Model-based Prediction and Optimal Control of Pandemics by Non-pharmaceutical Interventions

Reza Sameni[*], *Senior Member, IEEE*

Abstract—A model-based signal processing framework is proposed for pandemic trend forecasting and control by using non-pharmaceutical interventions (NPI) at regional and country levels worldwide. The control objective is to prescribe quantifiable NPI strategies at different levels of stringency, which balance between human factors (such as new cases and death rates) and cost of intervention per region/country. Due to the significant differences in infrastructures and priorities of regions and countries, strategists are given the flexibility to weight between different NPIs, and to select the desired balance between the human factor and overall NPI cost.

The proposed framework is based on a *finite-horizon optimal control* (FHOC) formulation of the bi-objective problem and the FHOC is numerically solved by using an ad hoc *extended Kalman filtering/smoothing* framework. The algorithm enables strategists to select the desired balance between the human factor and NPI cost with a set of weights and parameters. The parameters of the model, are partially selected by epidemiological facts from COVID-19 studies, and partially trained by using machine learning techniques. The developed algorithm is applied on real global data from the Oxford COVID-19 Government Response Tracker project, which has categorized and quantified the regional responses to the pandemic for more than 300 countries and regions worldwide, since January 2020. This dataset has been used for NPI-based prediction and prescription during the XPRIZE Pandemic Response Challenge. The source codes developed for the proposed method are provided online.

I. INTRODUCTION

The COVID-19 pandemic highlighted the fact that social life, as a *dynamic system*, has always been in a *metastable* condition, which is continuously prone to pandemic outbreaks (regardless of the severity or geographical origin of pandemics). Parallel to medical solutions and vaccinations against known viruses, the rapid and effective response to future pandemics requires proactive plans in various aspects and by different scientific communities. Specifically, some of the prominent contributions, which can be made by the signal processing and data science communities include: 1) developing accurate spatio-temporal forecasting models (at different levels of abstraction), which simulate pandemic outbreaks and trends; 2) identifying quantifiable non-pharmaceutical intervention (NPI) plans, with fact-based estimates of the effectiveness and cost of each NPI [1]; 3) simulated multi-objective pandemic response strategies, which balance between NPI cost and effectiveness, to help governments and decision-makers in resource allocation and fact-based decision making to control new pandemic waves.

In this context, non-pharmaceutical interventions refer to actions and policies adopted by individuals, authorities or governments that help slowing down the spread of epidemic diseases. Enforcement of social distancing, face covering, restrictions on social events and public transportation, etc., were among the NPIs that we have all experienced during the COVID-19 pandemic. NPIs are among the best ways of controlling pandemic diseases when vaccines or medications are not yet available[1].

During the COVID-19 pandemic, several attempts were made to categorize and quantify the various NPIs of different regions and nations. The quantification of the NPI is essential for comparing the effectiveness of regional policies in containing the pandemic spread. By using signal processing and machine learning techniques, the quantified NPI can be used to forecast the future trends of the pandemic and to simulate *"what if scenarios"* for the better management of human and medical resources, and to eventually prescribe appropriate NPI for controlling the pandemic [2]. The Oxford COVID-19 Government Response Tracker (OxCGRT) is one of the NPI tracking projects, which were launched and regularly updated during the COVID-19 pandemic [3]. Most recently, this project has been used in the data science community to launch data challenges for NPI-based prediction and prescription plans. Specifically, the XPRIZE Pandemic Response Challenge addressed the problem of predicting future trends of the pandemic in different regions and countries under NPIs, and prescribing *Pareto efficient* NPI that compromise between the number of new cases and the weighted-cost of intervention [4]. The challenge was motivated by the fact that due to the variations in infrastructures, available resources and priorities, policymakers worldwide tend to give different weights to each NPI and are interested to know the impacts and consequences of each policy in advance.

During this challenge, the Alphanumerics Team from the Department of Biomedical Informatics at Emory University, adopted a rigorous model-based signal processing approach, based on *estimation theory* and *finite horizon optimal control* to address the problem of weighted NPI prescription. The notion of multi-objective finite horizon pandemic control has also been considered by other researchers in simulated scenarios [5], [6].

Since the only globally registered NPIs in the OxCGRT data

[1]See Centers for Disease Control and Prevention guidelines on NPIs: https://www.cdc.gov/nonpharmaceutical-interventions/.

are the total confirmed cases, the total confirmed deaths, and the daily NPIs, we have adopted an extension of a generic susceptible-infected (SI) compartmental model from our previous work [7], as the base model for all regions/countries. The proposed model parameters are trained on historic data and used to predict future trends from input NPI by using an extended Kalman filter. It is further shown that the forecasting model can be integrated with a finite horizon optimal controller to find the optimal daily NPIs with arbitrary NPI cost weight vectors. It should be noted that since the XPRIZE Challenge was held before the global availability of COVID-19 vaccines, the interventions considered for all nations do not include vaccinations. This fact is also reflected in the compartmental model detailed in the sequel. Nevertheless, the proposed framework is generic and can be extended to various compartmental models, whenever more accurate data exist for a specific region/nation. The source codes for the developed models and algorithms are provided online [8].

Considering that the majority of recent data analysis research tend to use data-driven machine learning (ML) techniques, this research demonstrates how classical signal processing and optimal control theories can be used and combined with ML techniques for solving modern real-world problems. In Section II the OxCGRT NPI data is explained. Section III details the data model. The developed finite horizon optimal NPI prescription framework is elaborated in Section IV. This framework is combined with an extended Kalman filter/smoother for pandemic forecasting in Section V, followed by the details of model training and implementation in Section VI. The results on real data from the COVID-19 pandemic and a detailed discussion on the proposed method are presented in Sections VII and VIII, followed by concluding remarks and future directions.

II. THE NON-PHARMACEUTICAL INTERVENTIONS DATASET

To date, the Oxford COVID-19 Government Response Tracker (OxCGRT) is an ongoing project [1], which categorizes and quantifies the NPI policies of different regions/nations since the beginning of the pandemic. The dataset was used in the XPRIZE Pandemic Response Challenge [4], which addressed the problem of pandemic trend forecasting in different regions/countries under social interventions (e.g., social distancing, mandatory mask wearing, social gathering prohibitions, closure of schools and public transportation limitations), and prescribing efficient NPIs that compromise between human factors (infection and death rates) and a weighted cost of intervention. The subset of OxCGRT NPI codes used by the XPRIZE Challenge are listed in Table I. Note that the OxCGRT dataset uses the Johns Hopkins Coronavirus dataset for US states [9], which provides a state-level estimation/prescription for the US.

III. DATA MODEL

The two major classes of methods for epidemic disease spread modeling are:

1) *Compartmental models*, which split the total population of a region into various compartments (groups) such
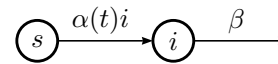

Fig. 1. The base susceptible-infected compartmental model with NPI-controlled infection rate

as susceptibles, exposed, infected, recovered, vaccinated, diseased, etc. These compartments are used to form differential/difference equations, which are fit on real data and are analytically or numerically solved to predict future trends of the disease spread.

