Information Theoretic Model Selection for Accurately Estimating Unreported COVID-19 Infections

Jiaming Cui[1], Arash Haddadan[2], A S M Ahsan-Ul Haque[3], Bijaya Adhikari[4], Anil Vullikanti[2,3], and B. Aditya Prakash[1,*]

[1]College of Computing, Georgia Institute of Technology, Atlanta, GA 30332
[2]Biocomplexity Institute, University of Virginia, Charlottesville, VA 22904
[3]Department of Computer Science, University of Virginia, Charlottesville, VA 22904
[4]Department of Computer Science, The University of Iowa, Iowa City, IA 52242

Abstract

Estimating the true magnitude of infections was one of the significant challenges in combating the COVID-19 outbreak early on. Our inability in doing so allowed unreported infections to drive up disease spread in numerous regions in the US and worldwide. Even today, identifying the true magnitude (the number of total infections) is still challenging, despite the use of surveillance-based methods such as serological studies, due to their costs and biases. This paper proposes an information theoretic approach to estimate total infections accurately. Our approach is built on top of ordinary differential equations based epidemiological models, which have been used extensively in understanding the dynamics of COVID-19, and aims to estimate the true total infections and a parameterization that "best describes" the observed reported infections. Our experiments show that the parameterization learned by our framework leads to a better estimation of total infections and forecasts of the reported infections compared to a "baseline" parameterization, which is learned via usual model calibration. We also demonstrate that our framework can be leveraged to simulate what-if scenarios with non-pharmaceutical interventions. Our results also support earlier findings that most COVID-19 infections were unreported and non-pharmaceutical interventions indeed helped mitigate the COVID-19 outbreak. Our approach gives a general method to use information theoretic techniques to improve epidemic modeling, which can also be applied to other diseases.

[*]To whom correspondence should be addressed. E-mail: baindityap@cc.gatech.edu

Introduction

The COVID-19 pandemic has emerged as one of the most formidable public health challenges in recent history. By Dec 1, 2021, it had already resulted in more than 48 million reported infections and 0.78 million deaths in the United States alone. Worldwide reported infections total at 262 million and deaths at 5.22 million [18]. The devastating effect of COVID-19 extends for the public health sector and the economy as well. For example, in the US, the unemployment rate peaked at 15.8 percent in April 2020 [6], and U.S. GDP contracted at a 3.5% annualized rate for 2020 [1]. Similar economic impacts have been observed worldwide.

One of the significant challenges in combating the COVID-19 pandemic early on was the inability to estimate the true magnitude of infections. As noted by many studies [13, 11, 55, 53, 39], a significant number of COVID-19 infections were unreported, due to various factors such as the lack of testing and asymptomatic infections. In the early stages, transmissions in previously unreported regions were facilitated by unreported infections before being detected. For example, phylogenetic studies revealed that COVID-19 had locally spread in Washington state before active community surveillance was implemented in early 2020 [14]. Similarly, there were only 23 reported infections in five major U.S. cities by March 1, 2020. However, it has been estimated that there were already more than 28000 total infections in those cities by then [5]. Similar trends were observed in Italy, Germany, and the UK [58], and this continues to be a challenge even today [8, 30].

More generally, accurate estimation of the true magnitude of infections is a fundamental epidemiological question and critical for pandemic planning and response. To capture it, we define the *reported rate* as the ratio of reported infections to total infections in this paper. This contrasts with the more extensively studied notion of *case ascertainment rate*, which is the ratio of reported symptomatic infections to the actual number of symptomatic infections [50] (i.e., this doesn't take asymptomatic infections into account). We focus on the reported rate because asymptomatic infections are also infectious and contribute substantially to the community transmission as many studies have shown [56, 41]. Estimation of the reported rate is challenging, and there are limitations in all current methods, as as discuss next.

One of the most effective methods to identify the reported rate in a region is through large-scale serological studies [54, 26, 62]. These surveys use blood tests to identify the prevalence of antibodies against SARS-CoV-2 in a large population. The CDC COVID Data Tracker portal [2, 26] summarizes the results of serological studies conducted by commercial laboratories at a national level as well as at 10 specific sites. For example, the estimated reported rate was at most 0.1 in Minneapolis and South Florida as of April 2020 (i.e., at least 10 times more total infections than reported infections). While serological studies can give an accurate estimation, they are expensive and are not sustainable in the long run [4]. It is even more challenging to perform them in real-time as there are unavoidable delays between sample collection and laboratory tests [2, 26]. Serological studies can also suffer from sampling biases, and heuristics need to be designed to account for them [9].

Other lines of work to estimate the reported rate for COVID-19 exploits existing influenza surveillance systems to estimate symptomatic infections [40], due to symptomatic similarities between the two diseases. However, they also suffer from ad-hoc corrections to account for the similarities between COVID-19 and influenza symptoms.

In the face of these challenges, data scientists and epidemiologists have devoted much time and effort to estimate the reported rate through epidemiological models [39, 50, 12, 48, 36, 38, 60, 25, 33, 61]. They calibrate these models to reported data (say the time series of reported infections D_{reported}) to learn parameterizations and use them to estimate the reported rate. For an epidemiological model, denoted by O_{M}, we refer to a parameterization learned from such model

Figure 1: The workflow of MDLINFER. (a) The baseline parameterization underestimates the unreported infections (Shown in light green in the iceberg plot in (c)). (b) Our MDL model selection approach MDLINFER estimates the magnitude of COVID-19 infections accurately by inferring unreported infections given the reported infections. In MDLINFER, the DATA is the reported infections D_{reported}, which a hypothetical Sender S wants to transmit to a Receiver R. At a high level, the sender S will send the DATA by sending the total infections D, the "baseline parameterization" $\mathbf{p}$, and the "candidate" parameterization $\mathbf{p}$'. (c) MDLPARAM reveals that a large majority of COVID-19 infections were unreported. We visualize the estimated reported rate using the iceberg. The height of the iceberg over water (in white) and the height below the water (in blue) are proportional to the reported and unreported infections estimated using MDLINFER. The number of COVID-19 reported infections is just the tip of an iceberg, while the whole iceberg corresponds to the total infections estimated by MDLINFER.

calibration as a "baseline parameterization" $\mathbf{p}$, and the estimated reported rate as α_{reported}; later we will consider specific choices of O_{M} from recent work. However, this "model estimated" reported rate is still problematic because of ad-hoc calibrating assumptions such as predefining a set of parameters [17] or heuristics for parameter estimation (e.g., limiting the search space to a narrow range [19]). As we will show later in our results, minor errors in estimates of reported rates may lead to very different forecasts of future trends and intervention policy recommendations. Hence a principled approach is needed in order to estimate the actual reported rate.

To achieve this goal, we propose a new information theory-based approach. For simplicity, let us assume the reported rate to be a constant (we discuss how we handle time-varying rates in the next section). Our central intuition is: imagine we are given D, a time series consisting of the accurate values of total infections within a given time period. By calibrating the epidemiological model O_{M} against D, the model can learn a "candidate" parameterization $\mathbf{p}$' and use it to calculate the corresponding reported rate $\alpha'_{\text{reported}}$. Intuitively, this parameterization $\mathbf{p}$' should be better than the baseline parameterization $\mathbf{p}$ because it is done using the total infections as well not just the reported ones. Using this D and $\alpha'_{\text{reported}}$, we should be able to consequently also "describe"

the reported infections D_{reported} "better" (as D_{reported} should be just $\alpha'_{\text{reported}} \times D$). One way to measure the quality of this description is to compute the number of bits needed to encode it. Hence our problem is one of developing a principled framework to choose among competing descriptions (i.e., perform "model selection") such that we can find the one that takes the fewest number of bits. So we develop an information theoretic framework to estimate the actual total infections - by posing an optimization problem which (roughly) aims to find the D which minimizes the number of bits required to encode the observed D_{reported}.

