncation are very common in epidemic data. Only under a thoughg investigation of the reporting system at hand, is that incidence data may be corrected to obtain a reasonable enough good proxy for estimating R_t (McGough, Johansson, Lipsitch, and Menzies, 2020; Salmon, Schumacher, and Höhle, 2016). As we say, only ideally is that $I_t = KI_t^*$. Even after standard data pre processing techniques are applied, conditions may change during an epidemic with the net result that K changes over time, for example. Undoubtedly, the use of raw incidence data, without proper pre processing, most likely will lead to incorrect/biased estimates of R_t (Gostic et al., 2020). We do not discuss data preprocessing any further and assume $I_t = KI_t^*$ throughout, with an unknown fixed K.

In Cori et al. (2013) the estimation of R_t, based on the proxy incidence I_t, uses (1) to form the statement that

$$E(I_t|I_{t-1}, I_{t-2}, \ldots, I_1) = R_t\Lambda_t.$$

A probability model is then proposed for the observed I_t conditional on $I_{t-1}, I_{t-2}, \ldots, I_1$ with the expected value $R_t\Lambda_t$. Cori et al. (2013) proposed the Poisson model $I_t|I_{t-1}, I_{t-2}, \ldots, I_1 \sim Po(R_t\Lambda_t)$. Incidence data is commonly overdisperse and the immediate improvement on this Poisson model would be to use a more general count data model as a Negative Binomial (Lindén and Mäntyniemi, 2011). Dispersion parameters may be fixed or estimated in the Negative Binomial (Flaxman et al., 2020). However, directly postulating a model for incidence data may result in estimates that depend
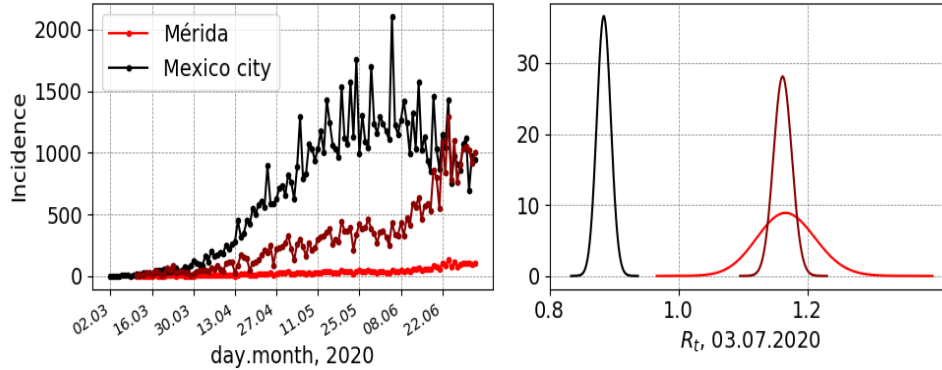
Figure 1: Incidence data for the ongoing COVID19 epidemic in two metropolitan areas of Mexico (left, confirmed cases, see Capistran et al., 2020). The corresponding posterior of R_t according to Cori et al. (2013) Poisson model (right). Also, the Mérida incidence data multiplied by 10, and its corresponding posterior (dark red). The uncertainty presented by the posterior responds to the incidence absolute values and not to the variation in reproduction numbers. The generation interval w_s used here is explained in Figure 2.

on the unknown constant K. Moreover, as far as estimating R_t is concern, the actual variability of I_t is not in principle relevant, only the variability observed in the ratios I_t/Λ_t.

The resulting inference of Cori et al. (2013), where they assume a constant R_t^τ of $\tau = 7$ days (with a gamma $Ga(5,1)$ conjugate prior) is illustrated in Figure 1. In particular, the coefficient of variation (CV) of the posterior distribution for R_t^τ is

$$\frac{1}{\sqrt{a + \sum_{t-\tau+1}^{t} I_t}}$$

(see Cori et al., 2013, p.3 web material). The CV will decrease if we increase the arbitrary constant K. It does depend only on the absolute current values for I_t and not on the intrinsic variation of the underlying R_t's, see Figure 1. Further elaboration on this Poisson model may be found in recent literature. Thompson et al. (2019) improves this Bayesian inference of R_t by also estimating w_s and adding the corresponding variability in the posterior of R_t.

