CoviHawkes: Temporal Point Process and Deep Learning based Covid-19 forecasting for India

Ambedkar Dukkipati, **Tony Gracious**, and **Shubham Gupta**

Department of Computer Science and Automation,
Indian Institute of Science, Bangalore, INDIA.
`[ambedkar, tonygracious, shubhamg]@iisc.ac.in`

Abstract

Lockdowns are one of the most effective measures for containing the spread of a pandemic. Unfortunately, they involve a heavy financial and emotional toll on the population that often outlasts the lockdown itself. This article argues in favor of "local" lockdowns, which are lockdowns focused on regions currently experiencing an outbreak. We propose a machine learning tool called CoviHawkes based on temporal point processes, called CoviHawkes that predicts the daily case counts for Covid-19 in India at the national, state, and district levels. Our short-term predictions (< 30 days) may be helpful for policymakers in identifying regions where a local lockdown must be proactively imposed to arrest the spread of the virus. Our long-term predictions (up to a few months) simulate the progression of the pandemic under various lockdown conditions, thereby providing a noisy indicator for a potential third wave of cases in India. Extensive experimental results validate the performance of our tool at all levels.

1 Introduction

We live in strange times. The Covid-19 pandemic has disrupted the lives and livelihood of people in unfathomable, if not unprecedented, ways. It has brought the devastating effects of social isolation to the center stage and challenged the very fabric of our society and its economy. While public safety measures like lockdowns help arrest the spread of the virus, unless imposed early on, their primary purpose is to delay the inevitable to allow healthcare systems to prepare themselves for the looming crises. As such, these lockdowns should be treated as strategic decisions, and data-aware methodologies should guide their implementation. This article provides technical details of a deep learning-based tool called CoviHawkes[1] that forecasts India's daily case counts at the national, state, and district levels. The predictions made by our tool may help the policymakers identify vulnerable regions to enact "local" lockdowns proactively.

Lockdowns have a disproportionate impact on different people. For example, while employees working in an IT firm can often continue their job remotely, taxi drivers and food delivery personnel lose their daily livelihood. The spending power of the consumers reduces as more people lose their income, which damages the overall economy and leads it into a vicious circle. The consequences of these lockdowns often outlast their duration by a wide margin. Combined with the fact that the infection outbreaks are almost always localized, it arguably makes more sense to enforce lockdowns in regions experiencing an outbreak instead of imposing nationwide or statewide lockdowns. We refer to these more focussed lockdowns as *local lockdowns.*

For local lockdowns to be effective, it is crucial to enforce them at an early stage while the infec-

---

[1] https://sml.csa.iisc.ac.in/covihawkes

tions are still contained within the region of interest. This requires not only continuous monitoring of the daily case counts but also accurate forecasts to help local administration take appropriate action in advance. CoviHawkes uses a deep learning model to make such predictions at the national, state, and district level. The proposed tool learns from: **(i)** historical patterns in the case counts, **(ii)** information about the average movement of people in the region (hereafter referred to as mobility features), and **(iii)** the region's demographic information. The data used by our tool is publicly available and aggregated at an appropriate level to ensure that it cannot be traced back to individuals, thus preserving their privacy.

At the heart of CoviHawkes is a well-known stochastic process known as the *Hawkes process* (Hawkes, 1971). Loosely speaking, a Hawkes process characterizes the rate at which such events happen. (Here, A person getting infected with the virus can be thought of as an event). It defines a probability distribution over the occurrences of these events in a way that makes future events more likely if the rate has been high in the recent past. In other words, it models a scenario where future infections are more likely if many people have caught the disease in the recent past. For this reason, the Hawkes process is known as a *self-exciting process*. Due to the nature of the spread of infectious diseases, a Hawkes process is a leading candidate for modeling such a phenomenon (Chiang et al., 2021; Garetto et al., 2021). Our tool CoviHawkes uses deep neural networks to model the parameters of the Hawkes process as a function of the input features mentioned in the previous paragraph.

The remainder of this article is organized as follows: Section 2 describes the underlying mathematical framework and our deep learning model in detail. Section 3 establishes the validity of short-term forecasts generated by our tool using standard machine learning procedures for model validation. Section 4 describes our approach for generating long-term forecasts and presents the corresponding results under various lockdown conditions. We wish to emphasize that CoviHawkes can notably predict short-term forecasts at the level of each district that makes it especially useful for policymakers in devising strategies for local lockdowns. We hope that such an approach would minimize the negative consequences of blanket lockdowns while still effectively containing the spread of the disease.

