

# Data-driven Simulation and Optimization for Covid-19 Exit Strategies

Salah Ghamizi  
SnT, University of Luxembourg  
salah.ghamizi@uni.lu

Renaud Rwemalika  
SnT, University of Luxembourg  
renaud.rwemalika@uni.lu

Maxime Cordy  
SnT, University of Luxembourg  
maxime.cordy@uni.lu

Lisa Veiber  
SnT, University of Luxembourg  
lisa.veiber@uni.lu

Tegawendé F. Bissyandé  
SnT, University of Luxembourg  
tegawende.bissyande@uni.lu

Mike Papadakis  
SnT, University of Luxembourg  
michail.papadakis@uni.lu

Jacques Klein  
SnT, University of Luxembourg  
jacques.klein@uni.lu

Yves Le Traon  
SnT, University of Luxembourg  
yves.letraon@uni.lu

## ABSTRACT

The rapid spread of the Coronavirus SARS-2 is a major challenge that led almost all governments worldwide to take drastic measures to respond to the tragedy. Chief among those measures is the massive lockdown of entire countries and cities, which beyond its global economic impact has created some deep social and psychological tensions within populations. While the adopted mitigation measures (including the lockdown) have generally proven useful, policymakers are now facing a critical question: *how and when to lift the mitigation measures?* A carefully-planned exit strategy is indeed necessary to recover from the pandemic without risking a new outbreak. Classically, exit strategies rely on mathematical modeling to predict the effect of public health interventions. Such models are unfortunately known to be sensitive to some key parameters, which are usually set based on rules-of-thumb.

In this paper, we propose to augment epidemiological forecasting with actual data-driven models that will learn to fine-tune predictions for different contexts (e.g., per country). We have therefore built a pandemic simulation and forecasting toolkit that combines a deep learning estimation of the epidemiological parameters of the disease in order to predict the cases and deaths, and a genetic algorithm component searching for optimal trade-offs/policies between constraints and objectives set by decision-makers.

Replaying pandemic evolution in various countries, we experimentally show that our approach yields predictions with much lower error rates than pure epidemiological models in 75% of the cases and achieves a 95%  $R^2$  score when the learning is transferred and tested on unseen countries. When used for forecasting, this approach provides actionable insights into the impact of individual measures and strategies.

## CCS CONCEPTS

• **Computing methodologies** → **Simulation types and techniques**; *Machine learning*.

## KEYWORDS

Covid19, Pandemic, Prediction, Exit strategies, SEIR, Deep Learning, Search-based Optimization

## ACM Reference Format:

Salah Ghamizi, Renaud Rwemalika, Maxime Cordy, Lisa Veiber, Tegawendé F. Bissyandé, Mike Papadakis, Jacques Klein, and Yves Le Traon. 2020. Data-driven Simulation and Optimization for Covid-19 Exit Strategies. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20), August 23–27, 2020, Virtual Event, CA, USA*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394486.3412863>

## 1 INTRODUCTION

Since the outbreak of the COVID-19 pandemic, the world has been facing a human tragedy with overwhelmed healthcare systems and fears of economic collapses. In the absence of vaccines to immunize the population rapidly at scale, governments have implemented various non-pharmaceutical public health interventions such as social distancing and lockdowns. Considering that the World Health Organisation (WHO) is foreseeing first clinical trials of vaccine for the end of the year 2020 [15], decision-makers must carefully plan their exit strategies: measures that were put in place to contain the coronavirus spread must be methodically lifted to avoid the risk of precipitating new outbreaks.

In this context, mathematical modelling offers public health planners with frameworks to make predictions about the spread of emerging diseases and assess the impact of possible mitigation strategies. This is particularly important when dealing with infectious diseases, such as COVID-19, where mass interventions (e.g., screening, social distancing, and vaccination) can lead to effects at a population level, including herd immunity, changes in the infection rate or even changes in the pathogen ecology as a consequence of selective pressure.

There are two main types of models: static cohort models and transmission dynamic models [3]. Static models, typically relying



This work is licensed under a Creative Commons Attribution International 4.0 License.  
KDD '20, August 23–27, 2020, Virtual Event, CA, USA  
© 2020 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-7998-4/20/08.  
<https://doi.org/10.1145/3394486.3412863>

on decision trees and Markov processes, assume a force of infection that is independent of the proportion of the population that is infected and therefore is of little use in response to highly infectious diseases like COVID-19. Transmission dynamic models, on the other hand, see a force of infection varying depending on the proportion of the population which is infected. Compared with static cohort models, transmission dynamic models are usually more complex to parameterize requiring epidemiological information on the infectious disease and demographic and economic information about the affected population.

Different techniques exist to implement dynamic approaches. Agent-Based Models (ABM) are simulations composed of agents that interact with each other and their environment. Because each agent can make its own rules, this type of approach can capture aggregate phenomena derived from the behavior of single agents. These models offer a great explainability of the root causes leading to the propagation of a disease but are computationally intensive to run and thus, hardly applicable to large populations. Indeed, the behaviour and the interaction of each type of agent needs to be fully defined in order for the model to be useful. These rules are case-specific and are not transferable from one population to another.

The most common approach to model the spread of infectious disease is the Susceptible-Infected-Removed (SIR) model and its extension *i.e* SEIR (Susceptible, Exposed, Infectious, and Recovered). This is a state-based model, every state expresses the degree of exposure of a population to the disease. It is also equation-based where each equation defines the rate to go from one state to the other. The SEIR model thus separates the population into four groups and simulates the evolution over time of each one of the subpopulations. The transition rates are defined by the time scale to which an individual can transmit the disease, the time to recovery, and the number of newly infected people due to an infected individual. The most varying parameter is effective reproduction number ( $R_t$ ), and expresses the number of people that can be contaminated by an infectious individual over a period of time.

