QUALITATIVE PHASE 2 DOCUMENT: VALENCIA IA4COVID TEAM

1. Actionability and Usability

Our goal in the Prescription phase of the competition is to develop an *interpretable*, data-driven and flexible prescription framework that would be usable by non machine-learning experts, such as citizens and policy makers in the Valencian Government. Our design principles are therefore driven by developing interpretable and transparent models. The interactive visualization of our prescriptor can be accessed via: https://public.tableau.com/profile/kristina.p8284#!/vizhome/Prescriptions_16117279637400/Visualize. A screenshot of our interactive visualization is found in Annex V.

Given the intervention costs, it automatically generates up to 10 Pareto-optimal intervention plans. For each plan, it shows the resulting number of cases and overall stringency, its position on the Pareto front and the activation regime of each of the 12 types of interventions that are part of the plan.

2. Explanation

The challenge entails finding the set of Pareto-optimal intervention policies with the best trade-off between their economic and social cost and their associated number of resulting COVID19 cases. As in the case of our predictor, we strongly believe in the strength of combining **complementary approaches** to have a more robust solution. Thus, we explored the combination of 3 methods. Fig 1 depicts the architecture of our approach.

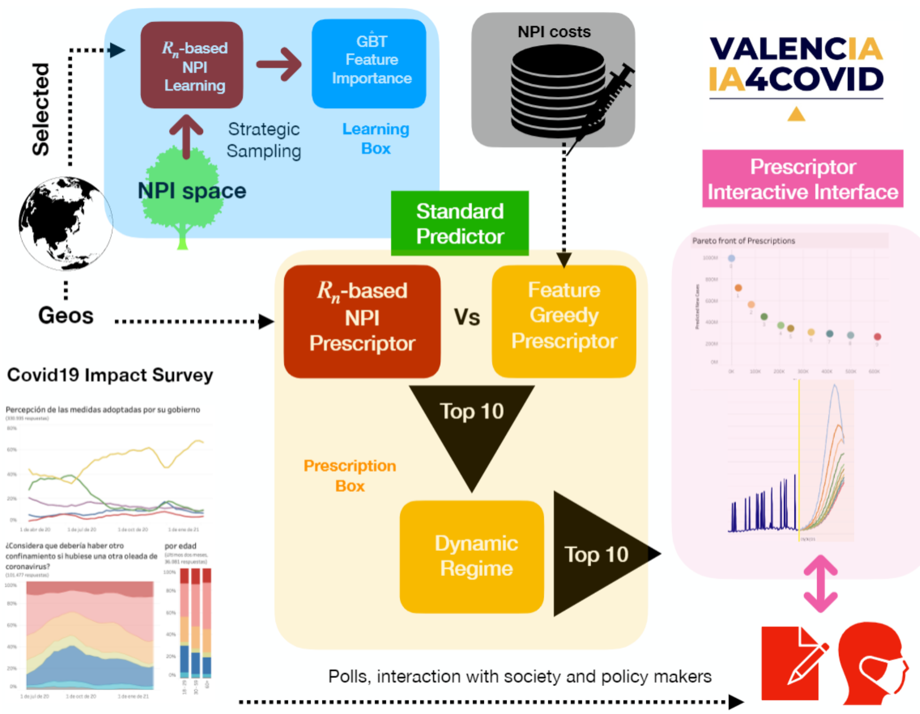

Fig 1. Architecture of the proposed prescription approach

2.1. Modeling the NPI - COVID19 cases space

Before building our model, we performed an exploratory data analysis of the problem space. Our goal was to shed light on the relationship between the Non-Pharmaceutical Interventions (NPIs) and the resulting number of COVID19 cases. An Intervention Policy (IP) consists of a sequence of daily 12-dimensional NPI vectors applied over a time period. Each dimension of the NPI vector corresponds to a different type of intervention, [C1, C2, C3, C4, C5, C6, C7, C8, H1, H2, H3, H6], where 'C' denotes confinement-based and 'H' denotes health-based interventions. Taking into consideration the possible values of each dimension of the NPI vector, there are 7,776,000 possible combinations of NPI vectors that could be applied at each time step. Annex I depicts the histogram of possible NPI combinations.

Each NPI vector, when applied during a minimum amount of time, would lead to a reduction/increase in the number of COVID19 cases in the country/region where it is applied. To better understand the impact that different NPI vectors have on the number of COVID19 cases, we ran numerous experiments where we called the predict function with different NPI scenarios on a sample of 21 countries over varying time periods of between 30 and 90 days. We obtained the resulting total number of cases, the total number of cases in the first 20 days after applying the intervention and the convergence Rn. For day n, Rn is defined as $R_n = PZ_n/(S_{n-1}Z_{n-1})$ where Z_n is the average number of daily cases during the last 7 days, P is the total population, and S_{n-1} is the of population that is susceptible to being infected by coronavirus (S_{n-1}=Population - Cumulative_number_cases up to day n-1).

In our experiments, we observed that the **same NPI vector** would lead to the **same convergence Rn** in **all the countries** and over **any time period** provided that the NPI was applied for long enough. The necessary time for the NPI to converge to its associated convergence Rn value is inversely proportional to the Rn: the larger the Rn, the faster the convergence. Also, the larger the Rn, the larger the number of resulting COVID19 cases. We refer to this finding as the **Rn synchronization principle.** Moreover, all countries underwent a **transitory** period of ~21 days since the application of a certain NPI before their Rn started converging towards its convergence value. The results of this analysis are described in Annex II.

