Open Data Science to Fight COVID-19: Winning the 500k XPRIZE Pandemic Response Challenge (Extended Abstract)*

Miguel Angel Lozano[1†], **Òscar Garibo-i-Orts**[2], **Eloy Piñol**[2], **Miguel Rebollo**[2], **Kristina Polotskaya**[3], **Miguel Angel Garcia-March**[2], **J. Alberto Conejero**[2], **Francisco Escolano**[1], **Nuria Oliver**[4]

[1]University of Alicante, Alicante, Spain
[2]IUMPA and VRAIN. Universitat Politècnica de València, València. Spain
[3]University Miguel Hernández, Elche, Spain
[4]ELLIS Alicante, Spain
malozano@ua.es, osgaor@upv.es, eloy.p.j@hotmail.com, mrebollo@upv.es, k.polotskaya@umh.es, garciamarch@mat.upv.es, aconejero@upv.es, sco@ua.es, nuria@ellisalicante.org

Abstract

We describe the deep learning-based COVID-19 cases predictor and the Pareto-optimal Non-Pharmaceutical Intervention (NPI) prescriptor developed by the winning team of the 500k XPRIZE Pandemic Response Challenge. The competition aimed at developing data-driven AI models to predict COVID-19 infection rates and to prescribe NPI Plans that governments, business leaders and organizations could implement to minimize harm when reopening their economies. In addition to the validation performed by XPRIZE with real data, our models were validated in a real-world scenario thanks to an ongoing collaboration with the Valencian Government in Spain. Our experience contributes to a necessary transition to more evidence-driven policy-making during a pandemic.

1 Introduction and Related Work

During a pandemic, predicting the number of infections under different circumstances is important to inform public health, healthcare and emergency system responses. Within the different approaches to predict the pandemics evolution, we find traditional compartmental meta-population models –such as SIR or SEIR [Hethcote, 2000], complex network [Pastor-Satorras *et al.*, 2015], agent-based individual [Ferguson *et al.*, 2005] and purely data-driven time series forecasting [Tayarani and Mohammad, 2020] and deep learning-based [Chatterjee *et al.*, 2020; Chimmula and Zhang, 2020; Arora *et al.*, 2020; Pereira *et al.*, 2020] models.

Given the exponential growth in the number of SARS-CoV-2 infections and the pressure in the health care systems, most countries have implemented non-pharmaceutical interventions (NPIs) during the current pandemic, designed to reduce human mobility and limit human interactions to contain the spread of the virus. How to model the impact that the applied NPIs have on the progression of the pandemic is a non-trivial task, particularly for traditional meta-population approaches. Moreover, the social and economic costs of applying NPIs for a sustained period have led to the largest global recession in history. Besides, the social cost of the pandemic is also staggering, preventing children and teenagers from attending schools, canceling cultural activities, and forbidding people to visit friends or relatives.

In view of these challenges, the XPRIZE foundation organized in November of 2020 a global competition, the 500K XPRIZE Pandemic Response Challenge sponsored by Cognizant [XPRIZE, 2021]. This 4-month challenge fostered the development of data-driven AI systems to *predict* COVID-19 infection rates and *prescribe* NPIs Plans that governments could implement to minimize harm when reopening their economies.

In this paper, we briefly summarize the predictor and prescriptor models developed by ValenciaIA4COVID, the competition's winning team. We direct the interested readers to [Lozano *et al.*, 2021] for a detailed description of both models. Our infections predictor is based on a novel architecture of Long Short Term Memory (LSTM) networks [Hochreiter and Schmidhuber, 1997]. Regarding the prescriptor part of our work, Miikkulainen *et al.* propose a neuroevolution approach to identify a Pareto-optimal set of NPIs [Miikkulainen *et al.*, 2021]. While this approach was recommended during the Challenge, our solution adopted a different methodology.

