

From Prediction to Prescription: AI-Based Optimization of Non-Pharmaceutical Interventions for the COVID-19 Pandemic

Risto Miikkulainen^{1,2}, Olivier Francon¹, Elliot Meyerson¹,
Xin Qiu¹, Elisa Canzani¹, and Babak Hodjat¹

¹Cognizant Technology Solutions and ²The University of Texas at Austin

Abstract

Several models have been developed to predict how the COVID-19 pandemic spreads, and how it could be contained with non-pharmaceutical interventions (NPIs) such as social distancing restrictions and school and business closures. This paper demonstrates how evolutionary AI could be used to facilitate the next step, i.e. determining most effective intervention strategies automatically. Through evolutionary surrogate-assisted prescription (ESP), it is possible to generate a large number of candidate strategies and evaluate them with predictive models. In principle, strategies can be customized for different countries and locales, and balance the need to contain the pandemic and the need to minimize their economic impact. While still limited by available data, early experiments suggest that workplace and school restrictions are the most important and need to be designed carefully. As more data becomes available, the approach can be increasingly useful in dealing with COVID-19 as well as possible future pandemics.

1 Introduction

The COVID-19 crisis is unprecedented in modern times, and caught the world largely unprepared. Since there is little experience and guidance, authorities have been responding in a variety of ways. Many different non-pharmaceutical interventions (NPIs) have been implemented at different stages of the pandemic and in different contexts. On the other hand, compared to past pandemics, for the first time almost real-time data is collected about these interventions, their economic impact, and the spread of the disease. These two factors create an excellent opportunity for computational modeling and machine learning.

Most of the modeling efforts so far have been based on traditional epidemiological methods, such as compartmental models [4]. Such models can be used to predict the spread of the disease, assuming a few parameters such as the basic reproduction number R_0 can be estimated accurately. New ideas have also emerged, including using cell-phone data to measure social distancing [66]. These models have been extended with NPIs by modifying the transmission rates: each NPI is assumed to reduce the transmission rate by a certain amount [60, 27, 15]. Such models have received a lot of attention: in this unprecedented situation, they are our only source of support for making informed decisions on how to reduce and contain the spread of the disease.

However, epidemiological models are far from perfect. Much about how the disease is transmitted, how prevalent it is in the population, how many people are immune and the strength of the immunity, is unknown, and it is difficult to parameterize the models accurately. Similarly, the

effects of NPIs are unpredictable in that their effect varies based on the cultural and economic environment, the stage of the pandemic, and above all, they interact in nonlinear ways. To overcome the uncertainty, data is crucial. Instead of estimating parameters, it is possible to fit the models to the data so that they predict already existing data more accurately. In the extreme, with enough data, it is possible to use machine learning simply to model the data. All the unknown epidemiological, cultural, and economic parameters and interactions are expressed in the time series of infections and NPIs. The fact that the epidemic spreads to different regions at a lag, and different geographies experience different stages of the spread at different times, and often react to it differently, gives us the opportunity to 'front-run' the epidemic with data-driven models. Machine learning can then be used to construct a model, such as a recurrent neural network, that accurately predicts the outcomes without having to understand precisely how they they emerge.

The data-driven modeling approach is implemented and evaluated in this paper. An LSTM neural network model [23, 19] is trained with publicly available data on infections and NPIs [22] in a number of countries and applied to predicting how the pandemic will unfold in them in the future. The predictions are cascaded one day at a time and constrained to a meaningful range. Even with current limited data, the predictions are surprisingly accurate and well-behaved. This result suggests that the data-driven machine learning approach is potentially a powerful new tool for epidemiological modeling. This is the first main contribution of the paper.

A more significant contribution, however, is to demonstrate that machine learning can also be used to take the next step, i.e. to extend the models from prediction to prescription. That is, given that we can predict how the NPIs affect the pandemic, we can also automatically discover effective NPI strategies. The technology required for this step is different from standard machine learning. The goal is not to model and predict processes for which data already exists, but to create new solutions that may have never existed before. In other words, it requires extending AI from imitation to creativity.

This extension is indeed underway in AI through several approaches such as reinforcement learning, Bayesian parameter optimization, gradient-based approaches, and evolutionary computation [44, 11, 39]. The approach taken in this paper is based on evolutionary surrogate-assisted prescription [16], a technique that combines evolutionary search with surrogate modeling (Figure 1).

In ESP, a predictive model is first formed through standard machine learning techniques, such as neural networks. Given actions taken in a given context (such as NPIs at a given stage of the pandemic), it predicts what the outcomes would be (such as infections, deaths, and economic cost). A prescriptive model, another neural network, is then formed to implement an optimal decision strategy, i.e. what actions should be taken in each context. Since optimal actions are not known, the Prescripator cannot be trained with standard supervised learning. However, it can be evolved, i.e. discovered through population-based search. Because it is often impossible or prohibitively costly to evaluate each candidate strategy in the real world, the Predictor model is used as a surrogate. In this manner, millions of candidate strategies can be generated and tested in order to find ones that optimize the desired outcomes.

ESP has been used in several real-world design and decision optimization problems, including discovering growth recipes for agriculture [29], and designs for e-commerce websites [45]. It often discovers effective solutions that are overlooked by human designers. Recently it has also been




Figure 1: *Elements of the ESP Decision Optimization Method.* A Predictor is trained with historical data on how given actions in given contexts led to specific outcomes. For instance in the NPI optimization problem, given the state of pandemic in a particular country and the NPIs in effect, it predicts the future number of cases and deaths and the cost of the NPIs. The Predictor can be any machine learning model trained with supervised methods, such as a random forest or a neural network, or even a simulation such as an epidemiological model. The Predictor is then used as a surrogate in order to evolve a Prescriptors, i.e. neural networks that implement decision policies (i.e. NPIs) resulting in best possible outcomes. With multiple conflicting objectives (such as cases/deaths and cost), evolution results in multiple Prescriptors, each representing a different tradeoff, from which a human decision maker can choose the ones that best matches their preferences.

applied to sequential decision-making tasks, and found to be more sample-efficient, reliable, and safe than standard reinforcement learning techniques [16]. Much of this performance is due to a surprising regularization effect that incompletely trained Predictors and Prescriptors bring about.

The ESP approach is applied in this paper to the problem of determining optimal NPIs for the COVID-19 pandemic. Using the data-driven LSTM model as the Predictor, a Prescriptor is evolved in a multi-objective setting to minimize the number of COVID-19 cases, as well as the number and stringency of NPIs (representing economic impact). In this process, evolution discovers a Pareto front of Prescriptors that represent different tradeoffs between these two objectives: Some utilize many NPIs to bring down the number of cases, and others minimize the number of NPIs with a cost of more cases. Therefore, the AI system is not designed to replace human decision makers, but instead to empower them: Humans choose which tradeoffs are the best, and the AI makes suggestions on how they can be achieved. It therefore constitutes a step towards using AI not just to model the pandemic, but to help contain it—which has so far been missing from the literature [8].

The current implementation should be taken as a proof of concept i.e. a demonstration of the potential power of the approach. The currently available data is still limited in quantity, accuracy, and detail. It is not yet possible to draw specific prescriptions reliably, such that e.g. in a particular country, a particular NPI can be safely lifted on a particular date. The results so far suggest that such prescriptions may become possible in the next few months, as the quality and quantity of the data improves. The experiments already point to two general conclusions. First, school and

workplace closings turn out to be the two most important NPIs in the simulations: they have the largest and most reliable effects on the number of cases compared to e.g. restrictions on gatherings and travel. Second, partial or alternating NPIs may be effective. Prescriptors repeatedly turn certain NPIs on and off over time, for example, schools opening and closing on a weekly basis seems to imply the need for restricting schools to be opened fewer days per week. This is a creative and surprising solution, given the limited variability of NPIs that is available to the Prescriptors. Together these conclusions already suggest a possible focus for efforts on establishing and lifting NPIs, in order to achieve maximum effect and minimal cost.

The paper begins with a background of epidemiological modeling and the general ESP approach. The datasets used, the data-driven predictive model developed, and the evolutionary optimization of the Prescriptor will then be described. Illustrative results will be reviewed in several cases, drawing general conclusions about the potential power of the approach. Future work on utilizing epidemiological models, supporting interactive exploration, modeling uncertainty, and generating explainable prescriptions will be discussed.

An interactive demo, allowing users to explore the Pareto front of Prescriptors on multiple countries, is available at <http://evolution.ml/esp/npis>.

2 Background: Modeling Epidemics

Different types of epidemiological models are briefly characterized, followed by a review of existing COVID-19 modeling efforts, and the emerging opportunity for machine learning models.

2.1 Types of Epidemiological Modes

Modern epidemic modeling started with the compartmental SIR model developed by McKendrick and Kermack at the beginning of the 20th century [34]. The SIR model assumes susceptible individuals (S) can get infected (I) and, after a certain infectious period, die or recover (R), becoming immune afterwards. The model then describes global transmission dynamics at a population scale as a system of differential equations in continuous time. Depending on disease characteristics, these compartments and the flow patterns between them can be further refined: For instance in the case of HIV, mixing and contact depends on age groups [43]. The main limitation of such metapopulation models is that random mixing is limited between individuals within population subgroups, i.e. compartments.

Newman [50] showed that contact patterns can be represented more accurately through a network topology, taking into account geography, demographics, and social factors, and thus overcoming limitations of compartmental models. Keeling and Eames [33] further highlighted the stochastic nature of transmissions by demonstrating SIR epidemics on five types of social and spatial networks for a population of 10,000 individuals. Later studies focused on evolutionary and adaptive networks [20], aiming to model the dynamics of social links, such as frequency, intensity, locality, and duration of contacts, which strongly influence the long term impacts of epidemics.

Indeed, it is now widely recognized that multiple factors at different levels influence how epidemics spread [62]. Models have become more detailed and sophisticated, relying on extensive

computational power now available to simulate them. In addition to compartmental and network models, agent-based simulations have emerged as a third simulation paradigm. Agent-based approaches describe the overall dynamics of infection as result of events and activities involving single individuals [65], resulting in potentially detailed but computationally demanding processes. For a more detailed review of the literature on epidemiological models and their mathematical insights see [3, 42].

2.2 Modeling the COVID-19 Pandemic

A variety of epidemiological simulations are currently used to model the COVID-19 pandemic. The focus is on simulating effects of NPIs in order to support decision making about response policies.

For instance, Stanford University [60] extended the SIR model up to nine compartments including susceptible, exposed, asymptomatic, presymptomatic, mild and severe symptomatic, hospitalized, deceased, and recovered populations, as well as a stochastic simulator for transitions between compartments, calibrated with MIDAS data [47]. Currently, the simulator supports up to three interventions at different times. In contrast, NPIs are implemented in a more granular fashion in the Bayesian inference model of Imperial College London [15]. Their model parameters are estimated empirically from ECDC data [13] for 11 countries. Deaths are then predicted as a function of infections based on the distribution of average time between infections and time-varying reproduction number R_t .

The Institute for Health Metrics and Evaluation (IHME) of Washington University [27] combined compartmental models and mixed-effects nonlinear regression to predict cumulative and daily death rate as a function of NPIs. They also forecast health service needs using a micro-simulation model. The IHME dashboard [64] compares estimated and confirmed daily infections upon testing rate for both U.S. and European countries up to state- and regional levels. Similarly, the University of Texas [66] developed a statistical model based on nonlinear Bayesian hierarchical regression with a negative-binomial model for daily variation in death rates. The novelty is to estimate social distancing using geolocation data from mobile phones, improving accuracy over the IHME model. So far, the forecasts rely on U.S. data and are not reliable beyond about weeks.

More broadly, the Centers for Disease Control and Prevention (CDC) [4] provides details about the different COVID-19 prediction models and their specific NPI assumptions. In fact, CDC works with partners to bring together forecasts for cumulative deaths over the following four weeks. Forecasting teams predict numbers of deaths using different types of data (e.g. WHO, COVID-19, demographic, mobility), methods, and estimates of the impacts of NPIs (e.g. social distancing and use of face coverings). In general, most of COVID-19 forecast approaches use curve-fitting and ensembles of mechanistic SIR models with age structures and different parameter assumptions. Social distancing and NPIs are usually not represented directly, but approximated as changes in transmission rates. The main advantage is that running simulations require a few input parameters based on average data at population scale.