Nouvellet et al. (2018) uses Cori et al. (2013) embedded in a more complex forecast system. The need for a statistical estimation of R_t is apparent, with a reasonable estimation of the uncertainty in all inferences.

Recently, and prompted by the current uses of R_t estimates during the COVID19 epidemic, Gostic et al. (2020) make a review of current methods and recommend the use of Cori et al. (2013), in contrast to other methods not based on (1), see Gostic et al. (2020) and references therein. Gostic et al. (2020) stress some other three points in the estimation of R_t, namely: 1) Data for estimating the infectivity function w_s is rarely available, and proxies for estimating w_s should be used with care. 2) Correct incidence data should always be used, representing the number of new infections every day, taking special care to correct for reporting delays and other issues. 3) Smoothing of R_t estimation is a matter of concern. Cori et al. (2013) uses a smoothed window of $\tau = 7$ days, but this should be used judiciously depending on the data at hand, although no definite guidelines are suggested in Gostic et al. (2020). However, Gostic et al. (2020) focus on other comparing issues in the estimation of R_t and did not concentrate on the adequateness of the uncertainty expressed by the posterior distributions of Cori et al. (2013).

Our aim here is to improve the Uncertainty Quantification (UQ) in the estimation of R_t, from the definition in (1), and using an analytic Bayesian approach (i.e. non-MCMC) that may be robustly embedded in other systems or used by non-experts.

We follow the same basic idea of Cori et al. (2013) but interpret (1) differently. Namely, define $\rho_t = log(R_t)$, then

$$E(i_t|i_{t-1}, i_{t-2}, \ldots, i_1, \rho_t, \rho_{t-1}, \rho_{t-2}, \ldots, \rho_1) = \rho_t + \lambda_t,$$

where $i_k = log(I_k) + log(K)$ and $\lambda_t = log(\Lambda_t) + log(K)$, for some unknown constant K. As usual, when taking logarithms of count data to postulate an observation model we assume that $i_t|\cdots \sim N(\rho_t + \lambda_t, \sigma)$, for some unknown variance σ, or equivalently

$$y_t = i_t - \lambda_t \sim N(\rho_t, \sigma_t^2).$$

As recommended by Cori et al. (2013), R_t should only be estimated for incidence with $I_t \geq 10$. This further justifies the use of the above Normal model. Therefore, with the observed log R_t's, i.e. the y_ts, we will try to estimate the $\rho_t = log(R_t)$s. Cori et al. (2013) assumed a constant R_t^τ over τ days, obtaining some smoothing. A further improvement on the approach

Table 1:   Updating formulas for the parameters of the non-central student-t posterior distribution for ρ_t, $\rho_t|D_t \sim T_{n_t}(m_t, c_t)$.

$$
\begin{array}{lll}
n_t = \delta n_{t-1} + 1 & a_t = g m_{t-1} & r_t^* = c_{t-1} + w^* \\
q_t^* = (r_t^* + 1) & q_t = s_{t-1} q_t^* & r_t = s_{t-1} r_t^* \\
e_t = (y_t - a_t) & s_t = \delta \frac{n_{t-1}}{n_t} s_{t-1} + \frac{s_{t-1}}{n_t} \frac{e_t^2}{q_t} & A_t = \frac{r_t^*}{q_t^*} \\[2mm]
m_t = a_t + A_t e_t & c_t = \frac{s_t}{s_{t-1}} \{ r_t - A_t^2 q_t \} &
\end{array}
$$

of Cori et al. (2013) is that we explicitly and formally model smoothing, or coherence, among the ρ_ts. This we do by postulating an autoregressive prior for the ρ_t's, modeling the fact that ρ_t should be *similar* (but not necessarily equal) to ρ_{t-1}. This creates a Dynamic Linear Model, and using Bayesian updating, forms a filtering type inference for the ρ_t sequence.