2 Data

The Challenge leveraged publicly available official COVID-19 case data together with the Oxford COVID-19 Government Response Tracker data set[1] as the main data sources to be used [Hale *et al.*, 2021]. It includes information for 186 countries and state/region-level data for the US, UK, Canada, and Brazil. The Challenge considered 182 countries[2], the 50

*This is an extended abstract for our ECML PKDD 2021 paper [Lozano *et al.*, 2021] (Best Applied Data Science Paper Award)

†Contact Author

[1]https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker

[2]Tonga, Malta, Turkmenistan & Virgin Islands were not included due to lack of reliable data.

US states and the 4 regions in the UK, yielding a total of 236 countries or regions. We will use GEO to denote the countries/regions.

The available data sources can be split into *case*-related data, *i.e.* number of daily confirmed COVID-19 cases, and *action or NPI*-related data, *i.e.* the NPIs and their level of activation each day for each GEO. In the Challenge, we considered 12 NPIs of two types: *confinement-based* and *public health-based*, with different possible levels of activation.

3 Predictors of COVID-19 Cases

This part of the Challenge required building a predictor of the number of confirmed COVID-19 cases in the 236 GEOs for up to 180 days into the future and considering the different NPIs implemented in each GEO. Next, we summarize our notation, followed by a description of our deep learning-based predictive model.

3.1 Notation

We use the following terms and notation:
1. GEO: We denote as GEO a country or a region (*e.g.* California). We use the index j to refer to each GEO.
2. Population (P^j)**:** P^j denotes the total population of GEO j, assuming that it is constant during the entire period of time.
3. NewCases (X_n^j)**:** The daily number of new cases on day n and GEO j is denoted by X_n^j starting on March, 11th 2020.
4. ConfirmedCases (Y_n^j)**:** The cumulative number of confirmed cases up to day n in GEO j is $Y_n^j = \sum_{i=1}^{n} X_i^j$.
5. SmoothedNewCases (Z_n^j)**:** We compute the average number of new cases between days $n-K+1$ and n in GEO j as $Z_n^j = \frac{1}{K}\sum_{i=0}^{K-1} X_{n-i}^j$. This prevents noise due to different imputation policies We use $K=7$ to smooth over one week.
6. CaseRatio (C_n^j)**:** The ratio of cases between two consecutive days is denoted by $C_n^j = Z_n^j / Z_{n-1}^j$.
7. Susceptible Population (S_n^j)**:** The number of susceptible individuals to be infected with coronavirus on day n and for GEO j is denoted by S_n^j.
8. ScaledCaseRatio (R_n^j)**:** It is the CaseRatio C_n^j divided by the proportion of susceptible individuals in GEO j, $R_n^j = C_n^j \frac{P^j}{S_n^j}$. It captures the effects of a finite population.
9. Action (A_n^j)**:** Vector of applied NPIs in GEO j on day n.
10. Stringency of A_n^j ($Str^j_{A_n}$)**:** The stringency of an NPI applied in GEO j on day n is $Str^j_{A_n} = \sum_{i=C1}^{H6} a_n^j(i) \cdot Cost^j(i)$, where $Cost^j$ is the cost vector of each of the 12 different types of NPIs ([C1...C8,H1,H2,H3,H6]) in GEO j.
11. Intervention Policy (IP): The sequence of daily 12-dimensional NPI or action vectors applied during period T.
12. Stringency of an Intervention Policy: The sum of stringencies of actions A_n^j applied each day n during T.

We denote estimations with a $\widehat{\ }$ symbol, *e.g.* $\widehat{X}_n^j$ is the estimated number new cases and $\widehat{R}_n^j$ the estimated scaled case ratio, both for GEO j and day n.