These methods are dependent on the validity of the input parameters like transition rates. While SEIR is a very powerful model, it presents a major limitation, it requires hyper-parameters that are hard to observe such as the infection rate of an individual. In practice, SEIR parameters are manually set to fit with the local observations to the considered population (e.g. country), and are not learnt from larger-scale observations. To circumvent the limitation of such epidemiological models, researchers started to take advantage of the advances made in Machine Learning (ML) in order to create models based on available large datasets [11, 14]. We name this family of approaches “ML-based”. and our own work falls in it.

Our first contribution is to devise a novel approach, *DN-SEIR*, that alleviates manual tuning of the SEIR model, by relying on large and trustable public datasets (large scale observations) and machine learning (to learn the parameters’ values for a given population). Our approach combines SEIR with a machine learning predictor, based on a deep learning model, to estimate the effective reproduction number ( $R_t$ ), over time. The machine learning predictor relies on demography and mobility features to predict an effective reproduction number. For each time increment,  $R_t$  is updated and used for the next day computation. We evaluated our approach on

twelve countries from all continents and showed that our approach that mixes demographics, mobility, and epidemiological data provides better forecasts for 9 out of 12 of the studied countries than a purely epidemiological modelling.

Our second contribution is to exploit this online prediction of effective reproduction number in a simulation tool<sup>1</sup> for policymakers which was recently advertised to the public<sup>2</sup>. Policymakers have to decide when to relax certain parts of society (workplaces, travels, schools... ) and to what extent it may create a new epidemic wave that would flood the hospitals with critical cases. The simulator enables one to make such a strategic exit plan for a certain country and predict its impact in terms of hospitalization, infected people and deaths. It is also designed to explore and optimize various exit strategies and constraints. We evaluate 3 common hand-crafted exit strategies and show that multi-objective genetic algorithms can find atypical strategies on the pareto-front that minimize both the death numbers and the economic impact.

## 2 RELATED WORK

Machine Learning approaches have been widely used to model and forecast former epidemics, especially to handle the large amount of data it involves and the increasing complexity of the epidemiological models underneath. Popular approaches remain regression trees and forests [8], and neural networks [4]. While there is a plethora of reports that tackle the COVID-19 forecasting, the peer-reviewed literature about ML and COVID-19 is rather scarce. Most approaches in public repositories tackle ML regression in combination with the SIR epidemiological model [1] or its SEIR extension [9, 12, 16].

Recent research about non-pharmaceutical interventions focused on mobility data to combine epidemiological models and learning algorithms. In [13], Vollmer et al. integrate mobility in a stochastic model. They focused on Italy and suggest that COVID-19 transmission rate and mobility metrics are closely related as well as mobility should be closely monitored in the next weeks and months..

Our approach relies on a similar intuition and combines learning from multiple countries with country-specific features. The transfer learning in our approach is able to separate the contribution of different measures like commuting to work, retail and recreation activities. We also show that using the output of our ML approach as a search fitness function leads to optimal exit strategies.

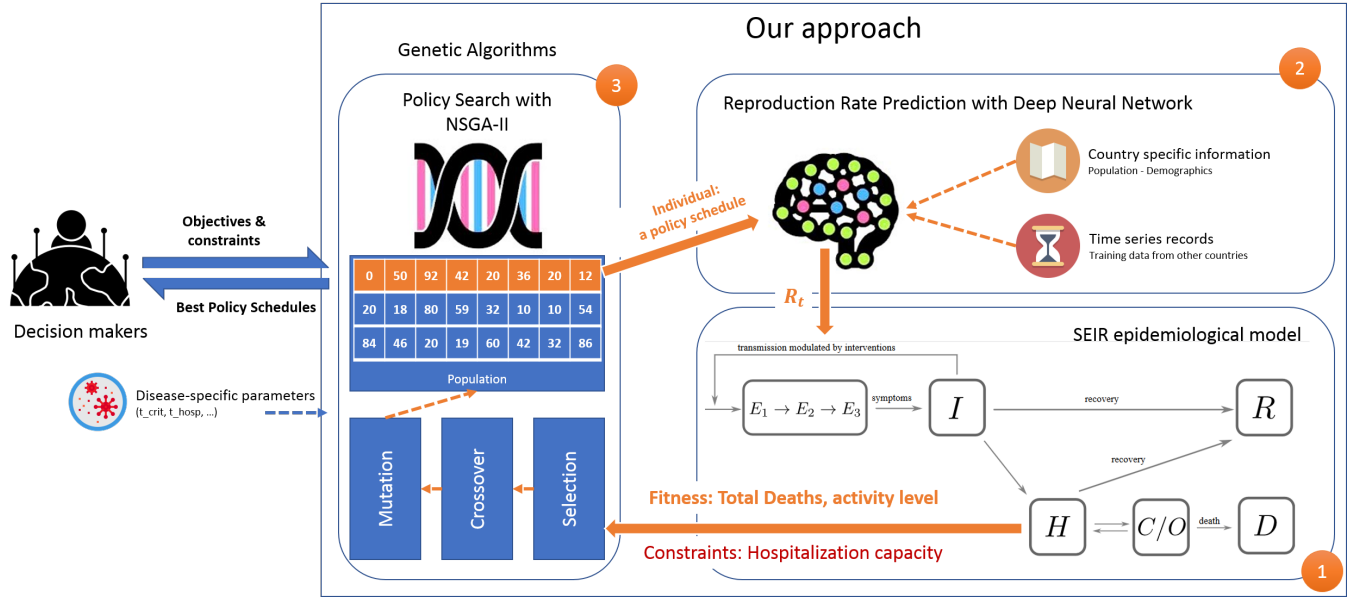
## 3 APPROACH

Our end goal is to provide policymakers with a tool to easily generate exit strategies and evaluate their impact. In particular, an exit strategy can be modeled like as a schedule of measures (*policy schedule*) that will impact the way the disease will spread. We restrict the policy schedule to mobility levels.

