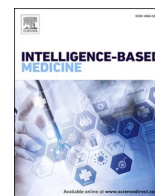




Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



VacSIM: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning

Raghav Awasthi^a, Keerat Kaur Guliani^b, Saif Ahmad Khan^a, Aniket Vashishtha^e,
 Mehrab Singh Gill^a, Arshita Bhatt^c, Aditya Nagori^d, Aniket Gupta^a,
 Ponnurangam Kumaraguru^a, Tavpritesh Sethi^{a,*},¹

^a Indraprastha Institute of Information Technology Delhi, India

^b Indian Institute of Technology Roorkee, India

^c Bhagwan Parshuram Institute of Technology, New Delhi, India

^d CSIR-Institute of Genomics and Integrative Biology, New Delhi, India

^e Maharaja Surajmal Institute of Technology, New Delhi, India

ARTICLE INFO

Keywords:

COVID-19

Vaccine distribution

Policy modeling

Reinforcement learning

Contextual bandits problem

ABSTRACT

A COVID-19 vaccine is our best bet for mitigating the ongoing onslaught of the pandemic. However, vaccine is also expected to be a limited resource. An optimal allocation strategy, especially in countries with access inequities and temporal separation of hot-spots, might be an effective way of halting the disease spread. We approach this problem by proposing a novel pipeline *VacSIM* that dovetails Deep Reinforcement Learning models into a Contextual Bandits approach for optimizing the distribution of COVID-19 vaccine. Whereas the Reinforcement Learning models suggest better actions and rewards, Contextual Bandits allow online modifications that may need to be implemented on a day-to-day basis in the real world scenario. We evaluate this framework against a naive allocation approach of distributing vaccine proportional to the incidence of COVID-19 cases in five different States across India (Assam, Delhi, Jharkhand, Maharashtra and Nagaland) and demonstrate up to 9039 potential infections prevented and a significant increase in the efficacy of limiting the spread over a period of 45 days through the *VacSIM* approach. Our models and the platform are extensible to all states of India and potentially across the globe. We also propose novel evaluation strategies including standard compartmental model-based projections and a causality-preserving evaluation of our model. Since all models carry assumptions that may need to be tested in various contexts, we open source our model *VacSIM* and contribute a new reinforcement learning environment compatible with OpenAI gym to make it extensible for real-world applications across the globe.²

1. Introduction

Vaccines have played a crucial role in combating infectious diseases for hundreds of years [1]. The successful eradication of smallpox, a transmittable disease responsible for causing huge casualties, happened due to the gradual and effective widespread use of vaccines [2]. Avoidance of millions of deaths and side effects caused by diseases like polio and tetanus was because of vaccination efforts which prove their impact and significance for global health. Despite this impact, if such a

vital resource is not optimally distributed and allocated to the regions and communities where its need is the most, this can lead to huge repercussions that could aggravate the situation and can increase the risk of escape variants in mixed and partially vaccinated communities [3]. All countries across the globe have eagerly waited for the launch of an effective vaccine against SARS-CoV-2. In India, the vaccine allocation process started on a first-come-first-serve basis, from January 2021 with the launch of the CoWIN³ portal by the Government of India. As potential candidates continue to introduce their products in the market,

* Corresponding author.

E-mail address: tavpriteshsethi@iiitd.ac.in (T. Sethi).

¹ Office: A-309, (R & D Block), Indraprastha Institute of Information Technology Delhi, 110020.

² <http://vacsim.tavlab.iiitd.edu.in:8000/>.

³ <https://www.cowin.gov.in/>.

there is an urgent need for optimal distribution strategies that would mitigate the pandemic at the fastest rate possible [4,5]. At the current rate of vaccination estimated to be 2.2 million doses per day, India may only be able to cover 30% of its population by end of this year, triggering concerns about a third wave hitting India, especially its rural parts, in the absence of herd immunity.⁴ Here we summarize the key factors that will need to be considered for effective mitigation:

- **Scarcity of supply:** Despite large scale production efforts, it is expected that the vaccine will still be a scarce resource as compared to the number of people who would need it. In addition to the vaccine itself, there may also be scarcity in the components leading to its delivery, e.g. syringes. Many researchers estimated that 60–70% of the population would require vaccination to achieve herd immunity. This implies that India would need 1.89 billion syringes to deliver the vaccine en-masse.⁵
- **Need for Equitable distribution and Digital inclusion:** A truly equitable distribution will not just be defined by the population or incidence of new cases alone, although these will be strong factors. Other factors ensuring equity of distribution include quantum of exposure, e.g. to the healthcare workforce that needs to be protected. In this paper, we assume that the exposure is proportional to the number of cases itself, although the proposed methodology allows more nuanced models to be constructed. There may also be unseen factors such as vaccine hoarding and political influence, which are not accounted for in this work. Another issue that has recently come into light is regarding the increased inaccessibility of vaccines in India, which are largely booked through the online CoWIN portal, owing to the significant digital divide in the country's population.
- **Transparent, measurable and effective policy:** The design of policy would need to be guided by data, before, during and after the vaccine administration to a section of the population. Since the viral dynamics are rapidly evolving, the policy should allow changes to be transparent and effects to be measurable in order to ensure maximum efficacy of the scarce vaccine resource. On the larger scale of states and nations, this would imply continuous monitoring of incidence rates vis-a-vis a policy action undertaken.

Although the aforementioned factors seem straightforward, the resulting dynamics that have emerged during the actual roll-out of the vaccine are far too complex for optimal decision making by humans. The daunting nature of such decision-making can be easily imagined for large and diverse countries such as India, especially where health is a state subject. Artificial intelligence for learning data-driven policies is expected to aid such decision making as there would be limited means to identify optimal actions in the real world. A 'near real-time' evaluation as per the demographic layout of states and consequent initiation of a rapid response to contain the spread of COVID-19 [6] will be required. Furthermore, these policies will need to be contextualized to the many variables governing demand or 'need' for the vaccine distribution to be fair and equitable [7]. Ground testing of these scenarios is not an option, and countries will have to face this challenge.

Reinforcement Learning (RL) algorithms have shown significant progress for decision making in medicine and public health [8–11]. In this paper, we introduce *VacSIM*, a novel feed-forward reinforcement learning approach for learning effective policy combined with near real-time optimization of vaccine distribution and demonstrate its potential benefit through the example of five states presenting contrasting contexts in India. Since real world experimentation was out of the question, the change in projected cases obtained via a standard epidemiological model was used to compare the *VacSIM* policy with a naive approach of incidence-based allocation. Finally, our novel model is

open-sourced and can be easily deployed by policymakers and researchers, thus can be used in any part of the world, by anyone, to make the process of distribution more transparent.

2. Methods

To get an optimal distribution of vaccine, we proposed a novel pipeline where we joined Deep Reinforcement Learning models with a Contextual Bandits model in a feed-forward way which makes the optimization process more robust, context-specific and easily deployable in real-time. In the following sections, we are adding some background information on Reinforcement Learning Algorithms in order to provide a sense of our pipeline.

