



Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



Full Length Article

Multitask learning and nonlinear optimal control of the COVID-19 outbreak: A geometric programming approach[☆]

Mikhail Hayhoe ^{a,*}, Francisco Barreras ^b, Victor M. Preciado ^a^a Department of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA, 19104, USA^b Department of Mathematics, University of Pennsylvania, Philadelphia, PA, 19104, USA

ARTICLE INFO

ABSTRACT

Keywords:

Epidemiology
Mathematical modeling
Multitask learning
Optimal control
Geometric programming

We propose a multitask learning approach to learn the parameters of a compartmental discrete-time epidemic model from various data sources and use it to design optimal control strategies of human-mobility restrictions that both curb the epidemic and minimize the economic costs associated with implementing non-pharmaceutical interventions. We develop an extension of the SEIR epidemic model that captures the effects of changes in human mobility on the spread of the disease. The parameters of the model are learned using a multitask learning approach that leverages both data on the number of deaths across a set of regions, and cellphone data on individuals' mobility patterns specific to each region. Using this model, we propose a nonlinear optimal control problem aiming to find the optimal mobility-based intervention strategy that curbs the spread of the epidemic while obeying a budget on the economic cost incurred. We also show that the solution to this nonlinear optimal control problem can be efficiently found, in polynomial time, using tools from geometric programming. Furthermore, in the absence of a straightforward mapping from human mobility data to economic costs, we propose a practical method by which a budget on economic losses incurred may be chosen to eliminate excess deaths due to over-utilization of hospital resources. Our results are demonstrated with numerical simulations using real data from the COVID-19 pandemic in the Philadelphia metropolitan area.

1. Introduction

Ever since the first COVID-19 case was reported on December 31st 2019 (World Health Organization, 2020a), the SARS-CoV-2 pandemic has spread world-wide, reaching alarming levels of spread and severity (World Health Organization, 2020b). The response to the first wave of COVID-19 by governments was the implementation of large scale non-pharmaceutical interventions (NPIs) ranging from contact tracing, quarantines and mask usage, to more aggressive measures like city wide shelter-in-place orders, air-travel restrictions and closures of non-essential businesses (Wu, Smith, Khurana, Siemaszko, & DeJesus-Banos, 2020). In the absence of pharmaceutical treatment, prevention, or herd immunity, NPIs remain the only tool to curb the spread of an epidemic in its earlier stages. In the case of COVID-19, governments across the world have implemented strategies to relax mobility restriction measures and reactivate the economy (Kaplan & Frias, 2020) while, at the same time, preventing the collapse of healthcare systems. However, relaxing mobility restrictions too fast or carelessly can result in

resounding waves, as we are currently observing for the case of COVID-19. In fact, as long as enough people in the population are susceptible, the danger of recurrent waves is not only real, but probable. In this situation, it is of utmost societal importance to develop reopening strategies in a principled manner utilizing the wealth of data readily available.

Several epidemic models have been proposed in the recent literature to simulate the effects of social distancing on the evolution of the pandemic; see, e.g., Achterberg et al. (2020), Bhouri et al. (2021), Chang et al. (2021). Although the majority of epidemic models in recent years are variations of the seminal mathematical models on theoretical epidemiology (see Nowzari, Preciado, and Pappas (2016) and references therein), the availability of rich datasets describing human mobility and behavior is rapidly changing the field of mathematical modeling of epidemics. Companies like Google, Foursquare, Safegraph, Baidu, and others, have provided public access to massive datasets describing human mobility, enabling the development of data-driven

[☆] This work was supported, in part, by the National Science Foundation, USA under awards NSF-TRIPODS-1934876, CAREER-ECCS-1651433, NSF-III-2008456, and the Rockefeller Foundation, USA.

* Corresponding author.

E-mail addresses: mhayhoe@seas.upenn.edu (M. Hayhoe), fbarrer@sas.upenn.edu (F. Barreras), preciado@seas.upenn.edu (V.M. Preciado).

epidemic models capturing the effects of mobility restrictions. Indeed, the choices faced by decision makers regarding disease management involve the use of multiple control actuators such as vaccination, quarantine, treatment or, as is the case for COVID-19, non-pharmaceutical interventions such as social distancing. These decisions must face the trade-off of minimizing the impact of the disease and the economic cost associated with their implementation.

In order to increase predictive power and utility for policy decision-makers, epidemic models have gradually increased their complexity to account for a multitude of features of real epidemics such as disease-specific compartmental models (Van den Broeck et al., 2011), resurgence (Watts, Muhamad, Medina, & Dodds, 2005), multi-scale effects (Hayhoe, Alajaji, & Gharesifard, 2018), seasonality (Balcan et al., 2010), differential risk structure in the population (Ferguson et al., 2020), healthcare system capacity (Piguillem & Shi, 2020), and uncertainty in testing data (Morato, Bastos, Cajueiro, & Normey-Rico, 2020), among others. This increased sophistication in the modeling often comes at the cost of mathematical intractability, a limitation often circumvented by formulating policies based on heuristics and/or simulations (Achterberg et al., 2020; Aleta et al., 2020; Balcan et al., 2010; Chang et al., 2021; Ferguson et al., 2020; Lorch et al., 2020). Although informative for certain scenarios, these proposed interventions are not the result of rigorously formulated optimal control problems and, thus, lack the guarantees and flexibility of mathematical optimization frameworks.

Conversely, the control of epidemics does not usually admit straightforward solutions from optimal control theory due to the presence of nonlinearities and/or the lack of convexity (Nowzari et al., 2016). There are important theoretical results in optimal control of epidemics which achieve mathematical tractability; for example, optimal resource allocation aiming to asymptotically drive the epidemic to extinction (Birge, Candogan, & Feng, 2020; Hota, Godbole, Bhariya, & Paré, 2020; Nowzari, Preciado, & Pappas, 2015; Preciado, Zargham, Enyioha, Jadbabaie, & Pappas, 2013, 2014; Van Mieghem, Omic, & Kooij, 2009; Wang, Chakrabarti, Wang, & Faloutsos, 2003). Other theoretical results are concerned with applications of Pontryagin's maximum principle (PMP), and find exact solutions to resource allocation problems under some variations of the SIS and SEIR, for example in Eshghi, Khouzani, Sarkar, and Venkatesh (2015), Khouzani, Venkatesh, and Sarkar (2011), Yan and Zou (2008). In contrast, data-driven control frameworks in real epidemics have found limited applicability due to the difficulty in incorporating real data and the challenges in solving non-convex programs exactly. Recent work Birge et al. (2020) has proposed a solution to the problem of optimal social distancing with economic constraints using a static convex program which seeks to ensure the decrease of the epidemic at all times; a condition that is too stringent and fails to consider more efficient dynamic strategies. In Piguillem and Shi (2020), an optimal quarantine and testing policy is derived in an optimal control framework, but the control actions and objective functions are restrictive and chosen ad-hoc for tractability. Other recent applications have tackled the non-convex optimal control of social distancing policies using model predictive control (MPC), an approach which has been applied to a plethora of non-linear control problems in industry (Garcia, Prett, & Morari, 1989). For example, in Köhler et al. (2020) the authors propose an optimal predictive control problem where a control input representing social distancing affects the infectivity rates directly and the number of fatalities is approximately minimized using a nonlinear program solver. A related application of MPC for optimal social distancing policies is found in Morato et al. (2020), in which, instead of controlling the infectivity rates directly, the control input represents a binary lockdown policy enacted by the government which has a delayed effect in the population's level of isolation, in turn affecting the infectivity rate.

A practical concern is whether it may be possible to design optimal control strategies based on mobility restrictions that are fully data-driven, in the sense that human mobility is measured and explicitly

incorporated in the epidemic model. In this paper, we propose a model of the spread of COVID-19 and a data-driven optimal control problem that directly minimizes the number of predicted cumulative deaths by implementing mobility restrictions in the population. We use real mobility data from Google (Google, 2020) to learn a nonlinear mapping representing the impact of human mobility on the parameters of a dynamical epidemic model and propose a novel nonlinear optimal control problem that can be solved efficiently using tools from geometric programming (Boyd, Kim, Vandenberghe, & Hassibi, 2007).

Our model consists of an extension of the classic SEIR model, augmented with compartments for asymptomatic and hospitalized agents. We assume that the rate at which agents become infected in any given day is a function of the mobility trends in the population for that same day, reflecting the fact that an increase in mobility leads to more infections. We rely on data regarding case counts and deaths from The New York Times (New York Times, 2020) and from Google's COVID-19 Community Mobility Reports (Google, 2020) to capture the changes in visitation patterns to different Places of Interest (POIs). The mobility data consists of several time series measuring visits to various categories of places such as Retail & Recreation, Grocery & Pharmacy, and Workplaces. The dataset is organized into separate time series for all counties in the United States, and measures visits to multiple categories of places against a benchmark established in January and February of 2020. The key observation is that a decision maker can enforce restrictions on visits to each of these categories to reduce the spread of the epidemic while incurring a cost to the economy. Hence, our objective is to design optimal strategies for mobility restrictions to contain the spread of COVID-19 while taking into account the associated economic cost.

Our approach is similar to that in Köhler et al. (2020), Morato et al. (2020) in that we solve the problem of optimal social distancing policies to curb an epidemic under state constraints, which can be implemented with a receding horizon. However, we stress two important differences: first, we explicitly model and learn the impact of human mobility on the evolution of the disease spread using real mobility data and can formulate granular continuous mobility restriction policies that vary across economic sectors; second, unlike Köhler et al. (2020), Morato et al. (2020) our optimization problem is reduced to a convex program which can be solved with great efficiency, has global optimality guarantees, and can accommodate a large number of variables and longer optimization horizons.

The structure of the paper is as follows. In Section 2 we introduce the notation used as well as some necessary background in geometric programming. In Section 3 we discuss the specifics of our data-driven model, consisting of a mobility layer and an epidemic layer. In Section 4 we discuss the details of our learning strategy to identify the parameters of the model. In Section 5 we present our optimal control framework and present simulations showing the effectiveness of our method. In Section 6 we conclude and discuss possibilities for further research. Appendix A includes additional results of practical interest, and Appendix B contains proofs of the results herein.

