

A Contact Network-Based Approach for Online Planning of Containment Measures for COVID-19

Guilherme Y. Thomaz, Denis D. Mauá, Leliane N. de Barros

¹Instituto de Matemática e Estatística, Universidade de São Paulo

gui.yambanis@gmail.com, ddm@ime.usp.br, leliane@ime.usp.br

***Abstract.** We use data from the 2017 Origin-Destination survey to build a representative contact network for the city of São Paulo, where individuals are connected by different social relations (school, work, neighborhood). The network is used to devise a stochastic discrete time and state compartmental model for the spread of the COVID-19. We employed the model to compare different mitigation strategies. The results show that even simple Monte Carlo planners greatly improve the performance over reactive strategies in terms of balancing the economical and health impacts of non-pharmaceutical interventions.*

1. Introduction

The ongoing COVID-19 disease pandemics caused by the SARS-CoV-2 virus has brought immense health, social and economic costs to societies around the globe. Worldwide local and national government responses to the disease have varied considerably, but generally consisted on applying non-pharmaceutical intervention measures such as traffic and travel restrictions, business and school closures, and even forced quarantine and isolation. The evolution and consequences of the disease have also been varied. While some countries have managed to maintain relatively low numbers of infections and deaths, other countries have experienced a collapse of health resources and a consequently relatively larger number of infections and deaths. At the heart of these different strategies lies the debate about the trade-offs between economic and health costs, as well as local society idiosyncrasies such as privacy and individual freedom concerns and public trust of authorities. Yet, there is a large consensus that preferred public policies should prioritize public health while minimizing economic impact. This is by far not an easy goal to meet, as it requires analyzing the long-term consequences of actions as well as the severe uncertainty involved.

The goal of this work is to study how automated planning techniques can contribute to the design of public mitigation and containment policies for COVID-19 at the city level, taking into account the interplay of different factors such as mobility and age-distribution patterns. We thus formulate the problem of control plan design as a finite-horizon discrete Markov Decision Process (MDP) with a network-based evolution model of the disease [Newman 2002, Xue 2020]. To make things more realistic, we use data from the 2017 São Paulo City Origin-Destination survey to extract a contact network that models the possible spread of the disease through work, school, shopping and transit relations, as well as public available data (number of hospital beds, disease transmission characteristics, etc). Unlike previous approaches based on probabilistic planning that considered only measures such as isolation and vaccination [Xue 2020, Kinathil et al. 2017], we consider control policies currently adopted by local authorities such as raising public awareness, banning large public gatherings, closing schools and business, and imposing

traffic restrictions. Our results show that optimized strategies deliver a much better compromise between health and economic impact.

2. Related Work

The number of papers modelling and reporting on the ongoing COVID-19 epidemics has been staggering. In a relatively short period, researchers have been able to identify key properties of the disease (infectiousness rate, fatality rate, risk factors, reproduction number, etc), in spite of the shortcomings of available data (under-reporting, asymptomatic individuals, etc). There has also been a sheer number of modelling proposals (many from non-experts/newcomers in the epidemiology field, such as the case of this work). Most of these proposals build on the well-known SIR model [Kermack and McKendrick 1927], a system of ordinary differential equations that describes the spread of the disease in terms of “population flows” among compartments (susceptible, infected, removed). While fairly well understood, the SIR model falls short of providing an accurate description due to assuming that individuals interact at random (with a so-called uniform mixing rate) and to not accounting for the effect of risk factors [Keeling and Eames 2005].

Current approaches differ in how they account for such non-heterogeneity by incorporating the impact of inter- and intra-city mobility, age, income, etc, and in how uncertainty about the disease characteristics is handled [Dureau et al. 2013, Kucharski et al. 2020]. Mellan et al. fitted a sophisticated stochastic transmission model from official data of COVID-19 in Brazil and from Google mobility trend data [Mellan et al. 2020]. The model accounts for mobility trends as well as age and income-related risk factors. They estimated that in the state of São Paulo the reproduction number went from around 3 at the onset of the outbreak on February 16 to roughly 1.5 on May 6th, following the enforcement of social distance measures by authorities. Silva et al. used a spatial SEIR model that accounts for the daily commute among main cities of the state of São Paulo, and assumed, as we do here, that the reproduction number of each city can be controlled at each time step. The control problem is then cast as a continuous constrained optimization. The main objective is to alternate the restrictions among different cities, bringing the disease spread to a desired level while minimizing economic impacts [Silva et al. 2020].

One interesting approach to building more realistic disease transmission models is to use contact networks (a.k.a. diffusion networks and mixing networks), which capture the interaction patterns at the individual level [Newman 2002, Keeling and Eames 2005, Bai et al. 2007]. Halloran et al. used contact networks to simulate the spread of airborne diseases among a small population [Halloran et al. 2002]. They built a network by randomly assigning individuals to age groups and households according to the US population distribution. Notably, children were assigned a school or day-care center where they closely interact with many other children, creating the main source of spread of the disease. This is supported by evidence that households and schools are the main routes of transmissions of several diseases such as smallpox and measles.

