422 - Reading 1 Summary Questions

Patient Type

- Would there have been a benefit to even further classifying of patient types? For example, instead of the four categories of IP, OP, CEP, and NCEP, what if they further segmented these classes into different priorities (as objectively as can be done)? **(Asked ~4 times)**

Penalty Cost

- How do you determine the penalty costs and how do you know that they are optimal for finding the optimal policy?

- How was the penalty cost for IPs decided in reference to the variable cost incurred by having the patient stay an extra day?

- The penalty score chosen for not scanning a guest should serve as a proxy for indirect impacts on revenue - has evaluation been done to justify the weight of the penalty chosen in relation to the reward?

- The method uses 'penalty cost/wait cost' to determine priority between different types of patients, but how specifically would these costs be quantified in a real world scenario?

- A base case penalty was evaluated as $500 by assumption, but is there a better way to evaluate this penalty? Assuming a penalty when it is part of the heuristic feels like a very large assumption to make, and should be calculated in a real-life situation

- What exactly are the waiting cost and penalty cost? How were these costs generated?

- Why are the penalty costs/revenues, shown in Table 2, assumed rather than obtained from real world observations? With real-world values, it's possible that the costs/revenues simplify the problem and therefore make this study effectively useless.

Applicability of Model to Other Domains

- Could the same MDP solution be applied to scheduling of other parts / functions of a hospital (such as scheduling nurses for work, etc). **(Asked ~3 times)**

- Can this model be extended to other imaging services as effectively? (E.g. MRI scans, PET scans etc.)

- Can this MDP model be used for scheduling problems in other areas such as banking, tutoring?

- Can this method be used for other hospital in other countries? Or what kind of hospital is best fit to this solution?

Maximizing Revenue

- Why did they choose to have net revenue as the parameter to maximize rather than minimizing the wait times of patients?

- How could the model be modified to balance maximizing revenue with the other patient-centric variables (like time waiting)?

- Why are the rewards purely monetary? Wouldn't patient satisfaction be a factor?

Experimental Setup Criticism

- Why they did not test their solution, the optimal schedule given by MDP, in practice as an actual schedule in a hospital for some duration of time to compare the results with other schedules in real world setting? **(Asked ~2 times)**

- The paper uses simulation in the evaluation. Should we be confident about their performance analysis?

Questionable Assumptions

- Why are they so certain that the people that operate and do the ct scans are always available, for instance many times there could be a shortage in staff resulting in fewer ct scans being able to be done. Which would make it scheduling by type of patients useless.

- How are they so sure that CT take the same duration of time for every patient, and how did they approximate this procedure? Maybe some people are more scared of the the procedure thus needing more time to get the procedure done.

- The paper seems to make an unmentioned assumption on the quality of the scanners, though the scenario where one of the scanners broke down would be disastrous on revenue and patient scheduling. How would that and the subsequent rescheduling of patients factor into the MDP?

- The values chosen are based off a single urban hospital. Is there reason to believe that this hospital is representative of a typical hospital? Because if not the model might be skewed.

- Is it reasonable to assume that scheduled OPs have a dedicated scanner, i.e. is it possible that a more optimal policy also schedules other patient types for said scanner?"

- Why was the arrival probability purely assumed instead of using a value that reflected the probability of outpatient arrivals from Harborview?

- Why is it assumed that at most one request of each type can arrive per slot?

- Why is it assumed that the slots have identical lengths?

Heuristics

- The heuristics used to find how effective an MDP-based optimal policy would be were relatively simple. Why not use a more complicated heuristic?

- "Heuristics were chosen as those that are 'common in the queuing literature', but are these known for efficiency? Why not test more varied heuristics"

- "What was the cause of the single and two scanner cases being so similar in the O-1 prioritization compared to other heuristics?"

Clarity Issues

- The authors repeatedly use the term ‘revenue’ in talking about the optimization target of their study, whereas another term ‘profit’, in describing that of a similar study by Green et al. But given that the authors take into account potential cost in their model, why do they not use the term ‘profit’? In particular, to what extent does this affect their claim as to uniqueness of their study with respect to that of Green et al in particular?

- What is the tornado chart and why is this chart being used in this research?

- Can you go over the basic idea for sensitivity analyses?

- "Administrators apparently us ad hoc or rule of thumb approaches, but what specifically are these approaches, and how efficient are they?"

