Predictive Monitoring of COVID-19

Jianxi Luo
Data-Driven Innovation Lab
Singapore University of Technology and Design

Working draft, updated on May 11, 2020

Abstract

During the current COVID-19 pandemic, there have been many efforts to forecast the infection cases, deaths and medical or economic indicators, with a variety of statistical or epidemiological models. Some of the forecasting projects have influenced the policies in some countries. However, the prediction of future is uncertain by nature, given the fundamental nature of the COVID-19 pandemic as a "wicked problem". The uncertainty is rooted in the many unknown unknowns about the contagious virus itself and the complex, heterogenous and dynamic human behaviours, government interventions and testing scenarios. The extreme uncertainty of this context makes the intent for prediction accuracy misleading. Herein, we do not aim to make "accurate" predictions about the future or to evaluate how accurate a prediction or a prediction model is. Instead, to address the uncertainty in dynamic real-world scenarios for which we do not have complete information and understanding, we explore the potentials of "predictive monitoring" with the aim to capture and make sense of the changes in theoretical predictions for meaningful signals of the uncertainty and changes in the real-world scenarios. Such signals from predictive monitoring are expected to make the planning, behaviours and mentality at the present time be more "future-informed" and possibly initiate and guide pre-cautionary actions now to shape the real future.

Introduction

Since the outbreak of COVID-19 in January 2020, researchers around the world have adopted classic or latest data science and AI techniques and applied them to the data available to predict the developments and trends of COVID-19 in different countries or regions. The noticeable efforts include the publicly available and continually updated forecasts by the Institute of Health Metrics and Evaluation (IHME) at University of Washington [1] and the MRC Centre for Global Infectious Disease Analysis at the Imperial College London [2], among others. Table 1 presents a list of publicly accessible COVID-19 forecasting efforts around the world using a variety of statistical or epidemic process models. Some forecasts focus on future deaths and hospitality needs [3,4,5] and infection cases and peaks [6,7,8], while others focus on the impact of social distancing, travel restrictions, and mitigation and suppression strategies [7,11,12].

Table 1. Public COVID-19 forecasting initiatives around the world, as of May 11, 2020

Organization	URL	Methods
Imperial College London	https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/	Mechanistic transmission models
University of Geneva, ETH Zürich & EPFL	https://renkulab.shinyapps.io/COVID-19-Epidemic-Forecasting/	Statistical models
Massachusetts Institute of Technology	https://www.covidanalytics.io/projections	Modified SEIR model
Los Alamos National Laboratories	https://covid-19.bsvgateway.org/	Statistical dynamical growth model
The University of Washington, Seattle	https://covid19.healthdata.org/projections	Statistical model
The University of Texas, Austin	https://covid-19.tacc.utexas.edu/projections/	Statistical model
Northeastern University	https://covid19.gleamproject.org/	Spatial epidemic model
University of California, Los Angeles	https://covid19.uclaml.org/	Modified SEIR model

Some published studies have attempted to validate the accuracy of specific prediction methods [3,4,5,6], some of which used COVID-19 data from China. However, even the most cited forecasting method from the IHME has been found with model design issues [5,13] and that for 70 percent of time the actual death numbers fell outside its next-day predictions' 95 percent confidence interval [14]. The IHME team later revised the model [4] but the prediction errors remain high. In any case, researchers are learning and improving the methods and tools on the go in order to make more and more accurate predictions on the next developments of the COVID-19 pandemic [13]. Despite the intrinsic uncertaint nature of COVID-19 predictions, some efforts have already influenced policies or informed policy makers to some extents and in certain ways [14][15].

Given the value of predictions but also the difficulty to do it well under extreme uncertainty and for such a wicked problem as the COVID-10 pandemic, we aim to explore the values and potentials of predictive monitoring to deal with the uncertainty of predictions and make use and make sense of prediction excises for only suitable goods. Predictive monitoring means the continual monitoring of the predictions of crucial future events, such as the bending and ending of the pandemic life cycle curve, together with the actual history data to date. In predictive monitoring, the fundamental assumption is that real-world scenarios are changing, so predictions are expected to change over time. Changes in predictions are not viewed as errors or inaccuracy, but valuable signals about the changes in the present real-world scenarios.