2.1. Q learning

Policy learning algorithms can be broadly categorized into two types: off-policy and on-policy. Off-policy learning algorithms are preferred as these allow the agent to learn the optimal policy regardless of the agent's actions (under the current policy). This not only facilitates faster learning (convergence) but also avoids learning of a sub-optimal policy by the agent as it can explore continuously. One of the most popular off-policy algorithms in RL is Q-learning. Q-Learning uses the Bellman Equation for iterative updating:

$$Q^*(s, a) = E[r + \gamma \max_{a'} Q^*(s', a')] \quad (1)$$

where r is the immediate reward and γ is the discount factor. Given that an agent follows a policy π , then $Q_\pi(s, a)$ denotes the expected return or reward obtained by the agent if it chooses action a at state s and then keeps on following the policy π . $Q(s, a)$ estimates how favourable it is to choose action a at state s according to the current policy. The optimal Q function, $Q^*(s, a)$ denotes the maximum return or reward that can be achieved by first choosing action a at state s and then following the optimal policy. As evident from equation (1), $Q^*(s, a)$ is the sum of the immediate reward and the maximum reward possible thereafter from state s' discounted by γ (same as γ discounted maximum Q value for state s'). However, Q-Learning becomes impractical for environments where there may be a large (potentially infinite) set of actions to choose from. An example of a large environment with continuous action space [11] is optimization of number of syringes or medication dosage.

2.2. Deep Q-network

Deep Q-Networks (DQNs) are useful in the presence of large environments with elaborate state and action spaces. Here Q-learning is extended to DQN by using a deep neural network used to map input states to (action, Q-value) pairs [12,13]. Most real-world scenarios such as healthcare resource allocation have complex state-action spaces, and thus need advanced approaches such as DQNs. A DQN minimizes the loss at each iteration i as given below:

$$L_i(\theta_i) = E_{s,a,r,s' \sim \rho(\cdot)} [(y_i - Q(s, a; \theta_i))^2] \quad (2)$$

where $y_i = r + \gamma \max_{a'} Q(s', a')$; θ_{i-1} is the Temporal Difference (TD) target; $y_i - Q$ is the TD error and ρ represents behaviour distribution.

In our work, we employ DQN to better parameterize real-world conditions and make our model more efficient in terms of time and memory requirements.

2.3. Actor-critic using Kronecker-Factored Trust Region (ACKTR)

Q-learning and DQN are both value-based algorithms, where the goal is to learn a single deterministic action from a discrete set of actions by finding the maximum Q value. There are no trainable parameters in value-based methods that control the probabilities of action. Thus, a

⁴ <https://api.covid19india.org/>.

⁵ <https://www.nature.com/articles/d41586-021-00728-2>.

value-based method cannot solve an environment where the optimal policy is stochastic requiring specific probabilities. To overcome this challenge, we need to learn a stochastic map from state to action. In Policy Gradients [14,15], we usually use a neural network (or other function approximators) to directly model the action probabilities $\pi_\theta(a|s)$. Each time the agent interacts with the environment, we tweak the parameters θ of the neural network so that “good” actions (actions giving maximum reward) will be sampled more likely in the future. We repeat this process until the policy network converges to the optimal policy π^* .

Formally, the objective of Policy Gradients is to maximize the total future expected rewards $E[R_t]$, where R_t represents the sum of future discounted rewards. Whereas, in Q-learning and DQN the “policy” followed the action that maximized the Q value at each step, actor-critic methods combine Policy Gradient methods with a learned value function. With actor-critic methods [16], we learn two different functions: the policy (or “actor”), and the value (or the “critic”). In this paper, we used the ACKTR algorithm (on-policy algorithm) [17] with the Policy Gradient method which follows an actor-critic architecture: one deep neural network to estimate the policy and another network to estimate the advantage function (a measure of how good or bad a certain action is at a given state). ACKTR uses the Kronecker-factored approximation [18,19] to optimize both the actor and critic.

2.4. Bandit algorithms: multi-armed bandits and contextual bandits problem

Reinforcement algorithms being computationally expensive and time-consuming, cannot be used in real-time decision making. Contextual Bandits [20] algorithms on the other hand, are space and time-efficient. These are extension of multi-armed bandits algorithms. Multi-armed bandits can be simply be understood as a gambler with a fixed amount of money taking a sequence of actions on multiple slot machines while trying to optimize the overall reward. The K machines have reward probabilities, $\theta_1, \dots, \theta_K$. Action(A),Reward(R) tuples describe a multi-armed bandit, where at each time step t , an agent takes an action a on one slot machine and receives a reward r . The goal is to maximize the cumulative reward. This also corresponds to minimization of the potential loss by not picking the sub-optimal action, in case optimal actions with best reward are known, e.g. vaccination of a high-risk individual. The optimal reward probability θ^* of the optimal action a^* is:

$$\theta^* = Q(a_*) = \max_{a \in S} \{Q(a)\} = \max_{1 < i < K} \{\theta_i\} \quad (3)$$

where $Q(a)$ is the expected reward.

The Loss function or *Regret* [21] is a conceptual function to understand and optimize the performance of the Bandits problem. Since we don't know if an action played was the most ‘reward-fetching’, rewards against all the actions that can be played are sampled, and the difference between the action chosen and the action against the maximum reward is defined as ‘regret’. Therefore, minimizing regret achieves the goal of maximizing reward. The total regret we might have by not selecting the optimal action up to the time step T :

$$L_t = E[\sum_{i=1}^{t=T} (\theta^* - Q(a))] \quad (4)$$

Thus, a multi-armed bandit algorithm returns the best set of actions, provided the context is fixed. However, a multi-armed bandit algorithm does not account information from the changing context, a common occurrence in real-world scenarios such as the evolving pandemic. Therefore the agents may end up playing the same action multiple times even though the context may have changed, thus getting stuck at a sub-optimal condition. To circumvent this problem, Contextual Bandits algorithm is proposed which contextualizes the choice of the bandit to its current environment. Contextual Bandits, represented as a tuple Action

(A),Context(C),Reward(R), play an action based on its current context, given a corresponding reward, hence are more relevant to real world environments such as the vaccine distribution problem addressed in this work. Given, for time $t = 1 \dots n$, a set of contexts C and a set of possible actions A and reward/payoffs R are defined. At a particular instant, based on a context $c_t \in C$, an action $a_t \in A$ is chosen and a reward or a payoff

$$R = E[r|c, a] \quad (5)$$

is obtained. For an optimal action $a^* \in A$ such that the expectation of reward against this action is maximum,

$$r^* = \max_{a_i \in A} E(r|c, a_i) \quad (6)$$

the regret can be expressed as

$$Z = [r^* - E(r|c, a_t)] \quad (7)$$

and *cumulative regret* can be expressed as

$$Z^* = \sum_{i=1}^n Z \quad (8)$$

Contextual Bandits model was implemented using the python package of Vowpal Wabbit [22].