Regarding the COVID-19 disease, Zhang et al. collected self-reported contact patterns before and after the outbreak from a survey with 636 inhabitants of the city of Wuhan, the initial epicenter of the outbreak, and with 557 inhabitants of Shanghai, one of the largest and densely populated city in China [Zhang et al. 2020]. The differences

in contact patterns at the two dates allowed researchers to measure the effect of non-pharmaceutical interventions in decreasing transmission through social contact. In short, the study observed that the mean number of contacts reduced from about 14 (in the baseline period) to 2 for the outbreak period. The latter approximates the mean number of household members in the region, reflecting the severity of the interventions (which included case-based individual isolation and contact tracing).

As mentioned in the introduction, there have been approaches for using planning techniques for designing control policies. Kinathil et al. used symbolic dynamic programming to optimally solve for an offline policy of a discrete-time SIR model with continuous state and actions [Kinathil et al. 2017]. The work that most closely resembles our approach is by Xue, who used a contact network-based compartmental transmission model to prescribe node-level interventions (viz. vaccination and isolation) in a “online” fashion [Xue 2020]. They approached the planning problem by a novel and sophisticated mixed-linear integer programming reformulation.

3. A Network-Based COVID-19 Evolution Model for the City of São Paulo

We now describe the main ingredients of our planning approach: a contact network based model of disease transmission, the action/interventions and the cost model.

3.1. A Contact Network for the City of São Paulo

We use data from the 2017 São Paulo Origin-Destination survey to extract a contact network [Newman 2002, Keeling and Eames 2005] for the city of São Paulo. That survey has been carried out every ten years since 1967 to support the implementation of public policies in the metropolitan area of the city of São Paulo.¹ The 2017 edition collected data related to mobility trends, including the place of residence, work and school/university, as well as the type of occupation, age, and income of level of nearly 87 thousand individuals and 32 thousand households. This provides us with a very representative picture of the interaction patterns of inhabitants of the city of São Paulo that combines spatial, demographic and regional information. The right plot in Figure 1 shows the distribution of the locations of surveyed individuals’ households in each city district.

We use the data to obtain an attributed *multigraph* $G = (V, E)$ where V denotes the set of surveyed individuals v_1, \dots, v_N and E contains triples $(r, \{v_i, v_j\})$ representing relations of type r between individuals v_i and v_j .² Each node is associated with available personal information such as age and type of work, although the latter is currently ignored by our model. We consider four types of relations between individuals (available from the survey data): *school*, *work*, *neighborhood* and *household*, denoting, respectively, interactions that take place when two individuals go to the same school/university, workplace, local shops and share a household. We obtained the edges in E as follows. The household relations were obtained by connecting individuals with the same household id in the data set. For the other types of relation, we used a k-means clustering algorithm to group individuals by proximity of the corresponding relation (e.g. to group individuals that study in a close location); we then connected every pair of individuals assigned to the same cluster. We set the parameters of the clustering algorithm differently for each relation, looking to

¹The raw data and a summary of results of the OD 2017 survey are available at <https://bit.ly/38viIse>.

²Alternatively, one can think of G as a collection of graphs $G_r = (V, E_r)$, one for each relation r .




Figure 1. Left: 2D Histogram of relations by district of individual's household. Right: Number of individuals with households in each district.

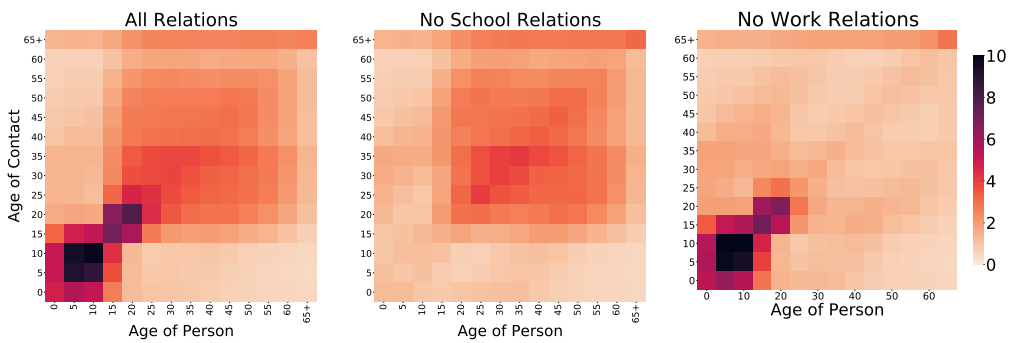


Figure 2. Age-related contact relation histograms for the network.