2.2. Prescription method 1: Rn-based NPI selection

Based on this finding, one could easily obtain the Pareto-optimal front of intervention policies if the mapping between the 7.78 million of possible combinations of the NPI vector and their associated convergence Rn is known. Unfortunately, generating such a matrix 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 follows and depicted in Annex II: we computed the NPI-Rn mapping for **all possible** combinations of NPI vectors with stringencies[1] between [0-6] and [34-28]. Next, we added a **random sample** of NPI combinations and **all combinations** of NPI vectors with 1 and 2 non-zero entries.

Using this NPI-Rn matrix, we trained different state-of-the-art machine-learning models to predict the Rn for any given NPI combination. The best performing and interpretable models were Gradient Boosted Trees, which obtained a MAE on the test set of 0.0003. While such MAE was still too high for us to be able to fill-in all the missing elements in the NPI-Rn 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 Rn and hence the resulting number of COVID19 cases (our feature importance analysis is described in Annex III).

Thus, we also included in our NPI-Rn matrix **all combinations** of the NPI vectors with non-zero values in their C1, C2, C4, C5 and H2 interventions and zero in the rest of dimensions.

As a result, we generate a matrix with the mapping between **54,652** different NPI combinations and 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 Rn. We did all computations for a sample of 21 countries/regions described in Annex I.

At run time, given an input cost vector, we compute the stringency of each row in the matrix and identify the NPI combinations that are in 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 Rn.

2.2. Prescription method 2: Feature-based greedy NPI selection

Given the feature importance analysis described above and given a cost vector, we rank each dimension of the NPI vector by its priority = feat_importance/cost. We then run a greedy algorithm that activates each dimension consecutively by order of its priority. This strategy is related to the greedy strategies developed to solve the knapsack problem[2].

2.3. Prescription method 3: Neuroevolution-based NPI selection

Third, we developed a neuroevolution approach using the NEAT framework. Given that we prioritize interpretability, we see the neuroevolution as a complement to our main, interpretable method. We experimented with niching and different fitness functions. We obtained the best results a fitness function that uses the Rn rather than the number of cases, given by $fitness = -((2 * Rn)^2 + (log(stringency/2000 + 10)^2)$. We set the niching parameter to 3.0 and the number of individuals to 10. Unfortunately, all the solutions proposed by this approach were dominated by the solutions proposed by the two methods previously described: the Rn-based and feature-based greedy NPI selection algorithms. Thus, we decided to exclude the neuroevolved solutions from the prescribe function that we uploaded in the sandbox.

2.4. Model combination

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

2.5. Dynamic policy definition

Finally, we need to identify a dynamic regime of applying the selected NPIs over the time 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;

[1] Computed with equal, unitary costs

[2] https://en.wikipedia.org/wiki/Knapsack_problem

(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 from a practical perspective.

3. Addressing the Challenge

We have followed all the rules and recommendations suggested by the organizers. The prescriptions are automatically provided by our model without any manual settings.

4. Inclusivity and Fairness

Diversity is a key pillar in our team and in our work. We are a diverse team which includes computer scientists, engineers, physicists, mathematicians, economists and public policy makers. Our team is one of the few teams that is co-lead by a female scientist. Moreover, our approach combines 3 different methods to ensure that a diverse set of solutions are considered. Our framework also leverages citizen data from our large-scale citizen survey called COVID19ImpactSurvey[3] to determine the social cost of different interventions in the real-world. With over 500,000 answers from Spain, Italy, Germany and Brazil, and covering a wide range of demographic groups across different geographies, our survey enables us to consider in an inclusive way the economic and social cost of the confinement interventions and their impact on people's lives. Via the survey, we know how supportive the population is of applying more interventions, how compliant they are with the already implemented measures, and the psychological, economic and labor impact of the interventions. Finally, our close collaboration with the Valencian Government[4] enables our model to include valuable information related to the real cost of different NPIs in our region. We plan to use our model to assist the Valencian Government in their decision making regarding the interventions to deploy in the months to come.

5. Generality

Our aim in the challenge has been to provide meaningful prescriptions for all the regions. We have not developed speciality regions within the challenge. However, our prescription model --similarly to our predictor-- will be used to assist the Valencian Government of Spain.

6. Consistency

We performed numerous tests of our prescriptor under different costs and time frames and both for a selection of the 20 most affected countries/regions and for all the countries/regions in the challenge.

7. Transparency and Trust

Transparency is a key design principle for us, given that we plan to use our prescriptor in the Valencian region. Therefore, we prioritized an interpretable approach to the challenge vs a deep learning-based approach, which we envisioned as a complement to our main approach. We have also devoted significant effort in building an intuitive visualization that would enable policy makers and citizens to easily understand the trade-offs made by each of the recommended policies. Finally, our citizen survey helps inform policy makers about the social cost of different interventions, the population's compliance with existing measures and their willingness to be subjected to additional measures.

8. Collaborative Contributions

One of the elements of the XPRIZE challenge that we really value is the opportunity to share knowledge and findings with the rest of participants via collaborative contributions. We believe in the power of collective intelligence and in the importance of open data/open science to inspire others and accelerate progress.