We will continue to assume the generation interval w_s fixed and dedicate our effort to improving the UQ in the Bayesian estimation of R_t. As mentioned above, we do not discuss data preprocessing and assume that I_t is a correct proxy for estimating R_t. We do not provide an analysis on the effects of incidence data anomalies on our estimation of R_t. The paper is organized as follows. We present the details of our approach in the next section and in section 3 we present two examples using recent COVID19 incidence data from Mexico. Finally, in Section 4 we make a discussion of our results.

2   Materials and methods

A Dynamic Linear Model (West and Harrison, 1997) is proposed for the estimation of the ρ_t's. Namely,

$$
\begin{aligned}
y_t &= \rho_t + \nu_t, & \nu_t &\sim N(0, \phi_t^{-1}). \\
\rho_t &= \rho_{t-1} + \omega_t, & \omega_t &\sim N(0, \phi_t^{-1} w^*). \\
\phi_t &= \gamma_t \phi_{t-1}/\delta, & \gamma_t &\sim Beta(\delta n_{t-1}/2, (1-\delta)n_{t-1}/2).
\end{aligned}
$$

With $\rho_0 \sim N(m_0, \phi_t^{-1} c_0^*)$ an autoregressive AR(1) (prior) model is proposed to estimate the log R_t's. Here we state that ρ_t is equal to ρ_{t-1} plus some noise ω_t. The precision ϕ_t process, i.e. $\sigma_t^2 = \phi_t^{-1}$, is the variance of the observed

log R_t's, y_t, and is assumed unknown. With $\gamma_1 \sim Beta(\delta n_0/2, (1-\delta)n_0/2)$ a second AR(1) is assumed to estimate these variances. n_t are the discounted degrees of freedom $n_t = \delta n_{t-1} + 1$, (see West and Harrison, 1997, sec. 10.8 for details). Here we also assume that the precision ϕ_t is similar to ϕ_{t-1} multiplied by the random innovation or shock γ_t/δ. Since $E(\gamma_t) = \delta$, γ_t/δ has a shifted *Beta* distribution with mean 1. This keeps all ϕ_t positive. The *discount factor* δ and the hyperparameters m_0, c_0^*, w^*, n_0 are part of the modeling and prior definition (see below) and considered known.

The marginal posterior of ρ_t may be calculated analytically and corresponds to a non-central student-t distribution with n_t degrees of freedom

$$\rho_t|D_t \sim T_{n_t}(m_t, c_t) \tag{2}$$

where $D_t = (y_t, D_{t-1})$; see Table 1 for the recursive calculation of the central parameter m_t and dispersion parameter c_t. This constitutes a filtering type estimation (Kalman, 1960) of the ρ_t's, with discounted variance (West and Harrison, 1997, sec. 10.8). The marginal posterior distribution of R_t may be obtained from (2), namely

$$f_{R_t}(r|D_t) = r^{-1} f_{\rho_t}(log(r)|D_t).$$

In particular, quantiles for this posterior may be calculated by transforming the corresponding quantiles of (2) since $R_t = exp(\rho_t)$ is monotonic.