obtain realistic groups. Classmates clusters (i.e., of the school relation type) were further subdivided using the individual reported education level, so as to avoid unrealistic largely clusters at large educational facilities (which contain students from several levels). We initialized cluster centroids according to the zone boundaries (which are roughly subdivisions of the districts shown in Figure 1), and set their number so that the mean size μ of the resulting clusters satisfied $\mu_{\text{household}} < \mu_{\text{neighborhood}} < \mu_{\text{work}} < \mu_{\text{school}}$. The resulting graph consists of 55492 nodes/individuals, of which 27001 have jobs and 12605 are students. The household relation has 61634 edges; the neighbor relation has 348835 edges; the work relation has 342525 edges, and the school relation has 209044 edges. The 2D histogram on the left in Figure 1 shows the number of contact relations (i.e., no. of edges) among surveyed residents in two city districts. We see that contacts are more frequent within individuals of the same district, and there are several districts well connected to each other. The distribution of contact relations segmented by age and type of relations are shown in Figure 2. Notably, these histograms are very similar to the ones observed by Zhang et al. [Zhang et al. 2020] in what regards the distribution of contacts by age, as well as the effect of workplace and school closures. The high symmetry indicates the importance of considering age in controlling the spread.

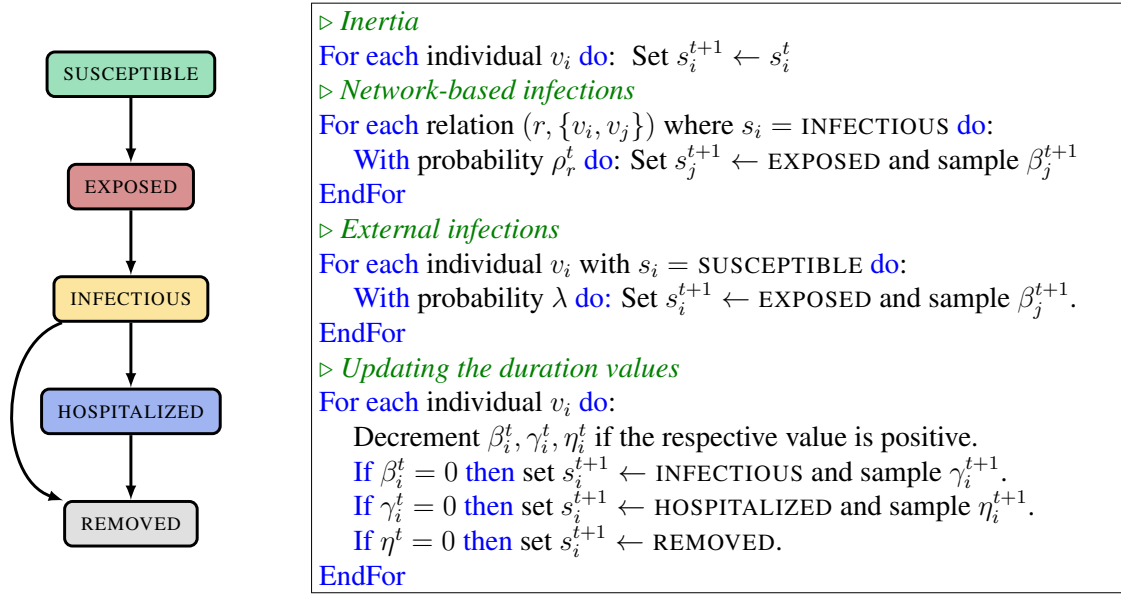


Figure 3. Simulating the COVID-19 disease evolution.

3.2. Disease Transmission Model

We use a stochastic compartmental epidemic model where each individual v_i is associated with a variable s_i^t that takes on the values SUSCEPTIBLE, EXPOSED (i.e., infected but not yet contagious), INFECTIOUS (infected and contagious), HOSPITALIZED and REMOVED (dead, or recovered and immune to the disease), corresponding to a stage of the disease at time t , as illustrated in the diagram in Figure 3 (self-loops are omitted for clarity).

The initial individual states s_i^0 are randomly sampled such that in certain districts (Morumbi, Vila Mariana and Jardins) 10 random individuals begin at the EXPOSED state, and the remaining individuals (in those districts and in all other districts) begin as SUSCEPTIBLE. We selected the districts with initial infections based on official data.³ The variable $\rho_r^t \in [0, 1]$ denotes the *infectiousness* of a relation r at step t ; If $\rho_r^t = 0$ then the relation is no longer active (i.e., it does not influence on the spread of the disease) and $\rho_r^t = 1$ denotes that any contact between an INFECTIOUS individual and a SUSCEPTIBLE individual results in the latter becoming EXPOSED. As we discuss later, the control actions (activation and deactivation of relations) modify the infectiousness values of the relations according to a predefined set of possible configurations. We assume the infectiousness of the household relation cannot be modified, and remains constant at $\rho_{\text{household}} = 0.06$. That value was chosen according to recent estimates [Wang et al. 2020, Li et al. 2020]. The infectiousness of the remaining relations are specified such that the basic reproduction number in our simulations approximates the estimates from the literature [Mellan et al. 2020]. This implies in having $\rho_{\text{work}} = 0.1\rho_{\text{household}}$, $\rho_{\text{school}} = 0.15\rho_{\text{household}}$, $\rho_{\text{neighbor}} = 0.1\rho_{\text{household}}$. There is also a “leak infection probability” $\lambda = 5e-5$ that models roughly 2 susceptible individuals per day getting infected “spontaneously”, i.e. irrespective of the states of their neighbors. This accounts for external infection sources such as infected individuals moving into the city. Thus, a SUSCEPTIBLE individual v_i is infected according to a Noisy-Or distribution where each INFECTIOUS neighbour independently “fails” to spread