In contrast, since contact dynamics in agent-based and network approaches results from events and activities of single individuals and their locations, they can be more accurate in modeling social distancing and NPIs. However, their parameters need to be calibrated appropriately, which is difficult to do with available data. Mechanistic transmission models can help overcome data col-

lection challenges when tracing a real network of individuals. Different tracing techniques, such as comprehensive diary-based studies (e.g. POLYMOD [48]), and recorded movements of individuals (e.g. transportation networks [25] and dollar bill tacking [2]) have been investigated in the past to sample real networks with limited resources and data availability. In the context of COVID-19, as mentioned above, a new opportunity is to use mobile phone data to support implementation of social distancing measures (e.g. [66]). Debates on value and ethics of tracking people movements to monitor the COVID-19 are still ongoing [10]. Another challenge for current epidemiological modeling methods is that they require significant computing resources and sophisticated parallelization algorithms. As a step to making them feasible, EpiFast [1] reduces the SIR epidemics simulation problem to a sequence of graph operations on distributed memory systems, and thereby significantly decreases the simulation cost. The COVID-19 pandemic has accelerated efforts to develop solutions to these challenges, and is likely to result in improved models in the future.

2.3 Opportunity for Machine Learning Models

Any of the above models that include NPI effects and generate long-term predictions could be used as the Predictor with ESP, even several of them together as an ensemble. Taking advantage of them is indeed a compelling direction of future work (Section 8). However, this paper focuses on evaluating a new opportunity: Building the model solely based on past data using machine learning. Given that data about the COVID-19 pandemic is generated, recorded, and made available more than any epidemic before, such a new approach may be feasible for the first time.

There is a lot of promise in this data-driven approach. The epidemiological models require several assumptions about the population, culture, and environment, depend on several parameters that are difficult to set accurately, and cannot take into account many possible nonlinear and dynamic interactions between the NPIs, and in the population. In contrast, all such complexities are implicitly included in the data. The data-driven models are phenomenological, i.e. they do not explain how the given outcomes are produced, but given enough data, they can be accurate in predicting them. This is the first hypotheses tested in this paper; as shown in Section 5, it turns out that even with the limited data available at this point, data-driven models can be useful.

All the models discussed so far are predictive: Based on knowledge of the populations and the epidemic, and the data so far about its progress in different populations and efforts to contain it, they estimate how the disease will progress in the future. By themselves, these models do not make recommendations, or prescriptions, of what NPIs would be most effective. It is possible to manually set up hypothetical future NPI strategies and use the models to evaluate how well they would work. However, only a few strategies can be tested in this manner, and the process is limited by the ability of human experts to think of promising strategies. Given past experience with surrogate modeling and population-based search, automated methods may be more effective in this process. A method for doing so, ESP, will be described next.

3 Method: Evolutionary Surrogate-assisted Prescription (ESP)

ESP is a continuous black-box optimization process for adaptive decision-making [16]. In ESP, a model of the problem domain is used as a surrogate for the problem itself. This model, called the Predictor (P_d), takes a decision as its input, and predicts the outcomes of that decision. A decision consists of a context (i.e. a problem) and actions to be taken in that context.

A second model, called the Prescriptor (P_s), is then created. It takes a context as its input, and outputs actions that would optimize outcomes in that context. In order to develop the Prescriptor, the Predictor is used as the surrogate, i.e. a less costly and risky alternative to the real world.

More formally, given a set of possible contexts \mathbb{C} and possible actions \mathbb{A} , a decision policy D returns a set of actions A to be performed in each context C :

$$D(C) = A, \quad (1)$$

where $C \in \mathbb{C}$ and $A \in \mathbb{A}$. For each such (C, A) pair there is a set of outcomes $O(C, A)$, and the Predictor P_d is defined as

$$P_d(C, A) = O, \quad (2)$$

and the Prescriptor P_s implements the decision policy as

$$P_s(C) = A, \quad (3)$$

such that $\sum_{i,j} O_j(C_i, A_i)$ over all possible contexts i and outcome dimensions j is maximized (assuming they improve with increase). It thus approximates the optimal decision policy for the problem. Note that the optimal actions A are not known, and must therefore be found through search.

In the case of the NPI optimization problem, context C consists of information regarding a region. This might include data on the number of available ICU beds, population distribution, time since the first case of the disease, current COVID-19 cases, and fatality rate. Actions A in this case specify whether or not the different possible NPIs are implemented within that region. The outcomes O for each decision measure the number of cases and fatalities within two weeks of the decision, and the cost of each NPI.

The ESP algorithm then operates as an outer loop that constructs the Predictor and Prescriptor models (Figure 2):

1. Train a Predictor based on historical training data;
2. Evolve Prescriptors with the Predictor as the surrogate;
3. Apply the best Prescriptor in the real world;
4. Collect the new data and add to the training set;
5. Repeat.

In the case of NPI optimization, there is currently no Step 3 since the system is not yet incorporated into decision making. However, any NPIs implemented in the real world, whether similar or dissimilar to ESP's prescriptions, will similarly result in new training data. As usual in evolutionary

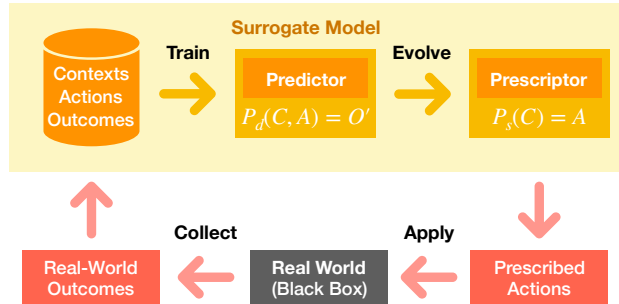


Figure 2: *The ESP Outer Loop*. The Predictor can be trained gradually at the same time as the Prescriptor is evolved, using the Prescriptor to drive exploration. That is, the user can decide to apply the Prescriptor’s outputs to the real world, observe the outcomes, and aggregate them into the Predictor’s training set. However, any new prescriptions implemented in the real world, whether similar to the Prescriptors or not, can be used to augment the training dataset.

search, the process terminates when a satisfactory level of outcomes has been reached, or no more progress can be made, or the system iterates indefinitely, continuously adapting to changes in the real world (e.g., adapting to the advent of vaccines or antiviral drugs).

The Predictor model is built by modeling a (C, A, O) dataset. The choice of algorithm depends on the domain, i.e. how much data there is, whether it is continuous or discrete, structured or unstructured. Random forests, symbolic regression, and neural networks have been used successfully in this role in the past [16, 29]. In some cases, such as NPI optimization, an ensemble of data-driven and simulation models may be useful, in order to capture expected or fine-grained behavior that might not yet have been reflected in the data (Section 8).

The Prescriptor model is built using neuroevolution: Neural networks because they can express complex nonlinear mappings naturally, and evolution because it is an efficient way to discover such mappings [61], and naturally suited to optimize multiple objectives [7, 12]. Because it is evolved with the Predictor, the Prescriptor is not restricted by a finite training dataset, or limited opportunities to evaluate in the real world. Instead, the Predictor serves as a fitness function, and it can be queried frequently and efficiently. In a multiobjective setting, ESP produces multiple Prescriptors, selected from the Pareto front of the multiobjective neuroevolution run. The Prescriptor is the novel aspect of ESP: it makes it possible to discover effective solutions that do not already exist, even those that might be overlooked by human decision makers.

In past work, ESP was found to be effective in standard reinforcement learning benchmarks [16], as well as in real-world applications of discovering recipes for growing plants in controlled environments [29] and in designing e-commerce websites [45]. The benchmarks were useful because ESP could be evaluated in them as an autonomous decision making system. However, in real-world applications, including NPI optimization, it is most naturally used as a system for augmenting human decision making. To support this role, it is possible to include a ‘scratchpad’ functionality, whereby the human decision maker can see the predicted outcomes of the ESP-prescribed actions, as well as modify the prescriptions and weigh the outcomes as part of their decision making. A

scratchpad will be an integral part of the NPI interactive tool as well (Section 8).

Another helpful extension of ESP is to automatically estimate the confidence in the predicted outcomes. With neural networks, softmax of the output is often used as an estimation of confidence, but this measure is often inaccurate [17]. A better approach is to build a certainty estimation model to complement the Predictor through a Gaussian process of training data residuals using input and output kernels [57]. This approach will be described in more detail in Section 5.3.

In the NPI optimization task, ESP is built to prescribe the NPIs for the current day such that the number of cases and cost that would result in the next two weeks is optimized. The details of the Predictor and Prescriptor in this setup will be described next, after an overview of the data used to construct them.

4 Data

Even though COVID-19 is not the first global pandemic, it is the first that is recorded in significant detail, providing data that is made publicly available daily as the pandemic unfolds.

The earliest data that became available included the number of confirmed cases, the number of deaths, and the number of recovered patients, per country, region, and day. A well known such source is the Johns Hopkins University (JHU [28]), updated daily and quoted widely in the press. Several predictive models have built based on this data, and there is even a Kaggle competition to predict daily confirmed cases using it [30]. The Kaggle site encourages companies and organizations to contribute other useful datasets. Other data is available about e.g. the population and medical system in each country, but they have not yet been used to inform the models.

Once there was enough data to model how the pandemic spreads, an important question started to emerge: what can we do to contain it? Pharmaceutical interventions such as treatments and vaccines will take time to develop, so the focus has been on implementing non-pharmaceutical interventions, i.e. NPIs. The goal is to "flatten the curve," i.e. limit the spread, gain time, and prevent hospitals from being overwhelmed until a vaccine can be developed [14, 15]. Augmenting the health data with data on NPIs, it would be possible to learn how the NPIs affect health outcomes. For instance, since the pandemic affected China and other parts of Asia before the rest of the world, it would be possible to learn from their examples.

However, compiling NPI data turned out to be more difficult, and took longer. Each country took different actions, at different levels, in different cities or regions. These decision were reported in the press, but initially were not aggregated and normalized, making it hard to form a dataset. For instance in the US, such datasets started to come out only in April [31, 35, 63]. Based on data from the Covid Tracking Project [63], University of Washington's Institute for Health Metrics and Evaluation (IHME) developed a dashboard that shows the NPI timeline [26]. Updated continuously, the dashboard shows a projection for daily deaths and estimated infections (estimated infections are higher than confirmed infections due to limited testing). The model includes a social distancing factor computed from the NPIs, held constant for the future. At a more global scale, Oxford University's Blavatnik School of Government provides a dataset of the number of cases, deaths and NPIs for most countries on a daily basis [22, 52, 56]. It initially included six 'Closure and Containment' NPIs, but on April 29 was extended to eight NPIs with more granular intervention

NPI name	Level 0	Level 1	Level 2	Level 3	Level 4
C1.School closing	no measures	recommend closing	require closing (only some levels or categories, e.g. just high school, or just public schools)	require closing all levels	
C2.Workplace closing	no measures	recommend closing (or recommend work from home)	require closing (or work from home) for some sectors or categories of workers	require closing (or work from home) for all-but-essential workplaces (eg grocery stores, doctors)	
C3.Cancel public events	no measures	recommend cancelling	require cancelling		
C4.Restrictions on gatherings	no restrictions	restrictions on very large gatherings (the limit is above 1000 people)	restrictions on gatherings between 101-1000 people	restrictions on gatherings between 11-100 people	restrictions on gatherings of 10 people or less
C5.Close public transport	no measures	recommend closing (or significantly reduce volume/route/means of transport available)	require closing (or prohibit most citizens from using it)		
C6.Stay at home requirements	no measures	recommend not leaving house	require not leaving house with exceptions for daily exercise, grocery shopping, and 'essential' trips	require not leaving house with minimal exceptions (e.g. allowed to leave once a week, or only one person can leave at a time, etc)	
C7.Restrictions on internal movement	no measures	recommend not to travel between regions/cities	internal movement restrictions in place		
C8.International travel controls	no restrictions	screening arrivals	quarantine arrivals from some or all regions	ban arrivals from some regions	ban on all regions or total border closure

Table 1: *Definitions of the NPI data used in the experiments.* The data consists of daily 'Closure and containment' measures in the 'Coronavirus Government Response Tracker' provided by Oxford University [22, 56, 51]. There are eight kinds of NPIs, each ranging in stringency from 0 (no measures) to 2, 3, or 4 (full measures). Together with data on daily cases, this data is used to train the Predictor model.

levels. A detailed explanation of the data is provided in a codebook [51].

The Oxford dataset was used as a source in the current ESP study. The models were trained using the 'ConfirmedCases' data for the cases and 'Closure and Containment' data for the NPIs. The other NPI categories in the dataset, i.e. 'Economic response', 'Public Health' and 'Miscellaneous' measures because they have less direct impact on the spread of the epidemic. A summary of these NPIs is given in Table 1.