2.1 Autoregressive modeling and prior distributions

The assumption of Cori et al. (2013) that R_t remains constant over τ days is relaxed here with the AR(1) prior model for the ρ_t's, $\rho_t = \rho_{t-1} + \omega_t, \omega_t \sim N(0, \phi_t^{-1}w^*)$. The variance discount factor δ is key for learning from recent variability in the ρ_t's and gradually cease to learn from more distant ρ_t's. It is proved that the degrees of freedom n_t, which we may visualize as the "effective" sample size, the number of sample points in the past actually used to estimate the variance component, has the property $n_t \to (1-\delta)^{-1}$ (West and Harrison, 1997, sec. 10.8). We postulate that $(1-\delta)^{-1} = 2\tau$, or $\delta = 1 - (2\tau)^{-1}$. That is, the limit to learning from the past is 2τ days. Regarding the prior for $\rho_0 \sim N(m_0, \phi_0^{-1}c_0^*)$, we center it at $m_0 = 0$ (i.e. $R_t = 1$) with the same variance for the observed y_t, that is $c_0^* = 1$. Regarding the prior for $\gamma_1 \sim Beta(\delta n_0/2, (1-\delta)n_0/2)$, note that by design it is centered
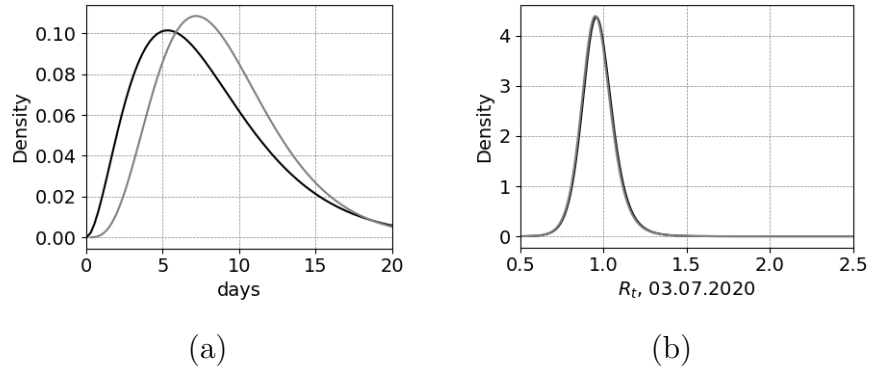
Figure 2: (a) Our postulated Erlang(3, 8/3) (black) and slightly shifted Erlang(5, 9/5) (grey), as the generation interval of COVID19. (b) The resulting posterior distributions of R_t on 03JUL2020 for our method, for the Mérida COVID19 incidence data presented in Figure 1.

at δ. Using $n_0 = 2$ provides a diffuse prior with large variance, and leads to the default non informative $Beta(1/2, 1/2)$ in the case of $\delta = 1/2$.

The variance for the AR(1) model for ρ_t is crucial, as it controls the memory and smoothness in our approach. Our a priori statement on this variance is the value of the variance multiplier w^*. The prior model for ρ_t is in fact a summation of Gaussian shocks (as in the Weinner process). It is clear then that $\rho_t|\rho_{t-k} \sim N(\rho_{t-k}, \phi_t^{-1}kw^*)$. If the variance is similar to, or larger than, the variance of the observational process, that is $kw^* \geq 1$, then little smoothing is obtained at lag k since ρ_t can vary as much as, or more, than y_{t-k}. Taking $k = \tau$, we set $\tau w^* = 2$, thus $w^* = 2/\tau$, to also limit the level of smoothing to less than τ days. We postulate $\tau w^* = 2$ as a pragmatic large enough variance multiplier, that leads to a quite similar smoothing obtained by Cori et al. (2013) using the same value for τ. This is illustrated in the examples presented in the next Section.

3   Results

A key input to calculating R_t is postulating w_s through an infectious disease generation interval. w_s may be seen as the probability ($\sum_{s=1}^{\infty} w_s = 1$) that an infected person infects other people s days from infection (Fraser, 2007).

A readily way to define w_s is by defining a pdf $f(s)$ and with its cdf let $w_s = F(s) - F(s-1)$. For the COVID19 epidemic, we postulate an Erlang(3, 8/3) pdf depicted in Figure 2(a). The expected value is at 8 days and the maximum infectivity is at 5.5 days, decaying near zero after 20 days. This is based on early (Team et al., 2020; Verity et al., 2020) as well as recent reports on the viral load and implied infectiousness of the decease (see Cevik, Kuppalli, Kindrachuk, and Peiris, 2020, fig. 2). The use of the Erlang distribution in epidemics has been suggested elsewhere (Champredon, Dushoff, and Earn, 2018).