2   CoviHawkes Model

In this section, we describe our method for a region of interest $\mathcal{R}$ that can be a district, state, or nation. Let $C(t) \in \mathbb{N}$ denote the number of new Covid-19 cases in the region $\mathcal{R}$ at time t. We use $\mathbf{m}(t) \in \mathbb{R}^{d_m}$ to denote a vector of mobility features for region $\mathcal{R}$ that describes the percentage change in various activities such as "going to work", "staying at home", "grocery shopping", etc, in this region on day t. Additionally, we will use $\mathcal{H}(t)$ to refer to the history of observed data till time t, i.e., $\mathcal{H}(t) = \{(C(s), \mathbf{m}(s))\}_{s \leq t}$.

Conditioned on the history $\mathcal{H}(t-1)$, our model assumes that $C(t)$ is a Poisson distributed random variable with mean $\lambda(t)$. Therefore, for all $c = 0, 1, 2, \ldots$,

$$\mathrm{P}\left(C(t) = c \mid \mathcal{H}(t-1)\right) = \frac{\lambda(t)^c \exp(-\lambda(t))}{c!}.$$

Following the discrete-time Hawkes process literature, we model $\lambda(t)$ as

$$\lambda(t) = \mu + \sum_{i=1}^{L} w(L-i)R(t-i)C(t-i).$$

Here, $\mu \in [0, \infty)$ is the base count rate for the region of interest, L is the length of the time window within which past case counts affect the present case count, $R(t)$ is the *reproduction number* of the virus, and $w(1), \ldots, w(L)$ are weights assigned to the previous L days such that $w(i) \geq 0$ and $\sum_{i=1}^{L} w(i) = 1$.

In other words, infections occur on day t with a base rate of μ. However, a person may also catch the virus from another person infected i days ago $(i \leq L)$. The probability of such a infection is proportional to the weight $w(L - i)$ assigned to that day and the reproduction number $R(t-i)$ of the virus i days ago. The parameters μ, $w(1), \ldots, w(L)$, and $R(t)$ are learned directly from the data by maximizing the log-likelihood of the observations. We model $R(t)$ as a function of the past mobility data using an LSTM as described in Section 2.1.

In the real world, a region has a finite population and it is rare for people to get infected twice.

Moreover, vaccinated people are also less likely to get infected, and hence can be removed from the set of susceptible population members for modeling purposes (Rizoiu et al., 2018). To consider these factors, we use a discounted value of $\lambda(t)$ as shown below:

$$\tilde{\lambda}(t) = \left(1 - \frac{N(t-1) + V(t-1)}{N}\right) \lambda(t).$$

Here, $N(t-1) = \sum_{s=1}^{t-1} C(s)$ is the total number of people that have already been infected before time t, $V(t-1)$ is the number of vaccinated people, and N is the total size of the population in the region. As $N(t-1)$ and $V(t-1)$ increase, the discounted rate $\tilde{\lambda}(t)$ decreases. All parameters including the parameters of the LSTM that computes $R(t)$ are learned via maximum-likelihood estimation.

2.1 Estimating the value of $R(t)$

$R(t)$ measures the rate at which the virus reproduces itself. A higher value of $R(t)$ entails faster spread of the disease and vice versa. Owing to the emergence of variants of the virus, the value of $R(t)$ changes over time and must be estimated periodically using recent trends in the data. To do so, we use an LSTM (Hochreiter and Schmidhuber, 1997). Let $\mathbf{x}(t) \in \mathbb{R}^{(L+1)d_m}$ be a vector obtained by concatenating the mobility vectors $\mathbf{m}(t-i-\Delta)$ for $i = 1, \ldots, L$, and the vector of case counts in the previous L days $[C(t-L), \ldots, C(t-1)]$. The mobility features used at time t at delayed by an additional gap $\Delta > 0$ to take the incubation period of the disease into account. We use $d_m = 6$, $L = 28$, and $\Delta = 14$ in our experiments. $R(t)$ is then computed as:

$$\mathbf{h}(t) = \text{LSTM}(\mathbf{h}(t-1), \mathbf{x}(t))$$
$$R(t) = \ln(1 + \exp(\mathbf{w}^\mathsf{T}\mathbf{h}(t) + b)).$$

Here, $\mathbf{h}(t) \in \mathbb{R}^d$ is the hidden state of the LSTM at time t, and $\mathbf{w} \in \mathbb{R}^d$ and $b \in \mathbb{R}$ are learnable parameters.

