Application of Probabilistic Sampling Algorithm in COVID-19 Medical Resources Allocation

Yue He *

Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, United States

* Corresponding Author Email: yhe23@uw.edu

Abstract. During the COVID-19 pandemic, the virus responsible for this global pandemic is highly contagious. This caused shortage in critical medical resources including vaccines, ventilators, and viral test kits. Thus, the health-care system around the world encountered heavy pressure. In the case of handling the surge in infection population and frequent changes in medical conditions, deterministic approaches cannot perform as well as probabilistic sampling algorithms that incorporate dynamic mechanisms to adapt to environments with high uncertainty. This paper reviewed several studies exploring the application of probabilistic sampling approaches – involving Thompson Sampling, Bayesian inference, and multi-armed bandits – in pandemic medical resource triage. By comparing with the data for static allocation policies, these studies revealed that such methods consistently made improvement in beneficial use of the medical resources among diverse types of resources. Though their developed approaches were constrained by common limitations for probabilistic frameworks, such as high reliance on the accuracy and computational efficiency of environmental estimations, these works suggested one promising direction for formulating allocation in the future pandemics: combining existing medical scoring systems with the probabilistic frameworks while exploring auxiliary tools, such as ones particularly enhancing interpretability or pre-training the data pipeline.

Keywords: Probabilistic Sampling Algorithm, Thompson Sampling, Multi-Armed Bandits, COVID-19 Resource Allocation, Adaptive Testing.

1. Introduction

During the COVID-19 pandemic, the high transmissibility of the SARS-CoV-2 virus and the resulting rapid surge in the infected population causes severe shortage in the medical resources such as ICU beds, mechanical ventilation, and other critical materials for medical device.

Although the ethical principles and allocation frameworks for scarce medical resources established in prior pandemics have been applied to the COVID-19 pandemic [1, 2], having the goal of maximizing the benefits of scarce medical resources in a short time, the diversity of involved factors, the broadness of impacted areas, and the necessity of ensuring fairness still posed significant challenges for the clinicians in making allocation decisions. Therefore, utilizing developed schemes to allocate those scarce medical resources is one effective and necessary approach to make better exploitation of the resources and save more lives in such a severe situation.

However, as one of the two fundamental problem-solving approaches, deterministic allocation could fail in the pandemic. Because of their no reliance on randomness, the algorithms of this type follow fixed execution steps, do not incorporate any random or pseudo-random components in their operation, and produce predictable output for similar input. This impedes it from adapting well to the evolving conditions during the pandemic.

Among the deterministic algorithms, the Explore-Then-Commit (ETC) algorithm operates in two distinct phases: it firstly samples each option a predetermined number of times, then commits to the option with highest estimated reward. The Upper Confidence Bound (UCB) algorithm dynamically balances exploration and exploitation through selecting the option with the highest upper confidence bound, which is computed using predefined function and accounts for both average reward and uncertainty. Regardless of the incorporation of dynamic update in estimated value for each option, the fixed execution components in the deterministic algorithms cause their unsuitability for environments with high instability and uncertainty. Consequently, deterministic allocation schemes often underperform in highly stochastic environments like a pandemic, where demand and outcomes evolve unpredictably.

In contrast, the probabilistic and Bayesian algorithms that can dynamically adapt to the changing environment while maintaining good performance are more suitable to implement during the pandemic. In this paper, we would study the application of probabilistic sampling algorithms - including Thomspon-Sampling (TS) algorithm, Bayesian inference, and multi-armed bandits - in the context of medical resource allocation in the COVID-19 pandemic. This paper aims to synthesis the related findings, compare the performance of the probabilistic sampling algorithms, and discuss their limitations or future improvements.

2. Background

Probabilistic sampling algorithms incorporate randomization into their decision-making process, dynamically balancing exploration and exploitation. The stochastic component helps reduce bias, improve representativeness, and achieve good performance even in environments with significant uncertainty or complicated factors. These algorithms typically combine prior knowledge with probabilistic models, which enables the weighing of multiple possibilities, and update their inference values with new observed evidence. By iteratively refining the understanding of the system, they attempt to build a statistical representation of potential outcomes and quantify uncertainty in their estimation procedure. This probabilistic characterization allows the decision-making to become more adaptive and data-driven, rather than relying on or sticking to fixed hypothesis.

Among the probabilistic approaches to solving sequential problems, one of the most widely applied and comprehensively studied is the Thompson-Sampling (TS) algorithm. TS samples from the posterior distribution for each option to make the maximizing-reward choice and updates the distribution (e.g., incrementing success/failure counts in a Beta-Bernoulli model) based on observed data, which naturally balances the exploration of uncertain options and the exploitation of high-performing ones. This scheme's simplicity for implementation and strong empirical performance makes it a strong candidate in real-world cases, including the COVID-19 pandemic medical resource allocation, where requires adaptive and uncertainty-aware decision-making.

In a pandemic resource allocation setting, probabilistic sampling methods allow the scarce resources, such as vaccines and ICU beds, to be adaptively assigned to patients based on evolving demand, supplies, and effectiveness. For instance, TS enables such adaptivity by continuously learning from new available outcomes and making choices based on samples generated from updated estimations.

