

Modeling, state estimation, and optimal control for the US COVID-19 outbreak

Calvin Tsay,^{1†} Fernando Lejarza,^{1†} Mark A. Stadtherr¹, Michael Baldea^{1,2*}

¹McKetta Department of Chemical Engineering, The University of Texas at Austin,
Austin, TX, USA

²Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin
Austin, TX, USA

[†]These authors contributed equally to this work

*To whom correspondence should be addressed; E-mail: mbaldea@che.utexas.edu

Abstract

The novel coronavirus SARS-CoV-2 and resulting COVID-19 disease have had an unprecedented spread and continue to cause an increasing number of fatalities worldwide. While vaccines are still under development, social distancing, extensive testing, and quarantine of confirmed infected subjects remain the most effective measures to contain the pandemic. These measures carry a significant socioeconomic cost. In this work, we introduce a novel optimization-based decision-making framework for managing the COVID-19 outbreak in the US. This includes modeling the dynamics of affected populations, estimating the model parameters and hidden states from data, and an optimal control strategy for sequencing social distancing and testing events such that the number of infections is minimized. The analysis of our extensive computational efforts reveals that social distancing and quarantining are most effective when implemented early, with quarantining of confirmed infected subjects having a much higher impact. Further, we find that “on-off” policies alternating between strict social distancing and relaxing such restrictions can be effective at “flattening” the curve while likely minimizing social and economic cost.

Introduction

Since its first reported case in early December 2019 in Wuhan, Hubei Province, China, the novel Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) and the resulting COVID-19

disease have reached 184 countries/regions causing a total of 115,225 deaths, as of this writing. The rate of spread of the virus is substantially higher than that of similar previously reported epidemics such as SARS and MERS [1]. Because of this world-wide and rapid spread, on March 11, 2020, the World Health Organization (WHO) declared the COVID-19 outbreak a global pandemic [2]. The exponential increase in the number of confirmed infectious cases and deaths has caused countries all around the world to respond with severe lock-down, quarantining, and social distancing measures to contain the spread of the disease. For example, northern Italy went into a state of emergency on March 8, imposing a complete lock-down that was expanded to the rest of the country three days later [3]. Following Italy's attempts to decrease the contagion rate, Spain enforced a nationwide lock-down on March 16 [4]. In contrast, the response in the US was significantly slower than those of the aforementioned (and most other European) countries, delaying widespread social distancing measures and disease screening initiatives.

Given the exponential growth of the virus, policy makers face the urgent challenge of determining the appropriate responses. For example, restrictions on population mobility, increased resource investments (e.g., personal protective equipment, ventilators, hospital beds in the ICU), as well as improved COVID-19 screening can greatly impact the spread of infection [5, 6]. Here, fundamental epidemiological models [7], typically comprising sets of coupled, nonlinear ordinary differential equations (ODEs), are valuable tools for simulating the dynamics of the epidemic and investigating suppression strategies. Several works (e.g., Anastassopoulou et al. [8], Peng et al. [9], Magal and Webb [10]) have addressed the development and fitting of such models for the COVID-19 outbreak using regression techniques. Models were subsequently used to characterize the infection [11, 12], predict its potential spread [13, 14, 15], and/or evaluate mitigation strategies [16, 17].

On the other hand, fewer studies have investigated the dynamic optimization, or *optimal control*, of epidemiological models in relation to the current COVID-19 outbreak. In this context, the aforementioned epidemiological models are used to determine the decisions, policies or inputs, such as the social distancing measures and testing rates, that lead to the best outcomes at the population level. Such outcomes include, for example, minimizing the peak number of infected people or total number of deaths, while also accounting for constrained resource availability and the extent to which social distancing is feasible. Among works of this type, Djidjou-Demasse et al. [18] investigated the optimal control of a single "intervention" input variable for an ODE model. Moore and Okyere [19] formulated a similar problem with additional inputs, including hospitalization rates and environmental spraying. Both studies solved the optimal control problem using an iterative forward-backward sweep method [20] and assumed the model parameters to be fully known. While optimal control of epidemiological models has been well-studied (e.g., Biswas et al. [21], Neilan and Lenhart [22]), a limiting factor for implementing modeling and optimization concepts during early stages of an outbreak is that obtaining accurate estimates of key model parameters can be challenging.

In this work, we report a novel and complete dynamic optimization-based approach to the entire epidemiological modeling and outbreak control workflow for the US COVID-19 outbreak. We first formulate a dynamic optimization strategy for identifying both the values of


time-invariant parameters and the historical trajectories of time-varying parameters (i.e., inputs) of an epidemiological model. We then investigate how optimal control of the inputs to the same model, which reflect social distancing and testing, relates to infection mitigation strategies. The optimal control problem is cast as a *simultaneous* dynamic optimization problem (i.e., reformulated as an algebraic problem) [23], enabling the natural use of deterministic global optimization technology [24]. This in turn provides solutions that are proven to be globally optimal, whereas iterative schemes such as the above [20] often only guarantee global optimality under certain conditions (the term “global” here refers to the notion that the best possible disease control policy is identified from a set of “local” solutions that otherwise satisfy optimality conditions. “Global” in this context should not be interpreted as “world-wide”). Furthermore, we show how state estimation should be used to update important hidden model states (e.g., the number unconfirmed infections) to reduce the impact of model inaccuracy. Implementations of the parameter estimation and optimal control problems in open-source software are provided freely at <https://github.com/Baldea-Group/covid-19>.

Results

Mathematical Modeling

We consider the compartmental model shown in Fig 1, which brings a few key modifications to the conventional SEIR (susceptible-exposed-infectious-recovered) structure. Importantly, the $a(t)$ state is added to account for infected subjects that are not included in the count of confirmed cases, e.g., because they are asymptomatic or there is insufficient testing. The virus has been found to be asymptomatic in up to 20–40% of cases [25, 26] and may be transmitted by asymptomatic carriers [27]. We also include a $p(t)$ state, which tracks the population that perishes due to the virus.

The modified SEAIR model has six compartments/states: $s(t)$ represents the number of subjects that are susceptible to infection, $e(t)$ the number that have been exposed to the virus, $a(t)$ the number that are infected but asymptomatic/unconfirmed, $i(t)$ the number with confirmed infections, and $r(t)$ the number that have recovered from infection. The states are assumed to sum to a known total population N , and $p(t)$ can be calculated algebraically at all times. The model structure is shown in Fig 1.



Input/parameter	Range	Description
$\alpha_a(t)$ [days ⁻¹]	[0.05, 0.5]	social distancing
$\alpha_i(t)$ [days ⁻¹]	[0.01, 0.3]	quarantining
$\kappa(t)$ [days ⁻¹]	[0.10, 0.3]	rate of testing
β [days ⁻¹]	[0.001, 0.05]	recovery rate
μ [days ⁻¹]	[0.001, 0.05]	death rate
e_0 [people]	[0, $N \times 10^{-6}$]	initial exposed

Figure 1: Digraph representation of SEAIR model; model inputs and parameters.

