

Robot Dance: a mathematical optimization platform for intervention against Covid-19 in a complex network

Luis Gustavo Nonato *

Pedro Peixoto †

Tiago Pereira ‡

Claudia Sagastizábal §

Paulo J. S. Silva ¶

Abstract

Robot Dance is a computational platform developed in response to the coronavirus outbreak, to support the decision making on public policies at a regional level. The tool is suitable for understanding and suggesting levels of intervention needed to contain the spread of diseases when the mobility of inhabitants through a regional network is a concern. Such is the case of a highly contagious virus like Covid-19, for which the epidemiological compartmental models describing the infection dynamics must consider the circulation of people. Robot Dance anticipates the spread of the transmission in a complex regional network, identifying fragile links where applying differentiated measures of containment, testing, and vaccination is the most effective. By solving a stochastic optimization problem on a complex network, the model determines optimal strategies on the basis of commuting of individuals and the situation of hospitals in each district. Uncertainty in the capacity of intensive care beds is handled by a chance-constraint approach. Some functionalities of Robot Dance are illustrated on the state of São Paulo in Brazil, using real data for a region with more than forty million inhabitants.

AMS Subject classification: 90C15, 49-04, 92C60, 90B10

Keywords: mathematics of Covid-19, stochastic optimization, mobility matrix, complex networks

1 Context and motivation

The impact that Covid-19 has had on our lives is no news nowadays. Driven by the needs of their population and Governments, different countries have reacted against the pandemic in different manners. The diversity observed worldwide in public policies is in stark contrast with the uniform and massive response that the international scientific community has given without distinction of disciplines or frontiers. Throughout the globe, researchers swiftly gathered forces to offer clues and responses to the challenges presented by the pandemic on various fronts. The state of affairs is no different in Brazil, a country with continental dimensions and very heterogeneous society. This last feature poses additional hurdles to an already difficult situation, as it puts in question the effectiveness of uniform policies of quarantine adopted in many regions to contain the spread of coronavirus.

Robot Dance is an integrated platform developed to assess and forecast the consequences of interventions taken on a regional level when there is a disease outbreak and the spread of the disease is affected by the circulation of people living in the region. The question that all Governments face worldwide is how to mitigate (present and future) pandemic waves while limiting collateral economic damage. Robot Dance assists the decision making, providing a response that is customized to the specific local needs. By “local” we mean any territorial division, of political or administrative nature, into which the area of interest is partitioned (such as health districts, cities,

*ICMC, USP, São Carlos, SP, Brazil (gnonato@icmc.usp.br).

†USP, São Paulo, SP, Brazil (ppeixoto@usp.br).

‡ICMC, USP, São Carlos, SP, Brazil (tiago@icmc.usp.br).

§IMECC/UNICAMP, Campinas, São Paulo, Brazil (sagastiz@unicamp.br).

¶IMECC/UNICAMP, Campinas, São Paulo, Brazil (pjssilva@unicamp.br).

neighborhoods). Robot Dance considers the region as a network with nodes through which the disease propagates, following patterns whose shape and intensity change from night to daytime, particularly during working hours.

The public health of the region is the result of the dynamic evolution of elements that are intertwined in a nontrivial architecture in the complex network. For Robot Dance these elements are the mobility of inhabitants, the epidemiological situation in each district, and the hospitals' capacity in the considered region. Having been designed in a mathematical optimization framework, Robot Dance makes an optimal assignment of resources, taking into account those three aspects on a regional level, over a temporal horizon of several months.

The state of São Paulo, the most populous and the epicenter of the Covid-19 outbreak in Brazil [Can+20b] is a good example of a complex network where the virus circulates. With a surface comparable to the United Kingdom, and a population similar to the one in Spain, São Paulo is responsible for about 33% of Brazil's GDP. As illustrated by Figure 1, many people travel to work, commuting daily from smaller towns to large urban centers. The figure also reveals several such hubs spread across the state.




Figure 1: Number of daily trips in a typical pre-pandemic week in São Paulo state, Brazil (February 2020). The star-shaped lines, in darker colors when the circulation between nearby cities is higher, reveals each star center as an important commuting hub for the network. [Pei+20].

The situation described by Figure 1 for São Paulo state is typical not only in Brazil, but also in countries that have most of its economic infrastructure concentrated in a few central points. To go to work, large portions of the population are forced to travel long distances every day between cities. Because of this feature, keeping a global view both in space and time is essential to suitably capture the disease transmission. Besides, those nodes are highly heterogeneous (in terms of wealth, of hospital capacity, of business activities), making the impact of mobility also noticeable beyond the daily scale. The phenomenon was indeed observed for Covid-19 in São Paulo, with the well-off population initially carrying the virus from the hubs with international airports to the rest of the state, progressively more and more to the inland, following a long-term circulation pattern [Can+20a; Car+20].

In order to obtain a rigorous mathematical model that represents well the Covid-19 situation, it was crucial to suitably merge epidemiological knowledge with tools from several areas in applied mathematics, including optimization, data science, visualization, and scientific computing. Furthermore, in the considered setting, there is an acute lack of reliable data on the progress of the pandemic, due to subnotification, asymptomatic cases, and changes on how cases are reported over time in the official databases. This handicap is a recurrent difficulty for all the mathematical developments linked to Covid-19. To mitigate the impact of this issue on the decision process, in addition to intensive backtesting and cross-validation of the parameters defining the problem, Robot Dance incorporates uncertainty in the use of intensive care unit (ICU) beds. The limit in ICU capacity, a constraint in the optimization problem, is then handled in the form of probabilistic constraints, [CC60; Pré95; Den06]. In this

manner, even if reliable data is scarce, the tool still provides useful insight to compare on a qualitative level the merits of different strategies of containment.

The platform can determine different mitigation protocols that make optimal use of the ICU beds. This reveals particularly useful for a network like the one represented in Figure 1, having a complex heterogeneous topology. The numerical results in Section 5 show that considering the mobility network in the epidemiological model allows Robot Dance to pinpoint critical locations where targeted policies, of local nature, are the most effective. For the example of São Paulo, this means that instead of putting in lockdown the whole state at once, different degrees of social distancing can be imposed in different districts, depending on the local situation of the hospitals, the severity of infections, and the exchanges that the district in question has with other nodes in the network, because of commuting. Robot Dance can also anticipate where the creation of a pool of ICU beds, to be shared by nearby districts, is most effective, considering not only the epidemiological state in the districts directly involved, but also their interaction with the whole network, through the mobility of their inhabitants. Those functionalities are illustrated at the scale of the twenty-two health districts that compose São Paulo state.

Before presenting details about Robot Dance’s mathematical formulation, we provide a succinct overview of other mathematical models developed in response to the Covid-19 pandemic.

Some related works, contribution, and terminology

The literature about computational mathematical models that simulate and analyze the spread of Covid-19 outbreak has grown quickly in the past months. We focus the discussion on some works that describe the pandemic dynamics by compartmental models, like SIR and SEIR [Bra08]. Rather than a thorough review (a formidable enterprise), we handpick a few works to contextualize Robot Dance’s features. From this perspective, there are two main categories of models in the literature, one considering the pandemic dynamics only, and another group that complements the temporal distribution of the disease with geographical considerations, like Robot Dance.

Temporal methods typically rely on the calibration of parameters of compartmental models to analyze and forecast the behavior of the virus outbreak, drawing conclusions based on such estimates. For example, [Agu+20] first uses an SEIR compartmental model to estimate the reinfection, transmission, and recovery rates; these parameters are afterwards plugged into simulations, to determine when herd immunity can be reached. In [Che+20] the number of asymptomatic cases in an SIR model is approximated by fitting time-dependent transmission and recovery rates. In a manner similar, the approach [ASV20] fits the time-dependent parameters of an SEIR model to determine the effect of distancing measures and the time horizon required for such measures to take effect. Optimal control formulations of SEIR models are considered in [MR16] and [BG20], for defining strategies of vaccination and isolation, respectively. The latter work defines differentiated isolation strategies, depending on age dependent controls. Among temporal models, [Duq+20], the closest to Robot Dance, uses time triggers to forecast the severity of containment measures and keep the health system from collapsing in the short term. Rather than counting lockdown days, Robot Dance offers the possibility of alternating severe confinement with relaxed social distancing. Such “dance” can be performed differently for different nodes, for instance keeping open together two nearby districts that have complementary economic activities (or one open and one closed, if better for the business of the region in terms of work patterns).

Spatio-Temporal methods consider the geographical dissemination of the disease due to human mobility. The typical approach is to distribute the population in groups, setting a compartmental model in each group, adapting the model to accommodate mobility between groups. This is the proposal for the SEIR variant in [WLL20], one of the first spatio-temporal models in the context of Covid-19. More sophisticated approaches were proposed later on, such as [Are+20] which segregates the population in patches and in age stratum to assess the performance of containment protocols. In [Pei+20], the basis for the illustration in Figure 1, geolocation mobile phone data is used to estimate the risk of infection in each city.

Robot Dance is based on a spatio-temporal model that splits the population into groups (nodes), representing the mobility as links in the network. The tool can anticipate the geographical evolution of the disease and evaluate the consequences of potential containment and prevention strategies. This is fundamental as a response to studies like [Kis+20], foreseeing resurgences in contagion for at least three more years. Having cast the epidemiological transmission in a complex network, Robot Dance reveals certain epidemiological roads, links that will dominate the spread of the disease, as well as nodes where the need for ICU beds will be more acute. In this way, it is possible to forecast the effect of different protocols in a manner that the surges can be handled by the health

system in each district, without collapsing. A distinctive feature of Robot Dance is the optimization of protocols that can *alternate* social distancing measures in neighboring nodes. Allowing for contiguous districts to face different restriction levels makes it possible not only to keep some nodes relaxed while others undergo more severe measures of social distancing, but also to change the configuration of which districts are open and which ones are closed along the time horizon. It is this automatic choreography, alternating a “hammer” of strict confinement with the “dance” of relaxation, that gave the name to the tool, inspired from the blog post [Pue20].