2) *Agent-based models*, which model the behaviors of individuals and their interactions at a simplified level of abstraction. Using these models, large groups of agents are generated in stochastic simulated environments as they randomly move, interact and probabilistically pass the infection to one another, recover, pass away, etc. The population-level properties are calculated by ensemble averaging over the entire population.

Each approach has its advantages and limitations. For large population sizes at regional or national levels— which is the scope of the current study— the first approach is asymptotically accurate and is more advantageous as it can be analytically studied in a rigorous mathematical framework and combined with state estimation techniques for forecasting, and optimal control theories for NPI prescription. Therefore, the first approach was adopted for this study, using a contact-controlled time-variant version of the so-called *susceptible-infected (SI)* compartmental model shown in Fig. 1, which is a simplified variant of the general multi-compartment models studied in our previous work [7]. Apparently, more accurate models can be used for the regions that further data such as the number of recovered, hospitalized, vaccinated, or the age pyramid of the population are available. However, for the current study, since the global data provided in the Oxford dataset were the number of daily confirmed cases, total death cases, and the regional NPIs, the same SI model is used for all regions and countries.

With this background, the nonlinear dynamic equations corresponding to the proposed compartmental model are:

$$\begin{aligned} \dot{s}(t) &= -\alpha(t)s(t)i(t) \\ \dot{i}(t) &= \alpha(t)s(t)i(t) - \beta i(t) \\ \dot{\alpha}(t) &= -\gamma\alpha(t) + \gamma h[\mathbf{u}(t)] \end{aligned} \quad (1)$$

where

- $s(t)$: the fraction of population in a region/country that is susceptible at time t (i.e., the susceptible population divided by the population size N);
- $i(t)$: the fraction of population that is infected and contagious at t (i.e., the infected contagious population divided by the regional population size N);
- $\mathbf{u}(t) \in \mathbb{R}^L$: the NPI vector considered as an exogenous control input ($L = 12$ for the list of NPI in Table I, used for the XPRIZE Challenge [4]). The full definition of the Oxford data NPI set are detailed in [3];
- $\alpha(t)$: the time-variant contagion rate with inverse time unit;
- $h[\mathbf{u}(t)]$: a causal monotonic function of the NPI, which maps the NPI to the contagion rate;

TABLE I
SUBSET OF NPI INDEXES FROM THE OXFORD COVID-19 GOVERNMENT RESPONSE TRACKER (OxCGRT) PROJECT USED IN THIS STUDY, ADOPTED FROM [1]

Index	Description	Values
C1	School closing	0: no measures 1: recommend closing or all schools open with alterations resulting in significant differences compared to non-COVID-19 operations, 2: require closing (only some levels or categories, e.g., only high school or public schools), 3: require closing all levels, Blank: no data
C2	Workplace closing	0: no measures, 1: recommend closing (or recommend work from home), 2: require closing (or work from home) for some sectors or categories of workers, 3: require closing (or work from home) for all-but-essential workplaces (e.g., grocery stores, doctors), Blank: no data
C3	Cancel public events	0: no measures, 1: recommend canceling, 2: require canceling, Blank: no data
C4	Restrictions on gatherings	0: no restrictions, 1: restrictions on very large gatherings (above 1000 people), 2: restrictions on gatherings (101–1000 people), 3: restrictions on gatherings (11–100 people), 4: restrictions on gatherings (up to 10 people), Blank: no data
C5	Close public transport	0: no measures, 1: recommend closing (or significantly reduce volume/route/means of transport available), 2: require closing (or prohibit most citizens from using it), Blank: no data
C6	Stay at home requirements	0: no measures, 1: recommend not leaving house, 2: require not leaving house with exceptions for daily exercise, grocery shopping, and 'essential' travels, 3: require not leaving house with minimal exceptions (e.g. allowed to leave once a week, or only one person at a time), Blank: no data
C7	Internal movement restrictions	0: no measures, 1: recommend not to travel between regions/cities, 2: internal movement restrictions in place, Blank: no data
C8	International travel controls	0: no restrictions, 1: screening arrivals, 2: quarantine arrivals from some or all regions, 3: ban arrivals from some regions, 4: ban on all regions or total border closure, Blank: no data
H1	Public information campaigns	0: no COVID-19 public information campaign, 1: public officials urging caution about COVID-19, 2: coordinated public information campaign (e.g. across traditional and social media), Blank: no data
H2	Testing policy	0: no testing policy, 1: only those who have symptoms AND meet specific criteria (e.g. key workers, admitted to hospital, came into contact with a known case, returned from overseas), 2: testing all symptomatic people, 3: open public testing (e.g. "drive through" testing available to asymptomatic people), Blank: no data
H3	Contact tracing	0: no contact tracing, 1: limited contact tracing (not done for all cases), 2: comprehensive contact tracing (done for all identified cases)
H6	Facial coverings	0: no policy, 1: recommended, 2: required in some specified shared/public spaces outside the home with other people present, or some situations when social distancing not possible, 3: required in all shared/public spaces outside the home with other people present or all situations when social distancing not possible, 4: required outside the home at all times regardless of location or presence of other people

- β: the rate of elimination from the contagious group (through quarantine, recovery, or death), assumed to be constant in the simplified case;
- γ: the *action to effect rate* (or the inverse of the NPI to individual contact rate lag), which accounts for the delay between adopting an NPI policy and the onset of its practical effectiveness in containing the pandemic. The third equation in (1) is equivalent to $\alpha(t) = \gamma \exp[-\gamma(t-t_0)] * h[\mathbf{u}(t)]$ (for $t \geq t_0$), which is a smoothed version of $h[\mathbf{u}(t)]$. As a corner case, $\gamma \to \infty$ represents zero latency between action and effect, resulting in $\alpha(t) = h[\mathbf{u}(t)]$.

The parameters β, γ and the function $h[\mathbf{u}(t)]$ require learning using the observed variables, as detailed in Section VI. Furthermore, in [7] we showed how the infection *reproduction rate* $\mathcal{R}_t$ can be calculated from $\alpha(t)$ and β. Specifically, using the eigenanalysis-based definition of the reproduction rate proposed in [7], during the pandemic outbreak (when only several percents of the population are infected and *herd immunity* has not been reached), we have:

$$\mathcal{R}_t \approx \exp[\Delta(\alpha(t) - \beta)] \qquad (2)$$

where Δ is the reproduction rate generation time unit.

Finally, for estimation purposes, the dynamic equations in (1) can be related to real-world reports of the fractions of *new cases*:

$$n(t) = \alpha(t)s(t)i(t) + v(t), \qquad (3)$$

or through the fraction of *total confirmed cases*:

$$c(t) = s(t_0) - s(t) + v(t), \qquad (4)$$

where $v(t)$ is measurement noise due to case report errors (which inevitably existed during the COVID-19 global reports), and $s(t_0)$ is the initial susceptible population fraction at the beginning of the pandemic.