To illustrate that our inferences are stable, we also present an Erlang generation interval shifted to the right, that may also be considered adequate for COVID19. The resulting posterior distributions for the R_t using our method, for Mérida, México, on 03JUL2020 are presented in Figure 2(b). Certainly, under broad regular circumstances, the Bayesian posterior operator is smooth with respect to changes in the prior (see Christen and Pérez-Garmendia, 2020, for example). We continue with an Erlang(3, 8/3) as our generation interval distribution for COVID19 in both examples below.

We use the same data sets presented in Figure 1 to illustrate our approach, plotting the posterior ρ_t time series from the last date in the time series to 60 days before, see Figure 3. Also plotted in Figure 3 are the posterior for the R_ts using Cori et al. (2013) and the observed $R_t = e^{y_t}$s. Regarding the Mérida data, note how Cori et al. (2013) posterior is basically insensitive to the variability observed in the data, the posterior inter-quantile ranges remain basically equal. For our approach, note from 05.05 to 19.05 the large variability in e^{y_t} leads to large posterior ranges whereas from 02.06 to 09.06 a more stable observed R_t results in far shorter posterior inter-quantile ranges. Later on, by the end of June, more spread posteriors are seen again given the observed second wave of large and unstable R_ts. A sharp increase in e^{y_t} follows a sharp decrease, from above 2 to nearly 1, within 10 days until 03.07. Our AR(1) modeling results in a smooth decay on the position of the posterior for R_t, and shorter spread, as a response to the short stable spell of observed e^{y_t}. For these same 10 days, a similar behavior is seen in the position of Cori et al. (2013)'s posterior, but this as a result of the moving window of τ days, as explained in the previous section. As already mentioned, the UQ resulting from the posterior of Cori et al. (2013) responds only to the absolute values of I_ts, which in this case are around 25 to 100 in May and June.

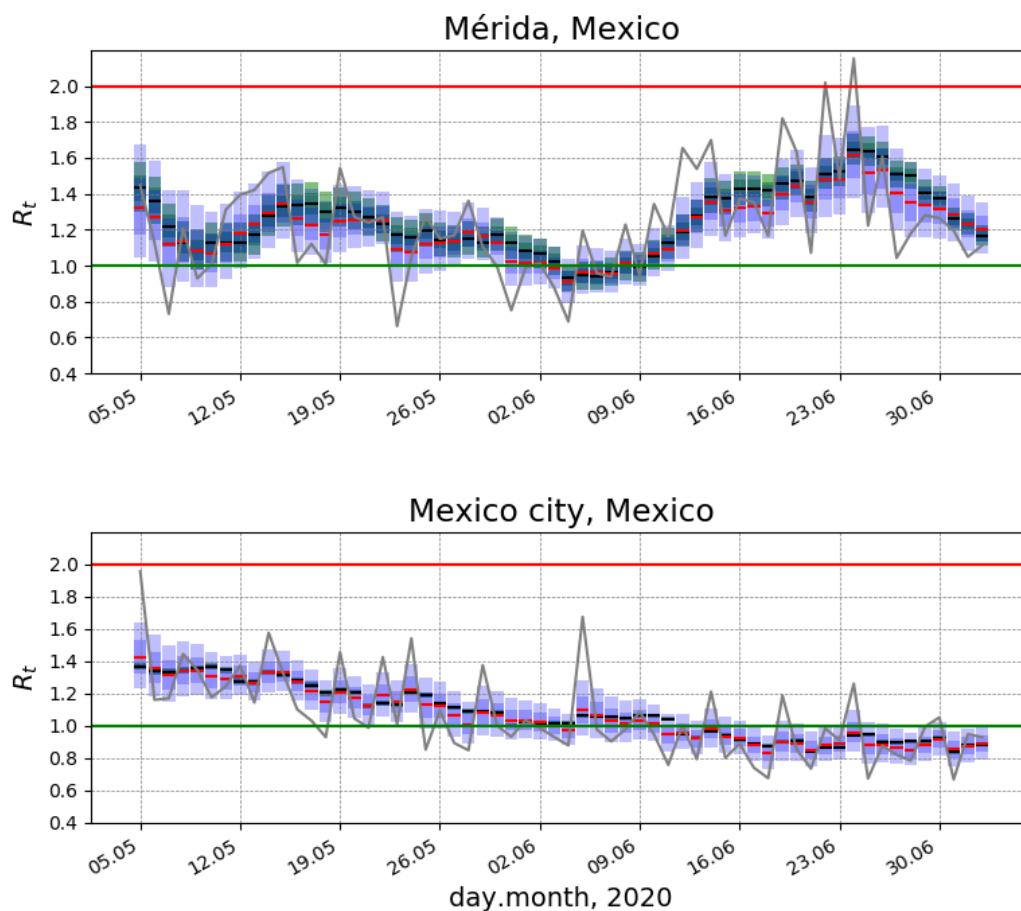Regarding the results from Mexico city, presented in Figure 3, note how