2. Background and notation

Throughout this paper bold characters are used to denote vectors and upper-case characters denote either matrices or compartments of the epidemic model. For the following definitions let $x_1, \dots, x_n \geq 0$ denote n non-negative variables, and let $\mathbf{x} = (x_1, \dots, x_n)$. When considering our epidemic model we use tildes for the states of the true (nonlinear) model, e.g. $\tilde{S}(t)$, and omit the tildes for the linearized model, e.g., $S(t)$.

Definition 1 (Monomial). A function $f(\mathbf{x})$ is called a *monomial* if it has the form

$$f(\mathbf{x}) = c x_1^{a_1} x_2^{a_2} \dots x_n^{a_n},$$

for $a_1, \dots, a_n \in \mathbb{R}$ and $c > 0$.

Definition 2 (Posynomial). A sum of one or more monomials is called a *posynomial*, that is, a function of the form

$$f(\mathbf{x}) = \sum_{i=1}^k c_i x_1^{a_{i,1}} x_2^{a_{i,2}} \dots x_n^{a_{i,n}}.$$

Since posynomials admit negative exponents but do not admit negative coefficients they are not necessarily polynomials, and vice versa. We remark that posynomials are closed under addition, multiplication, and positive scalar multiplication. This implies that if the entries of two matrices $A \in \mathbb{R}^{m \times k}$ and $B \in \mathbb{R}^{k \times n}$ are posynomials of the same variables, then so are the entries of their product AB , since $[AB]_{i,j} = \sum_{t=1}^k A_{i,t} B_{t,j}$, which is a sum of products of posynomials. This result extends trivially to the product of an arbitrary number of matrices with posynomial entries.

Definition 3 (Convex in log-scale). A function $f(\mathbf{x})$ is convex in log-scale if the function $F(\mathbf{y}) := \log f(\exp(\mathbf{y}))$ is convex (where the exponentiation is component wise).

A careful application of Hölder's inequality shows that posynomials are convex in log-scale.

We solve the epidemic control problems presented herein using a quasi-convex optimization framework called *geometric programming* (Boyd et al., 2007; Boyd & Vandenberghe, 2004), which has found wide applicability in fields such as communication systems (Chiang, 2005), epidemiology (Preciado et al., 2014), and control (Ogura, Kishida, & Lam, 2019), among others. A *geometric program* (GP) is a mathematical optimization program of the form

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} \quad f(\mathbf{x}) \\ & \text{subject to} \quad q_i(\mathbf{x}) \leq 1 \quad \text{for } i \in 1, \dots, m \\ & \quad h_i(\mathbf{x}) = 1 \quad \text{for } i \in 1, \dots, p, \end{aligned}$$

where $f(\mathbf{x}), q_1(\mathbf{x}), \dots, q_m(\mathbf{x})$ are posynomials and $h_1(\mathbf{x}), \dots, h_p(\mathbf{x})$ are monomials. Due to the convexity in log-scale, one can exactly transform¹ a GP to a convex program by means of the logarithmic change of variables $y_i = \log(x_i)$ and transforming the objective and constraints with the logarithmic transformations $F(\mathbf{y}) = \log f(\exp(\mathbf{y}))$, $Q_i(\mathbf{y}) = \log q_i(\exp(\mathbf{y}))$ and $H_i(\mathbf{y}) = \log h_i(\exp(\mathbf{y}))$ to obtain

$$\begin{aligned} & \underset{\mathbf{y}}{\text{minimize}} \quad F(\mathbf{y}) \\ & \text{subject to} \quad Q_i(\mathbf{y}) \leq 0 \quad \text{for } i \in 1, \dots, m \\ & \quad H_i(\mathbf{y}) = 0 \quad \text{for } i \in 1, \dots, p, \end{aligned}$$

which is a convex program that can be efficiently solved using, for example, primal-dual interior-point methods; see Dahl and Andersen (2021) for more details. In practice, geometric programs with tens of thousands of decision variables and constraints can be solved to find the global optimum in a matter of seconds on a standard laptop computer.

3. Model

We now describe the epidemiological model under analysis. We consider a region with N individuals and propose a population model with *homogeneous mixing*, i.e., every pair of individuals come in contact with a probability that depends on aggregated mobility variables. Population models are commonly used in the absence of granular data on the network of social contacts in a region (Nowzari et al., 2016). We assume that each individual can visit POIs belonging to different categories $1, 2, \dots, K$; let $\mathbf{m}(t) = (m_1(t), m_2(t), \dots, m_K(t))$ denote a vector of human mobility variables capturing the percentage change in volume of visits to each of those categories at a particular discrete time t




Fig. 1. Illustration of the epidemic model under consideration.

(e.g., days) relative to a pre-established baseline; hence, $m_k(t) \in [0, \infty)$. In particular, we use publicly available mobility data from Google's COVID-19 Community Mobility Reports (Google, 2020) which capture daily changes in visitation patterns to public places, for example, Retail & recreation, Grocery & pharmacy, and Parks, as well as time spent at work. For each category, Google reports the relative change in visits compared to a baseline of mobility measured before the lockdown measures took place. This baseline corresponds to the median daily visits to each category over the period comprising January 3 through February 6, 2020.

Our model consists of two layers: a *mobility layer* and an *epidemic layer*. The mobility layer captures the effect of human mobility on the spread of the disease and influences the dynamics taking place on the epidemic layer. In the epidemic layer, we consider an extension of the classic SEIR epidemic model (Brauer, Castillo-Chavez, & Feng, 2019; Martcheva, 2015) which explicitly accounts for asymptomatic hosts and hospitalizations, as in Giordano et al. (2020). Each of the individuals in a region belongs to one of seven possible compartments described below. In this model, $\tilde{S}(t)$ represents the number of individuals susceptible to becoming infected at a discrete time t . In our optimal control problems, we will consider a finite horizon T_c over which we can assume an almost constant number of susceptible individuals; hence, in these problems we assume $\tilde{S}(t) \approx S_0$ for all $0 \leq t \leq T_c$, which linearizes the model. The variable $\tilde{E}(t)$ represents the number of individuals who have contracted the virus (exposed) but are in an incubation period at time t . After the incubation period, agents can move to one of the infectious compartments; $\tilde{I}(t)$ represents the number of symptomatic individuals and $\tilde{A}(t)$ represents the number of asymptomatic individuals. The asymptomatic compartment is included since asymptomatic individuals play a crucial role in the spread of COVID-19, with transmission rates that are different from symptomatic individuals (Giordano et al., 2020; Gandhi, Yokoe, & Havlir, 2020). Asymptomatic individuals eventually recover on their own and move on to the recovered compartment, represented by the variable $R(t)$. Symptomatic individuals can recover on their own, or their symptoms can worsen and they subsequently require hospitalization, in which case they are moved into a hospitalized compartment, represented by the variable $\tilde{H}(t)$. Since hospital capacity is a principal concern with the treatment of COVID-19, we include this compartment to constrain the control problems described in Section 5. In particular, our mobility-based control input will be constrained to prevent a hospital capacity overflow. Finally, agents that are hospitalized may either recover and transition to the compartment represented by $\tilde{R}(t)$ or may die, and subsequently transition to the compartment represented by $\tilde{D}(t)$. With explicit data on the number of deaths in every region, we may train our model to predict the population of this compartment. We make the simplifying assumption that only individuals with severe symptoms are at a risk of dying and, hence, all of them have been previously hospitalized.

All parameters related to the dynamics of this model are summarized in Table 1. Using these parameters, the discrete-time evolution of

¹ In particular, we remark that the globally optimal solution to this convex program corresponds to the globally optimal solution of the original GP.

the number of individuals in each compartment, illustrated in Fig. 1, is given by:

$$\tilde{S}(t+1) = \tilde{S}(t) - \tilde{S}(t)\beta(t)(\gamma_A \tilde{A}(t) + \tilde{I}(t)), \quad (1)$$

$$\tilde{E}(t+1) = (1 - \rho_{EI} - \rho_{EA}) \tilde{E}(t) + \tilde{S}(t)\beta(t)(\gamma_A \tilde{A}(t) + \tilde{I}(t)), \quad (2)$$

$$\tilde{I}(t+1) = (1 - \rho_{IR} - \rho_{IH}) \tilde{I}(t) + \rho_{EI} \tilde{E}(t), \quad (3)$$

$$\tilde{A}(t+1) = (1 - \rho_{AR}) \tilde{A}(t) + \rho_{EA} \tilde{E}(t), \quad (4)$$

$$\tilde{H}(t+1) = (1 - \rho_{HR}) \tilde{H}(t) + \rho_{IH} \tilde{I}(t), \quad (5)$$

$$\tilde{R}(t+1) = \tilde{R}(t) + \rho_{IR} \tilde{I}(t) + \rho_{AR} \tilde{A}(t) + (1 - \alpha_D) \rho_{HR} \tilde{H}(t), \quad (6)$$

$$\tilde{D}(t+1) = \tilde{D}(t) + \alpha_D \rho_{HR} \tilde{H}(t). \quad (7)$$

In our model, we assume that susceptible individuals can transition into the exposed compartment when in contact with either Infected (symptomatic) or Asymptomatic individuals. We assume that the rate at which asymptomatic individuals infect others is weighted by a constant (unknown) parameter γ_A . We also assume that the portion of hospitalized individuals who die, relative to those that recover, is equal to an unknown constant α_D . In Section 4, we will introduce a methodology to learn these (and other) unknown parameters in our model. As mentioned above, this model is intended to solve optimal control problems over a finite time horizon T_c over which the number of new infected individuals is small compared to the entire population, so we can assume that the number of susceptible individuals at any time $0 \leq t \leq T_c$, $S(t)$, is well approximated with a constant; hence, in our control problems we set $\tilde{S}(t) = S_0$. Moreover, this assumption linearizes² the dynamics of the states; for notational clarity, we will omit the tildes when considering the states of the linear model. As we will see in Section 5, this linearization ensures that the entries of the state vector at any given time are posynomials on the parameter $\beta(t)$. However, we incorporate a non-linear dependency of the parameter $\beta(t)$ on the mobility restriction variables, which we will use as our external control variable, rendering the resulting model non-linear and multiplicative in the control input. Fortunately, the states of this linearized model upper-bound the states of the true model, as shown in the lemma below. Appendix A.2 contains simulation results to examine the tightness of these bounds in practice. Hence, by reducing the number of deaths in the linearized model, $D(t)$, we are guaranteed to reduce the number of deaths in the true nonlinear model, $\tilde{D}(t)$.