The SEAIR model is described by the following equations:

$$\frac{ds(t)}{dt} = \frac{-\alpha_a(t)}{N}s(t)a(t) - \frac{\alpha_i(t)}{N}s(t)i(t) + \gamma r(t) \quad (1)$$

$$\frac{de(t)}{dt} = \frac{\alpha_a(t)}{N}s(t)a(t) + \frac{\alpha_i(t)}{N}s(t)i(t) - t_{\text{latent}}^{-1}e(t) \quad (2)$$

$$\frac{da(t)}{dt} = t_{\text{latent}}^{-1}e(t) - \kappa(t)a(t) - \rho a(t) \quad (3)$$

$$\frac{di(t)}{dt} = \kappa(t)a(t) - \beta i(t) - \mu i(t) \quad (4)$$

$$\frac{dr(t)}{dt} = \rho a(t) + \beta i(t) - \gamma r(t) \quad (5)$$

$$\frac{dp(t)}{dt} = \mu i(t) \quad (6)$$

where $\alpha_a(t)$ and $\alpha_i(t)$ are the rates of exposure to the virus from the population of asymptomatic/unconfirmed and confirmed infected subjects, respectively. These two rates of exposure are modeled as time-varying and independent inputs to reflect different measures taken during the course of the pandemic. Specifically, $\alpha_a(t)$ corresponds to exposure from asymptomatic carriers (a) and reflects social distancing and/or shelter-in-place strategies. On the other hand, $\alpha_i(t)$ corresponds to exposure from infected subjects (i) and reflects quarantining of infected subjects. We also model $\kappa(t)$, or the rate at which unconfirmed cases become confirmed, as a time-dependent input to reflect varying levels of screening and testing.

The other parameters are assumed to be constant over the time horizons considered herein. t_{latent}^{-1} is the inverse of the latent period of the virus, or the time before an exposed subject

becomes infectious. We assume a value of $t_{\text{latent}}^{-1} = 0.5 \text{ days}^{-1}$ based on Peng et al. [9]. ρ describes the infectious period for subjects with unconfirmed infections, for which we assume a value of $\rho = 0.1 \text{ days}^{-1}$, based on Rocklöv et al. [28]. β and μ describe, respectively, the rates at which subjects with confirmed infections recover and perish. Finally, γ describes the rate at which recovered subjects become susceptible to the disease again. Since the immunity period for the virus is unknown, we assume a value of $\gamma = 0$. While these parameters largely describe the virus itself, they may change over longer time horizons, e.g., as new treatments are developed.

In summary, the SEAIR model (1)–(6) has three time-varying inputs ($\alpha_a(t)$, $\alpha_i(t)$, and $\kappa(t)$), three parameters (t_{latent}^{-1} , ρ , and γ) whose values are based on available literature, and two parameters whose values must be estimated (β and μ).

Parameter estimation results

This work primarily addresses the current COVID-19 situation in the US, but the relatively early stage of the outbreak in the US renders the available data insufficient to fit parameters reflective of prevention measures already in place. For this reason, we perform a comparative analysis by solving the parameter estimation problem for Italy, Spain, and Germany, where, while the epidemic has not been fully contained yet, the daily number of new confirmed cases has been on a steady decline.

Model parameters for the US, Italy, Spain, and Germany

Fig 2 shows the predicted values obtained by solving the parameter estimation problem and the historical data by country, retrieved by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (<https://github.com/CSSEGISandData/COVID-19>; accessed April 6, 2020). Day 1 of the dataset corresponds to January 22. The solid line in each state trajectory plot shows the mean predicted value. Standard deviations for the predictions were estimated using bootstrapping, with the time-invariant parameters sampled from a multivariate normal distribution. The means and covariances of the estimated parameters for each country are provided in the Supplementary Material. Plots corresponding to infected, recovered and perished (left plots in Fig 2) were obtained by simulating (1)–(6) using the estimated trajectories of the time-varying inputs, shown in the right-most column of Fig 2.

Parameter estimation results provide several interesting insights. First we note that the mortality rates, while higher for the European countries, are comparable in magnitude for the four cases considered and are similar to prior estimates [29]. The higher death rates in Italy and Spain are likely explained by demographic factors (e.g., age), as well as the medical resources available to treat the infected population. From the trajectories of α_a , α_i , and κ , we observe that social distancing measures (i.e., lower values of α_a) were first implemented in Europe, and more recently the US, which agrees with their true chronology. This insight is confirmed by the fitted values for e_0 , with Italy having the highest value and the US the lowest, which is representative

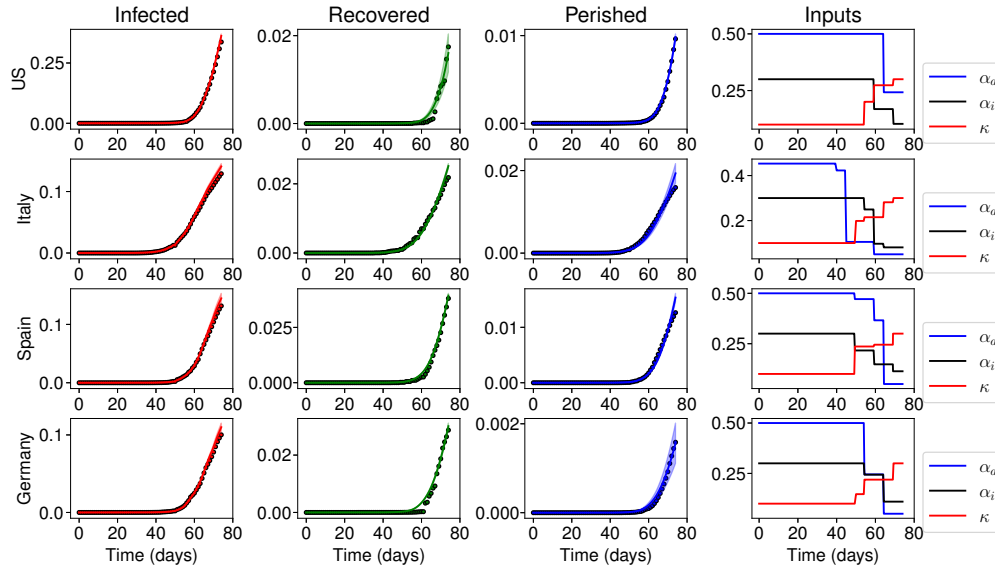


Figure 2: Historical data and model fit for infected ($i(t)$), recovered ($r(t)$), and perished ($p(t)$) subjects (in millions). Solid lines represent the mean of 500 Monte Carlo simulations, shaded areas represent two standard deviations from the mean, and circle markers are historical data. The right-most column shows the fitted trajectories of the time-varying inputs.

of the beginning of the outbreak in each country. In general terms, the identified evolution of containment and testing measures follow similar trends for all four countries, which supports that the model structure is correct and can appropriately reflect (see Methods) different dynamics of the same underlying disease. Further, we expect that the parameters obtained for the US, while being fitted using relatively premature information, are likely an adequate representation of the current COVID-19 situation.