As illustrated in Figure 1, we propose to combine a genetic algorithm (to search for policy schedules), a deep learning model (to predict the evolution of the effective reproduction number induced by a given policy schedule) and an epidemiological model (to forecast, based on the computed effective reproduction numbers, the effect of the scheduled policies on public health over time, e.g. deaths and hospitalization occupancy). Our three components work within a feedback loop. At each iteration, the genetic algorithm

<sup>1</sup>Open sourced and available on <https://github.com/yamizi/Covid19>

<sup>2</sup><https://t.co/FN5pn1dMOR>



**Figure 1: Our approach** relies on a feedback loop where the Genetic Algorithm searches for optimal exit strategies using a fitness function computed from the epidemiological model outputs. The epidemiological model’s parameters are learned with a Machine Learning algorithm that uses population mobility behaviours and demographics as input features.

builds a population of policy schedules, which the deep learning and epidemiological models allow to evaluate. In turn, this feedback is used to generate better schedules, optimizing health-related objectives (e.g. minimize total deaths) while satisfying hard constraints (e.g. never exceed hospitalization capacity).

### 3.1 Estimating Impacts on Public Health with Epidemiological Models

Epidemiological models predict the state of a population struck by a pandemic over time, based on state transition parameters and the evolution of the effective reproductive number,  $R_t$ , of the disease.

We use an extension of SEIR model, i.e. the SEI-HCRD compartmental model – Susceptible ( $S$ ) → Exposed ( $E$ ) → Infectious ( $I$ ) → Removed (Hospitalized ( $H$ ), Critical ( $C$ ), Recovered ( $Rec$ ), Dead ( $D$ )). Such model can be defined by the following system of ordinary differential equations:

$$\frac{dS}{dt} = -\frac{R_t}{t_{inf}} \cdot I \cdot S \quad (1)$$

$$\frac{dE}{dt} = \frac{R_t}{t_{inf}} \cdot I \cdot S - \frac{1}{t_{inc}} \cdot E \quad (2)$$

$$\frac{dI}{dt} = \frac{1}{t_{inc}} \cdot E - \frac{1}{t_{inf}} \cdot I \quad (3)$$

$$\frac{dH}{dt} = \frac{1-m}{t_{inf}} \cdot I + \frac{1-f}{t_{crit}} \cdot C - \frac{1}{t_{hosp}} \cdot H \quad (4)$$

$$\quad (5)$$

$$\frac{dC}{dt} = \frac{c}{t_{hosp}} \cdot H - \frac{1}{t_{crit}} \cdot C \quad (6)$$

$$\frac{dRec}{dt} = \frac{m}{t_{inf}} \cdot I + \frac{1-c}{t_{hosp}} \cdot H \quad (7)$$

$$\frac{dD}{dt} = \frac{f}{t_{crit}} \cdot C \quad (8)$$

such that  $S + E + I + H + C + Rec + D$  is equal to the total population and  $R_t$  denotes the effective reproduction number over time.

The SEI-HCRD involves several parameters.  $t_{suffix}$  is the transition time estimated to transit from one population to the other.  $t_{inc}$  is the average incubation period,  $t_{inf}$  is the average infectious period,  $t_{hosp}$  is the average hospitalization time in normal state (i.e. until any patient recovers or enters a critical state).  $t_{crit}$  is the average hospitalization time in a critical state (i.e. until death or recovery). The parameters  $m$ ,  $c$ ,  $f$  determine the severity of the infection:  $m$  is the percentage of infected individuals with non-severe symptoms (i.e. they are asymptomatic or have mild symptoms) and which, therefore, are not hospitalized.  $c$  is the percentage of hospitalized persons who will eventually enter a critical state. Finally,  $f$  denotes the percentage of persons in the critical state who will pass away.

To set these parameters, we lean on Liu et al.’s study [5] and assign them with the constant values reported in Table 1. Then, given the time series of effective reproduction numbers over time,  $\{R_t\}$ , the SEI-HCRD model computes the resulting impacts on the population, including the number of deaths. Instead of manually assigning fitted values to  $R_t$ , we propose to predict them from scheduled exit strategies using deep learning models.

Transition time	Value	Transition ratio	Value
$t_{inc}$	5.6 days	$m$	80%
$t_{inf}$	2.9 days	$c$	10%
$t_{hosp}$	4 days	$f$	30%
$t_{crit}$	14 days		

**Table 1: Parameter values used in the SEI-HCRD model.**

## 3.2 Predicting the Effective Reproduction Number over time with Deep Learning

**3.2.1 Feature engineering.** We start from Google’s Mobility Reports<sup>3</sup>, which track the mobility trends over time, for different categories of places in 97 different countries. Each feature corresponds to a category of places and its value captures the daily traffic of such places. More precisely, the value of the feature is the difference between the daily traffic and the traffic baseline (i.e. the median traffic for the same weekday during the 5-weeks ranging from January 3 to February 6, 2020). The reports include 6 categories: Grocery & pharmacy, park, transit stations, retail & recreation, residential and workplace.

Of course, the values of these features are country-specific and are largely impacted by the mitigation strategy of each country. For example, Figure 2 shows the evolution of all features for Luxembourg, Italy and Japan. We observe that Italy and Luxembourg have drastically reduced their activities whereas Japan does not exhibit a significant reduction, for the case of schools and international travels.