Therefore, predictive monitoring differs fundamentally from the traditional and common prediction or forecasting practices (some examples in Table 1) that attempt to make a prediction now that can be accurate about the real future. Such prediction practices might be more reasonable when the context is less uncertain. Such predictions subconsciously take the future as fixed, while we assume future is not fixed, a result of the happenings, changes and interventions from now to them, and uncertain. The traditional COVID-19 prediction practices are more like weather forecasts where the future weather is extrinsic to us and cannot be changed by us. However, the evolution of the pandemic is also affected by the evolving human behaviours and government interventions, etc. In addition, predictive monitoring also differs from the common monitoring practice that reports actual past cases of infection, recovery and death, which may stimulate reactive and responsive actions. By contrast, predictive monitoring may inform, initiate and guide more future-informed planning, policies and actions to shape the real future. Table 2 presents a taxonomy that explicates the differences of predictive monitoring from traditional prediction and monitoring.

Table 2. The taxonomy for predictive monitoring, traditional prediction, and monitoring

		What Value Does It Deliver?	
		Future-Informed	Past-Informed
When Is It Suitable?	High Environment Uncertainty	*Predictive Monitoring*	*Monitoring*
	Low Environment Uncertainty	*Prediction*	

The Predictive Monitoring Experiment

- Theory

We experimented predictive monitoring in the realistic context of the on-going COVID-19, in order to explore its potentials and develop specific guidelines and strategies for the right use of it. To run the experiment, the first is to choose a prediction model and data source, before we can update and monitor the predictions with daily new data coming in over time. The propagation of infectious diseases often follows a life cycle pattern, from the outbreak to the acceleration phase, inflection point, deacceleration phase and eventual ending. Such a life cycle is the result of the infection process, property of the virus, the nature of a population and the adaptive and countering behaviours of agents including individuals (avoiding physical contact) and governments (locking down cities) in the population.

However, the pandemic life cycles vary by countries (or regional populations), and different countries might be in different phases of the life cycles at a same point in time.

For instance, on April 21, in Singapore, Prime Minister Hsien-Loong Lee announced the extension of circuit breaker to June 1 in response to the spikes of COVID-19 cases in Singapore, on the same day when Prime Minister Giuseppe Conte announced Italy’s plan to reopen businesses in Italy from May 4. Ideally speaking, such decisions and planning can be rationalized by well knowing where our own country (together with other countries and the world as a whole) is in its own pandemic life cycle, when the turning point is coming if it has yet come, and when the pandemic will end. Adjustments may be made according to the changes in the estimations and predictions on these fronts. The basis for such actionable estimation is the pandemic’s overall life cycle.

- Model

The pandemic life cycle pattern is expected to appear as a S-shape curve when one plots the accumulative count of infection cases over time or equivalently as a “bell-shape” curve of the daily counts over time (see examples in Figure 1). Note that the bell here is not expected to be symmetrical with no expectation of a normal distribution, but a long tail to the right. Such patterns as well as the underlying dynamics have been well studied in various domains including population growth, diffusion of new technologies in the society and infectious diseases, and have theoretically established mathematical models, such as the logistic model that describes a general life cycle phenomenon and the SIR (susceptible-infected-recovered) model [17,18,19] that describes the dynamic epidemic process of the spread of infectious diseases.