Following the discussion above, to encode the D_{reported}, we intuitively need the variables D and $\alpha'_{\text{reported}}$. Let us denote DATA $= D_{\text{reported}}$ and MODEL $= (D, \alpha'_{\text{reported}})$. Note that this MODEL is meant for encoding purposes only and is different from the base epidemiological ODE model O_M. The optimization problem then is defined over all possible values of MODEL to encode DATA efficiently. Following the convention in information theory, we use $L(\text{MODEL})$ to denote the number of bits required to encode the MODEL; and $L(\text{DATA}|\text{MODEL})$ to denote the number of bits required to encode the DATA, D_{reported}, given the MODEL. The overall objective of our optimization problem is to infer an optimal MODEL*, which minimizes $L(\text{MODEL})+L(\text{DATA}|\text{MODEL})$. To be more specific, we in fact set our MODEL $= (D, \mathbf{p'}, \mathbf{p})$ in our complete formulation (see Methods section for details).

The optimization framework described above follows the popular two-part Minimum Description Length (MDL) framework which follows the general maxim of 'induction by compression'. The two-part MDL (aka sender-receiver framework) consists of hypothetical actors S and R. Sender S has DATA and wants to transmit it to receiver R using as few bits as possible [24]. Hence, sender S searches for the best possible MODEL, which minimizes the overall cost of encoding and transmitting both the MODEL and the DATA given the MODEL. Note that the two-part MDL (and MDL in general) does not assume the nature of the DATA or the MODEL, which exactly fits our objective. It has been widely used for numerous optimization problems ranging from network summarization [34], causality inference [16], and failure detection in critical infrastructures [10]. MDL has also previously been used for some epidemiological problems, mainly in inferring patient-zero and associated infections in cascades over contact networks [47]. However, our technique is the first to propose an MDL-based approach on top of ODE-based epidemiological models, which are harder to formulate and optimize.

We call our MDL-based optimization framework MDLINFER. Our experiments showcase our method using two different ODE-based epidemiological models as O_M: SAPHIRE model [25] and SEIR+HD model [33]. They both have previously performed well in fitting reported infections and highlighted beneficial insights for the COVID-19 response. The SAPHIRE model focuses on two key features of the outbreak: high covertness and high transmissibility that drove the outbreak of COVID-19 in Wuhan. The SEIR+HD model investigates how non-pharmaceutical interventions like social distancing will be needed to maintain epidemic control. These models are representative enough to show that MDLINFER gives consistent performance with multiple epidemiological models with different dynamics. As we will show later, we do not estimate an accurate reported rate and total infections in both models without our framework. We use BASEPARAM to represent the baseline parameterization $\mathbf{p}$ obtained via standard model calibration. Similarly, we use MDLPARAM to represent the optimal parameterization identified by MDLINFER. Our experiments show that MDLPARAM is clearly superior to BASEPARAM in estimating total infections. An example is shown in Figure 1. By March 11, 2020, the Minneapolis Metro Area had only 16 COVID-19 reported infections. The best version of BASEPARAM estimated 182 total infections (colored as light green in the iceberg). On the other hand, our MDLPARAM gave an estimate of 301 total infections shown below the sea level, which is closer to the total infections estimated from serological studies [26, 2]. Additionally, MDLPARAM also leads to better a fit and *future* projections on reported infections. We also demonstrate that MDLPARAM can aid policy making by analyzing

counter-factual non-pharmaceutical interventions, while inaccurate BASEPARAM estimates lead to wrong non-pharmaceutical intervention conclusions.

Results

Here, we present our empirical findings on a set of experiments at different geographical regions and time periods. We choose these regions and periods based on the severity of outbreak and the availability of the serological studies and symptomatic surveillance dataset. As mentioned above, we use the SAPHIRE model [25] and SEIR+HD model [33] as the base epidemiological models O_M. We show consistently good performance for MDLINFER on both models.

In each region, we divide the timeline into two time periods: (i) observed period, when only the number of reported infections are available, and the models are calibrated to learn the parameterizations, and (ii) future period, where we evaluate the forecasts generated by the epidemiological models and calibrated parameterizations learned in the observed period. To handle the time-varying reported rates, we divide the observed period into multiple sub periods and learn different rates for each sub period separately.

(A) Estimating Total Infections: MDLPARAM estimates total infections more accurately than BASEPARAM

Here, we use the point estimates of the total infections calculated from serological studies as the ground truth (black dots shown in Figure 2). We call it SEROSTUDY$_{\text{Tinf}}$. We also plot MDLPARAM's estimation of total infections, MDLPARAM$_{\text{Tinf}}$, in the same figure (red curve). To compare the performance of our approach and the baselines with SEROSTUDY$_{\text{Tinf}}$, we use the cumulative value of estimated total infections. However, serological studies numbers are not directly comparable with the total infections because of the lag between antibodies becoming detectable and infections being reported [2, 26]. In Figure 2, we have already accounted for this lag following the CDC serological studies [2, 26] (See Methods section for details). The vertical black line shows a 95% confidence interval for SEROSTUDY$_{\text{Tinf}}$. The blue curve represents BASEPARAM's estimation of total infections, BASEPARAM$_{\text{Tinf}}$. As seen in the figure, MDLPARAM$_{\text{Tinf}}$ falls within the confidence interval of the estimates given by serological studies. Significantly, in Figure 2 (c) and (d) for South Florida, the BASEPARAM for the SAPHIRE model [25] overestimates the total infections, while the SEIR+HD model underestimates the total infections. However, MDLPARAM consistently predicts the total infections correctly. This observation shows that as necessary, MDLPARAM$_{\text{Tinf}}$ can improve upon the BASEPARAM$_{\text{Tinf}}$ in either direction (i.e., by increasing or decreasing the total infections). Note that the MDLPARAM$_{\text{Tinf}}$ curve from both the models are similar even when the BASEPARAM$_{\text{Tinf}}$ curves are different. This shows that MDLINFER is consistently able to estimate total infections more accurately.

To quantify the performance gap between the two approaches, we first compute the root mean squared error (RMSE) between SEROSTUDY$_{\text{Tinf}}$ and BASEPARAM$_{\text{Tinf}}$. We then compute the same between SEROSTUDY$_{\text{Tinf}}$ and MDLPARAM$_{\text{Tinf}}$. We then compute the ratio, ρ_{Tinf}, of the two RMSE errors as $\frac{\text{RMSE}(\text{BASEPARAM}_{\text{Tinf}}, \text{SEROSTUDY}_{\text{Tinf}})}{\text{RMSE}(\text{MDLPARAM}_{\text{Tinf}}, \text{SEROSTUDY}_{\text{Tinf}})}$. Note that the values of ρ_{Tinf} being greater than 1 imply that the MDLPARAM$_{\text{Tinf}}$ is closer to SEROSTUDY$_{\text{Tinf}}$ estimates than BASEPARAM$_{\text{Tinf}}$. In Figure 2 (e), we plot ρ_{Tinf}. Overall, the ρ_{Tinf} values are greater in Figure 2 (a) to (d) (1.20, 2.62, 5.46, and 1.22), which indicates that MDLPARAM performs better overall. Note that even when the value of ρ_{Tinf} is 1.20, the improvement made by MDLPARAM$_{\text{Tinf}}$ over BASEPARAM$_{\text{Tinf}}$ in terms of RMSE is about **12091**. Hence, one can conclude that MDLPARAM$_{\text{Tinf}}$ is indeed superior to BASEPARAM$_{\text{Tinf}}$, when it comes to estimating total number of infections.

Figure 2: BASEPARAM's and MDLPARAM's estimation of total infections. (a)-(d) MDLPARAM's estimation of total infections is more accurate than BASEPARAM's. The grey dash line divides the observed and future periods (which was not accessible to the model while calibrating). The blue and red curves represent BASEPARAM's estimation of total infections, BASEPARAM$_{\text{Tinf}}$, and MDLPARAM's estimation of total infections, MDLPARAM$_{\text{Tinf}}$, respectively. The black point estimates and confidence intervals represent the total infections estimated by serological studies [2, 26], or SEROSTUDY$_{\text{Tinf}}$. Note that the plot corresponds to different geographic regions for SAPHIRE and SEIR+HD epidemiological models, and hence the scales differ. (e) The performance metric, ρ_{Tinf}, comparing MDLPARAM$_{\text{Tinf}}$ against BASEPARAM$_{\text{Tinf}}$ is shown for the regions in (a)-(d). Here, the values of ρ_{Tinf} are 1.20, 2.62, 5.46, and 1.22, implying that MDLPARAM performs better in the estimation of total infections than BASEPARAM for different geographic regions and epidemiological models.