The number of cases was selected as the target for the predictions (instead number of deaths, which is generally believed to be more reliable), because case numbers are higher and the data is smoother overall. The model also utilizes a full 21-day case history which it can use to uncover structural regularities in the case data. For instance, it discovers that many fewer cases are reported on the weekends in France and Spain. However, the data is still noisy for several reasons:

- There are other differences in how cases are reported in each country;
- Some countries, like the US, do not have a uniform manner of reporting the cases;

- Cases were simply not detected for a while, and testing policy still widely differs from country to country.
- Some countries, like China, US, and Italy, implemented NPIs at a state / regional level, and it is difficult to express them at the country level;
- As usual with datasets, there are mistakes, missing days, double-counted days, etc.

It is also important to note that there is roughly a two-week delay between the time a person is infected and the time the case is detected. A similar delay can therefore be expected between the time an NPI is put in places and its effect on the number of cases.

Despite these challenges, it is possible to use the data to train a useful model to predict future cases. This data-driven machine learning approach will be described next.

5 Data-Driven Predictive Model

With the above data sources, machine learning techniques can be used to build a predictive model. Good use of recent deep learning approaches to sequence processing can be made in this process, in particular recurrent neural networks. However, a method of cascading the predictions needs to be developed so that they can reach several steps into the future. Furthermore, methods are needed that keep the predictions within a sensible range even with limited data.

5.1 Predictor Model Design

This section describes the step-by-step design of the learned predictor. For a given country, let x_n be the number of new cases on day n . The goal is to predict x_n in the future. First, consider the minimal epidemic model

$$x_n = R_n x_{n-1} \implies R_n = \frac{x_n}{x_{n-1}}, \text{ for some } R_n \geq 0. \quad (4)$$

where the factor R_n is to be predicted. Focusing on such factors is fundamental to epidemiological models, and, when learning a predictive model from data, makes it possible to normalize prediction targets across countries and across time, thereby simplifying the learning task.

Training targets R_n can be constructed directly from daily case data for each country. However, in many countries case reporting is noisy and unreliable, leading to unreasonably high noise in daily R_n . This effect can be mitigated by instead forming smoothed targets based on a moving average z_n of new cases:

$$z_n = R_n z_{n-1} \implies R_n = \frac{z_n}{z_{n-1}}, \quad \text{where} \quad z_n = \frac{1}{K} \sum_{i=0}^{K-1} x_{n-i}. \quad (5)$$

In this paper, $K = 7$ for all models, i.e. prediction targets are smoothed over the preceding week.

To capture the effects of finite population size and immunity, an additional factor is included that scales predictions by the proportion of the population that could possibly become new cases:

$$z_n = \frac{P - y_{n-1}}{P} R_n z_{n-1} \implies R_n = \frac{P z_n}{(P - y_{n-1}) z_{n-1}}, \quad (6)$$

where P is the population size, and $y_n = \sum_{i=0}^n x_i$ is the total number of recorded cases by day n . Notice that, when evaluating a trained model, the predicted \hat{x}_n can be recovered from a predicted \hat{R}_n by

$$\hat{x}_n = (\hat{R}_n \frac{P - y_{n-1}}{P} - 1) K z_{n-1} + x_{n-K}. \quad (7)$$

Note that this formulation assumes that recovered cases are fully immune: When $P = y_{n-1}$, the number of new cases must go to 0. This assumption can be relaxed in the future by adding a factor to Equation 6 (either taken from the literature or learned) to represent people who were infected and are no longer immune.

The trainable function implementing \hat{R}_n can now be described. The prediction \hat{R}_n should be a function of (1) NPIs enacted over previous days, and (2) the underlying state of the pandemic distinct from the enacted NPIs. For the models in this paper, (1) is represented by the NPI restrictiveness values for the past $T = 21$ days over all $N = 8$ available NPIs, and (2) is represented autoregressively by the T previous values of R_n (or, during forecasting, by the predicted \hat{R}_n when the true R_n is unavailable). Formally,

$$\hat{R}_n = f(\mathbf{A}_n, \mathbf{r}_n), \quad \text{with} \quad \mathbf{A}_n \in \mathbb{N}_0^{T \times N} \quad \text{and} \quad \mathbf{r}_n \in \mathbb{R}_{\geq 0}^T. \quad (8)$$

In contrast to epidemiological models that make predictions based on today’s state only, this data-driven model predicts based on data from the preceding three weeks.

To help the model generalize with a relatively small amount of training data, the model is made more tractable by decomposing f with respect to its inputs:

$$\hat{R}_n = f(\mathbf{A}_n, \mathbf{r}_n) = (1 - g(\mathbf{A}_n)) h(\mathbf{r}_n), \quad \text{with} \quad g(\mathbf{A}_n) \in [0, 1] \quad \text{and} \quad h(\mathbf{r}_n) \geq 0. \quad (9)$$

Here, the factor $g(\mathbf{A}_n)$ can be viewed as the effect of social distancing (i.e. NPIs), and $h(\mathbf{r}_n)$ as the endogenous growth rate of the disease.

To make effective use of the nonlinear and temporal aspects of the data, both g and h are implemented as LSTM models [23], each with a single LSTM layer of 32 units, followed by a dense layer with a single output. To satisfy their output bounds, the dense layers of g and h are followed by sigmoid and softplus activation, respectively.

Importantly, the factorization of f into g and h makes it possible to explicitly incorporate the constraint that *increasing the stringency of NPIs cannot decrease their effectiveness*. This idea is incorporated by constraining g to be monotonic with respect to each NPI, i.e.

$$\min(\mathbf{A} - \mathbf{A}') \geq 0 \implies g(\mathbf{A}) \geq g(\mathbf{A}'). \quad (10)$$

This constraint is enforced by requiring all trainable parameters of g to be non-negative, except for the single bias parameter in its dense layer. This non-negativity is implemented by setting all trainable parameters to their absolute value after each parameter update.




Figure 3: *The Predictor Neural Network.* This diagram shows the Keras representation of the learnable predictor model. The previous 21 days of R_{n-t} are fed into the context_input; the previous 21 days of stringency values for the eight NPIs are fed into the action_input. The Lambda layer combines the context branch h and the action branch g as specified in Equation 9 to produce a prediction \hat{R}_n . The effects of social distancing and endogenous growth rate of the pandemic are processed in separate pathways, making it possible to ensure that stringency has a monotonic effect, resulting in more regular predictions.

Note that although the model is trained only to predict one day in the future, it can make predictions arbitrarily far into the future given a schedule of NPIs by autoregressively feeding the predicted \hat{R}_{n+t} back into the model as input.

For the experiments in this paper, the model for f was implemented in Keras [5]. The Keras diagram of the model is shown in Figure 3. The model is trained end-to-end to minimize mean absolute error (MAE) with respect to targets R_n using the Adam optimizer [37] with default parameters and batch size 32. MAE was used instead of mean squared error (MSE) because it is more robust to remaining structural noise in the training data. The last 14 days of data were withheld from the dataset for testing. For the remaining data, the R_n were clipped to the range $[0, 2]$ to handle extreme outliers, and randomly split into 90% for training and 10% for validation during training. The model was trained until validation MAE did not improve for 20 epochs, at which point the weights yielding the best validation MAE were restored. Since the model and dataset are relatively small compared to common deep learning datasets, the model is relatively inexpensive to train. On a 2018 MacBook Pro Laptop with six Intel i7 cores, the model takes 276 ± 31 seconds to train (mean and std. err. computed over 10 independent training runs).

5.2 Predictor Empirical Evaluation

To validate the factored monotonic LSTM (NPI-LSTM) predictor design described above, it was compared to a suite of baseline machine learning regression models. These baselines included linear regression, random forest regression (RF), support vector regression (SVR) with an RBF kernel, and feed-forward neural network regression (MLP). Each baseline was implemented with sci-kit learn, using their default parameters [54]. Each method was trained independently 10 times on the training dataset described in Section 5.1. The results on the test dataset (last $T^* = 14$ days of the $C = 20$ countries with the most cases) were evaluated with respect to four complementary performance metrics. In particular, for the comparisons in this section, training data consisted of data up until May 6, 2020, and test data consisted of data from May 7 through May 20, 2020.

Suppose training data ends on day n . Let \hat{R}_{n+t}^c and \hat{x}_{n+t}^c be the model output and the corresponding predicted new cases (recovered via Equation 7) for the c th country at day $n + t$. The

metrics were:

1-step \hat{R}_n MAE This metric is simply the loss the models were explicitly trained to minimize, i.e. minimize $|R_n - \hat{R}_n|$ given the ground truth for the previous 21 days:

$$\frac{1}{CT^*} \sum_{c=1}^C \sum_{t=1}^{T^*} |R_{n+t}^c - \hat{R}_{n+t}^c|. \quad (11)$$

The remaining three metrics are based not on single-step prediction, but the complete 14 day forecast for each country:

Raw Case MAE This is the most intuitive metric, included as an interpretable reference point. It is simply the MAE w.r.t. new cases over the 14 test days summed over all 20 test countries:

$$\sum_{c=1}^C \left| \sum_{t=1}^{T^*} x_{n+t}^c - \sum_{t=1}^{T^*} \hat{x}_{n+t}^c \right|. \quad (12)$$

Normalized Case MAE This metric normalizes the case MAE of each country by the number of true cases in the 14 day window, so that errors are in a similar range across countries. Such normalization is important for aggregating results over countries that have different population sizes, or are in different stages of the pandemic:

$$\frac{1}{C} \sum_{c=1}^C \frac{\left| \sum_{t=1}^{T^*} x_{n+t}^c - \sum_{t=1}^{T^*} \hat{x}_{n+t}^c \right|}{\sum_{t=1}^{T^*} x_{n+t}^c}. \quad (13)$$

Mean Rank This metric ranks the methods in terms of case error for each country, and then averages over countries. It indicates how often a method will be preferred over others on a country-by-country basis:

$$\frac{1}{C} \sum_{c=1}^C \text{rank} \left(\left| \sum_{t=1}^{T^*} x_{n+t}^c - \sum_{t=1}^{T^*} \hat{x}_{n+t}^c \right| \right), \quad (14)$$

where $\text{rank}(\cdot)$ returns the rank of the error across all five methods, i.e. the method with the lowest error receives rank of 0, the next-best method receives rank of 1, and so on.

Of these four metrics, Normalized Case MAE gives the most complete picture of how well a method is doing, since it combines detailed case information of Raw Case MAE with fairness across countries similar to Mean Rank. The results are shown in Table 2. NPI-LSTM outperforms the baselines on all metrics. Interestingly, although RF and SVR do quite well in terms of the loss on which they were trained (1-step \hat{R}_n MAE), the simple linear model outperforms them substantially on the metrics that require forecasting beyond a single day, showing the difficulty that off-the-shelf nonlinear methods have in handling such forecasting.

Method	Norm. Case MAE	Raw Case MAE	Mean Rank	1-step \hat{R}_n MAE
MLP	2.47±1.22	1089126±540789	3.19±0.09	0.769±0.033
RF	0.95±0.05	221308±8717	1.98±0.10	0.512±0.000
SVR	0.71±0.00	280731±0	1.76±0.09	0.520±0.000
Linear	0.64±0.00	176070±0	1.63±0.09	0.902±0.000
NPI-LSTM	0.42±0.04	154194±14593	1.46±0.08	0.510±0.001

Table 2: *Performance comparison of proposed predictor (NPI-LSTM) with baselines.* This table shows results along the four metrics described in Section 5.2 with mean and standard error over 10 trials. Interestingly, although RF and SVR do quite well in terms of the loss they were trained on (1-step \hat{R}_n MAE), the simple linear model outperforms them substantially on the metrics that require forecasting beyond a single day, showing the difficulty that off-the-shelf nonlinear methods have in handling such forecasting. In contrast, with the extensions developed specifically for the epidemic modeling case, the NPI-LSTM methods outperforms the baselines on all metrics.

To verify that the predictions are meaningful and accurate, four example scenarios, i.e. four different countries at different stages of the pandemic, are plotted in Figure 4 (active cases at each day is approximated as the sum of new cases over the prior 14 days). Day 0 represents the point in time when 10 total cases were diagnosed; in each case, stringent NPIs were enacted soon after. The predictor was trained on data up until April 17, 2020, and the predictions started on April 18, with 21 days of data before the start day given to the predictor as initial input. Assuming the NPIs in effect on the start day will remain unchanged, it will then predict the number of cases 180 days into the future. Importantly, during the first 14 days its predictions can be compared to the actual number of cases. For comparison, another prediction plot is generated from the same start date assuming no NPIs from that date on. As can be seen from the figure, (1) the predictions match the actual progress well, (2) assuming the current stringent NPIs continue, the cases will eventually go to 0, and (3) with no NPIs, there is a large increase of cases, followed by an eventual decrease as the population becomes immune. The predictions thus follow meaningful trajectories.