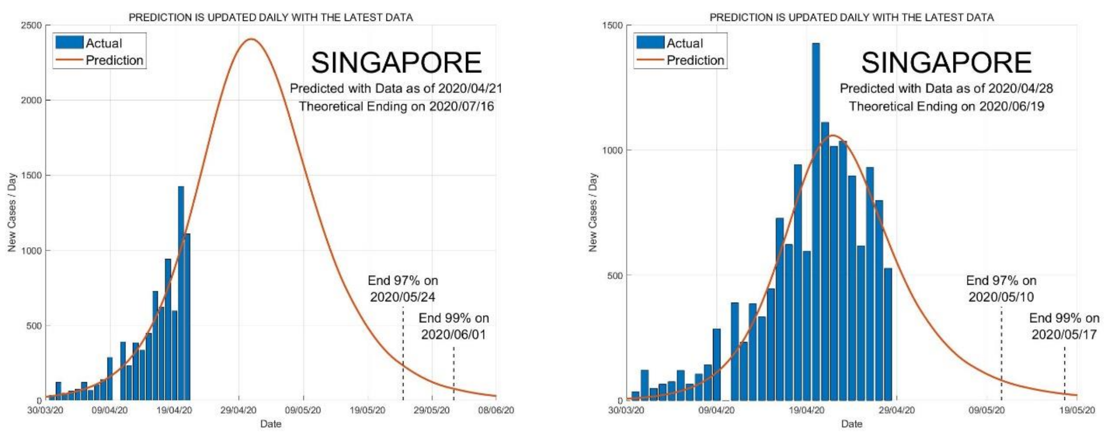

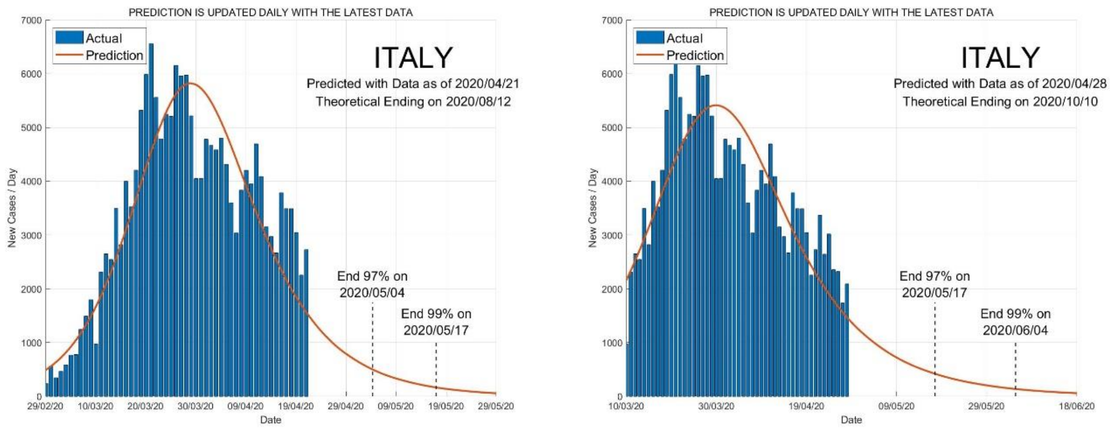

Figure 1. Continuous Data-Driven Estimations of COVID-19 Life Cycle, Turning and Ending Dates for Singapore and Italy as of April 21 versus April 28, 2020

The SIR model is employed in this experiment for a few reasons. One, it is context-specific and models the dynamic process of inflections in a population over time. Second, it requires simple data inputs that are publicly available. Third, there are open source computer codes available for quick adoption. Here we will not repeat the details of the SIR model in this paper, which can be easily found in many mathematics textbooks. Essentially, the SIR model use three ordinary differential equations to describe the dynamic flows of people between three compartments: S for the number of susceptible people, I for the number of infectious people, and R for the number of removed people (either recovered, died or immured) in the population. The SIR model incorporates two main parameters, beta and gamma. Gamma is the number of days one is contagious and a property of the virus. Beta is the average number of people infected by a previously infected person and is related to not only the interaction patterns of people in the society (which social distancing can influence) but also the infection process property of the virus [17,18,19].

- Implementation

The values of these two parameters determine the shape of an infectious disease's specific life cycle curve for a population. In particular, the model, which a system of three differential equations for S, I and R in its original form, can be reduced to one function about the total infection count, or equivalently the daily new infection counts. This key variable is the sum of I and R and has publicly available data reported by official channels every day. Please refer to the paper [20] by Milan Batista for the model reduction. Therefore, only the data of the accumulative infection cases over time (which can be also used to derive the daily new cases) is required to regress the key parameters and other constants and thus train a model to derive the overall pandemic life cycle curve.

Batista also developed open-source computer codes to implement the regression using the reduced function [21]. In our experiments, we applied the codes of Batisa to the COVID-19 accumulative infection data for each country from "Our World in Data" [22] to regress the parameters and constants of the basic SIR model. Note that, more sophisticated derivative versions of the SIR model with more compartments, such as the SEIR model, have also been used in COVID-19 forecasting (such as [6] and several listed in Table 1), but additional increased equations and parameters also require more sophisticated data inputs which we do not have. Regressions are run for individual countries and updated daily with the newest accumulative and daily infection count data becoming available daily. Not the data for all countries can produce statistically meaningful regression results. Only the countries with satisfactory goodness-of-fit between model and data as measured by R^2 greater than 0.8 are accepted, analysed and reported. For these countries, the regressed model for each of them is used to estimate the full pandemic life cycle and plot the life cycle curve.