State simulations for fitted parameters

For the estimated parameter values as described previously, we simulate the results of implementing two different simplistic control policies: (i) a moderate level of social distancing, quarantining, and testing, that results in the asymptomatic and infected exposures being halved from their terminal values shown in Fig 2; and (ii) an extreme policy by enforcing a complete lockdown and highest possible screening access (i.e., $\alpha_a(t)$ and $\alpha_i(t)$ are fixed at their lower bounds and $\kappa(t)$ is fixed at its upper bound). The population levels resulting from implementing these two policies are shown in Fig 3, where the numbers of recovered subjects are omitted for the sake of brevity. Since the three European countries showed very similar trends under the two policies, we only compare the results for the US to those obtained for Italy. On the one hand,

from Fig 3 it is evident that more moderate control policies can result in an alarming number of infected cases and deaths, particularly in the US. On the other hand, enforcing strict lock-downs and maximizing testing can result in earlier and substantially flatter pandemic peaks, with significantly lower numbers of casualties. While the latter approach is the most effective at preventing further exponential spread of the disease, it is also the most socially and economically disruptive policy. In light of this trade-off, we argue that by means of employing optimal control concepts it is possible to find effective policies that maintain the number of infected cases below a given threshold, while minimizing the extent of social and economic disruption.

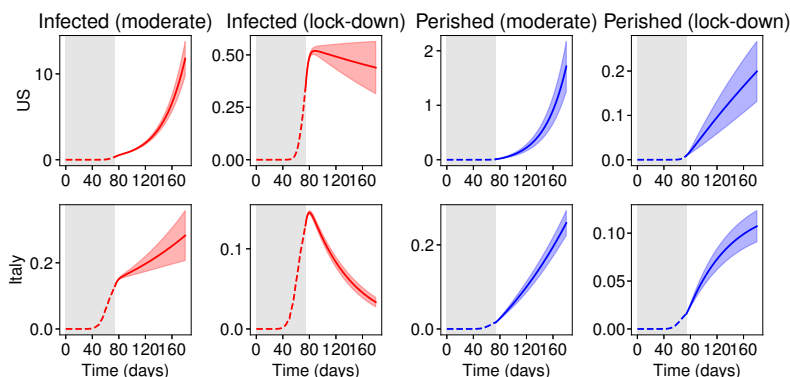


Figure 3: Simulation of future infected ($i(t)$), and perished ($p(t)$) subjects (in millions) for current control policies. Shaded grey area indicates historical data, color solid lines represent the mean of 500 Monte Carlo simulations, and color shaded area represents two standard deviations from the mean.

Dynamic optimization results

Here, we consider the optimal control of the COVID-19 infection in the US using the deterministic SEAIR model (1)–(6) with the mean values of the time-invariant parameters as described above. As such, the optimization problem only considers the mean predicted values (e.g., solid lines in Figs 2 and 3). Future work may consider a stochastic optimization approach that, at the expense of increased computational effort, leverages information about model uncertainty.

Optimization of future actions

We solve a min-max optimization problem that minimizes a measure of socioeconomic cost, subject to keeping the peak (max) value of the infected population below a given number i_{peak} . The lower bound of $\kappa(t)$ is raised to 0.15, from 0.10 in Fig 1. This change was made because (i) the historical value of κ for all countries analyzed quickly jumps to its maximum value 0.3,

perhaps reflecting public initiatives, and (ii) the $[0.15, 0.3]$ range translates to $\rho/(\rho + \kappa) \in [25\%, 40\%]$, which is reflective of current estimates of the asymptomatic ratio [25, 26].

Fig 4 depicts the solution of this problem for the cases where $i_{\text{peak}} = 500,000$ subjects and $i_{\text{peak}} = 1,000,000$ subjects. The optimization problem was solved for the 100 days following the end of the available dataset ($t_0 = 74$, $t_f = 174$), with the initial condition at t_0 computed by simulating days 1–74. Although the optimization problem only considers the mean values (solid lines), two standard deviations are shown in the shaded areas to reflect parametric model uncertainty for the given inputs, computed using the same bootstrapping approach as above.

Considering the current status of the pandemic in the US, keeping the peak below 500,000 infected subjects is very challenging. Note that the minimum feasible value of i_{peak} corresponds to a peak of 458,159 infected subjects. Therefore, maintaining $i(t) \leq 500,000$ requires immediately decreasing α_a and α_i to their lower bounds for approximately the next 25 days, during which the exposed population can decrease significantly. After this period, α_a is relaxed to its upper bound for five brief periods (Fig 4, left). However, α_i remains at its lower bound at all times, reflecting a strict quarantining policy.

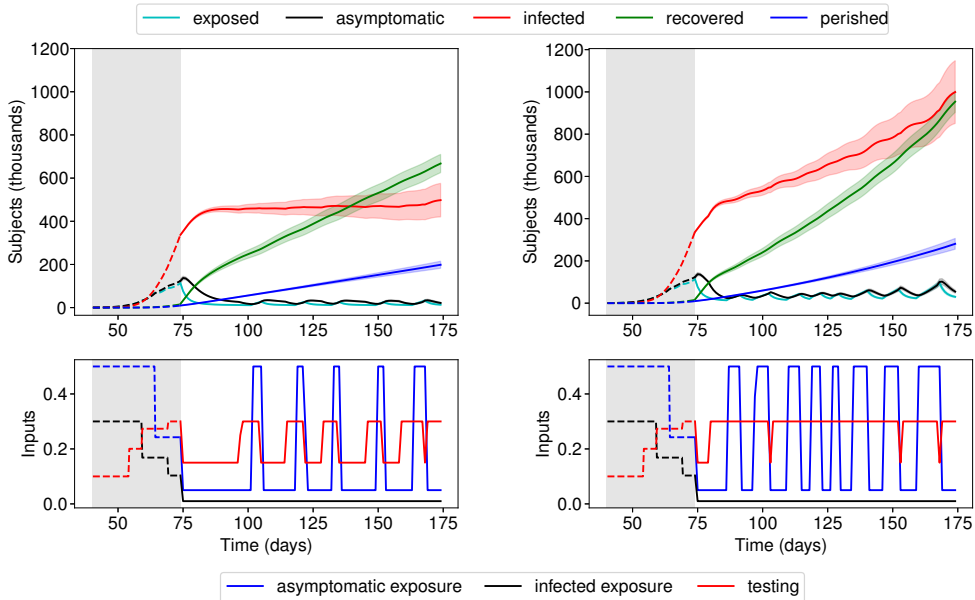


Figure 4: Optimal control policy to limit peak infections to 500,000 (left) or 1,000,000 (right) in the next 100 days. Top: population numbers, with two standard deviations shaded. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized).

This policy reveals that the impact of α_i on the infected population size is larger than that of α_a . In turn, this suggests that quarantining of infected people is more important than social

distancing, which mitigates exposure to unconfirmed cases. κ is decreased during periods of decreased α_a , suggesting that testing is less important during times of social distancing. Intuitively, when the exposure to asymptomatic subjects is already low due to social distancing, there is less benefit to testing asymptomatic subjects (and transferring them to the similarly quarantined infected population). On the other hand, κ is increased preemptively for each period of increased α_a , suggesting that testing is most important in the days leading up to a period of relaxed social distancing. This testing moves asymptomatic subjects to the confirmed-infected population, which can remain quarantined.