2.5. VacSIM: A feed forward approach

In **VacSIM**, we concatenated the Deep Reinforcement Learning models (i.e. ACKTR or DQN) and Contextual Bandits model in a feed-forward manner to create an off-policy learning framework. This was incorporated to select the optimal policy through the Contextual Bandits approach from the ones generated by the Deep Reinforcement Learning models, as shown in Fig. 1. This was done to address the following challenges that need to be tackled in real world problems such as optimal vaccine distribution:

- **Solving in real-time:** A vaccine distribution problem is expected to be fast-paced. Thus, an overwhelming amount of brainstorming with constrained resources and time would be required to develop an effective policy for the near future. This calls for the development of a prompt and an easily reproducible setup.
- **Lack of ground truth:** This is one of the key challenges in this paper. Since the roll-out of the vaccine will not give us the liberty of testing various hypotheses, a lack of ground truth data is generally expected. This is analogous to zero-shot learning problems, thus precluding a simple supervised learning-based approach.
- **Absence of evaluation with certainty:** Lack of ground testing naturally implies nil on-ground evaluation. In that case, it often becomes challenging to employ evaluation metrics that offer a significant amount of confidence in results. In order to solve this problem, we rely upon the simulated decrease in the number of susceptible people as the vaccine is administered.
- **Model scaling:** We have ensured that the learning process of the models simulates the relationship between different objects in the real world environment accurately, and at the same time can be scaled down in response to computational efficiency and resource utilization challenges. This is done by choosing the right set of assumptions that reflect the real world scenario.

We pipelined the Deep Reinforcement Learning models with a Contextual Bandits approach where recommendations for vaccine distribution policy(generated by the RL models) were used as training data for the Contextual Bandits. These models were trained separately using the same reward function defined in equation (6).Reinforcement Learning approaches replicate human decision-making processes. However, the absence of evaluation makes them less trustworthy, especially when real lives may be at stake. To this end, we also presented two different novel model evaluation techniques in this paper:

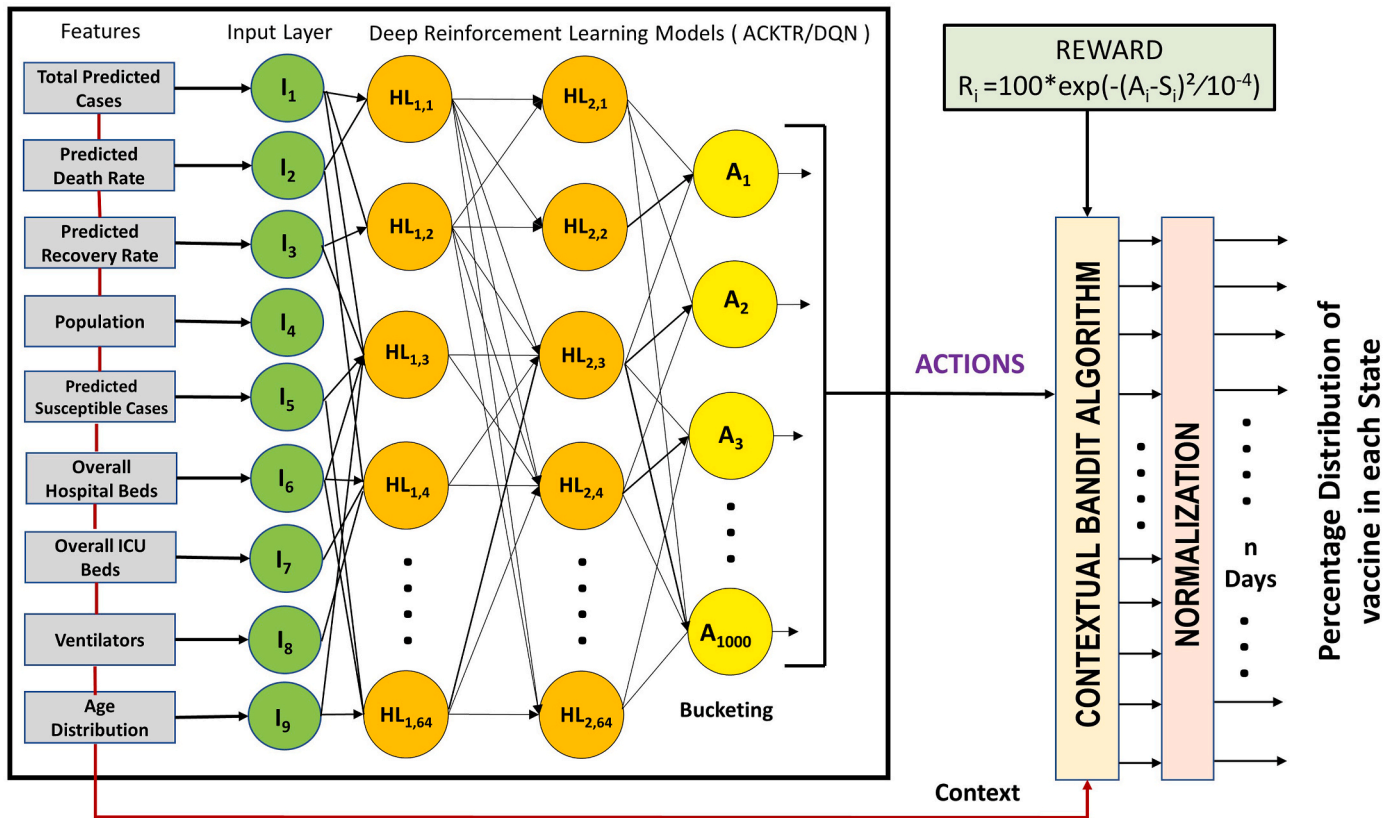


Fig. 1. VacSIM architecture: A novel feed forward pipeline for policy learning of COVID-19 vaccine optimal distribution. The actions and rewards obtained from the Deep Reinforcement Learning models were fed forward into the training of Contextual Bandit algorithm so that faster optimal online decisions could be calculated for further dates taking in context, the ever changing demographics of the states.

1. Estimating future impact of vaccination by projecting COVID-19 cases, using SEIR models [23].
2. Learning Bayesian Network to verify the underlying structure of COVID-19 trajectory indicators.

3. Results

3.1. Model setup

We extracted the State-wise time series data of COVID-19 characteristics from the website of Ministry of Health and Family Welfare of India.⁶ and used them in the experiments as described in Fig. 2. In order to determine vaccine allocation in the near future, we require future scenarios of COVID-19. Thus, in order to train VacSIM, we projected Susceptible, Confirmed, Recovered cases from 1 December 2020 to 26 December 2020 using SEIR model (refer supplementary for model details). The five States chosen for this study, i.e., Assam, Delhi, Jharkhand, Maharashtra and Nagaland are representative of possible scenarios for the vaccine distribution problem, i.e., high incidence (Maharashtra and Delhi), moderate incidence (Assam) and low incidence (Jharkhand and Nagaland). In choosing the five different States, we hope to generalize our predictions to other States across the spectrum while minimizing the bias introduced into the learning by a widely variant COVID-19 incidence across the country.