The significance of alternating hammer and dance is undeniable, as the policymaker then has a strategy that avoids shutting down the economic activities in a region, all at once. The joint examination of city commute, health infrastructure, and specific containment measures district-wise, gives a global insight of many different actions. For instance, it is possible to analyze the effect on the whole region of “surgical” local measures, such as increasing the number of ICU beds in some key nodes, critical for the network. In the case study in §§ 5.3, 5.4, once Robot Dance identifies nearby nodes that are on some critical epidemiological link, better containment policies are computed, putting in place a mechanism of cooperation, to share hospital infrastructure.

Having put the contributions of Robot Dance in context, we are now in a position to give the mathematical formulation of its key ingredients. The mathematical optimization problem under consideration minimizes an objective function over a feasible set with constraints describing the virus transmission, how people commute between cities for work, and the hospital capacity. The corresponding models are detailed below.

2 Basic epidemiological constraints and continuous-time problem

Recall that nodes in the network correspond to cities, or administrative health districts, or some subdivision of the area under consideration. The analysis is done for a network with nodes in a set I , with cardinality $|I|$.

Along the time horizon defined by given initial and final times \mathbb{T}_0 and \mathbb{T}_1 , the epidemiological state of the region is characterized by the population compartments, of Susceptible, Exposed, Infected and Recovered individuals, considered as percentages of the total population in each district. In a continuous time formulation, the regional state at time t is given by the vector function

$$x(t) := (S, E, I, R)(t) \in [0, 1]^{4 \times |I|}, \text{ for } t \in [\mathbb{T}_0, \mathbb{T}_1],$$

noting that $S + E + I + R = 1$ for each node and time. When convenient, an individual component of the state vector is referred to by a sub-index, putting into brackets its ordinal position. This convention is used in (1) below, where S and I , the first and third components of the vector x , are denoted by $x_{[1]}$ and $x_{[3]}$, respectively.

We shall also make use of a coefficient $\alpha(t) \in [0, 1]$ to weigh in the portion of the considered time t that corresponds to the night hours. In a classical SEIR model, without mobility, $\alpha(t) \equiv 1$, but in Robot Dance, no mobility corresponds to the pandemic evolution only during the night hours, when individuals remain in their own node (typically during one third of the day, so $\alpha(t) = 1/3$ in our setting). The distinction is made necessary by the commuting to work during the day, with individuals traveling between nodes, potentially carrying the disease. Figure 2 illustrates how this circulation changes the proportion of the susceptible and exposed compartments among nodes, the precise mathematical formulation is given by the relations in Section 3.

The effective infection rate $r(t) \in [0, \bar{r}] \subset \mathbb{R}^{|I|}$ defines the way the disease evolves, while T_{inc} and T_{inf} are the incubation and infection periods (the former taken as the time from exposure to becoming an active spreader). For each node i , given an initial condition at time \mathbb{T}_0 , that is $x_0 = (S_0, E_0, I_0, R_0)$ with $S_0 + E_0 + I_0 + R_0 = 1$, the pandemic evolution during the night is described by a system of ordinary differential equations:

$$\begin{aligned} \dot{S}^i(t) &= -\frac{\alpha(t)}{T_{\text{inf}}} r^i(t) S^i(t) I(t) & \dot{E}^i(t) &= \frac{\alpha(t)}{T_{\text{inf}}} r^i(t) S^i(t) I(t) - \frac{1}{T_{\text{inc}}} E^i(t) \\ \dot{I}^i(t) &= \frac{1}{T_{\text{inc}}} E^i(t) - \frac{1}{T_{\text{inf}}} I^i(t) & \dot{R}^i(t) &= \frac{1}{T_{\text{inf}}} I^i(t). \end{aligned}$$

In this dynamic system, nonlinearity stems from the product $S^i(t) I^i(t)$, quantifying the amount of susceptible population of node i that gets exposed to infection. Defining the intermediate variable

$$y^{ii}(t) := S^i(t) I^i(t) = x_{[1]}^i(t) x_{[3]}^i(t) \quad \text{for } i \in I, \quad (1)$$

the relations for susceptible and exposed compartments become

$$\dot{S}^i(t) = -\frac{\alpha(t)}{T_{\text{inf}}} r^i(t) y^{ii}(t) \quad \text{and} \quad \dot{E}^i(t) = \frac{\alpha(t)}{T_{\text{inf}}} r^i(t) y^{ii}(t) - \frac{1}{T_{\text{inc}}} E^i(t) \quad (2)$$

(the somewhat redundant notation with a double super-index for the new variable will reveal useful in Section 3).

The additional variable exhibits the portion in the dynamics structure that is linear on the state. Specifically, introducing the matrix and vector

$$\mathbb{M} := \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -\frac{1}{T_{\text{inc}}} & 0 & 0 \\ 0 & \frac{1}{T_{\text{inc}}} & -\frac{1}{T_{\text{inf}}} & 0 \\ 0 & 0 & \frac{1}{T_{\text{inf}}} & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{m} := \frac{1}{T_{\text{inf}}} \begin{pmatrix} -1 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad (3)$$

the system of ordinary differential equations for the night cycle can be written as follows:

$$\dot{x}^i(t) = \mathbb{M} x^i(t) + \mathbf{m} \alpha(t) r^i(t) y^{ii}(t) \quad \text{for } i \in I. \quad (4)$$

Incidentally note that, because the columns in the matrix add up to zero and $S_0 + E_0 + I_0 + R_0 = 1$, the components of both $x^i(t)$ and $y^{ii}(t)$ lie in the cube $[0, 1]^{|I|}$. Also, thanks to the introduction of the additional variable (1), the relations in (4) involve, at most, the product of two variables (r and y). Keeping the model in a quadratic format, similar to the one in (4), has a tremendous impact on the numerical solution. We shall come back to this point in Proposition 3.1, after having incorporated the mobility between nodes in Section 3.

In our optimization model the control variable is the infection rate in the network, $(r^i(t), i \in I)$, for which an upper bound $\bar{r} \in \mathbb{R}^{|I|}$ is known. The functional space for the control will be $L_2([\mathbb{T}_0, \mathbb{T}_1], \mathbb{R}^{|I|})$, the Banach space of Lebesgue measurable functions from $[\mathbb{T}_0, \mathbb{T}_1]$ to $\mathbb{R}^{|I|}$. States will be in the Sobolev space $W^{1,2}([\mathbb{T}_0, \mathbb{T}_1], \mathbb{R}^{4|I|})$ of absolutely continuous functions, endowed with the norm $\|x\|_W^2 := \|x\|_2^2 + \|\dot{x}\|_2^2$. Suppose for the moment there are no more constraints (this is not the case for Robot Dance). If φ and ψ denote two functions defining performance indicators for the state and control, respectively, the problem to be solved is

$$\left\{ \begin{array}{ll} \min_{r(t) \in [0, \bar{r}]} & \int_{\mathbb{T}_0}^{\mathbb{T}_1} (\varphi(x(t)) + \psi(r(t))) dt \\ \text{s.t.} & (x(t), y(t), r(t)) \quad \text{satisfy (3)-(4) for a.e. } t \in [\mathbb{T}_0, \mathbb{T}_1] \\ & y(t) \quad \text{satisfies (1) for a.e. } t \in [\mathbb{T}_0, \mathbb{T}_1] \\ & x(\mathbb{T}_0) = x_0, r(\mathbb{T}_0) = r_0 \quad \text{are given initial conditions.} \end{array} \right. \quad (5)$$

This is an optimal control problem with state constraints for which a solution exists, by compactness of both the state and control. Furthermore, since (4) depends on the control linearly, when the function ψ is linear, the necessary optimality conditions given by Pontryagin maximum principle ensure that any optimal control is a concatenation of bang-bang and singular arcs [Cla13]. The formulation (12) solved by Robot Dance has a structure similar to (5), with additional constraints involving only the states or only the control. The dynamics combining the night cycle in (4) with the day dynamics described in Section 3 below) is no longer linear on the control, but the setting still corresponds to a so-called pure control-state constrained problem, in which the only explicit relation between the state and the control appears in the dynamical system.

To solve infinite dimensional problems like (5) we adopt the approach said “discretize-and-optimize”. That is, we consider a discrete approximation and use nonlinear programming methods to solve the resulting optimization problem. Convergence of the solution of the discrete optimization problems to a reference solution of the continuous problem is shown in [DHM00], under rather general conditions. These conditions, which hold for Robot Dance, are related to smoothness, independence, controllability and coercivity holding at the reference solution of the continuous problem; for full details, we refer to [DHM00],

Discretization of the differential equations amounts to choosing an approximation scheme for the left-hand side in (4). In our model, central finite differences appeared as a good compromise,

$$\dot{x}^i(t) \approx \frac{x_{t+1}^i - x_{t-1}^i}{2}.$$

Notice that now the time dependence is a sub-index, because we are dealing with vectors, no longer with functions. Regarding the auxiliary variable defined in (1), this means that now one of its sub-indices refers to the vector component and another one to the time step. Accordingly, given x^i, y^{ii}, r^i with vectorial time components x_t^i, y_t^{ii}, r_t^i for $t \in \mathcal{T} := \{\mathbb{T}_0, \mathbb{T}_0 + 1, \dots, \mathbb{T}_1\}$, the discretized system has the form

$$\begin{cases} x_{t+1}^i &= x_{t-1}^i + 2\mathbb{M}x_t^i + 2\mathfrak{m}\alpha_t r_t^i y_t^{ii} & t \in \mathcal{T}, \quad i \in I \\ y_t^{ii} &= x_{[t,1]}^i x_{[t,3]}^i & t \in \mathcal{T}, \quad i \in I, \end{cases} \quad (6)$$

for suitable initial conditions x_{-1}^i, x_0^i with components adding up to 1, and an initial reproduction rate r_0^i .