Makes Sense of Prediction Changes

As shown in the examples in Figure 1, the initial segment of the curve is fitted with the data to date and the remaining segment of the curve is "predicted". With the estimated full life cycle curve, one can easily observe which phase of the pandemic life cycle a specific country is in (with actual data plotted together), when the inflection point (the peak in the bell-shape curve) is coming (for the interests of the countries still in the accelerating phase), and when the pandemic will end (for the interests of all countries). Therefore, our predictive monitoring is focused on such high-level transitioning characteristics of the pandemic's total life cycle [7], instead of the specific numbers of accumulative or daily cases on a specific day, which the traditional forecasting efforts ([3,4,5] and others in Table 1) try to forecast with confidence.

The inflection point of the pandemic life cycle curve is specific and appears as the peak in the bell-shape curve. However, estimating the "ending date" is not straight-forward and may be done differently for different considerations. Most theoretically, one can define the pandemic's end date as the day with the last infection case of the pandemic, and thus operationalize the estimation of the end date as the day with the last predicted infection at the right most end of the estimated pandemic life cycle curve. However, practically, estimation of the theoretical ending might not be useful to provide guidance for the planning of activities of governments, companies and individuals. One might consider an early date when predominately most predicted infections have been actualized and only a small portion of the total predicted epidemic population is left (e.g., the case of Australia as of now). The total predicted epidemic population size is the area under the entire curve. In our experiments, we monitored three alternative estimates of end dates in the order of conservativeness.

- The date to reach the last expected case;

- The date to reach 99% of the total expected cases;
- The date to reach 97% of the total expected cases.

In any case, specifying an end date is arbitrary in nature. For flexibility, one may simply exploit the estimated life cycle curve, especially its right most tail segment, to screen and sense when the pandemic gradually vanishes to which extent.

It is noteworthy that the bell-shape curve (of daily cases, instead of the S-shape of accumulative cases) is chosen to visualize the life cycle because it allows easy detection of the inflection point as the peak of the curve to distinguish countries in acceleration and deacceleration phases. For instance, Figure 1A visually reveals on April 21 Singapore was still in its acceleration phase, whereas Figure 1C shows Italy has passed its inflection point. At the time, the estimated "future" turning date (i.e., the inflection point of the curve) for Singapore would be May 1. However, as shown in Figure 1B, on April 28, Singapore has already past its inflection point based on the updated curve with newer data from April 21 to 28, earlier than the turning date predicted on 21 April (in Figure 1A). In contrast, from April 21 (Figure 1C) to April 28 (Figure 1D), the curves of Italy are slightly lifted, and the later predictions for Italy suggest consistently later 97%, 99% and 100% ending dates.

These changes are discovered through <u>predictive monitoring</u> of the actual developments and estimations together holistically. We continually monitor the predictions, not really hoping the previous predictions to be tested true or accurate later when the real "future" comes, but for detecting in the changes of the predictions over time. From a traditional perspective, the difference between a future prediction and a previous one on the same variable would be considered a bad thing and a proof of failure of the prediction model [13,14]. Instead, here we tend to make sense of such changes from the earlier to later predictions for meaningful signals as to what are happening in the dynamically changing real-world scenarios, based on the fundamental understanding that predictions made over time should be different when the real-world scenarios are changing.

For example, the changed predictions of the theoretical pandemic end dates of Singapore over time may reveal the effects of the recently strengthened measures of the Singapore government and more cautions of the local citizen from PM Lee's announcement of circuit breaker extension on 21 April. The changed predictions of pandemic end dates for Italy may result from the slightly relaxed government control measures and human behaviours in Italy in the past week. The pandemic curves of Singapore and Italy have shifted over time, as the real-world scenarios have dynamically changed. It would be wrong to expect the curve estimated with data from the previous scenario to represent the curve for a later scenario. Instead, the curves should be continually re-estimated with the latest data, the predictions based on these curves should be continually monitored, and the changes in the predictions may reveal changes in real-world scenarios over time. Monitoring and detecting such changes in the predictions provides the main value of predictive monitoring.