Lemma 1. *The states of the linearized model (wherein $S(t) = S_0$) upper-bound the states of the model in (1)–(7). Thus, for all $t > 0$,*

$$\tilde{S}(t) \leq S(t) = S_0, \quad \tilde{E}(t) \leq E(t), \quad \tilde{I}(t) \leq I(t),$$

$$\tilde{A}(t) \leq A(t), \quad \tilde{H}(t) \leq H(t), \quad \tilde{R}(t) \leq R(t), \quad \tilde{D}(t) \leq D(t).$$

Proof. See Appendix B.1. \square

The mobility layer of our model incorporates the effects of non-pharmaceutical interventions, such as social distancing and other forms of mobility restrictions, which a decision maker may employ to curb the spread of the epidemic. By reducing human mobility, a decision maker induces fewer contacts between susceptible and infected individuals and, thus, reduces the risk of infection. In particular, we relate the infection rate $\beta(t)$ to a time series $\mathbf{m}(t)$ of human mobility variables by means of an unknown function $f(\cdot)$. We choose f to be a parametric function whose parameters will be learned from data (described in detail in Section 4.1).

In order to employ non-pharmaceutical interventions, a decision maker designs a mobility control strategy to set the human mobility variables $\mathbf{m}(t)$ for some finite horizon $t \in \{0, \dots, T_c\}$. In mathematical

Table 1

Summary of parameters in epidemic model. We assume $\rho_{IR} + \rho_{IH} < 1$, $\rho_{EI} < 1$, $\rho_{AR} < 1$, and $\rho_{HR} < 1$, since these parameters capture the fraction of individuals in each compartment that transition to other compartments, which must be less than one for the model to be meaningful. We also assume $\gamma_A \in [0, 1]$, reflecting the fact that asymptomatic individuals are less likely to spread the disease.

Parameter	Description
$\beta(t)$	Rate at which susceptible individuals become infected due to contacts with infectious individuals at time t ; $\beta(t)$ is a function of the mobility variables $\mathbf{m}(t)$ at time t
γ_A	Weight representing lower risk of infection when infectious individuals are asymptomatic
ρ_{EI}	Rate at which exposed individuals become symptomatic
ρ_{EA}	Rate at which exposed individuals become asymptomatic
ρ_{IR}	Rate at which symptomatic individuals recover on their own
ρ_{IH}	Rate at which symptomatic individuals develop severe symptoms and become hospitalized
ρ_{AR}	Rate at which asymptomatic individuals recover on their own
ρ_{HR}	Rate at which hospitalized individuals recover
α_D	Proportion of hospitalized individuals that die

terms, the decision maker designs an input $\{\mathbf{u}(t)\}_{t=0}^{T_c}$ which affects future values of the mobility variables. For simplicity, we assume an identity mapping between mobility and the input, so that $\mathbf{m}(t) = \mathbf{u}(t)$, and $\mathbf{u}(t)$ is in some set of admissible actions \mathcal{U} . In other words, we assume that mobility patterns may be directly actuated to a desired level by a decision maker, subject to feasibility constraints.

Intuitively, lower values of $\mathbf{u}(t)$ correspond to more restrictions on human mobility. The decision maker may have fine-grained control over their control strategy, for example by closing individual establishments, imposing occupancy limits, or restricting hours of operation, and as such we treat the individual components of the control action $u_k(t)$ as continuous variables. Moreover, some categories may have different admissible actions, e.g., it may not be possible to close down all pharmacies but closing all gyms is reasonable. In mathematical terms, we will consider a set of allowable control actions \mathcal{U} that is described using posynomial inequalities and monomial equalities.

Furthermore, implementing mobility restrictions in this manner cannot be done without incurring a financial loss. Closure of businesses causes economic losses, which need to be taken into consideration when selecting an appropriate control strategy. In particular, applying the temporal control strategy $\mathbf{u}(t)$ of mobility restrictions incurs a cost $C_t(\mathbf{u}(t))$ which we assume to be monotonically decreasing with \mathbf{u} and convex in log-scale, reflecting that the costs on society of restricting mobility are marginally increasing.

4. Parameter estimation via multitask learning

Several recent epidemic prediction methods opt to set some (or all) of the parameters in their models to estimations from the medical and virology literature (e.g. Achterberg et al., 2020; Aleta et al., 2020; Bhouri et al., 2021; Birge et al., 2020; Chang et al., 2021). However, these parameters often have wide confidence intervals and are commonly inferred from statistical models that do not take into account the effects of social distancing and hospital capacity (Ali et al., 2020). Other recent works learn these parameters from data (Bhouri et al., 2021; Van den Broeck et al., 2011; Chang et al., 2021; Hota et al., 2020), but do not explicitly model the infectivity of the epidemic using mobility patterns, making the design of related control strategies difficult. In contrast to these approaches, our model is entirely data-driven in that all parameters used (including initial conditions) are learned directly from data, and we learn an explicit mapping between mobility patterns and the spread of the epidemic. In particular, we employ a multitask learning approach (Caruana, 1997) by leveraging both data on the number of deaths across a set of regions, and mobility data describing how often individuals in a region visit different points of interest. The key idea is to consider each region of interest as

² In general the resulting dynamics are bilinear due to the terms $\beta(t)\tilde{A}(t)$ and $\beta(t)\tilde{I}(t)$, but in the original SEIR model $\beta(t) = \beta \forall t$, rendering these dynamics linear in the states. For this reason we refer to these dynamics as linear.




Fig. 2. The learning pipeline for a given region. Global parameters, learned using all available data, are shown in magenta; initial conditions, learned for each region, are shown in blue; mobility mapping parameters, again different across regions, are shown in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

a separate prediction subtask, while many parameters of the models describing the evolution of the epidemic are shared across regions. Indeed, certain parameters of the epidemic model are intrinsic to the disease; hence, they do not depend on the geographical location from which data is collected. As such, when calibrating this model to a given region, we can benefit from the data from other regions by employing a multitask learning framework, providing better parameter estimates and avoiding overfitting. Towards this goal, we pool data from multiple regions (e.g., US counties) and minimize a global cost function in which the global parameters are shared across regions but the mobility parameters and initial conditions are specific to each region. This provides *statistical data amplification*, as the prediction subtask for each county is given additional data with which to learn the global epidemic parameters. Another relevant facet of multitask learning is *representation bias*; the learning procedure will avoid local minima intrinsic to only a few (or just one) of the counties. The learning pipeline from data to predictions for a single region is shown in Fig. 2. We fit our models using publicly available mobility data from Google's COVID-19 Community Mobility Reports (Google, 2020) which captures daily changes in visitation patterns to public places, for example, Retail & Recreation, Grocery & Pharmacy, and Parks, as well as time spent at Workplaces. This dataset is organized as different time series for six different categories. For each category, Google reports the relative change in visits compared to a baseline of mobility measured before the lockdown measures took place. This baseline corresponds to the median daily visits to each category over the period comprising January 3 through February 6, 2020. Furthermore, we use public data from The New York Times, based on reports from state and local health agencies (New York Times, 2020), consisting of daily and cumulative caseloads and deaths attributed to COVID-19 in The United States. To account for inconsistencies and lags in reporting, we compute a 7-day rolling average on the original time series for the calibration of our model.

As mentioned above, there are two layers to our model, namely (1) the *mobility layer*, which is a mapping from mobility data to the infection rate $\beta(t)$ (described in Section 4.1) and (2) the *epidemic layer*, which describes the dynamics of the disease itself (discussed in Section 4.2). Since parameters such as the latency period, ratio of infected individuals who develop symptoms, and case fatality ratio are intrinsic to the disease and should not vary greatly based on the geographical area being studied, we group these together across regions as global parameters and learn them jointly with all the available data. However, the mapping from mobility data to the infection rate and initial conditions of the regional epidemic are dependent on the locality, and thus they are learned using only the data from their region.

4.1. The mobility layer

To learn the function $f : \mathbf{m} \mapsto \beta$ from mobility data to the infection rate, we must first select an appropriate class of functions for such a

mapping. Although we could use any parametric family of functions, such as neural networks, to estimate f , not all choices are tractable. In particular, neural networks may provide great prediction performance but would render an intractable control problem. In order to obtain a tractable control problem, we chose to model the function f using a parametric posynomial function (Boyd et al., 2007). As we will show in Section 5, this choice allows us to use geometric programming to efficiently solve several nonlinear optimal control problems of interest. Thus, we model f in a parametric way as

$$f(\mathbf{m}; \theta, \alpha, b) = \sum_{k=1}^K \theta_k m_k^{\alpha_k} + b, \quad (8)$$

where we recall that $m_k \in [0, \infty)$. From a practical standpoint, this posynomial approximation is justified because β can be viewed as the product of the contact rate (the expected number of contacts an individual has with others) and the transmission risk, which is constant over time. Moreover, the number of contacts within a category should exponentially increase with the number of visits to points of interest in that category. Since the mobility data is stratified across K different categories, we allow the parameters to be different across the categories. Thus, in the parametric function $f(\mathbf{m}; \theta, \alpha, b)$ the probability of transmission is captured by $\theta = (\theta_1, \dots, \theta_K) \in \mathbb{R}_{\geq 0}^K$, the exponential growth of infectivity is captured by $\alpha = (\alpha_1, \dots, \alpha_K) \in \mathbb{R}^K$, and the bias term b accounts for potentially unmodeled infections.

