

Optimising Lockdown Policies for Epidemic Control using Reinforcement Learning

A summary and mathematical description

Harshad Khadilkar*, Tanuja Ganu⁺, Deva P Seetharam[#]

*IIT Bombay, TCS Research, ⁺Microsoft Research, [#]Independent consultant

Abstract—In the context of the ongoing Covid-19 pandemic, several reports and studies have attempted to model and predict the spread of the disease. There is also intense debate about policies for limiting the damage, both to health and to the economy. On the one hand, the health and safety of the population is the principal consideration for most countries. On the other hand, we cannot ignore the potential for long-term economic damage caused by strict nation-wide lockdowns. In this working paper, we present a quantitative way to compute lockdown decisions for individual cities or regions, while balancing health and economic considerations. Furthermore, these policies are *learnt* automatically by the proposed algorithm, as a function of disease parameters (infectiousness, gestation period, duration of symptoms, probability of death) and population characteristics (density, movement propensity). We account for realistic considerations such as imperfect lockdowns, and show that the policy obtained using reinforcement learning is a viable quantitative approach towards lockdowns.

I. EXECUTIVE SUMMARY

Important note: *This manuscript reports an early version of the model and results, while deeper analysis is in progress. We wish to disseminate this work as soon as possible, in the context of the current Covid-19 urgency. We will continue to add more detailed description and literature review in stages.*

The goal of this section is to present a non-technical summary of the methodology and the results of this study. More detailed mathematical description is available in Section II.

A. Contributions

The key contributions of this study are as follows.

- 1) We provide a methodology to compute the optimal lockdown/release policy for each node in a network, given disease characteristics and network properties.
- 2) This optimal policy is automatically learnt by the algorithm, using a technique known as reinforcement learning [1]. No explicit optimisation knowledge is required on the part of the user.
- 3) The policy is explicitly cognizant of health and economic costs, with the weight on each factor specified by the user's preferences.
- 4) The same algorithm can be used to compute the policy for any changes in (i) network data, (ii) disease parameters, and (iii) cost definitions. Only the relevant input values need to be updated by the user.

harshad.khadilkar@tcs.com, harshadk@iitb.ac.in,
tanuja.ganu@microsoft.com, deva.seetharam@gmail.com




Fig. 1. Network layout, with size of bubbles indicating population size. Black lines indicate strongest connections for each node, while blue lines indicate the next strongest set of connections. Note that each pair of nodes is potentially linked with each other.

- 5) The computed policies are fully explainable, which means that the users can ask useful questions about the interventions and the reasons behind each decision. In that sense, the optimisation algorithm is *not a black box*.

B. Disclaimers

Before proceeding further, we wish to make explicit the limitations of this study.

- 1) None of the authors are experts on communicable diseases. We have used disease models from available literature [2], which has a fair degree of consensus.
- 2) The network model used in this study is basic and not very sophisticated. It does account for population size and geography, but the effects are at a macro level. We do not model at the level of specific individuals and their movements.
- 3) So far, we have not used real data for the network model. We have tried to match the characteristics of Covid-19 as closely as possible, within the constraints of time and tractability.

C. Methodology in brief

We approach this study as a combination of three steps. These are described below.

Obtain or generate data: Before beginning the study, we need to define the network and disease parameters. We have




Fig. 2. Evolution of infection rates when the nodes are never locked down. The x-axis is in days, extending up to a year.

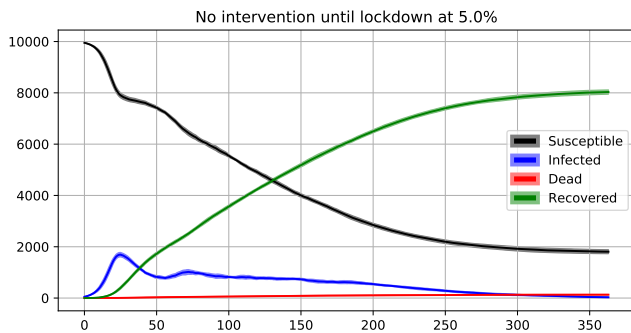


Fig. 3. Evolution of infection rates when each node is locked down when the fraction of symptomatic people in it exceeds 5%. The x-axis is in days, extending up to a year.

generated a network with 100 nodes and 10000 individuals. Each node contains a certain number of individuals, who need not be uniformly distributed. Strength of connections between each pair of nodes is (i) directly proportional to the product of the populations of the two nodes, and (ii) inversely proportional to the square root of the distance between them. The resulting network structure is shown in Fig. 1. Note that the figure only shows the two strongest connections for each node, but that all pairs of nodes are theoretically connected.