In other words, our default expectation in predictive monitoring is that predictions will change, especially when the real-world scenarios, such as government policies, testing protocols and human behaviours, are also rapidly changing. In such cases, we should not expect the model trained with data as of today to be true for a different scenario later. When considering the dynamics of human behaviours and government policies and other real-world scenarios that the mathematical model and training data cannot accurate represent, predictive monitoring would be a more valuable exercise, rather than making a prediction now to see if it is a hit or miss in future.

The changes in the predicted theoretical events, such as the theoretical ending dates, may also allow us to sense or measure the uncertainty rooted from the real-world scenarios on the ground. Therefore, we also report the standard deviations in N latest and connectively predicted theoretical end dates as an indicator of uncertainty. Such a measure is often called "volatility" in finance when used to evaluate the uncertainty associated with stock prices. If the standard deviation of the connective predicted ending dates is small (regardless of their accuracy), it indicates the real-world scenarios are not changing. If it is high, it might imply changes are happening in the real-world scenarios. Figure 2 reports the past 5-day volatilities of the estimates of the theoretical ending dates of some major countries with model-data (of daily new cases) fits R^2 > 0.8.

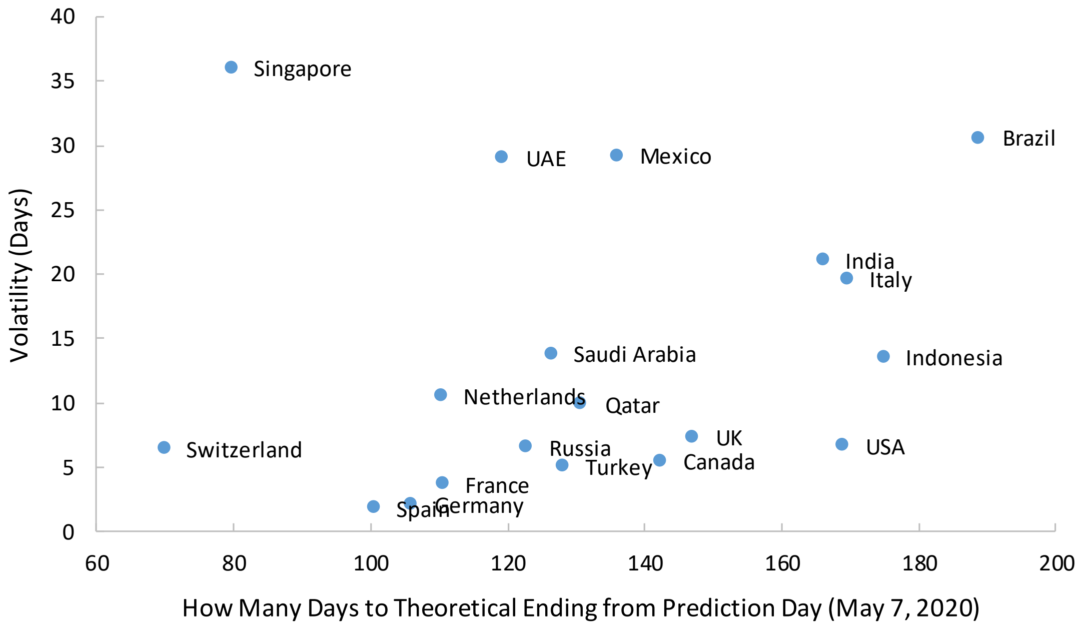

Figure 2. Prediction Volatility against Length of Time to Predicted Ending Dates

Some findings in Figure 2 are noteworthy. There appears a general correlation between the ending and volatility across countries, and also exist outliers, such as Singapore with a relatively close ending date but extremely high volatility, indicting uncertainty on the ground recently, which might be related to the testing protocols because the cases in Singapore are contained in the dormitories of foreign workers. Also, Brazil stands out with a very far ending date and high volatility, indicating an undesirable and uncertain real-world scenario on the ground in Brazil now and demand caution and pre-cautionary actions. USA has a quite long time to its theoretical ending but appears quite stable in predictions. Switzerland presents a desirable case with the closest ending date and also low volatility.