Fig 4 (right) depicts the solution obtained for a value of $i_{\text{peak}} = 1,000,000$ people. While α_i again remains at its lower bound at all times, re-emphasizing the importance of quarantining infected subjects, there are more frequent and longer periods of increased α_a . In this case, κ remains at its upper bound for most of the control horizon, and is only decreased slightly immediately before extended periods of social distancing. In both cases, the optimization problem does not account for the effect(s) of the control policy after 100 days, and the growth of the infected population near the end of the time horizon may be concerning. This can be addressed, e.g., using a moving horizon control strategy, where policy measures are revised periodically, as discussed below.

The lower bounds used for α_a and α_i are lower than estimated values for the US for the past (Fig 4), corresponding to new levels of social distancing and quarantining. To investigate the effect of not achieving these proposed levels, we solve the same problem with $i_{\text{peak}} = 1,000,000$ people, using different lower bounds for the inputs. When the lower bound of α_a is doubled (less social distancing is achieved), the same infected population peak can be maintained by increasing the frequency of social distancing periods. However, when the lower bound of α_i is doubled (less quarantining is achieved), maintaining the same peak requires social distancing at almost all times, closely resembling the control policy found for a lower peak (Fig 4, left). Keeping the peak below 1,000,000 infected subjects is infeasible if both bounds are doubled. These results are shown in the Supplementary Material and further confirm the relative importance of quarantining compared to social distancing.

Compared to the decreases seen in Fig 3 for a lock-down scenario, the solutions in Fig 4 seek to merely flatten the growth of the infected population. The associated optimal control policies resemble “bang-bang” control (the inputs are always at the lower or upper bound), which can be expected from a systems-theoretic point of view for (1)–(6), as the equations are linear functions of the inputs α_a , α_i , and κ . This result supports the strategy proposed by the Imperial College COVID-19 Response Team [30], which involves periodic suppression measures. Their strategy triggers the start of a social distancing period by when the number of weekly ICU cases increases past an “on” threshold, and the end of the period when the same number decreases below an “off” threshold. Fig 4 shows that this type of strategy maintains a low exposed population, and therefore flattens the growth of the infected population while still allowing periods of social mobility.

Revisiting past actions

We examine the same optimal control problem with a time horizon starting from day 40 (i.e., covering 34 days in the past). Rolling the time horizon of the optimization problem backwards to $t_0 = 40$, allows us to investigate the optimal inputs for days 40–74, corresponding roughly the month of March. Fig 5 (left) depicts the solution obtained for a value of $i_{\text{peak}} = 500,000$ people. The initial condition at $t_0 = 40$ was computed by simulating days 1–40. The minimum historical values for α_a and α_i were used as their respective lower bounds for days 40–70, such that the optimal control policy for days 40–70 only involves an extent of quarantining and social distancing already experienced by the general population.

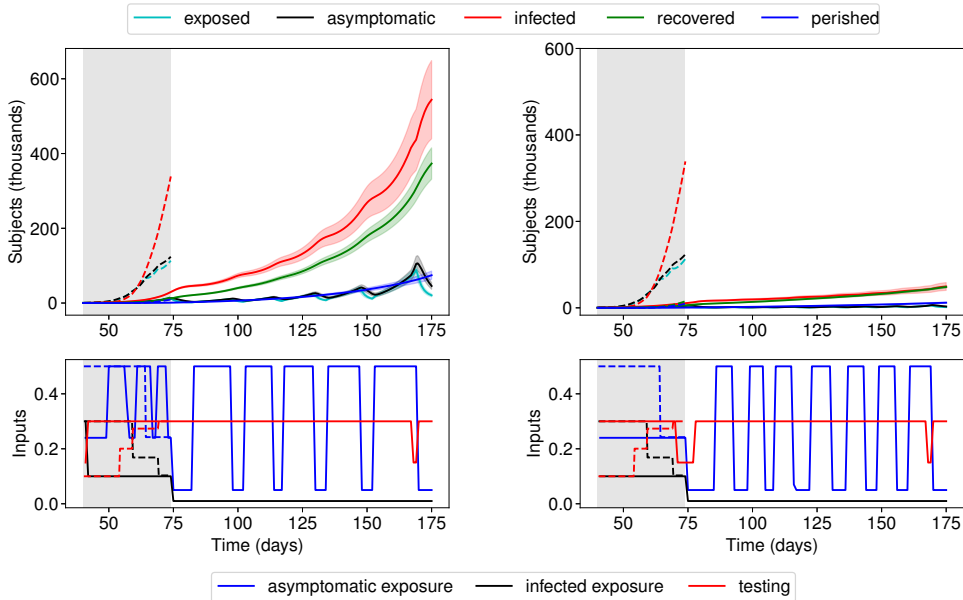


Figure 5: Optimal control policy to limit peak infections to 500,000 (left) or 50,000 (right) in the past 34 days and next 100 days. Top: population numbers, with two standard deviations shaded. Bottom: containment and testing profiles. The shaded grey area indicates past days, for which the historical inputs are shown as dashed lines and the optimized inputs are shown as solid lines.

In this case, maintaining the number of infected subjects below 500,000 appears easy. The minimum feasible peak of infected subjects is 16,110, compared to 458,159 when the optimization starts at day 74. Therefore, maintaining $i(t) \leq 500,000$ can involve extended periods of no social distancing $\alpha_a = 0.5$. Similar to the result in Fig 5 (right), testing and quarantine are important in this scenario: κ remains at its upper bound and α_i its lower bound. The solution for a minimum peak of 50,000 infected subjects, shown in Fig 5 (right), involves more frequent periods of social distancing. Early implementation of testing and quarantining strategies clearly

has an enormous effect, manifest in the large difference between the historical inputs (dashed) and optimal inputs (solid) in days 40–74 of Fig 5.

Estimation of hidden states and moving horizon control

The SEAIR model (1)–(6) includes two *hidden* states, $e(t)$ and $a(t)$, which are not measured in practice (note that widespread serological testing may eventually reveal the true levels of the asymptomatic population a). Nevertheless, solutions for the optimal control problem based on (1)–(6) are highly dependent on the values of $e(t)$ and $a(t)$, and therefore the initial conditions $e(t_0)$ and $a(t_0)$. Fig 6 (left) depicts the optimal control policies for $i_{\text{peak}} = 1,000,000$ people over the next 100 days (similar to Fig 4, right), where $e(t_0)$ and $a(t_0)$ are now underestimated by 50%. This is a practically motivated scenario, as the number of asymptomatic cases is considered largely uncertain [6, 31]. The dashed-dotted lines show the state profiles predicted by the optimization problem (with incorrect initial conditions), while the solid lines show their true evolution for the given inputs.

The error in the hidden states $e(t)$ and $a(t)$ causes the predicted and actual profiles to diverge over time. By the end of the 100 days, there are 1.26 million actual infected subjects, over 25% greater than the predicted one million. To account for discrepancies between the modeled and true systems, the optimal control inputs should be periodically updated. We propose to achieve this with a *moving horizon* decision-making strategy: after a given length of time, the dynamic optimization problem is re-solved, setting the measured values of the states as their initial conditions. The values of the un-measured, hidden states can be estimated using state estimation techniques, such as the Kalman filter (the optimal linear estimator) or its nonlinear extensions.

