Using Markov Decision Processes for Medical Resource Allocation Decisions

Dennis Li
University of Maryland, College Park
dli12348@umd.edu

Diana Chou
University of Maryland, College Park
dchou@cs.umd.edu

Abstract

This paper discusses an optimized approach to decision making in regards to resource allocation within health care systems and the medical industry. Decision making in medical resource management is often sequential and uncertain in nature. Markov Decision Processes (MDPs) are used to model decision-making in sequential environments and can be used as a technique for solving decision problems. MDPs allow a decision holder, to make the best possible decisions in the environment they inhabit, which changed state in response to choices of action made by the agent. We will be describing our MDP model and illustrating the performance of our design with data focusing on health care professionals as well as medical supplies. We will be analyzing our evaluations and offer a promising approach towards an optimal policy focusing on increased measures of health. We will also discuss the challenges in addressing complex resource allocation problems that arise in healthcare settings.

Keywords

Markov decision processes, health care systems, optimal resource allocation, medical decision making

1 Introduction

The health care system in the United States consists of several moving parts that are led by stochastic decision-making. Such a system often causes sequential decision problems to arise. When determining the selected amount and destination for a specific set of resources, a decision holder must consider factors such as the demographic and statistical analysis of the location, such as supply and demand, which can constantly change. One source of uncertainty is the level of health within a population. In the case of a hospital setting that has an average amount of daily foot traffic, there is some amount of chance where a sudden influx of admitted patients can occur at any given point in time. Other sources of uncertainty could deal with the availability of scarce resources or the limitations of resources set by standards and regulations. Imagine hospital patients all needing the same type of medical supplies or treatments over a short period of time, however, given the nature of the type of resource, the amount was not properly allocated to that region. In the current healthcare system, most of these allocation decisions are made using heuristic strategies, in which solutions are practical and sufficient for short-term goals but not necessarily optimal.

In this work, we explore Markov Decision Processes, a model used for sequential decision making, as an approach for developing an optimal policy in resource allocation. We are particularly interested in how allocating sets of resources to various geographical regions can impact different measures of health. We focus on life expectancy, years of potential life lost, child and infant mortality, as well as quality of life for optimization. We express this problem through a Markov Decision Process (MDP) with attributes for each state and action for the agent, or allocation decision-maker. Our steps towards this solution involve the following:

1. We determine the framework for our MDP, the components involving the agent and the environment, and the method for finding an optimal decision policy.
2. We develop a design for our MDP, where the agent is the decision maker and the states are the resources.
3. We evaluate our MDP model with existing data sets containing information on specific allocations and the resulting rewards. Experiments to be conducted will focus on a set of specific health measures.

The rest of the paper is organized as follows. Section 2 describes the framework for Markov Decision Processes and our MDP design. Section 3 describes the data we will be testing our model on. In section 4, we run our experiments and evaluate our results with different measures of health. We proceed to discuss our findings in section 5, as well as related works in section 6. We then conclude the paper is section 7.

2 Markov Decision Processes

In this section, we provide a brief overview of MDPs before detailing the structure of our design. A Markov Decision Process (MDP) is a decision making framework for modeling dynamic systems under stochastic and uncertain environments. An MDP binds an agent's decisions through the proper definition of system states, which are variables that contain the information for making future decisions. An MDP model can cover a wide range of applications set by finite and infinite horizons (the total number of decisions we can make) as well as discrete and continuous time (how often and randomly decisions are made).

2.1 MDP Overview

Markov Decision Processes assume a finite number of states and actions. At each step, an agent observes a state and executes an action, incurring rewards to be maximized (or costs to be minimized). The rewards and succeeding state only depend on the current state and chosen action, and the transition occurs according to a known probability distribution based on the level of uncertainty in the environment. As a standard assumption of a Markov process, we assume that we have perfect knowledge of the current state so that future transitions and rewards are independent of the past.

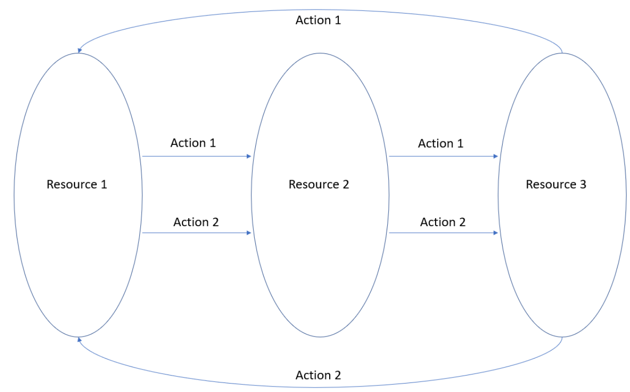

Fig. 1: Representation of the structure of our MDP with 3 states and 2 actions. Overall structure will remain the same while number of states and number of actions can be changed freely

2.2 MDP Definition

A Markov Decision Process is defined by a tuple (S, A, T, R), where:

1. S is the set of defined states and for every state $s \in S$
2. A is the set of all actions or decisions and A_s is the are the available actions at state s
3. T is a a known probability transition function: $T(s, a, s')$ where $s \in S, a \in A, a' \in S$ is $P(s'|s, a)$
4. R is the reward function $R(s, a, s')$ which is the reward at the current step
   A discount factor, γ, affecting how the reward is counted can also be factored at each step, to motivate the agent in favor of taking actions earlier.