Next, we clean the collected data. When some values (for a given feature) are missing, we fill the gap by interpolating between the closest days with available information. For each category of places, we smooth the corresponding feature over the 5, 10, 15, and 30 past days, resulting in four new features.

We complete our dataset with demographic features and with the corresponding day of the week to take into account the weekly fluctuations of the data.

Overall, our feature engineering process yields 4,625 inputs of 32 features each, which are recapitulated in Table 2.

**3.2.2 Training & validation.** The model can be seen as a supervised predictor, taking as input the mobility and demographic features to predict an effective reproduction number,  $R_t$ , for each time index  $t$ . Thus, in order to train it, we need to label the dataset with  $R_t$  value for each day of the training period.

Since the real effective reproduction numbers are not known, we estimate them by fitting the SEI-HCRD model to the real-world number of cases and deaths. Since the  $R_t$  are time-dependent, we represent them as a decaying function and seek the best parameters’ value for this function (which yield the fittest SEI-HCRD model). We opted for the Hill decay function because it showed good results in the literature [10]. Thus,  $R_t$  is given by  $\frac{1}{(1+(\frac{t}{L})^k)} \cdot R_0$  where  $R_0$ ,  $k$  and  $L$  are parameters. We use the L-BFGS optimization method to find the values of these parameters that minimize the Mean Square Error of the number of cases and deaths predicted by the SEI-HCRD model. Once these parameter values are found, they can be injected

<sup>3</sup><https://www.google.com/covid19/mobility/>

Feature category	Feature	Values
Mobility	Grocery & Pharmacy	[-100,50]
	Parks	[-100,50]
	Residential	[-100,50]
	Retail & Recreation	[-100,50]
	Transit stations	[-100,50]
	Workplace	[-100,50]
Demographics	GDP	$\mathbb{N}$
	Population	$\mathbb{N}$
	Density	$\mathbb{R}$
	Area	$\mathbb{N}$
	Proportion of population under 15yrs	[0,1]
	Proportion of population over 64yrs	[0,1]
	Day of week	[0-6]
	Region (Continent)	[0-10]

**Table 2: Features of the Feed Forward Neural Network. Mobility features are augmented by smoothing over 5, 10, 15 and 30 days, hence each mobility feature corresponds to 4 inputs.**

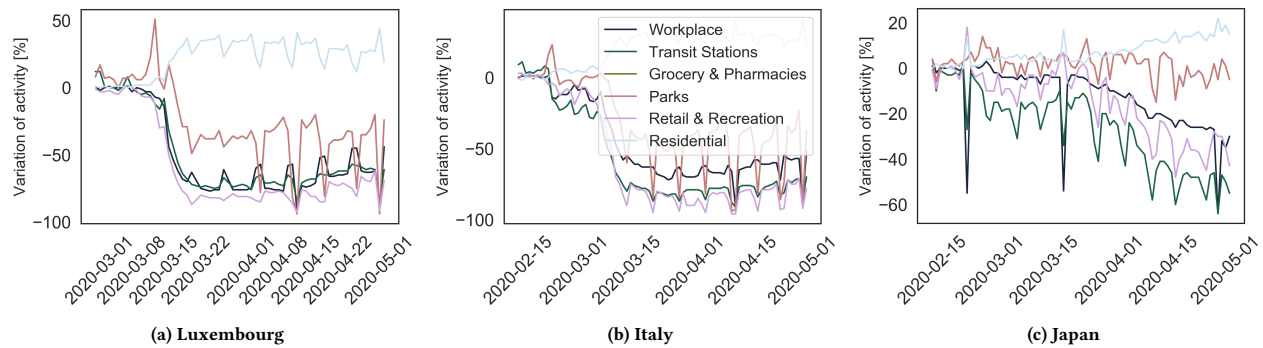
back in the Hill decay function to generate a time series for the past values of  $R_t$ .

The analysis of feature correlations (Figure 3a) shows that when working on a country-by-country basis, all the mobility features are highly correlated (over 0.75) for some countries as all the activity reduction and closure were enacted in most sectors at the same time. Co-linear features offering little to no information gain to the learner, we train the model on all countries to reduce any correlation between the different features. Thus, the model is trained using all the countries at once and the train/test split is done randomly.

Once a training set with expected values ( $R_t$  values computed from the estimated Hill decay function) is established, a supervised model can be trained to predict the future values of  $R_t$  using the mobility and demographic data as features.

To do so, we rely on a Feed-Forward Neural Network (FFNN). The architecture of the FFNN and its hyper-parameters are optimized using a grid search to minimize the mean square error with cross-validation. The search leads to an architecture with 2 fully-connected hidden layers with, respectively, 1000 and 50 neurons.

In addition to the FFNN, we also evaluate two other estimators. Since the problem takes the form of a time series, we investigate the performances of a Long Short-Term Memory (LSTM) network which allows sequence to sequence transformation, hence, learning from the features of the past days to predict the next days. We use a 15 days window to predict the values for the next 7 days. Note that in this case, at each step, we use the computed values of  $R_t$  from the past iterations as input for the model in addition to the rest of the features. Finally, we investigate one last approach, Gradient Boosting using 500 estimators. For each of the approaches, we evaluate the performances using two classical metric, i.e the coefficient of determination, ( $R^2$  score) and the Root Mean Square



**Figure 2: Evolution of the mobility indicators for Luxembourg, Italy, and Japan. A value of 0 means that the activity is at the same level as before the confinement, a value of -100% is a total stop of the activity and a positive values shows an increase of the activity compared to the reference value.**

Error (RMSE)<sup>4</sup>. The train:test splits were performed with (1) random split between train and test data (2) split based on the region the country is located in (i.e testing on unseen countries) (3) split based on the time, all values before a certain date are considered for the training set and the ones after are used to build the test set. In all instances, we keep a ratio train:test around 80:20.