IV. FINITE HORIZON OPTIMAL NPI CONTROL

A. Cost function and problem statement

From (1), the total number of new infections over an arbitrary time window $[t_0, t_1]$ is:

$$J_0(\mathbf{u}) = \int_{t=t_0}^{t_1} \alpha(t)s(t)i(t)dt, \qquad (5)$$

and the total weighted-cost of NPIs over the same time period is

$$J_1(\mathbf{u}) = \int_{t=t_0}^{t_1} \mathbf{w}(t)^T \mathbf{u}(t)dt \qquad (6)$$

where $\mathbf{w}(t)$ is the NPI weight vector given as input. The motivation for the user selected weight vector $\mathbf{w}(t)$ is that the cost of intervention is different across regions. A stereotypical example considered in the XPRIZE Challenge was that "closing public transportation may be much costlier in London than it is in Los Angeles. Such preferences are expressed as weights associated with each intervention plan dimension, given to the prescriptor as input for each region [4]."

With these assumptions, the optimal NPI prescription problem can be formulated as a bi-objective optimization problem, with total cost:

$$J(\mathbf{u}) = (1 - \epsilon)J_0(\mathbf{u}) + \epsilon J_1(\mathbf{u}) \quad \text{s.t. } \mathbf{u} \in \Gamma \quad (7)$$

where $\epsilon \in [0, 1]$ is a free parameter that compromises between the human factor (J_0) and the NPI cost (J_1), and Γ is the set of admissible inputs:

$$\Gamma = \{\mathbf{u} | \mathbf{u}^{\min} \leq \mathbf{u}(t) \leq \mathbf{u}^{\max}, \forall t \in [t_0, t_1]\} \quad (8)$$

where $\mathbf{u}^{\min}$ and $\mathbf{u}^{\max}$ are (element-wise) the minimum and maximum ranges of each NPI from the Oxford dataset ('Values' column in Table I). Accordingly, $\mathbf{u}^{\min} = \mathbf{0}$ and $\mathbf{u}^{\max} = [3, 3, 2, 4, 2, 3, 2, 4, 2, 3, 2, 4]^T$

For a given weight pair $\{\epsilon, \mathbf{w}(t)\}$, the objective is to find $\mathbf{u}^*(t)$ for all $t \in [t_0, t_1]$, such that:

$$J(\mathbf{u}^*) = \min_{\Gamma}(J(\mathbf{u})) \quad (9)$$

B. The Pareto optimal solution

The problem (9) can be solved by *finite horizon optimization* [10]. The inputs which satisfy this equation are known as *Pareto optimal (efficient)*, in optimal control theory. In fact, for an arbitrary weight vector $\mathbf{w}(t)$, by sweeping ϵ over $[0, 1]$, the *Pareto-optimal front* of the optimization problem is found, from which pandemic strategists can select the desired free parameter ϵ (which balances the desired operation point which balance between NPI effectiveness and cost) and its corresponding optimal NPI $\mathbf{u}^*(t)$, to be adopted by the country/region.

To solve (9), first the corresponding *Hamiltonian* function is formed [10, Ch. 2]:

$$\begin{aligned}\mathcal{H} = \ & (1 - \epsilon)\alpha(t)s(t)i(t) + \epsilon\mathbf{w}(t)^T\mathbf{u}(t) \\ & -\lambda_1(t)\alpha(t)s(t)i(t) \\ & +\lambda_2(t)[\alpha(t)s(t)i(t) - \beta i(t)] \\ & -\gamma\lambda_3(t)\{\alpha(t) - h[\mathbf{u}(t)]\}\end{aligned} \quad (10)$$

where $\lambda_1(t)$, $\lambda_2(t)$ and $\lambda_2(t)$ are known as *co-states*. According to *Pontryagin's minimum principle*, the co-states and the optimal solution $\mathbf{u}^*$ satisfy [10, Ch. 6]:

$$\begin{aligned}\dot{\lambda}_1(t) &= -\frac{\partial\mathcal{H}}{\partial s} = [\lambda_1(t) - \lambda_2(t) - (1-\epsilon)]\alpha(t)i(t) \\ \dot{\lambda}_2(t) &= -\frac{\partial\mathcal{H}}{\partial i} = [\lambda_1(t) - \lambda_2(t) - (1-\epsilon)]\alpha(t)s(t) + \beta\lambda_2(t) \\ \dot{\lambda}_3(t) &= -\frac{\partial\mathcal{H}}{\partial\alpha} = [\lambda_1(t) - \lambda_2(t) - (1-\epsilon)]s(t)i(t) + \gamma\lambda_3(t) \\ \mathcal{H}(\mathbf{u}^*) &\leq \mathcal{H}(\mathbf{u}), \quad \forall\mathbf{u} \in \Gamma\end{aligned} \quad (11)$$

When the inputs are unconstrained, the Hamiltonian minimizer input $\mathbf{u}^*$, in the last condition of (11) can be found by solving

$$\boldsymbol{\nabla}_{\mathbf{u}}\mathcal{H}(\mathbf{u}^*) = \mathbf{0}, \quad (12)$$

where $\boldsymbol{\nabla}_{\mathbf{u}}\mathcal{H}$ denotes the Hamiltonian gradient with respect to the input vector $\mathbf{u}$, and the condition should hold element-wise. In this case, a sufficient condition for the existence of a solution is to have $\boldsymbol{\nabla}_{\mathbf{u}}^2\mathcal{H}(\mathbf{u}^*) \succ 0$ (where $\boldsymbol{\nabla}_{\mathbf{u}}^2$ denotes the Hessian operator with respect to the input vector $\mathbf{u}$, and $\succ 0$ denotes positive-definiteness). In the constrained input case—

as in this problem— where the inputs are confined to the admissible set (8), while the global solution of (12) might not exist or belong to the admissible set (8), a Hamiltonian minimizer optimal input still exists. In either case, the optimal input is found as a parametric function of the costate $\lambda_3(t)$ and the other model parameters.

The parametric optimal input found from (12) is next combined with (11) and (1) to calculate the states, using the initial conditions and appropriate boundary conditions (also known as the *transversality conditions*) on the co-states and the Hamiltonian. The desired boundary conditions, which satisfy the pandemic control problem are:

$$\lambda_1(t_1) = 0, \quad \lambda_2(t_1) = 0, \quad \lambda_3(t_1) = 0. \quad (13)$$

The conditions in (13) are the general free end-point conditions of finite horizon optimization problems, which match the objectives of the pandemic control problem. Alternative transversality conditions that can be studied within the proposed framework are [10, Section 2.7]:

1) When the end-time t_1 is not fixed, but we require that $i(t_1)$ reaches below $i_{\max}$ by the end of the control period (*infinite horizon* scenario). This requires the additional condition: $\mathcal{H}(t_1) = 0$.
2) Assuming that the objective of any NPI policy over a reasonable time period $[t_0, t_1]$ (long enough to make the NPIs effective) is to bring the number of active cases down to $i(t_1) \leq i_{\max}$, where $i_{\max}$ is some target fraction of active cases (ideally zero), the second condition in (13) can be replaced by: $\lambda_2(t_1)[i(t_1) - i_{\max}] = 0$.
3) We require that $i(t_1)$ drops below $i_{\max}$ any time before a maximum end time t_f, which requires $(t_1 - t_f)\mathcal{H}(t_1) = 0$.