³We selected the 3 districts with higher case incidence in the first official report: <https://bit.ly/3e0tvvH>.

Age (years)	Probability of hospitalization
0 to 9	0.1%
10 to 19	0.3%
20 to 29	1.2%
30 to 39	3.2%
40 to 49	4.9%
50 to 59	10.2%
60 to 69	16.6%
70 to 79	24.3%
80+	27.3%

Parameters	Mean	Std. Dev.	Support
Incubation period (β)	5.6	2.8	[2, 15]
Infectious period (γ)	6.2	4.3	[2, 22]
Hospitalization length (η)	8.6	6.7	[2, 33]

Table 1. Left: Probability of an infectious individual requiring hospitalization. Right: Parameters of the truncated normal distributions (in days).

the disease to that individual with probability $1 - \rho_r^t$:

$$\Pr(s_i^{t+1} = \text{EXPOSED} \mid s_i^t = \text{SUSCEPTIBLE}, \mathbf{s}^t, \rho^t) = 1 - (1 - \lambda) \prod_{\substack{(r, \{v_i, v_j\}) \in E, \\ s_j = \text{INFECTIOUS}}} (1 - \rho_r^t),$$

where ρ^t is the vector of relation infectiousness, and the product runs over the incident edges on v_i whose neighbor $v_j = \text{INFECTIOUS}$. When an individual v_i is infected (i.e., changes from **SUSCEPTIBLE** to **EXPOSED**), an *incubation period* β_i^t is sampled from a truncated normal distribution and rounded down (we assume the variable takes on some arbitrary negative value before the individual is **EXPOSED**). The variable is decremented in the subsequent steps (as long as it is positive). If $\beta_i^t = 0$, the individual v_i transits to the **INFECTIOUS** state, and an *infectious period* γ_i^t is then sampled from another truncated normal distribution (and rounded down). The variable γ_i^t is also decremented at each time step, and when it reaches zero, the **INFECTIOUS** individual becomes **HOSPITALIZED** with the age-dependent probability given in Table 1(left), extracted from the estimates in [Ferguson et al. 2020]. If the individual becomes **HOSPITALIZED**, a value η_i^t is sampled from a truncated normal distribution determining the length of hospital stay. If the individual does not become **HOSPITALIZED**, we set $\eta_i^t = 0$. In either case, when $\eta_i^t = 0$, the individual becomes **REMOVED**. The mean and standard deviations of the normal distributions for the initial values of β_i^t , γ_i^t and η_i^t were selected to reflect the estimates in [Linton and et al. 2020], and are given in Table 1(right). The corresponding random variables are truncated below at two days, and above by the 99th percentile value according to the normal distribution, so as to avoid non-realistic scenarios. The entire procedure is described by the pseudo-code in Figure 3. Note that the pseudo-code describes one time step (one day) of evolution of the disease.

Figure 4 (left) shows the pointwise average epidemiological curve of our model for 10 simulations (note the different scales for exposed and hospitalized proportions). Notice the delay between the number of exposed and hospitalized individuals that is also observed in real-data, a behaviour that makes planning for intervention measures challenging. Figure 5 shows the number of exposed and/or infected individuals in each district at different time steps. The epidemic progresses from the more central districts which were seeded with the initial infections, towards the outskirts.

The number of available beds in the city is around 0.25 per 100 inhabitants, ac-

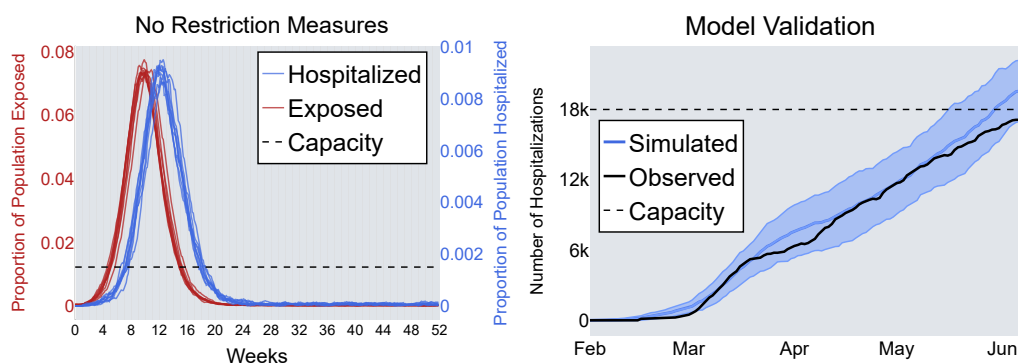


Figure 4. Left: Simulated epidemiological curve of the unrestricted model. Right: Observed and simulated number of hospitalizations.