In its current format, Robot Dance does not consider age groups. Further compartments, in particular discriminating age groups, could be treated similarly, expanding the state vector x and its companion additional variable y into the corresponding new components (keeping in mind that such extensions will increase the dimension of a problem which is already very large-scale).

3 Modeling daily commute

As shown by Figure 2, during the day, transit between nearby nodes changes the dynamics of the susceptible and exposed compartments.

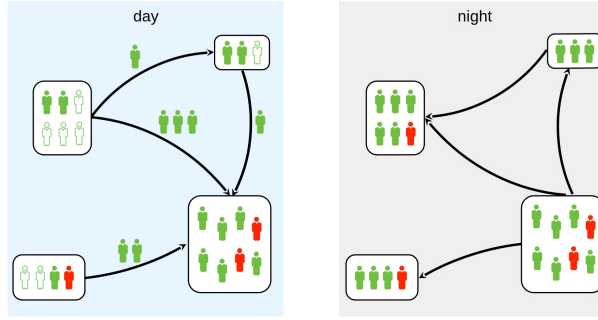


Figure 2: Commuting between nodes changes the proportions of susceptible and exposed individuals.

We make use of an $|I| \times |I|$ mobility matrix with entries $p^{ij} \in [0, 1]$, representing the percentage of inhabitants of node i traveling from node i to node j . Among several sources of mobility data that became available for Covid-19 research purposes, such as [URLa; URLb], we make use of high resolution smartphone geolocation provided by the Brazilian company InLoco [URLc]. The company gathers movement patterns of more than one-fifth of the inhabitants in São Paulo state, providing a social distance index with the so-called Hexagonal Hierarchical Spatial Index level 8 spatial resolution, accounting for hexagons of approximately 460 m edge length [URLd]. The index is the percentage of users that stayed home at a given day divided by the number of users in the company database for a given area. Also, from [Pei+20] and similarly to Figure 1, we have InLoco mobility data for São Paulo state with high resolution (approximately 1 km²), in the form of daily origin-destination movement counts.

Diagonal elements in the mobility matrix represent the proportion of people that did not leave the district. The values p^{ii} are given by the social distance index, added to the proportion of people that moved at least 1 km, but stayed in the district. The complement of those values, considering people who left the district i , is distributed among the off-diagonal terms p^{ij} , proportionally to the trip count recorded between nodes i and j .

As explained below, for the model to take into account mobility, the susceptible and exposed dynamics in (2) must incorporate commuting. First, for those nodes whose value $r^i(t)$ is smaller than the natural reproduction number of the disease without intervention, it is natural to assume that inbound travel will be discouraged, say by a factor $\zeta^i \in [0, 1]$. Letting N^i stand for the total population of node i , this is reflected by the entries of the *effective mobility matrix* and the *effective population*, defined by

$$\mathbb{P}^{ki}(t) := \zeta^i r^i(t) p^{ki} \quad \text{and} \quad \mathbb{N}^i(t) := \sum_{k \in I} \mathbb{P}^{ki}(t) N^k \quad \text{for } i \in I, \quad (7)$$

respectively. The second modification is related to the intermediate variable defined in (1) to represent the product

$$S^i(t)I^i(t),$$

the amount of susceptible population of node i that gets exposed to infection. Because commuting modifies the circulation flow, with mobility the dynamics must now consider the products

$$p^{ij}(t)S^i(t)I^j(t),$$

that is, the susceptible population of node i that gets exposed to infection at other nodes due to commuting there. The products involve the following *effective infection* ratio:

$$\mathbb{I}^j(t) := \frac{1}{\mathbb{N}^j(t)} \sum_{k \in I} p^{kj}(t) I^k(t) N^k \quad \text{for } j \in I. \quad (8)$$

Putting all these ingredients together gives the following ordinary differential equations for the susceptible and exposed compartments:

$$\begin{aligned} \dot{S}^i(t) &= -\frac{\alpha(t)}{T_{\text{inf}}} r^i(t) y^{ii}(t) - \frac{(1-\alpha(t))}{T_{\text{inf}}} \sum_{j \in I} r^j(t) p^{ij}(t) S^i(t) \mathbb{I}^j(t) \\ \dot{E}^i(t) &= \frac{\alpha(t)}{T_{\text{inf}}} r^i(t) y^{ii}(t) + \frac{(1-\alpha(t))}{T_{\text{inf}}} \sum_{j \in I} r^j(t) p^{ij}(t) S^i(t) \mathbb{I}^j(t) - \frac{1}{T_{\text{inc}}} E^i(t). \end{aligned}$$

where the factor $\alpha(t)$ is used to weigh the night cycle (2) with the mobility during day time (in fractions equal to 1/3 and 2/3, for example). The other equations, defining $\dot{I}^i(t)$ and $\dot{R}^i(t)$, are as before.

Since keeping a quadratic format as in (4) is crucial for implementation, we show below how to extend the device in (1) to suitably represent the 4th-degree monomial terms in the summations above.

Proposition 3.1 (Quadratic reformulation). *For all $i, j \in I$, the mobility elements $p^{ij}, \zeta^i \in [0, 1]$ are given parameters, as well as the night and day cycle weights $\alpha(t) \in (0, 1)$ for $t \in \mathcal{T}$.*

The SEIR model with mobility can be reformulated as follows:

$$\begin{aligned} \dot{S}^i(t) &= -\frac{\alpha(t)}{T_{\text{inf}}} r^i(t) y^{ii}(t) - \frac{(1-\alpha(t))}{T_{\text{inf}}} \sum_{j \in I} \zeta^j p^{ij} u^j(t) z^{ij}(t) \\ \dot{E}^i(t) &= \frac{\alpha(t)}{T_{\text{inf}}} r^i(t) y^{ii}(t) + \frac{(1-\alpha(t))}{T_{\text{inf}}} \sum_{j \in I} \zeta^j p^{ij} u^j(t) z^{ij}(t) - \frac{1}{T_{\text{inc}}} E^i(t) \\ \dot{I}^i(t) &= \frac{1}{T_{\text{inc}}} E^i(t) - \frac{1}{T_{\text{inf}}} I^i(t) \\ \dot{R}^i(t) &= \frac{1}{T_{\text{inf}}} I^i(t). \end{aligned} \quad (9)$$

In the dynamics (9) we introduced the additional state variables

$$y^{ik}(t) := S^i(t) I^k(t) = x_{[1]}^i(t) x_{[3]}^k(t) \quad \text{and} \quad z^{ij}(t) := \frac{\sum_{k \in I} p^{kj} y^{ik}(t) N^k}{\sum_{k \in I} p^{kj} N^k}, \quad (10)$$

defined for $i, j, k \in I$, as well as the following additional control variable, defined for $j \in I$:

$$u^j(t) := r^j(t) r^j(t). \quad (11)$$

Proof. We shall rewrite the rightmost terms in the susceptible and exposed compartments as follows

$$r^j(t) p^{ij}(t) S^i(t) \mathbb{I}^j(t) = \zeta^j p^{ij} u^j(t) z^{ij}(t).$$

Plugging the expression (7) for the effective mobility in both the numerator and denominator in (8) gives that

$$\mathbb{N}^j(t) = \zeta^j r^j(t) \sum_{k \in I} p^{kj} N^k \quad \text{and} \quad \sum_{k \in I} \mathbb{P}^{kj}(t) \mathbb{I}^k(t) N^k = \zeta^j r^j(t) \sum_{k \in I} p^{kj} \mathbb{I}^k(t) N^k.$$

In turn, the corresponding equivalent expression for the probability of local infections is

$$\mathbb{I}^i(t) = \frac{1}{\sum_{k \in I} p^{ki} N^k} \sum_{k \in I} p^{ki} \mathbb{I}^k(t) N^k.$$

Shortening the denominator to $\mathcal{D} := \sum_{k \in I} p^{ki} \mathbb{I}^k(t) N^k$ for the sake of clarity, the following algebraic manipulations gives the stated result:

$$\begin{aligned} r^j(t) \mathbb{P}^{ij}(t) \mathbb{S}^i(t) \mathbb{I}^j(t) &= r^j(t) \mathbb{P}^{ij}(t) \mathbb{S}^i(t) \frac{1}{\mathcal{D}} \sum_{k \in I} p^{kj} \mathbb{I}^k(t) N^k \\ &= r^j(t) \mathbb{P}^{ij}(t) \frac{1}{\mathcal{D}} \sum_{k \in I} p^{kj} \mathbb{S}^i(t) \mathbb{I}^k(t) N^k \\ &= r^j(t) \mathbb{P}^{ij}(t) \frac{1}{\mathcal{D}} \sum_{k \in I} p^{kj} y^{ik}(t) N^k \\ &= \zeta^j p^{ij} u^j(t) z^{ij}(t). \end{aligned}$$