We consider a relatively long time span of 25 days before each revision of policy decisions (and re-optimization), as planned suppression strategies likely cannot be altered quickly. Fig 6 (right) shows the optimal trajectories found for the same problem (initial conditions for e and a underestimated by 50%) using the moving horizon strategy. The vertical dotted lines indicate times when the optimization problem was resolved. Values of the hidden states were estimated daily with the (discrete) unscented Kalman filter [32], implemented using the pykalman library (<https://pykalman.github.io/>; Accessed April 7, 2020). Measurements for state estimation and control updates were simulated by adding independent, normally distributed noise to the true state values. A standard deviation of 5,000 people was used for the measured states (i , r , and p) based on the size of residuals during parameter estimation.

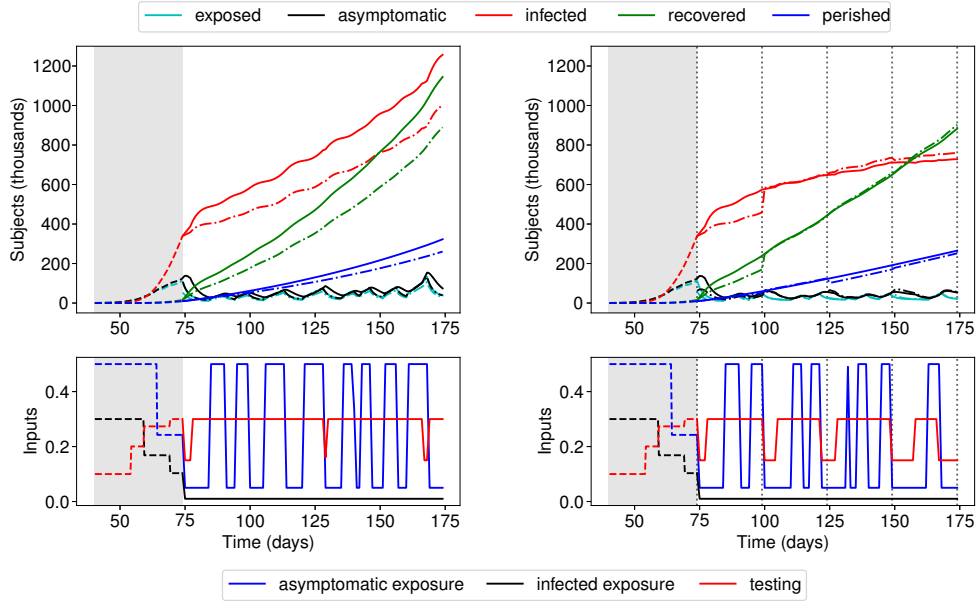


Figure 6: Optimal moving horizon control policy (right) to limit peak infections to 1,000,000, with e and a underestimated by 50% at $t = 74$, and comparison to the same situation without a moving horizon strategy (left). Top: predicted (dash-dotted) and true (solid) population numbers. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized). The policies are updated every 25 days with daily state estimation.

The predicted values (dash-dotted lines in Fig 6) are first updated when the optimization problem is re-solved at $t = 99$ days. The solutions to the re-optimization problems are shown in the Supplementary Material. The effect of this moving horizon strategy is quite significant, in comparison to the optimal policy shown in Fig 6 (left). Periods of social distancing are quickly increased after day 99 to account for the larger-than-expected increase in infected population. κ is also updated accordingly. The moving horizon approach exhibits three clear benefits:

- Errors between the model and reality are minimized by incorporating new measurements as they become available
- The infected population does not grow exponentially at the end of the time horizon, as is the case when considering only a single time horizon
- The control policies can be computed indefinitely into the future by shifting the time window forward in time, while the problem size remains tractable (only 100 days are considered during optimization)

The effect of state estimation is most notable in the first moving horizon window, between days 74–99. Here, while the values predicted by the optimization problem are consistently under-predicted, the estimated values approach the (unknown) true values as daily measurements are added. The new solution of the optimization problem at day 99 uses the population levels measured/estimated based on new data as initial conditions, and therefore predicts more accurate trajectories.

Discussion

This work investigated dynamic optimization strategies to characterize and control the US COVID-19 outbreak, by minimizing the socioeconomic cost associated with containment strategies and testing. The results provide several overarching conclusions.

The quarantining of infected subjects is the most important of the considered mitigation strategies and should be maximized at all times. Additionally, periods of social distancing help to flatten the peak by preventing exposure from asymptomatic and unconfirmed cases. Screening and testing for the disease are key immediately preceding periods of relaxed social distancing, in order to minimize the number of unconfirmed infections during periods of social mobility. Early action has much larger effects than later interventions, even as the later interventions are more drastic. Optimal policies are highly dependent on estimates of “hidden states,” i.e., the asymptomatic and unconfirmed cases. Moving horizon (periodically updated) policies and state estimation should be used to mitigate inaccuracies in the model and counts of asymptomatic/unconfirmed cases, by accounting for new data as they becomes available.

The “on-off” policies identified as optimal are likely the easiest to implement in practical policy and to convey to the general population (as opposed to a policy where the social distancing parameters would take values between their upper and lower bounds). Their implementation would simply alternate between the strictest possible limitations, followed by periods of relative freedom of movement.

The model used in this work does not include population influx to and/or outflux from the given system. Therefore, the optimal containment and testing strategies found do not account for new cases that may enter from outside the US. We also assumed that the recovery and death rates for the virus are constant and equal to their historical values for the next 100 days. Thus the results do not account for the possibility of improved medical treatments, vaccine development, viral mutation, etc. Similarly, the model does not account for possible surveillance/detection bias in the historical data (e.g., increased likelihood of testing for subjects with more severe symptoms).

Methods

Parameter estimation

We follow nonlinear (least-squares) regression for parameter estimation, which can be cast as an optimization problem, where the objective is to find the parameter values that minimize the mean squared error (MSE) between the predicted states and their measured values. The measured values correspond to data retrieved by the CSSE at Johns Hopkins University (<https://github.com/CSSEGISandData/COVID-19>; accessed April 6, 2020) and are provided in the Supplementary Material for several countries. The data consist of the total number of infected (\hat{i}_j), recovered (\hat{r}_j) and dead (\hat{p}_j) reported subjects, where j represents each day during the time period from January 22 to April 5, where the index j represents the time period each value was recorded. The MSE is given by:

$$L_0 = \sum_{j=0}^N (i(t=j) - \hat{i}_j)^2 + (r(t=j) - \hat{r}_j)^2 + (p(t=j) - \hat{p}_j)^2 \quad (7)$$

We note that when solving the parameter estimation problem, the contribution of the term $\rho a(t)$ to the recovered state $r(t)$, is removed from the system (1)-(6). This is because the data obtained only accounts for those recovered that were confirmed to have the disease, meaning the only contribution to the observed recovered is the term $\beta i(t)$.

The parameters with values/trajectories to estimate are $\alpha_a(t)$, $\alpha_i(t)$, $\kappa(t)$, β , and μ . Furthermore, we must estimate the initial condition of the unreported states e and a at $t = 0$. We assume that there are initially no asymptomatic infections. The “seed” of the outbreak is therefore the initial number of exposed subjects $e_0 = e(t = 0)$, which we include as a parameter in the least-squares problem. We include bounds on the possible parameter values, based on values reported for similar models fitted to data from other regions. [9] Note that, while κ reflects the level of testing, it also affects the predicted asymptomatic ratio, which cannot be controlled. Therefore, tighter bounds should be used for κ than for α_a and α_i . The estimated parameters and their bounds are summarized in Fig 1.