(B) Estimating Reported Infections: MDLPARAM leads to better fit and projection than BASEPARAM at different stages of the COVID-19 epidemic

Here, we use the same observed period to learn the parameterizations. We then forecast the reported infections in the future periods, which were not accessible to the model while training. The results are summarized in Figure 3. In (a) to (d) and (f) to (i), the vertical grey dash divides the observed and future period. The black plus symbols represent New York Times reported infections, NYT-Rinf. The blue curve represents BASEPARAM's estimation of reported infections, BASEPARAM$_{\text{Rinf}}$. Similarly, the red curve represents MDLPARAM's estimation of reported infections, MDLPARAM$_{\text{Rinf}}$. Note that the curves to the right of the horizontal grey line are the future predictions. As seen in the figure, MDLPARAM$_{\text{Rinf}}$ aligns more closely with NYT-Rinf than BASEPARAM$_{\text{Rinf}}$, indicating the superiority of MDLINFER in estimating reported infections for the future period.

We define a performance metric ρ_{Rinf} as $\frac{\text{RMSE}(\text{BASEPARAM}_{\text{Rinf}}, \text{NYT-Rinf})}{\text{RMSE}(\text{MDLPARAM}_{\text{Rinf}}, \text{NYT-Rinf})}$ to compare MDLPARAM$_{\text{Rinf}}$ against BASEPARAM$_{\text{Rinf}}$ in a manner similar to above. In Figure 3 (e) and (j), we plot the ρ_{Rinf} for the observed and the future period. In both (e) and (j), we notice that the ρ_{Rinf} is close to or greater than 1. This further shows that MDLPARAM$_{\text{Rinf}}$ has a better or at least closer fit for reported infections than BASEPARAM$_{\text{Rinf}}$. Additionally, the ρ_{Rinf} for the future period is even greater than ρ_{Rinf} for the observed period, which shows that MDLPARAM$_{\text{Rinf}}$ performs even better in forecast than BASEPARAM$_{\text{Rinf}}$.

Note that Figure 3 (a) to (d) correspond to the early state of the COVID-19 epidemic in spring and summer 2020, and Figure 3 (f) to (i) correspond to fall 2020. Here, the value of ρ_{Rinf} is greater than 1 indicating that MDLPARAM$_{\text{Rinf}}$ outperforms BASEPARAM$_{\text{Rinf}}$ and MDLINFER performs well in estimating temporal patterns at different stages of the COVID-19 epidemic.

Figure 3: BASEPARAM's and MDLPARAM's estimation of reported infections. (a)-(d) and (f)-(i) MDLPARAM's estimation of reported infections has better fit and projection than BASEPARAM's. The grey dash line divides the observed period and future period. Black plus symbols represent the infections reported by the New York Times (NYT-Rinf). The blue curve represents BASEPARAM's estimation of reported infections BASEPARAM$_{\text{Rinf}}$. Similarly, red curve represents MDLPARAM's estimation of reported infections MDLPARAM$_{\text{Rinf}}$. Each plot corresponds to a different geographic region for the SAPHIRE model or SEIR+HD model with different scales. The results show that MDLPARAM$_{\text{Rinf}}$ aligns closer with NYT-Rinf than BASEPARAM$_{\text{Rinf}}$. Besides, MDLPARAM$_{\text{Rinf}}$ has a better projection for the future period than BASEPARAM$_{\text{Rinf}}$. (e) and (j) The performance metric, ρ_{Rinf}, comparing MDLPARAM$_{\text{Rinf}}$ against BASEPARAM$_{\text{Rinf}}$, is shown for the regions for both the observed period (e), and the future period (j). MDLPARAM$_{\text{Rinf}}$ is close to or better than BASEPARAM$_{\text{Rinf}}$ in both the observed and future period.

(C) Estimating Symptomatic Rate: MDLPARAM estimates the symptomatic rate more accurately than BASEPARAM

We validate this observation using Facebook's symptomatic surveillance dataset [49]. Here we plot the estimated symptomatic rate over time and overlay the estimates and standard error from the symptomatic surveillance data (See Figure 4). The red and blue curves are the MDLPARAM's and BASEPARAM's estimation of symptomatic rates, MDLPARAM$_{\text{Symp}}$ and BASEPARAM$_{\text{Symp}}$ respectively. Note that the SAPHIRE model does not contain states corresponding to the symptomatic infections. Therefore, we only focus on the SEIR+HD model. We compare the trends of the MDLPARAM$_{\text{Symp}}$ and BASEPARAM$_{\text{Symp}}$ with the symptomatic surveillance results. We focus on trends rather than actual values because the symptomatic rate numbers could be biased [49] (see Methods section for a detailed discussion) and therefore cannot be compared directly with model outputs (like what we have done for serological studies). As seen in the figure, the MDLPARAM$_{\text{Symp}}$ captures the trends of the surveyed symptomatic rate RATE$_{\text{Symp}}$ (represented by black plus symbols) much better than BASEPARAM$_{\text{Symp}}$.

To summarize, these three sets of experiments in (A), (B) and (C) together demonstrate that BASEPARAM fails to accurately estimate all the infections including unreported ones. On the other hand, MDLPARAM estimates total infections closer to those estimated by serological studies and better fits reported infections and symptomatic rate trends.

Figure 4: MDLPARAM's estimation of symptomatic rate is more accurate than BASEPARAM's. The blue and red curves represent the BASEPARAM's estimation of symptomatic rate (BASEPARAM$_{\text{Symp}}$), and MDLPARAM's estimation of symptomatic rate (MDLPARAM$_{\text{Symp}}$) respectively. The black points and the shaded regions are the point estimate with standard error for RATE$_{\text{Symp}}$ (the COVID-related symptomatic rates derived from the symptomatic surveillance dataset [49, 51]).

Evaluating the effect of Non-pharmaceutical Interventions

(D) MDLPARAM reveals that a large majority of COVID-19 infections were unreported

We computed the cumulative reported rate MDLPARAM$_{\text{Rate}}$ measured by the ratio of the cumulative value of reported infections to the total infections estimated by MDLPARAM over time and plotted it in Figure 5 (a). The figure shows that the MDLPARAM$_{\text{Rate}}$ increases in early March, and then gradually decreases. This observation is explained by the community spread-driven COVID-19 outbreaks that were not reported until early March, which fits the earlier study [40].

(E) Non-pharmaceutical interventions on asymptomatic and presymptomatic infections are essential to control the COVID-19 epidemic

Our simulations show that non-pharmaceutical interventions on asymptomatic and presymptomatic infections are essential to control COVID-19. Here, we plotted the simulated reported infections of MDLPARAM in Figure 5 (c) (red curve). We then repeated the simulation of reported infections for 5 different scenarios: (i) isolate just the reported infections, (ii) just the symptomatic infections, and isolate symptomatic infections in addition to (iii) 25%, (iv) 50%, and (v) 75% of both asymptomatic and presymptomatic infections. In our setup, we assume that the infectivity reduces by half when a person is isolated. As seen in the figure, when only the reported infections are isolated, there is almost no change in the "future" reported infections. However, when we isolate both the reported and symptomatic infections, the reported infections decreases significantly. Even here, the reported infections are still not decreasing. On the other hand, non-pharmaceutical interventions for some fraction of asymptomatic and presymptomatic infections decrease reported infections. Thus, we can conclude that non-pharmaceutical interventions on asymptomatic infections are essential in controlling the COVID-19 epidemic.