Since susceptible individuals may become infected by either symptomatic or asymptomatic individuals, the mobility mapping f is incorporated into the epidemic layer in two terms, as seen in (2). Firstly, it is used to model new infections from symptomatic individuals via the term $\beta(\mathbf{m}(t))S_0 I(t)$; secondly, we include a weighting term γ_A in the term $\gamma_A \beta(\mathbf{m}(t))S_0 A(t)$ to model the rate of new infections from asymptomatic infectious individuals. Thus, altogether we denote the set of parameters corresponding to the mobility mapping function for region i as $\Psi_{mobility}^i := \{\theta, \alpha, b, \gamma_A\}$.

4.2. The epidemic layer

Our model is a latent-space model; hence, the states are not fully observable. Following common practice with these models, we treat the unobserved initial conditions as unknown parameters to be identified from the data. Since the dynamical trajectories of the epidemic are different across regions (e.g., US counties), for each region i we learn a set of initial conditions $\Psi_0^i := \{S_0^i, \dots, R_0^i\}$. As mentioned previously, we follow a multitask learning approach and, thus, the remaining global parameters, which we assume to be intrinsic to the disease, are shared across all regions and learned collectively using data available from all regions. The set of global parameters is denoted by $\Psi_{global} := \{\rho_{EA}, \rho_{PI}, \rho_{PAR}, \rho_{PIH}, \rho_{PIR}, \rho_{PHR}, \alpha_D\}$, which includes all clinical parameters that depend on the nature of the virus alone (i.e., they are not influenced by social mobility).




Fig. 3. Example of predicted deaths for several US counties, trained on 90 days of mobility and death count data, and tested on 30 days of the same. Predictions are shown in blue, and the rolling seven-day averages of real deaths data are shown as dashed lines. The area in white denotes training data, and the area in red denotes test data.

4.3. Simulated predictions for Philadelphia and surrounding counties

In order to validate the predictive accuracy of our data-driven model, we conducted a case study using counties from the Delaware Valley metropolitan statistical area (commonly known as the *Philadelphia metropolitan area*), which includes Philadelphia County and nine surrounding regions, including counties in Maryland, New Jersey, and Delaware. Using county-level data, we learned both the local mobility mapping functions and initial conditions, as well as the global clinical parameters of the epidemic. Due to the known inconsistencies and lags in reporting of cases and deaths, we use the rolling seven-day average of cumulative deaths for both training and prediction. Formally, if we have training data for M counties over T_d days, the training loss for the set of parameters $\Psi := \Psi_{\text{global}} \cup_{i=1}^M \Psi_{\text{mobility}}^i \cup_{i=1}^M \Psi_0^i$ is the following mean-squared error loss function:

$$\ell_{\text{train}}(\Psi) = \frac{1}{MT_d} \sum_{i=1}^M \sum_{t=1}^{T_d} \left(\frac{X^i(t)}{N_i} - \frac{\hat{X}^i(t; \Psi)}{N_i} \right)^2, \quad (9)$$

where N_i is the population of region i , $X^i(t)$ is the measured rolling 7-day average of cumulative deaths on day t for region i , and $\hat{X}^i(t; \Psi)$ is the predicted value for the same region for a set of parameters Ψ . We normalize the number of deaths by the population in each county to avoid biasing our predictions towards counties with a larger population. Notably, any differentiable loss function may be substituted in this approach; mean-squared error was chosen for simplicity.

Our model was trained by computing gradients of the loss function ℓ_{train} with respect to all parameters in Ψ via the automatic differentiation package `autograd` (Maclaurin, Duvenaud, & Adams, 2015) and running stochastic gradient descent using `adam` (Kingma & Ba, 2014) over several independent trials using different initial guesses to account for the non-convexity of the training problem. In particular, `autograd` performs reverse-mode differentiation (i.e., backpropagation) to compute numerical gradients; hence, differentiability of the loss function ℓ_{train} and the dynamical equations governing our system help to ensure the existence of such gradients. Global clinical parameters were initialized by using plausible values from the medical literature (Bhouri et al., 2021; Day, 2020; He et al., 2020; Lauer et al., 2020; Nishiura et al., 2020; Pei, Kandula, & Shaman, 2020; Rahmandad, Lim, & Sterman, 2020; Woelfel et al., 2020), while local parameters are initialized randomly in each trial. Different counties are split across random batches, allowing the global parameters (trained using all data

in our multitask approach) to converge more quickly, which in turn allows the local parameters to converge to values which fit more closely with the global parameters. From these trials, we select the set of parameters with lowest testing error, and present some examples of predictions using these chosen parameters in Fig. 3.

5. Optimal control using geometric programming

Traditional optimal control techniques are not directly applicable to compartmental epidemic models for a number of reasons. First, epidemic models are typically nonlinear and the effect of non-pharmaceutical interventions is not additive, but multiplicative. Therefore, standard techniques such as LQR cannot be readily employed. Furthermore, a direct application of Pontryagin's Maximum Principle results in a high-dimensional two-point boundary value problem for which numerical methods have no convergence guarantees and present scalability issues. Due to the modeling choices proposed in this paper, we obtain a mobility-driven epidemic dynamics amenable to solve certain optimal control problems using tools from geometric programming (Boyd et al., 2007). In particular, we can solve the data-driven optimal control problems aiming to minimize the final number of deaths while respecting budget constraints on the economic costs associated with implementing mobility restrictions, as well as avoiding hospital overflows, with guarantees of global optimality. In practice, all people may not exactly follow the desired mobility restrictions; thus, such a strategy may be implemented in receding horizon. For example, if a decision maker sets mobility restrictions on a weekly basis, they may plan ahead for several months, implement a week's restrictions, then re-plan the following week and implement the resulting strategy, and repeat.

In order to employ geometric programming in this setting, it is necessary that the states can be expressed as posynomial functions of the control input. The negative term in (1) bars us from directly using our model; however, the horizon $0 \leq t \leq T_c$ over which our control problem takes place is small enough that the number of new infected individuals is small relative to the total population, and so we may assume a nearly constant number of susceptible individuals exist, i.e., $\tilde{S}(t) \approx S_0$. This assumption linearizes our model, and allows us to express the states as posynomial functions of the control input; hence, we will consider this linearized model for the remainder of this section. Fortunately, this linearization has a modest effect on our predictions;

indeed, this linearized model upper-bounds the states of the true model, as illustrated in [Lemma 1](#).

We start our analysis by considering the minimization of the final number of deaths. Since we will minimize deaths in the linear model, $D(t)$, we remark by [Lemma 1](#) that we induce conservatism in our control framework by minimizing an upper bound on the number of deaths in the nonlinear model, $\tilde{D}(t)$. For completeness in our analysis, we allow the inclusion of a daily discount factor γ_D and a terminal cost γ_∞ on the number of deaths at the end of the horizon T_c . Hence, the decision maker aims to minimize the following objective function:

$$J := \sum_{t=1}^{T_c-1} \gamma_D^t \alpha_D \rho_{HR} H(t) + \gamma_\infty \alpha_D \rho_{HR} H(T_c). \quad (10)$$

We remark from [\(7\)](#) that the incident deaths on day t , i.e., the number of new deaths $D(t) - D(t-1)$, is exactly $\alpha_D \rho_{HR} H(t)$. The discount factor $\gamma_D \in (0, 1]$ is motivated by the uncertainty of deaths in the future which might be prevented by interventions not available in the present day. For example, the probability that a vaccine is widely available in the future increases as time passes. In particular, $1 - \gamma_D$ can be considered the probability of a vaccine being widely available on each day in the future, so the probability of no vaccine being widely available by day t is γ_D^t and, hence, deaths predicted at day t are only accounted for if there is not a vaccine that could prevent them. The terminal cost γ_∞ illustrates the desire to keep the number of deaths low beyond the time horizon in consideration. For example, let us assume that beyond the time horizon T_c the epidemic has been curbed and the number of new daily deaths falls below $\alpha_D \rho_{HR} H(T_c)$. In the worst case scenario we would have $H(t) = H(T_c)$ for $t \geq T_c$, hence,

$$\sum_{t=T_c}^{\infty} \gamma_D^t H(t) = H(T_c) \frac{\gamma_D^{T_c}}{1 - \gamma_D}.$$

Defining $\gamma_\infty := \gamma_D^{T_c} / (1 - \gamma_D)$, and $\gamma_\infty = 0$ if $\gamma_D = 1$, the discounted number of deaths beyond T_c is given by the terminal cost $\gamma_\infty \alpha_D \rho_{HR} H(T_c)$.

Given that the infection rate $\beta(t)$ depends on mobility, we assume that a decision maker can restrict mobility dynamically to curb the number of deaths by designing a mobility strategy $\mathbf{u}(t)$ so that $\beta(t) = f(\mathbf{u}(t))$. Furthermore, we assume that $\mathbf{u}(t)$ is constrained to be within a set \mathcal{U} reflecting that essential businesses cannot be severely restricted and that some mobility restrictions are only partially effective.

These mobility restrictions incur a cost which could be measured in terms of a pecuniary cost to the economy, absolute number of visits lost by businesses, or impact on the utility of citizens. In our framework, we quantify the economic cost of imposing a mobility control strategy $\mathbf{u}(t)$ using a cost function $C_t(\mathbf{u}(t))$ which, in general, can be time-varying; hence, for example, we can use different costs for mobility restrictions on workdays and weekends. Moreover, such a cost function may incorporate terminal costs to account for economic losses beyond the time horizon T_c . We choose to model $C_t(\mathbf{u}(t))$ as a posynomial on $\mathbf{u}(t)$, since this is amenable to a geometric programming approach. We investigate the problem of choosing an optimal mobility control strategy $\mathbf{u}^*(t)$ that minimizes the number of cumulative deaths while keeping the total cost of the intervention, given by $\sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t))$, below a pre-specified budget \mathcal{B} , while obeying a limit on hospitalizations τ_H .

As we will show, this problem can be expressed as a geometric program; this stems from the fact that the states $H(t)$ and $D(t)$ can be expressed as posynomial functions of the mobility control variables $\mathbf{u}(t)$, as shown in the following lemma.

