Team Metis2020

Ridwan Alam, Vanessa Hu, Ryan Lewis, Ramon Martin,
Neda Saleem, Brian Tam, Nick Wilders, Andrew Zhou

INTRODUCTION

This document serves as the qualitative description and analysis for the submission to XPrize's COVID Response Challenge Prescription Round by Team Metis2020. Because COVID-19 is a global issue like no other that affects every single individual on the planet, our team believes that the qualitative analysis and discussion of the humane aspects of this model are of utmost importance.

Here we provide an overview of our model, highlighting its strengths and describing areas with growth opportunities. We will also describe our diverse, qualified team along with our unique collaborative process, including our recent contributions to the larger conversation about COVID-19, the vaccine, and how data can be a vital tool in solving this problem.

ACTIONABILITY / USEABILITY

Our prescriptor distinguishes itself with its unparalleled customizability. Rather than adopt a one-size-fits-all approach, we train a separate neural network from scratch for each query, ensuring that each region's idiosyncrasies are fully accounted for. This approach would seem to pose scalability and practicality issues, but we leverage the power of TensorFlow static compilation to construct the network architecture only once, then reset its weights (without recompiling the graph) and reuse the graph for each subsequent problem.

Moreover, the overarching framework of our approach is extensible to a general family of predictors, beyond simply the standard predictor provided by the Pandemic Response Challenge Organizers. Any predictor whose gradient with respect to the NPIs can be feasibly calculated or estimated can be substituted into our model. A simple example would be a differentiable predictor. Fully black-box predictors are not supported, but it should be emphasized that virtually any useful predictor will have some discernible topography that can be exploited.

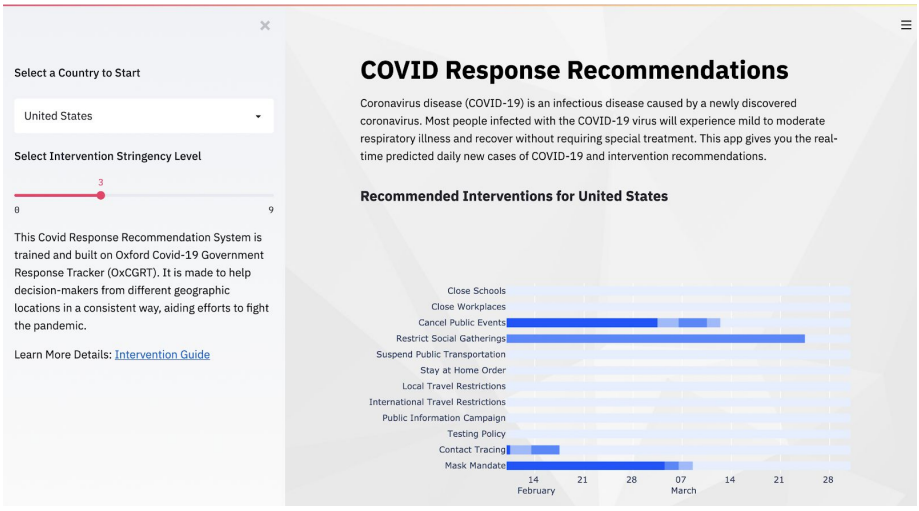

We built a website www.StopCovidNow.co to help decision-makers from different geographic locations in a consistent way, aiding efforts to fight the pandemic.

ADDRESSING THE CHALLENGE

One of the strengths of our predictor model was addressing the individualized regional qualities that lead to a full and effective prediction, utilizing a weighted MAE that takes the population of a given region heavily into account. This approach allowed for more accurate scoring for smaller countries where the penalty is weighted against fair and accurate predictions. We value realistic representation of this problem at all stages of the competition.

We have continued this emphasis on accurate and fair predictions by creating neural networks that can be quickly trained and customized for each region. This process, though it may appear time-consuming at the outset, is pivotal to create prescriptions that fully address the challenge at hand. Much of our model development, in fact, was geared toward leveraging the computational efficiency of the TensorFlow library to produce prescriptions in a practical amount of time.

GENERALITY & CONSISTENCY

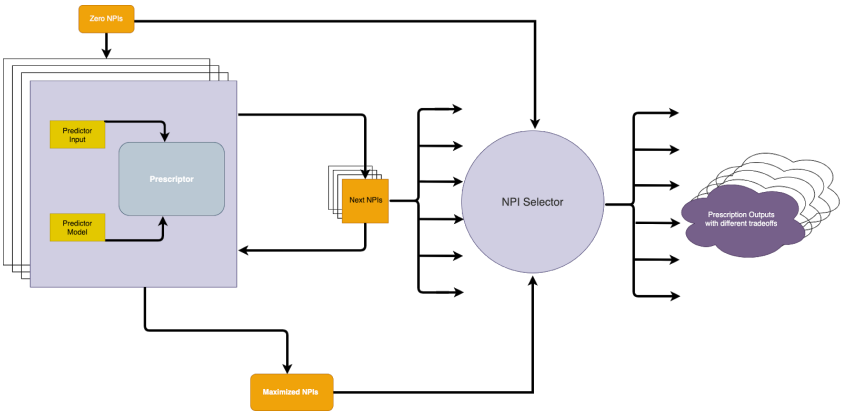

Because our model relies on case development and inputted parameters within each country, and the nuanced situation that each country faces in its COVID-19 outbreak, it is optimized to ensure an accurate prediction specific to the region in question. This generality comes from an unbiased neural network trained for each individual region that explicitly accounts for that region's on-the-ground situation.