In the following sections, this paper reviews recent studies that applied TS and other probabilistic sampling methods to allocate medical resources during the COVID-19 pandemic, focusing on their advantages, limitations, and lessons for future epidemic processing.

3. Application in COVID-19 Medical Resource Allocation

3.1. Vaccine Allocation

In pandemics, especially the high contagious ones, vaccines are important resources because they can safely train people's immune system and protect people from being infected by the illness, which is essential to the control of susceptible population and reduction in cases and mortality. For COVID-19, since the SARS-CoV-2 was one novel virus with high mutability, the lack of evidence in immune response to this paroxysm and the rapid emergence of viral variants with immune-escape capabilities, which largely reduced the effectiveness of first-generation vaccines, posed challenges in vaccine R&D. The resulting high complexity in allocation of vaccines motivated the use of adaptive, data-drive approaches such as probabilistic sampling algorithms.

Rey et al. (2021) proposed a Thompson Sampling-based optimization policy for vaccine allocation in the COVID-19 pandemic, which learns the mean effectiveness of vaccines and dynamically reallocated doses over time. In their constructed modeling, the authors set agents of the main countries as the interconnected nodes in a mobility network and collected testing results under multiple scenarios. Compared to population-proportional allocation strategies, their proposed method achieved great reduction in infections and deaths [3]. Thus, for probabilistic sampling algorithms like TS, the findings in this paper demonstrated their adaptivity to the evolving pandemic conditions, particularly the changes in vaccine efficacy.

Expanding on this, some studies committed to handling remaining issues for vaccine allocation, such as the time delay in vaccination feedback, which is a common challenge when time gap between vaccination and the outcome of viral testing occurs. As the disease prevalence was simultaneously evolving, to address the influences of such time delay, Wu and Wager (2022) suggested one Partial Likelihood Thompson Sampling (PLTS) method to determine most effective vaccine more quickly. According to their experiment results based on semi-synthetic COVID-19 data, PLTS exhibited better speed and accuracy in identifying the existing vaccine with highest efficacy over standard randomized trails [4]. In addition, Cimpean et al. (2023) proposed a Bayesian m-top exploration approach that learns multiple top-performing vaccination policies at the same time [5]. This framework enables faster and better matching between vaccine types and recipients. These works together indicated the potential of probabilistic sampling approaches for more accurate identification of most effective vaccination policy, generally better performance under uncertainty.

Supported by the improvement observed in epidemic data, these studies showed the value of probabilistic sampling algorithms in managing vaccine allocation. In the next section, we will examine the application of similar approaches in allocating other scarce and important medical resources.

3.2. Ventilator Allocation

Reviewing modern infectious disease history, in the pandemics caused by respiratory viruses, the shortage in ventilator supply was not rare. This is because the spread of these viruses can be easily achieved through the air and get aggravated in the airtight or densely populated areas, arising the need for increasing ventilation and avoiding crowds [6]. Due to the unprecedentedly rapid spread of the SARS-CoV-2 virus, which causes the COVID-19, the health-care systems were globally overwhelmed by the imbalance between the exponentially increased infections and the availability of medical ventilators and intensive care units (ICU) [7]. As a result, making proper rationing based on complicated clinical and ethical factors is vital to provide high-quality treatment to more patients and save more lives.

Anderson et al. (2023) integrated predictions of survival probability and ventilator length-of-use in their proposed probabilistic scheme for ventilator allocation, incremental survival probability per length-of-use (ISP-LU) [8]. This framework sampled from predicted outcome distributions and aimed to maximize the marginal survival benefit of ventilator use. It made a Pareto-improvement in the simulations over the methods strictly following Sequential Organ Failure Assessment (SOFA) score. In another study of Grand-Clément et al. (2021), the researchers formulated the ventilator triage problem as a Markov Decision Process (MDP) [9]. Since MDP involves decision-making when state-transition probability is uncertain, the built modeling incorporates uncertainty in patient outcome and medical system dynamics. It also uses dynamic programming recursion to refine estimates for each decision. This algorithm achieved significant reduction in the deaths due to ventilator shortage in the simulated surge scenarios.

These studies include uncertainty information, such as patient diversity and shifting demand, to propose probabilistic approaches. The developed methods update predictions of the medical system state continuously. They then adjust allocation policy to maximize the beneficial use of resources at each decision point. This presents the superior performance of probabilistic sampling algorithms in ventilator triage.

3.3. Testing Resource Allocation

Another critical challenge during the COVID-19 pandemic is the allocation of limited viral testing resources. Especially in the early outbreak phase of the pandemic, as viral test kits were undeveloped and cannot be mass produced, the detection of positive cases was even more difficult. To maximize outbreak detection and minimize resource waste, Grushka-Cohen et al. (2020) designed a Multi-Armed Bandit framework that made allocation based on risk scores of individuals and reserved a portion of capacity for randomly testing [10]. Within a predefined time period, the framework concluded shared features of the positive cases and modified risk scoring rules adaptively. Their simulation results revealed that this probabilistic approach detected the outbreak of pandemic earlier, and captured more positive cases than static schemes using equivalent testing resources.