Lemma 2. *The functions $H(t)$ and $D(t)$, representing the number of hospitalized individuals and deaths at time t , are posynomial functions of the entries of $\mathbf{u}(t)$ for $t = 0, 1, \dots, T_c$.*

Proof. See [Appendix B.2](#). \square

Since a positively weighted sum of posynomials is also a posynomial, we obtain our main result below.

Theorem 1. *If $C_t(\mathbf{u}(t))$ is a posynomial cost function for all t , and the set of admissible control actions \mathcal{U} is described by posynomial inequalities and monomial equalities, then the following is a geometric program:*

$$\begin{aligned} & \underset{\mathbf{u}(0), \mathbf{u}(1), \dots, \mathbf{u}(T_c-1)}{\text{minimize}} && \sum_{t=1}^{T_c-1} \gamma_D^t H(t) + \gamma_\infty H(T_c) \\ & \text{subject to} && \sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t)) \leq \mathcal{B} \\ & && H(t) \leq \tau_H, \quad t = 1, \dots, T_c, \\ & && \mathbf{u}(t) \in \mathcal{U}, \quad t = 0, \dots, T_c - 1. \end{aligned} \quad (11)$$

Proof. See [Appendix B.3](#). \square

The choice of the cost function and the budget \mathcal{B} may be difficult to discern. It may be unclear how reductions of mobility in a given category impact the economy, and the choice of a budget for those mobility restrictions engenders an implicit trade-off between minimizing deaths and economic costs. The choice of these cost functions is up to the decision maker and, in practice, could leverage available economic data about employment, revenue losses and social-wellness. To choose such a budget \mathcal{B} once a cost function has been defined, we propose a principled approach to obtain the minimal budget under which hospitals remain below capacity, preventing the steep increase in deaths due to an overwhelmed healthcare system. This is achieved by solving an auxiliary optimal control problem that aims to find the minimum cost required to keep the number of hospitalized individuals $H(t)$ below a threshold τ_H at all times. This auxiliary optimal control problem is also a geometric program as long as the cost functions $C_t(\mathbf{u}(t))$ are posynomials (for example, found via posynomial fitting ([Boyd et al., 2007](#)) on economic data); this is stated in [Theorem 2](#) and the proof follows from [Lemma 2](#).

Theorem 2. *If $C_t(\mathbf{u}(t))$ is a posynomial cost function for all t , and the set of admissible control actions \mathcal{U} is described by posynomial inequalities and monomial equalities, then the minimal budget required to keep hospitalizations below a given threshold τ_H is given by*

$$\mathcal{B}^* = \sum_{t=0}^{T_c-1} C_t(\mathbf{u}^*(t)), \quad (12)$$

where \mathbf{u}^* is the solution to the geometric program

$$\begin{aligned} & \underset{\mathbf{u}(0), \dots, \mathbf{u}(T_c-1)}{\text{minimize}} && \sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t)) \\ & \text{subject to} && H(t) \leq \tau_H, \quad t = 1, \dots, T_c, \\ & && \mathbf{u}(t) \in \mathcal{U}, \quad t = 0, \dots, T_c - 1. \end{aligned} \quad (13)$$

Proof. See [Appendix B.4](#). \square

The budget \mathcal{B}^* obtained from [Theorem 2](#) can be seen as a conservative cost which only guarantees that hospital operations remain within capacity, avoiding overflow. The decision maker should then use a budget $\mathcal{B} \geq \mathcal{B}^*$ in implementing the results from [Theorem 1](#) to obtain a less conservative control input $\mathbf{u}(t)$.

5.1. Practical considerations

As mentioned previously, in practice some individuals may not follow guidance regarding mobility restrictions; thus, our control strategy may be implemented in receding horizon. Since a decision maker may wish to set mobility restrictions on a rolling basis (e.g., weekly), they may plan ahead for several months, implement restrictions within the




Fig. 4. Minimal-cost control strategy $\mathbf{u}^*(t)$ for Philadelphia County, PA, obtained by solving the geometric program (13). For comparison, we plot the true mobility data $\mathbf{m}(t)$ over the same time period in (a). As shown in Lemma 1, the number of hospitalized individuals in the linear model, $H(t)$, upper-bounds the number of hospitalized individuals in the nonlinear model, $\hat{H}(t)$. We also include the hospitalizations predicted based on the true mobility data $\mathbf{m}(t)$ as a baseline. Parameters of the models used herein were learned as described in Section 4.3. We obtain the budget \mathcal{B}^* to use in geometric program (11) by taking the total cost of the minimal-cost control strategy, i.e., $\mathcal{B}^* = \sum_{t=0}^{T_c-1} C_t(\mathbf{u}^*(t))$. Control actions in (b) for categories except transit stations are similar, and are overlaid.

time frame of interest, then re-plan at the end of the time period, and repeat. Since the decision maker may tune the different categories of mobility variables independently, we allow each of the components $k = 1, \dots, K$ of the mobility control action $\mathbf{u}(t)$ to vary between a lower bound $\underline{u}_k > 0$ (representing full lockdown) and an upper bound \bar{u}_k (representing no restrictions at all). Moreover, to reflect the fact that decision makers want to avoid dramatic changes to restrictions on a consistent basis, we enforce that the policy must remain fixed for periods of one week (seven days), and that each component of mobility may change no more than $\Delta\%$, i.e., $(1-\Delta)u_k(t-1) \leq u_k(t) \leq (1+\Delta)u_k(t-1)$, $\forall k, t \in \{1, \dots, T\}$. However, it may be the case that categories of mobility are dependent, and cannot be tuned arbitrarily; we explore this idea in detail in Appendix A.1. These conditions lead to the set of admissible control actions

$$\mathcal{U} = \{ \mathbf{u}_1(0), \dots, \mathbf{u}_K(0), \dots, \mathbf{u}_1(T_c), \dots, \mathbf{u}_K(T_c) \in \mathbb{R}^{K \times T_c} : \mathbf{u}_k(t) \in [\underline{u}_k, \bar{u}_k] \cap [(1-\Delta)u_k(t-1), (1+\Delta)u_k(t-1)], \mathbf{u}(s) = \mathbf{u}(s+1) = \dots = \mathbf{u}(s+6), s \in \{0, 7, \dots\} \}. \quad (14)$$

In practice the cost function $C_t(\mathbf{u}(t))$ may be supplied by a decision maker, or may be found via posynomial fitting (Boyd et al., 2007) using economic data from a region. In the absence of such data, we choose a time-invariant cost function $C(\mathbf{u}(t))$ that satisfies the requirements of being convex in log-scale and decreasing, given by

$$C(\mathbf{u}(t)) = \sum_{k=1}^K c_k \frac{u_k(t)^{-1} - \bar{u}_k^{-1}}{\underline{u}_k^{-1} - \bar{u}_k^{-1}}, \quad (15)$$

where $\mathbf{c} = (c_1, \dots, c_K)$ is a relative cost weighting of the mobility categories. This relative cost could reflect that visits lost to, for example, healthcare facilities are more costly than visits lost to retail venues. As

shown in the lemma below, this selection of cost function and set of allowable control actions renders (11) and (13) geometric programs.

Lemma 3. *The set of admissible control actions in (14), where $0 < \underline{u}_k < \bar{u}_k$ for all $k \in \{1, \dots, K\}$, $\Delta > 0$, and the cost function in (15), where $c_k \geq 0$ for all $k \in \{1, \dots, K\}$, are sufficient for (11) and (13) to be geometric programs.*

Proof. See Appendix B.5. \square

5.2. Control simulations

We demonstrate the effectiveness of our control approach with a case study for the counties in the greater Philadelphia area, illustrated in Fig. 4 and Fig. 5. All simulations were performed on a laptop computer with 16GB of RAM and a 2.2 GHz Intel Core i7 CPU. The parameters of our compartmental model as well as the mobility mapping $\beta(\mathbf{u}(t))$ are learned from data as described in Section 4.3 using the proposed multitask learning framework. Our models are trained using recent data before the widespread usage of vaccines, from August 1st to October 31st, 2020, and tested from November 1st–30th, 2020. In particular, in Fig. 5(c) we illustrate the number of cumulative deaths predicted using the optimal control strategy $\mathbf{u}^*(t)$ as compared to the number of cumulative deaths predicted based on the true mobility data $\mathbf{m}(t)$ from Philadelphia. We invoke Lemma 3 to select the cost function $C_t(\mathbf{u}(t))$ and set of allowable control actions \mathcal{U} . To elucidate the set of allowable control actions \mathcal{U} , we select the values \underline{u}_k and \bar{u}_k independently for each category based on the mobility data used in our multitask learning framework. We set the limit on the relative change in mobility, Δ , to 10%. We further re-scale the mobility data to be in




Fig. 5. Optimal minimal-death control strategy $u^*(t)$ for Philadelphia County, PA, obtained by solving the minimal-death GP (11), from August 1st to November 30th, 2020. In Figure (c) we compare the cumulative deaths predicted using the optimal control strategy $u^*(t)$ in the nonlinear model, $\tilde{D}(t)$ (blue line with circle markers), cumulative deaths predicted in the linear model, $D(t)$ (red line with star markers), and the cumulative deaths predicted based on the true mobility data $m(t)$ as a prediction baseline (green line with cross markers). Parameters of the models used herein were learned as described in Section 4.3, with the budget B^* taken from the solution to the minimal-cost GP (13). Control actions in Figure (a) for categories except transit stations are similar, and are overlaid.

$[0, 1]^K$, which has been shown to speed up convergence of learning methods (Bishop, 2006). For simplicity, we assign equal costs to each category of mobility, so that $c_k = 1$ for all $k \in \{1, \dots, K\}$.