C. The NPI to inter-human contact map

The solution of the NPI optimization problem depends on the choice of $h[\mathbf{u}(t)]$, i.e. the NPI to inter-human contact mapping model. By common sense, $h[\mathbf{u}(t)]$ is expected to be a monotonically decreasing function of the input NPI vector $\mathbf{u}(t)$. In other words, more strict restrictions on social contact (corresponding to the higher values in Table I) should reduce the person-to-person contact rates population-wise (this was the globally accepted rationale behind the social restrictions during the COVID-19 pandemic). However, the exact shape of $h[\mathbf{u}(t)]$ generally requires *learning* from historic data, where the monotonic decreasing assumption acts as a constraint during learning.

Based on this empirical assumption, we study the following two cases, which lead to closed form solutions for the optimal input as functions of the model co-states.

1) Linear regression model: Let us take:

$$h[\mathbf{u}(t)] = b + \mathbf{a}^T[\mathbf{u}^{\max} - \mathbf{u}(t)] \quad (14)$$

where $\mathbf{a}$ is a vector of *input influence weights* and b is a constant bias (intercept value). The LASSO model falls into this category. Adding the constraint $\mathbf{a} \geq \mathbf{0}$ guarantees the monotonically decreasing relationship between the NPI and

α. In other words, more stringent NPI policies have a non-increasing effect on the human interactions parameter α (i.e. the NPI do not have any counter-impacts on the contact rates). Inserting $h[\mathbf{u}(t)]$ in (10) we find:

$$\nabla_{\mathbf{u}}\mathcal{H}(\mathbf{u}) = \epsilon\mathbf{w}(t) - \gamma\lambda_3(t)\mathbf{a} \qquad (15)$$

Now since $\nabla_{\mathbf{u}}\mathcal{H}(\mathbf{u})$ is independent of $\mathbf{u}$, depending on its sign, the Hamiltonian (which is a linear function of $\mathbf{u}$), admits its minimum at one of the extreme ends of the admissible input ranges (8). This eventually results in

$$u_k^*(t) = \begin{cases} u_k^{\min}: & \epsilon w_k(t) > \gamma\lambda_3(t)a_k \\ u_k^{\max}: & \epsilon w_k(t) < \gamma\lambda_3(t)a_k \end{cases} \qquad (16)$$

for $k = 1, \ldots, L$.

2) Quadratic regression: In the second case, we assume

$$h[\mathbf{u}(t)] = b + \mathbf{a}^T[\mathbf{u}^{\max} - \mathbf{u}(t)] + \frac{1}{2}[\mathbf{u}^{\max} - \mathbf{u}(t)]^T\mathbf{S}[\mathbf{u}^{\max} - \mathbf{u}(t)] \qquad (17)$$

where $\mathbf{a} \geq \mathbf{0}$ and $\mathbf{S} \in \mathbb{R}^L$ is a positive definite matrix. These assumptions guarantee the monotonically decreasing relationship between the NPI and $\alpha(t)$. Multivariate constrained polynomial fitting can be used to find $\mathbf{S}$, $\mathbf{a}$ and b from historic NPI and case-report data[2]. Therefore, the required constraint polynomial fitting is straightforward to implement by conventional least squares solvers. In this case, we have:

$$\nabla_{\mathbf{u}}\mathcal{H}(\mathbf{u}) = \epsilon\mathbf{w}(t) - \gamma\lambda_3(t)\{\mathbf{a} + \mathbf{S}[\mathbf{u}^{\max} - \mathbf{u}(t)]\} \qquad (18)$$

and setting $\nabla_{\mathbf{u}}\mathcal{H}(\tilde{\mathbf{u}}) = \mathbf{0}$ gives

$$\tilde{\mathbf{u}} = \mathbf{u}^{\max} - \mathbf{S}^{-1}[\frac{\epsilon\mathbf{w}(t)}{\gamma\lambda_3(t)} - \mathbf{a}] \qquad (19)$$

From the *second partial derivative test*, since $\nabla_{\mathbf{u}}^2\mathcal{H} = \gamma\lambda_3(t)\mathbf{S}$, and the fact that $\mathbf{S}$ is assumed to be positive definite, three cases may occur: 1) if $\lambda_3(t) > 0$, $\tilde{\mathbf{u}}$ is a local minimum, 2) if $\lambda_3(t) < 0$, it is a local maximum, and 3) $\lambda_3(t) = 0$ results in a saddle point. Therefore, by applying Pontryagin's minimum principle and considering the admissible input range (8), after some algebraic simplifications we find:

$$u_k^*(t) = \begin{cases} u_k^{\min}: & \epsilon w_k(t) > \gamma\lambda_3(t)(a_k + s_k) \\ \tilde{u}_k: & \gamma\lambda_3(t)a_k < \epsilon w_k(t) < \gamma\lambda_3(t)(a_k + s_k) \\ u_k^{\max}: & \epsilon w_k(t) < \gamma\lambda_3(t)a_k \end{cases} \qquad (20)$$

where $\tilde{u}_k$ and s_k are the kth entries of the vectors $\tilde{\mathbf{u}}$ and $\mathbf{S}(\mathbf{u}^{\max} - \mathbf{u}^{\min})$, respectively. Comparing (20) and (16), it is clear how the quadratic case simplifies to the linear case when $\mathbf{S} \to 0$. It is also seen that in the linear case, the "optimal NPI" is always one of the extreme cases $u_k^{\min}$ (no action) or $u_k^{\max}$ (maximum stringency). But when the NPI to contact rate map $h(\cdot)$ is nonlinear, intermediate interventions may also be in the optimal NPI set.

---

[2]Despite the quadratic form of (17), since it is linear in parameters $\mathbf{S}$, $\mathbf{a}$ and b, multivariate constrained polynomial fitting is applicable to find the unknown parameters by using historic data (model fitting over previous NPI actions adopted by different nations/states, since the beginning of the pandemic).

V. A UNIFIED PANDEMIC TREND PREDICTOR AND NPI PRESCRIPTOR

For an ideal model, the state and co-state dynamic equations detailed in Section IV can be solved with numerical toolboxes for finite horizon control (cf. [11] for a MATLAB-based solution). However, in practice, there are some major issues, which limit the numerical performance, including: 1) model inaccuracies, 2) noisy observations (inaccurate case reports), 3) missing reports (during holidays), 4) unknown or variable parameters (which is inevitable for a highly dynamic multi-aspect complex system, such as a global pandemic), 5) the difficulty of incorporating the start- and end-point boundary conditions from (13).