□

Notice that in (9) the equality $\dot{\mathbb{S}}^i(t) + \dot{\mathbb{E}}^i(t) + \dot{\mathbb{I}}^i(t) + \dot{\mathbb{R}}^i(t) = 0$ holds for all i and t . This ensures that the sum of these state variables is invariant in time (we omit writing the superfluous constraint $\mathbb{S}(t) + \mathbb{E}(t) + \mathbb{I}(t) + \mathbb{R}(t) = 1$). The full optimal control problem with mobility has the following format

$$\left\{ \begin{array}{ll} \min_{r(t) \in [0, \bar{r}]} & \int_{\mathbb{T}_0}^{\mathbb{T}_1} (\varphi(x(t)) + \psi(r(t))) dt \\ \text{s.t.} & (x(t), y(t), z(t), r(t), u(t)) \quad \text{satisfy (9) for a.e. } t \in [\mathbb{T}_0, \mathbb{T}_1] \\ & (y(t), z(t)) \quad \text{satisfy (10) for a.e. } t \in [\mathbb{T}_0, \mathbb{T}_1] \\ & u(t) \quad \text{satisfies (11) for a.e. } t \in [\mathbb{T}_0, \mathbb{T}_1] \\ & x(\mathbb{T}_0) = x_0, r(\mathbb{T}_0) = r_0 \quad \text{are given initial conditions.} \end{array} \right. \quad (12)$$

Given an initial state, the discretization of these epidemiological constraints using central finite differences involves at most the product of two variables. The feasible set in the optimization problem solved by Robot Dance is therefore nonconvex and quadratic. This structure is often exploited by modeling languages and optimization solvers like JuMP [DHL17] and IPOPT [WB05], the packages used by Robot Dance for modeling and optimization respectively. Keeping the quadratic structure for the optimization problem was crucial for efficiency.

4 Representing hospital capacity

In addition to the epidemiological constraints with mobility, Robot Dance considers the capacity that hospitals in the network have to attend individuals sick with the virus. As mentioned, a critical question is the quality of the available data, mostly with respect to the number of infections. In order to somehow mitigate this issue, Robot Dance sets a probabilistic constraint for the use of ICU beds in each district. Thanks to the structure of the considered uncertainty, the chance constraint can be cast into a deterministic equivalent reformulation, see [CC62; CC63]. The procedure is explained below.

In the network, the ICU bed capacity are values v_t^i , given for each $i \in I$ and $t \in \mathcal{T}$. Let icu_t denote the percentage of infected population that needs intensive care. Suppose for the moment the ratio is a fixed percentage given in input, considered uniform across the region under analysis. To define the capacity constraint for the i th

district, we consider a new variable V_t^i , representing an accumulation of sick individuals, over τ days prior to t :

$$V_t^i := \sum_{k=t-\tau}^t \mathbf{I}_k^i =: \mathbb{V}_t^i{}^\top x, \quad (13)$$

where the vector $\mathbb{V}_t^i \in \mathbb{R}^{|I||\mathcal{T}|}$ is suitably defined. The value of τ , usually ranging in $[7, 10]$, corresponds to the average number of days infected individuals typically spend in the ICU, see [Sou+20].

The deterministic constraint limiting the ICU attention has the expression

$$\text{icu}_t V_t^i \leq v_t^i. \quad (14)$$

In this simple version, the ratio icu_t is considered constant, say $\mathbb{E}[\text{icu}_t]$, the mean value over the nodes in the network. To make a more realistic modeling of this constraint, crucial for the problem under consideration, we consider the ratio is a stochastic process that we shall approximate by a time series, [DK12].

In order to determine the parameters of the time series, we make an estimation using historical records of intensive care unit beds and new infected individuals in the region, see § 5.2. The data was scaled using the minimum and maximum records, respectively denoted by \min and \max , so that a time-series model is adjusted for values in the range $[0, 1]$:

$$\mathcal{R}_t := \frac{\text{icu}_t - \min}{\Delta} \quad \text{for } \Delta := \max - \min. \quad (15)$$

By examining several indicators of model adequacy, the best fit was an autoregressive model of order p :

$$\mathcal{R}_t = c_0 + c_1 t + \sum_{j=1}^p \phi_j \mathcal{R}_{t-j} + \omega_t, \quad \text{where } \omega_t \sim \text{iid } \mathcal{N}(0, \sigma_\omega^2).$$

In the rightmost expression, the white noise ω_t is a random variable that is independent and identically distributed according to a normal distribution with zero mean and variance given by σ_ω^2 . In order to obtain the probabilistic expression for the constraint (14), some algebraic manipulations need to be done, to write the recursion as a function of the starting values. This is a given data, called the tendency of the series. Written in vectorial form,

$$\text{the tendency is } \mathcal{S}_0 := \begin{pmatrix} \mathcal{R}_0 \\ \mathcal{R}_{-1} \\ \dots \\ \mathcal{R}_{p-1} \end{pmatrix} \quad \text{and the time series is } \mathcal{S}_t := \begin{pmatrix} \mathcal{R}_t \\ \mathcal{R}_{t-1} \\ \dots \\ \mathcal{R}_{t-p} \end{pmatrix} \quad \text{for } t \geq 1.$$

$$\text{Extending the intercept, drift and noise with zeros, to vectors in } \mathbb{R}^p: \quad C_0 := \begin{pmatrix} c_0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad C_1 := \begin{pmatrix} c_1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \varepsilon_t := \begin{pmatrix} \omega_t \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

we can write the whole series as a vector:

$$\mathcal{S}_t = C_0 + C_1 t + A \mathcal{S}_{t-1} + \varepsilon_t, \quad \text{where } A := \begin{pmatrix} \phi_1 & \phi_2 & \dots & \phi_p \\ 1 & 0 & \dots & 0 \\ \vdots & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \end{pmatrix} \quad \text{is a } p \times p \text{ matrix.}$$

A recursive application for $t \geq 1$ results in the identity

$$\mathcal{S}_t = \sum_{k=0}^{t-1} A^k C_0 + \sum_{k=0}^{t-1} (t-k) A^k C_1 + A^t \mathcal{S}_0 + \sum_{k=0}^{t-1} A^k \varepsilon_{t-k}.$$

The explicit calculation of entries a_{ij}^k , for any power matrix A^k can be done efficiently before the optimization process starts. In particular, since the extended noise has mean 0, we obtain the expression for the expected value of the first component of the scaled series, which is the component of interest in our case:

$$\mathbb{E}[\mathcal{R}_t] := \sum_{k=0}^{t-1} a_{11}^k c_0 + \sum_{k=0}^{t-1} a_{11}^k (t-k) c_1 + \sum_{j=1}^p a_{1j}^t \mathcal{R}_{-j}, \text{ for } \mathcal{R}_t \text{ defined in (15).}$$

Scaling back using (15), yields the expected value of the time series for the original data for $t \geq \mathbb{T}_0 + 1$:

$$\mathbb{E}[\text{icu}_t] := (1 - \sum_{j=1}^p a_{1j}^t) \text{min} + \Delta \sum_{k=0}^{t-1} a_{11}^k c_0 + \Delta \sum_{k=0}^{t-1} a_{11}^k (t-k) c_1 + \sum_{j=1}^p a_{1j}^t \text{icu}_{\mathbb{T}_0-j}. \quad (16)$$

By a similar procedure, the recursion and the identity in (15) yield the expression for the time series

$$\text{icu}_t(\omega) = \mathbb{E}[\text{icu}_t] + \Delta \sum_{k=0}^{t-1} a_{11}^k \omega_{t-k}, \text{ for } t \geq \mathbb{T}_0 + 1. \quad (17)$$

In turn, plugging this relation in constraint (14), we see that

$$\left(\mathbb{E}[\text{icu}_t] + \Delta \sum_{k=0}^{t-1} a_{11}^k \omega_{t-k} \right) V_t^i \leq v_t^i.$$

As a result, given a sufficiently small value $\mathbf{p} \in (0, \frac{1}{2})$, requiring that the constraint holds with high probability (equal to $1-\mathbf{p}$) amounts to ensuring, for each $t \in \mathcal{T}$,

$$\mathbb{P} [\xi_t \leq v_t^i - \mathbb{E}[\text{icu}_t] V_t^i] \geq 1 - \mathbf{p}, \quad (18)$$

for the normally distributed random variable

$$\xi_t := \Delta \left(\sum_{k=0}^{t-1} a_{11}^k \omega_{t-k} \right) V_t^i \text{ where } \omega_t \sim \text{iid } \mathcal{N}(0, \sigma_\omega^2).$$

The chance-constraint model adopted for uncertainty disregards potential correlations between time steps for the ratio. A joint constraint, setting the probability over the whole time horizon, would be preferable, but estimating temporal correlations to implement that model is beyond the predictive power of the data at hand.