The FFNN provided a  $R^2$  score of 0.95 with a random split and 0.97 with a region split, while the gradient boosting could only achieve a  $R^2$  score of 0.83. The LSTM offered slightly better performances with a  $R^2$  score of 0.95 when splitting on a region base to over 0.99 when splitting on a time base. We use the FFNN model in our following experiments. We refer to the combination of FFNN and SEI-HCRD as **DN-SEIR**.

**3.2.3 Interpretability.** Although Machine Learning algorithms provide increased accuracy in a wide variety of domains, their black-box nature makes them inherently non-interpretable. Indeed, in our case, a multi-layer neural network solely contains information in the form of numerical weights and connections. The model reasoning from input to output remains opaque. Nevertheless, interpretability can be reached as a post-hoc analysis through an independent interpretability framework. We choose Shapley Additive exPlanations, *SHAP* [7] as it provides intuitive visualization-based explanations which can be incorporated in the simulator. The *SHAP* framework is based on game theory and Shapley value. In game theory, Shapley values indicate how to fairly distribute a ‘payout’ among players. A model can be thought of as a game where each feature value for each instance is a player, while a prediction or model output is a payout. In practice, the Shapley value of a feature value is the contribution of this value for this particular exemplar compared to the average prediction for the specific dataset. *SHAP* provides several advantages. First, this framework has different types of explainers to optimally provide explanations to different models, whether tree-based or kernel-based. Moreover, *SHAP* exhibits an Efficiency property through the Shapley values. Indeed this component guarantees that the difference between prediction

and average prediction is fairly distributed among the values of the features of this particular prediction.

### 3.3 Optimization of Policy Schedules with Genetic Algorithms

Our search method uses NSGA-II [2], an established Genetic Algorithm (GA) for multi-objective optimization that uses non dominated sorting to find pareto-optimal solutions.

**Solution space:** Any solution generated by NSGA-II is a policy schedule. A schedule consists of a list of vectors, each of which is associated with a mobility feature and encodes the value of the feature for each time index  $t$ . A value ranges from 0 (no restriction) to 100 (full lockdown). The indices  $t$  go from April 30 to September 30, with steps of 2 weeks.

**Objectives:** We use 2 fitness functions which represent the compromising health and societal impacts of policy scheduled. We quality these impacts with the total number of total deaths between April 30 and December 30 and the mean of the mobility feature values over the same period. The first objective must be minimized while the second is maximized.

**Constraints:** For a policy schedule to be an acceptable solution, we require that the number of critical cases never exceeds the hospitalization capacity (ICU) of the country of interest. This is an important requirement for policymakers, as critical cases which may not be correctly hospitalized likely result in additional deaths.

**Selector:** Current Pareto-front solutions are selected in priority. If there are more than the population size, they are filtered based on a crowding distance (here, we use Manhattan distance in the objective space). Otherwise, we fill the population with non-optimal solutions, selected using a binary tournament selection: Pareto-dominant solutions are retained in priority; in case of non-dominance, crowding distance to the Pareto front is used for tie-breaking.

We rely on Pymoo<sup>5</sup> to implement our NSGA-II search and use the library’s default values for the remaining parameters like mutation rate or crossover.

<sup>4</sup>The closer  $R^2$  score to 1 the better and the closer RMSE to 0 the better

<sup>5</sup>[www.pymoo.org](http://www.pymoo.org), the most starred python GA library on Github



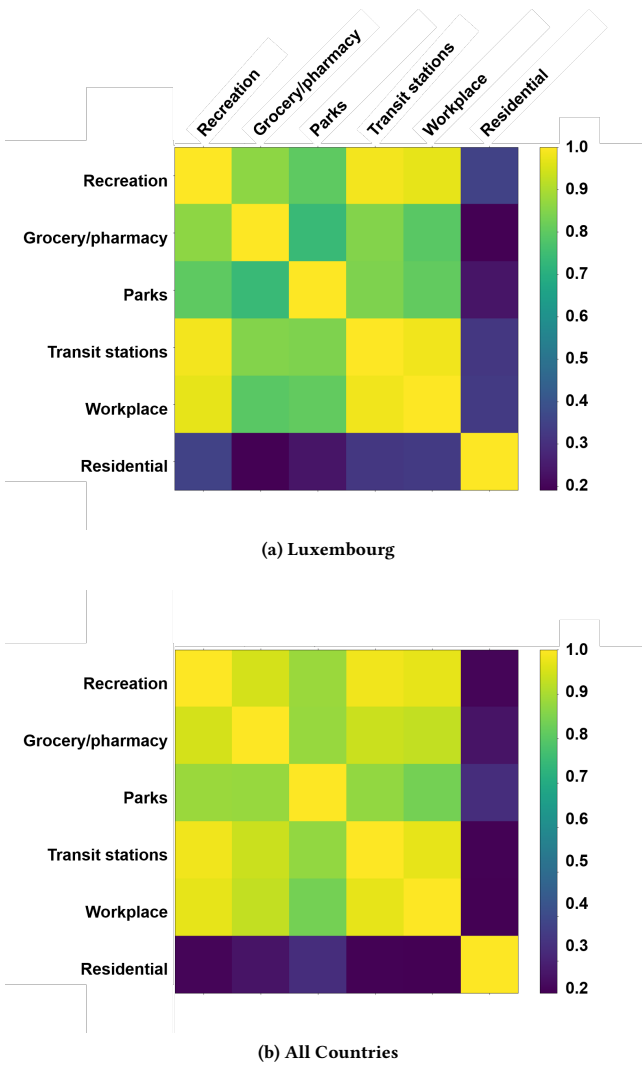


Figure 3: Mobility feature correlations.