We show the minimal-cost control action $u^*(t)$ computed by solving (13) in Fig. 4. As mentioned previously, interior point methods solve these problems very efficiently; full discussions on the computational complexity of our approach are in Appendix A.3. Fig. 4(a) illustrates the measured mobility data $m(t)$ in Philadelphia County in order to compare with the designed control strategies. In Fig. 4(b) we show the individual mobility control strategies for each category; recall that a lower value of $u_k^*(t)$ corresponds to a lower value of mobility allowed for category k , which incurs a higher cost. In order to trade off between reduction in infections and economic costs, mobility in parks is allowed to be larger than all other categories. Mobility is kept consistently low across the time horizon to keep the number of hospitalized individuals under control, but allowed to increase later on, reducing the overall cost as an economic trade-off. In Fig. 4(d) we see the result of Lemma 1, namely that the number of hospitalizations in the linear model $H(t)$, which is used in the GP (13), upper-bounds the hospitalizations in the nonlinear model $\tilde{H}(t)$. Indeed, we see that our GP is providing a conservative estimate for the control action $u^*(t)$: even though $H(t)$ is near the hospitalization threshold τ_H , the predicted hospitalizations under the nonlinear model $\tilde{H}(t)$ are approaching zero at the end of the time horizon; we explore this phenomenon further in Appendix A.2.

In Fig. 5 we show the results of Theorem 1 by solving the GP (11). While this problem is related to the GP (13), the objective in this case is to minimize deaths subject to the budget B^* computed by solving (13). Hence, while the mobility control strategy $u^*(t)$ and corresponding cost are similar to what is shown in Fig. 4, there is a sharper initial restriction to mobility which is lifted in a more gradual manner. Fig. 5(c) shows the number of cumulative deaths using the mobility control action $u^*(t)$ in blue, as compared to the deaths incurred without any control, i.e., with the measured mobility pattern data $m(t)$. To confirm the accuracy of these predictions, we also show the number of deaths recorded in the same time horizon as the dashed black line. Indeed, by solving the GP (11) and implementing the associated control strategy, our model predicts over 200 lives could have been saved between August 1st and November 30th. As before we show the predictions using the linear model, which upper-bound the predictions in the nonlinear model; this further illustrates the conservatism of our control approach.

6. Conclusion and future work

In this paper we presented a multitask learning and nonlinear optimal control framework that aims to bridge the gap between optimal control theory of epidemic models and applicable data-driven models for analyzing the spread of COVID-19 and other future epidemics. To

identify the parameters of our model, we propose a multitask learning approach that leverages mobility and epidemic data from multiple regions to capture how daily changes in mobility patterns affect the spread of the disease, and to accurately predict the resulting daily and cumulative deaths. Using this data-driven model we present a nonlinear optimal control framework using geometric programming to efficiently design non-pharmaceutical interventions to limit the spread of the epidemic while obeying a budget constraint on the economic loss incurred. Furthermore, we present a principled method for determining such a budget based on eliminating excess deaths due to over-utilization of hospital resources. We validate both our model and our control framework in a case study on the greater Philadelphia area.

In the future, this work could be extended to accommodate for robustness considerations as well as stochastic transitions in the epidemic layer, which introduce the additional challenge of expressing chance constraints as posynomial functions. The success of geometric programming in our work comes from expressing states of the system as posynomials on the mobility variables, allowing for the potential extension to models with more sophisticated mappings from human-mobility to epidemic dynamics. For example, generalized geometric programming admits functions that are max-monomials or posynomials with fractional exponents (Boyd et al., 2007), which opens the door to modeling epidemic dynamics and human-mobility using ReLU neural networks or posynomial approximations to arbitrary functions. Furthermore, since geometric programs can be solved efficiently, our approach could be applied to models with higher complexity; for example, networked metapopulation models which can make use of more granular datasets.


Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Additional practical results

A.1. Dimensionality reduction

While mobility data may be expressed in terms of several different categories, in practice these categories are highly correlated; the correlation between mobility categories in Philadelphia is shown in Fig. A.6(a). As such, it may not be reasonable to assume that a decision maker may tune these categories independently and arbitrarily. However, it may also be difficult to elucidate the explicit relationships present from data alone. As a compromise, we may perform dimensionality reduction on the measured mobility pattern data via principal



(a) Pairwise correlation coefficients between different categories of mobility data.
 (b) Percentage of explained variance in principal component analysis.

Fig. A.6. Exploration of mobility pattern data in Philadelphia county. Analysis of pairwise correlation suggests two correlated groupings, which is supported by two principal components explaining over 98% of observed variance in the data.




Fig. A.7. Example of optimal minimal-death control strategy $u^*(t)$ with lower-dimensional mobility data for Philadelphia County, PA, obtained by solving the minimal-death GP (11), from August 1st to November 30th, 2020. In (c) and (d), we plot the prediction using the linear model with the strategy $u^*(t)$, $H(t)$ and $D(t)$, predictions using the nonlinear model with the strategy $u^*(t)$, $\hat{H}(t)$ and $\hat{D}(t)$, and the predictions using the true mobility data $m(t)$ as a baseline for no intervention, $\hat{H}(t)$ and $\hat{D}(t)$. Parameters of the models used herein were learned as described in Section 4.3 using the principal component projections of mobility data $\tilde{m}(t)$, with the budget B^* taken from the solution to the minimal-cost GP (13).

component analysis (PCA). In particular, since our data is in the form of time series, we take the first differences of the mobility data $\mathbf{m}_d(t) := \mathbf{m}(t) - \mathbf{m}(t-1)$ to ensure stationarity, find the $K \times K$ covariance matrix C of the K -dimensional observations $\mathbf{m}_d(1), \dots, \mathbf{m}_d(T_c)$, and compute the principal components \mathbf{w}_k of C (Bishop, 2006). We then project into the space of the mobility pattern data as $\tilde{\mathbf{m}}(t) := [\mathbf{w}_1, \dots, \mathbf{w}_l]^\top \mathbf{m}(t)$ for some number of principal components l . As shown in Fig. A.6(b), the first two principal components are enough to describe over 98% of the variance in mobility data in the greater Philadelphia region over the time period of interest. For this reason, we present additional results in which this two-dimensional principal component projection, $\tilde{\mathbf{m}}(t)$, is used in place of the full mobility data from the Google Mobility

Report, $\mathbf{m}(t)$. In particular, our multitask learning approach uses $\tilde{\mathbf{m}}(t)$ to learn the mobility mapping $\beta(\cdot)$, as well as all other parameters in our model. We then use the mapping and parameters to design control strategies via the GPs (11) and (13). A reasonable budget B^* is deduced by solving (13) using this two-dimensional mobility data, and then a control strategy to minimize deaths $u^*(t)$ is designed by solving (11). These results are illustrated in Fig. A.7. Compared to using the full-dimensional mobility data, this approach does not control the pandemic as quickly and, hence, saves fewer lives. These results illustrate that, should it be possible, it is advantageous to independently control different categories of mobility. However, even if it is not possible to execute fine-grained control across many categories of mobility, this

multitask learning-driven nonlinear optimal control framework is still able to reduce the number of deaths while respecting limits on the economic costs incurred.

A.2. Relationship between the linear and nonlinear model

As shown in [Lemma 1](#), the linear epidemic model is always an upper bound to the nonlinear model. However, it is of practical importance to investigate how tight this bound may be. As such, we present results on the gap between the final states in the linear and nonlinear model in [Fig. A.8](#). For each value of the hospitalization threshold τ_H , we solve the minimal cost GP [\(13\)](#) to find the minimal budget $\mathcal{B}(\tau_H)$. As mentioned in [Section 5](#), this budget is minimal in the sense that it is the lowest cost incurred while respecting the threshold on the number of hospitalizations, τ_H , but in practice a larger budget is desirable. As such, we vary the budget for the minimal death GP [\(11\)](#) between $\mathcal{B}(\tau_H)$ and $10\mathcal{B}(\tau_H)$, and record the final number of hospitalizations in the linear model, $H(T_c)$, and the nonlinear model, $\tilde{H}(T_c)$, for $T = 61$ days. We also record the cumulative number of deaths in the linear and nonlinear models ($D(T_c)$ and $\tilde{D}(T_c)$, respectively). We then plot the difference between the quantities for the linear and nonlinear models as functions of the hospitalization threshold τ_H and the respective budgets $\mathcal{B}(\tau_H)$. Lower values of this difference correspond to a smaller gap between the linear and nonlinear models. As we can see in [Fig. A.8](#), budgets closer to $\mathcal{B}(\tau_H)$, i.e., the minimal feasible budget to respect τ_H , lead to a much larger gap between the models; this trend is consistent across values of τ_H . We observe threshold behavior, whereby a moderate increase in the budget leads to a dramatic decrease in the gap between the models. This threshold behavior explains the difference between the linear and nonlinear models in the simulations of [Section 5](#) and [Appendix A.1](#), as we employ a budget close to the minimal cost obtained from solving the GP [\(13\)](#). Hence, there exists a moderate trade-off between the prescribed budget and the tightness of the upper bounds given by the states of the linear model. Fortunately, these results illustrate that implementing the designed strategies on the true nonlinear system can lead to better outcomes than expected using any budget.

A.3. Time complexity

In order to validate the efficiency and scalability of our approach, we present data on the size and time complexity of the GPs constructed herein in [Table A.2](#). These results are not exhaustive, but meant to illustrate the practical applicability of our approach. For example, computational complexity may be reduced by dimensionality reduction techniques as in [Appendix A.1](#). All experiments were performed on a laptop with an Intel Core i7 processor running at 2.2 GHz and 16GB of RAM. These results illustrate the tractability of using geometric programming to solve problems with a large number of decision variables and constraints (see [Table A.2](#)).