Due to these issues, we propose a novel technique, based on optimal state estimation. Accordingly, we have integrated the finite horizon NPI optimizer and the new-case predictor in classical extended Kalman filter (EKF) and extended Kalman smoother (EKS) schemes [12]. Using (1), (3) and (11), the state-augmented dynamic equations for the EKF are:

$$\begin{aligned}
\dot{s}(t) &= -\alpha(t)s(t)i(t) + w_s(t) \\
\dot{i}(t) &= \alpha(t)s(t)i(t) - \beta i(t) + w_i(t) \\
\dot{\alpha}(t) &= -\gamma\alpha(t) + \gamma h[\mathbf{u}^*(t)] + w_\alpha(t) \\
\dot{\lambda}_1(t) &= [\lambda_1(t) - \lambda_2(t) - (1-\epsilon)]\alpha(t)i(t) + \eta_1(t) \\
\dot{\lambda}_2(t) &= [\lambda_1(t) - \lambda_2(t) - (1-\epsilon)]\alpha(t)s(t) + \beta\lambda_2(t) + \eta_2(t) \\
\dot{\lambda}_3(t) &= [\lambda_1(t) - \lambda_2(t) - (1-\epsilon)]s(t)i(t) + \gamma\lambda_3(t) + \eta_3(t) \\
n(t) &= \alpha(t)s(t)i(t) + v(t)
\end{aligned} \qquad (21)$$

where the first six equations are the state and co-state dynamics, the last equation is the observation equation and $h[\mathbf{u}^*(t)]$ is the impact of the optimal control calculated from (16) or (20). The terms $w_s(t)$, $w_i(t)$, $w_\alpha(t)$, $\eta_1(t)$, $\eta_2(t)$ and $\eta_3(t)$ in (21) represent process noises, and $v(t)$ is observation noise. Note that for an estimation based on the total number of confirmed cases (instead of the new cases), the last observation equation in (21) can be replaced with (4). The dynamic system (21) may now be numerically solved by using standard EKF and EKS equations. The discretized version of (21), which is required for the discrete-dime implementation of the EKF and EKS are detailed in the Appendix.

The finite horizon optimal solution and the EKF/EKS-based forecasting model can now be unified for optimal NPI prescription and pandemic trend forecasting. The overall proposed method which unifies both objectives is summarized in Algorithm 1. The MATLAB implementation of the overall developed algorithm is available in our online repository.

VI. MODEL TRAINING

The model parameters $h[\mathbf{u}(t)]$, β, γ and the EKF/EKS parameters require region-wise training or fact-based selection. For this study, we used classical techniques for *Kalman filter engineering*, based on monitoring the properties of the *innovations process* of the Kalman filter to select and automatically adapt the Kalman filter parameters (initial/final states and covariance matrices) over time [13, Ch. 8]. The parameters related to the social and epidemic aspects of the model are explained in the sequel.

Algorithm 1 Summary of the proposed algorithm

Input: Historic case reports and NPI files (or an arbitrary scenario file from the *standard predictor model*)
Input: The NPI weights $\mathbf{w}(t)$, per region/country
Input: The Pareto front tuning parameter $\epsilon \in [0, 1]$.

1: **for all** Regions **do**
2:    Train the compartmental model parameters over historic NPI and case reports (or the standard predictor scenario file).
3:    Use EKF and EKS for prediction and prescription of finite horizon optimal control inputs $\mathbf{u}^*(t)$.
4: **end for**

The mapping $h[\mathbf{u}(t)]$ was trained over the Oxford dataset historic cases and NPIs (as precised in the later demonstrated results [3]) and tested over the end of training date up to the current date. For this, the developed EKS was first applied to the historic data by neglecting the explicit relation between the NPI and contact rates (equivalent to $h[\mathbf{u}(t)] = 0$). Referring to (21), this assumption is equivalent to considering the input-driven fluctuations of $\alpha(t)$ inside the process noise $w_\alpha(t)$. Therefore, the entry of the process noise covariance matrix, which corresponds to $w_\alpha(t)$ is selected to be higher, to include the inaccuracy of the model due to neglecting $h[\mathbf{u}(t)]$. The resulting EKS gives a primary estimate of $\alpha(t)$ over the training period, which in the next stage is given to a constrained LASSO or quadratic polynomial fitter (for the linear and quadratic forms presumed in Section IV-C), to estimate $h[\mathbf{u}(t)]$ using the historic NPI data. In the next phase, the trained model $h[\mathbf{u}(t)]$, together with the historic data is used in a second round of EKS; this time by using the historic NPI and apparently a smaller *a priori* assumption for the variance of $w_\alpha(t)$ (as it no longer accounts for $h[\mathbf{u}(t)]$). After the secondary EKS, the new estimates of $\alpha(t)$ are once more used to refine the model parameters of $h[\mathbf{u}(t)]$. The refined parameters are stored per country/region for utilization during the prescription phase (over real or synthetic scenarios).

The *action to effect rate* parameter γ was selected intuitively. From various social experiences, it is reasonable to expect a smooth transition in $\alpha(t)$ due to any change in the NPI. This is based on the social experience that imposing any policy on a complex social system is rarely abrupt. Although the transition is region and NPI dependent, in order to reduce the model complexity, we have fixed $\gamma=1/(7 \text{ days})=0.1429 \text{ days}^{-1}$, for all regions/countries.

The recovery parameter β was selected by educated guesses from the CDC reports regarding recovery and contagion periods[3]. Accordingly, multiple scientific studies worldwide have reported that an exposed subject is no longer *infectious* after three to four weeks. This is evidently a stochastic range. To clarify, with an exponential model such as the SI mode, in absence of new infected cases ($\alpha = 0$), we find the ratio $i(t_0 + T)/i(t_0) = \exp(-\beta T)$, which can be considered as an

---

[3]Refer to CDC guidelines for Interim Guidance on Ending Isolation and Precautions for Adults with COVID-19: https://www.cdc.gov/coronavirus/2019-ncov/hcp/duration-isolation.html

---

exponential law for the *probability of infectiousness after T time units (days)*. Combining the model with the CDC reports, we derive the following rule for setting β:

$$\beta = \frac{-\log(\text{probability of contagion after } T \text{ time units})}{T} \tag{22}$$

For the later presented results, we have set the probability of contagion to 0.01 and T=21 days, resulting in β=0.2193 days^{-1}.

Following recent studies [14], the reproduction rate of the pandemic during outbreak was taken to be $\mathcal{R}_0$=2.5, which using (2) together with β were used to initialize $\alpha(t_0)$, which is the contact rate during outbreak.