#### 4 RESEARCH QUESTIONS

Our end goal is to provide decision-makers with a tool allowing to easily generate exit strategies in order to evaluate their impact. To achieve this, we use a deep neural network as a proxy to evaluate the hyper-parameters of a SEI-HCRD model based on mobility and demographic data. However, to be useful, the neural network needs to be able to capture enough information in the mobility and demographic data alone to make accurate predictions. Hence, we formulate the first research question as follow:

**RQ1:** Can we predict the effective reproduction number based on mobility and demographic data?

Ultimately the proposed approach is intended to allow policy-makers to evaluate and select exit strategies by analysing their impact on multiple aspects such as the number of death, possible overflow of healthcare capacities or the perturbation of economic

activities. To evaluate the capacity of our approach to model such strategies, we evaluate it by predicting the impact of various popular scenarios. Thus, we ask the following question:

**RQ2:** How does our approach react under different exit strategies?

We conclude our investigation by a comparison of the impact of the popular scenarios with the impact of the one proposed by the search algorithm. The algorithm minimizes the number of deaths and the socio-economical impacts generated by a diminution of activities (mobility) while avoiding the over-saturation of healthcare capacities. To evaluate the results, we compared them to the “naive” scenarios formulated in the previous research question and thus ask:

**RQ3:** How do the exit strategies proposed by the search algorithm perform against popular ones?

### 5 RESULTS

#### 5.1 Predicting the Effective Reproduction Number

*Comparison of a fitted SEIR model and our DN-SEIR model.* We estimate the confidence interval by evaluating the mean and standard deviation of a Bayesian Ridge Regressor on each element of the test set (using the same training set as the FFNN). We use a grid search over 3 values for each of its 2 hyper-parameters  $\alpha_{init}$  (0.1, 1, 1.9) and  $\lambda_{init}$  (1, 0.1, 0.01).


We compare in Table 3 the predicted cases of the time-dependant SEIR (Fitted on past cases/deaths) and the DN-SEIR approach. The DN-SEIR approach has a lower error to the ground truth values of cases for 9 over 12 countries in comparison with a time regression approach. Besides, the ground truth falls within the confidence interval of the DN-SEIR model for 10 of the 12 countries evaluated. It is worth noting that while 5 of the 12 countries lie under 5% error, 9 over 12 countries do not exceed 15% error.

In Figure 4, we evaluate the RMSE between the predicted cases of the SEIR model and the DN-SEIR approach (with its optimistic and pessimistic boundaries) across all the countries of the dataset. The simulation spans on 7 days (instead of 12 days for the results presented in table 3) and we compare the number of cases on April 29th, the last day of available mobility data for all countries. We use the Wilcoxon signed-rank test to compare if two distributions are equal. The results show that for all three prediction DN-SEIR, DN-SEIR max and DN-SEIR min we can reject the hypothesis ( $p$ -values  $\ll 0.05$ ) that they are equal to the SEIR prediction. We then perform a Vargha and Delaney’s  $A_{12}$  test to analyze the effect size. We see that our approach generates a lower RMSE with a small effect size. These results indicate that even in very short term (7 days), relying only on the mobility data yields better results than applying a regression over the past values.

*Model interpretation at the global level.* We fit SHAP on the full dataset and show in Figure 5 the summary of its impact analysis where the features are ordered in decreasing order of influence. The distribution of each feature spans horizontally to inform on the impact of the feature on the final decision (positive on the right side), while the colour of each point of the distribution provides

Country	SEIR cases	DN-SEIR cases	True cases	$\epsilon_r$	$\epsilon_{rSEIR}$
Belgium	47,219	51,532 [45,130-60,464]	47,859	0.07	0.01
France	213,103	191,188 [166,043-224,554]	128,442	0.33*	0.66
Germany	161,383	181,438 [163,333-207,151]	157,641	0.13	0.02
Greece	2,429	2,478 [2,343-2,666]	2,576	0.03*	0.06
Italy	6 201,802	202,369 [187,205-223,194]	203,591	<0.01*	<0.01
Luxembourg	3,668	3,724 [3,578-3,936]	3,769	0.01*	0.03
Spain	209,646	230,794 [211,299-257,782]	240,743	0.04*	0.13
Brazil	60,714	68,271 [55,268-86,164]	78,162	0.14*	0.22
Cameroon	1,432	1,645 [1,375-2,003]	1,832	0.11*	0.22
Canada	77,614	63,520 [51,977-78,913]	51,597	0.19*	0.5
Japan	12,250	11,353 [10,010-13,288]	14,088	0.24	0.13
United Kingdom	201,701	160,442 [134,251-188,088]	165,221	0.03*	0.22

**Table 3: Total cases as of 29/04 as predicted by a time-regression SEIR, and by DN-SEIR model. Both models trained until 11/04.  $\epsilon_r$  and  $\epsilon_{rSEIR}$  are the absolute relative error of the DN-SEIR model and time-regression SEIR respectively. (\*) indicates that the DN-SEIR yields less error than the time-regression SEIR.**



**Figure 4: RMSE between true cases and predicted cases by our DN-SEIR model and a fitted SEIR model.**

information on the range of values of the feature that have an impact (redder color codes for higher values of the feature). We can see for instance for the feature **transit station** that higher values have a positive impact on prediction, i.e higher values of **transit station** translate into higher values of  $R_t$ . This insight is common across the three countries. The same goes for the features related to **retail and recreation** activities.