In the case of infinite utilities, there are a couple of solutions that allow the system to converge. Applying a finite horizon allows for termination after a fixed T steps and gives non-stationary policies (π depends on the time left). Discounting $(0 < \gamma < 1)$ affects rewards to lose life over time (a smaller γ means a smaller horizon). Also, absorbing states can guarantee that for every policy, a terminal state will eventually be reached.

2.3 MDP Design

The states of our MDP will be the resources for which an allocation decision will be made. Each state will correspond to a different resource. So for a set of resources $R = \{R_1, R_2, R_3, ..., R_n\}$, the states of our MDP will be $S = \{S_1, S_2, S_3, ..., S_n\}$ where S_i is the state corresponding to the resource R_i.

For the actions, we define each action to instead represent a class of smaller individual actions, which we will hereafter refer to as action classes. We define action classes to mean a range which encompasses some amount of smaller individual actions. This was decided as a way to be able to reduce the total number of actions in our MDP. Had we attributed a single action to represent each data point as they were provided in the original data, we would have wound up with thousands, if not tens of thousands of different actions. So as a remedy to this, our classes of actions serve as a range to reduce the total size of our actions set. We had to further manipulate our set of actions as different resources had data points of different magnitudes, so for the same set of actions to be applicable to all states (resources), the final result for the actions became percentiles. To account for the various different magnitudes across resources, the raw values that a percentile correlates to will be based on the specific values for a resource. In other words, the action classes, and thus the corresponding percentiles, will be the same for all the resources, but what the percentiles are actually representing will be based on the data values of the specific resource. Put another way, the 100th percentile for any resource R will always be the maximum data value for that resource, with all other percentiles also being adjusted accordingly. More formally, we define our action set $A = \{(P_{1a}, P_{1b}), (P_{2a}, P_{2b}), (P_{3a}, P_{3b}), ...(P_{Xa}, P_{Xb})\}$, where P_{ia} represents the lower limit percentile for action class i inclusive, P_{ib} represents the upper limit percentile for action class i exclusive except the very last action class which will contain the upper limit which is the 100th percentile, and X represents the total number of action classes the MDP will have. Consequently, $\forall i \in [X], P_{ib} - P_{ia} = (100/X)$; so if we have $X = 5$ action classes, then every action class will represent a 20th percentile range, $[0th, 20th)$,$[20th, 40th)$, and so on. If we have $X = 100$ action classes, then every action class will represent a single percentile range, $[0th, 1st), [1st, 2nd), [2nd, 3rd)$, and so on. For making sure the action classes are representative of the actual resource data numbers, $\forall$ resources r, the 100th percentile will be max(all data for r), the 99th percentile will be 0.99$*$max(all data for r), and so on.

Figure 1 provides a simple representation of what our MDP looks like. All the actions for any state s leads to the same state s'. In other words, there is always only one possible state that can be reached from any given state.

Table 1: Sample reward data

	Resource 1	Resource 2	Total Reward
Geographical Region 1	Accounts for 20% of total resources in this region	Accounts for 80 % of total resources in this region	T1
Geographical Region 2	Accounts for 30% of total resources in this region	Accounts for 70% of total resources in this region	T2

We define our transition function such that $\forall s, s' \in S, a \in A, T(s,a,s') = [1]$. We define all actions to be deterministic based on the idea that when an agent makes a decision to allocate a amount for a resource x, they can be sure that taking such an action will guarantee a amount of resource x to be allocated. The loop structure of our MDP was simply a result of more convenient implementation.

The reward measure will be based on some measure of health. Different measures can be substituted in as the reward to see if using different measures results in different policies. For any measure, the reward for each action will be the average of the weighted values of that measure for their respective geographical regions. The weights will be based on what the percentage of the specific resource in question is in relation to all the resources under consideration. Consider the example in table 1. For an action acting on resource 1, the reward for this action will be $r_1 = avg(0.2*T1, 0.3*T2)$, and the reward for the same action on resource 2 would be $r_2 = avg(0.8*T1, 0.7*T2)$, assuming there are only 2 resources under consideration and the amounts of each resource across both regions fall into the same action class. This means that taking the given action class on resource 1 will result in a reward of r_1 and a reward of r_2 when taken on resource 2. This approach depends on the assumption that larger proportions of a resource produce larger effects in a geographical region and thus should carry more weight.