By contrast, we purposefully avoid such metrics as “margin of errors” and “confidence intervals”, because our assumption is the pandemic’s real-world scenarios are uncertain and evolving by nature and thus there is no target value to define an “error”. In contrast, the confidence intervals have been called quantified uncertainty bands in some of the COVID-19 forecasts [3,5]. Such calling might be theoretically questionable. In complexity science, uncertainty is defined as unknown unknown and not quantifiable. It is the case of the COVID-19. Only the known unknown, which is often called” risk”, is quantifiable. Most of the forecasting uncertainty measures are in fact risk measures. However, it is also a question whether risk measurement is meaningful in the uncertain case of COVID-19. Risk measurement is suitable when there is a fixed target, but it is just that we are unsure to what extent we can hit the target.

In sum, the foregoing examples and elaboration are aimed to explicate the importance of predictive monitoring or continually monitoring predictions to address uncertainty, detect and evaluate changes (such as human behavers and government control measures) made in the dynamic real-world scenarios in real time. It also allows the estimation of the volatility of the predictions as an indicator of the uncertainty of the underlying real-world scenarios. Thus, predictive monitoring differs from making a one-shot prediction for it to become true in the future and differs from the monitoring of actual cases every day (see Table 2 again).

Broader Discussion

Predictive monitoring for each country should be read and interpreted together with what are happening in the real world and government policy changes. For instance, Singapore government’s strengthened restrictions in April may have bended its curve earlier than the previously predicted ones, and the early relaxation of social distancing and lockdown in Italy and Germany might increase infection rates and thus delay the pandemic ending as predicted now. Also, the predictive monitoring of a country should not be read in isolation, but together with the predictions and real time situations of other countries. No country is in isolation in the world today. The monitoring and control of one country must be coupled with the monitoring and control of other countries.

For example, while the predictive monitoring shows the pandemic has “theoretically” ended in China, South Korea and Australia (despite a small number of domestic cases reported daily), it also shows many other countries (such as Brazil, USA) and the world as a whole will still suffer till the end of 2020 if we remain in our present trajectories of government policies and individual behaviours and without medical cures and vaccines for COVID-19. Therefore, the governments of China, South Korea and Australia may not want to open their international ports so soon and lift the domestic restrictions so quickly, until the pandemic nears its end in the world as whole. Although it is the time for all of us to isolate and distance physically from each other, it is also the time that needs more sharing of data, information and knowledge and more close coordination.

For countries that are still early in their own pandemic life cycles (such as Brazil still in the acceleration phase as of May 8 based on our predictive monitoring), the prediction of the rest of the curve, inflection point and ending dates will be more teasing, but also inherent less relevant to the “real future” to come given that the actual data only cover a smaller and early portion of the total life cycle and many real-world scenarios that the model cannot describe are expected to change. By contrast, for countries that have passed their inflection points and been approaching ending phases, prediction is less useful. When uncertainty is low, it is more likely that we can derive and approve a highly predictive model. However, in such cases, the trained model is more about explaining the history and less about predicting the future. For those countries, uncertainty still exists, for example, a new epidemic wave might come if the governments and individuals lift controls and disciplines too early, especially when the pandemic is still prevalent in other countries.

Summary

Predictive monitoring may completement the traditional historical case monitoring practice and the traditional accuracy-oriented prediction practices to deliver complementary values. The value of continuous predictive monitoring might be greater when the real-world scenarios that the models cannot describe are inherently dynamic and more uncertain. We will continually monitor the estimated pandemic life cycle curves and end dates and explore valuable insights from the monitored prediction changes, as an experiment to explore the potentials of as well as develop guidelines and strategies for valuable predictive monitoring practices.

In the meantime, readers must take any prediction, regardless of the model and data, with caution. Over-optimism based on some predicted end dates is dangerous because it may loosen our disciplines and controls and cause the turnaround of the virus and infection. Although prediction based on science and data is aimed to be objective, it is uncertain by nature. One thing that is certain is that the model, data and prediction are inaccurate and insufficient to fully represent the complex, evolving, and heterogeneous realities of our world. The model we use in the experiment is only theoretically suitable for one stage or wave of the epidemic evolution, and relatively more meaningful when applied to data for each single stage if the country has experienced multiple stages (such as Singapore). The prediction is also conditioned by the quality of the data. The data publicly available today is based on tests, which are done differently in different countries and over time periods. They do not necessarily represent the total infection account which is the theoretical input of the model. One should expect changes in the continually monitored predictions, instead of fixed expectations.