In order to avail vaccine for the highly vulnerable population and to combat COVID-19 infection spread, we selected the context comprising of features listed in Table 1.

The implementation of VacSIM pipeline is detailed henceforth.

3.2. Deep Reinforcement Learning models

Open-AI [26] stable-baselines framework was used to construct a novel and relevant environment suited to our problem statement for ACKTR and DQN to learn in this environment.

3.2.1. Input

We considered the observation space(s) as a grid over context vectors with the latter describing the context vector of states. These include Population, Age Distribution, Ventilators, Overall ICU Beds, Overall Hospital Beds, Predicted Susceptible Cases, Total Predicted Cases, Predicted Death Rate, and Predicted Recovery Rate at a given time. The Action space describes the quantum of vaccines distributed to each of the states across India, represented as discrete buckets. For example, a bucket size of 1000 implies a batch of 1000 vaccines rolled out to a particular state of India. Formally, the action space a was defined as, $a = \{b_0, b_1, b_2, b_3, \dots, b_b, \dots, b_{n-1}\}$

$$b_i = i * \frac{\text{Total Available Vaccine}}{n} \tag{9}$$

where n is defined as *bucket size* and considered as a hyper parameter.

3.2.2. Training

Following are the assumptions used while building the environment:

- The nature of the vaccine is purely preventative and not curative, i.e., it introduces immunity when administered to a susceptible person against future COVID infection but plays no role in fighting against existing infections (if any).

⁶ <https://mohfw.gov.in/>.

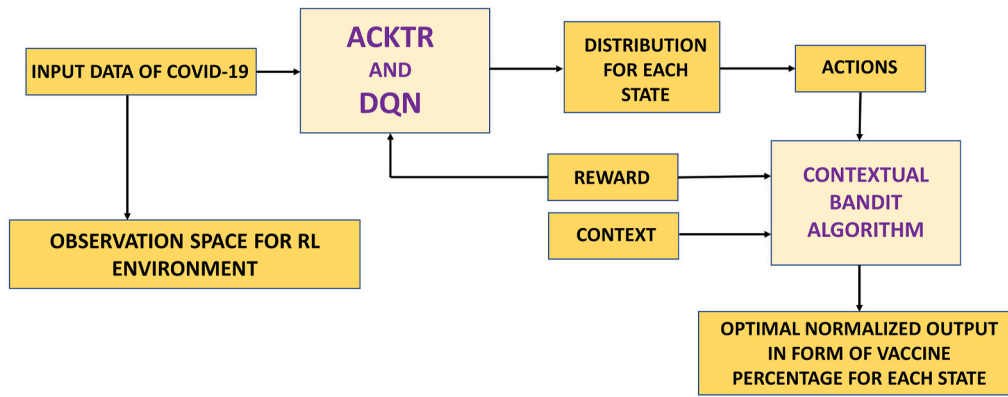


Fig. 2. Flexible model setup for optimization of vaccine distribution policy in India using the VacSIM approach where components can be easily replaced with alternatives and adopted for diverse settings.

Table 1 Features present in context space used in our models.

S. No.	Feature	Description
1	Predicted Death Rate	The percentage ratio of the predicted deaths in the State to the total predicted cases in that State calculated using projections obtained from a fitted standard <i>Compartmental model</i> , i.e. a Susceptible, Exposed, Infected and Recovered (SEIR) model.
2	Predicted Recovery Rate	The percentage ratio of the predicted recoveries in the State to the total predicted cases in that State using projections obtained from the SEIR model.
3	Population of a State	We extracted the population for each State from the 2011 census data conducted by the government of India [24].
4	Predicted Susceptible Cases of a State	We estimated the susceptible population as the difference between the population of a particular State and the total number of predicted infected cases of that State.
5	Hospital Facilities in a State	We used overall Hospital Beds, overall ICU Beds and Ventilators data in our models [25].
6	Age Distribution of a State	In order to prioritize vulnerable population, we considered people with age greater than 50.

- The vaccine has 100% efficacy, i.e. all people who are vaccinated will necessarily develop immunity in 45 days [27] against future SARS-CoV-2 infection. This assumption is easily modifiable and is expected to have little consequence in deciding the optimal allocation unless data reveal differential immunity response of certain populations within India. However, we leave scope for it to be modified as per the situation by taking this as a hyper-parameter for VacSIM.

Reward Function: The reward function was designed to maximize the decrease in the susceptible population count with the minimum amount of vaccine, considering it to be a limited resource.

$$R_i = 100 * exp(-(A_i - S_i)^2 / 10^{-4}) \tag{10}$$

where R_i is the reward given by the environment; S_i is the fraction of susceptible population in a particular State and A_i is the fraction of vaccine distributed to the State by the model.

The Explore-Exploit Trade-off: We tap into the explore-vs-exploit dilemma, allowing the model to reassess its approach with ample opportunities and accordingly, redirect the learning process. We set the exploration rate at 10% for DQN. However, this too is flexible and can be treated as a hyper-parameter. Here we reason that the local context is most strongly influenced by its own past values, hence we trained our

model independently for the selected 5 States from 1 December to 26 December, in order to maintain spatial independence. 10,000 and 2000 iterations Fig. 3 were used to learn DQN and ACKTR respectively. Iterations for training were decided based on the stability of the reward function curve. Where Iterations refers to one complete exploration of environment by a model agent.

Hyper-parameters: A complete list of the hyper-parameters is given in Table 2.

3.2.3. Output

The output of the Deep Reinforcement Learning models was a distribution set dictating the share of each recipient in the batch of vaccines dispatched at the time of delivery.

3.3. Contextual bandits model

The output of the Deep Reinforcement Learning model is spanned over 26 days (1 December-26 December 2020). The distribution sets so obtained were scaled to bucket sizes ranging from 100 to 1000 with an increment step of 100. The actions of the five States associated with each bucket size were normalized to get the percentage distribution ratio for all States. Normalized here refers to the percentage ratio of a given distribution set of a State for a given bucket size and the sum of the distribution sets of the five States over that bucket size for a given date.

3.3.1. Training

We considered 1 December-26 December data along with distribution sets and the corresponding set of rewards obtained from the Deep Reinforcement Learning models as training data for Contextual Bandits.

3.3.2. Number of actions and features

The action space and the features in the Context space are the same as those in the Deep Reinforcement Learning model.

3.3.3. Testing

Using the Context, actions and the corresponding set of rewards as the training dataset, we tested the model day-wise for a period of five days (27 December-31 December) with each day having 10 possible bucket sizes, i.e. 100 to 1000 for each State as output.

3.3.4. Output

The unadjusted actions (which were not normalized) obtained after testing the model were first adjusted bucket-wise for each State for a particular date by taking the percentage ratio of the unadjusted action of that State and the sum of unadjusted actions of all five States.