3.2 SIR Epidemiological Model

The predictors model the dynamics of the epidemics in each GEO j using an underlying basic SIR compartmental meta-

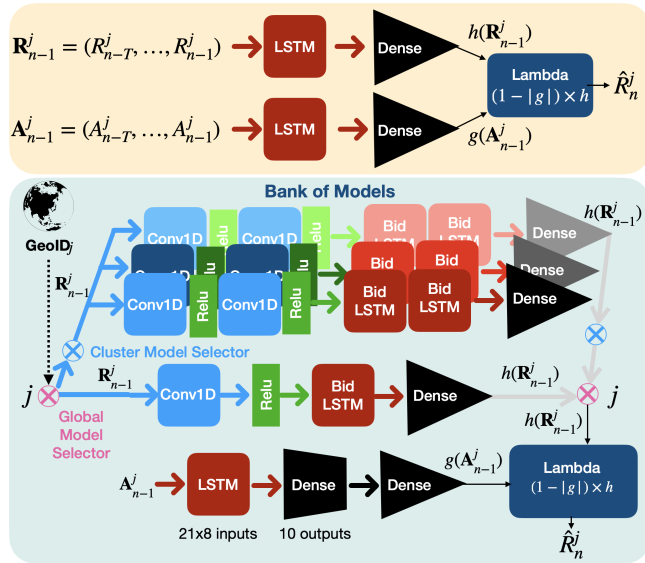

Figure 1: Top: Baseline LSTM-based predictor; Bottom: ValenciaIA4COVID predictor.

population model [Allen, 1994]. Here, the population is divided into three different states: S (Susceptible), Z (Infected), and D (Removed, due to recovery or death). The evolution of the number of infected individuals is given by $\frac{dZ^j}{dt} = \beta \frac{S^j}{P_j} Z^j - \mu Z^j$, where β is the infection rate which controls the probability of transition between the S and Z; and μ is the recovery or removal rate, controlling the probability of transition between the Z and D states. The goal of the predictors is to estimate $\widehat{R}_n^j$ given the data up to day $n-1$, where $R_n^j = \frac{(1-\mu)P_j}{S_n^j} + \beta = \frac{Z_n^j}{Z_{n-1}^j}\frac{P^j}{S_n^j}$. Note that $\widehat{X}_n^j$ can be computed from $\widehat{R}_n^j$ [Lozano *et al.*, 2021]. Since R_n^j depends on the transmission rate, which in turn depends on the NPIs, the predictors consider both the number of COVID-19 infections (*context*) and the applied NPIs (*actions*) each day in each GEO.

3.3 ValenciaIA4COVID (V4C) Predictor

The Challenge organizers provided the baseline, or standard predictor [Miikkulainen *et al.*, 2021]. It consists of two parallel LSTMs, one to model the *context* – given by the R_n^j – and the other to model the *actions* (A_n^j) applied on day n in GEO j. Figure 1 (top) depicts the architecture of this baseline model. It uses the context and action data to obtain predictions separately, joining both outputs via a lambda merge layer.

Similarly to the baseline predictor, we implemented an architecture with 2 LSTM-based branches: a *context* branch, where we modeled the R_n time series and an *action* branch, where we modeled the time series of the eight confinement-based ([C1...C8]) Non-pharmaceutical Interventions. While we did not consider public health-based NPIs, we improved the baseline predictor in several ways. We denote this improved model as the ValenciaIA4COVID or V4C predictor.

Context branch. Given the large variability in the time series of confirmed COVID-19 cases depending on the GEO, we developed a bank of LSTM context models, shown in Figure 1 (bottom). We created such a bank by clustering the GEOS via a K-means algorithm applied to the time series of the reported number of COVID-19 cases per 100K inhabitants. We optimized the number of clusters using the Elbow method, obtaining 15 different clusters. We implemented two different LSTM-based architectures, as depicted in Figure 1 (bottom): one for the *reference* model and the other for each of the eight *cluster* models. The *reference* model includes a convolutional layer with ReLu activation function and a bidirectional LSTM followed by a dense layer. Each convolutional layer has 64 filters of size 8. This reference model empirically generalized well for 135 GEOs. The *cluster* models consist of a stacked version of the architecture of the reference model, with two convolutional layers and two stacked bidirectional LSTMs.