Figure 5: Cumulative reported rate and non-pharmaceutical intervention simulation results. (a) The cumulative reported rate trend during COVID-19 outbreak. The blue and red curve represent the cumulative reported rate of BASEPARAM, $\text{BASEPARAM}_{\text{Rate}}$, and of MDLPARAM, $\text{MDLPARAM}_{\text{Rate}}$, respectively. The black point estimate and its confidence interval represent the reported rate $\text{SEROSTUDY}_{\text{Rate}}$ estimated by serological studies [2, 26]. (b) The grey dash line divides the observed period and future period. The blue curve represents the BASEPARAM's estimation of reported infections, $\text{BASEPARAM}_{\text{Rinf}}$. The other five curves represent the simulated reported infections for the 5 scenarios described in the Results section. (c) Non-pharmaceutical interventions on asymptomatic and presymptomatic infections are essential to control the COVID-19 epidemic. The red curve represents the MDLPARAM's estimation of reported infections, $\text{MDLPARAM}_{\text{Rinf}}$. The other five curves represent the simulated reported infections for the same 5 scenarios as in (b).

(F) Accuracy of non-pharmaceutical intervention simulations relies on the good estimation of unreported infections

Here, we also plot the simulated reported infections generated of BASEPARAM in Figure 5 (b) (blue curve). As seen in the figure, based on BASEPARAM, we can infer that only non-pharmaceutical interventions on symptomatic infections are enough to control the COVID-19 epidemic. However, this has been proven to be incorrect by prior studies and real-world observations [41]. Therefore, we can conclude that the accuracy of non-pharmaceutical intervention simulation relies on the quality of the learned parameterization.

Discussion and Future Work

This study proposes MDLINFER, a data-driven model selection approach that automatically estimates total infections based on the optimal parameterization of epidemiological models. Our approach leverages the information theoretic Minimum Length Description (MDL) principle, to select total infections that "best describe" the observed outbreak. Our approach addresses several gaps in current practice including the long-term infeasibility of serological studies [26], and ad-hoc assumptions of epidemiological models. MDLINFER employs a principled method that selects the total infections which "describes" the reported infections best. At the same time, the existing epidemiological models [33, 39, 43, 25] typically rely on the baseline parameterization learned by ad-hoc heuristics. The MDLINFER framework can also be adapted to work on a set of calibrated parameterizations to generate uncertainty estimates.

Overall, our results show that MDLINFER estimates total infections at various geographical

locations and different epidemiological models more accurately than the baseline parameterization from both directions, i.e., it corrects both over- and under-estimates. For example, compared to the baseline parameterization, we correctly estimate 55719 more infections by April 1 for the SEIR+HD model (Figure. 2 (d)), and 87636 fewer infections for the SAPHIRE model (Fig. 2 (c)) for South Florida. We also show that MDLINFER leads to a better fit of the reported infections in the observed period and more accurate forecasts for the future period than the baseline parameterization. We reveal that a large majority of COVID-19 infections were unreported, where non-pharmaceutical interventions on unreported infections can help to mitigate the COVID-19 outbreak. We also show that MDLINFER estimates a more accurate symptomatic rate than the baseline parameterization. Additionally, our results show consistent performance with respect to the reported infections and serological studies on both the SAPHIRE and SEIR+HD models.

The MDLINFER framework is likely to be helpful in the surveillance of COVID-19 in the near future, and for future epidemics. Even with the U.S. returning to normalcy, surveillance of the pandemic is still essential for public health. The daily incidence of COVID-19 has decreased from early 2021 to summer 2021, according to the CDC COVID Data Tracker portal [7, 35]. However, new variants of the SARS-CoV-2 (e.g., the Delta and Omicron variants) have been spreading rapidly [37, 46, 21]. Testing for these new variants and large-scale surveillance via laboratory tests may be limited and less systematic than what was done for COVID-19 before. In such settings, using our MDLINFER framework, epidemiologists and policymakers can improve the accuracy of estimates of total infections (without large-scale serological studies), as well as forecasts of their models.

One of the limitations of our work is that the benefits of using MDLINFER depends on the suitability of the epidemiological model. If the epidemiological model is not expressive enough for the observed data, then the gains from MDLINFER may not be significant. As future work, it may be helpful to adapt MDLINFER to measure the quality of the base epidemiological model. We also note that MDLINFER is built on ODE-based epidemiological models; other kinds of epidemic models, e.g., agent-based models [42, 28, 57, 17, 44], are more suitable in some settings. It would be interesting to extend MDLINFER to incorporate such models. Finally, there is significant population heterogeneity in disease outcomes [15, 31], e.g., differences in severity rate or mortality rate, when infected with COVID-19, for different age group [22, 27], which has not been considered in our work.

To summarize, MDLINFER is a robust data-driven method to accurately estimate total infections, which will help data scientists, epidemiologists, and policy-makers to further improve existing ODE-based epidemiological models, make accurate forecasts, and combat the ongoing COVID-19 pandemic. More generally, MDLINFER opens up a new line of research in epidemic modeling using information theoretic frameworks.

Materials and Methods

Data

Datasets

We use the following publicly available datasets for our study:

1. **New York Times reported infections [3]:** This dataset (NYT-Rinf) consists of the daily time sequence of reported COVID-19 infections D_{reported} and the mortality $D_{\text{mortality}}$ (cumulative values) for each county in the US starting from January 21, 2020 to current.

2. **Serological studies [26, 2]:** This dataset consists of the point and 95% confidence interval estimates of the prevalence of antibodies to SARS-CoV-2 in 10 US locations every 3–4 weeks

from March to July 2020. For each location, CDC works with commercial laboratories to collect the blood specimens in the population and test them for antibodies to SARS-CoV-2. Each specimen collection period ranges from 6 to 14 days. As suggested by prior work [32, 45], these serological studies have high sensitivity to antibodies for 6 months after infections. Hence, using the prevalence and total population in one location, we can compute the estimated total infections $\text{SEROSTUDY}_{\text{Tinf}}$ for the past 6 months (i.e., from the beginning of the pandemic since January 2020). However, this $\text{SEROSTUDY}_{\text{Tinf}}$ can not be compared with the epidemiological model estimated total infection numbers directly. The reasons are (i) the antibodies may take 10 to 14 days delay to be detectable after infection [63, 52] and (ii) the 6-14 range period for specimen collection as mentioned before. To account for this, we compare the $\text{SEROSTUDY}_{\text{Tinf}}$ numbers with the model estimated total infections of 7 days prior to the first day of specimen collection period (as also suggested by the CDC serological studies work [26]).

3. **Symptomatic surveillance [49, 51]:** This dataset consists of point estimate $\text{RATE}_{\text{Symp}}$ (and standard error) of the COVID-related symptomatic rate for each county in the US starting from April 6, 2020 to date. The survey asks a series of questions on randomly sampled social media (Facebook) users to estimate the percentage of people who have a COVID-like symptoms such as as the fever along with cough or shortness of breath or difficulty breathing on a given day. However, there are several caveats such as they could not cover all symptoms of COVID-19 and these symptoms can be also caused by many other conditions, due to which they are not expected to be unbiased estimates for the true symptomatic rate [49]. Hence we focus on trends instead of the exact numbers. As the original symptomatic surveillance data is at a county level, we weighted the numbers using the total population of each county to compute the $\text{RATE}_{\text{Symp}}$.