We will use a discount factor of $\gamma = 1$. This choice stems from the assumption that all decisions, no matter what point in time they were made, should have the same influence/importance. With $\gamma = 1$ and a loop structure, we resort to using a finite horizon with $T = |S|$ steps to guarantee convergence.

3 Data

The way our MDP determines the best allocation policy for a set of resources is by examining the impact that existing allocation decisions of those resources have on various measures of health. Given this, testing our MDP would require data that includes information on the specific allocation amounts of various different resources as well as information on various measures of health in the places that were allocated those resources. To this end, we utilized publicly available data sets which provided this information. The original resources data was distributed from [1], but a reorganized version was made available via [2], which was what we ultimately based our MDP experiments on. The health measure data was obtained via [3], but since the data was not provided in a manner that was readily usable for our MDP, we compiled the separate sets on the site to create our own working data set. One point to note is that this site is not the original source of the information; it is more an intermediary site which has aggregated data from various sources that include government agencies, medical associations, local reporting sites, and more. The specific sources for all the original information are available at [3]. Another facet of our data that is worth noting, and which will become more apparent in our experiments, is the imbalance of human vs non-human data. More specifically, there is simply much more data included on human resources. As a result, the policies we obtain will be much more reflective and applicable to human resources.

4 Experiments

We run our MDP through various allocation decisions of different groups of resources along with different measures of health. For each experiment, our MDP produces an allocation policy which, according to our design, defines the optimal allocation of each resource in consideration while maximizing the benefits for the given health measure. We set up simulations to obtain both national and state-level policies; national-level policies are based on an overview over all the available data with no consideration for origin of any data point, while state-level policies are achieved by separately considering the data in each state. In each experiment, the set of resources for which a policy is obtained was decided based on certain previous studies as well as our own judgements of which resources would, or should, have the most effect on the specific health measure. For example, if life-expectancy were the health measure in consideration, it would make more sense that the policy should deal with Geriatrics, the health care of the elderly, more so than with Neonatology, the medical care of newborn infants; the opposite would be true should the health measure in consideration be infant mortality.

4.1 Life Expectancy

We start our experiments with life-expectancy as the health measure as it is a common and important health measure that is easy to interpret and understand but which also has considerable implications in many other areas [4]. Most of the resources included in this experiment were based on our own judgement, with the major exception being that of Primary Care Physicians, where a previous study [5] suggests a positive association between increased primary care physician supply and life expectancy.