To prevent over-fitting the available data, we restrict the time-varying inputs (α_a , α_i , and κ) to be piece-wise constant over five-day intervals. These constraints reflect the fact that policies implemented have significant time delays in steering the states towards the desired values, and therefore should not and practically cannot be manipulated too frequently before their effect is observed on the population. Additionally, we constrain $\alpha_a(t)$ and $\alpha_i(t)$ to be monotonically decreasing over time, reflecting the enforcement of increasingly stricter disease control measures. Similarly, $\kappa(t)$ was constrained to be monotonically increasing over time to account for increased COVID-19 screening availability.

To easily interface with raw data sources, the least-squares regression problem was implemented in Python using the Pyomo modeling and optimization package [33]. We discretized the dynamical system (1)-(6) with respect to the time domain using orthogonal collocation on finite elements [23], with one finite element per day. We used IPOPT [34] to solve the resulting

nonlinear dynamic optimization problem to local optimality. Standard deviations for the time-invariant parameters were approximated by fixing the time-varying inputs to their estimated trajectories, and solving a maximum likelihood problem for the time-invariant parameters using the Model Validation tool in gPROMS (general PROcess Modeling System) v5.1.4 [35]. Residuals were assumed independent and normally distributed, with variances estimated in the same maximum likelihood problem. The parameter covariances (see SI) are then computed from the diagonal entries of the Hessian of the objective function.

Optimal control

The aim of the optimal control problem is to find the trajectories of “handles” $\alpha_a(t)$, $\alpha_i(t)$ and $\kappa(t)$, that minimize (or maximize) the value of a certain objective function. Most response measures for the COVID-19 outbreak seek to “flatten” the epidemic curve, that is, to contain the growth rate of the number of infected subjects via a combination of social distancing and testing. Clearly, social distancing and isolation/quarantining carry significant social and economic costs. We thus formulate the optimal control problem as a dynamic optimization problem, aiming to minimize a measure of social and economic cost subject to ensuring that the maximum number of infected subjects remains under a given peak value, i_{peak} . The optimization problem is expressed mathematically as:

$$\begin{aligned}
 & \min_{\alpha_a(t), \alpha_i(t), \kappa(t)} \int_{t_0}^{t_f} C(t) dt \\
 & \text{s.t. SEAIR model (1)-(6)} \\
 & \quad \max_t (i(t)) \leq i_{\text{peak}} \\
 & \quad C(t) = -\alpha_a(t) - \alpha_i(t) + \lambda_\kappa \kappa(t) \\
 & \quad \alpha_a(t) \in [0.05, 0.5] \\
 & \quad \alpha_i(t) \in [0.01, 0.3] \\
 & \quad \kappa(t) \in [0.15, 0.3]
 \end{aligned} \tag{8}$$

where $C(t)$ is the cost function, and λ_κ is the relative cost of testing (increasing κ). Testing costs are assumed to be relatively small in comparison to isolation measures, and we selected $\lambda_\kappa = 0.1$. Problem (8) can be solved over a longer time horizon (e.g., $t_f = 100$ days) as in the case of the initial studies presented in this work, or over a shorter horizon (e.g., $t_f = 25$ days) as in the case of the moving horizon approach. In the latter case, the time window considered, $[t_0, t_f]$ is “shifted right” at a pre-determined time interval (25 days in the case considered earlier). Note that t_f is typically longer than the frequency at which the optimization problem is recomputed, and only the solution for the first time step(s) is implemented.

While the problem (8) identifies the minimum societal cost required to achieve a certain peak value of infections, a similar problem could be formulated by minimizing i_{peak} subject to an upper bound on $C(t)$, i.e., finding the lowest achievable peak for a given total societal cost.

For this study we investigate the relationship between C and i_{peak} by solving (8) for varying values of i_{peak}

The optimization problem (8) was again solved by discretization of the time domain using orthogonal collocation, with one finite element per day. To report the best possible solutions, the resulting algebraic optimization problem was implemented in GAMS and solved using the commercial global optimization solver BARON v18.5.8 [36]. Problems were solved to a 0.1% optimality gap (i.e., the reported solution is proven to have an objective function value within 0.1% of that of the best possible solution). We additionally provide an implementation of the same problem using open-source tools (Pyomo [33] and IPOPT [34]) in the Supplementary Material, which can be used to solve (8) to local optimality.

References

- [1] Noah C Peeri, Nistha Shrestha, Md Siddikur Rahman, Rafdzah Zaki, Zhengqi Tan, Saana Bibi, Mahdi Baghbanzadeh, Nasrin Aghamohammadi, Wenyi Zhang, and Ubydul Haque. The SARS, MERS and novel coronavirus (COVID-19) epidemics, the newest and biggest global health threats: what lessons have we learned? *International Journal of Epidemiology*, 2020. doi:10.1093/ije/dyaa033.
- [2] Catrin Sohrabi, Zaid Alsafi, Niamh O'Neill, Mehdi Khan, Ahmed Kerwan, Ahmed Al-Jabir, Christos Iosifidis, and Riaz Agha. World Health Organization declares global emergency: A review of the 2019 novel coronavirus (COVID-19). *International Journal of Surgery*, 76:71–76, 2020.
- [3] M Paterlini. Lockdown in Italy: personal stories of doing science during the COVID-19 quarantine. *Nature: Career Feature*, 2020. doi:10.1038/d41586-020-01001-8.
- [4] Aurelio Tobías. Evaluation of the lockdowns for the SARS-CoV-2 epidemic in Italy and Spain after one month follow up. *Science of The Total Environment*, page 138539, 2020.
- [5] Moritz UG Kraemer, Chia-Hung Yang, Bernardo Gutierrez, Chieh-Hsi Wu, Brennan Klein, David M Pigott, Louis du Plessis, Nuno R Faria, Ruoran Li, William P Hanage, et al. The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science*, eabb4218, 2020.
- [6] Katelyn Gostic, Ana CR Gomez, Riley O Mummah, Adam J Kucharski, and James O Lloyd-Smith. Estimated effectiveness of symptom and risk screening to prevent the spread of COVID-19. *eLife*, 9:e55570, 2020.
- [7] Matt J Keeling and Pejman Rohani. *Modeling infectious diseases in humans and animals*. Princeton University Press, 2011.