The model can also take predictions from any predictor that may replace the provided LSTM in the future. The predictor need only adhere to the following structure: given a region, a period of time, and the NPIs from the present day up to the end of that period, output the total number of new cases in that region over the period. This structure is, in fact, a slightly relaxed version of the Round 1 Predictor API, which additionally requires case predictions for each day. See the below *Innovation* section for an explanation of our customized gradient descent algorithm.

TRANSPARENCY AND TRUST

Recognizing that the model's mathematical abstractions may not be helpful to the layperson or non-technical policymaker, we have developed a responsive web application that allows users to explore our prescriptions in real-time. For each region, we provide a set of up to ten prescriptions for a specific timeframe, with each prescription optimized for a different caseload-stringency tradeoff, and graphically display their associated caseloads and stringencies. The user may inspect these prescriptions for different areas, modify the costs of different interventions to yield a different set of prescriptions, and explore the differences between the differently-optimized prescriptions for a given region.

Additionally, our prescriptions are accumulative, or tiered: each more stringent prescription in a set includes or intensifies all interventions from prescriptions in lower tiers. Thus users can see in what order interventions will be implemented, giving them insight into the nature of our process that mathematical details might not provide.

COLLABORATIVE CONTRIBUTIONS

Our code is also designed to be interpreted by and available on an open-source platform. As described above, our general framework is adaptable to take in output from alternative predictive models, allowing continual updates to be easily and accurately updated. Additionally, our GitHub repository and Tableau Dashboards (1 | 2) are publicly accessible. We used various open-source modules, libraries, and tools, such as Python, Plotly, NumPy, TensorFlow, Keras, and Streamlit.

We continue to stay involved with larger conversations concerning a data scientist's role in the global pandemic; our team leader, Vanessa Hu, served as a panelist on January 22, 2021, in an open discussion with Moderna's CEO, Mr. Stéphane Bancel, regarding the mRNA Covid-19 vaccine, business model transformation, and other efforts to fight the pandemic.

INNOVATION

Our innovative approach is highlighted above, but it bears repeating that the generalized blank neural network, ready to factor in the predictions from any model (be it the provided LSTM or another model developed in the future), allows for quick and effective updates to provide predictions and treatment recommendations that could save lives.

The crucial innovation in our prescriptor is its customized use of gradient descent. The essential iterative loop of our model is as follows: begin with a prescription of across-the-board zero NPIs, use gradient descent to identify the most helpful NPIs, as weighted by NPI cost, and update the prescription by incrementing those NPIs by one. The number of incremented NPIs per iteration is an adjustable parameter that controls how finely graduated the process is; a lower number will result in more precise prescriptions but will also increase the running time of the algorithm. Once our current prescription reaches the maximum possible set of NPIs, we end our algorithm and select a set of prescriptions from the full set.

Our decision to eschew floating-point arithmetic and snap all NPIs to their integer values is a key design choice for our algorithm. Though gradient descent is a floating-point operation by definition, we use the gradient simply to identify the most promising NPI(s) in each iteration, which we increment by a fixed amount that is not necessarily proportional to the gradient magnitude. Such an approach allows for great algorithmic transparency while also yielding an accumulative set of prescriptions: each builds on the last, so a layperson can see the consequences of increasing stringency in an interpretable way.

We also select our set of prescriptions with an innovative decision process. We first note that caseload as a metric for the success of a prescriptor is useless in a vacuum: the specific number of cases must be applied in reference to some standard. Thus, we precompute the maximum and a minimum number of cases for the given time period by forecasting cases for a prescription with zero NPIs and a prescription with maximum NPIs across the board.

The performance of a given prescriptor, then, is the number of cases it prevents over the maximum number of preventable cases, or the proportion of cases solved. We use this metric to select a set of prescriptions that is suitably spaced out among this axis, so we avoid selecting solutions that correspond to very similar caseload-stringency tradeoffs. We might fall into this pitfall if we used stringency to select prescriptors, as the relationship between prescriptors' stringency and caseloads is unpredictable. For instance, prescriptors with 10 stringencies and 30 stringencies might solve 95% and 96% of the total reducible cases, in which case they would represent very similar tradeoff balances despite appearing to be very different.

It should be noted that this approach relies on the assumption that zero NPIs leads to the maximum case scenario and maximum NPIs leads to the minimum case scenario. Though this assumption may not be strictly true for every predictor, it should nonetheless be approximately correct based on our contextual knowledge.