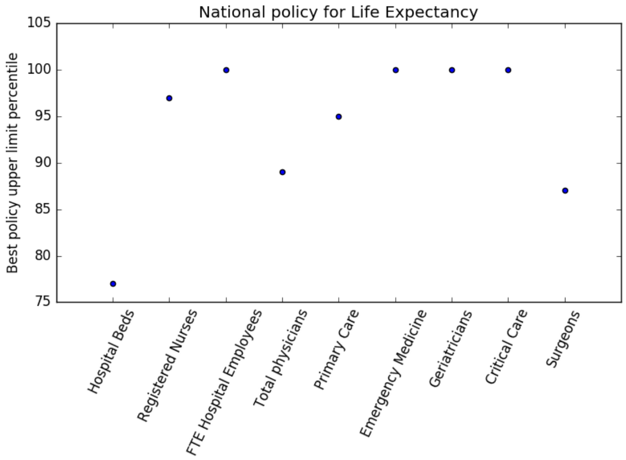

Fig. 2: National policy for life expectancy

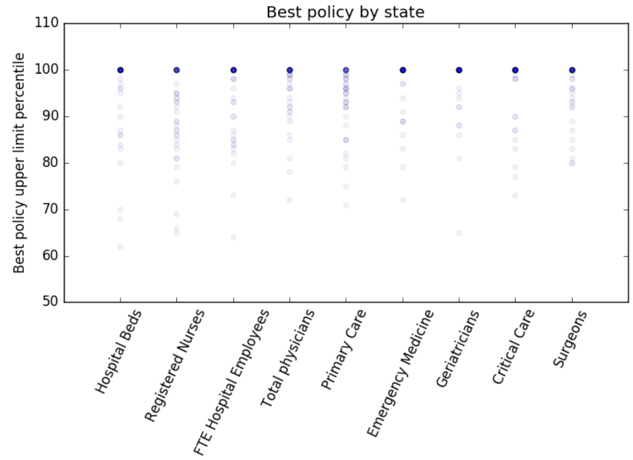

Fig. 3: Best policies that maximizes life expectancy for each state

A quick analysis of the national policy shown in figure 2 reveals the national policy suggests that maximal life expectancy is achieved when primary care physicians are allocated at the 95th percentile of what is currently being allocated, which seems to stay in line with the results obtained by [5]. Though this policy isn't exactly suggesting the same thing, it would stand to reason that the policy implies that places which are allocating, or have been allocated, below the 95th percentile would be able to increase life expectancy should the allocation be closer to the 95th percentile. On the other hand, this policy would also mean that places with allocations over the 95th percentile do not do any better, if not worse, which stands in contrast to the increasing life expectancy with increasing primary care physician supply of [5]. As such, while our findings would seem to confirm the positive association, it also suggests that there is a limit to this association. Such reasoning can also be applied to the other resources that have a recommended best policy of less than the 100th percentile. In short, these resources hit a limit in terms of how much they can help increase life expectancy as their allocation amounts are increased.

Conversely, this policy would suggest that in order to maximize life expectancy, resources such as FTE (Full Time Equivalent) Hospital Employees, Emergency Medicine Physicians, Geriatricians, and Critical Care Physicians should be maximized. In other words, as far as the data used in this paper goes, increasing the allocations for these resources will have a corresponding increase of life expectancy with no observable limit.

With that being said, it's also worth mentioning that because this policy is proposed at a national level, this kind of policy is subject and susceptible to large variations in measurements that may affect the final resulting policies. More specifically, because of how rewards are calculated, the rewards for national policies will be a weighted average over quite a large number of data points, and it is much easier for the many data points in such an average to skew the resulting average in one way or another. This may produce results that highlight negligible or nonexistent factors while ignoring other more pertinent factors. Such a possibility is not impossible and is something we explore by comparing this policy with one that would result from a more fine-grained examination, namely at the state level instead of the national level.

If we examine figure 3, it is immediately evident that there is no one-size-fits-all policy, which may have been the conclusion based on the national-policy. If we consider the state policies for primary care physicians in particular, the state policies would suggest something closer with the findings of [5] than what may be suggested by the national policy. For this specific resource, it seems more likely that the 95th percentile allocation as suggested by the national policy was merely a product of how it was generated, culminating with the 95th percentile having the highest average, which can also be interpreted as the 95th percentile having the highest expected gain for life expectancy. Instead, it seems the most common best policy for a state would be to maximize the allocation of primary care physicians, something which holds true for all the other resources as well.

However, even though the most common best policy for states in any given resource would be to maximize the allocation of that resource, this is clearly not the case for *every* state; in fact, there are quite a fair number of states for all the resources that achieved their maximal life expectancy with less, and some considerably less, than the maximal allocation of the resource, suggesting there is quite the difference across states in life expectancy gains for any percentile. Such differences could very well influence a policy obtained through a process like the national-level policy, which would explain the noticeable difference between the national policy and the individual state policies. But returning to the original observation that there is quite a fair amount of variation in the states' best policies, an immediate resulting inquiry would be why this is the case, and why is it as extensive as it is. While this is a very relevant consideration, our MDP, and the purposes of this paper, are not suited to researching the likely causes of this phenomenon. We do possess data on various potential causes but such investigation would require modifications to our MDP design or the construction of other models entirely, so we leave this as a future work.

The other major observation from comparing the state level policies with the national policy is the considerable disconnect between the two types of policies for certain resources. For this experiment in particular, this disconnect seems most pronounced between the national and state allocation policies for hospital beds. Looking more closely at the two types of policies for this resource, we see that the national policy for hospital beds is just above the 75th percentile, maybe around 76 or 77. However, from the states' policies, we see that virtually every single state has a policy above the national policy. Based on what we defined our rewards system to be and how rewards are calculated as specified in section 3, it would seem most intuitive that the national level policy would reflect an average, more or less, of the various state policies, as the actions and rewards for both types of policies are the same, only the rewards for national policies are obtained from averaging a larger set, one that spans more than one state, supposedly, as opposed to just averaging over a single state for state policies. Because of this, the fact that the national policy for this resource is nowhere close to being an average of the state policies is understandably suspicious. This is definitely something worth addressing, but we leave the specific discussion of this to the next experiment.

4.2 YPLL (Years of Potential Life Lost)

A second experiment we conduct uses Years of Potential Life Lost rate as the health measure to consider. This specific health measure brings with it a more representative measure to analyze premature mortality. The relevance and importance of this specific measure is in its ability to address the inability of other mortality data in painting a wholesome picture of mortality trends; specifically the fact that such mortality data are inevitably dominated by deaths of the elderly [6], and thus perhaps not as pertinent when attempting to focus on aspects such as preventable deaths.