Note that one of the advantages of the EKF/EKS framework is that the model parameters can also be considered as state variables and be *state augmented* with the other equations to be estimated (or updated over time). This approach can be used for both γ and β to refine the initial educated guesses.

Finally, the regional/national population sizes, as required for normalizing the total and new contaminated cases to the normalized variables of the SI model were obtained from public global population datasets and assumed to remain fix over the study (i.e., immigration, inter-border travels, natural birth/deaths have been neglected throughout the study).

VII. RESULTS

The OxCGRT dataset has above 300 countries and regions (states). Due to inconsistencies in the reported COVID-19 cases some of the countries/regions were omitted from the study during the XPRIZE Challenge and the proposed prediction-prescription algorithm was trained and applied to a total number of 235 regions/countries, with arbitrary NPI cost weights. The training period for the model was from January 1, 2020 to Feb 7, 2021, and the test phase was from Feb 8, 2021 up to the current date (May 7, 2021 in the presented results). As proof of concept, the bi-objective optimization space of J_1 (NPI cost) vs J_0 (human factor) are shown in Fig. 2 for several countries worldwide. Accordingly, each point in this figure corresponds to a (J_0, J_1) pair for a sequence of NPI scenarios over the test phase. In Fig. 2, the red points are the result of the proposed method for different values of ϵ (the bi-objective optimization free parameter). The black crosses correspond to continuing the latest NPI policy of each government (at the end of the training date). Finally the blue points correspond to a pool of random constant stringency $\mathbf{u}(t) = \boldsymbol{\kappa}$, $\boldsymbol{\kappa} \in [\mathbf{u}^{\min}, \mathbf{u}^{\max}]$, and random variable stringency $\mathbf{u}(t) \in [\mathbf{u}^{\min}, \mathbf{u}^{\max}]$. In all cases, the user defined NPI weight vector was chosen to be equal for all NPI ($\mathbf{w}(t) = \mathbf{1}$), i.e., the NPI were considered equally weighted for the policymaker.

As a bi-objective problem, the Pareto efficient front comprises of the NPI points which either have a smaller value of J_0 or J_1, while the non-efficient solutions are the ones for which there exists at least a point that gives a smaller cost of both J_0 and J_1. In other words, a Pareto efficient solution should outperform any other solution either in its cost or efficiency. As trivial cases, the maximum stringency case $\epsilon = 0$ (maximal enforcement of social limitations, to minimize

human cost) and the minimum stringency case $\epsilon = 1$ (no social constraints, to minimize costs of intervention) are both Pareto efficient; the former minimizes the human losses and the latter minimizes the socioeconomic cost of intervention. Apparently, policymakers prefer a balance between the these two objectives. Therefore, the Pareto efficient NPI policies should be close to the origin or along one of the left or bottom axes. From Fig. 2, it is seen that none of the current NPI policies adopted by countries/regions are optimal (assuming equal NPI weights), and despite the significant differences between the different studied countries, the Pareto optimal points all belong to the proposed algorithm. Note that the "optimal point" is clearly a function of the bi-objective parameter ϵ, eventually selected by the policymakers.

Similar performance was obtained for all the 235 studied countries and regions.

A sample result of tracking the trend of new cases and the exponential growth using the hereby proposed extended Kalman smoother on real daily reported cases of the US since the 100th case report is shown in Fig. 3. For simplicity, the effect of NPIs have not been considered in this example.

The proposed algorithm is computationally very efficient. The MATLAB version of the codes applied to all regions and countries (235 in total), takes less than 30 s to train over the historic cases on a MacBook Pro laptop with 2.3 GHz Quad-Core Intel Core i7 and 32GB of memory, without notable optimizations. The run-time on the test scenarios takes about 15 s in total for all regions (as it contains only one EKS stage during the test period, per region/country).

VIII. DISCUSSION

The highlights of the developed algorithm for predicting pandemic trends and prescribing Pareto optimal NPI policies include:

- The method is based on theoretical derivations and within the scope of the proposed compartmental model accuracy (which is asymptotically accurate for region/country-level population sizes), it gives accurate *Pareto efficient* solutions.
- The operation point on the Pareto front can be selected by a single parameter $\epsilon \in [0, 1]$ selected by the policymaker, where the corner case $\epsilon = 0$ neglects the NPI cost (in favor of the human factor) and $\epsilon = 1$ neglects the human factor (in favor of intervention cost).
- The prediction and prescription problems are integrated in a unified framework. Nevertheless, the method is applicable to both real-world data and any other machine-learning based technique, which accurately predicts pandemic trends from historic data (see for example [2]).
- This framework can be used for *targeted pandemic control*, where strategists can target specific infection bounds that match the medical resources of a country/region, over a fixed or maximally bounded period of time. Therefore, apart from the optimal NPI and fatality rate objectives, such scenarios can also be considered: *"how to bring the pandemic reproduction rate below 0.8 by a specific time?"*, or *"how to bring the new cases below 200 per day in less than two months?"* The training phase of the pandemic over historic data together with the forecasting model, can be used to study the feasibility of such scenarios and the prescription of the required NPI policies that would achieve these objectives.

- Both the model parameters and NPI cost weights can be updated over time. Specifically, we do not assume the NPI cost weights to be constant over time. Therefore, unprecedented events such as vaccination or virus mutation effects can be integrated in the model with appropriate training. In fact, according to the so-called *principle of optimality*, "any portion of an optimal control trajectory is optimal [10, Sec 6.4]", which implies that optimality of future actions is independent of the past. Therefore, the prescribed optimal control strategy may be adopted at any point, regardless of the past actions of a region/nation.
- Since the Pareto front solutions are found by mathematical derivation (rather than trial and error or cumbersome searches), the proposed framework is extremely computationally efficient and the run-time computational load for testing arbitrary scenarios is minimal. This permits the combination of the proposed method with other machine learning methods to reduce the search space and to improve the accuracy on other datasets and under more complicated models such as the Long short-term memory (LSTM), as in [2].
- The predictor part of the model gives confidence intervals during both the prediction and prescription steps of the algorithm. Therefore, the performance and well-function of the algorithm can be continuously monitored and adapted.
- The proposed framework is extendable to pharmaceutical intervention plans and vaccinations, whenever sufficient data is available to train alternative compartmental models.
- Finally, the hereby proposed method which utilized EKF and finite-interval EKS to resolve the numerical problem of finite horizon optimal control can be used for other optimal prediction and control problems on real-world data.

IX. CONCLUSION AND FUTURE WORK

In this research, a model-based approach was used for the prediction and prescription of NPI that best balance between an arbitrary weighted-cost of interventions and the human factors (number of new cases) during a pandemic. The proposed algorithm and the prescribed NPI were proved to be Pareto optimal, to the extent of the accuracy of the utilized compartmental model. Software implementations of the proposed algorithms are online available at [8].