- What is add-on outpatients?

Live Implementation/Future Direction

- Was the policy implemented in the workplace/hospital after the study was concluded? **(Asked ~2 times)**

- In the conclusion, they say that "[their] work has several limitations owing to the complexity of the schedule problem with multiply scanners and multiple patient types". This was written several years ago. Has there been any progress done in finding a solution that requires no simplifying for this type of scheduling problem? If so, how much?

- This paper was from 6 years ago; are such methods commonplace in scheduling tasks now and have better methods been discovered since then? If not, what do you think is the obstacle towards implementing such methods in real world systems?

- Has any hospital attempted to implement MDPs in their scheduling rather than simply simulating, and did they report findings that were similar to the paper's?

- How does the model perform compared to the other models previously suggested? The comparisons shown in the paper are only against various naive models.

- What are the total estimated costs of implementing this method in practice?

- How exactly is this approach going to be implemented?

- The proposed solution does make some assumptions which greatly simplify the problem. What would be the next steps and how would the model be different to account for the complexity of patients needs?

- How can the solution actually be rolled out in a way that will minimize risk, in case it does not perform as well as it does in theory.

- What factors are limiting the application of this study to improving patient service?

- How can this mathematical model and MDP be changed to automatically schedule patients in real life and in real time?

- Can they automate the entire decision process where a CT technologist would not be required?

- In abstract's conclusions it is written that the optimal policy can be implemented relatively easily and then in the next sentence it is stated that priority heuristics are much easier to implement. Does it mean that the optimal policy is not easy enough to implement? **(Asked ~3 times)**

- How often do medical providers use smart policies like Markov decision process versus simple heuristic policies in practice?

Optimal Policy

- What if we were to disregard the economic part and look for the most efficient solution? Would the answer still be the optimal policy?

- Since the optimal policy is focused on maximizing the daily net revenue, is it fair to evaluate it in a different performance metrics?

- Would the optimal policy generate more revenue if it was compared with a better/tighter heuristic decision?

- In class we see how the value iteration works to find an optimal policy, they used instead Approximate dynamic programming. (I know that I can find the paper of Powell, read it and maybe understand the procedure) Which is the difference between those two approaches?

- Why is the decision of the optimal policy mostly influenced by costs and revenue vs quality of healthcare?

- How long did it take for the optimal policy to form?

- Since there are so many variables to account for (type of patient, type of scan, number of scanners, number of scans per day), how does the optimal policy know which rewards are linked to which variables?

- Since every patient is so different, shouldn't the optimal policy for scheduling take in account every factor of the patient as opposed to only the factors listed in this experiment?

- Would adding more variables such as month, gender of patient, or weather affect the optimal policy?

- How will having more than 2 CT scanners affect the optimal policy (what about 3,4,5,6,7, etc), as we saw that that there was some notable difference in the optimal policy in 1 and 2 scanner cases, does there exist a trend for MDP with regards to this? (I'm not sure

how many CT scanners a hospital has, but It seems like at ~$2.5 million a big enough hospital will have more than 2). **(Asked ~3 times)**

- How much influence would increasing the stochasticity (patients not arriving on time, patients that need to be immediately scanned) have on the optimal policy?

- What would be the improvement of the optimal policy be if it changed the prioritization of patients at certain parts of day, based on data shown in Fig.1? For example, the policy could prioritize emergency patients at 7-8:30am and after 1:30pm while prioritizing inpatients at all other remaining times.

- Why did the optimal policy sometimes lead to more patients waiting than if an heuristic was used?

- Should the optimal policy just be focused on least waiting times, revenue or a mix of both?

- In the optimal policy should the “emergency” aspect have a much bigger part because patients that are more unwell need to have lesser waiting times?

- How well does the optimal policy hold against more complex heuristics? The heuristics that were compared against seem very simplistic and not likely to fully reflect scheduling in all hospitals.

- Would a FCFS with weighted priority policy do better than the current optimal policy?

Ethics

- Would the model cause ethical problems in hospitals? Is revenue a good metric of performance if we put life and death situations into consideration? **(Asked ~5 times)**

MDP

- What’s the advantage of using finite-horizon MDP than infinite horizon MDP? **(Asked ~6 times)**

- Has machine learning been considered as an approach to increase the performance of the MDP?

- Could the MDP model be easily extended to account for the limitations/weaknesses of this work, or would it be much more difficult to do so?