Action branch. We used an LSTM followed by two dense layers to smooth the output and hence better capture non-linearities.

Merge function. The two branches use the data from the last 21 days that are combined into a final dense layer to get the predicted $\widehat{R}_n$. The outputs of each branch (h and g) are merged by the lambda function $\widehat{R}_n^j = f(\mathbf{A}_{n-1}^j, \mathbf{R}_{n-1}^j) = (1 - g(\mathbf{A}_{n-1}^j))h(\mathbf{R}_{n-1}^j)$. Thus, the predicted $\widehat{R}_n$ provided by the *context* branch is modified by the output from the *action* branch. The stricter the NPIs, the larger the output from the action layer, thus reducing the context layer's output. Finally, once the model gives the predicted $\widehat{R}_n$, the predicted number of new infections for the day n, $\widehat{X}_n$, is obtained.

4 Prescriptor of Intervention Policies

The final phase of the XPRIZE competition required building a *prescriptor* which would recommend for each GEO and for any period, up to 10 different Intervention Policies (IP) with the best balance between their economic/social cost and the resulting number of COVID-19 cases.

Thus, it entailed solving a two-objective optimization problem by identifying the set of solutions that would be on the Pareto front [Belakaria *et al.*, 2019; Lu *et al.*, 2020; Miikkulainen *et al.*, 2021]. On the one hand, there is the *stringency* of a certain IP which captured the sum of the costs of implementing such a policy. On the other hand, there is the number of COVID-19 cases per 100K inhabitants which would result from applying such IP. Given that this is a hypothetical scenario, the number of COVID-19 infections under the IPs was estimated by the *baseline* or standard predictor provided by the XPRIZE Challenge organizers. All the teams used the same predictor to enable the judges to properly compare the prescriptors from different teams.

Our goal in the Prescription phase of the competition was to develop an *interpretable*, data-driven, and flexible prescription framework that would be usable by non-machine-learning experts, such as citizens and policymakers in the Valencian Government. An intervention policy IP_1 *dominates* another intervention policy IP_2 if the stringency(IP_1) $\leq$ stringency(IP_2) and the resulting number of COVID-19 cases under IP_1 $<$ than under IP_2. The goal was to find up to 10 IPs for each GEO, for any period, and any costs that would dominate the rest of possible IPs. As in the case of our predictor, we decided to combine complementary approaches to have a more robust solution, shown in Figure 2.

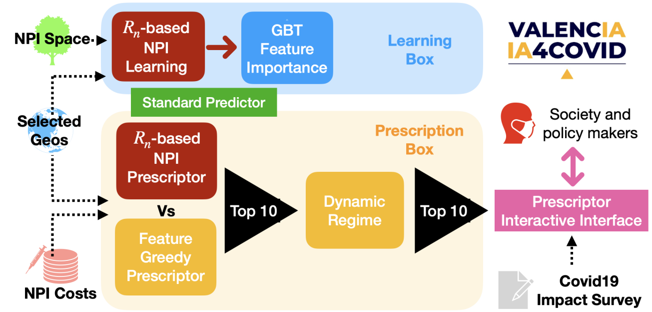

Figure 2: Valencia4COVID Prescriptor.

We first modeled the NPI - COVID-19 cases space. Considering all the possible values of each dimension of the NPI or action vector, 7,776,000 possible combinations of NPI vectors could be applied at each time step. In our experiments, we observed that the **same NPI vector** would lead to the **same convergence** $\widehat{R}_n$ in **all the GEOs** and over **any time period** provided that the NPI was applied for long enough ($\sim$21 days). We refer to this phenomenon as the R_n *synchronization principle*.