The main conclusion from these experiments is that the data-driven approach works surprisingly well, even with limited data. As will be discussed in Section 8, it should be possible to improve it in the future, when more and better data becomes available. It might also be possible to combine the strengths of the different approaches to prediction and epidemiological modeling. In any case, the current predictive model already makes it possible to build Prescriptors for ESP, as will be discussed in Section 6. The way confidence can be estimated in it will be described next.

5.3 Modeling Uncertainty in Predictions

An important aspect of any decision system is to estimate confidence in its outcomes. In prescribing NPIs, this means estimating uncertainty in the Predictor, i.e. deriving confidence intervals on the predicted number of future cases. In simulation models such as those reviewed in Section 2, variation is usually created by running the models multiple times with slightly different initial conditions or parameter values, and measuring the resulting variance in the predictions. With neural network

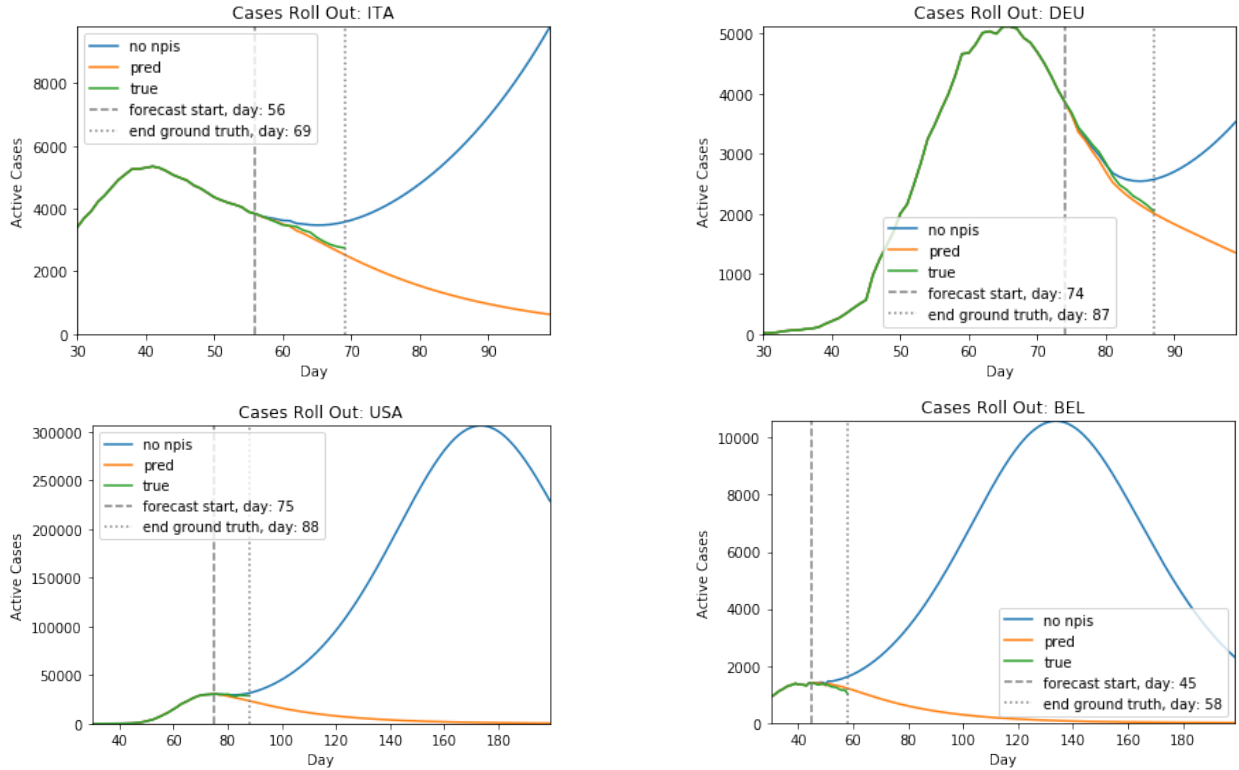


Figure 4: *Illustrating the predictive ability of the NPI-LSTM model.* Actual and projected cases are shown for four sample countries. The model predicts the number of cases accurately for the first 14 days where it can be compared with the actual future data (between the vertical lines). The prolonged 180-day predictions are also meaningful, reducing the number of cases to zero with stringent NPIs, and predicting a major increase followed by an eventual decrease with less stringent NPIs. Thus, with proper constraints, data-driven machine learning models can be surprisingly accurate in predicting the progress of the pandemic even with limited data.

predictors, it is possible to measure uncertainty more directly by combining a Bayesian model with it [49, 18, 36]. Such extended models tend to be less accurate than pure predictive models, and also harder to set up and train [17, 38].

A recent alternative is to train a separate model to estimate uncertainty in point-prediction models [57]. In this approach, called RIO, a Gaussian Process is fit to the original residual errors in the training set. The I/O kernel of RIO utilizes both input and output of the original model so that information can be used where it is most reliable. In several benchmarks, RIO has been shown to construct reliable confidence intervals. Surprisingly, it can then be used to improve the point predictions of the original model, by correcting them towards the estimated mean. RIO can be directly applied to any machine learning model without modifications or retraining. It therefore forms a good basis for estimating uncertainty in the COVID-19 Predictor as well.

Algorithm 1 Procedure of applying RIO to the COVID-19 Predictor

Require:

- $\{((\mathbf{A}_i, \mathbf{r}_i), R_i)\}_{i=1}^n$: training data
- $\hat{\mathbf{R}} = \{\hat{R}_i\}_{i=1}^n$: original predictions on training data
- $(\mathbf{A}_*, \mathbf{r}_*)$: data to be predicted
- \hat{R}_* : original prediction on $(\mathbf{A}_*, \mathbf{r}_*)$

Ensure:

- $\hat{R}'_* \sim \mathcal{N}(\hat{R}_* + \bar{\hat{e}}_*, \text{var}(\hat{e}_*))$: a distribution of calibrated prediction

Training Phase:

- 1: Calculate residuals $\mathbf{e} = \{e_i = R_i - \hat{R}_i\}_{i=1}^n$
- 2: Obtain hidden states of the two LSTM models in original predictor $\{g(\mathbf{A}_n), h(\mathbf{r}_n)\}_{i=1}^n$
- 3: **for** each optimizer step **do**
- 4: Calculate covariance matrix $\mathbf{K}_c(\{g(\mathbf{A}_n), h(\mathbf{r}_n)\}_{i=1}^n, \hat{\mathbf{R}}), (\{g(\mathbf{A}_n), h(\mathbf{r}_n)\}_{i=1}^n, \hat{\mathbf{R}})$, where each entry is given by $k_c(((g(\mathbf{A}_i), h(\mathbf{r}_i)), \hat{R}_i), ((g(\mathbf{A}_j), h(\mathbf{r}_j)), \hat{R}_j)) = k_{\text{in}}((g(\mathbf{A}_i), h(\mathbf{r}_i)), (g(\mathbf{A}_j), h(\mathbf{r}_j))) + k_{\text{out}}(\hat{R}_i, \hat{R}_j)$, for $i, j = 1, 2, \dots, n$
- 5: Optimize GP hyperparameters by maximizing log marginal likelihood
 $\log p(\mathbf{e} | \{g(\mathbf{A}_n), h(\mathbf{r}_n)\}_{i=1}^n, \hat{\mathbf{R}}) = -\frac{1}{2} \mathbf{e}^\top (\mathbf{K}_c(\cdot) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{e} - \frac{1}{2} \log |\mathbf{K}_c(\cdot) + \sigma_n^2 \mathbf{I}| - \frac{n}{2} \log 2\pi$

Deployment Phase:

- 6: Calculate residual mean $\bar{\hat{e}}_* = \mathbf{k}_*^\top (\mathbf{K}_c(\cdot) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{e}$ and residual variance $\text{var}(\hat{e}_*) = k_c(((g(\mathbf{A}_*), h(\mathbf{r}_*)), \hat{R}_*), ((g(\mathbf{A}_*), h(\mathbf{r}_*)), \hat{R}_*)) - \mathbf{k}_*^\top (\mathbf{K}_c(\cdot) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*$
 - 7: Return distribution of calibrated prediction $\hat{R}'_* \sim \mathcal{N}(\hat{R}_* + \bar{\hat{e}}_*, \text{var}(\hat{e}_*))$
-

In order to extend RIO to time-series predictions, the hidden states of the two LSTM models (before the lambda layer in Figure 3) are concatenated and fed into the input kernel of RIO. The original predictions of the predictor are used by the output kernel. RIO is then trained to fit the residuals of the original predictions. During deployment, the trained RIO model then provides a Gaussian distribution for the calibrated predictions. The details of this process are presented in Algorithm 1.

To validate this process empirically, RIO was trained and tested on COVID-19 data from 21 selected countries. The data was preprocessed in four steps: (1) Within the 30 most affected countries in terms of total confirmed cases, select the countries on which the original predictor gives MAE less than 0.07. (2) Remove the outlier days that have an R larger than 2.0. (3) Remove the data of the earliest 10 days (after the first 21 days), so that the training data are closer to the recent situations. (4) For each country, randomly select 14 days as the testing data. The hyperparameters in these steps were found to be appropriate empirically. All remaining days were then used as the training data. Table 3 shows the results of applying RIO in this manner to the original predictor. The conclusion is that RIO constructs reasonable confidence intervals at several confidence levels, and slightly improves the prediction accuracy. It can therefore be expected to work well in estimating confidence in the NPI prescription outcomes as well.

However, RIO will first need to be extended to model uncertainty in time series. Because NPI-LSTM forecasts are highly nonlinear and autoregressive, analytic methods are intractable.

Table 3: Results after applying RIO to predictor

Dataset	original MAE	MAE with RIO	95% CI Coverage	90% CI Coverage	68% CI Coverage
Training dataset	0.0379	0.0342	0.932	0.909	0.761
Testing dataset	0.0471	0.0467	0.884	0.847	0.687

CI coverage means the percentage of testing outcomes that are within the estimated confidence intervals (CIs).

Instead, given that the predictor model with RIO returns both the mean and standard deviation for \hat{R}_n , the standard deviation after t days in the future can be estimated via Monte Carlo rollouts. Specifically, for each step in each rollout, instead of predicting \hat{R} and feeding it back into the model to predict the next step, \hat{R} is sampled from the Gaussian distribution returned by RIO, and this sample is fed back into the model. Thus, after T^* steps, a sample is generated from the forecast distribution. Given several such samples, standard deviations are computed empirically for all forecasted days $1 \leq t \leq T^*$. Standard deviations are computed empirically over many samples (100 in the experiments in this paper).

Thus, RIO makes it possible to estimate uncertainty in the predictions, which in turn helps the decision maker interpret and trust the results, i.e. how reliable the outcomes are for the recommendations that the Prescriptors generate. The method for discovering good Prescriptors will be described next.

6 Evolutionary Prescriptive Model

Whereas many different models could be used as a Predictor, the Prescriptor is the heart of the ESP approach, and needs to be constructed using modern search techniques. This section describes the process of evolving neural networks for this task. A number of example strategies are presented from the Pareto front, representing trade-offs between objectives, as well as examples for countries at different stages of the pandemic, and counterfactual examples comparing possible vs. actual outcomes. General conclusions are drawn on which NPIs matter the most, and how they could be implemented most effectively.

Any of the existing neuroevolution methods [61] could be used to construct the Prescriptor as long as it evolves the entire network including all of its weight parameters: Neural architecture search cannot be used easily since there are no targets (i.e. known optimal NPIs) with which to train it with gradient descent. The most straightforward approach of evolving a vector of weights for a fixed topology is therefore used and found to be sufficient in this case. The Prescriptor model (Figure 5) is a neural network with one input layer of size 21, corresponding to case information R_{n-t} (as defined in Equation 6) for the prior 21 days. This input is the same as the context_input of the Predictor. The input layer is followed by a fully-connected hidden layer of size 32 with the tanh activation function, and eight outputs (of size one) with the sigmoid activation function. The outputs represent the eight possible NPIs, which will then be input to the Predictor. Each output is further scaled to represent the corresponding NPI stringency levels: three for 'Cancel public events', 'Close public transport', and 'Restrictions on internal movement'; four for 'School closing', 'Workplace closing', and 'Stay at home'; five for 'Restrictions on gatherings' and 'International travel controls'.