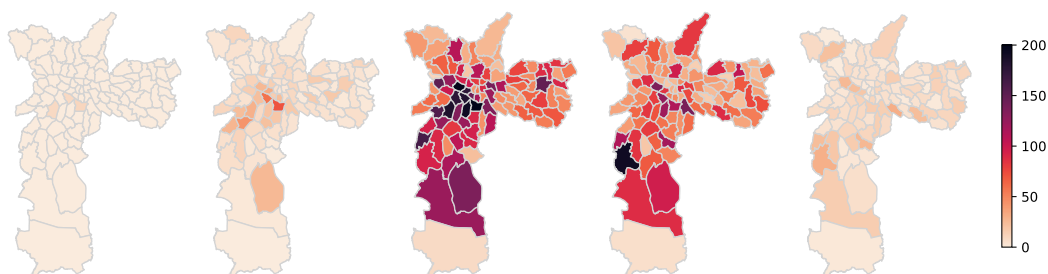


Figure 5. Number of infected individuals per district after 0, 4, 8, 12 and 18 weeks.

ording to official data.⁴ At any time, a large proportion of those beds are occupied by cases not related to COVID. On the other hand, new hospital beds are usually created when in times of epidemics, and social distancing measures often lead to a decrease of other causes of hospitalizations (e.g. car accidents).⁵ As a rule-of-thumb, we consider the number of hospital beds available to COVID patients to be 60% of the total capacity (about 0.15 per 100 inhabitants). This value is indicated by the dashed line in Figure 4 (and in the following similar graphics). One can observe that the number of hospitalizations largely surpasses the capacity when no interventions are enforced.

To ascertain the validity of the proposed model, we simulated the disease using real isolation indices extracted from cell phone data.⁶ The right plot in Figure 4 (right) compares the observed number of hospitalizations due to SARS (Severe Acute Respiratory Syndrome) in the given period (the black line), with the mean number and standard deviation (blue line/strip) of hospitalized individuals in 10 simulations of our model, with initial conditions similar to the real observations.⁷

⁴This was previous to the onset of the outbreak, according to <https://bit.ly/2AAx0vc>.

⁵According to official data, the number of ICU beds for COVID patients in the city of São Paulo has doubled since the onset of the outbreak (see <https://bit.ly/2NYMqMQ>), while the number of traffic accidents has reduced 30%, according to local news (viz. <https://glo.bo/3f2J2MO>).

⁶At each day, we scale the value of infectiousness of each relation by the respective reduction as reported by the Google Mobility Trends Report (<https://bit.ly/2NX40kj>). For the school relations, we used the baseline values during weekdays and reduced them during the weekends prior to March 23th, when classes were suspended (<https://glo.bo/2VQ0IJg>). Afterwards, the infectiousness of the relation goes to zero.

⁷We used the estimates by the Observatório COVID-19 (<https://bit.ly/2O0CrX6>) as observations.

3.3. Interventions

As we commented earlier, we assume the decision maker acts by means of modifying the infectiousness of the relations ρ_r^t using population-level interventions such as encouraging social distancing, raising public awareness, closing schools, universities or workplaces, restricting city traffic, banning public gatherings, or imposing city-wide lockdown. These actions take place immediately and last for a fixed number of days (one week in our experiments). We assume that those actions are applied to all individuals, indifferently, and do not modify the infectiousness of the household. For the remaining relations, we consider the following five interventions/actions, selected to simulate the possible mitigation measures announced by the city authorities:

- **Unrestricted:** The infectiousness of all relations are kept at their initial values.
- **Social distancing:** We scale down by 20% the infectiousness of the work, school and neighbourhood relations with respect to their unrestricted values. This simulates raising public awareness and banning large public gatherings, while leaving schools and retail/services open.
- **Light Quarantine:** We scale down by 40% the infectiousness of work and neighborhood relations, and set the infectiousness of the school relations to zero. This simulates closing school, universities and retail/services, banning large public gatherings and raising public awareness.
- **Hard Quarantine:** We scale down by 70% the infectiousness of the work and neighborhood relations, and set the infectiousness of the school relation to zero. This is a stricter version of the previous intervention where only essential workers commute, and city traffic is constrained.
- **Lockdown:** We scale down infectiousness of the work and neighbourhood relations by 90% and set the infectiousness of the school relation to zero. This models a situation where only extremely essential workers are allowed to transit.