Figure 3: R_t's for incidence data presented in Figure 1 for the last 60 days. From the posterior marginal of each ρ_t in (2) we calculate its 10%, 25%, 50%, 75% and 90% quantiles and these are transformed back to obtain quantiles of R_t. These are plotted vertically with blue shades and the median with a red line. The Poisson model posterior of R_t of Cori et al. (2013) is also plotted, using the same quantiles with green shades and a black line for the median. The posterior of Cori et al. (2013) for Mexico city is so sharp that the quantile shades are hardly seen. The gray lines are the observed $R_t = e^{y_t}$s.

the posterior of Cori et al. (2013) basically collapses around the median, given the large I_t counts of this large metropolitan area (with a population of more than 21 million, I_t varies around 1000 to 2000 in May and June). On the contrary, our UQ provided by the posterior of R_ts presents larger spreads in early May, when an unstable spell is observed for e^{y_t}. Later on, after most of June with a more stable period, the uncertainty decreases with far sharper posteriors by the end of the month and the firsts days of July.

Regarding the smoothing obtained in the estimation of R_t with respect to the observed values, that is, gray lines in Figure 3, see how sharp changes in the latter are reasonably smoothed out by the former. Our strategy to setting the variance multiplier hyperparamenter w^* in the AR(1) model for ρ_t (i.e. $w^* = 2/\tau$, with $\tau = 7$ days, see Section 2.1) results in a smoothing similar to that of Cori et al. (2013)'s, although with larger and adaptive inter-quantile ranges, as already commented. The smoothness obtained by Cori et al. (2013) is a result of assuming a fixed R_t for $\tau = 7$ days and calculating the posterior independently for each R_t. Ours is the result of the DLM and the autoregressive prior model on ρ_t.

4   Discussion

We present a new approach for estimating the effective reproduction number R_t with the following improvements: 1) It does not depend on a compart-mental model, specific to a particular disease. 2) Using previous ideas (Cori et al., 2013; Fraser, 2007) we construct a novel statistical model and Bayesian inference approach to estimate R_t. 3) The resulting UQ is better suited re-sponding to the recent variation of observed R_ts. 4) An AR(1) process is used to promote smoothness on the estimation and 5) all calculations are simple, leading to a Bayesian updating (filtering) type analysis that may be used by non-experts. In fact, a simple to use Python program with an im-plementation of our method (and also the original of Cori et al., 2013) may be downloaded from https://github.com/andreschristen/RTs. We are working also on an alternative implementation in a spreadsheet.