Method 1: R_n-Based NPI Selection

Based on the R_n *synchronization principle*, one could easily obtain the Pareto-optimal front of intervention policies if the mapping between the 7.8 million of possible combinations of the NPI vector and their associated convergence $\widehat{R}_n$ were to be known. Unfortunately, generating such a mapping was not feasible in the time frame provided by the Challenge as it would require making millions of calls to the predict function. Hence, we opted for computing a sample of such a matrix, obtained as (1) all the NPI vectors with stringencies [0 to 6] and [28 to 34]; (2) all NPI vectors with one and two non-zero entries; and (3) a random sample of 10,000 NPIs.

For each NPI in the sample, we computed the convergence $\widehat{R}_n$, and the resulting total number of COVID-19 cases in 20 and 60 days. Using this NPI-$\widehat{R}_n$ matrix, we trained state-of-the-art machine-learning models to predict the $\widehat{R}_n$ for any given NPI vector. The best performing and explainable model were Gradient Boosted Trees. While their MAE was still too large to be used to fill-in all the missing elements in the NPI-$\widehat{R}_n$ matrix, we carried out a feature importance analysis and discovered that the C2, C1, H2, C4, and C5 interventions are, in this order, the most important to predict their associated $\widehat{R}_n$ and hence the resulting number of cases. Thus, we also included in our NPI-$\widehat{R}_n$ matrix all the NPI vectors with non-zero values in their C1, C2, C4, C5, and H2 interventions and zero in the others. This led to 54,652 NPI vectors.

As a result, we generated a matrix with the mapping between these different NPI vectors, their associated stringencies (at cost 1), the number of cases that they would lead to at 20 days and at 60 days, and their convergence $\widehat{R}_n$. At run time, given an input cost vector, the prescriptor computes the stringency of each row in the NPI-$\widehat{R}_n$ matrix. With this information, it identifies the NPI combinations that are on the Pareto front by selecting those that lead to the best trade-off between their stringencies, their associated number of cases at 20 and 60 days, and their convergence $\widehat{R}_n$.

Method 2: Feature Greedy NPI Selection
As per the feature importance analysis described above and given a cost vector, we developed a greedy NPI prescriptor as follows: each dimension of the NPI vector is ranked by its *priority*, computed as its feature importance divided by its cost. This prescriptor consists of a greedy algorithm that consecutively activates each NPI dimension to its maximum value ordered by its priority.

Prescriptor Combination
Each of the methods above provides a set of NPI recommendations for each GEO for each day. From such a set, we select the 10 best NPIs that: (1) are not dominated by any other NPI; and (2) contribute to having a diverse set of NPIs that cover the full range of possible stringency values.

Intervention Policy Definition
Finally, the prescriptor needs to provide a set of up to 10 Intervention Policies, *i.e. dynamic* regimes of applying the selected NPIs over the period of interest. To do so, we compute all possible combinations of subsequently applying the selected NPIs in chunks of minimum 14 days (to enable the NPIs to act) and identify the Pareto-front set of combinations that would yield the optimal trade-off between stringency and number of cases. The total number of chunks is dynamically determined. From this set of combinations, we again select the 10 that (1) are not dominated by any other policy; (2) contribute to having a diverse set of policies along the stringency axis and (3) minimize the changes in NPIs, as every NPI change has a social cost.

5 Experimental Results

5.1 Predictor

Table 1 displays the MAE per 100K inhabitants and the Mean Rank of the proposed model when compared to the baseline model provided by the XPRIZE organizers. We also include the results of only using our *reference* context model without the clusters. Our model outperforms the baseline model in all evaluation scenarios in terms of MAE and Mean Rank.

All the models were trained with data from [Hale *et al.*, 2020], from March 11th to December 17th 2020, for the 20 most affected countries in terms of confirmed cases. As consistency is important, we evaluated the models both in short-term (3 weeks) and long-term predictions (180 days).

In our collaboration with the President of the Valencian Government in Spain, we were able to share the predictions of our predictor during the 3rd wave of the COVID-19 pandemic that started right after Christmas of 2020. Our predictor was very accurate in predicting the evolution of the pandemic while taking into account the NPIs implemented at the time, providing valuable input to the Government in its decision-making.