For this experiment, we decided to use the same set of resources as the first experiment as we believe the same set of resources are also rather influential when addressing preventable deaths. We present the national level policy in figure 4 mostly to provide a sense of what allocations have the highest expected reward, which in this case is minimizing YPLL, due to the reasons specified in the previous experiment. The state-level policies demonstrate a comparable situation with its corresponding national-level policy as seen in the previous experiment. Considering how both life expectancy and YPLL involve preventing or delaying death in one way or another, it seems plausible that the policies for the same set of resources, at least at the state level, are rather similar. However, the main discussion point which will be the primary focus of this experiment discussion is the disconnect between national and state policies that was mentioned previously. The reason this discussion was left to this section despite being introduced earlier simply rests on the fact that

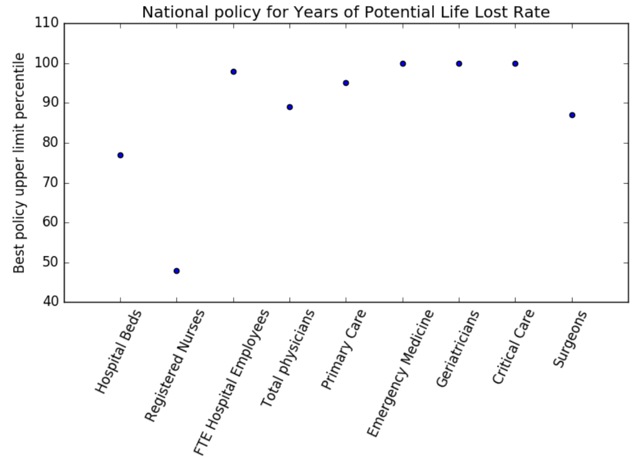

Fig. 4: National Policy to minimize Years of Potential Life Lost

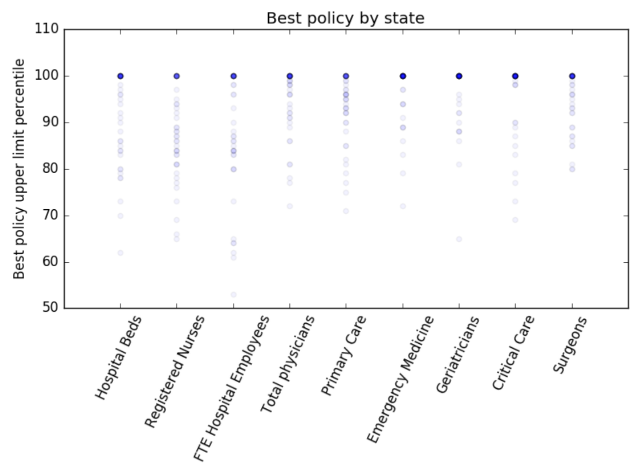

Fig. 5: Best policies that minimize YPLL for each state

the occurrence is much more pronounced in this experiment.

While the life expectancy experiment highlighted this fact with the hospital beds resource, for this experiment it is most apparent with the registered nurses resource, and to a higher degree as well. Looking at the national policy for this resource, we see that the recommended allocation percentile is just under 50. However, when looking at the state policies, there isn't even a single state that is recommended at or below the 60th percentile. We will show how this is possible before discussing potential ramifications.

Consider the example data set presented in table 2. When finding the national policy, the calculated rewards for the 100th, 90th, and 80th percentiles would all be $(100 + 10 + 10)/3 = 40$, while the calculated reward for the 50th percentile would be $(45 + 45 + 45)/3 = 45$. Afterwards, when choosing the percentile for best policy, the choices would be 100th percentile: reward 40, 90th percentile: reward 40, 80th percentile: reward 40, and 50th percentile: reward 45. As a result, the national policy for this resource would be the 50th percentile. And yet, when the states are

Table 2: Sample data set

State	Alloc. percentile	Reward Amount
1	100	100
1	90	10
1	80	10
1	50	45
2	100	10
2	90	100
2	80	10
2	50	45
3	100	10
3	90	10
3	80	100
3	50	45

considered individually, the best policy for state 1 would be the 100th percentile, state 2 would be 90th percentile, and state 3 would be 80th percentile, none of which are the same as the national policy. From this example, we see how variation across states can heavily influence national policy, painting quite the different picture and resulting in allocation decisions that would, in certain cases such as this example, clearly not be optimal. While such skew can also impact state-level policies, the variation would need to be much more extensive, but at that point, if such an allocation decision (percentile) carries with it that much uncertainty, perhaps it would be better to stick with a decision that may have a lower absolute gain but a much higher guarantee.