We have made the following eight contributions since the start of the XPRIZE: We shared (1) a .csv file with demographic information that we had collected for each country/region split in 5 buckets of ages [0-4 5-14 15-34 35-64 65+]; (2) information relative to the *COVID19impactsurvey*, including a visualization of all the responses to date; (3) a clarification relative to whether there was going to be an evaluation of the models with historic data; (4) a visualization of the historic NPIs and COVID-19 cases for all the regions; (5) an analysis of the communities of regions that are identified by analyzing the flights between them as a potentially useful metric to cluster regions and to better model the impact of NPIs; (6) the clusters of regions identified by our algorithm which have proven useful in the prediction phase; (7) a bug that we found in the standard predictor which made it ignore time periods from the past; and (8) a histogram of all possible combinations of NPIs.

9. Innovation

Our approach is innovative in several ways: (1) it combines different approaches to identify the set of 10 Pareto front optimal solutions; (2) it prioritizes interpretable models; (3) it is computationally efficient, generating up to 10 prescriptions for each region/country and for a 90-day period in less than 2 hours on a server/workstation and 2 h 40' 36'' on a MacBook Pro 2,7 GHz Intel Core i7 4 cores 16 GB 2133 MHz LPDDR3; (4) it leverages data coming from a citizen survey to inform policy makers in their decision making related to which COVID19 interventions to apply. In fact, the survey answers have been extensively used by the Valencian Government since it was launched on March 28th, 2020. They plan to use our prescriptor as a key tool to support their decision making in the months to come.

[3] https://ellisalicante.org/en/covid19impactsurvey

[4] http://infocoronavirus.gva.es/es/grup-de-ciencies-de-dades-del-covid-19-de-la-comunitat-valenciana

ANNEX I. NON-PHARMACEUTICAL INTERVENTION ANALYSIS

In the challenge, we studied the space of all possible combinations of the 12-dimensional NPI vectors. Table I.1 depicts each dimension of the vector and the possible values that it may adopt, where 0 denotes that such intervention is not applied. Each NPI vector has a stringency at cost one which is given by the sum of the value of each of its dimensions.

Dimension of the NPI vector	**Possible values**
C1 School closing	[0, 1, 2, 3]
C2 Workplace closing	[0,1,2,3]
C3 Cancel public events	[0,1,2]
C4 Restrictions on gatherings	[0,1,2,3,4]
C5 Close public transport	[0,1,2]
C6 Stay at home requirements	[0,1,2,3]
C7 Restrictions on internal movements	[0,1,2]
C8 International travel controls	[0,1,2,3,4]
H1 Public information campaigns	[0,1,2]
H2 Testing policy	[0,1,2,3]
H3 Contact tracing	[0,1,2]
H6 Facial coverings	[0,1,2,3,4]

Table I.1. Possible values of each of the 12 dimensions of the NPI vector

Given such 12-dimensional NPI vectors, there are 7,776,000 possible combinations, which, when grouped by stringency (at cost one), result in the histogram shown in Fig I.1.

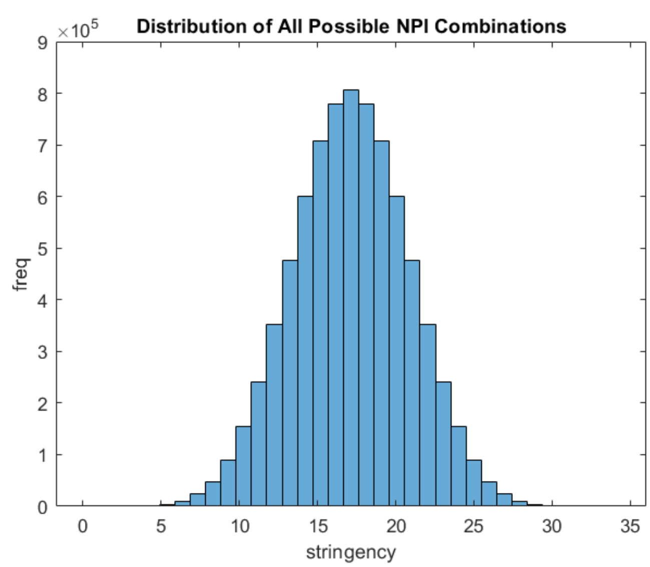

Figure I.1. Histogram of all possible NPI vectors based on their stringency at cost one

Ideally, we would compute the associated convergence Rn for each of these 7.78 million of possible NPI combinations. However, due to time and computation constraints, performing such a mapping was not feasible. Hence, we sampled the space of possible NPIs, such that we computed the convergence Rn for a subset of **54,652 NPI** combinations obtained as follows:

- All the NPI combinations with stringency [0 to 6] and [28 to 34].
- All the NPI combinations where the most important dimensions as per our feature importance analysis (see Annex III) were non-zero, namely: C1, C2, C4, C5, C8 and H2, with the rest of dimensions set to 0. These combinations range from stringency 0 to stringency 19.
- A random sample of the rest of the unexplored stringencies

We did our experiments in the following 21 representative countries/regions: United States, Brazil, India, Mexico, Italy, China, United Kingdom, France United Kingdom/England, Russia, Iran, Spain, Argentina, Colombia, United States/New York, Peru, Germany, Poland, South Africa, United States/Texas, and United States/California. In addition to the convergence Rn, we computed the total number of cases in 20 and 60 days.