In future studies, different aspects of this framework can be extended and improved, including:

- Using advanced machine learning algorithms for learning the NPI to contact rate function $\mathbf{h}(\cdot)$. The LSTM is specifically a promising approach.
- For regions which have access to additional data (e.g., the number of hospitalized, number of vaccinated, fatality
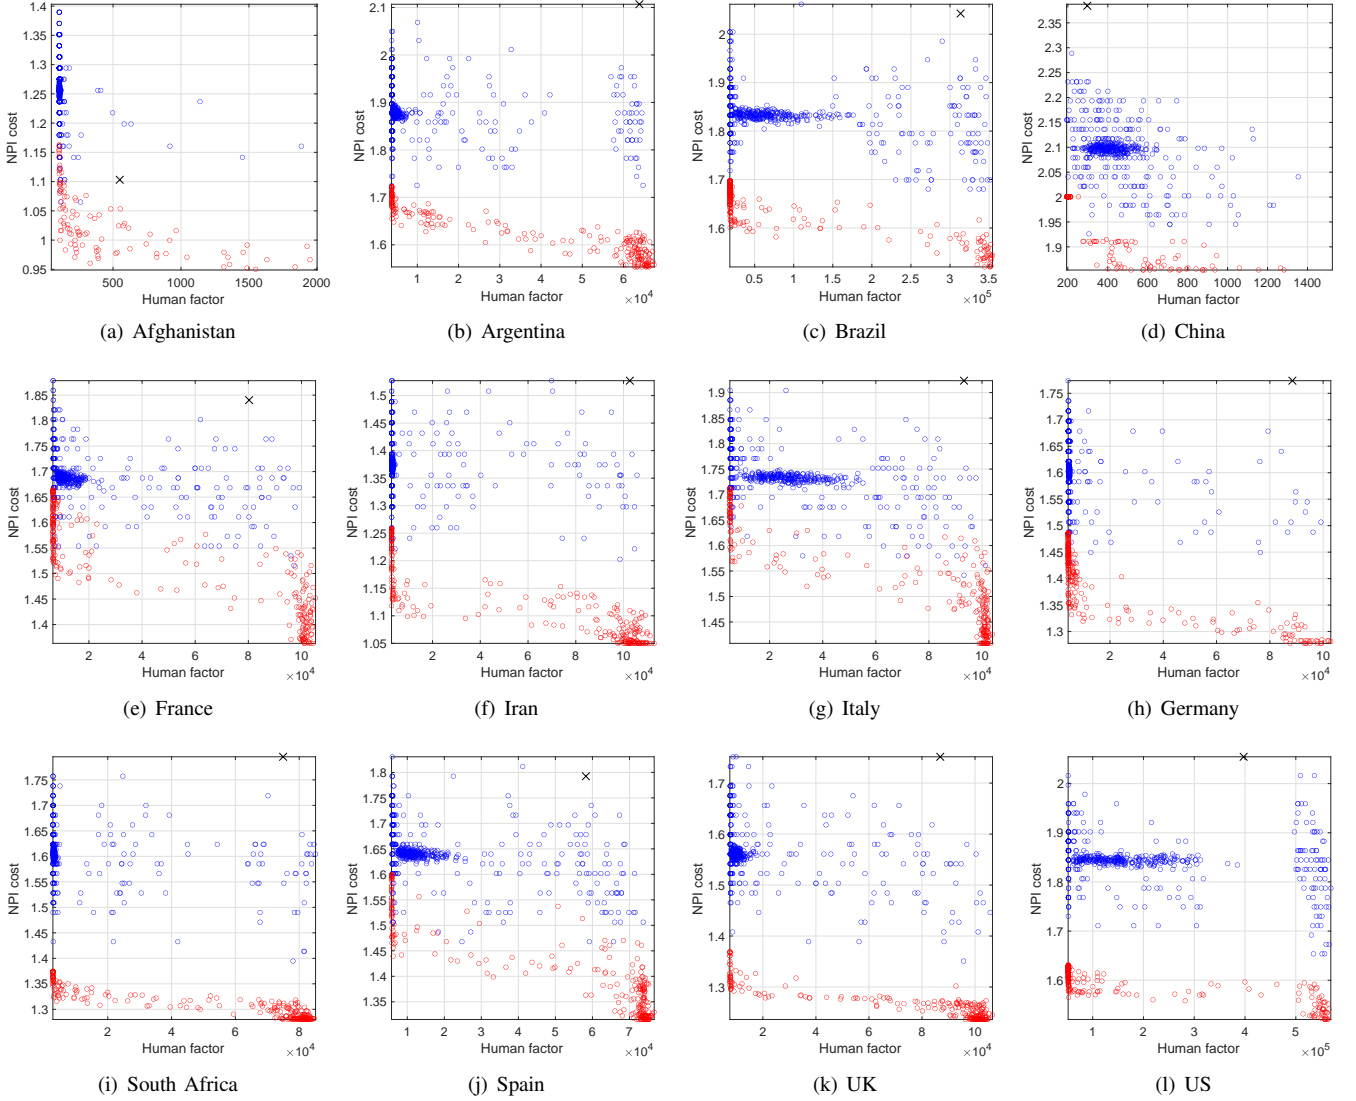
Fig. 2. Biobjective optimization space for sample countries. Black cross: fixed NPI (continuing current policies); Blue: random NPI inputs (both constant and variable over time); Red: finite-horizon optimal input for 250 $\epsilon \in [0, 1]$. $h[\mathbf{u}(t)]$ was found by linear regression over historic NPI from Jan 1, 2020 to Feb 7, 2021, using a LASSO with positive coefficients constraint.

rate of the virus, the population age pyramid, etc.), more accurate models such as the fatal susceptible-exposed-infected-recovered (SEIR) can be used [7]. Other indexes such as the daily death reports can be augmented as additional observation equations in the dynamic model (21), and will help to increase the EKF/EKS accuracies.

- Theoretical aspects of the proposed EKF/EKS frameworks, including stability conditions, parameter identifiability and robustness to parameter and modeling errors require further studies.