From this discussion, we believe it would be more prudent to consider allocation policies at a state-level while leaving the national level policy just as a means to designate the allocation with the highest expected gain. As such, for the remaining two experiments, we still present both types of policies, but analyses provided will take this review into account.

4.3 Other Experiments

We perform a few other experiments, all displayed in figure 6, as evidence of our MDPs ability to operate on a variety of different health measures. Our experiments on Child and Infant Mortality were motivated in part as a response to how YPLL is calculated [6], where certain measures such as deaths of younger age groups can have a sizeable impact. The resources we decided to consider for these two experiments were chosen based on observed leading causes of death worldwide for children under 5 [7]. From these experiments, we once again see a wide variety of allocation percentiles with the best expected results in the national policies while the state policies once again display a tendency to suggest there exists a reason a sizeable number of states are able to achieve maximal results with fewer resources.

Our final experiment involves using a self reported quality of life health measure, which seems to provide a considerable indication of mortality risk [8], a common theme in the other health measures we investigate. This time, there doesn't seem to be that large of a discrepancy between the national and state-level policies as compared to some of our other experiments, so although there is still quite a fair bit of variation between states as seen in the state policies, perhaps the variation across states for any given allocation percentile is not as extensive as other health measures. Although the retest reliability of self reporting [9] could be a potential reason, we will have to leave this as a future work as mentioned earlier.

5 Discussion

For the state-level policies across all the different health measures, two commonalities were the persistent nature of policies to cluster at the 100th percentile for all resources and the considerable variability of the policies despite this cluster. As common intuition would dictate that higher allocations of resources leads to better results, which in this case would mean larger improvements for any given health measure, the clustering of 100th percentile policy recommendations does not seem that surprising. But the same intuition would conflict with the large, and consistent, variation of allocation policies across states for all the resources across all health measures. If the intuition really was true, it alone would be hard pressed to explain such extensive varying resulting state policies. These state policies our MDP produced would suggest deciding on an allocation policy goes much farther beyond just "more is better". While there can be a whole host of different potential reason as to why some states are able to achieve the same results, if not better results, with considerably less resources, at this point we can only speculate as to what those causes may be. In fact, it may even be possible that the relatively lower recommended allocations is no more than a result of that specific state being exceptionally better in terms of the health measure to begin with, and thus not requiring as much resources to be allocated to it.

There are also some important points to consider, as discussed in [10], which cause us to contemplate priority in decision-making. The first point mentioning that efficiency has both ethical and economic importance. In the case where certain health systems are more efficient than others, if we want systems to meet more health needs, then we prefer more efficient health systems. A second point is that efficiency is not the only goal of health policy; health policy is not only concerned with the improvement of population health as a whole, but also with fairness in that distribution of health. Since it depends on subjective levels of reasoning, it is complex to determine how to favor better outcomes vs. fair chances. These cases bring up current policies that