The distribution of the explored combinations are shown in Figure I.2 and the details of the total number of NPI combinations for each value of stringency is summarized in Table I.2. We show both the number of possible combinations and the number of combinations for which we computed the convergence Rn.

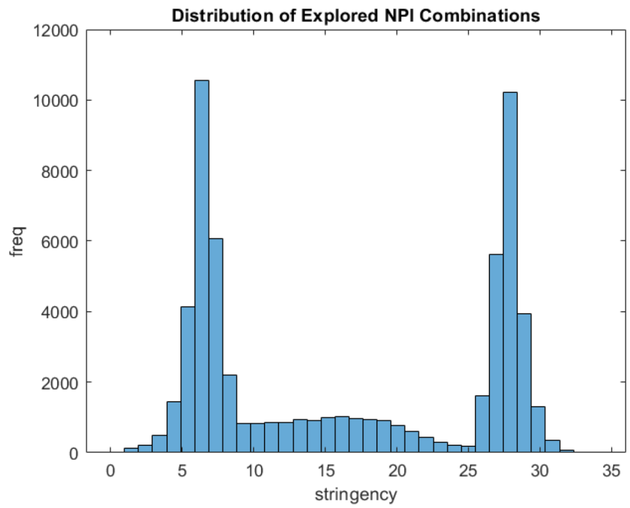

Figure I.2. Distribution of explored NPI combinations.

stringency	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
total	1	12	78	359	1,301	3,927	10,218	23,449	48,231	89,967	153,543	241,374
generated	1	12	78	359	1301	3,927	10,218	5,743	2,059	821	836	857
%	1	1	1	1	1	1	1	0.24	0.04	0.009	0.005	0.003
stringency	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**
total	351,326	475,347	599,643	706,818	779,659	805,494	779,659	706,818	599,643	475,347	351,326	241,374
generated	868	945	924	1,003	1,012	977	948	906	758	613	445	296
%	0.002	0.0019	0.0015	0.0014	0.0013	0.0012	0.0012	0.0012	0.0012	0.0012	0.0012	0.0012
stringency	**24**	**25**	**26**	**27**	**28**	**29**	**30**	**31**	**32**	**33**	**34**	
total	153,543	89,967	48,231	23,449	10,218	3,927	1,301	359	78	12	1	
generated	199	187	1,621	5,628	10,218	3,927	1,301	359	78	12	1	
%	0.0013	0.002	0.03	0.24	1	1	1	1	1	1	1	

Table I.1. Volume of existing NPI combinations and the explored ones

From these **54,652 NPI** combinations, we need to identify the subset that dominates the rest. Figure I.3 depicts each NPI combination as a point in the stringency-number of cases space. The color of each dot in the figure represents the expense of the NPI, computed as:

$$expense = Stringency \times \frac{total\ Cases - min(total\ Cases)}{max(total\ Cases) - min(total\ Cases)}$$

We see how the NPIs with the smallest expense are on the Pareto front.

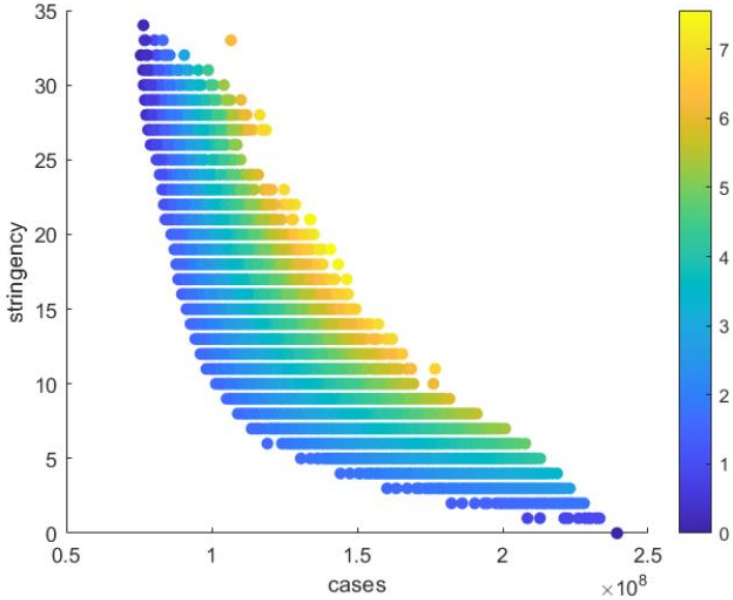

Figure I.3 54,652 NPI combinations with their associated stringencies (at cost 1), number of cases and expense.

ANNEX II. MAPPING BETWEEN NPIs AND CONVERGENCE Rn

As previously described, we discovered that the **same NPI would lead to the same convergence Rn when applied for sufficiently long time period (> 21 days),** independently of the country and time period of application. We refer to this finding as the **Rn synchronization principle.**

Based on this finding, we empirically computed the mapping between **54,652** NPI vectors and their convergence Rn, together with the resulting number of cases in 20 and 60 days for the 21 representative countries as described in Annex I.

To illustrate this finding, the figures below show a few examples of the convergence Rn for several values of the NPI vector.

