LUND UNIVERSITY

Identifiability issues in estimating the impact of interventions on Covid-19 spread

Gustafsson, Fredrik; Jaldén, Joakim; Bernhardsson, Bo; Soltesz, Kristian

Identifiability issues in estimating the impact of interventions on Covid-19 spread [★]

Fredrik Gustafsson [*] **Joakim Jaldén** [**]
Bo Bernhardsson [***] **Kristian Soltesz** [***]

[*] *Linköping University, Div. of Automatic Control, Linköping, Sweden*
(e-mail: fredrik.gustafsson@liu.se)
[**] *KTH Royal Institute of Technology, Div. Division of Information*
Science and Engineering, Stockholm, Sweden (e-mail:jalden@kth.se)
[***] *Lund University, Dept. Automatic Control, Lund, Sweden (e-mail:*
{bo.bernhardsson,kristian.soltesz}@control.lth.se)

Abstract: The Covid-19 pandemic has spawned numerous dynamic modeling attempts aimed at estimation, prediction, and ultimately control. The predictive power of these attempts has varied, and there remains a lack of consensus regarding the mechanisms of virus spread and the effectiveness of various non-pharmaceutical interventions that have been enforced regionally as well as nationally. Setting out in data available in the spring of 2020, and with a now-famous model by Imperial College researchers as example, we employ an information-theoretical approach to shed light on why the predictive power of early modeling approaches have remained disappointingly poor.

Keywords: Epidemiology, Covid-19, Sensitivity, Identifiability, Bayesian methods

1. INTRODUCTION

We illustrate how the combination of poor excitation in modelling data and opaque modelling easily lead to the loss of practical identifiability of the resulting model. While being well-known within system theory, these issues have been neglected in remarkably many epidemiological models used to estimate or predict aspects of the ongoing SASRS-CoV-2 pandemic. Several such models have had impact on policy, and ultimately on people's lives. In March 2020 the Imperial College Covid-19 Response Team (ICCRT) issued a report [Ferguson et al. (2020)], based on an agent model, that resulted in an overnight change of the UK response policy, manifested by a nation-wide lockdown [Boseley (2020)]. Later the same month, the ICCRT published a subsequent report [Flaxman et al. (2020b)], in which a Bayesian model ascribed almost all reduction of the viral reproduction number R to lockdowns in all modeled European countries, except for Sweden. The non-pharmaceutical intervention (NPI) that brought $R < 1$ in Sweden was in contrast ascribed to have had only negligible effect on R in all the other countries. A revision of the second model [Flaxman et al. (2020b)] was published in Nature [Flaxman et al. (2020a)]. The model code, including revision history, has been made publicly available [ICCRT (2020)]. When executing the original code, we found that the model suffered sensitivity

issues, leading to practical loss of identifiability under the data available at the time of publication [Soltesz et al. (2020)]. This prompted us to conduct a systematic investigation of the model. Here we provide a summary of that investigation.

2. THE ICCRT NPI MODEL

The basic dynamics of how NPIs affect reported data is illustrated in Fig. 1. Each NPI is assumed to have an effect on the society that affects the spread, through (presumably) decreasing the effective reproduction number R. Since the infected population is not directly measurable, an observation model for some measurable data is needed.

Fig. 1. Block diagram for the ICCRT NPI model.

2.1 NPI Definitions

The binary indicator function $\sigma_{k,t,m}$ models NPI $k = 1, \ldots, K$ effectuated at time $t = 1, \ldots, T$ in country $m = 1, \ldots, M$, where $\sigma_{k,t,m} = \infty$ denotes that NPI k was not effectuated in country m during the studied time interval $t \in [1, T]$).

The modeled NPI categories are:

- Lockdown;
- Public events ban;
- School closure;
- Self isolation;

- Social distancing.

In addition to these, the model ascribes a country-specific effect for the last NPI enacted in each country, modeled as a virtual country-specific NPI.

It is worthy of notice that all five NPI categories were enacted in 10 of the eleven Europeans countries included in [Flaxman et al. (2020a)]. In the eleventh country, Sweden, four of the categories were enacted, but there was no lockdown. All NPI enactments took place in March 2020, most of them within a single week.

2.2 Basic NPI Model

The model relies on three distributions of times between:

- becoming infected and infecting the next person is called the serial interval and is here denoted $p_{II}(\tau)$;
- infection and the onset of symptoms is denoted $p_{IO}(\tau)$;
- symptoms onset and death is denoted $p_{OD}(\tau)$.

The distribution of time between infection and death, $p_{ID}(\tau) = p_{IO} * p_{OD}(\tau)$, is used to relate the number of infected people to the reported number of deaths.