Being a linear combination of independent white noises, the Gaussian variable ξ_t has zero mean and variance

$$\sigma^2[\xi_t] := \sigma_\omega^2 (\Delta V_t^i)^2 \sum_{k=0}^{t-1} a_{11}^k{}^2. \quad (19)$$

The explicit deterministic equivalent formulation of the chance constraint (18) uses the inverse cumulative distribution function of the standard Gaussian distribution, denoted by F^- . Specifically, see for instance [Pré95; Ack+10], given a scalar upper bound g and a normally distributed scalar random variable ξ ,

$$\mathbb{P}[\xi \leq g] \geq 1 - \mathbf{p} \iff g \geq \mu[\xi] + F^-(1 - \mathbf{p})\sigma[\xi]. \quad (20)$$

In our case $\xi = \xi_t$ and $g = v_t^i - \mathbb{E}[\text{icu}_t] V_t^i$, which together with (18) and (20), yield the constraint

$$v_t^i \geq \mathbb{E}[\text{icu}_t] V_t^i + F^-(1 - \mathbf{p})\sigma[\xi_t],$$

because $\mu[\xi_t] = 0$. To derive the expression for the standard deviation from the variance in (19), recall that both

$\Delta, V_t^i \geq 0$ and rearrange terms, so that

$$\mathbb{E}[\text{icu}_t] \mathbb{V}_t^{i\top} x \leq v_t^i - F^-(1-p)\sigma_\omega \Delta \mathbb{V}_t^{i\top} x \sqrt{\sum_{k=0}^{t-1} a_{11}^k}^2 \quad \text{with } \mathbb{E}[\text{icu}_t] \text{ from (16), for } i \in I, t \in \mathcal{T}, \quad (21)$$

and where we replaced V_t^i by the rightmost expression in (13). This “robustified” inequality, to be compared with the initial deterministic inequality (14), is also a linear constraint. Probabilistic constraints are appealing in real-life problems, because of their straightforward and natural interpretation. The computational solution of probabilistically constrained optimization problems, exploiting structural properties in different settings, has been addressed in various forms in the stochastic programming literature, [KV06; AB16]. The model (21) is related to the approach in [GS12], see also [AS14]. Works based on sampling average approximations dealt with by integer programming techniques are [LA08] and [PAS09]. Augmented Lagrangians that effectively compute p -efficient points were considered in [DM12]. For other methodologies and solvers, we refer to [PP11] and [Bra12].

5 Robot Dance in action

Robot Dance is written in Julia [Bez+17] using the JuMP modeling language [DHL17]. The source code, data, and scripts used in our experiments are available in <https://github.com/pjssilva/Robot-dance>. We present two case studies that show the versatility of the platform for analyzing the social and economic impact of different configurations. The first experiment illustrates how enforcing alternation between loose and strong containment measures modifies the disease dynamics and improves the level of economic activity in the network. The second case study focuses on mobility and ICU beds. Before reporting these results, we provide some details on the computational implementation of the model.

5.1 Objective functions

Given initial conditions x_0 and r_0 , an upper bound $\bar{r} \in \mathbb{R}^{|I|}$, and a time horizon $\mathcal{T} = \{\mathbb{T}_0, \dots, \mathbb{T}_1\}$ typically formed by days, the mathematical optimization problem solved by Robot Dance has the form

$$\left\{ \begin{array}{ll} \min_{r_t \in [0, \bar{r}]} & \sum_{t \in \mathcal{T}} (\varphi_t(x_t) + \psi_t(r_t)) \\ \text{s.t.} & (x, y, z, r, u) \quad \text{satisfy versions of (9)-(11) in discrete time} \\ & x \quad \text{satisfies (21)} \\ & r \in \mathcal{R} \quad \text{for some polyhedral set } \mathcal{R} \subset \mathbb{R}^{|I| \times (\mathbb{T}_1 - \mathbb{T}_0)}. \end{array} \right. \quad (22)$$

The set \mathcal{R} may impose some temporal patterns on the control variable, for example, if considered advantageous for the region, to prevent too abrupt changes in consecutive weeks (switching from free circulation to lockdown).

In (22) the objective functions φ and ψ can be chosen by the user among several possibilities. In our runs, the state is not considered in the objective function ($\varphi_t \equiv 0$). As for the terms ψ_t , typical measures for assessing the control performance are often a combination of the following functions.

If ensuring a maximal circulation is a priority then, given weights $w_i := \frac{N^i}{\sum_{j \in I} N^j}$,

$$\psi_t^{\text{MAXCIRC}}(r_t) := \sum_{i \in I} w_i (\bar{r} - r_t^i)$$

keeps the rate as close as possible to the upper bound \bar{r} , understood as the one with free circulation under a “new normal”, see Section 5.2 below.

Given a minimum achievable reproduction number \underline{r} , the function

$$\psi_t^{\text{PATTERN}}(r_t) := - \sum_{i \in I} \delta_t^i w_i (\bar{r} - r_t^i) (\underline{r} - r_t^i)$$

follows the pattern of confinement and relaxation induced by the scalars δ_t^i .

To prevent lockdowns from being too long in a single node:

$$\psi_t^{\text{SHORT}}(r_t) := - \sum_{i \in I} \delta_t^i w_i (r_t^i - r_{t-1}^i)^2.$$

Finally, to encourage alternation of the containment levels between specific nodes,

$$\psi_t^{\text{ALTERN}}(r_t) := - \sum_{i \in I} \sum_{j \neq i \in I} \min(\delta_t^i, \delta_t^j) \min(w_i, w_j) (r_t^i - r_t^j)^2.$$

The last objective function is necessary to enforce controls of bang-bang type because, with mobility, the SEIR dynamic (9) is no longer linear on the control (the new control variable u in (11) is quadratic on r). Additionally, notice that with this configuration the objective function is nonconvex quadratic, and when solving (22) the nonlinear programming methods employed as solvers provide only local minima.

5.2 Benchmark information, parameters, and data fitting

The state of São Paulo (SP) is partitioned in 22 health districts, represented in Robot Dance by a complex network with 22 nodes intertwined by the mobility links shown in Figure 1. For all the experiments the optimization is performed for the whole SP state, over the period July 1, 2020 – July 28, 2021. Following São Paulo’s public policy, the controls r_t^i of each node are only changed every two weeks. In (21), the ICU capacity constraints hold in 90% of the cases ($p = 0.1$). With this setting, the largest nonconvex quadratic optimization instance of (22) in the benchmark had 275,033 variables, 274,452 equality, and 7,980 inequality constraints. Such problem was solved in 9 min and 34 s of wall time using Ipopt with the MA97 parallel linear solver [Gro13] on a desktop computer with an AMD Ryzen 1700X processor, that has 8 cores running at 3.4 GHz, and 64 GB of RAM.

We now describe briefly some basic parameters used in the implementation. We used daily time steps in the discretization of the differential equations (9). Such granularity in the time steps, together with the central difference scheme, proved sufficient to ensure good accuracy of the results in all of our experiments. Following [WLL20], the incubation and infected times were set to $T_{\text{inc}} = 5.2$ and $T_{\text{inf}} = 2.9$, respectively. In (22) the upper bound on the control, set to $\bar{r} = 1.8$, represents the virus reproduction number when there are no stringent restrictions on circulation, but considering a “new normal” setting. An estimate equal to 2.5, from [Liu+20], was suitable for the first months of the pandemic, before the month of July. Robot Dance employs an smaller value to capture the decrease in the reproduction rate that was observed as a result of behavioral changes in the society (stricter hygiene habits, face shields and other protective measures not associated with quarantines). On the other end, the lower bound $\underline{r} = 0.8$ is in line with the lowest achievable reproduction number observed under a severe lockdown, in European countries like Germany, Italy, and Spain [Hot+20].

Official data is well known to suffer from severe subnotification in Brazil [Sil+20]. To find a suitable value for the compensation, we first used public records for the city of São Paulo to define a trajectory for the R-compartment. The value for July 29th, 2020 was compared with the one estimated for the same date in a serological inquiry made in the whole city [URLe]. We observed that to arrive at the same number it was necessary to multiply the official data by $\Gamma := 11.6$. Such astonishingly high factor was confirmed experimentally by the field research [Hal+20]. The correcting factor was used in all the calculations, including the initial conditions x_0 , estimated running a variation of Robot Dance to fit the data publicly available for São Paulo state [URLf].

The parameters of the time series in (17) were calibrated from historical records of ICU beds occupancy and new cases from [URLg], discriminating the metropolitan area and the interior of the state, [Hal+20]. The ratio was defined taking moving averages as follows:

$$\rho_t := \frac{\text{Moving Average}[(\text{ICU beds})_t]}{\text{Moving Average}\left[\sum_{\ell=t-7}^t (\text{Corrected new cases})_\ell\right]}. \quad (23)$$

In the denominator, $\tau = 7$ as suggested by [Azi+20], using a proxy of (13) that corrects the available data with the

factor Γ and taking a mean over three days to smooth the intermittency observed in the records, after a weekend:

$$(\text{Corrected new cases})_\ell := \frac{\Gamma}{3} \sum_{k=\ell-3}^{\ell} (\text{new cases})_k. \quad (24)$$

For the calibration, only the records prior to $\mathbb{T}_0 = \text{July 1, 2020}$ were considered, excluding the first days, whose standard deviation exhibits an unusual cusp. Using a training set with 75% of the remaining data, various statistical tests indicated that adjustment with a Box-Jenkins auto-regressive process was suitable. We estimated models with lag, differencing and moving average parameters $(p, d, q) \in \{1, 4\} \times \{0, 1\} \times \{0, 1, 2\}$. The best fit was obtained with a pure autoregressive model of lag 2 ($p = 2$ and $d = q = 0$).