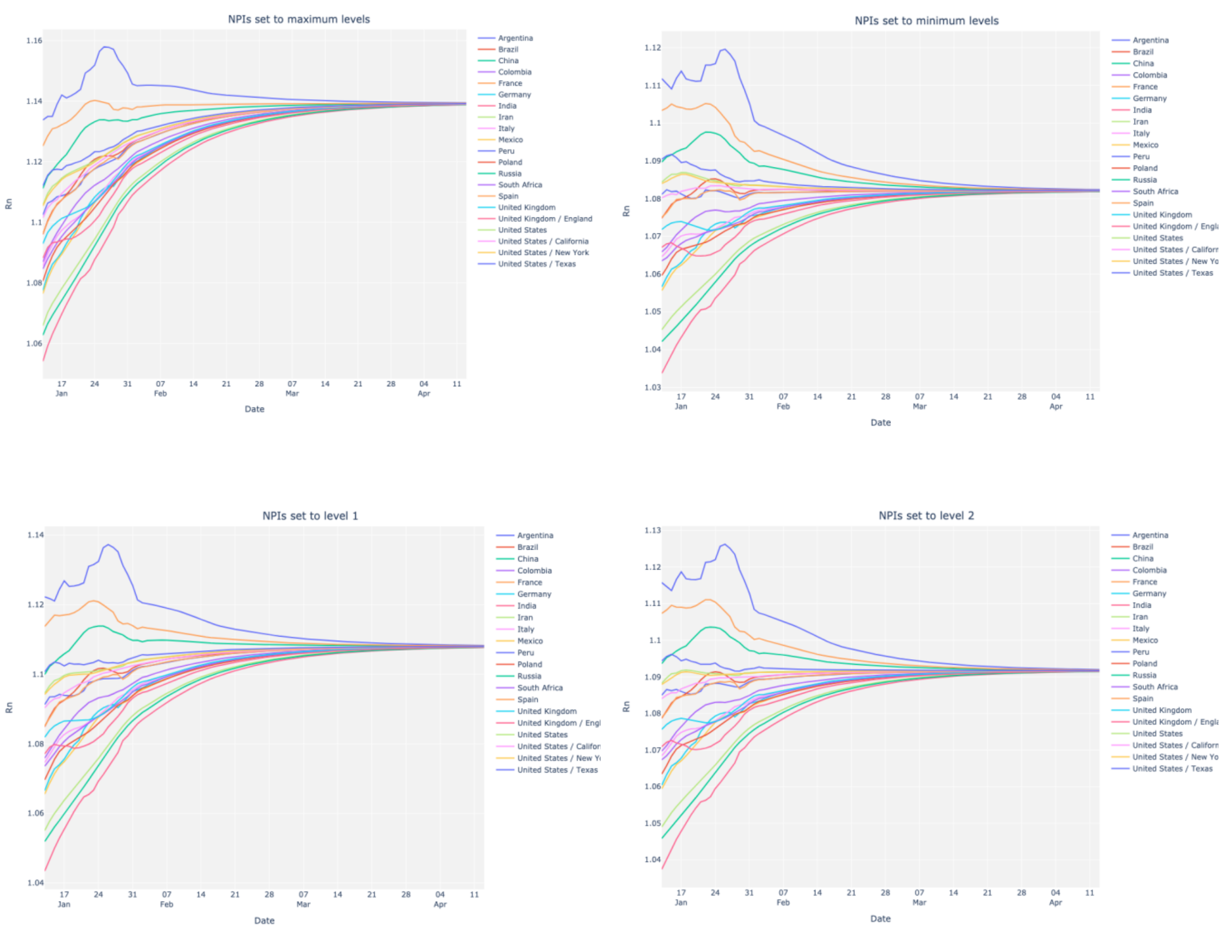

Figure II.1 Exemplary convergence Rn for 4 different combinations of the NPI vector.

ANNEX III. NPI FEATURE IMPORTANCE ANALYSIS

Given the previously described **synchronization principle** (convergence in all countries/regions to the same Rn under the same NPI when applied for long enough), we have now a measure that *quantitatively characterizes each NPI*. We use the set of **54,652 NPIs** for which we have their convergence Rn as ground truth to train a machine learning model to automatically infer the resulting convergence Rn for a given instance of the NPI vector. For this purpose, we trained both an MLP regressor and Gradient Boosted Trees. In both cases, we obtained a MAE of 3*10e-4 (with hyper-parameter optimization using the AX platform in the case of the MLP). Unfortunately, such MAE is too high given that we needed a precision in the estimation of the Rn of 10e-7.

Even though we could not use a machine learning model to compute the mapping between the NPI vector and the convergence Rn, we could perform a **feature importance analysis** to determine the impact that each dimension of the NPI vector has on determining the convergence Rn and hence the resulting number of cases. Figure III.1 depicts the feature and permutation importance of each of the 12 dimensions of the NPI vector. The feature importances for the [C1, C2, C3, C4, C5, C6, C7, H1, H2, H6, H8] dimensions are:
[**0.16253311 0.33043205** 0.02189676 0.06947346 0.05884782 0.03694285 0.0307605 0.07734699 0.05564001 0.10437437 0.0270645 0.02468759]

As seen in the vector above and the Figure below, C2 (workplace closing) followed by C1 (school closing) are the most important interventions to drive the number of cases up/down according to our model. The least important are C3 (cancel public events) and H8 (facial coverings), which makes intuitive sense, as they seem to be interventions that are implemented in all countries/regions and hence do not seem to provide differential value in making the number of cases increase or decrease.