Our Approach

Minimum Length Description Principle

We formulate the problem of estimating total COVID-19 infections using the Minimum Length Description (MDL) principle [10, 24]. We use the two-part sender-receiver framework. The conceptual goal of the framework is to transmit the DATA from the possession of the hypothetical sender S to the hypothetical receiver R. We assume the sender does this by first sending a MODEL and then sending the DATA under this MODEL. In the MDL framework, we want to minimize the number of bits for this process. We do this by identifying the MODEL that encodes the DATA such that the total number of bits needed to encode both the MODEL and the DATA is minimized. Hence our cost function (the total number of bits needed) is composed of two parts: (i) model cost $L(\text{MODEL})$: The cost in bits of encoding the MODEL and (ii) data cost $L(\text{DATA}|\text{MODEL})$: The cost in bits of encoding the DATA given the MODEL. Intuitively, the idea is that a good MODEL will lead to a fewer number of bits needed to encode both MODEL and DATA. We formulate the general MDL optimization problem as follows: Given the DATA, $L(\text{MODEL})$, and $L(\text{DATA}|\text{MODEL})$, find MODEL^* such that

$$\text{MODEL}^* = \arg\min_{\text{MODEL}} L(\text{MODEL}) + L(\text{DATA}|\text{MODEL}) \tag{1}$$

In our situation, the DATA is the reported COVID-19 infections D_{reported}: it is the only real-world data given to us (indeed note that total infections are not directly observed). As described in the introduction section, the MODEL is intuitively $(D, \alpha'_{\text{reported}})$. Here D refers to a candidate total infections time series, and $\alpha'_{\text{reported}}$ is the corresponding reported rate from D. To compute

$\alpha'_{\text{reported}}$, we need the candidate parameterization **p'** (after O_M is calibrated using D). Further, we choose to also add **p**, making our MODEL to be $(D, \mathbf{p'}, \mathbf{p})$. There are alternative MODELs that can be considered, but we choose this MODEL $= (D, \mathbf{p'}, \mathbf{p})$ and explain more in the Supplementary Information. Next, we will introduce how we formulate our problem of estimating total infections using the MDL principle.

MDL Formulation

First, we need to introduce some notations used for MDL Formulation. By running the epidemiological model with **p**, O_M will output reported infections. However, this output is an estimate of the reported infections by O_M rather than the exact D_{reported}. This is because even though the model may have been calibrated using D_{reported}, the calibration cannot be perfect, and there will be differences between these estimated reported infections and the D_{reported}. Therefore, we term this estimated reported infections as $D_{\text{reported}}(\mathbf{p})$. Similarly, we have the $D_{\text{reported}}(\mathbf{p'})$ and estimated total infections $D(\mathbf{p'})$ for **p'**. As described in the introduction section, we can also calculate the estimated reported rate α_{reported} and $\alpha'_{\text{reported}}$ using **p** and **p'**. With these notations, next we will formulate the space of all possible MODEL and give the equation for the cost in bits of encoding MODEL and DATA.

The Space of MODEL

We have MODEL $= (D, \mathbf{p'}, \mathbf{p})$ as described above. Hence our MODEL space will be all possible daily sequences for D and all possible parameterizations for **p** and **p'**. The MDL framework will search in this space to find the MODEL*.

Model Cost

With MODEL $= (D, \mathbf{p'}, \mathbf{p})$, we conceptualize the model cost by imagining that the sender S will send the MODEL $= (D, \mathbf{p'}, \mathbf{p})$ to the receiver R in three parts: (i) first send the **p** (sent by encoding **p** directly) (ii) next send the **p'** given **p** (sent by encoding $\mathbf{p'} - \mathbf{p}$) and (iii) then send D given **p'** and **p** (sent by encoding $\alpha'_{\text{reported}} \times D - D_{\text{reported}}(\mathbf{p})$). We term the model cost as $L(D, \mathbf{p'}, \mathbf{p})$, with three components: $\text{COST}(\mathbf{p})$, $\text{COST}(\mathbf{p'}|\mathbf{p})$, and $\text{COST}(D|\mathbf{p'}, \mathbf{p})$. Hence,

$$L(D, \mathbf{p'}, \mathbf{p}) = \text{COST}(\mathbf{p}) + \text{COST}(\mathbf{p'} - \mathbf{p}|\mathbf{p}) + \text{COST}(\alpha'_{\text{reported}} \times D - D_{\text{reported}}(\mathbf{p})|\mathbf{p'}, \mathbf{p}) \qquad (2)$$

For Equation 2, the $\text{COST}(\cdot)$ function gives the total number of bits we need to spend in encoding each term. The details of the encoding method can be found in the Supplementary Information.

Data Cost

We need to send the DATA $= D_{\text{reported}}$ next given the MODEL. Given MODEL $= (D, \mathbf{p'}, \mathbf{p})$, we send DATA by encoding $\frac{D - D_{\text{reported}}}{1 - \alpha'_{\text{reported}}} - D(\mathbf{p'})$. Since D, **p** and **p'** have already been sent, the receiver R could recover back to D_{reported} from $\frac{D - D_{\text{reported}}}{1 - \alpha'_{\text{reported}}} - D(\mathbf{p'})$. We term data cost as $L(D_{\text{reported}}|D, \mathbf{p'}, \mathbf{p})$ and formulate it as Equation 3.

$$L(D_{\text{reported}}|D, \mathbf{p'}, \mathbf{p}) = \text{COST}(\frac{D - D_{\text{reported}}}{1 - \alpha'_{\text{reported}}} - D(\mathbf{p'})|D, \mathbf{p'}, \mathbf{p}) \qquad (3)$$

Total Cost

With $L(D, \mathbf{p'}, \mathbf{p})$ as in Equation 2 and $L(D_{\text{reported}}|D, \mathbf{p'}, \mathbf{p})$ as in Equation 3 above, the total cost $L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})$ is:

$$
\begin{aligned}
L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p}) &= L(D, \mathbf{p'}, \mathbf{p}) + L(D_{\text{reported}}|D, \mathbf{p'}, \mathbf{p}) \\
&= \text{COST}(\mathbf{p}) + \text{COST}(\mathbf{p'} - \mathbf{p}|\mathbf{p}) \\
&+ \text{COST}(\alpha'_{\text{reported}} \times D - D_{\text{reported}}(\mathbf{p})|\mathbf{p'}, \mathbf{p}) \\
&+ \text{COST}(\frac{D - D_{\text{reported}}}{1 - \alpha'_{\text{reported}}} - D(\mathbf{p'})|D, \mathbf{p'}, \mathbf{p})
\end{aligned}
\tag{4}
$$

Problem Statement

Note that our main objective is to estimate the total infections D. With $L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})$ formulated in Equation 4, we can state the problem as: Given the time sequence D_{reported} and epidemiological model O_{M}, find D^* that minimizes the MDL total cost i.e.

$$
D^* = \arg \min_{D} \; L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})
\tag{5}
$$

Algorithm

Next, we will present our algorithm to solve the problem in Equation 5. Note that directly searching D^* naively is intractable since D^* is a daily sequence not a scalar. Instead, we propose first finding a "good enough" reported rate $\alpha^*_{\text{reported}}$ quickly with the constraint $D = \frac{D_{\text{reported}}}{\alpha^*_{\text{reported}}}$ to reduce the search space. Then with this $\alpha^*_{\text{reported}}$, we can search for the optimal D^* in Equation 5. Hence we propose a two-step algorithm: (i) do a linear search to find a good reported rate $\alpha^*_{\text{reported}}$ (ii) given the $\alpha^*_{\text{reported}}$ found above, use an optimization method to find the D^* that minimizes $L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})$ with $\alpha^*_{\text{reported}}$ constraints.

Step 1: Find the $\alpha^*_{\text{reported}}$

In step 1, we do a linear search on α_{reported} to find the $\alpha^*_{\text{reported}}$. As stated before, we use $\frac{D_{\text{reported}}}{\alpha_{\text{reported}}}$ as the D in $L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})$ to help reduce the search space. Here, we formulate step 1 algorithm as Equation 6.

$$
\alpha^*_{\text{reported}} = \arg \min_{\alpha_{\text{reported}}} L(D_{\text{reported}}, \frac{D_{\text{reported}}}{\alpha_{\text{reported}}}, \mathbf{p'}, \mathbf{p})
\tag{6}
$$

Step 2: Find the D^* given $\alpha^*_{\text{reported}}$

With the $\alpha^*_{\text{reported}}$ found in step 1, we next find the D^* that minimizes the $L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})$. Note that we have already found a good $\alpha^*_{\text{reported}}$, we can constrain the D^* to ensure that the sum of D^* equals to the sum of $\frac{D_{\text{reported}}}{\alpha^*_{\text{reported}}}$. We use the Nelder-Mead method [20] solve this constrained optimization problem for D^*. Here, we formulate step 2 algorithm as Equation 7.

$$
D^* = \arg \min_{D} L(D_{\text{reported}}, D, \mathbf{p'}, \mathbf{p})
\tag{7}
$$

We describe the two-step algorithm in more detail in the Supplementary Information.