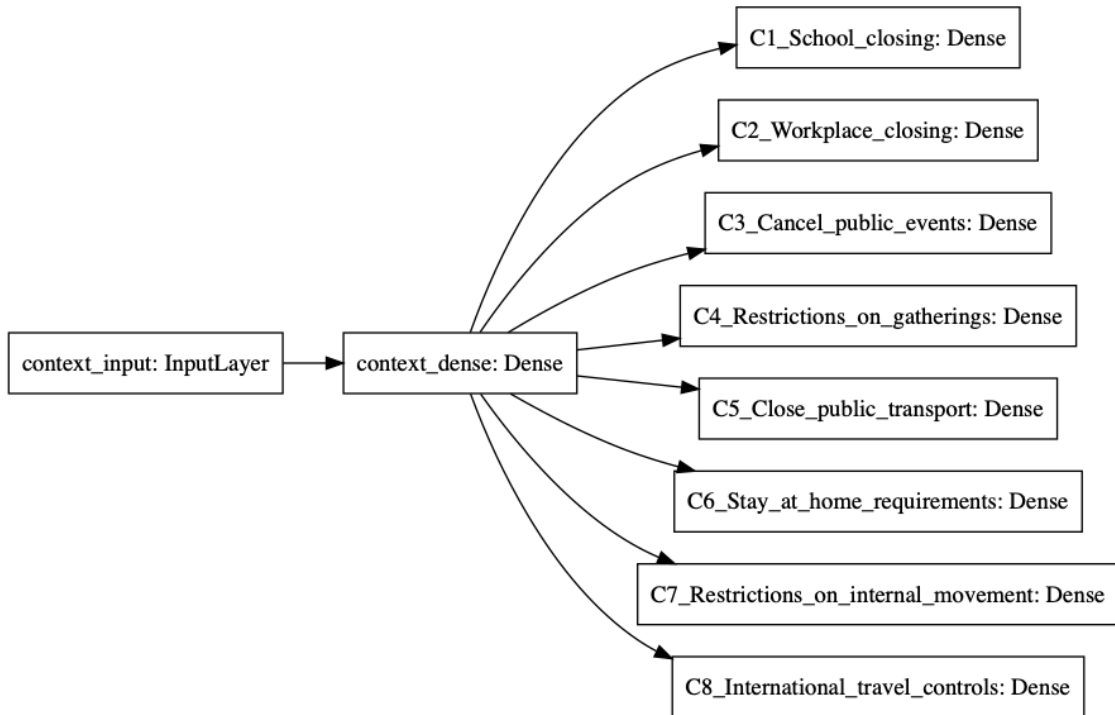


Figure 5: *The Prescriptor Neural Network*. Given 21 past days of case information (R_{n-t} in Equation 6) as input (context_input), the network generates recommended stringency values for each of the eight NPIs. The network is fully connected with one hidden layer. Because there are no targets, i.e. the optimal NPIs are not known, gradient descent cannot be used; instead, all weights and biases are evolved based on how well the network’s NPI recommendations work along the cases and cost objectives, as predicted by the Predictor.

The initial population uses orthogonal initialization of weights in each layer with a mean of 0 and a standard deviation of 1 [58]. The population size is 250 and the top 6% of the population is carried over as elites. Parents are selected by tournament selection of the top 20% of candidates using the NSGA-II algorithm [9]. Recombination is performed by uniform crossover at the weight-level, and there is a 20% probability of multiplying each weight by a mutation factor, where mutation factors are drawn from $\mathcal{N}(1, 0.1)$.

Prescriptor candidates are evaluated according to two objectives: (1) the expected number of cases over the next 180 days according to the prescribed NPIs, and (2) the total stringency of the prescribed NPIs, serving as a proxy for the economic cost of the NPIs. Both objectives have to be minimized. The evaluation is done on the 20 countries with the most deaths in the historical data: United States, United Kingdom, Italy, France, Spain, Brazil, Belgium, Germany, Iran, Canada, Netherlands, Mexico, China, Turkey, Sweden, India, Ecuador, Russia, Peru, Switzerland.

On the evaluation start date, each Prescriptor is fed with the last 21 days of case information. Its outputs are used as the NPIs at the evaluation start date, and combined with the NPIs for the




Figure 6: *Fitness of the Final Population along the Case and Cost Objectives.* The candidates at the lower left side are on the Pareto front, representing the best tradeoffs. Those in red are used in the examples below and in the interactive demo (numbered 0 to 19 from left to right). They are the 20 candidates with the highest crowding distance in NSGA-II. The other candidates in the Pareto front are in dark blue and other final population candidates in light blue. An animation of how this population evolved can be seen at <http://evolution.ml/esp/mpi>.

previous 20 days. These 21 days of case information and NPIs are given to the Predictor as input, and it outputs the predicted case information for the next day. This output is used as the most recent input for the next day, and the process continues for the next 180 days. At the end of the process, the average number of predicted new cases over the 180-day period is used as the value of the first objective. Similarly, the average of daily stringencies of the prescribed NPIs over the 180-day period is used as the value for the second objective.

After each candidate is evaluated in this manner, the next generation of candidates is generated. Evolution is run for 110 generations, or approximately 72 hours, on a single CPU host. During the course of evolution, candidates are discovered that are increasingly more fit along the two objectives. In the end, the collection of candidates that represent best possible tradeoffs between objectives (the Pareto front, i.e. the set of candidates that are better than all other candidates in at least one objective) is the final result of the experiment (Figure 6). From this collection, it is up to the human decision maker to pick the tradeoff that achieves a desirable balance between cases and cost. Or put in another way, given a desired balance, the ESP system will find the best to

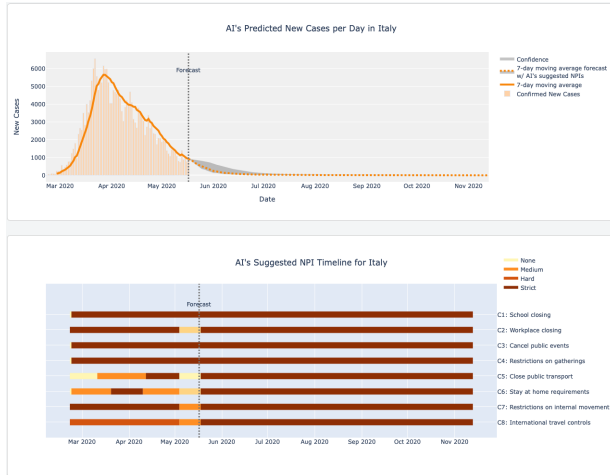
achieve it (i.e. with the lowest cost and the lowest number of cases).

To illustrate these different tradeoffs, Figure 7 shows the NPI Prescriptions and the resulting forecasts for four different Prescriptors from the Pareto front for one country (Italy). The Prescriptor that minimizes cases prescribes the most stringent NPIs across the board, and as a result, the number of cases is minimized effectively. The Prescriptor that minimizes NPI stringency lifts all NPIs right away, and the number of cases explodes as a result. The third Prescriptor was chosen from the middle of the Pareto front, and it represents one particular way to balance the two objectives. It lifts most of the NPIs, allows some public events, and keeps the schools and workplaces closed. As a result, the number of cases is still minimized, albeit slightly slower than in the most stringent case. Lifting more of the NPIs, in particular workplace restrictions, would cause the number of cases to start climbing. In this manner, the decision maker may explore the Pareto front, finding a point that achieves the most desirable balance of cases and cost for the current stage of the pandemic.

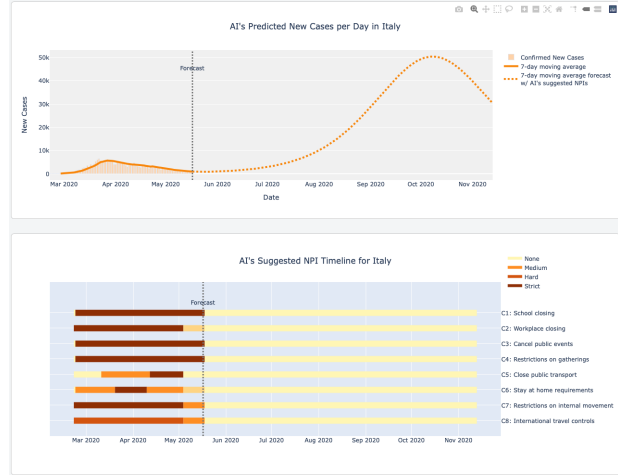
The shadowed area in Figures 7(a) and (c) represents the uncertainty of the prediction, i.e. \pm one standard deviation of the 100 Monte Carlo rollouts under uncertainty estimated through RIO. It is often asymmetric because there is more variance in how the pandemic can spread than how it can be contained. While uncertainty is narrow with stringent Prescriptors (Figure 7(a)), with less stringent ones it often increases significantly with time. The increase can be especially dramatic with Prescriptors with minimal NPIs, such as those in Figures 7(b) and (d), where the uncertainty actually exceeds the display area. The reason is that there has not yet been much training data in this area, i.e. the stage where countries are lifting most NPIs after the peak of the pandemic has passed. The model’s suggestions at this point should be taken as indicative only; in the future these confidence estimates are likely to become more narrow and more reliable conclusions can be drawn then. However, this result can already be interpreted to suggest that such minimal-NPI prescriptions are fragile, and make the country vulnerable to subsequent waves of the pandemic (see also Figures 8(a) and (b)).

To illustrate this process, Figure 8 shows possible choices for three different countries at different stages of the pandemic. For Brazil, where the pandemic is still spreading rapidly, a relatively stringent Prescriptor 4 allows some freedom of movement without increasing the cases much compared to full lockdown. For US, where the number of cases has been relatively flat, a less stringent Prescriptor 7 may be chosen, limiting restrictions to schools, workplaces, and public events. However, if NPIs are lifted too much, opening up the workplaces completely, high numbers of cases are likely to return. For Iran, where there is a danger of a second wave, Prescriptor 6 provides more stringent NPIs to prevent cases from increasing, still limiting the restrictions to schools, workplaces and public events.

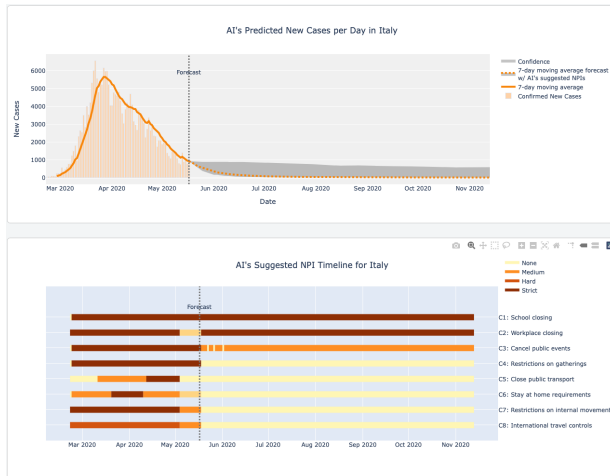
Interestingly, across several countries at different stages of the pandemic, a consistent pattern emerges: in order to keep the number of cases flat, other NPIs can be lifted gradually, but workplace and school restrictions need to be in effect much longer. Indeed these are the two activities where people spend a lot of time with other people indoors, where it is possible to be exposed to significant amounts of the virus [32, 53, 41]. In other activities, such as gatherings and travel, they may come to contact with many people briefly and often outdoors, mitigating the risk. Therefore, the main conclusion that can already be drawn from these preliminary prescription experiments is that it is



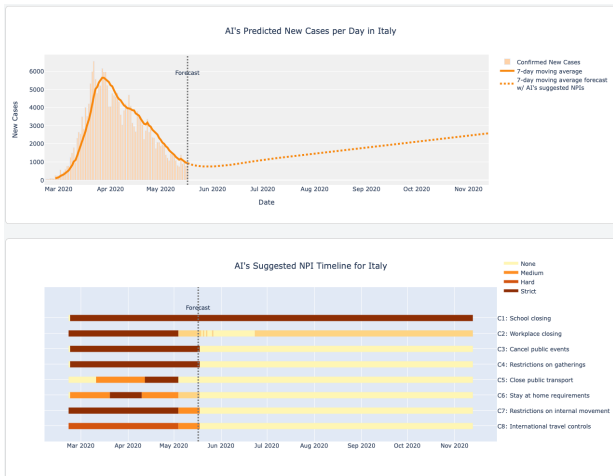
(a) Prescriptor 0:
Minimize cases



(b) Prescriptor 19:
Minimize Stringency



(c) Prescriptor 6:
An example of stringent tradeoff



(d) Prescriptor 8:
An example of non-stringent tradeoff

Figure 7: *Comparison of Different Prescriptors from the Pareto Front.* The recommendations of four different Prescriptors are shown for Italy. Daily cases are shown as orange vertical bars and their seven-day moving average as the orange line. The vertical line indicates the start of the forecast, and the gray area represents uncertainty around the prediction (not shown in (b) and (d) where they are wider than the display). The NPI prescriptions are shown below the case plot as horizontal bars, with color representing stringency. (a) The Prescriptor that minimizes the number of cases recommends a full set of NPIs at their maximum level of stringency. (b) The Prescriptor that minimizes the NPI stringency recommends lifting all NPIs. (c) A Prescriptor that tries to minimize the number of cases while lifting as many NPIs as possible recommends keeping restrictions mostly on schools and workplaces. (d) A Prescriptor that tries to reduce the cost more by opening up workplaces completely may result in cases climbing up. In this manner, the human decision maker can explore the tradeoffs between cases and cost, and the ESP system will recommend the best ways to achieve it.