We also model the disease parameters as available online, for Covid-19. Specifically, the incubation period is 5-10 days, the actual infected period is 7-14 days, exposed people can begin transmitting the disease immediately, only about 80% of people actually show visible symptoms, and the death rate for infected and symptomatic people is 2%. We also assume that if an exposed or infected person comes into contact with a susceptible person, the transmission probability is 100%. All the parameters described here are probabilistic, and actual values are generated during simulation. We run multiple simulations to get reliable statistics.

Effect of lockdown: Throughout this study, we assume that an open node allows people to travel to/from other open nodes in the network. However, people showing symptoms are not allowed to travel to other nodes (but asymptomatic and exposed people can do so). Furthermore, symptomatic people are quarantined within the node, but a small fraction may still circulate within the node by breaking the quarantine.




Fig. 4. Evolution of reward during RL training.

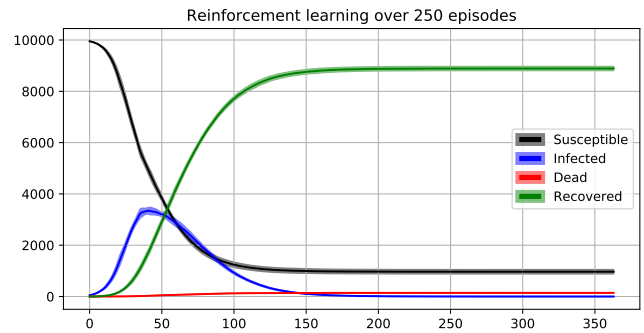


Fig. 5. Evolution of infection rates when each node is locked down according to the final reinforcement learning policy. The x-axis is in days, extending up to a year.

When a node is locked down, all travel to/from the node is blocked. However, a small fraction of people may still circulate within the node (especially ones who are not showing symptoms).

Baseline lockdown policies: In terms of intervention, we assume that each node has the option to be locked down or opened, once per week. The epidemic evolves on a daily basis, and after each week, the authorities can take this binary decision. As a reference, we define a set of policies that lock down any given node if the fraction of symptomatic people in that node cross a predefined threshold. We used threshold values of 5%, 10%, 20%, 50%, and 100% (the latter value ensures that nodes are never locked down). The first key question to be asked is, can we see the *flattening the curve* effect with these policies? The answer is yes: observe the evolution of infected people in Fig. 2 (with no lockdowns) and Fig. 3 (with lockdowns at 5% symptomatic people). We can see that the peak (blue curve) for the latter policy is lower, but a much longer period of infection is observed. The former policy also results in a higher number of deaths.

Train the reinforcement learning (RL) algorithm: RL works by running a large number of simulations of the spread of the disease, while attempting to find the optimal policy for lockdowns. The chief requirement is to quantify the cost of each outcome of the simulation. In this study, we place a weight of 1.0 on each day of lockdown, 1.0 on each person infected, and 2.5 on each death. A *reward* is defined as the negative of these costs (higher the reward, lower the cost).




Fig. 6. Comparison of infection rates across several policies.

As shown in Fig. 4, the RL algorithm is able to improve the reward over approximately 75 simulations, eventually performing as best as it can. The evolution of infection rates in Fig. 5 shows that the policy falls somewhere between the extremes shown in Fig. 2 and Fig. 3. A combined comparison for several policies is shown in Fig. 6, where we note that RL has the smallest peak that still has smooth evolution. The policies with 5% and 10% lockdowns have lower peaks, but they also show a kink when the locked-down nodes and opened up prematurely, and a second wave of infections is generated.

D. Analysis

The RL algorithm uses the following input features to take the lockdown decision for each node once per week:

- 1) Population of own node,
- 2) Fraction infected (symptomatic) in own node,
- 3) Fraction infected (symptomatic) in overall population,
- 4) Fraction of population recovered in own node,
- 5) Fraction dead in own node,
- 6) Potential infectors from rest of population,
- 7) Fraction increase in symptomatic population in own node in the last few days, and
- 8) Fraction increase in symptomatic population overall in the last few days.

We can ask interesting questions of this policy, as shown in Fig. 7. In this figure, we plot the actions taken as a function of two features at a time. All the other features are averaged over their observed values.