Calibrate Procedure

Our idea is to use the baseline parameterization **p** and candidate parameterization **p'** for the MDL formulation. To learn the parameterization **p** and **p'**, we need a calibrate procedure to learn them. Typically, there are many calibration procedures in literature [29, 59], and any calibrate procedure is acceptable. Here, we assume the existence of the calibrate procedure for O_M and term it as CALIBRATE. The CALIBRATE takes two inputs: O_M and the dataset. We further assume the dataset contains at least $D_{reported}$. This CALIBRATE uses loss function to learn parameterizations, usually root mean squared error (RMSE) [23] or log likelihood [33, 25] loss functions. Here, the baseline parameterization **p** is learned as Equation 8

$$\mathbf{p} = \text{CALIBRATE}(O_M, \{D_{reported}, \text{others}\}) \tag{8}$$

We are also trying to find a better candidate parameterization **p'** by calibrating on D (and $D_{reported}$). Our idea is to calibrate **p'** on $D_{reported}$ and the daily sequence of unreported infections ($D_{unreported}$). We use the same CALIBRATE procedure but just add another term on $D_{unreported}$ in the loss function. The candidate parameterization **p'** is learned as Equation 9.

$$\mathbf{p'} = \text{CALIBRATE}(O_M, \{D_{reported}, D_{unreported}, \text{others}\}) \tag{9}$$

SEIR+HD model and SAPHIRE model

Next, we describe how we use the two epidemiological models: SEIR+HD model and SAPHIRE model. As described in the main article, we use BASEPARAM and MDLPARAM to represent baseline parameterization and optimal parameterization identified by MDLINFER respectively. Note that as our MDLINFER is a principled and general framework, other epidemiological models or calibration procedures are also acceptable.

SEIR+HD model

The SEIR+HD model [33] consists of 10 states: Susceptible S, exposed E, pre-symptomatic I_P, severe symptomatic I_S, mild symptomatic I_M, asymptomatic I_A, hospitalized (eventual death) H_D, hospitalized (eventual recover) H_R, recovered R, and dead D. The calibration procedure described in Kain et al [33] only infers the transmission rate β_0 (the transmission rate in the absence of interventions), σ (the proportional reduction on β_0 under shelter-in-place), and E_0 (number of initial infections). The other parameters are fixed. They assume the importations only happen at the beginning of the pandemic (captured by E_0), and the total population N remains constant. Note that our work can be also extended to time-varying importations as well easily. We also extend the calibration procedure to infer two more parameters: α (proportion of asymptomatic infections) and α_1 (proportion of new symptomatic infections that are reported). We compute the newly reported infections and unreported infections as follows:

1. New reported infections = $\alpha_1 \times (dI_P I_S + dI_P I_M)$: Here $dI_P I_S + dI_P I_M$ is the number of new symptomatic infections everyday. $dI_P I_S$ is the number of patients switching their state from I_P to I_S (and similarly for $dI_P I_M$). We assume α_1 proportion of new symptomatic infections every day are reported.

2. New unreported infections = $(1 - \alpha_1) \times (dI_P I_S + dI_P I_M) + dE I_A$.

SAPHIRE model

The SAPHIRE model [25] consists of 7 states: Susceptible S, exposed E, pre-symptomatic P, ascertained infectious I, unascertained infectious A, hospitalized D, and recovered R. The calibration procedure [25] only infers the transmission rate β and reported rate r while keeping other parameters fixed. We also compute the newly reported infections and unreported infections as follows:

1. New reported infections $= \frac{\alpha P}{D_p}$: Here $\frac{P}{D_p}$ is the number of new infections from pre-symptomatic every day. D_p is the parameter for the presymptomatic infectious period and is fixed. α is the reported rate estimated by the epidemiological model.

2. New unreported infections $= \frac{(1-\alpha)P}{D_p}$.

BASEPARAM and MDLPARAM

Finally, we describe how we get the results from BASEPARAM and MDLPARAM. Here, we use the baseline parameterization **p** and BASEPARAM as examples (this can also be repeated for **p'** and MDLPARAM). Using both **p** and the epidemiological model O_M, we can calculate the BASEPARAM's estimation of total infections $\text{BASEPARAM}_{\text{Tinf}}$ as the cumulative values of $D(\mathbf{p})$ from pandemic's beginning. $D_{\text{reported}}(\mathbf{p})$ can be directly used as the BASEPARAM's estimation of reported infections. For the cumulative reported rate $\text{BASEPARAM}_{\text{Rate}}$, we calculate it as the cumulative values of NYT-Rinf divided by $D(\mathbf{p})$. For the symptomatic rate, the SEIR+HD model [33] could estimate the number of symptomatic rate $\text{BASEPARAM}_{\text{Symp}}$ by dividing the number of infections in state I_S and I_M by the population number. However, SAPHIRE model [25] does not contain states that correspond to the symptomatic cases, so we cannot estimate the symptomatic rate using this model.

Acknowledgements

This paper was partially supported by the NSF (Expeditions CCF-1918770 and CCF-1918656, CAREER IIS-2028586, RAPID IIS-2027862, Medium IIS-1955883, Medium IIS-2106961, IIS-1931628, IIS-1955797, IIS-2027848), NIH 2R01GM109718, CDC MInD program U01CK000589, ORNL and funds/computing resources from Georgia Tech and GTRI. B. A. was in part supported by the CDC MInD-Healthcare U01CK000531-Supplement. A.V.'s work is also supported in part by grants from the UVA Global Infectious Diseases Institute (GIDI).

References

[1] 4q gdp: Economy expands at a 4.0% annualized rate. https://finance.yahoo.com/news/4q-gdp-2020-us-economy-coronavirus-pandemic-180133456.html.

[2] Commercial laboratory seroprevalence survey data. https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/commercial-lab-surveys.html.

[3] Coronavirus in the u.s.:latest map and case count. https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html.

[4] Covid-19 test price information, https://www.questdiagnostics.com/business-solutions/health-plans/covid-19/pricing.

[5] Hidden outbreaks spread through u.s. cities far earlier than americans knew, estimates say. https://www.nytimes.com/2020/04/23/us/coronavirus-early-outbreaks-cities.html.

[6] Interim economic projections for 2020 and 2021. https://www.cbo.gov/publication/56351.

[7] United states covid-19 cases, deaths, and laboratory testing (naats) by state, territory, and jurisdiction. https://covid.cdc.gov/covid-data-tracker/#cases.

[8] Us covid cases likely more than double official count, experts say. https://www.cidrap.umn.edu/news-perspective/2021/07/us-covid-cases-likely-more-double-official-count-experts-say, 2021.

[9] ACCORSI, E. K., QIU, X., RUMPLER, E., KENNEDY-SHAFFER, L., KAHN, R., JOSHI, K., GOLDSTEIN, E., STENSRUD, M. J., NIEHUS, R., CEVIK, M., ET AL. How to detect and reduce potential sources of biases in studies of sars-cov-2 and covid-19. *European Journal of Epidemiology* (2021), 1–18.

[10] ADHIKARI, B., RANGUDU, P., PRAKASH, B. A., AND VULLIKANTI, A. Near-optimal mapping of network states using probes. In *Proceedings of the 2018 SIAM International Conference on Data Mining* (2018), SIAM, pp. 108–116.

[11] AGUILAR, J. B., FAUST, J. S., WESTAFER, L. M., AND GUTIERREZ, J. B. Investigating the impact of asymptomatic carriers on covid-19 transmission. *MedRxiv* (2020).

[12] ANGELOPOULOS, A. N., PATHAK, R., VARMA, R., AND JORDAN, M. I. On identifying and mitigating bias in the estimation of the covid-19 case fatality rate. *arXiv preprint arXiv:2003.08592* (2020).