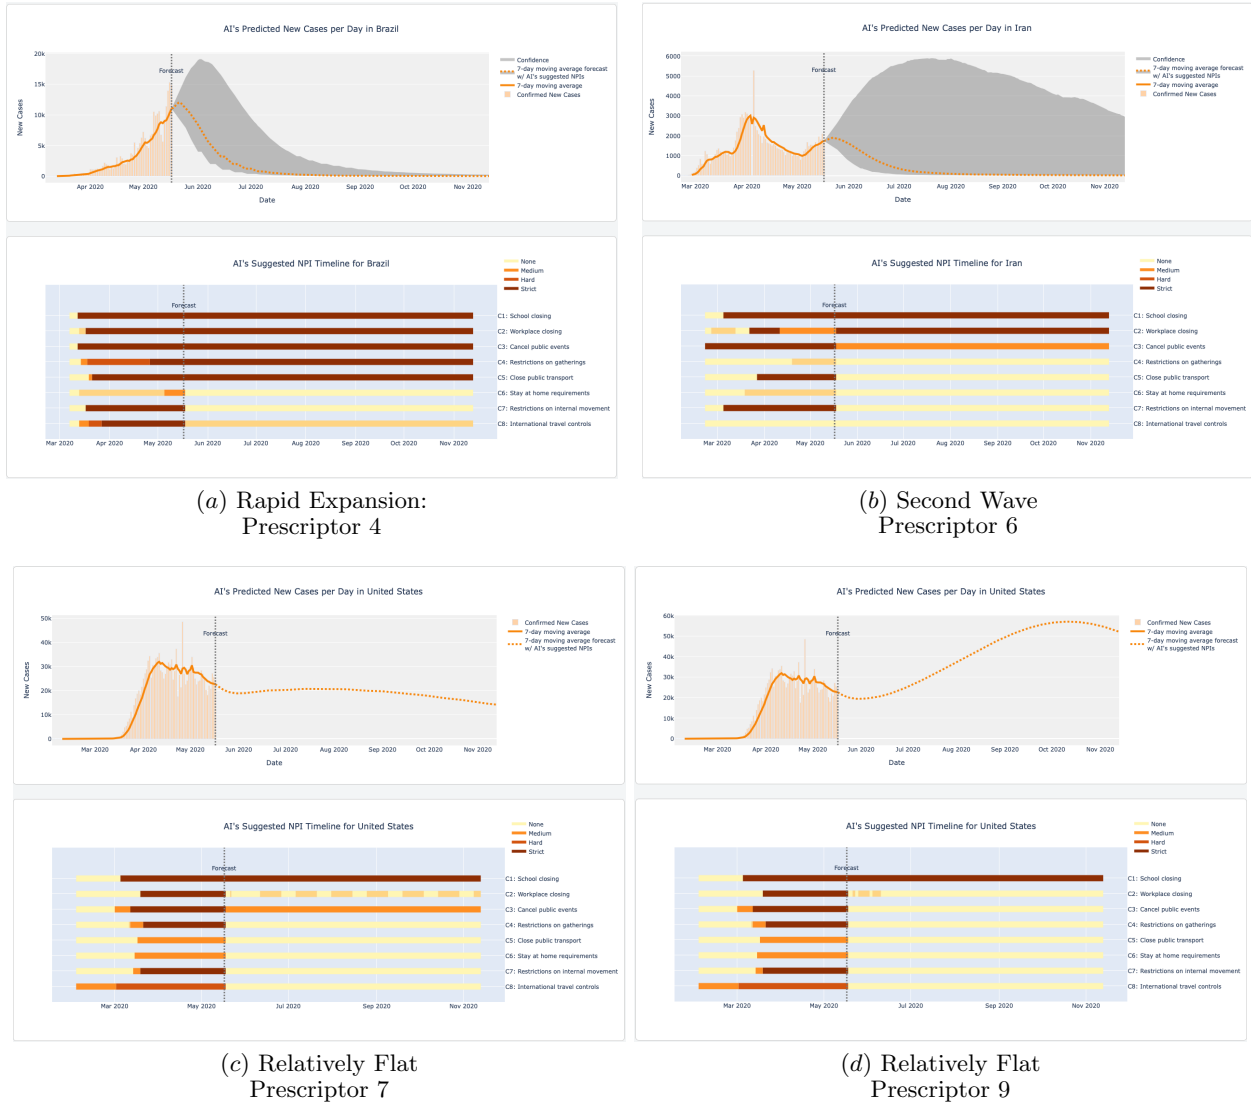


Figure 8: *Comparison of Tradeoff Prescriptions for Countries at Different Stages of the Pandemic.* The Prescriptors chosen represent a midrange in the balance between cases and cost, similar to that of Figure 7(c). (a) For Brazil, where the pandemic is spreading rapidly, Prescriptor 4 minimizes cases effectively while allowing some freedom of movement. (b) For Iran, where the pandemic appears to be entering a second wave, a more stringent Prescriptor 6 strikes a similar balance, but as the area of uncertainty suggests, this balance is risky. (c) For the US, where cases are relatively flat, a less stringent Prescriptor 7 allows reducing cases gradually with minimal cost. (d) In contrast, an even slightly less stringent Prescriptor such as 9 would allow a high number of cases to return. (Uncertainties exceed the display area for (c) and (d).) Interestingly, in all these cases as well as in Figure 7(c), schools and workplaces are subject to restrictions while others are lifted. Also, Prescriptor 7 often includes an alternation of stringency levels, suggesting a way to reduce the cost of the NPI while potentially keeping it effective. Thus, evolution discovers where NPIs may have the largest impact, and can suggest creative ways of implementing them.

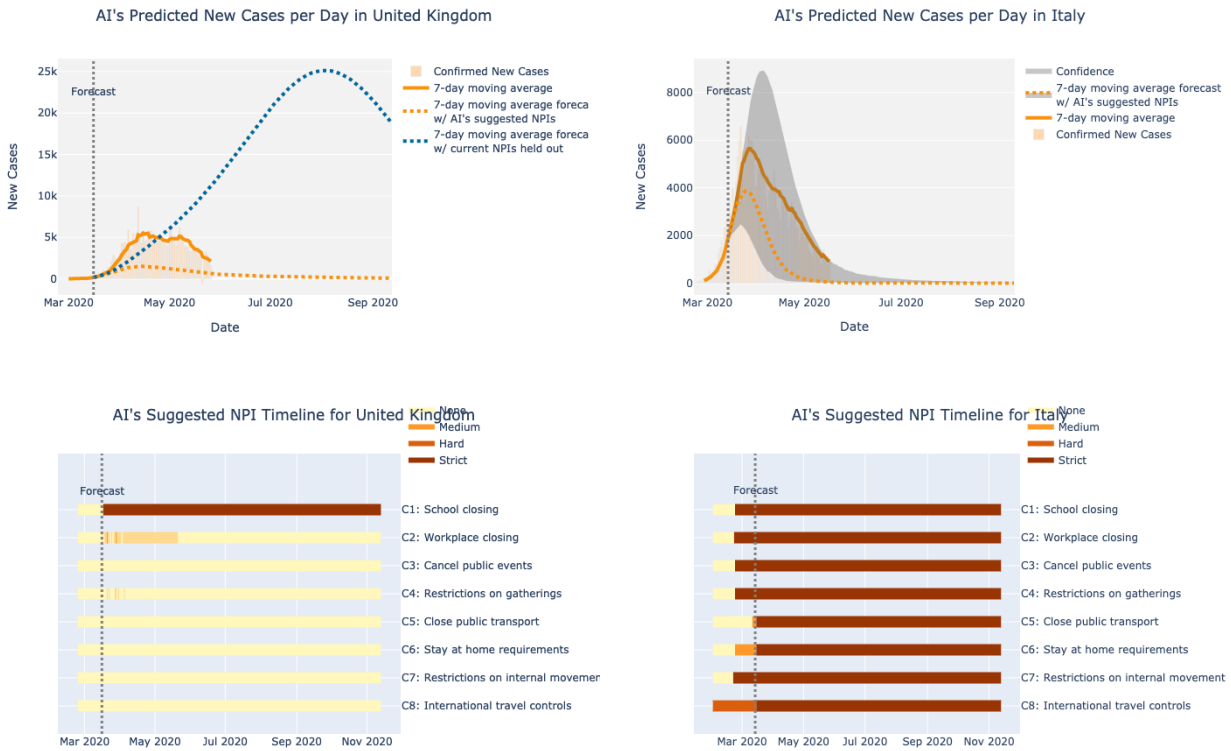
not the casual contacts but the extended contacts that matter. Consequently, when planning for lifting NPIs, attention should be paid in particular to how workplaces and schools can be opened safely.

Another interesting conclusion can be drawn from Figure 8(c): Alternating between weeks of opening workplaces and partially closing them may be an effective way to lessen the impact on the economy while reducing cases. This solution is interesting because it shows how evolution can be creative and find surprising and unusual solutions that are nevertheless effective. There is of course much literature documenting similar surprising discoveries in computational evolution [29, 45, 39], but it is encouraging to see that they are possible also in the NPI optimization domain. While on/off alternation of school and workplace closings may sound unwieldy, it is a real possibility [6]. Note also that it is the only creative solution available to the Prescripator: there are no options in its output for e.g. alternating remote and in-person work, extending school to wider hours, improving ventilation, moving classes outside, or other ways of possibly reducing exposure. How to best implement such distancing at school and workplace is left for human decision makers at this point; the model, however, makes a suggestion that coming up with such solutions may make it possible to lift the NPIs gradually, and thereby avoid secondary waves of cases.

Thus, in the early stages, the ESP approach may suggest how to “flatten the curve”, i.e. what NPIs should be implemented in order to slow down the spread of the disease. At later stages, it may recommend how the NPIs can be lifted and the economy restarted safely. A third role for the approach is to go back in time and evaluate counterfactuals, i.e. how well NPI strategies other than those actually implemented could have worked. In this manner, it may be possible to draw conclusions not only about the accuracy and limitations of the modeling approach, but also lessons for future waves of the current pandemic, for new regions where it is still spreading (like Africa), as well as for future pandemics.

For instance in the UK on March 16th, the NPIs actually in effect were the mild ‘recommend work from home’ and ‘recommend cancel public events’. With only these NPIs, the predicted number of cases could have been quite high (Figure 9(a)). A lockdown was implemented on March 24th, and the actual case numbers were significantly smaller. However, it is remarkable that Prescripator 8 would have required closing schools already on March 16th, and the predicted cases could have been much fewer even without a more extensive lockdown. Thus, the model suggests that an early intervention is crucial, and indeed other models have been used to draw similar conclusions [55]. What is interesting is that ESP suggests that it may be possible to control the pandemic with less than full lockdowns if acted early enough. Of course, the fully trained model was not available at that point, however these lessons may still be useful for countries, such as those in Africa, that are still in early stages, as well as for future pandemics.