Appendix B. Proofs

B.1. Proof of [Lemma 1](#)

Clearly, since in our model there is no possibility of reinfection, we have that $\tilde{S}(t) \leq S_0 = S(t)$. We now proceed by induction on the time step t . For the base case $t = 1$, recall that the initial conditions are identical for both the linearized and the true model. Thus, clearly from [\(2\)–\(7\)](#), the compartments of both the linearized and true models have identical values for $t = 1$. Now, assuming the bounds hold for $t = k - 1$, we have

$$\begin{aligned} \tilde{E}(k) &= (1 - \rho_{EI} - \rho_{EA})\tilde{E}(k-1) + \tilde{S}(k-1)\beta(t)(\gamma_A\tilde{A}(k-1) + \tilde{I}(k-1)) \\ &\leq (1 - \rho_{EI} - \rho_{EA})E(k-1) + S_0\beta(t)(\gamma_A A(k-1) + I(k-1)) = E(k), \end{aligned}$$

$$\begin{aligned} \tilde{I}(k) &= (1 - \rho_{IH} - \rho_{IR})\tilde{I}(k-1) + \rho_{EI}\tilde{E}(k-1) \\ &\leq (1 - \rho_{IH} - \rho_{IR})I(k-1) + \rho_{EI}E(k-1) = I(k), \\ \tilde{A}(k) &= (1 - \rho_{AR})\tilde{A}(k-1) + \rho_{EA}\tilde{E}(k-1) \\ &\leq (1 - \rho_{AR})A(k-1) + \rho_{EA}E(k-1) = A(k), \\ \tilde{H}(k) &= (1 - \rho_{HR})\tilde{H}(k-1) + \rho_{IH}\tilde{I}(k-1) \\ &\leq (1 - \rho_{HR})H(k-1) + \rho_{IH}I(k-1) = H(k-1), \\ \tilde{R}(k) &= \tilde{R}(k-1) + \rho_{IR}\tilde{I}(k-1) + \rho_{AR}\tilde{A}(k-1) + (1 - \alpha_D)\rho_{HR}\tilde{H}(k) \\ &\leq R(k-1) + \rho_{IR}I(k-1) + \rho_{AR}A(k-1) + (1 - \alpha_D)\rho_{HR}H(k) = R(k), \\ \tilde{D}(k) &= \tilde{D}(k-1) + \alpha_D\rho_{HR}\tilde{H}(k-1) \\ &\leq D(k-1) + \alpha_D\rho_{HR}H(k-1) = D(k), \end{aligned}$$

and the result follows.

B.2. Proof of [Lemma 2](#)

We can rewrite Eqs. [\(2\)–\(5\)](#) in matrix form by defining a state vector $\mathbf{x}(t) = [E(t), I(t), A(t), H(t)]^\top$ to obtain the dynamics

$$\mathbf{x}(t+1) = \begin{bmatrix} 1 - \rho_{EI} - \rho_{EA} & S_0\beta(\mathbf{u}(t)) & \gamma_A S_0\beta(\mathbf{u}(t)) & 0 \\ \rho_{EI} & 1 - \rho_{IR} - \rho_{IH} & 0 & 0 \\ \rho_{EA} & 0 & 1 - \rho_{AR} & 0 \\ 0 & \rho_{IH} & 0 & 1 - \rho_{HR} \end{bmatrix} \mathbf{x}(t) \\ =: M_t \mathbf{x}(t). \quad (\text{B.1})$$

It follows that

$$H(t) = [0, 0, 0, 1] M_{t-1} \cdots M_1 M_0 \mathbf{x}(0) =: f_H^t(\{\mathbf{u}(s)\}_{s=0}^{t-3}). \quad (\text{B.2})$$

Recalling that the mobility mapping $\beta(\mathbf{u}(s))$ is a posynomial, each of the matrices M_s have posynomial entries on $\beta(\mathbf{u}(s))$, and thus, on $\mathbf{u}(s)$; thus, it follows that the function $f_H^t(\{\mathbf{u}(s)\}_{s=0}^{t-3})$ in [\(B.2\)](#) is a posynomial on $\{\mathbf{u}(s)\}_{s=0}^{t-3}$ as it is the product of matrices with entries that are posynomials on $\mathbf{u}(s)$ for $s \in \{0, 1, \dots, t-3\}$. Moreover, from [\(7\)](#) we can see that $D(t)$ is simply a sum of positive constants and positive scalar multiples of $H(t)$ and is also a posynomial.

B.3. Proof of [Theorem 1](#)

By [Lemma 2](#), the objective function J in [\(10\)](#) is posynomial on $\{\mathbf{u}(t)\}_{t=0}^{T_c-4}$ as it is the sum of posynomials, since $\gamma_D \in (0, 1]$ and thus $\gamma_\infty \in [0, \infty)$ (defining $\gamma_\infty = 0$ when $\gamma_D = 1$). Moreover, the constraint $H(t) \leq \tau_H$ is clearly posynomial on $\{\mathbf{u}(s)\}_{s=0}^{t-3}$ for each $1 \leq t \leq T_c$ by [Lemma 2](#). By assumption the cost function $C_t(\mathbf{u}(t))$ is posynomial on the decision variables $\mathbf{u}(t)$ for all $0 \leq t \leq T_c - 1$, and thus the constraint $\sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t)) \leq \mathcal{B}$ is posynomial on $\{\mathbf{u}(t)\}_{t=0}^{T_c-1}$. Finally, by assumption the set \mathcal{U} is described via posynomial inequalities and monomial equalities on $\mathbf{u}(t)$ for every $0 \leq t \leq T_c - 1$. Thus, [\(11\)](#) is a geometric program, and the globally optimal solution \mathbf{u}^* may be found efficiently.

B.4. Proof of [Theorem 2](#)

Similarly to the proof of [Theorem 1](#), we may invoke [Lemma 2](#) to show that the constraint $H(t) \leq \tau_H$ is posynomial on $\{\mathbf{u}(s)\}_{s=0}^{t-3}$ for each $1 \leq t \leq T_c$. Again similarly to the proof of [Theorem 1](#), by assumption on $C_t(\mathbf{u}(t))$ the sum $\sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t))$ and, hence, the objective function, is posynomial on $\{\mathbf{u}(t)\}_{t=0}^{T_c-1}$, and the set \mathcal{U} is described via posynomial inequalities and monomial equalities on $\mathbf{u}(t)$ for every $0 \leq t \leq T_c - 1$. Hence, [\(13\)](#) is a geometric program, whose globally optimal solution \mathbf{u}^* may be found efficiently. Using this solution, we may compute the minimal budget to prevent hospital overflow as $\mathcal{B}^* = \sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t))$.




Fig. A.8. Difference between final hospitalizations (resp. cumulative deaths) in the linear model, $H(T_c)$ (resp. $D(T_c)$), and in the nonlinear model, $\tilde{H}(T_c)$ (resp. $\tilde{D}(T_c)$) as a function of the hospitalization threshold τ_H and the budget $B(\tau_H)$. For each value of τ_H , we solve the minimal-cost GP (13) to find the budget $B(\tau_H)$, then take multiples of this budget from $1\times$ to $10\times$. Trends appear consistent across hospitalization thresholds τ_H , but the difference between the models quickly shrinks as the budget $B(\tau_H)$ grows.

Table A.2

Space and time complexity of the minimal-cost GP (13) and minimal-deaths GP (11) instances solved herein, over multiple time horizons T_c . Space complexity is measured in terms of the number of decision variables and constraints of the GP in posynomial form and in the equivalent convex form, the latter of which the solver uses to find the optimal solution. Time complexity measures the running time for the solver to instantiate and solve the program.

Instance	T_c	Posynomial variables	Posynomial constants	Convex variables	Convex constants	Time (s)
GP (11)	30	16,365	7,291	52,445	29,841	18.2
GP (13)	30	17,287	7,693	55,287	31,443	14.7
GP (11)	60	69,220	31,506	219,660	125,531	104.8
GP (13)	60	71,112	32,338	225,472	128,813	86.4
GP (11)	90	158,570	72,686	501,490	287,011	379.1
GP (13)	90	161,402	73,938	510,162	291,913	380.8
GP (11)	120	294,055	135,296	928,247	531,666	2,204.0
GP (13)	120	297,960	137,017	940,208	538,422	737.2

B.5. Proof of Lemma 3

By definition in (14), for all $\{(u_1(t), \dots, u_K(t))\}_{t=0}^{T_c} \in \mathcal{U}$ we have $u_k(t)^{-1} \leq \underline{u}_k^{-1}$, $u_k(t) \leq \bar{u}_k$, and we may express the other inequalities as $u_k(t)^{-1}u_k(t-1) \leq (1-\Delta)^{-1}$, and $u_k(t)u_k(t-1)^{-1} \leq 1+\Delta$ for each $k \in \{1, \dots, K\}$. Moreover, the equality constraints may be satisfied by considering only one decision variable for each seven day period. Hence, \mathcal{U} may be described by posynomial inequalities. Next, notice that

$$C_t(\mathbf{u}(t)) = \sum_{k=1}^K c_k \frac{u_k(t)^{-1} - \bar{u}_k^{-1}}{\underline{u}_k^{-1} - \bar{u}_k^{-1}} = \sum_{k=1}^K \frac{c_k}{\underline{u}_k^{-1} - \bar{u}_k^{-1}} u_k(t)^{-1} - \sum_{k=1}^K \frac{\bar{u}_k^{-1}}{\underline{u}_k^{-1} - \bar{u}_k^{-1}}.$$

Clearly, since the values c_k , \underline{u}_k , and \bar{u}_k are constant, and $\underline{u}_k < \bar{u}_k$ implies $\underline{u}_k^{-1} - \bar{u}_k^{-1} > 0$, the left-hand term above is a posynomial in the entries of $\mathbf{u}(t)$, and the right-hand term is constant. Hence, the function $C_t(\mathbf{u}(t))$ is a posynomial in $\mathbf{u}(t) = (u_1(t), \dots, u_K(t))$ shifted by a constant (which does not affect the optimal values of $\mathbf{u}(t)$ in (13)). Considering the budget-constraint inequality in (11), from above we have

$$\sum_{t=0}^{T_c-1} C_t(\mathbf{u}(t)) \leq B \Leftrightarrow \sum_{t=0}^{T_c-1} \sum_{k=1}^K \frac{c_k}{\underline{u}_k^{-1} - \bar{u}_k^{-1}} u_k(t)^{-1} \leq B + \sum_{t=0}^{T_c-1} \sum_{k=1}^K \frac{\bar{u}_k^{-1}}{\underline{u}_k^{-1} - \bar{u}_k^{-1}},$$

which is clearly a posynomial inequality amenable to geometric programming.