[13] BAI, Y., YAO, L., WEI, T., TIAN, F., JIN, D.-Y., CHEN, L., AND WANG, M. Presumed asymptomatic carrier transmission of covid-19. *JAMA 323*, 14 (2020), 1406–1407.

[14] BEDFORD, T., GRENINGER, A. L., ROYCHOUDHURY, P., STARITA, L. M., FAMULARE, M., HUANG, M.-L., NALLA, A., PEPPER, G., REINHARDT, A., XIE, H., ET AL. Cryptic transmission of sars-cov-2 in washington state. *Science 370*, 6516 (2020), 571–575.

[15] BI, Q., WU, Y., MEI, S., YE, C., ZOU, X., ZHANG, Z., LIU, X., WEI, L., TRUELOVE, S. A., ZHANG, T., ET AL. Epidemiology and transmission of covid-19 in 391 cases and 1286 of their close contacts in shenzhen, china: a retrospective cohort study. *The Lancet infectious diseases 20*, 8 (2020), 911–919.

[16] BUDHATHOKI, K., AND VREEKEN, J. Origo: causal inference by compression. *Knowledge and Information Systems 56*, 2 (2018), 285–307.

[17] CHANG, S., PIERSON, E., KOH, P. W., GERARDIN, J., REDBIRD, B., GRUSKY, D., AND LESKOVEC, J. Mobility network models of covid-19 explain inequities and inform reopening. *Nature 589*, 7840 (2021), 82–87.

[18] DONG, E., DU, H., AND GARDNER, L. An interactive web-based dashboard to track covid-19 in real time. *The Lancet infectious diseases 20*, 5 (2020), 533–534.

[19] EDELING, W., ARABNEJAD, H., SINCLAIR, R., SULEIMENOVA, D., GOPALAKRISHNAN, K., BOSAK, B., GROEN, D., MAHMOOD, I., CROMMELIN, D., AND COVENEY, P. V. The impact of uncertainty on predictions of the covidsim epidemiological code. *Nature Computational Science 1*, 2 (2021), 128–135.

[20] GAO, F., AND HAN, L. Implementing the nelder-mead simplex algorithm with adaptive parameters. *Computational Optimization and Applications 51*, 1 (2012), 259–277.

[21] GEERS, D., SHAMIER, M. C., BOGERS, S., DEN HARTOG, G., GOMMERS, L., NIEUWKOOP, N. N., SCHMITZ, K. S., RIJSBERGEN, L. C., VAN OSCH, J. A., DIJKHUIZEN, E., ET AL. Sars-cov-2 variants of concern partially escape humoral but not t-cell responses in covid-19 convalescent donors and vaccinees. *Science Immunology 6*, 59 (2021).

[22] GOLDSTEIN, J. R., AND LEE, R. D. Demographic perspectives on the mortality of covid-19 and other epidemics. *Proceedings of the National Academy of Sciences 117*, 36 (2020), 22035–22041.

[23] GOPALAKRISHNAN, V., PETHE, S., KEFAYATI, S., SRINIVASAN, R., HAKE, P., DESHPANDE, A., LIU, X., HOANG, E., DAVILA, M., BIANCO, S., ET AL. Globally local: Hyper-local modeling for accurate forecast of covid-19. *Epidemics 37* (2021), 100510.

[24] GRÜNWALD, P. D. *The minimum description length principle.* MIT press, 2007.

[25] HAO, X., CHENG, S., WU, D., WU, T., LIN, X., AND WANG, C. Reconstruction of the full transmission dynamics of covid-19 in wuhan. *Nature 584*, 7821 (2020), 420–424.

[26] HAVERS, F. P., REED, C., LIM, T., MONTGOMERY, J. M., KLENA, J. D., HALL, A. J., FRY, A. M., CANNON, D. L., CHIANG, C.-F., GIBBONS, A., ET AL. Seroprevalence of antibodies to sars-cov-2 in 10 sites in the united states, march 23-may 12, 2020. *JAMA internal medicine 180*, 12 (2020), 1576–1586.

[27] HO, F. K., PETERMANN-ROCHA, F., GRAY, S. R., JANI, B. D., KATIKIREDDI, S. V., NIEDZWIEDZ, C. L., FOSTER, H., HASTIE, C. E., MACKAY, D. F., GILL, J. M., ET AL. Is older age associated with covid-19 mortality in the absence of other risk factors? general population cohort study of 470,034 participants. *PloS one 15*, 11 (2020), e0241824.

[28] HOERTEL, N., BLACHIER, M., BLANCO, C., OLFSON, M., MASSETTI, M., RICO, M. S., LIMOSIN, F., AND LELEU, H. A stochastic agent-based model of the sars-cov-2 epidemic in france. *Nature medicine 26*, 9 (2020), 1417–1421.

[29] IONIDES, E. L., BRETÓ, C., AND KING, A. A. Inference for nonlinear dynamical systems. *Proceedings of the National Academy of Sciences 103*, 49 (2006), 18438–18443.

[30] IRONS, N. J., AND RAFTERY, A. E. Estimating sars-cov-2 infections from deaths, confirmed cases, tests, and random surveys. *Proceedings of the National Academy of Sciences 118*, 31 (2021).

[31] JAY, J., BOR, J., NSOESIE, E. O., LIPSON, S. K., JONES, D. K., GALEA, S., AND RAIFMAN, J. Neighbourhood income and physical distancing during the covid-19 pandemic in the united states. *Nature human behaviour 4*, 12 (2020), 1294–1302.

[32] JONES, J. M., STONE, M., SULAEMAN, H., FINK, R. V., DAVE, H., LEVY, M. E., DI GERMANIO, C., GREEN, V., NOTARI, E., SAA, P., ET AL. Estimated us infection-and vaccine-induced sars-cov-2 seroprevalence based on blood donations, july 2020-may 2021. *JAMA 326*, 14 (2021), 1400–1409.

[33] KAIN, M. P., CHILDS, M. L., BECKER, A. D., AND MORDECAI, E. A. Chopping the tail: How preventing superspreading can help to maintain covid-19 control. *Epidemics 34* (2021), 100430.

[34] KOUTRA, D., KANG, U., VREEKEN, J., AND FALOUTSOS, C. Vog: Summarizing and understanding large graphs. In *Proceedings of the 2014 SIAM international conference on data mining* (2014), SIAM, pp. 91–99.

[35] KRAEMER, M. U., SCARPINO, S. V., MARIVATE, V., GUTIERREZ, B., XU, B., LEE, G., HAWKINS, J. B., RIVERS, C., PIGOTT, D. M., KATZ, R., ET AL. Data curation during a pandemic and lessons learned from covid-19. *Nature Computational Science 1*, 1 (2021), 9–10.

[36] KRAEMER, M. U., YANG, C.-H., GUTIERREZ, B., WU, C.-H., KLEIN, B., PIGOTT, D. M., DU PLESSIS, L., FARIA, N. R., LI, R., HANAGE, W. P., ET AL. The effect of human mobility and control measures on the covid-19 epidemic in china. *Science 368*, 6490 (2020), 493–497.

[37] KUSTIN, T., HAREL, N., FINKEL, U., PERCHIK, S., HARARI, S., TAHOR, M., CASPI, I., LEVY, R., LESHCHINSKY, M., KEN DROR, S., ET AL. Evidence for increased breakthrough rates of sars-cov-2 variants of concern in bnt162b2-mrna-vaccinated individuals. *Nature medicine 27*, 8 (2021), 1379–1384.

[38] LAI, S., RUKTANONCHAI, N. W., ZHOU, L., PROSPER, O., LUO, W., FLOYD, J. R., WESOLOWSKI, A., SANTILLANA, M., ZHANG, C., DU, X., ET AL. Effect of non-pharmaceutical interventions to contain covid-19 in china. *Nature 585*, 7825 (2020), 410–413.

[39] LI, R., PEI, S., CHEN, B., SONG, Y., ZHANG, T., YANG, W., AND SHAMAN, J. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (sars-cov-2). *Science 368*, 6490 (2020), 489–493.