Predictor	**Short-term**		**Long-term**	
	MAE	**Rank**	**MAE**	**Rank**
XPRIZE baseline	157.92	2.11	935.34	2.30
V4C (w/o clusters)	138.21	2.14	825.38	**1.83**
V4C with clusters	**126.33**	**1.75**	**803.59**	1.87

Table 1: Predictor results (MAE and Mean Rank) in the 236 GEOS.

Prescriptor	**Window size of prescription**		
	31-days	**61-days**	**180-days**
Greedy	127 / 1814	130 / 1829	163 / 1839
Feature greedy	921 / 114	930 / 117	**986** / 163
V4C prescriptor	**927 / 47**	**934 / 48**	**986 / 137**

Table 2: Prescriptor results: # of dominating / # of dominated prescriptions for 5-day (Aug 1st-Aug 5th, 2020), 31-day (Jan 1st-Jan 31st, 2021) and 90-day (Jan 1st-Mar 31st, 2021) time periods.

5.2 Prescriptor

Given the hypothetical nature of the prescriptor, we were not able to quantitatively evaluate its performance against ground truth. However, we carried out domination tests between the IPs recommended by our model compared to a greedy algorithm for the 236 GEOs in the Challenge and under both unitary and random costs policies for a period of 60 days into the future. Table 2 shows the number of times the IPs recommended by our prescriptor dominated and were dominated by the IPs suggested by the greedy approach for all GEOs. Moreover, our prescriptor provided the IP recommendations in under 2 hours for all GEOs in the Challenge, well below the maximum allowed limit of 6 hours.

6 Conclusions and Future Work

We have described the models developed by the winning team of the 500K XPRIZE Pandemic Response Challenge. In this prediction phase, we developed an LSTM-based bank of models which outperformed the baseline model provided by the Challenge organizers and yielded the third-best Mean Rank amongst all the teams. The Presidency of the Valencian Government successfully used our model in Spain during the third wave of COVID-19 infections in Dec. - Feb. 2021. In the prescription phase, we proposed a solution that leveraged the R_n *synchronization principle* to provide Pareto-optimal IPs that clearly dominated other approaches. Our work contributes to the necessary transition to more evidence-driven policy-making, particularly during a pandemic. Future lines of work include developing the intervention prescriptor within the Valencian Government, developing a theoretical proof of the R_n synchronization principle, and including the impact of vaccinations in our model.

Acknowledgements

The authors have been partially supported by grants FONDO SUPERA COVID-19 Santander-CRUE (CD4COVID19 2020–2021), Fundación BBVA for SARS-CoV-2 research (IA4COVID19 2020-2022) and the Valencian Government, grant VALENCIA IA4COVID (GVA-COVID19/2021/100). Francisco Escolano and Miguel Angel Lozano are supported by the project RTI2018-096223-B-I00 of the Spanish Government, Miguel Angel Garcia-March is supported by MEFP Beatriz Galindo program (BEAGAL18/00203) and Nuria Oliver is partially supported by a collaboration agreement signed in 2021 (Convenio Singular 2021) between the ELLIS Unit Alicante Foundation and the Valencian Government (Generalitat Valenciana, Conselleria d'Innovació, Universitats, Ciència i Societat Digital, Dirección General para el Avance de la Sociedad Digital). We thank the University of Alicante's Institute for Computer Research for their support with computing resources, co-financed by the EU and ERDF funds through IDIFEDER/2020/003.

References

[Allen, 1994] L.J.S. Allen. Some discrete-time SI, SIR, and SIS epidemic models. *Math. Biosci.*, 124(1):83–105, 1994.

[Arora *et al.*, 2020] Parul Arora, Himanshu Kumar, and Bijaya Ketan Panigrahi. Prediction and analysis of covid-19 positive cases using deep learning models: A descriptive case study of india. *Chaos, Solitons & Fractals*, 139:110017, 2020.