3 Short-term predictions

For validating our short-term predictions, we use case counts $C(t)$ and mobility features $\mathbf{m}(t)$ collected between March 2, 2020, and July 20, 2021. We divide this data into two subsets: training data
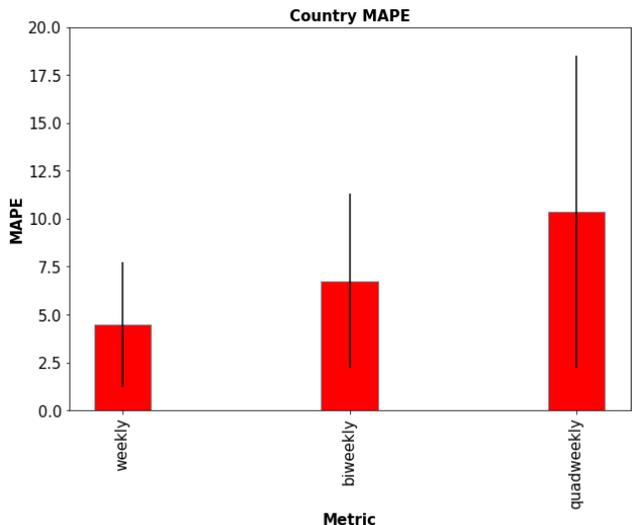
Figure 1: Country-Level MAPE scores for different sizes of forecasting window.

(March 2, 2020 to April 27, 2021) and validation data (April 27-July 20, 2021), and validate our model for forecasting windows of three sizes: 7, 14, and 28 days.

The validation period has 84 days. To validate the model for a forecasting window of size w, we divide the validation period into n possibly overlapping intervals of size w each. Let $I_1^w, I_2^w, \ldots, I_n^w$ be these intervals, then:

$$I_i^w = \{t_\text{s} + 7(i-1) + k\}_{k=0}^{w-1},$$

where t_s marks the start of the validation period. To make predictions for interval I_i^w, we independently train a model using all available data till the start of I_i^w and compute our predictions for this interval. Let $\hat{C}_i^w(t)$ be the prediction made by such a model at time t. The error for interval I_i^w is calculated as:

$$\psi(I_i^w) = \frac{|\sum_{t \in I_i^w} C(t) - \sum_{t \in I_i^w} \hat{C}_i^w(t)|}{\sum_{t \in I_i^w} C(t)} \times 100.$$

The error metric above is known as the Mean Absolute Percentage Error (MAPE). The process above is repeated for all $i = 1, \ldots, n$, and the error of the model on a window of size w is given by:

$$E(w) = \frac{1}{n} \sum_{i=1}^{n} \psi(I_i^w).$$

Figure 1 shows $E(w)$ for $w = 7$, 14, and 28 for the nation-level model that forecasts case counts

(a) $w = 7$ (b) $w = 14$ (c) $w = 28$
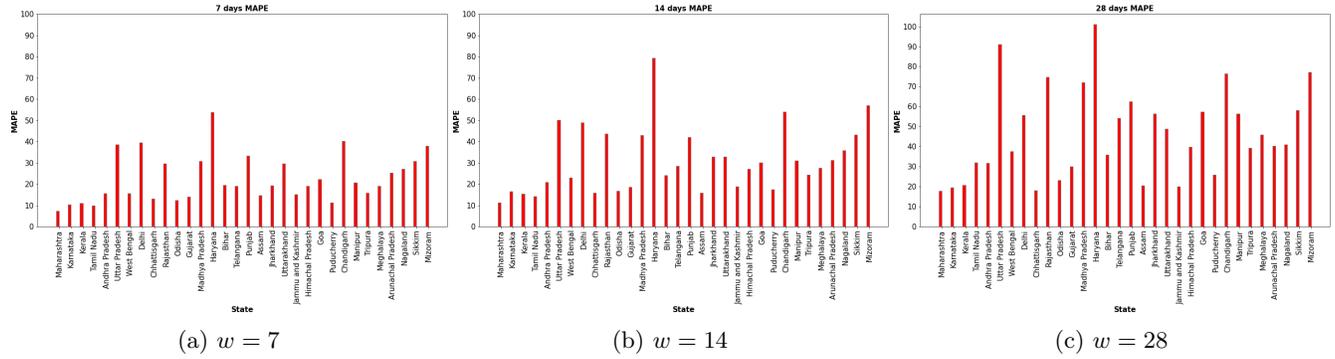
Figure 2: MAPE for models trained at the state level for different states and values of forecast window size w.