3.4. Cost model

Building a cost model for the pandemic is a daunting task, as it should encourage the minimization of infections while not imposing an excessive burden on the economy. Here we place special focus on preventing a collapse of the healthcare system as measured by the proportion of available hospital beds. Recall that we assume the availability of 0.15 hospital beds per 100 individuals (for COVID patients). Accordingly, we specify the cost of a percentage h of the population in the HOSPITALIZED state in any given day as the sigmoid curve $1/(1 + \exp(-44h + 6))$, shown in Figure 6 (left).

For the actions, we adopt a pragmatic/utilitarian approach where the cost of an intervention reflects the number of hospitalizations that we are willing to exchange for that level of economic constraints. We set the costs of Social Distancing, Light Quarantine, Hard Quarantine and Lockdown interventions, respectively, to match the costs (as given by the sigmoid function) of 0.05%, 0.1%, 0.125% and 0.15% of the population being hospitalized. The cost of an Unrestricted intervention is zero. Finally, we specify the total cost of an intervention at a given state as the maximum between the cost of the intervention and the cost of the number of hospitalized individuals. We also experimented with adding these two costs, but we observed better results with the maximum.

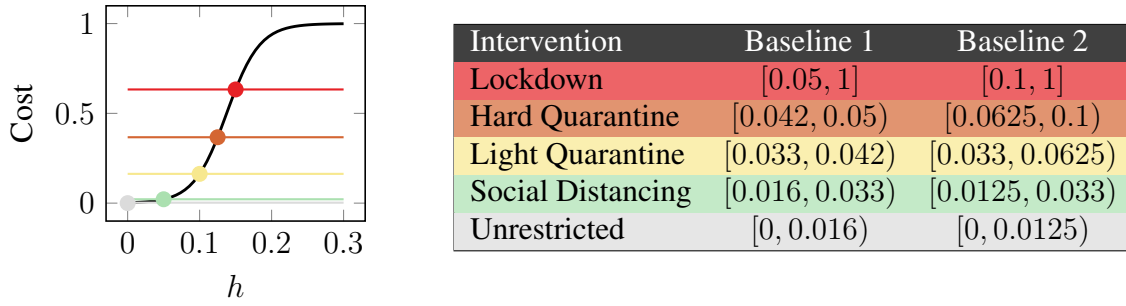


Figure 6. Left: Cost of intervention measures and hospitalizations. Right: Baseline mitigation (reactive) strategies.

4. Optimizing for Interventions

We are now all set for comparing different strategies for mitigating the outbreak according to our model. We assume that a decision step takes seven days of simulation, meaning that an intervention plan, once selected, remains in place for a week.

We first consider two baseline (time-independent) reactive policies, which adopt handcrafted rules of the form “If the percentage of hospitalized individuals is in a given interval, then adopt the following intervention measure”, as shown in Figure 6 (right). The corresponding intervals were defined based on the cost function and some trial and error to maintain the number of hospitalization within reasonable limits.

Figure 7 presents the curves of the proportion of hospitalizations per day of 10 simulations of the baseline strategies over 52 weeks (1 year). The black curve denotes the simulation with the median cost. For these simulations, we draw the background of each week to indicate the corresponding action (indicated by the color code used in Fig. 6). Both baseline strategies lead to an oscillatory behavior of the disease, with several peaks and valleys in the number of hospitalizations. They also lead to a number of hospitalizations above the capacity for a short period of time for some simulations, with Baseline 1 proposed for maintaining the median case below the capacity at all times, and Baseline 2 proposed for exceeding that capacity for the median curve. The results show that Baseline 1 plan is a more aggressive control measure, with Lockdowns and Hard quarantines being used even after a long period of the start of the outbreak (when the number of susceptible reduces considerably); it also has more periods with restrictions lifted (unrestricted). Baseline 2 plan, on the other hand, manages to reduce the number of peaks while also using less strict contingency measures (Lockdown and Hard quarantine).

We then considered optimized strategies using a simple online Monte Carlo approach with receding horizon H , as follows. At each decision step (i.e., start of a week) we consider all possible H -look head plans (i.e., all sequences of H interventions) where each subsequent intervention maintains, increases or reduces the restriction level of the previous intervention (e.g. a hard quarantine is followed by lockdown, hard quarantine or light quarantine). We then run S simulations of each plan starting from the current state, and we select the first intervention from the least expensive cost plan. Note that the cost of a plan sums over the costs of the days in the periods (as each decision step runs over a week). In short, we perform an online breadth-first search over H -step plans which are evaluated according to Monte Carlo cost estimates (over S simulations). In our