We can now define the NPI model (1). The infected population $I_{t,m}$ at time t in country m having population N_m is given by the recursion

$$I_{t,m} = R_{t,m} \frac{N_m - \sum_{\tau=0}^{t-1} I_{\tau,m}}{N_m} \sum_{\tau=0}^{t-1} I_{\tau,m} p_{II}(t - \tau). \quad (1a)$$

To get some intuition for this formula, consider the special case where $p_{II}(k) = \delta(k)$, that is, every infected person is infectious exactly after k days, and only that day, and then spreads the disease to $R_{t,m} I_{t,m}/N_m$ susceptible. The effect of NPI k on $R_{t,m}$ is modeled by:

$$R_{t,m} = R_{0,m} \exp\left(-\sum_{k=1}^{6} \alpha_k \sigma_{k,t,m}\right). \quad (1b)$$

Note that the parameter α_k is the same for all countries. This is what pools the M countries together within the model.

Finally, we have an observation model for the mortality data $y_{t,m}$,

$$y_{t,m} = \text{IFR} \sum_{\tau=0}^{t-1} I_{\tau,m} p_{ID}(t - \tau), \quad (1c)$$

where IFR denotes the infection-fatality ratio.

3. NPI ESTIMATION APPROACHES

In [Flaxman et al. (2020a)] a Bayesian MCMC framework was used to estimate the parameters $R_{0,m}$, $m = 1, \ldots, M$ and α_k, $k = 1, \ldots, K$. Our contribution has been to recast the model in a linear regression framework where the least squares (LS) method can be used.

3.1 Bayesian Approach

The Bayesian model [Flaxman et al. (2020a)] relies on the following prior distributions:

$$p_{II}(\tau) = \Gamma(6.5, 0.62), \quad (2a)$$
$$p_{IO}(\tau) = \Gamma(5.1, 0.86), \quad (2b)$$
$$p_{OD}(\tau) = \Gamma(18.8, 0.45). \quad (2c)$$

Here, $\Gamma(a, b)$ has mean a, coefficient of variation b, and standard deviation ab.

Priors for the NPI parameters were chosen as $\alpha_k \sim \Gamma(0.5, 1)$ in [Flaxman et al. (2020b)], and subsequently changed to

$$\alpha_k \sim \Gamma(1/6, 1) - \frac{\log(1.05)}{6}, \quad (3)$$

in [Flaxman et al. (2020a)], with the motivation that $\sum \alpha_k \sim U(0, 1.05)$ when all NPIs are effectuated. That is, there is a possibility that the interventions increase R by a factor $e^{1.05}$, but most of the prior is assigned to a significant decrease. As was noted by Nicholas Lewis in [Lewis (2020)], the Γ prior has a regularization effect. Suppose that the model needs $\sum_k \alpha_k = 1.75$, corresponding to a reduction of $1 - e^{1.75} = 0.83$ in R_t which was the median reduction in the MCMC samples of the Nature model [Flaxman et al. (2020a)]. Then the prior for $\alpha_i = 1.76/6$ for all i is 0.0023. If one $\alpha_i = 1.71$ dominates and the others are insignificant, $\alpha_j = 0.01$, then the prior is 64.3. That is, the prior gives a strong bias to having one or a few NPIs that explain data.

Further, $R_{0,m} \sim \mathcal{N}(2.4, |\kappa|)$ where $\kappa \sim \mathcal{N}(0, 0.5)$. The ICCRT model approach is based on sampling from the parameter priors, evolving (1a) and (1b), and then evaluating the likelihood for each sample in a Bayesian MCMC framework. A simplified description of the methodology essentially consists of the following steps:

(1) Generate random numbers of the $K + M$ involved parameters $\theta = (\alpha_:, R_{0,:})$. This is the most important step and there are many sampling strategies to achieve it.
(2) Simulate the model (1).
(3) Compute the likelihood for the observed mortality data.
(4) Generate a random number $u = U(0, 1)$.
(5) If the log likelihood ratio has increased more than u, the simulation is accepted, otherwise it is rejected.
(6) Continue until a pre-defined number of samples has been obtained, excluding the burn in time before the MCMC has converged to stationarity.

3.2 Fisherian Approach

It turns out that the estimation problem can be cast in a Fisherian LS framework. This explicitly reveals identifiability issues of the model, obscured by ambiguous prior choices in the Bayesian MCMC framework.