- In the first plot, we show the decisions as the fraction of symptomatic people increase within the same node (x-axis) and in the population overall (y-axis). Clearly, the policy is wary of increasing values in either direction.
- The second plot shows the actions for the population of a given node (x-axis) and its symptomatic population (y-axis). We can see that larger nodes are locked down earlier, once infection starts spreading.
- The third plot shows the actions for the fraction of symptomatic population in the node (x-axis) and the potentially infectious people arriving from other nodes (y-axis). If the potential for outside infection is higher, we see that the node is locked down as soon as infection starts spreading within the node.

II. DETAILED DESCRIPTION

We now proceed to give a detailed description of the mathematical model, for those who wish to go deeper into the technical details.

A. Generation of network and population structure

This portion is not necessary when the network data is derived from real sources. However, the data input needs to include, (i) the (x, y) location of each node, (ii) the population of each node in the network, and (iii) optionally, the strength of connectivity between each pair of nodes in normal times. Note that the third input is optional because one can use gravitational models [3] to approximate these connections.

In order to generate randomised data, we first define the number of nodes and the total population. The first step generates (x, y) locations for each node from a uniform distribution over a fixed range $[0, 100]$. The second step iteratively assigns a random number of people to each node, starting from small numbers (1 or 2). In successive iterations, the additional number of people attached to each node is proportional to the number of people already attached to that node. This process approximates the building up of large nodes (cities) with dense populations. We define the final population in node n as ρ_n . In the final step, we generate the strength of connectivity $\kappa_{n1, n2}$ between each pair of nodes. The strength is defined to be directly proportional to the product of the populations, and inversely proportional to the square root of distances between the nodes. This last portion deviates from traditional gravitational models. However, the square root ensures that large cities remain strongly connected (for example, Mumbai and Delhi) even though their distances may be large.

We also need to model the circulation of people within nodes and to other nodes. The parameters of this process are listed in Table I. Instead of physically modelling the movement of individuals, we compute the number of external (other node) population exposed by,

$$\text{Ext. pop. exposure for } n1 = \psi_{\text{open}} \sum_{n2 \neq n1} \kappa_{n1, n2} \cdot u_{n2, d} \cdot \rho_{n2}^*$$

if node $n1$ is currently open. Note that ρ_{n2}^* is the number of people in $n2$ who are *not* showing symptoms; we assume that symptomatic people are not allowed to travel at any time. The number of people within $n1$ who come into contact with each other is ψ_{open} for those not showing symptoms, and $\psi_{\text{quar, leak}}$ for those who are showing symptoms.

If node $n1$ is locked, we assume that there is no exposure to external population. All new exposures are produced by the small proportion $\psi_{\text{int, lock}}$ of people still circulating within the node.

B. Defining and modelling disease parameters

The parameters of infectious diseases that we have captured in this model are listed in Table II. The stages of the disease that we consider are shown in Fig. 8, based on prior literature [2] and known characteristics of Covid-19.




Fig. 7. Policy outcomes as a function of two features at a time.

TABLE I
POPULATION CIRCULATION AND CONTACT PARAMETERS.

Notation	Value	Explanation
-	50	Number of people initially infected and randomly dispersed in network
λ	$0 \leq \lambda \leq 1$	Fraction of population in a node, above which node must be locked
D_{lock}	7	Number of days which for which a lockdown is compulsorily enforced
ψ_{open}	0.1	What fraction of the population comes into contact with each person on a given day, when node is open
$\psi_{int,lock}$	0.01	What fraction of the population comes into contact within the node on a given day, when node is locked down
$\psi_{quar,leak}$	0.1	Fraction of symptomatic population that breaks quarantine and is at large
$u_{n,d}$	0 or 1	Whether a node n is open = 1 or locked = 0 on day d
ρ_n	$\rho \geq 0$ integer	Population of node n (after subtracting any that are dead)
$\kappa_{n1,n2}$	$\kappa_{n1,n2} > 0$	Connectivity strength between nodes $n1$ and $n2$

TABLE II
PROPERTIES OF THE DISEASE.

Notation	Value	Explanation
D_{cub}	10	Incubation days
D_{inf}	14	Days for which symptoms are shown
D_{trf}	1	Days after exposure when transmission can begin
P_{EI}	0.8	Probability of showing symptoms
P_{ID}	0.02	Prob. of symptomatic person dying
P_{trf}	1	Probability of infection of a susceptible person who comes into contact with an infected person




Fig. 8. Change in state for each person. S is susceptible, E is exposed (virus in the body but not yet affecting the immune system), I is infected (showing symptoms), A is asymptomatic carrier, D is dead, and R is recovered. Note that numbers are capture probabilities and not rates of change.