*SHAP* shows how the trends in the feature, modelled in our approach using smoothed features, play a significant role in the final prediction, and shows an opposing contribution to its associated daily feature. For transit, retail, and park features, the trends values over 5, 10 and 15 days counter the impact of their respective daily values, yet on a lower scale. This indicates an inertia phenomenon that can be explained by the actual delay between the actual numbers ( $R_t$ , cases, deaths) and the reported one and also the delay inherent to the epidemiological model.

The comparison of the three countries also shows that mobility features have a much higher impact on the prediction in Italy and

Luxembourg than Japan as their SHAP impact is much wider. This hints that other social distancing features in Japan could reduce the impact of mobility (masks for instance).

**RQ1 Answer:** Our approach yields predictions with much lower errors than pure epidemiological models in 75% of the cases and achieves a 95%  $R^2$  score when the learning is transferred and tested on unseen countries.

## 5.2 Mid-term predictions with exit strategies

In this section, we investigate four prediction strategies for an exit from the lockdowns that were taking place all over the world. The goal of the exit strategies is to allow a return to normal activities while minimizing the impact on the number of deaths and avoiding peaks in hospitalization that would saturate healthcare facilities. The strategies that we are investigating are the following:

- **Hard exit:** In this strategy, all mobility activities are resumed to normal on May 11, 2020.
- **Progressive exit:** Mobility activities are gradually restored, with an increase of 15% of the activity every 2 weeks until the pre-lockdown activity level is reached.
- **Cyclic exit:** Every two weeks, activity is resumed to normal then brought back to lock down situation. The process is repeated for 4 cycles, thus ending on 03/08/2020.
- **Status Quo:** The current situation (as of April 30th) is maintained for the entirety of the period.

Figure 6 shows the evolution of  $R_t$  values for Luxembourg, Italy and Japan. As expected we see no evolution from the initial value in the no exit case, a cyclic fluctuation in the case of a cyclic exit, a soft increase of the  $R_t$  values when applying a progressive exit and finally, a brutal jump in the case of a hard exit.  $R_t$  reaches a plateau typically quite rapidly after strategies are applied. The plateau depends on the mobility condition, therefore, we see two plateaus in the results, one with all activities remaining closed (no exit), and one where all the other strategies reach the same plateau when the mobility levels are restored to their pre-lockdown baseline.




Figure 5: Global interpretation of the model for Luxembourg, Italy and Japan.




Figure 6: Evolution of the  $R_t$  values for the four exit strategies modelled, i.e a hard exit, a progressive exit, a cyclic exit and status quo for Luxembourg, Italy and Japan.

We evaluate these strategies for the three countries Luxembourg, Italy and Japan and obtain the results depicted by Table 4. We choose those three countries because they present difference with respect to their demography, mitigation strategies and number of deaths attributed to COVID-19. Furthermore, they are amongst the few to provide reliable data about hospitalization capacities, hence allowing us to incorporate this information when looking for optimal policy schedules.

We compute for each strategy the Area Under Curve (AUC) of each mobility metric and provide the mean across all mobility values, and compare these hand-crafted strategies with the ones found by our Genetic Algorithm search. The search is run on 100 generations with a size of population of 100 and the hard constraint that critical hospitalizations should not exceed the country’s ICU capacity (2,054 for Italy, 1,822 for Japan and 42 for Luxembourg, [6]). All strategies are evaluated between April 30th and September 30th.

We report 2 metrics, the total deaths on September 30th and the mean Area Under Curve across all the mobility features over the 5 months. The later reflects an economic objective that we

need to maximize while the former is a healthcare objective. We state in the table three strategies found on the pareto. S1 is the pareto solution with the lowest death toll, S3 is the strategy with the highest mobility activity (and hence highest death toll) and S2 is the median death toll.

Our study shows that progressive lift strategies yield a similar economic footprint as 2-weeks cyclic strategies with fewer casualties (7% fewer deaths for Italy and 10 times less for Luxembourg).

**RQ2 Answer:** Our approach allows to see drastic changes based on different exit strategies. The progressive strategy offers in our experiments a better outcome than a hard or a cyclic strategy.

The search-based strategies (S1, S2, S3) perform better than manual strategies on the death metric for Italy and Japan, and for Luxembourg. For Luxembourg, S1 performs as good as the progressive strategy. S1 in particular performs on a par with the Progressive lift strategy.



Overall, our results show that the search for exit strategies can be guided and restricted to the policy-makers constraints (i.e. hospital capacity) and yields actionable strategies within constrained computation time. We ran the experiments 100 times (x100 generations each) and the hypervolume of the pareto-solutions converges within 80 generations.

**RQ3 Answer:** The search algorithm yields better strategies than the popular ones both in term of impact on the activity and number of deaths while ensuring that the healthcare facilities are not overwhelmed.

Country	Strategy	Mobility AUC	Deaths
Luxembourg	Status Quo	-10,721.6	108
	Hard	-91.9	2,763
	Progressive	-2,774.43	114
	Cyclic	-2,487.02	2,002
	Pareto-S1	-2,381.45	165
	Pareto-S2	-2,370.7	635
	Pareto-S3	-2,289.7	697
Italy	Status Quo	-6,006.13	32,015
	Hard	-57.53	37,377
	Progressive	-4,124.93	31,987
	Cyclic	-3,689.13	34,427
	Pareto-S1	-8,412.125	29,449
	Pareto-S2	-7,570.0	29,450
	Pareto-S3	-7,275.38	29,452
Japan	Status Quo	-3,431.7	709
	Hard	14.52	710
	Progressive	-1,005.3	708
	Cyclic	-896.03	709
	Pareto-S1	-2,106.21	654
	Pareto-S2	-1,170.33	660
	Pareto-S3	-1,106.33	671

**Table 4: Exit strategies comparison. Higher AUC and lower deaths are better.**

## 6 CONCLUSION

In this paper, we studied *DN-SEIR*, a data-driven approach to evaluate the effective reproduction number of the COVID-19 epidemic. In particular, we considered both manual and search-based mitigation strategies, with the aim to help decision-makers in the evaluation and selection of exit strategies. To this end, we evaluated the state-of-the-art compartment model (i.e. SEIR) and shew that our approach yields predictions closer to the ground truth. We also demonstrated that learning can transfer across different countries and a simple FFNN provides accurate and interpretable predictions. Finally, we proposed a search-based approach to evaluate and find optimal strategies that satisfy the constraints of the health facilities and achieve a quick economic recovery with limited casualties.