(a) $w = 7$ (b) $w = 14$ (c) $w = 28$
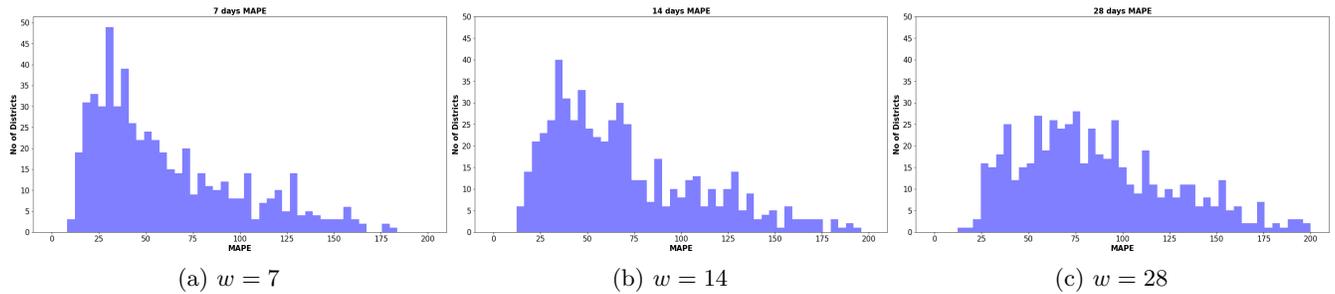
Figure 3: Histogram of MAPE scores for models trained at district level for different values of forecast window size w.

4

for the country as a whole. Figure 2 shows the same data for state-level models. As the number of districts is very large (> 500), we use histograms in Figure 3 to show the number of districts on the y-axis that have the corresponding MAPE scores on the x-axis. As before, Figure 3 has histograms for different values of w. The latest forecasts based on this model are available on our website https://sml.csa.iisc.ac.in/covihawkes.

4   Long-term predictions

Recall that our model uses mobility features $\{\mathbf{m}(t - \Delta - i)\}_{i=1}^{L}$ and case counts $\{C(t - i)\}_{i=1}^{L}$ from the last few days to forecast the future value at time t. Thus, we can only generate a forecast at time t if these feature values are available. To make predictions over a longer time horizon, we need appropriate proxies for the values of these features. The case counts can be boot-strapped, i.e., the model can treat its own past predictions as observed ground truth values. However, simulating the mobility values is much harder. Instead, we take a simplified approach as explained below.

We first identify four time intervals in the past that correspond to different lockdown conditions. The first interval, $I_s = $ [March 25-April 14, 2020], covers the time when there was a strict nationwide lockdown. The second interval, $I_m = $ [December 13-19, 2020], falls in the seventh *unlock* phase in India where most, though not all, of the restrictions, were lifted. The third interval, $I_n = $ [February 15-March 3, 2020] corresponds to the time before the pandemic when there was no lockdown. Finally, $I_c = $ [August 13-19, 2021], corresponds to the current mobility conditions. Table 1 summarizes the conditions under these intervals.

We generate four long-term forecasts, one each for the mobility conditions from the intervals I_s, I_m, I_n, and I_c. To generate forecasts corresponding to I_{x} ($\mathrm{x} \in \{s, m, n, c\}$), we compute the average value of mobility features for each weekday $i = 1, \ldots, 7$ in I_{x}. Let $D(i)$ be the set of all i^{th} weekdays between February 14, 2020 and August 20, 2021. For example, $D(1)$ will be the set of all Sundays in this interval. Then, define $\mathbf{m}_{\mathrm{x}}(i)$ as,

$$\mathbf{m}_{\mathrm{x}}(i) = \frac{1}{|I_{\mathrm{x}} \cap D(i)|} \sum_{t \in I_{\mathrm{x}}} \mathbf{1}\{t \in D(i)\}\mathbf{m}(t).$$

Whenever we require the mobility value for a day
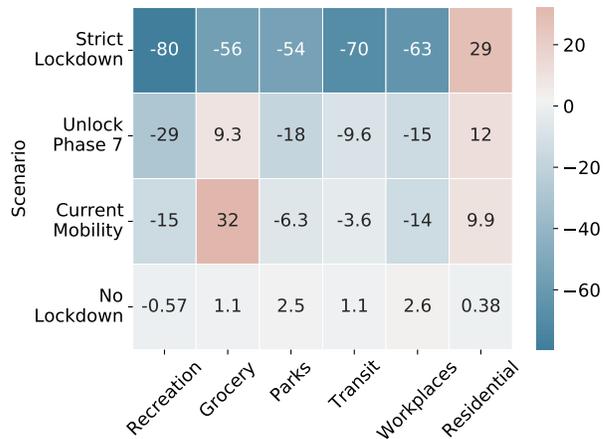
Figure 4: Heatmap showing average mobility values under various lockdown conditions. The columns correspond to the six mobility features that we use. Negative values indicate less mobility (the values represent percentage change as compared to a pre-pandemic baseline).

t in the future that has not been observed, we use $\mathbf{m}_{\mathrm{x}}(i)$ for an appropriate i chosen based on the day of the week at time t. This, together with bootstrapped count values, enables us to use our model to generate long-term forecasts.