We assume that people exposed to the disease will eventually show symptoms with a probability $P_{EI} = 0.8$. The other 20% of people will be asymptomatic carriers. All people in states E, I, and A are capable of transmission to others. People in the I state die with a probability $P_{ID} = 0.02$, while the rest go the Recovered state R. We assume that asymptomatic carriers do not die, but go to R. The rate at which changes in state occur are inversely proportional to the stage durations as listed in Table II. For example, the probability of an E state person moving to I or A is equal to $1/D_{cub}$ on any given day. The number of such transitions on a day is produced by the requisite number of draws (equal to the population in a node in that state) from a multinomial distribution, with the relevant probabilities ($[1/D_{cub}, 1 - 1/D_{cub}]$ in the example given above). New exposures are produced by the circulation model as described previously.

C. Reinforcement learning model

We use an RL algorithm known as Deep Q Network [4] to compute the optimal lockdown policies. At the beginning

of each week, the model observes the relevant states for each node (described in Section I-D) and produces a binary decision: $u_{n,d} = 1$ for open and $u_{n,d} = 0$ for lockdown. This decision is derived from a fully connected neural network.

We use a single network for computing the values of both actions, in each state. The input size is 8 (the number of features). There are 2 hidden layers with 12 and 8 neurons respectively, each with \tanh activation. The output layer has 2 *linear* neurons, producing the long-term value estimates for either action. Once an action is computed for all nodes, the simulation rolls forward for 7 days (one day at a time) before calling the model for a new set of decisions.

Each simulation lasts for 52 weeks (364 days). However, the terminal rewards are computed as soon as the number of active infections in the whole network goes to 0 (only S, R, D people: no further state changes are possible after this point). The terminal reward for node n is given by,

$$R_{n,term} = a - b (c_{dead} x_{n,dead} + c_{inf} x_{n,inf} + c_{lock} d_{n,lock}),$$

where $c_{\text{dead}} = 25$ is the penalty multiplier for number of deaths, $c_{\text{inf}} = 10$ is the penalty multiplier for number of infections, $c_{\text{lock}} = 10$ is the penalty multiplier for number of days of lockdown, $x_{n,\text{dead}}$ is the final number of deaths in node n , $x_{n,\text{inf}}$ is the final number of infections in node n , and $d_{n,\text{lock}}$ is the number of days for which node n was locked down. The constants $a = 2$ and $b = 0.001$ are user-defined parameters. Note that the terminal reward for each node is independent of the performance of other nodes. This value is backpropagated up to the start of the episode using a discount factor of 0.99.

In addition to the terminal rewards, we also define step rewards as given by,

$$R_{n,\text{step}} = -\frac{c_{\text{dead}}\Delta x_{n,\text{dead}} + c_{\text{inf}}\Delta x_{n,\text{inf}} + c_{\text{lock}}(1 - u_{n,d})}{\rho_n},$$

where the Δ terms are changes in state over the past week, the last term penalises a lockdown, and the normalisation factor is the population of the node. The target reward for each step is the average of the discounted terminal reward and the step reward, and is applied to the value output for

the chosen action (0 or 1). The predicted reward for the other action is left untouched. The loss is computed using mean squared error (*mse*). Since we use the full Monte-Carlo reward at the end of the episode, we do not need to set up a Bellman equation that factors in the next state value. Training is carried out after each simulation using stochastic gradient descent optimiser in *keras* with a learning rate of 0.001 and a momentum of 0.8, for 5 epochs. The collected memory is discarded after every simulation episode.

REFERENCES

- [1] R. Sutton and A. Barto, *Reinforcement learning: An introduction*, 2nd ed. MIT Press, 2012.
- [2] L. Perez and S. Dragicevic, "An agent-based approach for modeling dynamics of contagious disease spread," *International journal of health geographics*, vol. 8, no. 1, p. 50, 2009.
- [3] M. Allamanis, S. Scellato, and C. Mascolo, "Evolution of a location-based online social network: analysis and models," in *Proceedings of the 2012 Internet Measurement Conference*, 2012, pp. 145–158.
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep RL," *Nature*, vol. 518, no. 7540, p. 529, 2015.