The generation interval w_s is a crucial input to this and other similar methods to estimate R_t. Note that based on preliminary data on infected patients, some basic knowledge can be gained to postulate a reasonable w_s. This was the case since early stages of the COVID19 epidemic (Team et al., 2020). It is easier to gain information on w_s since it represents information on

the infectivity, bound to an average infected person. Social contact patters, lockdown measures etc do not play a role in w_s (from the underlying assumption that $\beta(t, s)$, in the renewal equation $I(t) = \int \beta(t, s)I(t - s)ds$, has the form $\beta(t, s) = R(t)w(s)$, see Fraser, 2007). The social/epidemiological factors influencing contagion patters are coded in R_t, an appealing feature of this type of approaches to estimate the latter, to monitor changes in epidemic data from the onset.

Estimating the infectious disease generation interval w_s requires time consuming clinical studies that are commonly not available during an ongoing epidemic with a new virus as Sars-Cov-2. Thompson et al. (2019) present an attempt to include uncertainty in the estimation of w_s into the posterior distribution of R_t of Cori et al. (2013). Including an estimate of w_s in our method, along with its corresponding uncertainty, is formally straightforward within this Bayesian framework, either by including a likelihood with additional observables or by postulating a prior process for w_s. Indeed, this will complicate the approach and most likely MCMC methods will need to be used.

As already mentioned, a reliable estimation of R_t is quite relevant, to be used as such, or as part of other larger forecasting or monitoring systems. Bettencourt and Ribeiro (2008) elaborate on a simple relation of future incidence with predicted R_ts. A reliable forecasting system of R_t could then be used for short term forecasting of I_t. A related idea is used in Nouvellet et al. (2018) for incidence forecast and other our purposes. Being at time t, the process $\rho_{t+k}, k = 1, 2, \ldots, p$ can be easily simulated, by simulating ρ_t and other parameters from the posterior at D_t and in turn simulating $\rho_{t+k}|\rho_{t+k-1}$ until ρ_{t+p}, to produce MC samples of the posterior predictive distribution for R_t. A feasible path for incidence forecasting based on forecasts of R_t is envisaged, but we need to leave this possibility for future research.

5 Acknowledgments

References

Bettencourt, L. M. A. and R. M. Ribeiro (2008, 05). Real time bayesian estimation of the epidemic potential of emerging infectious diseases. *PLOS ONE 3*(5), 1–9. 2, 12

Capistran, M. A., A. Capella, and J. A. Christen (2020). Forecasting hospital demand in metropolitan areas during the current covid-19 pandemic and estimates of lockdown-induced 2nd waves. *medRxiv*. 4

Cauchemez, S., P.-Y. Boëlle, G. Thomas, and A.-J. Valleron (2006, Sep). Estimating in real time the efficacy of measures to control emerging communicable diseases. *Am J Epidemiol 164*(6), 591–597. 2

Cevik, M., K. Kuppalli, J. Kindrachuk, and M. Peiris (2020). Virology, transmission, and pathogenesis of sars-cov-2. *BMJ 371*. 9

Champredon, D., J. Dushoff, and D. J. Earn (2018). Equivalence of the erlang-distributed seir epidemic model and the renewal equation. *SIAM Journal on Applied Mathematics 78*(6), 3258–3278. 9

Christen, J. and J. L. Pérez-Garmendia (2020). Weak and tv consistency in bayesian uncertainty quantification using disintegration. *Boletín de la Sociaded Matemática Mexicana (accepted)*. 9

Cintrón-Arias, A., C. Castillo-Chávez, L. M. A. Bettencourt, A. L. Lloyd, and H. T. Banks (2009, Apr). The estimation of the effective reproductive number from disease outbreak data. *Math Biosci Eng 6*(2), 261–282. 2

Cori, A., N. M. Ferguson, C. Fraser, and S. Cauchemez (2013, 09). A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology 178*(9), 1505–1512. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12