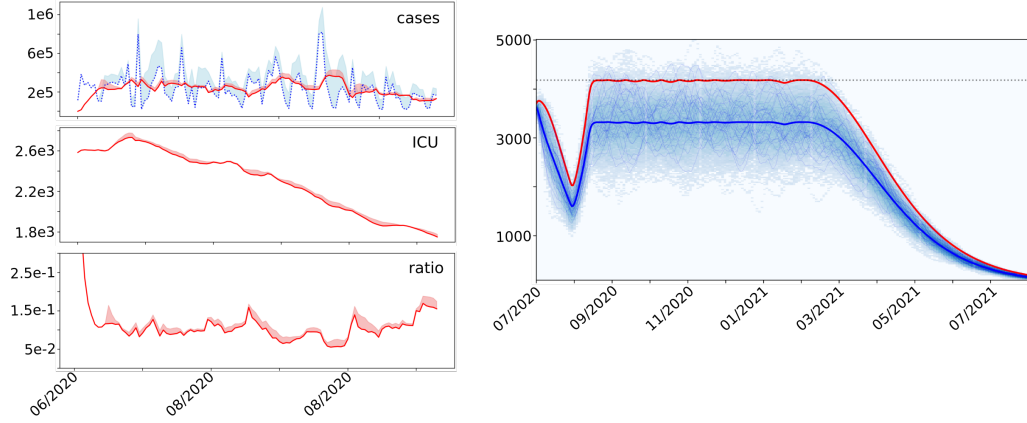


Figure 3: On the left, moving average and standard deviation of data defining the denominator (top), numerator (middle), and ratio (bottom) in (23). The blue line in the top plot corresponds to the original data (scaled with a factor 7Γ), shadowed by its standard deviation. On the right, trajectories forecasting the use of ICU beds in São Paulo city, using the time series (17) calibrated with the data on the left. The expected value and 90% quantile of the series are the two thicker lines, the dotted horizontal line represents the capacity.

The three graphs on the left in Figure 3 show the data available for defining the ratio ρ_t from (23), approximately 100 days since May 19, 2020. The red lines correspond to the denominator (top), numerator (middle) and quotient (bottom). Each line is shadowed by the standard deviation of the rolling mean. The blue line in the top represents the true new cases (without the correction (24) and without accumulation over seven days, scaled using a factor of 7Γ , to represent them together with the denominator in (23)). The fact that new cases are incorporated in the official database by “clumps” is very clear in the graph, as well as the significant variation this data exhibits from day to day. Incidentally, the big jump by the end of July was produced by a change of the testing policy in São Paulo state. Thanks to an increased investment of the Government in testing, all of a sudden a big number of new cases was uncovered in the state. The new numbers were all incorporated as individuals infected the same day in the database. This change of paradigm produced an abnormal increase in the record of that day.

The right plot in Figure 3 shows different trajectories of use of ICU beds in São Paulo city, forecasted from the data series (17), with tendency taken from the history of records.

5.3 Alternation of distancing protocols

In the network from Figure 1 considered by Robot Dance, SP city is the state capital, and has a population of 11.9 million inhabitants. The rest of the SP metropolitan area, with more than 10 million inhabitants, gathers a considerable number of small and medium-sized cities. Some of those cities host important factories and financial centers while others are essentially dormitory towns. Because of these features, many people commute daily in the metropolitan SP area, often traveling long time every day. In order to best capture the interaction of all the nodes in the state, Robot Dance considers its whole, with 22 health districts. For the sake of conciseness, we only report the results for the six districts that form the metropolitan SP area: SP city and five clusters of cities located

to the east, west, north, southeast, and southwest of the capital, denoted by E, W, N, SE, and SW, respectively.

To illustrate different choices of the decision maker, two different configurations were considered, one ensuring maximal circulation and another one enforcing alternation of the containment levels, every two weeks. The first option can be seen as a policy that favors a gradual relaxation of distancing measures. The second one, by contrast, switches from free circulation to lockdowns. This policy might be applicable in populations willing to endure two weeks of severe lockdown, as long as the confinement is planned in advance (for example for locations where, after enjoying two-weeks of fully open business, shop owners or factories can organize themselves to replenish stock or to do inventory during the lockdown periods). The respective objective functions in (22) are ψ^{MAXCIRC} and a weighed sum of all terms described in Section 5.1.

In order to ease the interpretation of the results, which vary both in time and space, the output of Robot Dance optimization process was organized in a diagram condensing all the information. In the visualizations shown on Figure 4, each row corresponds to one of the six health districts in the metropolitan SP area, using a color pattern for the severity of confinement. The consecutive rectangular blocks report the level (color) and duration (width) of the distancing protocols in the considered district. The black curve therein displays the dynamics of infected individuals. On the right of the diagram, a column indicates the highest percentage of the population that was sick, after an initial lockdown. This device, a “hammer”, ensures nonemptiness of the feasible set in (22).

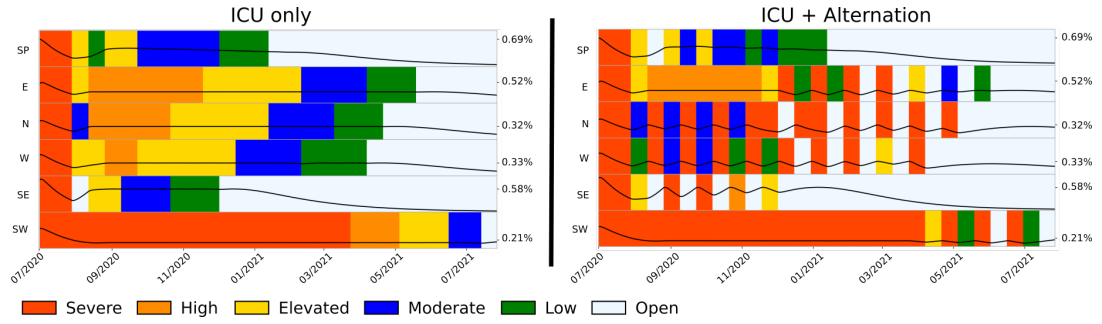


Figure 4: Comparison of two dynamics of distancing protocols in SP metropolitan area. On the left, when ensuring maximal circulation without exceeding the ICU capacity, the distancing protocols are progressively relaxed, but remain between severe and elevated levels for about six months in four regions (E, W, N, SW). By contrast, on the right, a protocol of alternation that also respects the ICU capacity has a clear positive effect, most notably in the west suburbs (W), during the months 9 to 12 in 2020. The column on the right of each diagram shows the maximal percentage of sick people attended by the health system, after an initial lockdown phase used to reduce the number of infected individuals, if their proportion was too high for the hospital capacity in the district.

On the left visualization in Figure 4, distancing protocols are progressively relaxed over time in almost all regions. However, four out of six regions (E, W, N, SW), remain between “severe” and “elevated” containment over almost six months. Such measures are very negative for non-essential business activities such as bars and restaurants and also for the educational and health systems. The situation improves on the right, where the alternation mechanism was put in place. We note that the city of São Paulo is not largely affected when the intermittent mechanism is turned on. However, now the regions N, W, and E greatly benefit from the “dance” that switches from two-week periods with loose protocols (fully open in the case of the SE) followed by 14 days of more rigid containment. Such a swing clearly mitigates the economic impact of the pandemic in those regions, as activities can be planned to happen during the relaxation periods, including replenishing stocks or promoting sales to take advantage of the open business. Finally, when comparing the black lines for the SE district in both graphs, we notice that alternation also relieves the stress on the medical staff (the health system is less pressured during the red block periods on the right).

5.4 Identifying critical links in the complex network

In Figure 4 we see that, even though the E, W, N, and SE benefit from alternation, neither SP city nor the SW district showed improvements. This phenomenon goes somehow against the expectation and was only detectable after running Robot Dance and analyzing carefully its results. This fact reveals an interesting interplay between the

two nodes in question. We focus on those two districts for our second case study. Our objective is twofold. First, to determine the impact of mobility, and second, to determine if the link SP-SW region is a crucial epidemiological road in the complex network.

To compute the output shown in Figure 4, Robot Dance considers that each region treats infected cases using only their own hospitals. We shall see that, even though the SW’s population is small (about one million, less than one twentieth of the total in the metropolitan area), its specific commuting patterns and poor hospital capacity has a significant impact on the SP city (with almost twelve million inhabitants, half of the total population in the metropolitan area). This happens because there is a big flux of people that commute from SP to SW during the day (corresponding to 12.5% of SW population) potentially carrying the disease. The left map in Figure 5 shows SP city in the green area, connected with red-yellow arcs of mobility with cities in the SW region, the blue area.

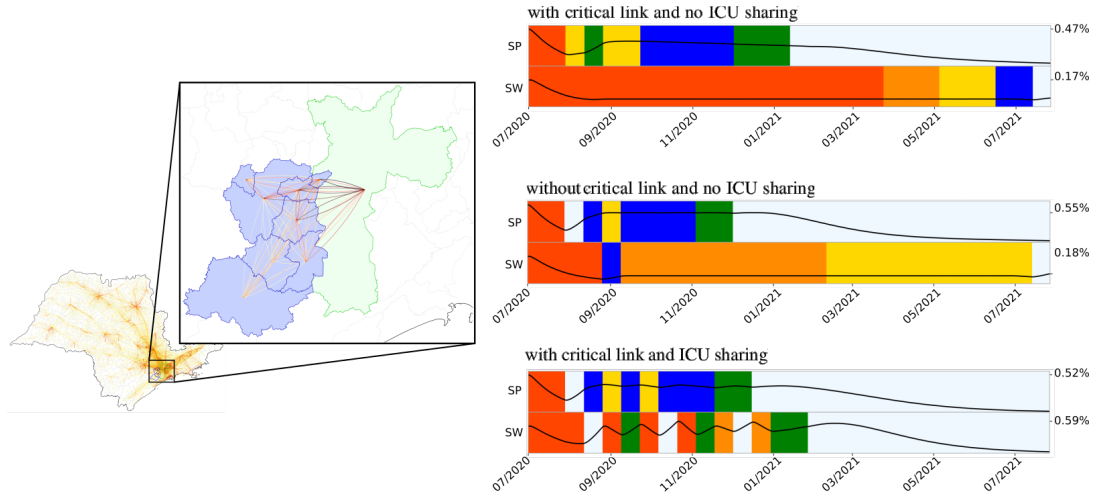


Figure 5: The left map zooms in the connections between the SP city and the SW district. The link SP-SW is critical for the network, but preventing commuting (eliminating the link) is not the solution. The middle graph shows that cutting the circulation in the link has practically no effect on the situation in SW/SP. By contrast, as shown in the bottom graph, allowing for commuting between the nodes (keeping the link) while implementing a pool of ICU beds improves significantly the overall state for SW without affecting much the situation in SP.