- What is a potential adjustment you can do to the MDP that will account for the variability in time taken to perform a scan, or multiple scans for a particular type of patient?

- In lecture, MDP specification includes a stopping state. In the case of this paper, does the stopping state correspond to the nth stage (nth time slot in a work day)? Is the number of stages the same everyday?

- If the researchers make assumptions more realistic, for example, patients can make more than one request of each type per time slot, how will the model be extended? Can they do an infinite-horizon MDP version?

- How can the MDP take into account the "perception of service quality" by the outpatients?

- The paper says they used a finite-horizon MDP. What was the horizon value?

- Why was an MDP used, and not some other type of algorithm/reinforcement learning? **(Asked ~3 times)**

- I notice that in this paper, they do not consider the discount factor in the reward function. Would the result change if we add the discount factor?

- How would the MDP account for human error? Assuming that hospital personnel are the ones ferrying patients between scanners, let's say one chose the wrong patient to bring in on accident, how would that affect the MDP?

- How effective is Value iteration in this process?

- How do we represent the states of an MDP if each type of patient requires different durations with the CT scanners?

- The paper answers why a finite MDP was explored, rather than an infinite MDP, but are there any cases that an infinite MDP would be suitable?

- The problem evaluated 32 scenarios. Is a Markov decision process the optimal algorithm to consider when the numbers of scenarios is relatively low?

Methodology

- The author mentioned that most of their data was obtained from a single academic hospital (Harborview Medical Centre). Doesn't this make the overall method biased? **(~Asked 4 times)**

- Is using only two levels for each of the parameters really sufficient for the analysis?

- Given that their study seems to rely on the data collected from a single hospital, could the set of parameter values they used be considered to generalize to the conditions at other hospitals? To what extent, does this affect the scope of their claims?

- The Table 2 (77) presenting their 32 combinations of parameter values shows the probability only for arrival of outpatients. Did they take into account the corresponding probabilities for the other categories of patients in testing their model with respect to the others?

- What was the rationale for using patients scanned in order of arrival as a heuristic decision rather than the urgency of medical attention that each patient required?

- Arrival probabilities are non-stationary but only limited to the time of day. Could the model be improved by considering seasonal trends? For example, there may be more arrivals in the winter than the summer due to falls and accidents from slippery conditions.

Scanners

- What if a patient has not been scanned for many days? The model doesn't concern about this.

- Why did the author only use single-scanner and two-scanner options? They did mention that most hospitals/facilities have "multiple scanners". So, why stop at two? **(Asked ~7 times)**

- How would the results change in the case that there are more than 2 scanners?

- How could we change the model if we have different types of scanners like MRI, CT…

- What would be the best number of scanners to have, but they have penalty costs of their own?

Model

- How does the model expect to change as one adds more categories to EPs, in terms of severity of the illness?

- "The article implies that unscanned patients by the end of the day is implicitly incorporated in the model. How is this done (e.g. start with a different initial state for next day)?"

- Is it possible to compare this model to the models developed in previous work? ie. Is this the "best strategy" explored for this problem?

- What needs to be changed in the model to be able to accommodate situations where more than one service request will arrive during a slot?

- The model assumed that non-critical emergency patients had the same clinical urgency as inpatients and outpatients. What kind of changes would be made if non-critical emergency patients were given priority over the latter two (but still under the critical emergency patients)?

- What happens if you add a sudden influx of patients? Example, due to a nearby accident. Will it still perform better than the heuristics?

- With a bigger hospital, comes more equipment but also more patients and variability. Can the model be extended to handle that?

- The parameters of the model are not directly learned from data. Would that be appropriate?

- Does the absolute values of rewards matter as long as the relative values stay the same, in terms of affecting the outcome in this experiment or in general?

- Why wasn't a partially observable Markov decision process used instead of a finite-horizon Markov decision process since the agent not being to observe the underlying state would have been more realistic?

- How would the model change if more than one patient of each type could arrive in a period?

- In a real world, CT scans could take varying times (10~60 min) rather than fixed time slots. Will this make the model completely useless?

Complexity

- "What is the time complexity of computing the scheduling using the proposed solution?"

- "What is the space complexity of computing the scheduling using the proposed solution?"

- What was the average runtime and performance costs to calculate the optimal policy when a new patient made a CT request?