Equation (1b) is a linear regression in the log domain:

$$\log R_{t,m} = \log R_{0,m} - \sum_{k=1}^{6} \alpha_k \sigma_{k,t,m}. \quad (4)$$

Taking the logarithm also of (1a), and eliminating $R_{t,m}$, gives

$$\log I_{t,m} = \log R_{0,m} - \sum_{k=1}^{6} \alpha_k \sigma_{k,t,m} \qquad (5)$$

$$+ \log\left(1 - \frac{\sum_{\tau=0}^{t-1} I_{\tau,m}}{N_m}\right) + \log\left(\sum_{\tau=0}^{t-1} I_{\tau,m} p_{II}(t-\tau)\right).$$

This fits the classical linear regression framework, with

$$z_{t,m} = x_{t,m}\theta, \qquad (6a)$$

where

$$z_{t,m} = \log I_{t,m} - \log\left(\sum_{\tau=0}^{t-1} I_{\tau,m} p_{II}(t-\tau)\right) \qquad (6b)$$

$$+ \log\left(1 - \frac{\sum_{\tau=0}^{t-1} I_{\tau,m}}{N_m}\right), \qquad (6c)$$

$$x_{t,m} = \left(\sigma_{:,t,m}, \ e_m^T\right), \qquad (6d)$$

$$\theta = \left(\alpha_{:,}, \log\left(R_{0,:}\right)\right)^T. \qquad (6e)$$

Here, e_m is the m^{th} unit vector. Vectorizing the data for all M countries, we get the more compact form

$$Z_t = X_t \theta, \qquad (7a)$$

$$Z_t = (z_{t,1}, \ldots, z_{t,M})^T, \qquad (7b)$$

$$\bar{X}_t = (\sigma_{:,t,:}, \ \mathrm{I}). \qquad (7c)$$

The solution is given by

$$\hat{\theta} = \left(\sum_t X_t^T X_t\right)^{-1} \left(\sum_t X_t^T Z_t\right)^{-1}, \qquad (8a)$$

$$\mathrm{Cov}(\hat{\theta}) = \lambda \left(\sum_t X_t^T X_t\right)^{-1}. \qquad (8b)$$

Here, $\lambda = \mathrm{Var}(z_{t,m})$ denotes the variance of the transformed data, assuming it to be the same for all times and countries.

3.3 Inverse Convolution

The above linear regression model requires knowledge of the number of infected people $I_{t,m}$ to compute $z_{t,m}$ of (6b). It can be estimated by inverse convolution from the model

$$D_{t,m} = \mathrm{IFR} \sum_{\tau=0}^{t-1} I_{\tau,m} p_{ID}(t-\tau). \qquad (9)$$

From a statistical viewpoint, this two-step procedure corresponds to a certainty equivalence assumption that does not attempt to give a correct description of the observation noise. However, it is useful for analyzing identifiability of the model.

4. IDENTIFIABILITY

The lack of excitation in the NPIs constitutes a fundamental limitation. It can be formalized through the Cramér-Rao lower bound (CRLB) that limits the achievable estimation accuracy of any unbiased estimator

$$\mathrm{Cov}(\hat{\alpha}) \geq \lambda \mathrm{J}(\theta^o)^{-1}, \qquad (10)$$

where $\mathrm{J}(\theta^o)$ is the Fisher information matrix (FIM) evaluated at the true parameter values θ^o, and λ is the variance

of the noise on $z_{t,m}$. Since we have a linear regression, the bound simplifies into

$$\mathrm{Cov}(\hat{\alpha}) \geq \lambda \left(\sum_t \bar{X}_t^T \bar{X}_t\right)^{-1}. \qquad (11)$$

4.1 Condition number of the FIM

One can note that X_t is a matrix with binary entries, rendering a FIM $\mathrm{J} = \sum_t X_t^T X_t$ over $\mathbb{Z}^+$. Identifiability is directly related to the condition number of the FIM. The Nature version of the model corresponds to $\mathrm{cond}(J) > 1\,000$, indicating a very high model sensitivity.

4.2 Eigen Decomposition of the FIM

Assume we apply all available data up to May 9, 70 days after March 1, and that we can compute the virtual measurement $z_{t,m}$, having Gaussian noise with variance λ. In this idealistic case, the CRLB dictates

$$\mathrm{Cov}(\hat{\alpha}) \geq \begin{bmatrix} 0.010 & -0.004 & -0.002 & -0.000 & -0.003 \\ -0.004 & 0.029 & -0.013 & -0.004 & -0.008 \\ -0.002 & -0.013 & 0.039 & -0.005 & -0.019 \\ -0.000 & -0.004 & -0.005 & 0.020 & -0.010 \\ -0.003 & -0.008 & -0.019 & -0.010 & 0.039 \end{bmatrix}. \quad (12)$$

That is, the variance of each parameter lies in the range $0.010 - 0.039$. Applying an SVD to the inverse FIM gives

$$D = \mathrm{diag}(0.0588, 0.0409, 0.0261, 0.0116, 0.0003), \quad (13)$$

$$U = \begin{bmatrix} 0.02 & -0.02 & 0.17 & 0.90 & 0.40 \\ -0.15 & 0.72 & -0.49 & -0.10 & 0.46 \\ 0.72 & -0.38 & -0.32 & -0.17 & 0.46 \\ 0.09 & 0.21 & 0.78 & -0.35 & 0.46 \\ -0.67 & -0.54 & -0.13 & -0.17 & 0.46 \end{bmatrix}. \quad (14)$$