Our approach paves the ways to automated strategy simulation and search and provides a simple, yet, powerful tool for policy makers to tailor exit strategies to their context and priorities. We can go further than our approach with better feature engineering or neural architecture search (with CNN or RNN). We can also extend the data-driven prediction of hyper-parameters not only

to the effective reproduction number but also to all the epidemiological parameters like hospitalization rate. This would require having access to accurate hospitalization data across a large pool of countries and can be achieved in the close future as more countries are sharing such data. Finally, we could extend our technique to a more-grained approach that takes into account age-specific or location-specific epidemiological models.

## ACKNOWLEDGMENT:

This work is mainly supported by the Luxembourg National Research Funds (FNR) through CORE project C18/IS/12669767/ STELLAR/LeTraon and covid19 projects COVID-19/2020-2/14863123/ PILOT/LeTraon and COVID-19/2020-1/14701707/ REBORN/Klein.

Special thanks to Guillaume HABEN from the University of Luxembourg for his help on the web simulator and his insights, and the AUF for their financial support through the CoronaSaaS project.

## REFERENCES

- [1] Sina F Ardabili, Amir Mosavi, Pedram Ghamisi, Filip Ferdinand, Annamaria R Varkonyi-Koczy, Uwe Reuter, Timon Rabczuk, and Peter M Atkinson. 2020. Covid-19 outbreak prediction with machine learning. *Available at SSRN 3580188* (2020).
- [2] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- [3] Mark Jit and Marc Brisson. 2011. Modelling the Epidemiology of Infectious Diseases for Decision Analysis. *Pharmacoeconomics* 29, 5 (may 2011), 371–386.
- [4] Fumito Koike and Nobuo Morimoto. 2018. Supervised forecasting of the range expansion of novel non-indigenous organisms: Alien pest organisms and the 2009 H1N1 flu pandemic. *Global Ecology and Biogeography* (04 2018).
- [5] Ying Liu, Albert A. Gayle, Annelies Wilder-Smith, and Joacim Rocklöv. 2020. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine* 27, 2 (mar 2020), 1–4.
- [6] Esteban Ortiz-Ospina, Max Roser, Hannah Ritchie, and Joe Hasell. 2020. Coronavirus Pandemic (COVID-19). *Our World in Data* (2020). <https://ourworldindata.org/coronavirus>.
- [7] Christoph Molnar. 2019. *Interpretable Machine Learning*.
- [8] Olav Titus Muurlink, Peter Stephenson, Mohammad Zahurul Islam, and Andrew W Taylor-Robinson. 2018. Long-term predictors of dengue outbreaks in Bangladesh: A data mining approach. *Infectious Disease Modelling* 3 (2018), 322–330.
- [9] Gaurav Pandey, Poonam Chaudhary, Rajan Gupta, and Saibal Pal. 2020. SEIR and Regression Model based COVID-19 outbreak predictions in India. *arXiv preprint* (2020).
- [10] T. Smith, N. Maire, A. Ross, M. Penny, N. Chitnis, A. Schapira, A. Studer, B. Genton, C. Lengeler, F. Tediosi, and et al. 2008. Towards a comprehensive simulation model of malaria epidemiology and control. *Parasitology* 135, 13 (2008), 1507–1516.
- [11] Nicholas Soares, David Chambers, Zachariah Carmichael, Anurag Daram, Dimpy P Shah, Kal Clark, Lloyd Potter, and Dhireesha Kudithipudi. 2020. *SIRNet: Understanding Social Distancing Measures with Hybrid Neural Network Model for COVID-19 Infectious Spread*. Technical Report. arXiv:2004.10376
- [12] Nicholas Soares, David Chambers, Zachariah Carmichael, Anurag Daram, Dimpy P. Shah, Kal Clark, Lloyd Potter, and Dhireesha Kudithipudi. 2020. *SIRNet: Understanding Social Distancing Measures with Hybrid Neural Network Model for COVID-19 Infectious Spread*.
- [13] M Vollmer, S Mishra, H Unwin, A Gandy, T Melan, V Bradley, H Zhu, H Coupland, I Hawryluk, M Hutchinson, et al. 2020. Report 20: A sub-national analysis of the rate of transmission of Covid-19 in Italy. (2020).
- [14] Michaela A C Vollmer, Swapnil Mishra, H Juliette T Unwin, Axel Gandy, et al. 2020. *Report 20 : Using mobility to estimate the transmission intensity of COVID-19 in Italy : A subnational analysis with future scenarios*. Technical Report May. Imperial College COVID-19 Response Team. 35 pages.
- [15] World Health Organisation. 2020. *a Coordinated Global Research Roadmap: 2019 Novel Coronavirus*. Number March.
- [16] Zifeng Yang, Zhiqi Zeng, Ke Wang, Sook-San Wong, Wenhua Liang, Mark Zanin, Peng Liu, Xudong Cao, Zhongqiang Gao, Zhitong Mai, et al. 2020. Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. *Journal of Thoracic Disease* (2020).