Ferguson, N. M., D. A. T. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley, and D. S. Burke (2006). Strategies for mitigating an influenza pandemic. *Nature 442*(7101), 448–452. 2

Flaxman, S., S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, et al. (2020). Estimating the effects of non-pharmaceutical interventions on covid-19 in europe. *Nature 584*(7820), 257–261. 3

Fraser, C. (2007, 02). Estimating individual and household reproduction numbers in an emerging epidemic. *PloS one 2*, e758. 2, 8, 11, 12

Fraser, C., S. Riley, R. M. Anderson, and N. M. Ferguson (2004). Factors that make an infectious disease outbreak controllable. *Proceedings of the National Academy of Sciences 101*(16), 6146–6151. 2

Gostic, K. M., L. McGough, E. Baskerville, S. Abbott, K. Joshi, C. Tedijanto, R. Kahn, R. Niehus, J. A. Hay, P. M. De Salazar, J. Hellewell, S. Meakin, J. Munday, N. Bosse, K. Sherratt, R. M. Thompson, L. F. White, J. Huisman, J. Scire, S. Bonhoeffer, T. Stadler, J. Wallinga, S. Funk, M. Lipsitch, and S. Cobey (2020). Practical considerations for measuring the effective reproductive number, rt. *medRxiv*. 3, 5

Kalman, R. E. (1960, 03). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering 82*(1), 35–45. 7

Lindén, A. and S. Mäntyniemi (2011). Using the negative binomial distribution to model overdispersion in ecological count data. *Ecology 92*(7), 1414–1421. 3

McGough, S. F., M. A. Johansson, M. Lipsitch, and N. A. Menzies (2020). Nowcasting by bayesian smoothing: A flexible, generalizable model for real-time epidemic tracking. *Plos-one*. 3

Nouvellet, P., A. Cori, T. Garske, I. M. Blake, I. Dorigatti, W. Hinsley, T. Jombart, H. L. Mills, G. Nedjati-Gilani, M. D. Van Kerkhove, C. Fraser, C. A. Donnelly, N. M. Ferguson, and S. Riley (2018). A simple approach to measure transmissibility and forecast incidence. *Epidemics 22*, 29 – 35. The RAPIDD Ebola Forecasting Challenge. 5, 12

RKI, R. K. I. (2020). Coronavirus disease 2019 (covid-19) - daily situation report of the robert koch institute. 2

Salmon, M., D. Schumacher, and M. Höhle (2016). Monitoring count time series in r: Aberration detection in public health surveillance. *Journal of Statistical Software, Articles 70*(10), 1–35. 3

Team, E. E. et al. (2020). Updated rapid risk assessment from ecdc on the novel coronavirus disease 2019 (covid-19) pandemic: increased transmission in the eu/eea and the uk. *Eurosurveillance 25*(10). 9, 11

Thompson, R., J. Stockwin, R. van Gaalen, J. Polonsky, Z. Kamvar, P. Demarsh, E. Dahlqwist, S. Li, E. Miguel, T. Jombart, J. Lessler, S. Cauchemez, and A. Cori (2019). Improved inference of time-varying reproduction numbers during infectious disease outbreaks. *Epidemics 29*, 100356. 2, 4, 12

Verity, R., L. C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, G. Cuomo-Dannenburg, H. Thompson, P. Walker, H. Fu, et al. (2020). Estimates of the severity of covid-19 disease. *medRxiv*. 9

Wallinga, J. and P. Teunis (2004, 09). Different Epidemic Curves for Severe Acute Respiratory Syndrome Reveal Similar Impacts of Control Measures. *American Journal of Epidemiology 160*(6), 509–516. 2

West, M. and J. Harrison (1997). *Bayesian Forecasting and Dynamic Models (2nd Ed.)*. Berlin, Heidelberg: Springer-Verlag. 6, 7