For this experiment, Robot Dance considers again all the 22 health districts in the state, using the objective function that maximizes circulation and again a probability level $p = 0.1$ in (21). There are three different configurations, shown in the visualizations on the right of Figure 4. The first one (top right), with free circulation, is referred to as “with critical link and no ICU sharing”. The second model (middle right) “without critical link and no ICU sharing”, forbids commuting between SP and the SW region (there is still mobility in the whole state). Finally, the third option (bottom right), “with critical link and ICU sharing”, allows commuting but defines a pool of ICU beds, to be shared among all the nodes in the metropolitan SP area.

The first diagram reproduces the left one in Figure 4, extracting the SP and SW rows. This setting allows for commuting, but infected cases of each node are attended locally. The comparison of the percentages on the right reveal a poor hospital capacity in the SW. In this node, only 0.17% of its own population can be sick (SP city dealt with a much higher value of 0.47% of its own population). This explains the more than eight months of lockdown, until the end of March 2021, that is forced onto the SW district in the top diagram (otherwise the health system collapses). Such a severe protocol is not realistic from an economic point of view, especially considering that the nearby SP city, less than 80 km away, fully opens its activities on the beginning of January 2021.

Having identified the node SW as critical from an epidemiological point of view, an alternative is to forbid commuting between the SW and SP city. The output of Robot Dance, shown in the visualization “without critical link and no ICU sharing” in Figure 5, confirms the crucial role that the SW region has in the network. Namely, without the flow of workers from the SW, the city of SP can fully open one month earlier, from December 2020, and its health capacity increases to 0.55% (from 0.47% in the top). The SW region can also relax a bit the strength of its measures from November 2020, facing one and a half month of high distancing followed by almost five

months of elevated measures before starting the alternation of hammer (red) and dance (white). The SW capacity to attend surges in its own hospitals still remains low, 0.19%.

Clearly, cutting the circulation for long periods is impractical, if at all feasible. Moreover, eliminating the critical link in the network is not the best option. The third policy addresses the drawbacks in the critical link with a pool of ICU beds, rather than only looking at mobility in the network. The percentages on the right in the bottom visualization in Figure 5 speak by themselves: SP maintains its high values (0.52%) while the SW maximal capacity jumps from 0.18% to 0.59%.

The use of ICU beds is shown in Figure 6 on the left, SP city in red and the suburbs in blue (the other 5 districts together, including the critical SW). The dotted curve shows the mean use, the thicker one the beds “reserved” by Robot Dance (the 90% quantile considered in the chance constraint (21) with $p = 0.1$). By comparing with the horizontal lines, with the total capacity in each node, we see that SP city large availability of ICU beds is never reached, while the SW remains close to saturation for months (the red curve is stuck against the horizontal line).




Figure 6: Having access to a pool of beds is beneficial for the whole SP metropolitan region. The plots show the average of ICU beds (dotted lines), the reserve induced by the chance constraint (red) and the maximum capacity in the region (black). On the left, a selfish configuration, with the suburb hospitals reaching the maximum capacity (within the chance constraint) for more than 3 months. On the right the result of sharing beds in a pool. Thanks to the beds made available from SP city, the suburbs have access to the extra capacity represented by the light red area. This additional reserve allows those districts to go sometimes beyond their own maximum level (dark red area). Robot Dance achieved all this gain without affecting much the situation in SP city.

In order to let the spare capacity in SP absorb part of the patients from the suburbs, Robot Dance was run with the constraints (21) grouped in a single pool of ICU beds for the SP metropolitan area. The positive impact of this measure is clear in the visualization “with critical link and ICU sharing” in the bottom right of Figure 5. Notice that, while the protocol does not change substantially in SP city (the three less severe containment measures only alternate more), having access to a pool of ICU beds has a tremendous positive impact on the SW region. Now the duration of containment measures drops to less than half, and SW is fully open from early January 2021. Moreover, severe distancing measures alternate with more relaxed protocols during the whole control period, thereby maintaining alive the economic and social activities in the district.

The fact of sharing ICU beds is beneficial for the whole metropolitan area, not only the critical SW node: with the pool of ICU beds, all the regions are fully open from January 2021 (we do not include the diagram for the whole state, to save space). The impact of the pool is also clear when comparing the left and right plots in Figure 6. The use of beds in SP is not much altered by the sharing mechanism, but the improvement for the suburbs is significant. The area shadowed in light red on the right plot represents the capacity gained by the suburbs (about 300 beds). This additional reserve of beds is actually employed by the suburbs, at those times where the dotted curve (mean use) gets higher than the horizontal line (local capacity). Those points are represented by the darker shadowed areas in Figure 6. Note in passing that the pool induces an effect similar to the alternation mechanism discussed in § 5.3 (this phenomenon is perceptible on the right plots in Figure 3, where the lines go up and down, making peaks).

To finish, recall that Robot Dance applies an initial “hammer” to guarantee nonemptiness of the feasible set in (22). This action corresponds to an initial short period of severe confinement that brings the percentage of infected population down to levels that are acceptable for the hospital capacity, in all the regions. The hammer is noticeable in the visualizations in Figure 4, all having red blocks in the first columns. The hammer materializes in the initial gully shape of all the trajectories forecasting the ICU bed use in the right of Figure 3, and also in the

curves in Figure 6.

6 Final comments

Social distancing measures have proven effective to contain the spread of Covid-19, but there is no universal recipe that can be applied throughout the globe. In each country, and for each region in a country, idiosyncrasies that depend on social and economic factors, individual to each place, must be taken into account. This issue, particularly important in Brazil, has increased relevance when the economic infrastructure is concentrated in megalopolis, forcing a large portion of the population to commute daily to work.

Robot Dance was formulated keeping these important considerations in mind. Its open source code has the versatility required to analyze and compare several configurations without much ado. To assess the robustness and identify possible limitations of the framework we conducted sensitivity analyses. On a qualitative level the model is relatively robust, proposing similar combinations of “hammer” and “dance” for variations of the initial conditions (up to 10%). Thanks to its SEIR representation with network mobility, Robot Dance can analyze interventions in specific nodes. Nodes can be neighborhoods in a city, or cities in a state, or states in a country. Section 5.4, comparing the model with and without a mechanism for sharing ICU beds in the SP metropolitan area, while analyzing the state as a whole, illustrates well this feature.

The problem solved by Robot Dance takes into account effects that are not directly observable, or intuitive, at first sight. This is explained by the mathematical optimization setting, that goes beyond simulations over a geographical network. Once again Section 5.4 is a good example of this situation: some crossed spatio-temporal effects make critical the link between the state capital and one much smaller suburban district; but simply eliminating the link, removing the circulation between the two nodes, does not resolve the issue.

The adaptability of Robot Dance is currently being extended to determine how to best deploy testing (or vaccination) campaigns, considering that in many countries nowadays, including Brazil, there is a rather low daily capacity for performing those tests (or vaccination). This goal is in syntony with the proposal in [Kis+20], that calls for an urgent need of projecting how the coronavirus will unfold in the coming years, to prevent recurrent wintertime outbreaks. As opposed to ad hoc simulations, the optimization framework of Robot Dance provides output that is reproducible, identifying possible critical links in a consistent manner on a regional level.

Acknowledgments Teseach supported by the following Brazilian agencies: CNPq (303552/2017-4, 304301/2019-1, 301778/2017-5, 306090/2019-0), FAPESP (2013/07375-0, 2016/18445-7, 2016/04190-7, 2018/24293-0), CEPID CeMEAI, FAPERJ (E-26/202.540/2019), and Instituto Serrapilheira.

References

- [Ack+10] W. v. Ackooij, R. Henrion, A. Möller, and R. Zorgati. “On probabilistic constraints induced by rectangular sets and multivariate normal distributions”. In: *Mathematical Methods of Operations Research* 71.3 (June 2010), pp. 535–549.
- [AS14] W. v. Ackooij and C. Sagastizábal. “Constrained Bundle Methods for Upper Inexact Oracles with Application to Joint Chance Constrained Energy Problems”. In: *SIAM Journal on Optimization* 24.2 (Jan. 2014), pp. 733–765.
- [ASV20] M. A. Acuña-Zegarra, M. Santana-Cibrian, and J. X. Velasco-Hernandez. “Modeling behavioral change and COVID-19 containment in Mexico: A trade-off between lockdown and compliance”. In: *Mathematical Biosciences* (2020), p. 108370.
- [AB16] L. Adam and M. Branda. “Nonlinear Chance Constrained Problems: Optimality Conditions, Regularization and Solvers”. In: *Journal of Optimization Theory and Applications* 170.2 (May 2016), pp. 419–436.
- [Agu+20] R. Aguas, R. M. Corder, J. G. King, G. Goncalves, M. U. Ferreira, and M. G. M. Gomes. “Herd immunity thresholds for SARS-CoV-2 estimated from unfolding epidemics”. In: *medRxiv* (2020).