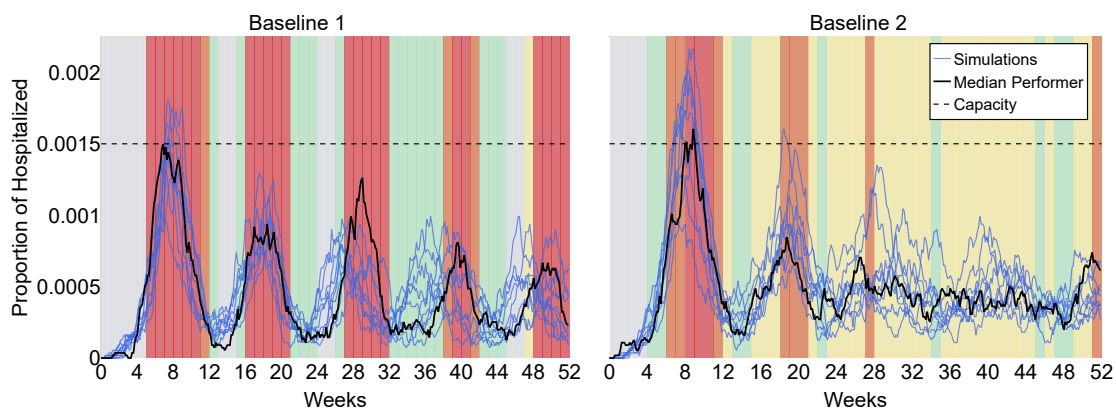


Figure 7. Hospitalization curves for the baseline strategies.

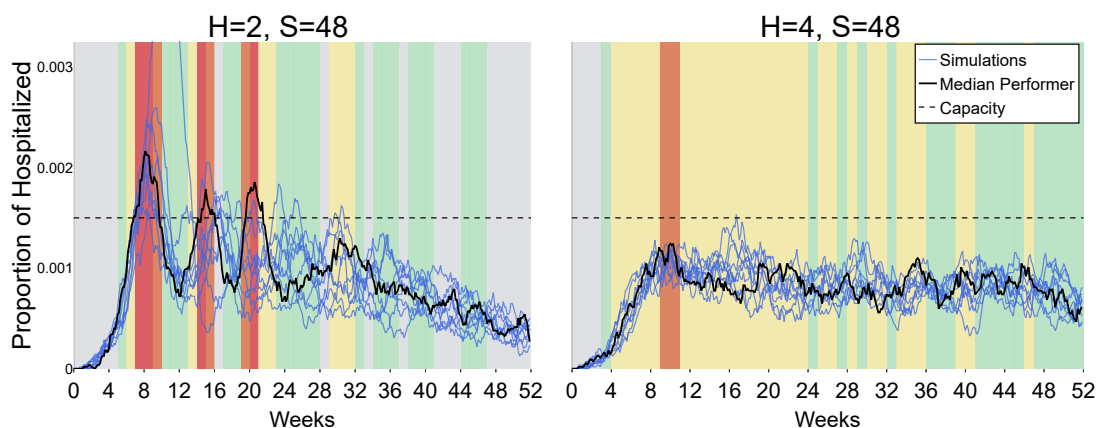


Figure 8. Hospitalization curves for the optimized strategies.

experiments, we use $S = 48$ with $H = 2, 3, 4$ weeks. Increasing H further was computationally too costly. With $S = 48$ and $H = 4$, a complete simulation of 52 weeks (with online planning) took on average 5 hours to complete on a high-end machine with AMD Threadripper 3960x 24-core CPU. All code and data will be released publicly upon acceptance.

Figure 8 presents the curves of hospitalizations per day for 10 simulations of the optimized strategies (in blue), with $S = 48$ and planning horizons $H = 2$ and $H = 4$ (we omit the case $H = 3$ for space). As before, the black line shows the simulation with the median cost, and the background is painted according to the actions used in that simulation. The first thing to notice is that the larger horizon H leads to much better performance, as expected. The $H = 4$ strategy is able to maintain the hospitalizations below the capacity for nearly all simulations at all days (note the small excess at a single curve for a single day), without resorting to Lockdown measures, and applying the Hard quarantine only once. Increasing the value of H had also the effect of acting earlier. The optimized strategies with $H = 2, 3, 4$ took on average 5, 4.2 and 3.3 weeks, respectively, to impose the first restriction measure. A longer horizon also leads to more sustained restriction measures. The first restriction measure remained active on average for 12.3, 16.1, and 22.2 weeks for $H = 2, 3$ and 4, respectively.

To have a more objective comparison among the different strategies, we computed

Strategy	Utilization	Capacity Overload	Cost
Unrestricted	1.85±0.027	293.8±9.0	81.2±1.9
Baseline 1	1.07±0.012	0.586±1.04	95.2±3.7
Baseline 2	1.01±0.02	1.775±1.882	69.5±4.0
H=2, S=48	0.73±0.102	14.7±18.2	68.8±5.3
H=4, S=48	0.74±0.029	0.003± 0.01	49.2±2.2

Table 2. Performance of the mitigation strategies.

the mean utilization of the health system, given by the absolute difference between the number of hospitalizations and the capacity, as well as the capacity overload, given by the area of the curve above the capacity threshold. These measures, which are not directly optimized by the Monte Carlo planners, measure the balance between the strictness of the intervention measures and the impact in the number of available beds. In Table 2 we show the results and the total cost of the executed actions. Relative to the baselines, we see that the optimized strategies obtain a much better compromise of economic costs and number of hospitalizations. The $H = 4$ planner, in particular, obtained a pronounced decrease in the mean utilization, and a sharp decrease in capacity overload. Note that the small standard deviations suggest that these results are robust across different simulations. Finally, the $H = 4$ optimized strategy minimized the cost (as expected).