We consider that the vaccine is distributed on 31 December 2020. The consequent distribution of vaccines among the five States for each

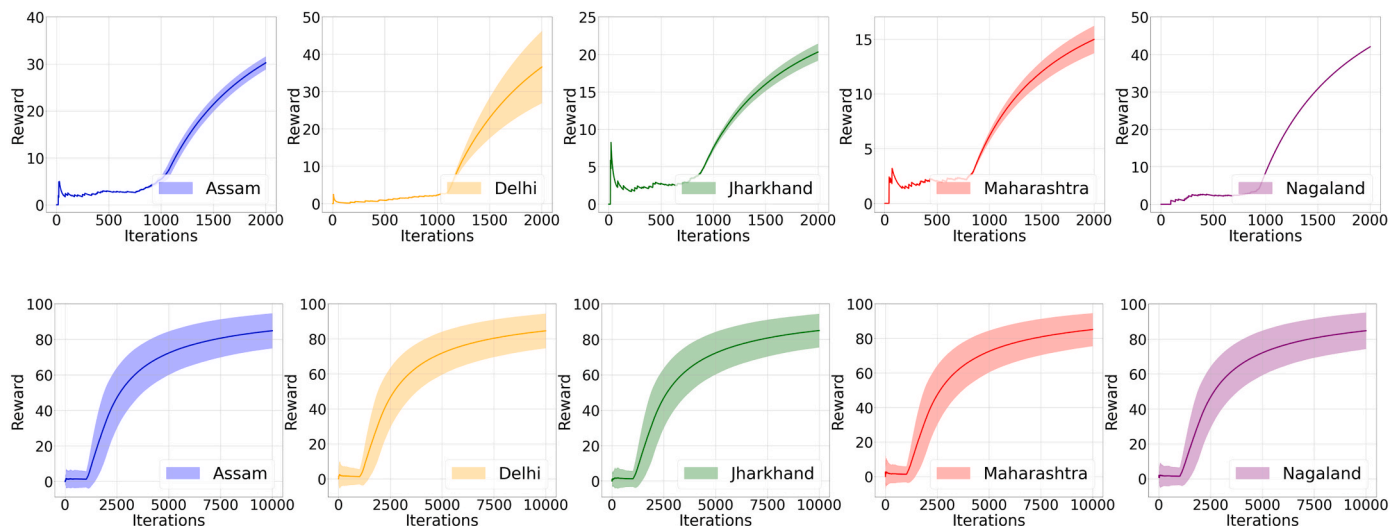


Fig. 3. Day-wise rewards at every iteration were recorded for all the states. The Mean and SD obtained over 26 day projections are shown as confidence bands. (Top) Smoothed and Increasing Reward Curve of DQN (Bottom) Smoothed and Increasing Reward Curve of ACKTR.

Table 2
Hyper-parameters used during Policy learning.

S. No.	Hyperparameter Name	Hyperparameter Value
1	Batch Size: number of vials in one round of distribution.	1000000
2	Exploration Rate of DQN	10%
3	Vaccine efficacy	100%
4	Number of days to reach full efficacy	45
5	Bucket size	1000
6	Number of recipients per day	5

bucket size was then evaluated. Our models learn well for the bucket range 200–500. The vaccine distribution for the same in case of the VacSIM policies is shown in Table 3. It was observed that both ACKTR + CB and DQN + CB yielded similar results, which is conceptually expected as both were aimed at reducing the overall susceptible population as optimally as possible.

3.4. Model evaluation through projection scenario using SEIR models

SEIR model is a differential equation-based epidemiological model used to simulate infection future trajectory. In the SEIR model, the total population N divides into 4 compartments S,E,I & R. Term S represents the susceptible individuals, E includes the latent phase of covid where the individual is infected but not infectious. I represent infected individuals and R represents the recovered individuals. Since there is no way that the evaluation of distribution policy can be done in the absence of a real world distribution, we defined the Naive baseline distribution policy as % of vaccine given to a State = % of Infected People in that State and compared it with our model’s learned distribution. With 10,00,000 doses and 5 States, we simulated the distribution of the available vaccine on 31 December for the Naive and VacSIM policies. The number of resulting (projected) infections for 45 days after the vaccine distribution was calculated using the SEIR Model. For each bucket size (200,300,400,500), day-wise total cases of all five States were summed up for both models (ACKTR/DQN) and then their differences were measured with total cases for all five States arising after Naive distribution as shown in Fig. 4. Our results indicate that the ACKTR based policy additionally reduces a total of 8845, 7787, 4703 and 5103 projected infected cases, with bucket sizes 200, 300, 400 and 500 respectively, in the next 45 days. Likewise, DQN based policy additionally reduces a total of 9039, 6686, 5355 and 4698 projected

infected cases, with bucket sizes 200, 300, 400 and 500 respectively.

3.5. Model evaluation through learning the causal structure of simulated data obtained

The ultimate goal of vaccine distribution is to reduce mortality and morbidity. Since our model relies entirely upon simulations, in the absence of a real-time vaccine data, we checked if the data generated by such an approach follow the cause-and-effect relationships as expected in the real world data. A Bayesian network is a special class of probabilistic graphical models which is useful for modeling causal relationships. Being a graph with directed edges without cycles, the parent-child relationships can be interpreted as cause and effect relationships under strong constraints. The directed acyclic graph (DAG) G is defined as a triple, $N = \{X, G, P\}$ defined that can be over a set of random variables, X . It encapsulates the underlying joint probability distribution P that can be factored as the product of probabilities of each node v conditioned upon its parents $par(v)$.

$$P(X) = \prod_{v \in V} P(x_v | x_{par(v)}) \tag{11}$$

where v corresponds to the random variable v and $par(v)$ are the set of parent variables.

Structure-learning of Bayesian network was carried out using Hill Climbing algorithm [28] with Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) as scoring functions, and ensemble averaging over 501 bootstrapped networks. These models were learned using wiseR package [29].

State-wise time series data of death, recovery, infected people, susceptible people and the amount of vaccine obtained from our model were used to learn the structure. Blacklisting among nodes was done such that vaccine percentage cannot be the child node of COVID-19 trajectory indicators (Susceptible, Recovery, Infected People, Death). The resulting structure shows a causal relationship between the vaccine amount (parent node) and susceptible count (child node), thus confirming the technical correctness of the VacSIM model through an external evaluation approach Fig. 5.

4. Discussion

Researchers worldwide have been working round-the-clock to find a vaccine against SARS-CoV-2, the virus responsible for the COVID-19 pandemic. Now that it is available, the distribution of the vaccines has

Table 3
VacSIM model output on different bucket size.