The diagonal matrix reveals that one linear combination of the parameter vector can be estimated with a variance bound of only 0.0003, which is 200 times better than the worst linear combination

$$\mathrm{Var}\left(\sum_k \hat{\alpha}_k\right) \geq 0.0015. \qquad (15)$$

The conclusion is that the sum of the parameters has a CRLB that is 6 to 20 times smaller than its individual components.

4.3 Practical implications

The high sensitivity of the ICCRT model is not merely concerning from a theoretical point of view. In [Soltesz et al. (2020)], we have analyzed how seemingly subtle changes in NPI category definitions result in the ICCRT model [Flaxman et al. (2020a)] ascribing wildly varying effectivnessess to different NPI categories. The two changes we considered were introduced by the ICCRT modellers in subsequent code versions. They redefined the crowd size associated with a public events ban, and whether a school closure should also encompass high schools and universities.

The perhaps most remarkable artefact of the identifiability issues is that the ICCRT model ascribes almost all reduction of virus spread during the spring of 2020 to the lockdown NPI in the 10 modeled European countries that implemented one. In tho only country that (according to

the model) did not implement a lockdown, another NPI category provided an almost as large spread reduction as that caused by lockdowns elsewhere. This other NPI category is either of school closure, public events ban, or the encouragement of social distancing, depending on which of the above-mentioned definition variants is adopted. Notably, neither of these three NPI categories had substantial impact in the other 10 countries. According to the model it was thus fortunate that a lockdown was implemented in all countries where it was effective, and omitted in the single country, where another, and elsewhere ineffective NPI, turned out crucial to bring down R.

5. DISCUSSION

Using established principles from systems theory, we have demonstrated that even a seemingly simplistic and data-driven phenomenological model can suffer from severe input sensitivity and identifiability issues.

There are indeed additional properties of the ICCRT model that would justify scrutiny. For instance, the base assumption that changes in R are solely driven by NPIs constrain the model to explain any reduction in R through the NPIs. Another problematic aspect is the stochastic delay, with a mean of about 3-4 weeks, between infection and the death report, which is what is typically directly measured. This adds further uncertainty to the (pseudo) observations of the regression model, and thus the parameter estimates. The perhaps most remarkable model property is, however, the special role ascribed to the last NPI enacted in a particular country. This prohibits prospective use of the model, since it is impossible to tell whether the latest enacted NPI was also the last.

Unfortunately, high sensitivity cannot solely be ascribed to questionable choices underlying the studied model. Instead, practical parameter identifiability is to a large extent inherently limited by available data. All NPIs were effectuated during a period of 22 days, most of them in one single week. The regression tensor therefore has many common leading zeros as well as many common trailing ones for all parameters, indicating poor excitation. Furthermore, several of the involved countries had very little virus spread when the NPIs were effectuated, leading to an almost zero output.

6. CONCLUSION

With much at stake during all phases of a pandemic, we conclude that it is crucial to thoroughly scrutinise any SARS-CoV-2 estimation or prediction model, prior to considering its use as decision support in policy-making. Such scrutiny relies on modellers following the practice used by the ICCRT in sharing open source code.

ACKNOWLEDGEMENTS

REFERENCES

Boseley, S. (2020). New data, new policy: why UK's coronavirus strategy changed. *The Guardian*. URL https://www.theguardian.com/world/2020/mar/16/new-data-new-policy-why-uks-coronavirus-strategy-has-changed. Cited 2020, May 3.

Ferguson, N., Laydon, D., et al. (2020). Report 9 - Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. doi:10.25561/77482. Preprint.

Flaxman, S., Mishra, S., et al. (2020a). Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, (584), 257–261. doi:10.1038/s41586-020-2405-7.

Flaxman, S., Mishra, S., et al. (2020b). Report 13 - Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries. doi:10.25561/77731. Preprint.

ICCRT (2020). ICCRT Github repository. doi:10.5281/zenodo.3888697.

Lewis, N. (2020). Did lockdowns really save 3 million COVID-19 deaths, as Flaxman et al. claim? URL https://judithcurry.com/2020/06/21/did-lockdowns-really-save-3-million-covid-19-deaths-as-flaxman-et-al-claim. Blog post.

Soltesz, K., Gustafsson, F., et al. (2020). Sensitivity analysis of the effects of non-pharmaceutical interventions on COVID-19 in Europe. doi:10.1101/2020.06.15.20131953. Preprint.