[40] LU, F. S., NGUYEN, A. T., LINK, N. B., MOLINA, M., DAVIS, J. T., CHINAZZI, M., XIONG, X., VESPIGNANI, A., LIPSITCH, M., AND SANTILLANA, M. Estimating the cumulative incidence of covid-19 in the united states using influenza surveillance, virologic testing, and mortality data: Four complementary approaches. *PLOS Computational Biology 17*, 6 (2021), e1008994.

[41] MOGHADAS, S. M., FITZPATRICK, M. C., SAH, P., PANDEY, A., SHOUKAT, A., SINGER, B. H., AND GALVANI, A. P. The implications of silent transmission for the control of covid-19 outbreaks. *Proceedings of the National Academy of Sciences 117*, 30 (2020), 17513–17515.

[42] NANDE, A., SHEEN, J., WALTERS, E. L., KLEIN, B., CHINAZZI, M., GHEORGHE, A. H., ADLAM, B., SHINNICK, J., TEJEDA, M. F., SCARPINO, S. V., ET AL. The effect of eviction moratoria on the transmission of sars-cov-2. *Nature communications 12*, 1 (2021), 1–13.

[43] PEI, S., KANDULA, S., AND SHAMAN, J. Differential effects of intervention timing on covid-19 spread in the united states. *Science advances 6*, 49 (2020), eabd6370.

[44] PEI, S., TENG, X., LEWIS, P., AND SHAMAN, J. Optimizing respiratory virus surveillance networks using uncertainty propagation. *Nature communications 12*, 1 (2021), 1–10.

[45] PELUSO, M. J., TAKAHASHI, S., HAKIM, J., KELLY, J. D., TORRES, L., IYER, N. S., TURCIOS, K., JANSON, O., MUNTER, S. E., THANH, C., ET AL. Sars-cov-2 antibody magnitude and detectability are driven by disease severity, timing, and assay. *Science advances 7*, 31 (2021), eabh3409.

[46] Planas, D., Veyer, D., Baidaliuk, A., Staropoli, I., Guivel-Benhassine, F., Rajah, M. M., Planchais, C., Porrot, F., Robillard, N., Puech, J., et al. Reduced sensitivity of sars-cov-2 variant delta to antibody neutralization. *Nature 596*, 7871 (2021), 276–280.

[47] Prakash, B. A., Vreeken, J., and Faloutsos, C. Spotting culprits in epidemics: How many and which ones? In *2012 IEEE 12th International Conference on Data Mining* (2012), IEEE, pp. 11–20.

[48] Press, W. H., and Levin, R. C. Modeling, post covid-19. *Science 370*, 6520 (2020), 1015–1015.

[49] Reinhart, A., Brooks, L., Jahja, M., Rumack, A., Tang, J., Agrawal, S., Al Saeed, W., Arnold, T., Basu, A., Bien, J., et al. An open repository of real-time covid-19 indicators. *Proceedings of the National Academy of Sciences 118*, 51 (2021).

[50] Russell, T. W., Golding, N., Hellewell, J., Abbott, S., Wright, L., Pearson, C. A., van Zandvoort, K., Jarvis, C. I., Gibbs, H., Liu, Y., et al. Reconstructing the early global dynamics of under-ascertained covid-19 cases and infections. *BMC medicine 18*, 1 (2020), 1–9.

[51] Salomon, J. A., Reinhart, A., Bilinski, A., Chua, E. J., La Motte-Kerr, W., Rönn, M. M., Reitsma, M. B., Morris, K. A., LaRocca, S., Farag, T. H., et al. The us covid-19 trends and impact survey: Continuous real-time measurement of covid-19 symptoms, risks, protective behaviors, testing, and vaccination. *Proceedings of the National Academy of Sciences 118*, 51 (2021).

[52] Sethuraman, N., Jeremiah, S. S., and Ryo, A. Interpreting diagnostic tests for sars-cov-2. *JAMA 323*, 22 (2020), 2249–2251.

[53] Shaman, J. An estimation of undetected covid cases in france. *Nature 590* (2020), 38–39.

[54] Sood, N., Simon, P., Ebner, P., Eichner, D., Reynolds, J., Bendavid, E., and Bhattacharya, J. Seroprevalence of sars-cov-2–specific antibodies among adults in los angeles county, california, on april 10-11, 2020. *JAMA 323*, 23 (2020), 2425–2427.

[55] Stockmaier, S., Stroeymeyt, N., Shattuck, E. C., Hawley, D. M., Meyers, L. A., and Bolnick, D. I. Infectious diseases and social distancing in nature. *Science 371*, 6533 (2021).

[56] Subramanian, R., He, Q., and Pascual, M. Quantifying asymptomatic infection and transmission of covid-19 in new york city using observed cases, serology, and testing capacity. *Proceedings of the National Academy of Sciences 118*, 9 (2021).

[57] Tian, Y., Sridhar, A., Yağan, O., and Poor, H. V. Analysis of the impact of mask-wearing in viral spread: Implications for covid-19. In *2021 American Control Conference (ACC)* (2021), IEEE, pp. 3132–3137.

[58] Tiwari, S., Vyasarayani, C., and Chatterjee, A. Data suggest covid-19 affected numbers greatly exceeded detected numbers, in four european countries, as per a delayed seiqr model. *Scientific reports 11*, 1 (2021), 1–12.

[59] TOTH, D. J., TANNER, W. D., KHADER, K., AND GUNDLAPALLI, A. V. Estimates of the risk of large or long-lasting outbreaks of middle east respiratory syndrome after importations outside the arabian peninsula. *Epidemics 16* (2016), 27–32.

[60] WELLS, C. R., SAH, P., MOGHADAS, S. M., PANDEY, A., SHOUKAT, A., WANG, Y., WANG, Z., MEYERS, L. A., SINGER, B. H., AND GALVANI, A. P. Impact of international travel and border control measures on the global spread of the novel 2019 coronavirus outbreak. *Proceedings of the National Academy of Sciences 117*, 13 (2020), 7504–7509.

[61] WILDER, B., CHARPIGNON, M., KILLIAN, J. A., OU, H.-C., MATE, A., JABBARI, S., PERRAULT, A., DESAI, A. N., TAMBE, M., AND MAJUMDER, M. S. Modeling between-population variation in covid-19 dynamics in hubei, lombardy, and new york city. *Proceedings of the National Academy of Sciences 117*, 41 (2020), 25904–25910.

[62] ZHANG, W., GOVINDAVARI, J. P., DAVIS, B. D., CHEN, S. S., KIM, J. T., SONG, J., LOPATEGUI, J., PLUMMER, J. T., AND VAIL, E. Analysis of genomic characteristics and transmission routes of patients with confirmed sars-cov-2 in southern california during the early stage of the us covid-19 pandemic. *JAMA network open 3*, 10 (2020), e2024191.

[63] ZHAO, J., YUAN, Q., WANG, H., LIU, W., LIAO, X., SU, Y., WANG, X., YUAN, J., LI, T., LI, J., ET AL. Antibody responses to sars-cov-2 in patients with novel coronavirus disease 2019. *Clinical infectious diseases 71*, 16 (2020), 2027–2034.

(a) Baseline Parameterization

$S \rightarrow E \rightarrow I_r \rightarrow R$
$\quad\quad\quad I_u \rightarrow R$

(b) MDL Parameterization Selection

Sender S

Candidate Parameterizations

DATA
MODEL
$S \rightarrow E \rightarrow I_r \rightarrow R$
$\quad\quad I_u$

DATA

DATA

Best Parameterization

DATA
MODEL
MDL Cost
$S \rightarrow E \rightarrow I_r \rightarrow R$
$\quad\quad\quad I_u$

Receiver R

(c) Iceberg Plot (example)

Reported infections

MdlParam unreported infections

BaseParam unreported infections

Minneapolis had only 16 reported COVID-19 infections by Mar 11, 2020. BasePuram estimated 182 total infections. However, MdlParam estimates that there were actually 301 total infections.

(a)

(b)

(c)