Model	Bucket Size	Assam	Delhi	Jharkhand	Maharashtra	Nagaland
ACKTR + CB	200	15.5	8	17.5	58	1
	300	15.5116	7.9208	17.4917	57.7558	1.3201
	400	15.5172	8.1281	17.734	57.3892	1.2315
	500	15.5206	8.055	17.6817	57.3674	1.3752
DQN + CB	200	15.8974	8.2051	16.9231	58.4615	0.5128
	300	15.9864	8.1633	17.0068	58.1633	0.6803
	400	16.1616	8.3333	17.1717	57.5758	0.7576
	500	16.129	8.2661	17.1371	57.6613	0.8065

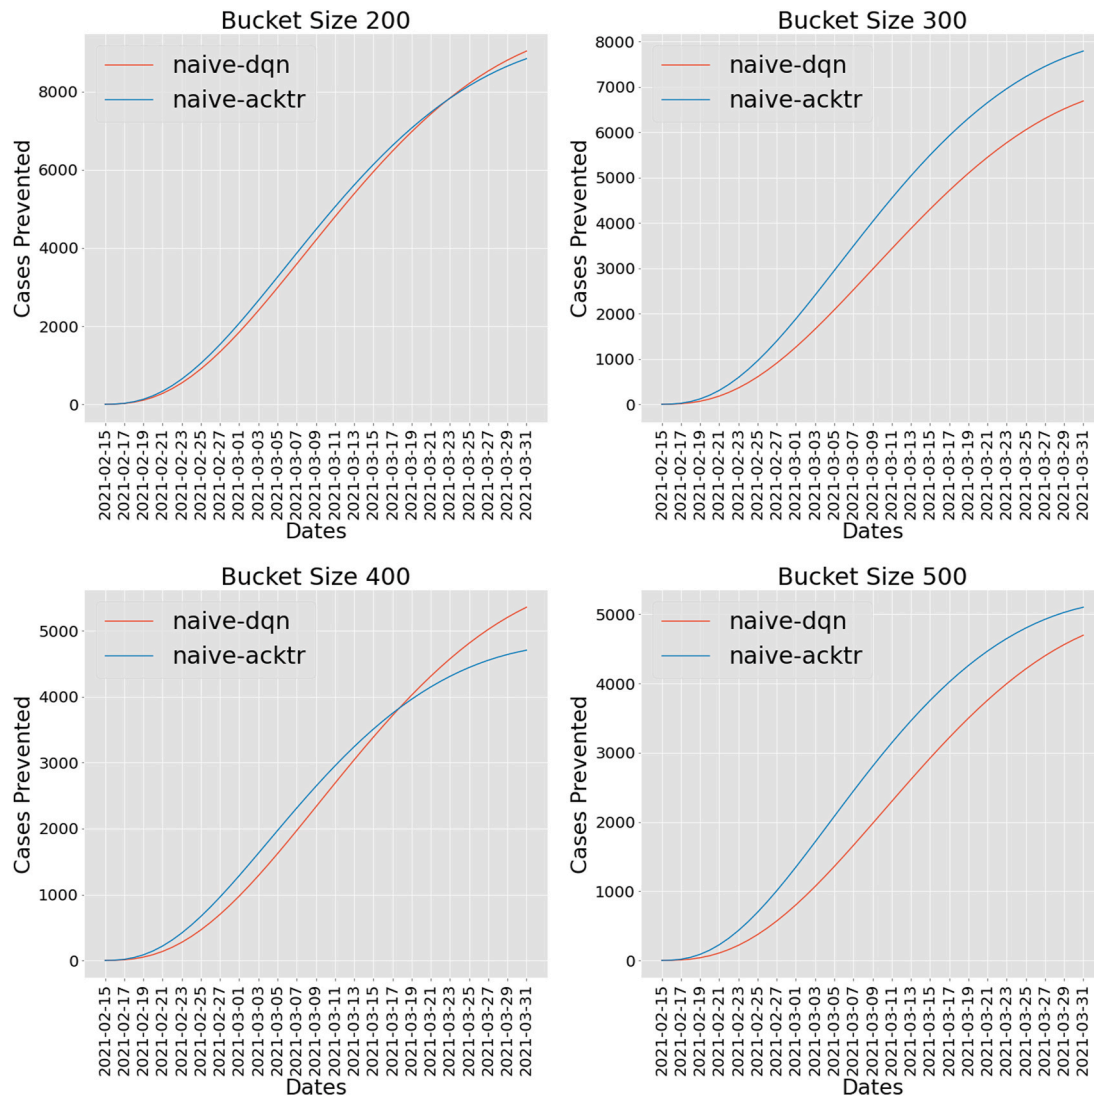


Fig. 4. Additional projected infection cases prevented in next 45 days by following VacSIM driven approach instead of Naive approach.

numerous logistical challenges (supply chain, legal agreements, quality control and application to name a few) which might slow down the distribution process. To circumvent these challenges, a strong decision support system is needed [30]. Reinforcement Learning (RL) based decision models [31] have been built during COVID-19. In this paper, we have developed a novel distribution policy model, *VacSIM* using Reinforcement Learning. We have pipelined an Actor-Critic using Kronecker-Factored Trust Region (ACKTR) model/Deep Q-Networks (DQN) model and Contextual Bandit model in a feedforward way such that the output (Action and Rewards) of ACKTR/DQN model are fed into the Contextual Bandit model in order to provide a sensible Context

comprising of actions and rewards. Contextual Bandits then optimize the policy considering demographic metrics such as the population of the State, Overall Hospital Beds, Overall ICU Beds, Ventilators and time series-based characteristics of the *COVID-19* spread (susceptible population, recovery rate, death rate, total infected cases) as Context. While distributing the vaccine, identifying the part of the population who needs it the most is a challenging task, and in our case, we addressed it by the usage of the aforementioned Context. Rather than using the present-day count of infected and susceptible people, we have used SEIR-based projection, which makes our predicted policy more robust and trustworthy. The percentage of vaccine distribution was the final

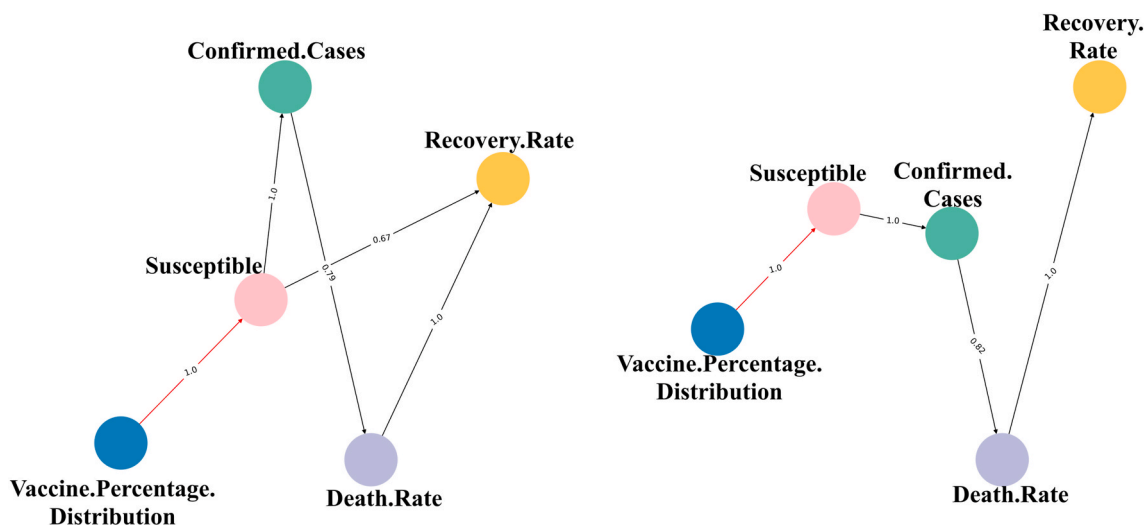


Fig. 5. Ensemble averaged causal structure of the Bayesian network obtained from 501 bootstraps, using Hill Climbing optimizer for AIC (left) and BIC (right) as scoring functions. Vaccine Percentage obtained from the model was observed as a parent node of Susceptible cases, thus indicating the causality preserving nature of VacSIM simulations.