Some of the limitations of the data-driven approach also become evident in retrospective studies. For instance Italy, where the pandemic took hold before most of the rest of the world, was supposed to be in a lockdown on March 16th (which started already on February 23rd). Yet, the model predicts that under such a lockdown (suggested e.g. by Prescripator 0 for that date), the number of cases should have been considerably smaller than they actually were (Figure 9(a)). Part of the reason may be that the population did not adhere stringently to the NPIs at that point; after all, the full scale of the pandemic was not yet evident. The population in Italy may also have been older



(a) Earlier NPIs in the UK
Prescriptor 0

(b) Predicted vs. actual in Italy
Prescriptor 0

Figure 9: *Evaluating the Model with Counterfactuals*. Going back in time to make prescriptions makes it possible to evaluate how accurate the model is and to draw lessons for the remainder of this pandemic and future pandemics. (a) After an initial phase of mild NPIs, a lockdown in the UK averted a sharp rise in cases. However, Prescriptor 8 would have recommended earlier NPIs that could have resulted in an even better outcome without a full lockdown. (b) In Italy, a full lockdown in effect on March 16th should have resulted in much fewer cases than it did, suggesting that cultural and demographic factors were different than in other countries, and that the implementation of NPIs need to take such factors into account in order to be effective.

and more susceptible. The data used to train the model comes from 20 different countries and at a later stage of the pandemic spread, and these populations may have followed social distancing more carefully—therefore, the model’s predictions on the effectiveness of lockdowns are too optimistic for Italy. To improve, local factors like culture, economy, population demographics, density, and mobility, may need to be taken into account in the models. Also, this result suggests that the implementation of NPIs need to be sensitive to such factors in order to be effective.

Retrospective studies also show that the epidemic needs to be well underway for the predictions to be reliable. For instance, Italy only had 0 to three cases per day until February 22, when they jumped to 17 in one single day. Trying to predict the pandemic e.g. on March 1st does not lead accurate results. The spread of the virus at that stage depends on individual events, like a church gathering, an executive business meeting in a hotel, or even a soccer match between Liverpool and Atletico Madrid [21], that are hard to predict.

Overall, however, the data-driven ESP approach works surprisingly well even with the current limited data, and can be a useful tool in understanding and dealing with the pandemic. An interactive demo, available on the web, that makes it possible to explore prescriptions and outcomes of the ESP model like those reviewed in this section, will be described next.

7 Interactive Demo

To help understand the mechanisms and possibilities of ESP models, an interactive demo of the current state of the approach to NPI optimization is available at <http://evolution.ml/esp/npis> (Figure 10). This demo will change as the models improve and new functionality is added.¹ At the time of this writing, the following interactions are possible:

The user can select a country by clicking on the map, and a Prescriptor from the Pareto front by clicking on the slider between Cases and NPIs. At the very left, the Prescriptors prefer to minimize cases and therefore usually recommend establishing nearly all possible NPIs. At the very right, the Prescriptors prefer to minimize NPIs and therefore usually recommend lifting nearly all of them—usually resulting in an explosion of cases. The most interesting Prescriptors are therefore somewhere in the middle-left of this range. Some of them are able to keep the cases flat while lifting most of the NPIs, as was discussed in Section 6.

The cases are plotted over time in the middle of the page. The bars indicate past history, used to initialize the Predictor, and future predictions as a line plot with confidence bounds towards the right. The prescribed NPIs are shown over time in the chart at the bottom of the page, with darker colors indicating more stringent version of each NPI. For the Prescriptors that balance the cases and number of NPIs, it is often possible to see an alternating pattern of stringency over time.

With the demo it is possible to explore the options for different countries at different stages of the pandemic. However, it is important to keep in mind that the demo is only a demonstration of the potential of the ESP approach: With current limited data it is not yet possible to make reliable recommendations in a particular case and especially with less stringent NPIs. In the aggregate, it

¹The examples presented in this paper can be replicated by with the appropriate snapshot of the demo: http://evolution.ml/demos/npidashboard/?forecast_folder=20200526_000000 for Figures 7, 8 and 10, and http://evolution.ml/demos/npidashboard/?forecast_folder=20200317_000001 for Figure 9.

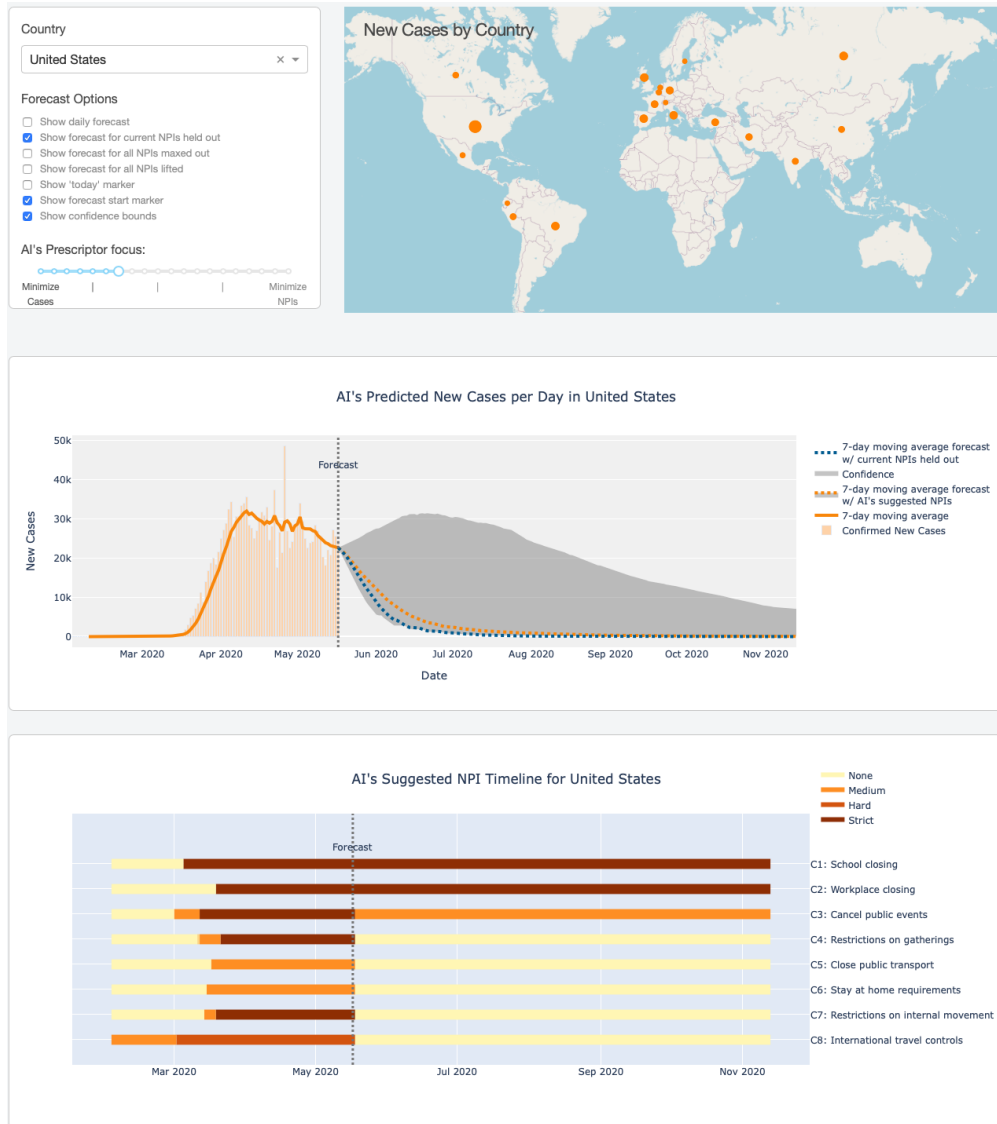


Figure 10: *An Interactive Demo of ESP in the NPI Optimization Problem.* Selecting the country from the map and a Prescripator with the desired case/cost trade-off from the Pareto front slider, the demo shows the predicted cases as a plot over time, and the prescribed NPIs with color-coded stringency over time. The demo can currently be used to understand the potential for the approach, and possibly in the future to augment human decision making in the pandemic. The demo is available at <http://evolution.ml/esp/mpi>.

is possible to draw general conclusions, as was done above. With more and better data and further development, the demo may eventually develop into a tool that can be used to augment human decision making in the pandemic.

8 Future Work

Given the encouraging results in this paper, the most compelling direction of future work consists of updating the model with new data as it becomes available. As NPIs are gradually lifted in many countries, the volume of data will increase, but data will also be more relevant for making decisions in the future. COVID-19 testing will hopefully improve as well so that the outcome measures will be more reliable. The models can be extended to predicting and minimizing deaths as well as cases. Such a multi-task learning environment should make predictions in each task more accurate [40]. Instead of the current six NPIs with a few stringency levels, data on more fine-grained and detailed NPIs may become available, as well as data on more fine-grained locations, such as US counties. In other words, data will improve in volume, relevance, accuracy, and detail, all of which will help make the predictors more precise, and thereby improve prescriptions.

Technically the most compelling direction is to take advantage of multiple prediction models, in particular more traditional compartmental or network models reviewed in Section 2. General assumptions about the spread of the disease are built in to these models, and they can thus serve as a stable reference when data is otherwise lacking in a particular case. On the other hand, it is sometimes hard to estimate the parameters that these models require, and data-driven models can be more mode accurate in specific cases. A particularly good approach might be to form an ensemble from these models (as is often done in machine learning; [67, 46]), and thereby combine their predictions systematically to maximize accuracy.

Another way to make the system more accurate and useful is to improve the outcome measures. Currently the cost of the NPIs is proxied based on how many of them are implemented and at what stringency level. Economic impact is difficult to measure, and the current approach, however approximate, already works relatively well. However, it may be possible to develop more accurate measures based on a variety of economic indicators, such as unemployment, consumer spending, and GNP. They need to be developed for each country separately, given different social and economic structures. With such measures, ESP would be free to find surprising solutions that, while stringent, may not have as high an economic impact.

The retrospective example of Italy in Figure 9(b) suggests that it may be difficult to transfer conclusions from one country to another, and to make accurate recommendations early on in the epidemic. An important future analysis will be to analyze systematically how much data and what kind of data is necessary. For instance, if the model had been developed based on the data from China, would it have worked for Iran and Italy? Or after China, Iran, and Italy, would it have worked for the US and the rest of Europe? That is, how much transfer is there between countries and how many scenarios does the model need to see before it becomes reliable? The lessons can be useful for the rest of the COVID-19 pandemic, as well as for future pandemics.

An important aspect of any decision system is to make it trustworthy, i.e. estimate confidence in its decisions and predictions, allow users to utilize their expert knowledge and explore alternatives,

and explain the decision recommendations. The first step was already taken in this study by applying the RIO uncertainty estimation method (Section 5.3) to the predictions. This approach may be improved in the future by grouping the countries according to original predictor performance, then training a dedicated RIO model for each group. In this way, each RIO model focuses on learning the predictive uncertainty of countries with similar patterns, so that the estimated confidence intervals become more reliable. As a further step, the estimated uncertainty can be used by the Prescripator to make safer decisions.

Second, a prescription “scratchpad” can be included, allowing the user to not only see the prescription details (as shown in Section 7), but also modify them by hand. In this manner, before any prescriptions are deployed, the user can utilize expert knowledge that may not be available for ESP. For instance, some NPIs in some countries may not be feasible or enforceable at a given time. The interface makes it possible to explore alternatives, and see the resulting outcome predictions immediately. In this manner, the user may find more refined prescriptions than those proposed by ESP, or convince him/herself that they are unlikely to exist.

Third, currently the prescriptions are generated by an evolved neural network, which may perform well in the task, but does not provide an explanation of how and why it arrived at a given prescription. In the future, it may be possible to evolve explicit rule sets for this task as well, or instead [59, 24]. Rule sets are readable, specifying which feature values in the context lead to which prescriptions. They can therefore be used to generate explanations of the learned decision strategies, and thereby make it easier for human decision makers to understand and trust them.

While the current ESP model of NPI optimization is a promising demonstration, if enough better data becomes available in the next few months, it may be possible to use the tool during the current COVID-19 pandemic to make useful recommendations. The general approach can also be developed further in the long term, eventually allowing decision makers to minimize the impact of future pandemics.

9 Conclusion

Recent advances in AI have made it possible to not only predict what would happen, but also prescribe what should be done. Also, widely available data has recently made it possible to build data-driven models that are surprisingly accurate in their predictions. This paper puts these two advances together to derive recommendations for NPIs in the current COVID-19 crisis. While preliminary, the model already leads to insights in which NPIs are most important to get right, as well as how they might be implemented most effectively. With further data and development, the approach may become a useful tool for policy makers, helping them to minimize impact of the current as well as future pandemics.

References

- [1] Bisset, K. R., Chen, J., Feng, X., Kumar, V. A., and Marathe, M. V. EpiFast: A fast algorithm for large scale realistic epidemic simulations on distributed memory systems. In *Proceedings of the 23rd international conference on Supercomputing*, pages 430–439, 2009.