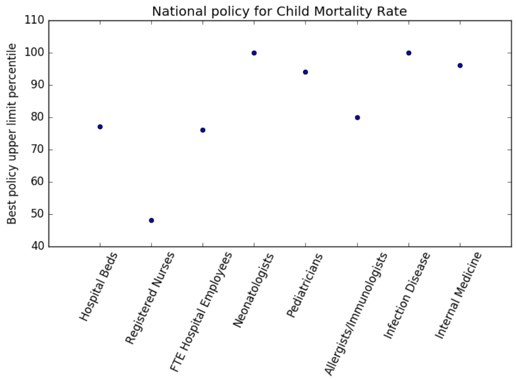

(a) National Policy for Child Mortality Rate

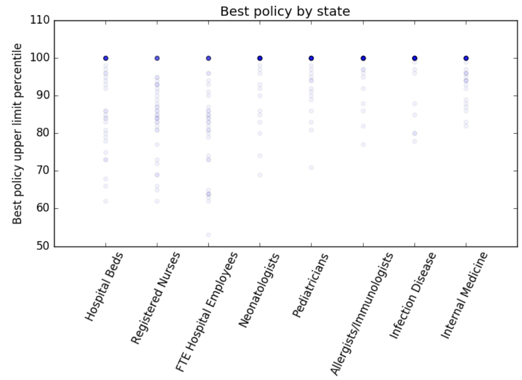

(b) State policies for Child Mortality Rate

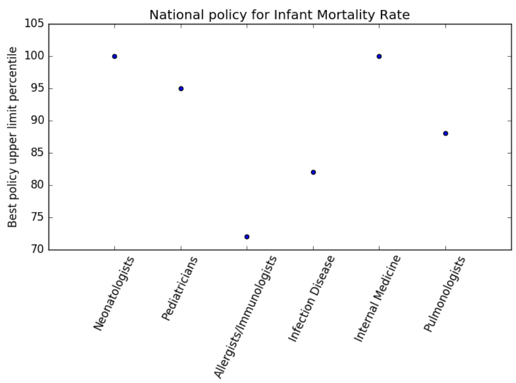

(c) National Policy for Infant Mortality Rate

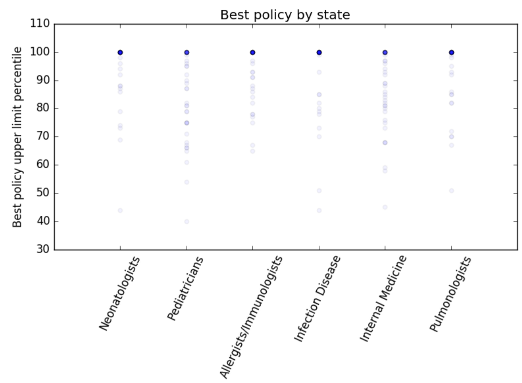

(d) State policies for Infant Mortality Rate

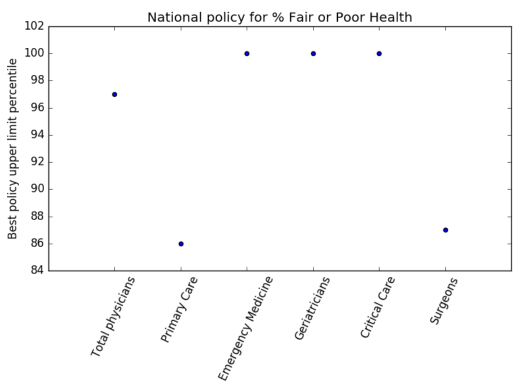

(e) National Policy for self reported health

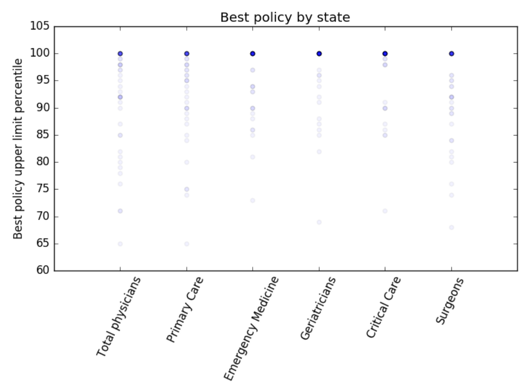

(f) State policies for self reported health

Fig. 6: National and state level policies for self reported health, child mortality, and infant mortality

have taken place to attempt for better medical resource allocation, including comparative effectiveness research (CER), and cost-effective analysis (CEA), both which aim to aid in the decision-making process by looking at effectiveness and fairness. CEA also incorporates a disability-adjusted life year (DALY) or a quality-adjusted life year (QALY) for measuring health outcomes.

Other current suggested policies aim to compromise between priority rules. For instance, the Triage-Treat-and-Release (TTR), a program developed in 2010 by the Lutheran Medical Center (LMC) in Brooklyn, New York , aimed to provide a solution for cases of overcrowded patients in emergency departments [11]. At LMC, differences between acuity in patients determined the admission and treatment process, and physicians as well as physician assistants (PAs) and nurse practitioners (NPs) are both used for both phases of serving low-acuity patients (whereas traditionally, NPs usually triage patients while physicians and PAs are responsible for treating patients.

6 Related Works

In recent literature, various studies have approached the problem of resource allocation with policies that propose how to make the best use those resources. The study in [11] looks closer into finding a solution for appropriately allocating medical resources that have already been distributed between low-acuity patients and patients in more critical conditions. In their study, they build off the TTR program with the development of a K-level threshold policy that would prioritize treatment unless K or more patients are in triage. Thus, is a class of policies that capture a decision-maker's valuation of each activity's importance, where lower K values signifies triage priority. Their approach is focused toward aiding physicians with an effective and simple method to allocate their time between triage and treatment. The research in [12] examines the current and historical trends in health resource distribution in the U.S. with respect to hospital beds and physicians. Their method involved using the Gini Coefficient (a statistical measure of distribution used to gauge economic inequality among a population) to evaluate variations in distribution at a county level during three decades, and their results demonstrated that physician distribution became less equitable, while hospital bed equity increased. Their findings confirmed positive associations between current physician and hospital bed distribution and suggested an increase in inequality of physician manpower.