- [8] Cleo Anastassopoulou, Lucia Russo, Athanasios Tsakris, and Constantinos Siettos. Data-based analysis, modelling and forecasting of the COVID-19 outbreak. *PloS ONE*, 15(3):e0230405, 2020.
- [9] Liangrong Peng, Wuyue Yang, Dongyan Zhang, Changjing Zhuge, and Liu Hong. Epidemic analysis of COVID-19 in China by dynamical modeling. *arXiv preprint arXiv:2002.06563*, 2020.
- [10] Pierre Magal, , and Glenn Webb. Predicting the number of reported and unreported cases for the COVID-19 epidemic in South Korea. *medRxiv 2020.03.21.20040154*, 2020.
- [11] Adam J Kucharski, Timothy W Russell, Charlie Diamond, Yang Liu, John Edmunds, Sebastian Funk, Rosalind M Eggo, Fiona Sun, Mark Jit, James D Munday, et al. Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The Lancet Infectious Diseases*, 2020. doi:10.1016/S1473-3099(20)30144-4.
- [12] Ying Liu, Albert A Gayle, Annelies Wilder-Smith, and Joacim Rocklöv. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine*, 27(2):taaa021, 2020.
- [13] Zhihua Liu, Pierre Magal, Ousmane Seydi, and Glenn Webb. Predicting the cumulative number of cases for the COVID-19 epidemic in China from early data. *arXiv preprint arXiv:2002.12298*, 2020.
- [14] Sang Woo Park, Daniel M Cornforth, Jonathan Dushoff, and Joshua S Weitz. The time scale of asymptomatic transmission affects estimates of epidemic potential in the COVID-19 outbreak. *medRxiv 2020.03.09.20033514*, 2020.
- [15] Joseph T Wu, Kathy Leung, and Gabriel M Leung. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in wuhan, china: a modelling study. *The Lancet*, 395(10225):689–697, 2020.
- [16] Reza Sameni. Mathematical modeling of epidemic diseases; a case study of the COVID-19 coronavirus. *arXiv preprint arXiv:2003.11371*, 2020.
- [17] Jinhua Pan, Ye Yao, Zhixi Liu, Mengyi Li, Ying Wang, Weizhen Dong, Haidong Kan, and Weibing Wang. Effectiveness of control strategies for Coronavirus Disease 2019: a SEIR dynamic modeling study. *medRxiv 2020.02.19.20025387*, 2020.
- [18] Ramses Djidjou-Demasse, Yannis Michalakis, Marc Choisy, Micea T Sofonea, and Samuel Alizon. Optimal COVID-19 epidemic control until vaccine deployment. *medRxiv 2020.04.02.2004918*, 2020.
- [19] Stephen E Moore and Eric Okyere. Controlling the transmission dynamics of COVID-19. *arXiv preprint arXiv:2004.00443*, 2020.

- [20] Suzanne Lenhart and John T Workman. *Optimal control applied to biological models*. CRC press, 2007.
- [21] Md Haider Ali Biswas, Luís Tiago Paiva, and MDR De Pinho. A SEIR model for control of infectious diseases with constraints. *Mathematical Biosciences and Engineering*, 11(4):761–784, 2014.
- [22] Rachael Miller Neilan and Suzanne Lenhart. An introduction to optimal control with an application in disease modeling. In *Modeling Paradigms and Analysis of Disease Trasmision Models*, pages 67–82, 2010.
- [23] Lorenz T Biegler. An overview of simultaneous strategies for dynamic optimization. *Chemical Engineering and Processing: Process Intensification*, 46(11):1043–1053, 2007.
- [24] Christodoulos A Floudas. *Deterministic global optimization: theory, methods and applications*, volume 37. Springer Science & Business Media, 2013.
- [25] Kenji Mizumoto, Katsushi Kagaya, Alexander Zarebski, and Gerardo Chowell. Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan, 2020. *Eurosurveillance*, 25(10):2000180, 2020.
- [26] Hiroshi Nishiura, Tetsuro Kobayashi, Takeshi Miyama, Ayako Suzuki, Sungmok Jung, Katsuma Hayashi, Ryo Kinoshita, Yichi Yang, Baoyin Yuan, Andrei R Akhmetzhanov, and Natalie M Linton. Estimation of the asymptomatic ratio of novel coronavirus infections (COVID-19). *International Journal of Infectious Diseases*, 2020. doi:10.1016/j.ijid.2020.03.020.
- [27] Yan Bai, Lingsheng Yao, Tao Wei, Fei Tian, Dong-Yan Jin, Lijuan Chen, and Meiyun Wang. Presumed asymptomatic carrier transmission of COVID-19. *JAMA*, 2020. doi:10.1001/jama.2020.2565.
- [28] Joacim Rocklöv, Henrik Sjödin, and Annelies Wilder-Smith. COVID-19 outbreak on the Diamond Princess cruise ship: estimating the epidemic potential and effectiveness of public health countermeasures. *Journal of Travel Medicine*, 2020. doi:10.1093/jtm/taaa030.
- [29] Chih-Cheng Lai, Tzu-Ping Shih, Wen-Chien Ko, Hung-Jen Tang, and Po-Ren Hsueh. Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) and corona virus disease-2019 (COVID-19): the epidemic and the challenges. *International Journal of Antimicrobial Agents*, page 105924, 2020.
- [30] Neil M Ferguson, Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunuba Perez, Gina Cuomo-Dannenburg, Amy Dighe, Ilaria Dorigatti, Han Fu, Katy Gaythorpe, Will Green,

Arran Hamlet, Wes Hinsley, Lucy C Okell, Sabine van Elsland, Hayley Thompson, Robert Verity, Erik Volz, Haowei Wang, Yuanrong Wang, Patrick GT Walker, Caroline Walters, Peter Winskill, Charles Whittaker, Christi A Donnelly, Steven Riley, and Azra C Ghani. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Technical report, Imperial College London, 2020.

- [31] Roy M Anderson, Hans Heesterbeek, Don Klinkenberg, and T Déirdre Hollingsworth. How will country-based mitigation measures influence the course of the COVID-19 epidemic? *The Lancet*, 395(10228):931–934, 2020.
- [32] Eric A Wan and Rudolph Van Der Merwe. The unscented Kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*, pages 153–158. IEEE, 2000.
- [33] William E Hart, Carl D Laird, Jean-Paul Watson, David L Woodruff, Gabriel A Hackebeil, Bethany L Nicholson, and John D Sirola. *Pyomo-optimization modeling in python*, volume 67. Springer, 2017.
- [34] Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
- [35] Process Systems Enterprise. general PROcess Modeling System (gPROMS). www.psenderprise.com/products/gproms, 1997-2020.
- [36] Mohit Tawarmalani and Nikolaos V Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103(2):225–249, 2005.

Acknowledgments

Partial financial support from the National Science Foundation (NSF) through CAREER Award 1454433 is acknowledged with gratitude. Data and open-source optimization models from this work are available at <https://github.com/Baldea-Group/covid-19>.

Supplementary Material

A least-squares regression problem was solved for the values of the parameters in the SEAIR model (1)–(6). The estimated means of the parameters for the USA, Italy, Spain, and Germany are given below in Table 1, and the covariance matrices are given in Table 2. These covariance matrices were used to generate samples for bootstrapping.

	USA	Italy	Spain	Germany
μ [days ⁻¹]	4.103×10^{-3}	10.329×10^{-3}	11.301×10^{-3}	1.512×10^{-3}
β [days ⁻¹]	6.696×10^{-3}	13.423×10^{-3}	$29.010 \pm \times 10^{-3}$	28.729×10^{-3}
e_0 [people]	1.511×10^{-1}	32.797×10^{-1}	2.419×10^{-1}	1.934×10^{-1}

Table 1: Mean values of time-invariant model parameters.