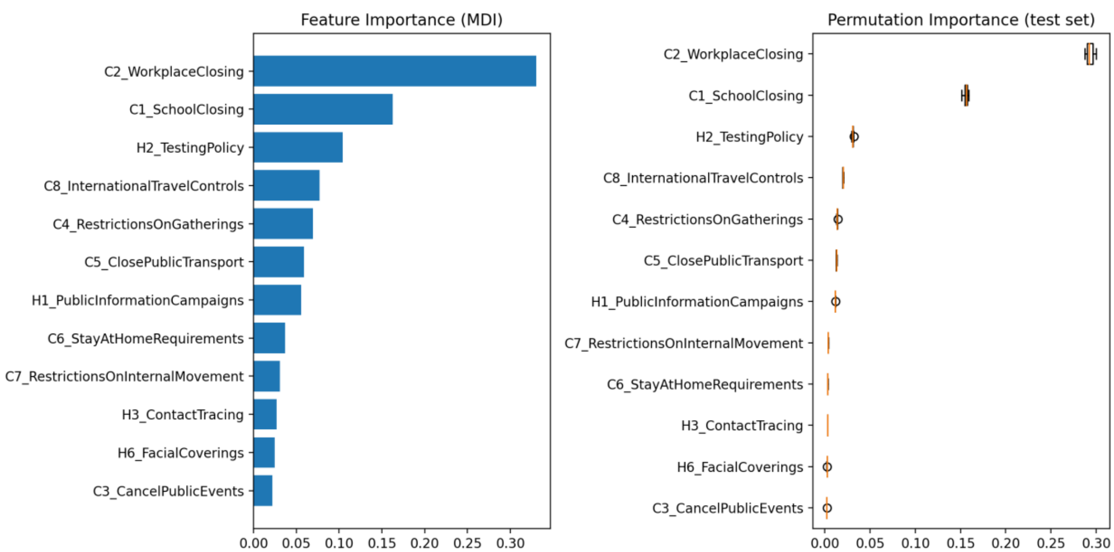

Figure III.1 Feature importance analysis using Gradient Boosted Trees.

We leverage this feature importance analysis in a greedy prescription algorithm which selects the top 10 prescriptions by activating to the maximum each dimension of the NPI vector in order of priority, given by: priority = feature_importance / cost.

ANNEX IV. DYNAMIC NPI POLICY OPTIMIZATION

Finally, we need to identify a dynamic regime for applying the selected NPIs over the time 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.

Figure IV.1 depicts an example of such Pareto front computed from an initial set of 19 candidate NPIs (the NPIs that are on the Pareto front as per the analysis previously described). In the Figure, we use temporal chunks of 21 days before changing NPIs and compute the NPI dynamic policy for a period of 60 days. There are 6,859 possible dynamic regimes of applying the NPIs. Of those, 225 dominate the rest and constitute a Pareto front of the dynamic NPI policy selection. The color of each dot represents again the expense of such dynamic NPI combination.

From such Pareto front, we select the 10 prescriptions 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 from a practical perspective.

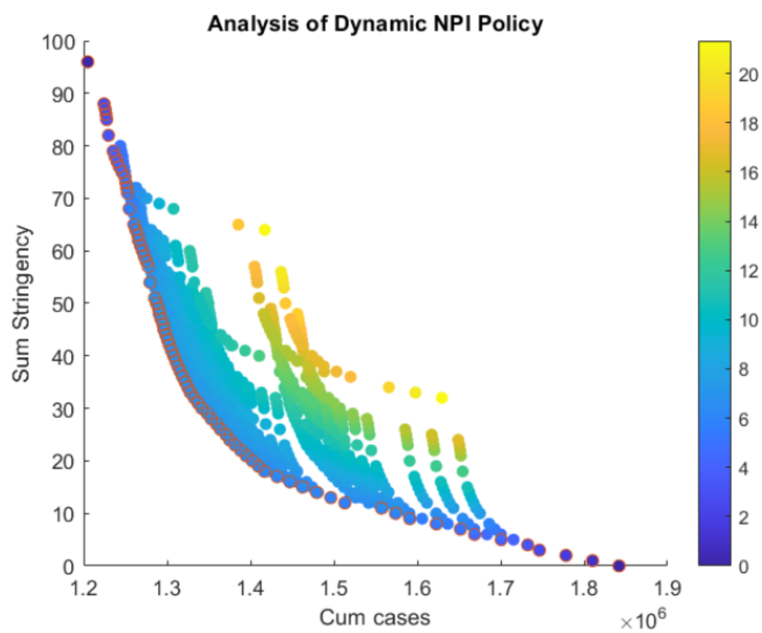

Figure IV.1 . Example of dynamic NPI policy optimization: from 6,859 possible NPI regimes, 225 are on the Pareto front.

ANNEX V. INTERACTIVE PRESCRIPTION VISUALIZATION

We have developed an interactive visualization of our prescriptor using Tableau, to enable its use by policy makers in the Valencian Government and other relevant stakeholders. The visualization can be found here: https://public.tableau.com/profile/kristina.p8284#!/vizhome/Prescriptions_16117279637400/Visualize.