output of VacSIM. From VacSIM, Maharashtra, the most vulnerable state for COVID-19, has received the maximum amount of vaccine, which showed the model could understand the context. Finally we proposed two approaches to evaluate such models 1) SEIR-based Simulation-based approach, 2) Bayesian Network-based Causality driven approach. In simulation-based approaches, we have projected the COVID-19 cases after the vaccination. Our results indicate that the VacSIM additionally reduces the 9039 cases compared to the naive-based approach. And from the Bayesian network, we found that vaccine distribution is most causally associated with susceptibility, which again showed that VacSIM output followed the design of model set-up and variability of context. Our model design and evaluation are entirely novel. Any other resource allocation algorithm like Genetic Algorithm, Simulated annealing, Hill climbing is reward independent. A comparison of policy with these algorithms are shown in our dashboard (<http://vacsim.tavlab.iitd.edu.in:8000/>). VacSIM outperformed these models when compared using SEIR based simulation approach. It is necessary to update models with new influences on a regular basis in order to capture varying trends occurring with emerging variants of COVID-19. In order to do so for different settings, our model is extendable and context independent, therefore will require re-parameterization to adapt to the new setting and learn optimal policy accordingly. Finally, we have provided open-source code for our pipeline which will enable testing of our claims by other researchers.

5. Limitations and future work

VacSIM may have some limitations shared by all Reinforcement Learning models, i.e. the transparency of their learning process and the explainability of their decisions. The development of VacSIM has been carried out while observing the pandemic in the past few months. However, the dynamic nature of pandemic may require a change in actions, thus calling for common sense working alongside artificial intelligence. A third limitation is the potential dependency between COVID-19 trajectory variables across states as these may not be completely independent especially after the lifting of lockdown. However, parametrizing these dependencies is a challenge and we reason that local context of the states is most strongly influenced by its own past values. As mentioned before, the availability of authentic data that is granular in terms of features covered is a major challenge in any kind of modelling work associated with the pandemic. By virtue of the same, we have given utmost priority to achieving herd immunity and decreasing

the susceptible population as much as possible, given that proper data capturing of those features is available. The same is reflected through our choice of reward function in this work. For future work, if a more reliable and extensive data source is available describing the information of different cohorts of the Indian population on a more atomic level (such as information regarding the comorbidities of a certain region), then such relevant features should be taken into account while designing the reward function. Similarly, designing reward functions with features which are important for the efficient distribution of scarce resources like ventilators, drugs on the availability of extensive reliable data sources can be done through our framework depending on the observation and action space.

It should also be noted that while we assume certain vaccine characteristics, such as efficacy and gap between both doses, these are hyperparameters that can be changed depending on the availability of exact reliable data. To date, there does not exist a single well-proven statistic for the dose-wise efficacy of Covaxin and Covishield (vaccines available in India currently). Even the gap between both doses has been changed repeatedly by the Government, the latest being 12–16 weeks⁷ for Covishield. In the emergence and spread of various mutations of COVID-19 virus, and responding variants of vaccines entering the market, such assumptions have been taken to make the pipeline as generalised as possible, so that when required, different pieces of information can be plugged into our framework and adjusted to identify the efficiency for different resources and context. Given these limitations, we have attempted to carefully design the VacSIM pipeline and associated reward function by making realistic assumptions which can be helpful in testing such a framework before moving to on-ground deployment.

To truly facilitate our ultimate goal of on-ground deployment of VacSIM, Plug-n-Play(PnP) pipeline can be created that gives policy-makers the flexibility to provide varying input contexts, such as those collected on the district level. This platform would be helpful in testing our framework's efficiency in a variety of settings, across many data points, observation spaces and even reward functions. In conclusion, we believe that artificial intelligence has a role to play in the optimal distribution of resources such as vaccines, syringes, drugs, personal protective equipment (PPEs), ICUs etc. The second wave of COVID-19 in India has also highlighted the importance of the same and our open

⁷ <https://www.bmj.com/content/372/bmj.n18>.

source VacSIM framework is a potential step towards enabling use of artificial intelligence by researchers and policymakers.

6. Conclusion

We provide a novel, open-source and extensible solution vaccine allocation problem that policymakers and researchers may refer to while making decisions. This feedforward network can be adapted and used for optimal allocation of various essential resources in varying scenarios and context as the Reinforcement Learning part of VacSIM can always be retrained to adapt to spatial and temporal variations by accommodating the latest values of variables like recovery rate and age distribution. Since there is some evidence that COVID-19 infections tend to show oscillatory patterns on shorter timescales [32], the key goal of our approach to inform an agile response to the problem of vaccine distribution has been fulfilled through this work. Although the present results are on five states, VacSIM will be extended to all the states of India.

Availability of data and material

All the data and source code are available at <https://github.com/tavlab-iiitd/VacSIM>.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

Not Applicable.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.ibmed.2022.100060>.