- [Are+20] A. Arenas, W. Cota, J. Gomez-Gardenes, S. Gómez, C. Granell, J. T. Matamalas, D. Soriano-Panos, and B. Steinegger. “A mathematical model for the spatiotemporal epidemic spreading of COVID19”. In: *MedRxiv* (2020).
- [Azi+20] S. Aziz, Y. M. Arabi, W. Alhazzani, L. Evans, G. Citerio, K. Fischkoff, J. Salluh, G. Meyfroidt, F. Alshamsi, S. Oczkowski, E. Azoulay, A. Price, L. Burry, A. Dzierba, A. Benintende, J. Morgan, G. Grasselli, A. Rhodes, M. H. Møller, L. Chu, S. Schwedhelm, J. J. Lowe, D. Bin, and M. D. Christian. “Managing ICU surge during the COVID-19 crisis: rapid guidelines”. In: *Intensive Care Medicine* 46.7 (June 2020), pp. 1303–1325.
- [Bez+17] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. “Julia: A fresh approach to numerical computing”. In: *SIAM review* 59.1 (2017), pp. 65–98.
- [BG20] J. F. Bonnans and J. Gianatti. “Optimal control techniques based on infection age for the study of the COVID-19 epidemic”. Preprint HAL-02558980. Sept. 2020.
- [Bra12] M. Branda. “Stochastic programming problems with generalized integrated chance constraints”. In: *Optimization* 61.8 (2012), pp. 949–968.
- [Bra08] F. Brauer. “Compartmental models in epidemiology”. In: *Mathematical epidemiology*. Springer, 2008, pp. 19–79.
- [Can+20a] D. D. S. Candido, A. Watts, L. Abade, M. U. Kraemer, O. G. Pybus, J. Croda, W. de Oliveira, K. Khan, E. C. Sabino, and N. R. Faria. “Routes for COVID-19 importation in Brazil”. In: *Journal of Travel Medicine* 27.3 (2020).
- [Can+20b] D. S. Candido et al. “Evolution and epidemic spread of SARS-CoV-2 in Brazil”. In: *Science* (2020).
- [Car+20] R. F. Carmo, B. E. Nunes, M. F. Machado, A. C. Armstrong, and C. D. Souza. “Expansion of COVID-19 within Brazil: the importance of highways”. In: *Journal of Travel Medicine* 27.5 (2020).
- [CC62] A. Charnes and W. Cooper. “Chance constraints and normal deviates”. In: *J. Amer. Statist. Assoc.* 57 (1962), pp. 134–148.
- [CC60] A. Charnes and W. Cooper. “Chance-constrained programming”. In: *Management Sci.* 6 (1959/1960), pp. 73–79.
- [CC63] A. Charnes and W. Cooper. “Deterministic equivalents for optimizing and satisficing under chance constraints”. In: *Operations Res.* 11 (1963), pp. 18–39.
- [Che+20] Y.-C. Chen, P.-E. Lu, C.-S. Chang, and T.-H. Liu. “A Time-dependent SIR model for COVID-19 with undetectable infected persons”. In: *arXiv preprint arXiv:2003.00122* (2020).
- [Cla13] F. Clarke. *Functional Analysis, Calculus of Variations and Optimal Control*. Springer London, 2013.
- [Den06] D. Dentcheva. “Optimization Models with Probabilistic Constraints”. In: *Probabilistic and Randomized Methods for Design under Uncertainty*. Springer-Verlag, 2006, pp. 49–97.
- [DM12] D. Dentcheva and G. Martinez. “Regularization methods for optimization problems with probabilistic constraints”. In: *Mathematical Programming* 138.1-2 (Apr. 2012), pp. 223–251.
- [DHM00] A. L. Dontchev, W. W. Hager, and K. Malanowski. “Error bounds for Euler approximation of a state and control constrained optimal control problem”. In: *Numerical Functional Analysis and Optimization* 21.5-6 (Jan. 2000), pp. 653–682.
- [DHL17] I. Dunning, J. Huchette, and M. Lubin. “JuMP: A Modeling Language for Mathematical Optimization”. In: *SIAM Review* 59.2 (Jan. 2017), pp. 295–320.
- [Duq+20] D. Duque, D. P. Morton, B. Singh, Z. Du, R. Pasco, and L. A. Meyers. “Timing social distancing to avert unmanageable COVID-19 hospital surges”. In: *Proceedings of the National Academy of Sciences* (2020).
- [DK12] J. Durbin and S. J. Koopman. *Time Series Analysis by State Space Methods*. Oxford University Press, May 2012.
- [Gro13] C. M. Group. *HSL(2013). A Collection of Fortran Codes for Large Scale Scientific Computation*. <http://www.hsl.rl.ac.uk>. 2013.

- [GS12] V. Guigues and C. Sagastizábal. “Exploiting the structure of autoregressive processes in chance-constrained multistage stochastic linear programs”. In: *Operations Research Letters* 40.6 (Nov. 2012), pp. 478–483.
- [Hal+20] P. C. Hallal, B. L. Horta, A. J. D. Barros, O. A. Dellagostin, F. P. Hartwig, L. C. Pellanda, C. J. Struchiner, M. N. Burattini, M. F. da Silveira, A. M. B. Menezes, F. C. Barros, and C. G. Victora. “Evolução da prevalência de infecção por COVID-19 no Rio Grande do Sul, Brasil: inquéritos sorológicos seriados”. In: *Ciência & Saúde Coletiva* 25.suppl 1 (June 2020), pp. 2395–2401.
- [Hot+20] T. Hotz, M. Glock, S. Heyder, S. Semper, A. Böhle, and A. Krämer. “Monitoring the spread of COVID-19 by estimating reproduction numbers over time”. In: *arXiv preprint arXiv:2004.08557* (2020).
- [Kis+20] S. M. Kissler, C. Tedijanto, E. Goldstein, Y. H. Grad, and M. Lipsitch. “Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period”. In: *Science* 368.6493 (Apr. 2020), pp. 860–868.
- [KV06] W. Klein Haneveld and M. van der Vlerk. “Integrated chance constraints: Reduced forms and an algorithm”. In: *Comput. Manag. Sci.* 3 (2006), p. 245.
- [Liu+20] Y. Liu, A. A. Gayle, A. Wilder-Smith, and J. Rocklöv. “The reproductive number of COVID-19 is higher compared to SARS coronavirus”. In: *Journal of travel medicine* (2020).
- [LA08] J. Luedtke and S. Ahmed. “A sample approximation approach for optimization with probabilistic constraints”. In: *SIAM J. Optim.* 19 (2008), p. 674.
- [MR16] H. Maurer and M. do Rosario de Pinho. “Optimal control of epidemiological SEIR models with L_1 -objectives and control-state constraints”. In: *Pacific Journal of Optimization* 12.2 (2016), pp. 415–436.
- [PAS09] B. Pagnoncelli, S. Ahmed, and A. Shapiro. “Sample average approximation method for chance constrained programming: Theory and applications”. In: *J. Optim. Theory Appl.* 142 (2009), p. 399.
- [Pei+20] P. S. Peixoto, D. Marcondes, C. Peixoto, and S. M. Oliva. “Modeling future spread of infections via mobile geolocation data and population dynamics. An application to COVID-19 in Brazil”. In: *PloS one* 15.7 (2020), e0235732.
- [PP11] G. C. Pflug and A. Pichler. “Approximations for Probability Distributions and Stochastic Optimization Problems”. In: *Stochastic Optimization Methods in Finance and Energy*. Springer New York, 2011, pp. 343–387.
- [Pré95] A. Prékopa. *Stochastic programming*. Vol. 324. Mathematics and its Applications. Dordrecht: Kluwer Academic Publishers Group, 1995, pp. xviii+599.
- [Pue20] T. Pueyo. “Coronavirus: The Hammer and the Dance”. <https://medium.com/@tomaspueyo/coronavirus-the-hammer-and-the-dance-be9337092b56>. March 2020.
- [Sil+20] M. Silveira, A. Barros, B. Horta, L. Pellanda, G. Victora, O. Dellagostin, C. Struchiner, M. Burattini, A. Valim, E. Berlezi, et al. “Repeated population-based surveys of antibodies against SARS-CoV-2 in Southern Brazil”. In: *medRxiv* (2020).
- [Sou+20] K. V. M. de Souza Noronha, G. R. Guedes, C. M. Turra, M. V. Andrade, L. Botega, D. Nogueira, J. A. Calazans, L. Carvalho, L. Servo, and M. F. Ferreira. “Pandemia por COVID-19 no Brasil: análise da demanda e da oferta de leitos hospitalares e equipamentos de ventilação assistida segundo diferentes cenários”. In: *Cadernos de Saúde Pública* 36.6 (2020).
- [URLa] URLdata-online. URL: <https://www.google.com/covid19/mobility/>.
- [URLb] URLdata-online. URL: <https://covid19.apple.com/mobility>.
- [URLc] URLdata-online. URL: <https://mapabrasileirodacovid.inloco.com.br>.
- [URLd] URLdata-online. URL: <https://h3geo.org/>.
- [URLe] URLdata-online. URL: https://www.prefeitura.sp.gov.br/cidade/secretarias/upload/saude/17_9_2020_PPT_COLETIVAADULTO_FASE%5C%205.pdf.

- [URLf] URLdata-online. URL: <https://raw.githubusercontent.com/seade-R/dados-covid-sp>.
- [URLg] URLdata-online. URL: <https://www.seade.gov.br/coronavirus/>.
- [WB05] A. Wächter and L. T. Biegler. “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming”. In: *Mathematical Programming* 106.1 (Apr. 2005), pp. 25–57.
- [WLL20] J. T. Wu, K. Leung, and G. M. Leung. “Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study”. In: *The Lancet* 395.10225 (2020), pp. 689–697.