References

Achterberg, M. A., Prasse, B., Ma, L., Trajanovski, S., Kitsak, M., & Van Mieghem, P. (2020). Comparing the accuracy of several network-based COVID-19 prediction algorithms. *International Journal of Forecasting*.

Aleta, A., Martín-Corral, D., y Piontti, A. P., Ajelli, M., Litvinova, M., Chinazzi, M., et al. (2020). Modelling the impact of testing, contact tracing and household quarantine on second waves of COVID-19. *Nature Human Behaviour*, 4(9), 964–971.

Ali, S. T., Wang, L., Lau, E. H. Y., Xu, X.-K., Du, Z., Wu, Y., et al. (2020). Serial interval of SARS-CoV-2 was shortened over time by nonpharmaceutical interventions. *Science*, 369(6507), 1106–1109.

Bai, Y., Yao, L., Wei, T., Tian, F., Jin, D.-Y., Chen, L., et al. (2020). Presumed asymptomatic carrier transmission of COVID-19. *Jama*, 323(14), 1406–1407.

Balcan, D., Gonçalves, B., Hu, H., Ramasco, J. J., Colizza, V., & Vespignani, A. (2010). Modeling the spatial spread of infectious diseases: The GLobal Epidemic and Mobility computational model. *Journal of Computational Science*, 1(3), 132–145.

Bhouria, M. A., Costabal, F. S., Wang, H., Linka, K., Peirlinck, M., Kuhl, E., et al. (2021). COVID-19 dynamics across the US: A deep learning study of human mobility and social behavior. *Computer Methods in Applied Mechanics and Engineering*, 382, 113891.

Birge, J. R., Candogan, O., & Feng, Y. (2020). Controlling epidemic spread: Reducing economic losses with targeted closures. *Becker Friedman Institute for Economics Working Paper No. 2020-57*, University of Chicago.

Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.

Boyd, S., Kim, S.-J., Vandenberghe, L., & Hassibi, A. (2007). A tutorial on geometric programming. *Optimization and Engineering*, 8(1), 67.

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Brauer, F., Castillo-Chavez, C., & Feng, Z. (2019). *Mathematical models in epidemiology*. Springer.

Van den Broeck, W., Gioannini, C., Gonçalves, B., Quaggiotto, M., Colizza, V., & Vespignani, A. (2011). The GLEaMviz computational tool, a publicly available software to explore realistic epidemic spreading scenarios at the global scale. *BMC Infectious Diseases*, 11(1), 37.

Caruana, R. (1997). Multitask learning. *Machine Learning*, 28(1), 41–75.

Chang, S. Y., Pierson, E., Koh, P. W., Gerardin, J., Redbird, B., Grusky, D., et al. (2021). Mobility network models of COVID-19 explain inequities and inform reopening. *Nature*, 589(7840), 82–87.

Chiang, M. (2005). *Geometric programming for communication systems*. Now Publishers Inc.

Dahl, J., & Andersen, E. D. (2021). A primal-dual interior-point algorithm for nonsymmetric exponential-cone optimization. *Mathematical Programming*, 1–30.

Day, M. (2020). COVID-19: four fifths of cases are asymptomatic, China figures indicate. British Medical Journal Publishing Group.

Eshghi, S., Khouzani, M. H. R., Sarkar, S., & Venkatesh, S. S. (2015). Optimal patching in clustered epidemics of malware. *IEEE Transactions on Network*, 24(1), 283–298.

Ferguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., et al. (2020). *Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand*.

Gandhi, M., Yokoe, D. S., & Havlir, D. V. (2020). *Asymptomatic transmission, the Achilles' heel of current strategies to control COVID-19*. Mass Medical Soc.

Garcia, C. E., Prett, D. M., & Morari, M. (1989). Model predictive control: Theory and practice—A survey. *Automatica*, 25(3), 335–348.

Giordano, G., Blanchini, F., Bruno, R., Colaneri, P., Di Filippo, A., Di Matteo, A., et al. (2020). Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. *Nature Medicine*, 26(6), 855–860.

Google (2020). COVID-19 community mobility reports. Available at <https://google.com/covid19/mobility/>. (Accessed 31 December 2020).

Hayhoe, M., Alajaji, F., & Gharesifard, B. (2018). A Polya contagion model for networks. *IEEE Transactions on Control of Network Systems*, 5(4), 1998–2010.

He, X., Lau, E. H. Y., Wu, P., Deng, X., Wang, J., Hao, X., et al. (2020). Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine*, 26(5), 672–675.

Hota, A. R., Godbole, J., Bhariya, P., & Paré, P. E. (2020). A closed-loop framework for inference, prediction and control of SIR epidemics on networks. ArXiv preprint [arXiv:2006.16185](https://arxiv.org/abs/2006.16185).

Kaplan, J., & Frias, L. (2020). Our ongoing list of how countries are reopening, and which ones remain under lockdown. *Business Insider*, <https://www.businessinsider.com/countries-on-lockdown-coronavirus-italy-2020-3>.

Khouzani, M. H. R., Venkatesh, S. S., & Sarkar, S. (2011). Market-based control of epidemics. In *2011 49th annual Allerton conference on communication, control, and computing (Allerton)* (pp. 314–320). IEEE.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. ArXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).

Köhler, J., Schwenkel, L., Koch, A., Berberich, J., Pauli, P., & Allgöwer, F. (2020). Robust and optimal predictive control of the COVID-19 outbreak. *Annual Reviews in Control*.

Lauer, S. A., Grantz, K. H., Bi, Q., Jones, F. K., Zheng, Q., Meredith, H. R., et al. (2020). The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. *Annals of Internal Medicine*, 172(9), 577–582.

Lorch, L., Trouleau, W., Tsirtsis, S., Szanto, A., Schölkopf, B., & Gomez-Rodriguez, M. (2020). A spatiotemporal epidemic model to quantify the effects of contact tracing, testing, and containment. ArXiv preprint [arXiv:2004.07641](https://arxiv.org/abs/2004.07641).

Maclaurin, D., Duvenaud, D., & Adams, R. P. (2015). Autograd: Effortless gradients in NumPy. In *ICML 2015 AutoML workshop: Vol. 238* (pp. 5).

Martcheva, M. (2015). *An introduction to mathematical epidemiology: Vol. 61*. Springer.

Morato, M. M., Bastos, S. B., Cajueiro, D. O., & Normey-Rico, J. E. (2020). An optimal predictive control strategy for COVID-19 (SARS-CoV-2) social distancing policies in Brazil. *Annual Reviews in Control*, 50, 417–431.

New York Times (2020). Coronavirus (COVID-19) data in the United States. Available at <https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html>. (Accessed 31 December 2020).

Nishiura, H., Kobayashi, T., Miyama, T., Suzuki, A., Jung, S.-m., Hayashi, K., et al. (2020). Estimation of the asymptomatic ratio of novel coronavirus infections (COVID-19). *International Journal of Infectious Diseases*, 94, 154.

Nowzari, C., Preciado, V. M., & Pappas, G. J. (2015). Optimal resource allocation for control of networked epidemic models. *IEEE Transactions on Control of Network Systems*, 4(2), 159–169.

Nowzari, C., Preciado, V. M., & Pappas, G. J. (2016). Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems Magazine*, 36(1), 26–46.

Okura, M., Kishida, M., & Lam, J. (2019). Geometric programming for optimal positive linear systems. *IEEE Transactions on Automatic Control*.

Pei, S., Kandula, S., & Shaman, J. (2020). Differential effects of intervention timing on COVID-19 spread in the United States. *Science advances*, 6(49), eabd6370.

Piguillem, F., & Shi, L. (2020). *Optimal COVID-19 quarantine and testing policies*. CEPR Discussion Paper No. DP14613.

Preciado, V. M., Zargham, M., Enyioha, C., Jadbabaie, A., & Pappas, G. J. (2013). Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks. In *52nd IEEE conference on decision and control* (pp. 7486–7491). IEEE.

Preciado, V. M., Zargham, M., Enyioha, C., Jadbabaie, A., & Pappas, G. J. (2014). Optimal resource allocation for network protection against spreading processes. *IEEE Transactions on Control of Network Systems*, 1(1), 99–108.

Rahmandad, H., Lim, T. Y., & Sterman, J. (2020). *Estimating COVID-19 under-reporting across 86 nations: implications for projections and control*. Available at SSRN 3635047.

Van Mieghem, P., Omic, J., & Kooij, R. (2009). Virus spread in networks. *IEEE/ACM Transactions on Networking*, 17(1), 1–14.

Wang, Y., Chakrabarti, D., Wang, C., & Faloutsos, C. (2003). Epidemic spreading in real networks: An eigenvalue viewpoint. In *22nd international symposium on reliable distributed systems, 2003. Proceedings*. (pp. 25–34). IEEE.

Watts, D. J., Muhamad, R., Medina, D. C., & Dodds, P. S. (2005). Multiscale, resurgent epidemics in a hierarchical metapopulation model. *Proceedings of the National Academy of Sciences*, 102(32), 11157–11162.

Woelfel, R., Corman, V. M., Guggemos, W., Seilmaier, M., Zange, S., Mueller, M. A., et al. (2020). Clinical presentation and virological assessment of hospitalized cases of coronavirus disease 2019 in a travel-associated transmission cluster. *MedRxiv*.

World Health Organization (2020a). *Novel Coronavirus (2019-nCoV): Situation Report 1: Technical documents*. World Health Organization.

World Health Organization (2020b). *Novel Coronavirus (2019-nCoV): Weekly Epidemiological Update, 21 September 2020: Technical documents*. World Health Organization.

Wu, J., Smith, S., Khurana, M., Siemaszko, C., & DeJesus-Banos, B. (2020). Stay-at-home orders across the country. Available at <https://www.nbcnews.com/health/health-news/here-are-stay-home-orders-across-country-n1168736>.

Yan, X., & Zou, Y. (2008). Optimal and sub-optimal quarantine and isolation control in SARS epidemics. *Mathematical and Computer Modelling*, 47(1–2), 235–245.