References

- [1] Greenwood B. The contribution of vaccination to global health: past, present and future. *Phil Trans Biol Sci Jun.* 2014;369 [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4024226/>. 1645.
- [2] [Online]. Available: Smallpox vaccines <https://www.who.int/news-room/feature-stories/detail/smallpox-vaccines>.
- [3] Rella SA, Kulikova YA, Dermitzakis ET, Kondrashov FA. Rates of SARS-CoV-2 transmission and vaccination impact the fate of vaccine-resistant strains. *Sci Rep Jul.* 2021;11(1):15729. bandiera abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary atype: Research Publisher: Nature Publishing Group Subject_term: Population genetics;Viral infection Subject_term_id: population-genetics;viral-infection. [Online]. Available: <https://www.nature.com/articles/s41598-021-95025-3>.
- [4] White DB, Angus DC. A proposed lottery system to allocate scarce COVID-19 medications: promoting fairness and generating knowledge. *JAMA Jul.* 2020;324(4):329–30. publisher: American Medical Association. [Online]. Available: <https://jamanetwork.com/journals/jama/fullarticle/2767751>.
- [5] Khamsi R. If a coronavirus vaccine arrives, can the world make enough? *Nature Apr.* 2020;580(7805):578–80. number: 7805 Publisher: Nature Publishing Group. [Online]. Available: <https://www.nature.com/articles/d41586-020-01063-8>.
- [6] D. Foster, C. McGregor, and S. El-masri, "A survey of agent-based intelligent decision support systems to support clinical," in *Management and research*, 1st intl. Workshop on multi-agent systems for medicine, computational biology, and bioinformatics.
- [7] Deo S, Manurkar S, Krishnan S, Franz C. COVID-19 vaccine: development. *Access Distr. Indian Context* 2020;378:16.
- [8] Liu S, See KC, Ngiam KY, Celi LA, Sun X, Feng M. Reinforcement learning for clinical decision support in critical care: comprehensive review. *J Med Internet Res Jul.* 2020;22(7) [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7400046/>.
- [9] Saria S. Individualized sepsis treatment using reinforcement learning. *Nat Med Nov.* 2018;24(11):1641–2. number: 11 Publisher: Nature Publishing Group. [Online]. Available: <https://www.nature.com/articles/s41591-018-0253-x>.
- [10] A. Raghu, M. Komorowski, L. A. Celi, P. Szolovits, and M. Ghassemi, "Continuous state-space models for optimal sepsis treatment: a deep reinforcement learning approach," p. 17.
- [11] Komorowski M, Celi LA, Badawi O, Gordon AC, Faisal AA. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nat Med Nov.* 2018;24(11):1716–20.
- [12] Hester T, Vecerik M, Pietquin O, Lanctot M, Schaul T, Piot B, Horgan D, Quan J, Sendonaris A, Osband I, Dulac-Arnold G, Agapiou J, Leibo J, Gruslys A. Deep Q-learning from demonstrations. *Proc AAAI Conf Artif Intell Apr.* 2018;32(1). number: 1. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/11757>.
- [13] arXiv:1509.06461 [cs] van Hasselt H, Guez A, Silver D. Deep reinforcement learning with double Q-learning. arXiv: 1509.06461. [Online]. Available: <http://arxiv.org/abs/1509.06461>; Dec. 2015.
- [14] Williams RJ. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach Learn May* 1992;8(3):229–56. <https://doi.org/10.1007/BF00992696> [Online]. Available:..
- [15] arXiv:1506.02438 [cs] Schulman J, Moritz P, Levine S, Jordan M, Abbeel P. High-dimensional continuous control using generalized advantage estimation. arXiv: 1506.02438. [Online]. Available: <http://arxiv.org/abs/1506.02438>; Oct. 2018.
- [16] arXiv:1602.01783 [cs] Mnih V, Badia AP, Mirza M, Graves A, Lillicrap TP, Harley T, Silver D, Kavukcuoglu K. Asynchronous methods for deep reinforcement learning. arXiv: 1602.01783. [Online]. Available: <http://arxiv.org/abs/1602.01783>; Jun. 2016.
- [17] arXiv:1708.05144 [cs] Wu Y, Mansimov E, Liao S, Grosse R, Ba J. Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation. arXiv: 1708.05144. [Online]. Available: <http://arxiv.org/abs/1708.05144>; Aug. 2017.
- [18] arXiv:1705.01064 [math, stat] Ly A, Marsman M, Verhagen J, Grasman R, Wagenmakers E-J. A tutorial on Fisher information. arXiv: 1705.01064. [Online]. Available: <http://arxiv.org/abs/1705.01064>; Oct. 2017.
- [19] arXiv:1412.1193 [cs, stat], Jun. 2020 Martens J. New insights and perspectives on the natural gradient method. arXiv: 1412.1193. [Online]. Available: <http://arxiv.org/abs/1412.1193>.
- [20] arXiv:1807.09809 [cs, stat] Collier M, Llorens HU. Deep contextual multi-armed bandits. arXiv: 1807.09809. [Online]. Available: <http://arxiv.org/abs/1807.09809>; Jul. 2018.
- [21] arXiv:1802.09127 [cs, stat] Riquelme C, Tucker G, Snoek J. Deep bayesian bandits showdown: an empirical comparison of bayesian deep networks for thompson sampling. arXiv: 1802.09127. [Online]. Available: <http://arxiv.org/abs/1802.09127>; Feb. 2018.
- [22] [Online]. Available: Vowpal Wabbit <https://vowpalwabbit.org/>.
- [23] Li MY, Muldowney JS. Global stability for the SEIR model in epidemiology. *Math Biosci Feb.* 1995;125(2):155–64 [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0025556495927565>.
- [24] Census of India Website. Office of the registrar general & census commissioner, India [Online]. Available: <https://censusindia.gov.in/2011-common/introductionToNpr.html>.
- [25] [Online]. Available: Development Data Lab <http://www.devdata.org/shrug>.
- [26] OpenAI. Gym: a toolkit for developing and comparing reinforcement learning algorithms [Online]. Available: <https://gym.openai.com>.
- [27] Polack FP, Thomas SJ, Kitchin N, Absalon J, Gurtman A, Lockhart S, Perez JL, Pérez Marc G, Moreira ED, Zerbini C, Bailey R, Swanson KA, Roychoudhury S, Koury K, Li P, Kalina WV, Cooper D, Frenck RW, Hammitt LL, Türeci z, Nell H, Schaefer A, Ünal S, Tresnan DB, Mather S, Dormitzer PR, Şahin U, Jansen KU, Gruber WC. Safety and efficacy of the BNT162b2 mRNA covid-19 vaccine. *N Engl J Med* 2020;383(27):2603–15. <https://doi.org/10.1056/NEJMoa2034577>. publisher: Massachusetts Medical Society eprint, [Online]. Available:..
- [28] Gámez JA, Mateo JL, Puerta JM. Learning Bayesian networks by hill climbing: efficient methods based on progressive restriction of the neighborhood. *Data Min Knowl Discov* 2011;22(1):106–48. <https://doi.org/10.1007/s10618-010-0178-6> [Online]. Available:..
- [29] Sethi T, Maheshwari S. wiseR: a shiny application for end-to-end bayesian decision network analysis and web-deployment. Nov. 2018. [Online]. Available: <https://CRAN.R-project.org/package=wiseR>.
- [30] Avorn J, Kesselheim A. Regulatory decision-making on COVID-19 vaccines during a public health emergency. *JAMA Aug.* 2020 [Online]. Available: <https://jamanetwork.com/journals/jama/fullarticle/2770244>.
- [31] Kwak GH, Ling L, Hui P. Deep reinforcement learning approaches for global public health strategies for COVID-19 pandemic. *PLoS One May* 2021;16(5):e0251550. publisher: Public Library of Science. [Online]. Available: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0251550>.
- [32] Bukhari Q, Jameel Y, Massaro JM, D'Agostino RB, Sr, Khan S. Periodic oscillations in daily reported infections and deaths for coronavirus disease 2019. *JAMA Netw Open Aug.* 2020;3(8):e2017521. <https://doi.org/10.1001/jamanetworkopen.2020.17521> [Online]. Available:..