[Belakaria *et al.*, 2019] S. Belakaria, A. Deshwal, and J.R. Doppa. Max-value entropy search for multi-objective bayesian optimization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *NeurIPS*, volume 32, 2019.

[Chatterjee *et al.*, 2020] A. Chatterjee, M.W. Gerdes, and S.G. Martinez. Statistical explorations and univariate timeseries analysis on COVID-19 datasets to understand the trend of disease spreading and death. *Sensors*, 20(11):3089, 2020.

[Chimmula and Zhang, 2020] Vinay Kumar Reddy Chimmula and Lei Zhang. Time series forecasting of covid-19 transmission in canada using LSTM networks. *Chaos, Solitons & Fractals*, 135:109864, 2020.

[Ferguson *et al.*, 2005] N.M. Ferguson, D.A.T. Cummings, S. Cauchemez, C. Fraser, and *e*t al. Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature*, 437(7056):209–214, 2005.

[Hale *et al.*, 2020] Thomas Hale, Sam Webster, Anna Petherick, Toby Phillips, and *e*t al. Oxford covid-19 government response tracker. https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker, 2020. Last accessed 29 Jan 2021.

[Hale *et al.*, 2021] T. Hale, N. Angrist, R. Goldszmidt, B. Kira, and *et al.* A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nat. Hum. Behav.*, pages 1–10, 2021.

[Hethcote, 2000] H.W. Hethcote. The mathematics of infectious diseases. *SIAM Rev.*, 42(4):599–653, 2000.

[Hochreiter and Schmidhuber, 1997] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.

[Lozano *et al.*, 2021] Miguel Angel Lozano, Òscar Garibo i Orts, Eloy Piñol, Miguel Rebollo, Kristina Polotskaya, Miguel Ángel García-March, J. Alberto Conejero, Francisco Escolano, and Nuria Oliver. Open data science to fight COVID-19: winning the 500k XPRIZE pandemic response challenge. In Yuxiao Dong, Nicolas Kourtellis, Barbara Hammer, and José Antonio Lozano, editors, *Machine Learning and Knowledge Discovery in Databases. Applied Data Science Track - European Conference, ECML PKDD 2021, Bilbao, Spain, September 13-17, 2021, Proceedings, Part IV*, volume 12978 of *Lecture Notes in Computer Science*, pages 384–399. Springer, 2021.

[Lu *et al.*, 2020] Z. Lu, I. Whalen, Y. Dhebar, K. Deb, and *et al.* NSGA-Net: Neural architecture search using multi-objective genetic algorithm (extended abstract). In Christian Bessiere, editor, *Proc. of the 29th Int. J. Conf. on AI, IJCAI-20*, pages 4750–4754, 2020.

[Miikkulainen *et al.*, 2021] R. Miikkulainen, O. Francon, E. Meyerson, X. Qiu, and *et al.* From prediction to prescription: Evolutionary optimization of nonpharmaceutical interventions in the COVID-19 pandemic. *IEEE Trans. Evol. Comput.*, 25(2):386–401, 2021.

[Pastor-Satorras *et al.*, 2015] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani. Epidemic processes in complex networks. *Rev. Mod. Phys.*, 87:925–979, Aug 2015.

[Pereira *et al.*, 2020] I.G.a Pereira, J.M. Guerin, Andouglas G. Silva J., G.S. Garcia, P. Piscitelli, A. Miani, C.o Distante, and L.M.G. Gonçalves. Forecasting COVID-19 dynamics in Brazil: a data driven approach. *Int. J. Environ. Res. Public Health*, 17(14):5115, 2020.

[Tayarani and Mohammad, 2020] N. Tayarani and H. Mohammad. Applications of artificial intelligence in battling against COVID-19: A literature review. *Chaos, Solitons & Fractals*, page 110338, 2020.

[XPRIZE, 2021] XPRIZE. 500k XPRIZE Pandemic Response Challenge, sponsored by Cognizant. https://www.xprize.org/challenge/pandemicresponse, 2021. Accessed: 2022-02-01.