- [2] Brockmann, D., Hufnagel, L., and Geisel, T. The scaling laws of human travel. *Nature*, 439(7075):462–465, 2006.
- [3] Canzani, E. and Lechner, U. Insights from modeling epidemics of infectious diseases - a literature review. In *Proceedings of the 12th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, 2015.
- [4] Centers for Diseases Control and Prevention. CDC - working to bring together forecasts for COVID-19 deaths in the us, 2020. <https://www.cdc.gov/coronavirus/2019-ncov/covid-data/forecasting-us.html>, accessed on 2020-05-22.
- [5] Chollet, F. et al. Keras, 2015. <https://github.com/keras-team/keras>, accessed 5/26/2020.
- [6] Chowdhury, R., Heng, K., Shawon, M. S. R., Goh, G., Okonofua, D., Ochoa-Rosales, C., Gonzalez-Jaramillo, V., Bhuiya, A., Reidpath, D., Prathapan, S., Shahzad, S., Althaus, C. L., Gonzalez-Jaramillo, N., Franco, O. H., and The Global Dynamic Interventions Strategies for COVID-19 Collaborative Group. Dynamic interventions to control COVID-19 pandemic: A multivariate prediction modelling study comparing 16 worldwide countries. *European Journal of Epidemiology*, 2020. 10.1007/s10654-020-00649-w.
- [7] Coello Coello, C. A. A comprehensive survey of evolutionary-based multiobjective optimization techniques. *International Journal of Knowledge and Information Systems*, 1:269–308, 1999.
- [8] Darby, C., Louie, G., and Matheny, J. Mitigating economic impacts of the COVID-19 pandemic and preserving u.s. strategic competitiveness in artificial intelligence. White Paper Series on Pandemic Response and Preparedness 2, National Security Commission on Artificial Intelligence, 2020.
- [9] Deb, K., Pratab, A., Agrawal, S., and Meyarivan, T. A fast and elitist multi-objective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6:181–197, 2002.
- [10] Dubov, A. and Shoptaw, S. The value and ethics of using technology to contain the COVID-19 epidemic. *The American Journal of Bioethics*, pages 1–5, 2020.
- [11] Elsken, T., Metzen, J. H., and Hutter, F. Neural architecture search: A survey. *Journal of Machine Learning Research*, 20:1–21, 2019.
- [12] Emmerich, M. and Deutz, A. A tutorial on multiobjective optimization: Fundamentals and evolutionary methods. *Natural Computation*, 17:585–609, 2018.
- [13] European Centre for Disease Prevention and Control. ECDC publications and data, 2020. <https://www.ecdc.europa.eu/en/publications-data>, accessed on 2020-05-22.
- [14] Ferguson, N. M. and et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Technical report, Imperial College COVID-19 Response Team, 2020.

- [15] Flaxman, S., Mishra, S., Gandy, A., Unwin, H. J. T., Coupland, H., Mellan, T. A., Zhu, H., Berah, T., Eaton, J. W., Guzman, P. N. P., Schmit, N., Callizo, L., Team, I. C. C.- R., Whitaker, C., Winskill, P., Xi, X., Ghani, A., Donnelly, C. A., Riley, S., Okell, L. C., Vollmer, M. A. C., Ferguson, N. M., and Bhatt, S. Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in European countries: Technical description update. *arXiv:2004.11342*, 2020.
- [16] Francon, O., Gonzalez, S., Hodjat, B., Meyerson, E., Miikkulainen, R., Qiu, X., and Shahrzad, H. Effective reinforcement learning through evolutionary surrogate-assisted prescription. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2020)*, 2020.
- [17] Gal, Y. and Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1050–1059, 2016.
- [18] Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. M. A. Conditional neural processes. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, pages 1704–1713, 2018.
- [19] Greff, K., Srivastava, R. K., Koutnk, J., Steunebrink, B. R., and Schmidhuber, J. LSTM: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28:222–2232, 2017.
- [20] Gross, T., DLima, C. J. D., and Blasius, B. Epidemic dynamics on an adaptive network. *Physical Review Letters*, 96(20):208701, 2006.
- [21] Guardian. Officials to investigate potential COVID-19 link with liverpool match. *The Guardian*, April 24 2020.
- [22] Hale, T., Webster, S., Petherick, A., Phillips, T., and Kira, B. Oxford COVID-19 government response tracker. Blavatnik School of Government, 2020. <https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker>, accessed 5/26/2020.
- [23] Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9:1735–1780, 1997.
- [24] Hodjat, B., Shahrzad, H., and Miikkulainen, R. PRETSL: Distributed probabilistic rule evolution for time-series classification. In R. Riolo, B. G., B. Worzel and Tozier, B., editors, *Genetic Programming Theory and Practice XIV*. Springer, New York, 2018.
- [25] Hufnagel, L., Brockmann, D., and Geisel, T. Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences*, 101(42):15124–15129, 2004.

- [26] IHME. COVID-19 projections, 2020. <https://covid19.healthdata.org/united-states-of-america/california>, accessed 5/26/2020.
- [27] IHME COVID-19 health service utilization forecasting team and Murray, C. J. L. Forecasting the impact of the first wave of the COVID-19 pandemic on hospital demand and deaths for the usa and european economic area countries. *medRxiv*, 10.1101/2020.04.21.20074732, 2020.
- [28] Johns Hopkins University. New cases of COVID-19 in world countries, 2020. <https://coronavirus.jhu.edu/data/new-cases>, accessed 5/26/2020.
- [29] Johnson, A. J., Meyerson, E., de la Parra, J., Savas, T. L., Miikkulainen, R., and Harper, C. B. Flavor-cyber-agriculture: Optimization of plant metabolites in an open-source control environment through surrogate modeling. *PLOS ONE*, 2019. <https://doi.org/10.1371/journal.pone.0213918>.
- [30] Kaggle. COVID19 global forecasting, 2020. <https://www.kaggle.com/c/covid19-global-forecasting-week-5/data>, accessed 5/26/2020.
- [31] Kaiser. State data and policy actions to address coronavirus, 2020. <https://www.kff.org/health-costs/issue-brief/state-data-and-policy-actions-to-address-coronavirus/>, accessed 5/26/2020.
- [32] Kay, J. COVID-19 superspreader events in 28 countries: Critical patterns and lessons. *Quillette*, April 23 2020.
- [33] Keeling, M. J. and Eames, K. T. Networks and epidemic models. *Journal of the Royal Society Interface*, 2(4):295–307, 2005.
- [34] Kermack, W. O. and McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London Series A*, 115(772):700–721, 1927.
- [35] Keystone. Coronavirus city and county non-pharmaceutical intervention rollout date dataset, 2020. <https://www.keystonestrategy.com/coronavirus-covid19-intervention-dataset-model>, accessed 5/26/2020.
- [36] Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, S. M. A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. *arXiv:1901.05761*, 2019.
- [37] Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- [38] Lakshminarayanan, B., Pritzel, A., and Blundell, C. Practical bayesian optimization of machine learning algorithms. In I. Guyon, S. B. H. W. R. F. S. V., U. V. Luxburg and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, page 64026413. Curran Associates, Inc., 2017.

- [39] Lehman, J., Clune, J., and Misevic, D. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial Life*, pages 1–33, 2020.
- [40] Liang, J., Meyerson, E., and Miikkulainen, R. Evolutionary architecture search for deep multitask networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2018.
- [41] Lu, J., Gu, J., Li, K., Xu, C., Su, W., Lai, Z., and et al. COVID-19 outbreak associated with air conditioning in restaurant, Guangzhou, China, 2020. *Emerging Infectious Diseases*, 2020.
- [42] Luke, D. A. and Stamatakis, K. A. Systems science methods in public health: Dynamics, networks, and agents. *Annual Review of Public Health*, 33:357–376, 2012.
- [43] May, R. M. and Anderson, R. M. Transmission dynamics of HIV infection. *Nature*, 326(6109):137–142, 1987.
- [44] Miikkulainen, R. Creative AI through evolutionary computation. In Banzhaf, W., Cheng, B., Deb, K., Holekamp, K., Lenski, R. E., Ofria, C., Pennock, R., Punch, B., and Whittaker, D., editors, *Evolution in Action: Past, Present and Future*. Springer, New York, 2020.
- [45] Miikkulainen, R., Brundage, M., Epstein, J., Foster, T., Hodjat, B., Iscoe, N., Jiang, J., Legrand, D., Nazari, S., Qiu, X., Scharff, M., Schoolland, C., Severn, R., and Shagrin, A. Ascend by evolv: AI-based massively multivariate conversion rate optimization. *AI Magazine*, 42:44–60, 2020.
- [46] Miikkulainen, R., Liang, J., Meyerson, E., Rawal, A., Fink, D., Francon, O., Raju, B., Shahrzad, H., Navruzyan, A., Duffy, N., and Hodjat, B. Evolving deep neural networks. In Morabito, C. F., Alippi, C., Choe, Y., and Kozma, R., editors, *Artificial Intelligence in the Age of Neural Networks and Brain Computing*. Elsevier, New York, in press.
- [47] Models of Infectious Disease Agent Study. Midas - online portal for COVID-19 modeling research, 2020. <https://midasnetwork.us/covid-19/>, accessed on 2020-05-22.
- [48] Mossong, J., Hens, N., Jit, M., Beutels, P., Auranen, K., Mikolajczyk, R., Massari, M., Salmaso, S., Tomba, G. S., Wallinga, J., et al. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS medicine*, 5(3), 2008.
- [49] Neal, R. M. *Bayesian Learning for Neural Networks*. Springer, Berlin, 1996.
- [50] Newman, M. E. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002.
- [51] Oxford. Codebook for the oxford covid-19 government response tracker, 2020. <https://github.com/OxCGRT/covid-policy-tracker/blob/master/documentation/codebook.md>, accessed 5/26/2020.

- [52] Oxford. Coronavirus government response tracker, 2020. <https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker>, accessed 5/26/2020.
- [53] Park, S. Y., Kim, Y. M., Yi, S., Lee, S., Na, B. J., Kim, C. B., and et al. Coronavirus disease outbreak in call center, South Korea. *Emerging Infectious Diseases*, <https://doi.org/10.3201/eid2608.201274>, 2020.
- [54] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [55] Pei, S., Kandula, S., and Shaman, J. Differential effects of intervention timing on COVID-19 spread in the united states. *medRxiv*, 10.1101/2020.05.15.20103655, 2020.
- [56] Petherick, A., Kira, B., Angrist, N., Hale, T., Phillips, T., and Webster, S. Variation in government responses to COVID-19. Blavatkin School of Government Working Paper Series BSG-WP-2020/032, Oxford University, 2020.
- [57] Qiu, X., Meyerson, E., and Miikkulainen, R. Quantifying point-prediction uncertainty in neural networks via residual estimation with an I/O kernel. In *Proceedings of the International Conference on Learning Representations*, 2020.
- [58] Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural network. In *Proceedings of the Second International Conference on Learning Representations (ICLR)*, 2014.
- [59] Shahrzad, H., Hodjat, B., Dolle, C., Denissov, A., Lau, S., Goodhew, D., Dyer, J., and Miikkulainen, R. Enhanced optimization with composite objectives and novelty pulsation. In *Genetic Programming Theory and Practice XVII*. Springer, New York, in press.
- [60] Stanford University. Potential long-term intervention strategies for COVID-19, 2020. <http://covid-measures.stanford.edu/>, accessed on 2020-05-21.
- [61] Stanley, K. O., Clune, J., Lehman, J., and Miikkulainen, R. Designing neural networks through evolutionary algorithms. *Nature Machine Intelligence*, 1:2435, 2019.
- [62] Swarup, S., Eubank, S. G., and Marathe, M. V. Computational epidemiology as a challenge domain for multiagent systems. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, pages 1173–1176, 2014.
- [63] The Atlantic. The covid tracking project, 2020. <https://covidtracking.com>, accessed 5/26/2020.
- [64] University of Washington Institute for Health Metrics and Evaluation. IHME - COVID-19 projection, 2020. <https://covid19.healthdata.org/united-states-of-america>, accessed on 2020-05-21.

- [65] Venkatramanan, S., Lewis, B., Chen, J., Higdon, D., Vullikanti, A., and Marathe, M. Using data-driven agent-based models for forecasting emerging infectious diseases. *Epidemics*, 22:43–49, 2018.
- [66] Woody, S., Garcia Tec, M., Dahan, M., Gaither, K., Lachmann, M., Fox, S., Meyers, L. A., and Scott, J. G. Projections for first-wave COVID-19 deaths across the us using social-distancing measures derived from mobile phones. *medRxiv*, 10.1101/2020.04.16.20068163, 2020.
- [67] Zaremba, W., Sutskever, I., and Vinyals, O. Recurrent neural network regularization. *arXiv:1409.2329*, 2014.