USA			
μ [days ⁻¹]	7.41×10^{-7}		
β [days ⁻¹]	9.58×10^{-8}	4.76×10^{-6}	
μ [people]	-2.20×10^{-10}	9.72×10^{-8}	4.87×10^{-8}

Italy			
μ [days ⁻¹]	1.07×10^{-8}		
β [days ⁻¹]	-1.48×10^{-6}	2.18×10^{-3}	
μ [people]	-2.64×10^{-8}	2.78×10^{-5}	3.70×10^{-7}

Spain			
μ [days ⁻¹]	-7.99×10^{-7}		
β [days ⁻¹]	-1.18×10^{-6}	3.87×10^{-6}	
μ [people]	2.97×10^{-7}	-3.51×10^{-7}	1.94×10^{-7}

Germany			
μ [days ⁻¹]	4.34×10^{-7}		
β [days ⁻¹]	9.58×10^{-8}	4.76×10^{-6}	
μ [people]	2.20×10^{-10}	9.72×10^{-8}	4.87×10^{-8}

Table 2: Covariance matrices of the time-invariant model parameters.

Given the parameters for the US outbreak, the optimization problem (8) can be solved for the next 100 days. Fig 4 shows the solutions for peak infected population i_{peak} values of 500,000 and 1,000,000 people. Here, we solve (8) for various values of i_{peak} to investigate the relationship between i_{peak} and the associated socioeconomic cost C . The results are given in Fig 7. As the cost function is in arbitrary units, we report the cost here as the percentage increase over the cost function value for normal societal operation, i.e., with no mitigation/testing strategies ($\alpha_a = 0.5$, $\alpha_i = 0.3$, $\kappa = 0.1$).

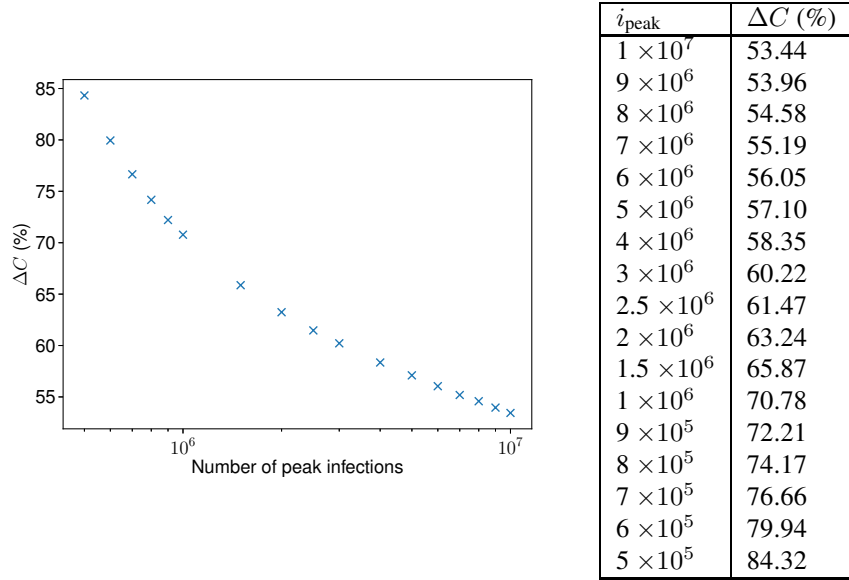


Figure 7: Cost required to limit peak infections to various values. Costs given as percentage increases from no-action taken scenario.

The solutions in Fig 4 are dependent on the lower bounds of α_a and α_i , which reflect levels of social distancing and quarantining higher than those experienced in the past. We therefore investigate (i) the effects of doubling the lower bound of α_a and (ii) the effects of doubling the lower bound of α_i . Results from these two scenarios are shown in Fig 8. The problem is infeasible if both bounds are simultaneously doubled.

The moving horizon strategy shown in Fig 6 involves updating control policies in 25-day intervals (after day 74). The problem solved at each update is identical, except that the initial conditions are updated based on the current measured/estimated populations. The solutions obtained at days 99, 124, and 149 are shown, respectively, in Fig 9, Fig 10, and Fig 11.

The moving horizon strategy (Fig 6) relies on state estimation for values of $e(t)$ and $a(t)$. While the control policies are only updated every 25 days, the state estimates are updated daily as new measurements are made available. The values of $e(t)$ and $a(t)$ were estimate using an unscented Kalman filter for moving horizon control. The estimates over time are shown in Fig 12.




Figure 8: Optimal input trajectories to keep peak infections below 1,000,000 in the next 100 days for different bounds. Top: $\alpha_a(t)$. Bottom: $\kappa(t)$. The shaded grey area indicates the solution found using the normal bounds, replicated from Fig 4 (right).

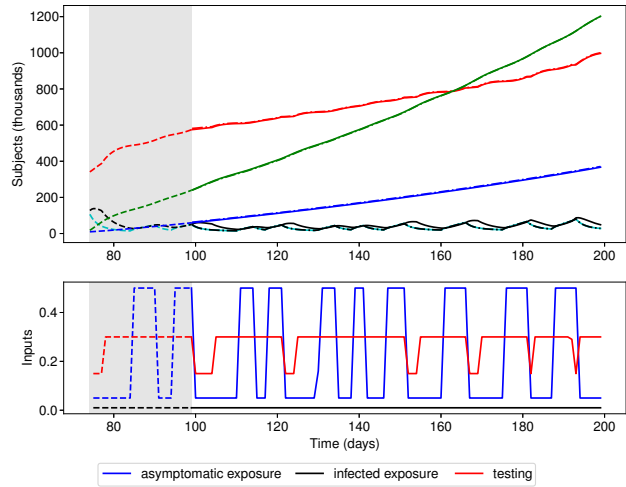


Figure 9: Optimal control policy found at day 99 in moving horizon scenario. Top: predicted (dash-dotted) and true (solid) population numbers. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized).




Figure 10: Optimal control policy found at day 124 in moving horizon scenario. Top: predicted (dash-dotted) and true (solid) population numbers. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized).




Figure 11: Optimal control policy found at day 149 in moving horizon scenario. Top: predicted (dash-dotted) and true (solid) population numbers. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized).

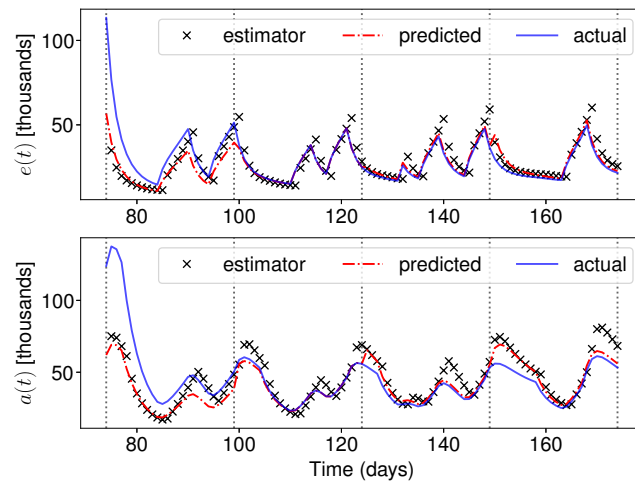


Figure 12: Estimation of hidden states during moving horizon control. Top: estimated, predicted (during optimization), and actual values of $e(t)$. Bottom: estimated, predicted, and actual values of $a(t)$. Re-optimization is performed every 25 days, while state estimation is performed daily.