There is also existing research in the use of Markov Decision Process to find solutions for other applications in the medical and health industry. Decisions in medical treatment is also a sequential process, as the course of a treatment must consider factors such as a patient's current health as well as the best treatment decisions for that patient in the future. As such, Physicians need to make subjective judgements about treatment strategies, and so [13] offers the use of MDPs as a technique for finding and optimal policy to solving these decisions. Their analysis refers successful applications of MDPs to medical treatments encompassing epidemic control, drug infusion, kidney and liver transplantation, and treatment for Spherocytosis, Ischemic heart disease, and breast cancer. Despite this wealth of potential applications, factors including heavy data requirements and computational limitations can cause for few successful applications of the decision-making model.

7 Conclusion

We explore a way to systematically determine the best course of action to take in allocating any given set of resources that will maximize the benefit it has on a specific health measure. Our approach involves constructing a Markov Decision Process which uses data on current and past allocations for those resources along with data on the desired health measure to maximize. We call attention to the most apparent shortcomings of our model, mainly the fact that the policies become increasingly dubious as the MDP is subject to data with increasing variability. Nevertheless, we show that our MDP is still able to operate and produce relevant results so long as the variability is controlled to an extent. Our MDP has a strong focus on quantitative measurements and bases all calculations wholly on such data, so there is no consideration for more qualitative aspects such as overall fairness. Furthermore, our MDP can only produce results which highlight the apparent effects, without being able to investigate the potential underlying causes; such examination is very relevant with the focus of this paper, but as it requires more work on the design of the MDP, it remains at large for the time being.

Acknowledgement

The authors wish to acknowledge the support of John P. Dickerson, Ph.D., Assistant Professor in the Department of Computer Science at the University of Maryland, College Park, for his teachings and guidance over the course of this project.

References

[1] Dartmouth Atlas Data. (n.d.). Retrieved May 17, 2020, from https://atlasdata.dartmouth.edu/static/research_data_archive?tab=35

[2] Hospital and Physician Capacity Measures by HSA Level - dataset by dartmouthatlas. (2017, July 25). Retrieved May 17, 2020, from https://data.world/dartmouthatlas/hospital-and-physician-capacity-measures-by-hsa-level

[3] Measures & Data Sources: 2020 Measures. (n.d.). Retrieved May 17, 2020, from https://www.countyhealthrankings.org/explore-health-rankings/measures-data-sources/2020-measures

[4] Health, D. of. (n.d.). Important Facts for Life Expectancy at Birth. Retrieved May 17, 2020, from https://ibis.health.utah.gov/ibisph-view/indicator/important_facts/LifeExpect.html

[5] Basu, S., Berkowitz, S. A., Phillips, R. L., Bitton, A., Landon, B. E., & Phillips, R. S. (2019). Association of Primary Care Physician Supply Withcc Population Mortality in the United States, 2005-2015. JAMA Internal Medicine, 179(4), 506. doi: 10.1001/jamainternmed.2018.7624

[6] Premature Mortality in the United States: Public Health Issues in the Use of Years of Potential Life Lost. (n.d.). Retrieved May 17, 2020, from https://www.cdc.gov/mmwr/preview/mmwrhtml/00001773.htm

[7] Child Mortality. (n.d.). Retrieved May 17, 2020, from https://data.unicef.org/topic/child-survival/under-five-mortality/

[8] DeSalvo K, Bloser N, Reynolds K, He J, Muntner P. Mortality prediction with a single general self-rated health question. J Gen Intern Med. 2006;21:267-275

[9] Andresen EM, Catlin TK, Wyrwich KW, Jackson-Thompson J. Retest reliability of surveillance questions on health related quality of life. J Epidemiol Community Health. 2003;57:339-343

[10] Daniels N. Resource Allocation and Priority Setting. 2016 Apr 13. In: H. Barrett D, W. Ortmann L, Dawson A, et al., editors. Public Health Ethics: Cases Spanning the Globe [Internet]. Cham (CH): Springer; 2016. Chapter 3. Available from:

https://www.ncbi.nlm.nih.gov/books/NBK435786/ doi: 10.1007/978-3-319-23847-0_3

[11] Zayas-Caban, G., Xie, J., Green, L., & Lewis, M. E. (2019). Policies for Physician Allocation to Triage and Treatment in Emergency Departments. IISE Transactions on Healthcare Systems Engineering, 1-33. doi:10.1080/24725579.2019.1620384

[12] Horev, T., Pesis-Katz, I., & Mukamel, D. B. (2004). Trends in geographic disparities in allocation of health care resources in the US. Health Policy, 68(2), 223–232. doi:10.1016/j.healthpol.2003.09.011

[13] Schaefer, A. J., Bailey, M. D., Shechter, S. M., & Roberts, M. S. (n.d.). Modeling Medical Treatment Using Markov Decision Processes. International Series in Operations Research & Management Science, 593–612. doi:10.1007/1-4020-8066-2_23