Future is always uncertain. We must keep this in mind when doing and reading any prediction. No one predicted the COVID-19 outbreak beforehand. With acknowledging the uncertain nature of the ongoing COVID-19 pandemic and our growing inter-connected and complex world, what are eventually and fundamentally needed are the flexibility, robustness and resilience of people, organizations and governments, as well as sharing and coordination, to deal with unpredictable and unwanted future events.

References

[1] Institute of Health Metrics and Evaluation (IHME) at University of Washington https://covid19.healthdata.org/united-states-of-america

[2] MRC Centre for Global Infectious Disease Analysis at the Imperial College. https://mrc-ide.github.io/covid19-short-term-forecasts/index.html

[3] IHME, COVID-19 Health Service Utilization Forecasting Team and Christopher JL Murray. Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months, 2020. Posted March 30, 2020. doi: https://doi.org/10.1101/2020.03.27.20043752

[4] IHME COVID-19 health service utilization forecasting team, Christopher JL Murray. Forecasting the impact of the first wave of the COVID-19 pandemic on hospital demand and deaths for the USA and European Economic Area countries. MedRxiv 2020.04.21.20074732; https://doi.org/10.1101/2020.04.21.20074732

[5] Woody S. et al. Projections for first-wave COVID-19 deaths across the US using social-distancing measures derived from mobile phones. MedRxiv 2020.04.16.20068163; https://doi.org/10.1101/2020.04.16.20068163

[6] Yang Z. et al. Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. J Thorac Dis 2020, 12(3): 165-174. doi: 10.21037/jtd.2020.02.64

[7] Ferguson NM, et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College London; 16 March 2020. Accessed at www.imperial.ac.uk/media/imperial-college/medicine/mrc-gida/2020-03-16-COVID19-Report-9.pdf on 31 March 2020.

[8] Petropoulos F, Makridakis S (2020) Forecasting the novel coronavirus COVID-19. PLoS ONE, 15(3): e0231236. https://doi.org/10.1371/ journal.pone.0231236

[11] Kissler SM. Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. Science, 14 April 2020 DOI: 10.1126/science.abb5793

[12] Matteo Chinazzi, et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. Science, 24 April 2020: 395-400

[13] Jewell NP, et al. Caution Warranted: Using the Institute for Health Metrics and Evaluation Model for Predicting the Course of the COVID-19 Pandemic. Ann Intern Med. 14 April 2020]. doi: https://doi.org/10.7326/M20-1565

[14] Roman Marchant, Noelle I. Samia, Ori Rosen, Martin A. Tanner, and Sally Cripps. Learning as we go: An examination of the statistical accuracy of covid19 daily death count predictions, 2020. doi: https://doi.org/10.1101/2020.04.11.20062257

[15] Grey S, MacAskill A. Special Report: Johnson listened to his scientists about coronavirus - but they were slow to sound the alarm. Reuters, April 7, 2020. https://reut.rs/2Rk3sa7

[16] Resnick B. The White House projects 100,000 to 200,000 Covid-19 deaths. Vox, March 31, 2020. https://www.vox.com/science-and-health/2020/3/31/21202188/us-deaths-coronavirus-trump-white-house-presser-modeling-100000

[17] Hethcote H.W. The mathematics of infectious diseases, SIAM Review 42, 599–653 (2000).

[18] Smith D, and Moore L. The SIR Model for Spread of Disease - The Differential Equation Model. Convergence (December 2004)

[19] Newman, MEJ. The structure and function of complex networks. SIAM Review, 45, 167-256 (2003)

[20] Milan Batista (2020) Estimation of the final size of the COVID-19 epidemic. Medarxiv https://www.medrxiv.org/content/10.1101/2020.02.16.20023606v5. doi: https://doi.org/10.1101/2020.02.16.20023606

[21] Milan Batista (2020). fitVirusCOVID19 (https://www.mathworks.com/matlabcentral/fileexchange/74658-fitviruscovid19), MATLAB Central File Exchange. Retrieved May 1, 2020.

[22] Our World in Data https://ourworldindata.org/coronavirus-source-data