Figures V.1 and V.2 show the two windows in our visualization. The main window is called "Prescriptors" and enables users to select the country/region of interest before it shows the 10 recommended prescriptions with their associated stringencies, total number of predicted COVID19 cases and levels of activation of each of the 12 dimensions of the NPI vector. It also shows the Pareto front of all prescriptions. The secondary window, called "Compare prescriptions" enables users to select 2 prescriptions to compare.

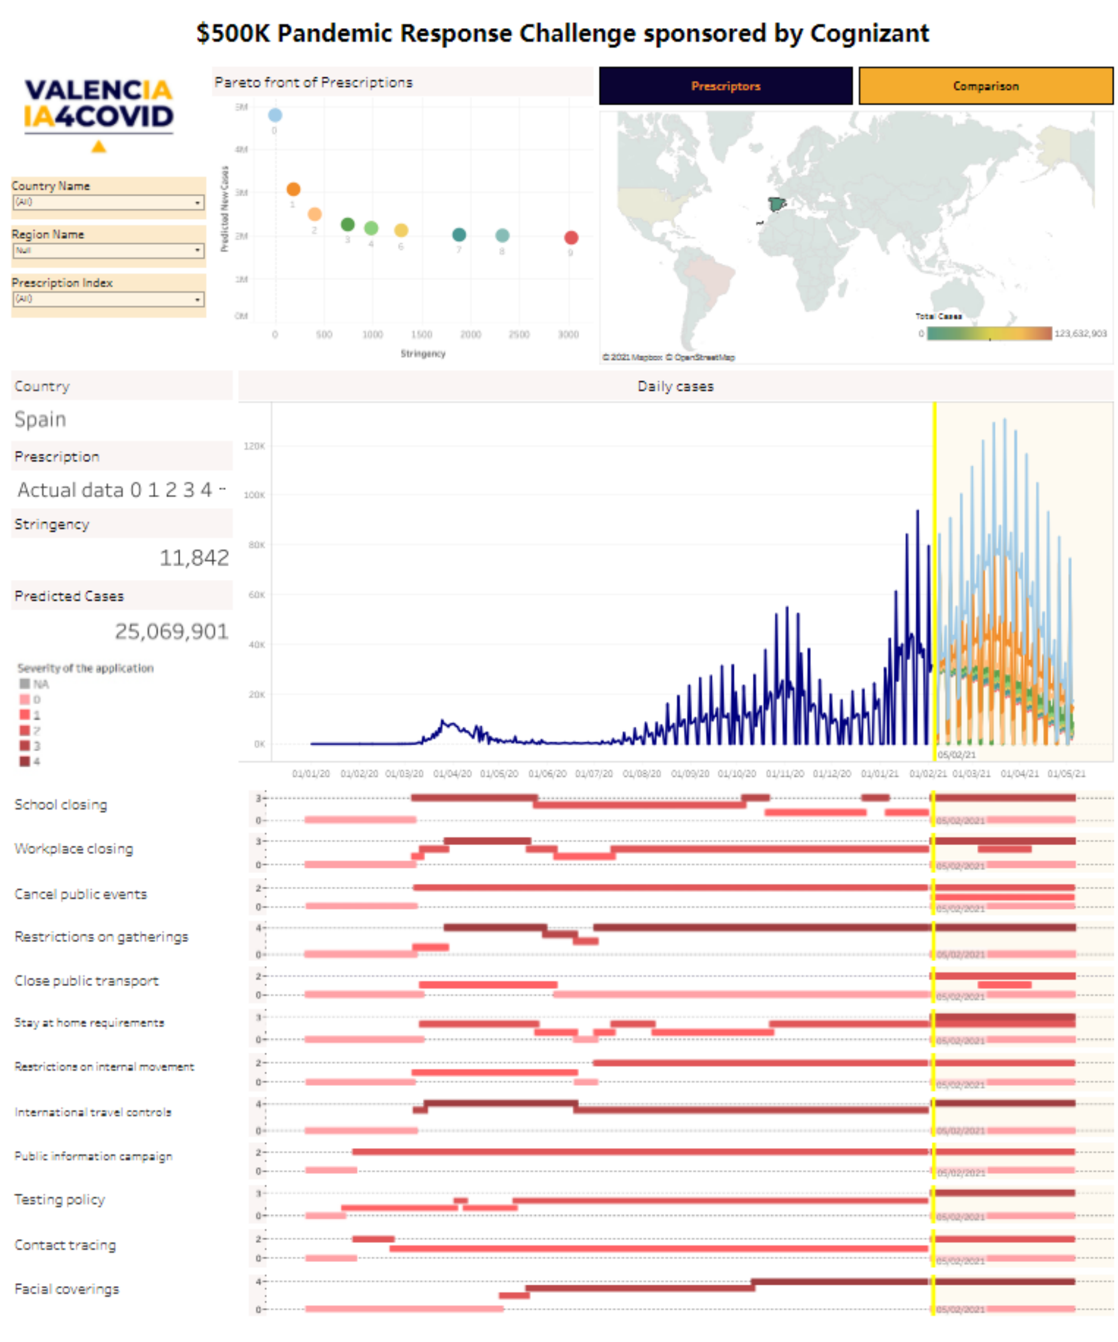

Figure V.1 . Prescriptor visualization interface. In this example, the system shows the 10 selected prescriptions for Spain