5. Conclusions and Future Work

In this work we used data from the Origin-Destination survey to build a realistic network-based epidemiological model of the spread of COVID-19 in the city of São Paulo. We then used to model to compare rule-based (reactive) mitigation strategies against online plans optimized by a Monte Carlo approach. Our results showed that optimized policies manage to control the spread of the disease within the health care capacity while minimizing the economic impacts of social distancing policies. To the best of our knowledge, this is the first work to apply automated planning techniques for the control of COVID-19 in São Paulo city. A major caveat of our approach is to assume the individuals' states (susceptible, infected, infections, removed) are observable. That can however be approximated by regularly testing a sample of the population (as it was the case), and then randomly assigning individuals in our network accordingly (following age, geographic and other observed characteristics).

There are several improvements that can be done to make our model more realistic. For example, we can incorporate the distinction between symptomatic and asymptomatic individuals, condition the infectiousness of a contact on age, and incorporate income information into the effects of hospitalizations. We can also refine our cost model, incorporating additional resources (no. of available tests, UCI beds, protective equipment, staff, etc); we can distinguish work relations by sectors (essential/non-essential, or retail, industry, etc). Finally, we can use more sophisticated online planning techniques [Xue 2020], or extract offline policies [Kinathil et al. 2017] and be able to use larger horizon values.

Acknowledgements

This work was partially supported by the CNPq grants #304012/2019-0 and 307979/2018-0, and CAPES Finance code 001.

References

- Bai, W., Zhou, T., and Wang, B. (2007). Immunization of susceptible–infected model on scale-free networks. *Physica A: Statistical Mech. and Its Applications*, 384:656–662.
- Dureau, J., Kalogeropoulos, K., and Baguelin, M. (2013). Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics*, 14:541–55.
- Ferguson, N. M., Laydon, D., and et al. (2020). Report 9: Impact of non-pharmaceutical interventions to reduce COVID-19 mortality and healthcare demand. Technical report, Imperial College London.
- Halloran, M., Longini, I., Nizam, A., and Yang, Y. (2002). Containing bioterrorist smallpox. *Science*, 298:1428–1432.
- Keeling, M. and Eames, K. (2005). Networks and epidemic models. *Journal of the Royal Societcity Interface*, 2:295–307.
- Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society A*, 115(772):700–721.
- Kinathil, S., Soh, H., and Sanner, S. (2017). Nonlinear optimization and symbolic dynamic programming for parameterized hybrid markov decision processes. In *AAAI-17 Workshop on Symbolic Inference and Optimization*.
- Kucharski, A. J., Russell, T. W., and et al. (2020). Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The Lancet Infectious Diseases*, 20(5):553–558.
- Li, W., Zhang, B., Lu, J., Liu, S., Chang, Z., Cao, P., Liu, X., Zhang, P., Ling, Y., Tao, K., and Chen, J. (2020). The characteristics of household transmission of COVID-19. *Clinical infectious diseases*.
- Linton, N. M. and et al. (2020). Incubation period and other epidemiological characteristics of 2019 novel coronavirus infections with right truncation: A statistical analysis of publicly available case data. *MDPI, Multidisciplinary Digital Publishing Institute*.
- Mellan, T. A., Hoeltgebaum, H. H., Mishra, S., and et al. (2020). Estimating COVID-19 cases and reproduction number in Brazil. Technical report, Imperial College London.
- Newman, M. (2002). Spread of epidemic disease on networks. *Physical Review E*, 66.
- Silva, P. J. S. S., Pereira, T., and Nonato, L. G. (2020). Robot dance: a city-wise automatic control of COVID-19 mitigation levels. Technical report, Unicamp. medRxiv 2020.05.11.20098541.
- Wang, Z., Ma, W., Zheng, X., Wu, G., and Zhang, R. (2020). Household transmission of SARS-CoV-2. *The Journal of infection*, 81(1):179–182.
- Xue, S. (2020). *Scheduling and Online Planning in Stochastic Diffusion Networks*. PhD thesis, Oregon State University.
- Zhang, J., Litvinova, M., Liang, Y., Wang, Y., Wang, W., Zhao, S., Wu, Q., Merler, S., Viboud, C., Vespignani, A., Ajelli, M., and Yu, H. (2020). Changes in contact patterns shape the dynamics of the COVID-19 outbreak in China. 368(6498):1481–1486.