Figure 4 shows the average mobility values $\frac{1}{7}\sum_{i=1}^{7}\mathbf{m}_{\mathrm{x}}(i)$ for various lockdown conditions and Figure 5 shows the long-term forecast at the nation-level under these conditions. One can see from Figure 4 that the current conditions have a higher mobility value as compared to the seventh unlock phase, despite the restrictions being similar in these two phases (see Table 1). Consequently, the model predicts a sharper rise in cases under I_c as compared to I_m. One can also see that the model predicts a significant third wave if the mobility returns to the pre-pandemic conditions (I_n). Interestingly, the model also indicates that a mild lockdown (as in I_m) would be almost as good as a very strict lockdown (as in I_s).

As is the case with every statistical model, Covi-iHawkes also makes several assumptions. For example, we train models independently for states and districts and do not consider the movement of people across these regions due to the lack of this data. Similarly, the model does not consider breakthrough infections, reinfections, or emergence of new variants of the virus in the long term. While our short-term forecasts are rigorously validated,

Interval	Allowed	Not Allowed
I_s	Essential government services like defence, police, and power generation. Healthcare systems, related industries, and emergency services. Banks. Shops selling ration, meat, dairy, and animal fodder. Internet services. E-commerce. Media.	All non-essential offices, shops, and industries. All transport (except for essential goods). All educational institutes, places of worship, and gatherings of any kind. Hospitality services.
I_m	All businesses, some at limited capacity. Graded reopening of schools. Cinema/Theaters at 50% capacity. Gatherings of up to 100 people in closed spaces with maximum occupancy of 50%. Transportation.	Large gatherings in indoor spaces. Different states have different restrictions.
I_c	Same as I_m, but with higher mobility values (see Figure 4)	
I_n	Everything open, business as usual	

Table 1: Lockdown conditions. I_s: Strict lockdown, I_m: Unlock Phase 7, I_c: Current conditions, I_n: No lockdown. See Section 4 for details about these conditions.
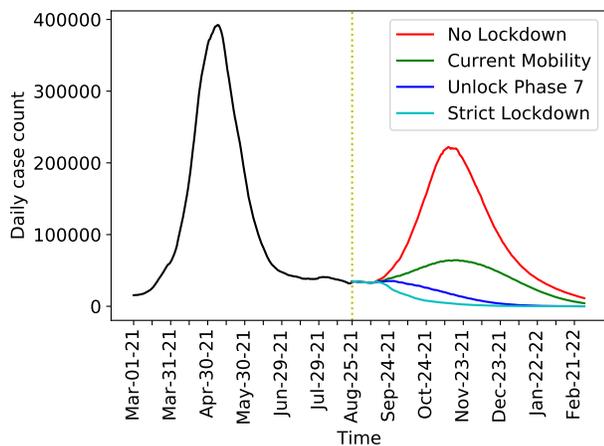
Figure 5: Nation-level long-term forecast.

validating long-term forecasts can be challenging. In particular, long-term forecast uses proxy values for the features as described above, and hence the errors compound over time. Moreover, usage of average mobility features ignores events like festivals and/or other unforeseen mass gatherings that are likely to trigger an uptick in the infection rate. As such, one must keep these caveats in mind while using these long-term predictions.

5 Conclusion

In this article, we introduced an AI tool for forecasting Covid-19 case counts in India at nation, state, and district levels. This model is based on Hawkes process, which is suitable for modeling the spread of infectious diseases due to its self-exciting nature. We rigorously validated the short-term predictions made by our model using standard validation procedures for time-series data. We also used our model to generate long-term forecast at the national level to provide an indication of case counts under different lockdown conditions. We hope that this model, especially its short-term forecasts at the district level, will be useful for policymakers in developing strategies for local lockdowns.

Acknowledgements

References

Chiang, W.-H., Liu, X., and Mohler, G. (2021). Hawkes process modeling of covid-19 with mobility leading indicators and spatial covariates. *International Journal of Forecasting*.

Garetto, M., Leonardi, E., and Torrisi, G. L. (2021). A time-modulated hawkes process to model the spread of covid-19 and the impact of countermeasures. *Annual Reviews in Control*, 51:551–563.

Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Rizoiu, M.-A., Mishra, S., Kong, Q., Carman, M., and Xie, L. (2018). Sir-hawkes: Linking epidemic models and hawkes processes to model diffusions in finite populations. International World Wide Web Conferences Steering Committee.