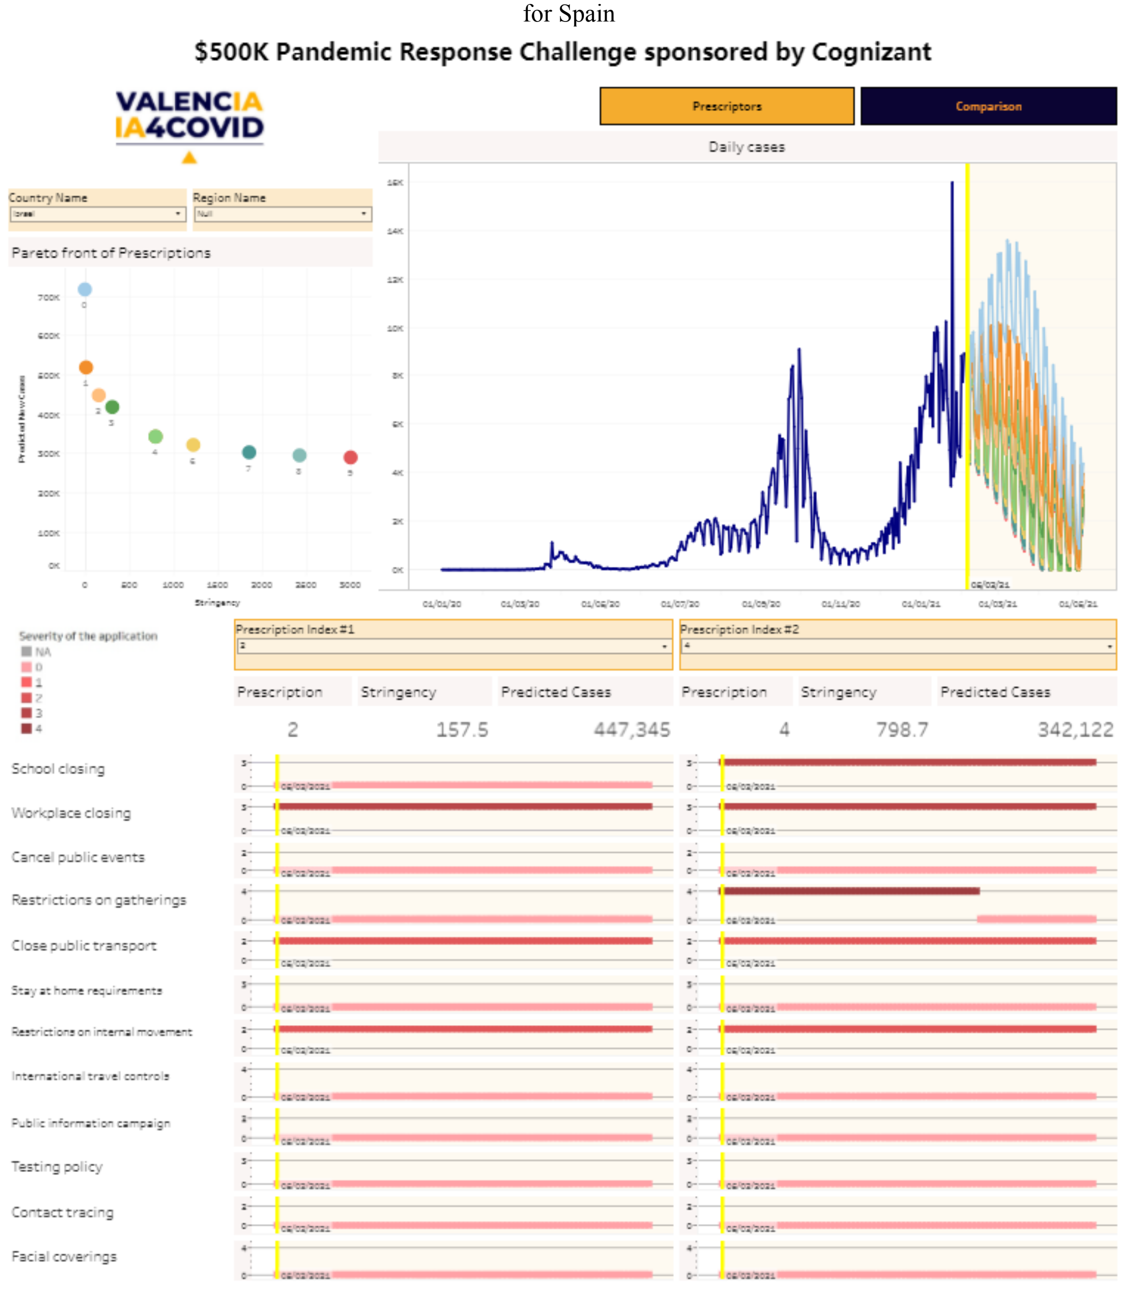

Figure V.1 . Prescriptor comparison interface. In this example, the system shows side to side two prescriptions from the 10 prescriptions proposed for Israel. The prescription on the left (#2) has an overall stringency of 157.5 and would result in 447,345 cases. The prescription on the right (#4) has an overall stringency of 798.7 and would result in 342,122 cases.