Combing the findings of these studies on vaccine, ventilator, testing resource allocation, probabilistic sampling algorithms consistently outperform static policies across diverse medical resource allocation.

4. Discussion

In this paper, the reviewed studies show probabilistic sampling algorithms can well allocate multiple types of medical resources during COVID-19. This is because they have two major advantages: adaptivity to changes in medical conditions, and exploration-exploitation balance in decision-making. When allocating vaccine, Thompson Sampling algorithm and its variants quickly identified the most effective vaccine and avoided as much waste as possible. The same results can be found in allocating ventilator and testing resource. The ISP-LU approach suggested by Anderson et al. increased average daily survivors by 23.15% over fixed SOFA-based schemes [8]. Also, the Multi-Armed Bandit model offered by Grushka-Cohen et al. detected over 90% positive cases using 70% of the test cost [10]. Its identification is much more efficient than the uniform testing policy. Such results verify the strong performance of probabilistic sampling algorithms in uncertain environment like a pandemic.

Simultaneously, the algorithms still face some challenges in application. One is their high reliance on the accuracy of input data. For example, Thompson Sampling algorithm highly depends on the prior distributions for each option. Incomplete case counts would make allocation decisions ineffective. Another challenge is from the frequent re-evaluation of pandemic values in these algorithms. This requires the estimation methods to be computationally efficient. Moreover, the actual deployment requires coordination among multiple sectors. The policy execution efficiency would be greatly delayed by esoteric output.

Even having the limitations, these studies still give some directions to make probabilistic sampling algorithms perform better. Firstly, the research by Anderson et al. offers one strategy: combining established scoring systems with probabilistic frameworks [8]. This method takes advantages of both existing rules and dynamic mechanisms. Secondly, as the ethical principles have not been fully implemented in the frameworks, their extensions that take extra factors into account should be explored. Thirdly, it's feasible to build data pipelines and simulation environments in advance. In the future pandemics, effective allocation algorithms can be rapidly developed and utilized to respond the potential scarcity in medical resources.

5. Conclusion

This paper reviewed the application of probabilisitc sampling algorithms, including Thompson Sampling, Bayesian inference, and multi-armed bandits, for allocating multiple types of scarce medical resources during the COVID-19 pandemic. Across the cases, these adaptive and uncertainty-aware approaches consistently exhibited better performance over static policies, through updating inference values based on real-time data and making benefit-maximizing decisions. In particular, these methods led to faster identification of effective vaccines, higher survival rates for ventilator patients, and earlier outbreak detection through targeted testing.

While challenges remain—such as the need for accurate data, computational efficiency, and interpretable outputs—the reviewed studies highlight a promising direction for future epidemic response. Combining established clinical scoring systems with adaptive frameworks, adding fairness-aware mechanisms, and preparing simulation environments in advance can assist the deployment of these startegies in the future pandemics.

References

[1] Kirkpatrick James N, et al. Scarce-resource allocation and patient triage during the COVID-19 pandemic: JACC review topic of the week. Journal of the American College of Cardiology, 2020, 76(1): 85-92.

[2] Emanuel Ezekiel J, et al. Fair allocation of scarce medical resources in the time of COVID-19. The New England Journal of Medicine, 2020, 382(21): 2049-2055.

[3] Rey David, Hammad Ahmed W, Saberi Meead. Vaccine allocation policy optimization and budget sharing mechanism using Thompson sampling. arXiv:2109.10004, 2021.

[4] Wu Han, Wager Stefan. Partial likelihood Thompson sampling. arXiv:2203.00820, 2022.

[5] Cimpean Alexandra, Verstraeten Timothy, Willem Lander, Hens Niel, Nowé Ann, Libin Pieter. Evaluating COVID-19 vaccine allocation policies using Bayesian m-top exploration. arXiv:2301.12822, 2023.

[6] United States Environmental Protection Agency. Ventilation and respiratory viruses. EPA Indoor Air Quality Resources, 2021.

[7] Santini A, Messina A, Costantini E, Protti A, Cecconi M. COVID-19: Dealing with ventilator shortage. Current Opinion in Critical Care, 2022, 28(6): 652-659.

[8] Anderson R, Orenstein E, Farias V F. Rationing scarce healthcare capacity: A study of the ventilator allocation guidelines during the COVID-19 pandemic. Production and Operations Management, 2023.

[9] Grand-Clément Julien, Goh You Hui, Chan Carri, Goyal Vineet, Chuang Elizabeth. Interpretable machine learning for resource allocation with application to ventilator triage. arXiv:2110.10994, 2021.

[10] Grushka-Cohen Hagit, Cohen Raphael, Shapira Bracha, Moran-Gilad Jacob, Rokach Lior. A framework for optimizing COVID-19 testing policy using a multi-armed bandit approach. arXiv:2007.14805, 2020.