APPENDIX
THE DISCRETE-TIME MODEL

For a discrete-time implementation of the EKF/EKS, the discrete form of the dynamic system (21) is required. Accord-

ingly, we define

$$\boldsymbol{s}_k = [s(k\Delta), i(k\Delta), \alpha(k\Delta)]^T$$
$$\boldsymbol{w}_k = [w_s(k\Delta), w_i(k\Delta), w_\alpha(k\Delta), \eta_1(k\Delta), \eta_2(k\Delta), \eta_3(k\Delta)]^T$$
$$n_k = n(k\Delta), \quad c_k = c(k\Delta), \quad v_k = v(k\Delta)$$

where Δ is the discretization time unit. Assuming that Δ is small as compared with the variations of the pandemic trends, a first order discrete approximation of (21) is found as follows:

$$s_{k+1} = s_k - \Delta\alpha_k s_k i_k + \Delta w_{sk}$$
$$i_{k+1} = i_k + \Delta\alpha_k s_k i_k - \Delta\beta i_k + \Delta w_{ik}$$
$$\alpha_{k+1} = \alpha_k - \Delta\gamma\alpha_k + \Delta\gamma h[\mathbf{u}_k^*] + \Delta w_{\alpha k}$$
$$\lambda_{1,k+1} = \lambda_{1k} + \Delta[\lambda_{1k} - \lambda_{2k} - (1-\epsilon)]\alpha_k i_k + \Delta\eta_{1k}$$
$$\lambda_{2,k+1} = \lambda_{2k} + \Delta[\lambda_{1k} - \lambda_{2k} - (1-\epsilon)]\alpha_k s_k + \Delta\beta\lambda_{2k} + \Delta\eta_{2k}$$
$$\lambda_{3,k+1} = \lambda_{3k} + \Delta[\lambda_{1k} - \lambda_{2k} - (1-\epsilon)]s_k i_k + \Delta\gamma\lambda_{3k} + \Delta\eta_{3k}$$
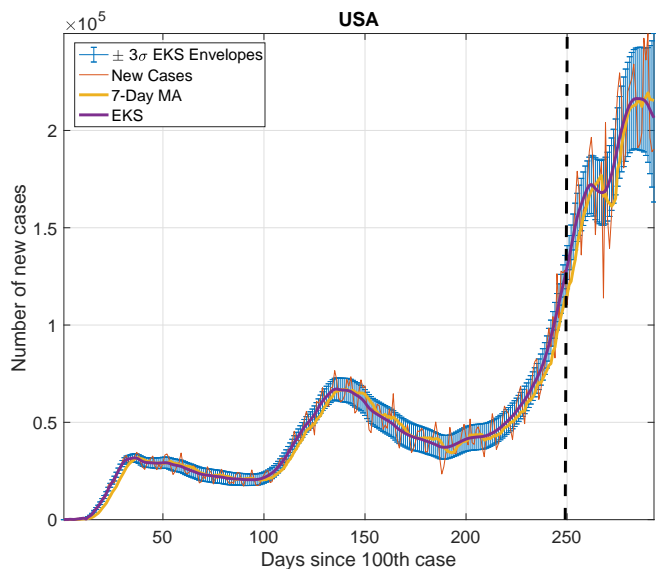$$n_k = \alpha_k s_k i_k + v_k$$

$$(23)$$

Fig. 3. Tracking the trend of new cases using the proposed extended Kalman smoother on daily reported cases of the US, since the 100th case report. The raw noisy daily reports have been adopted from the Oxford COVID-19 Government Response Tracker (OxCGRT) project [3]. The training period is up to day 250 and used to forecast the trends thereafter.

which can be formulated in a compact form:

$$\begin{aligned} \boldsymbol{s}_{k+1} &= \mathbf{f}(\boldsymbol{s}_k, \boldsymbol{w}_k; h(\mathbf{u}_k)) \\ n_k &= \mathbf{g}(\boldsymbol{s}_k) + v_k \end{aligned} \quad (24)$$

where $\mathbf{f}(\cdot)$ and $\mathbf{g}(\cdot)$ represent the nonlinear equations in (23). Note that following (4), if the number of confirmed cases is used as the observation, the second equation in (24) is replaced with

$$c_k = s_0 - s_k + v_k \quad (25)$$

which is a linear function of the state vector.

It is straightforward to linearize the discrete-time dynamic model (23) by calculating its Jacobian matrices, as required for the implementation of the EKF/EKS. An alternative approach is to use a *continuous-dynamics discrete-observations* approach, which is a classical method in optimal state estimation. Accordingly, for implementing the EKF/EKS, the state equations can be updated by using the continuous version of the dynamic model (21), while the observations are only updated on discrete-time intervals (e.g., on a daily basis).

ACKNOWLEDGMENT

REFERENCES

[1] T. Hale, N. Angrist, R. Goldszmidt, B. Kira, A. Petherick, T. Phillips, S. Webster, E. Cameron-Blake, L. Hallas, S. Majumdar, and H. Tatlow, "A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)," *Nature Human Behaviour*, Mar. 2021. [Online]. Available: https://doi.org/10.1038/s41562-021-01079-8

[2] R. Miikkulainen, O. Francon, E. Meyerson, X. Qiu, D. Sargent, E. Canzani, and B. Hodjat, "From prediction to prescription: Evolutionary optimization of nonpharmaceutical interventions in the covid-19 pandemic," *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 2, pp. 386–401, 2021.

[3] Thomas Hale and Sam Webster and Anna Petherick and Toby Phillips and Beatriz Kira, *Oxford COVID-19 Government Response Tracker*, 2020, Blavatnik School of Government. [Online]. Available: https://github.com/OxCGRT/covid-policy-tracker

[4] XPRIZE, *The XPRIZE Pandemic Response Challenge*, Oct 2020 – Feb 2021. [Online]. Available: https://xprize.org/pandemicresponse

[5] A. Mallela, "Optimal control applied to a SEIR model of 2019-nCoV with social distancing," *medRxiv*, apr 2020.

[6] L. Guan, C. Prieur, L. Zhang, C. Prieur, D. Georges, and P. Bellemain, "Transport effect of covid-19 pandemic in france," *Annual Reviews in Control*, vol. 50, pp. 394–408, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1367578820300663

[7] R. Sameni, "Mathematical Modeling of Epidemic Diseases; A Case Study of the COVID-19 Coronavirus," 2020. [Online]. Available: https://arxiv.org/abs/2003.11371

[8] Reza Sameni, *Open-access codes for the mathematical modeling of epidemic diseases*, 2021. [Online]. Available: https://github.com/alphanumericslab/EpidemicModeling.git

[9] E. Dong, H. Du, and L. Gardner, "An interactive web-based dashboard to track COVID-19 in real time," *The Lancet Infectious Diseases*, vol. 20, no. 5, pp. 533–534, May 2020. [Online]. Available: https://doi.org/10.1016/s1473-3099(20)30120-1

[10] D. S. Naidu, *Optimal control systems*. CRC press, 2003.

[11] X. Wang, "Solving optimal control problems with MATLAB: Indirect methods," *ISE Dept., NCSU, Raleigh, NC*, vol. 27695, 2009. [Online]. Available: http://solmaz.eng.uci.edu/Teaching/MAE274/SolvingOptContProb_MATLAB.pdf

[12] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Dover Publications, Inc., 1979.

[13] M. S. Grewal, L. R. Weill, and A. P. Andrews, *Global positioning systems, inertial navigation, and integration*. John Wiley & Sons, 2007.

[14] E. Petersen, M. Koopmans, U. Go, D. H. Hamer, N. Petrosillo, F. Castelli, M. Storgaard, S. A. Khalili, and L. Simonsen, "Comparing SARS-CoV-2 with SARS-CoV and influenza pandemics," *The Lancet Infectious Diseases*, vol. 20, no. 9, pp. e238–e244, sep 2020.