Scheduling

Patient type

- Is there a better way to classify patients than CEP, NCEP, OP, and IP?
- Patients are incredibly heterogeneous, and don't fit easily into the suggested types. How might patients be split into sub-categories?
- What are the ethical concerns in using this type of procedure to prioritize certain patients?
- In the priority-based heuristic rule, why was OP chosen to be prioritised instead of NCEP or IP?
- Why are patients organized by EP, IP and OP and what is the criterial for priority? Why aren't IP or OP categorized as critical and non-critical as well?
- Why would an IP not have a higher priority than it does? It seems as tho getting someone out as quickly as possible would be beneficial for efficiency levels and revenue
-
- It seems that patient priority was only considered in the case of CEPs, as for the remaining types of patients only number of requests was considered, does this mean we are only scheduling which patient type gets a slot? How do we schedule a patients within a group? (how do you determine which IP is scheduled over other IPs, first come first serve?)
-
- Why would an urgent OP not just be categorized as an EP? It seems as though this would be much more beneficial for all parties
- They mentioned that OPs are scanned by a dedicated scanner but I'm sure many hospitals cannot afford a separate scanner for OPs. How would the model be different if non-add-on OPs needed to be included?

- Why do you not want to have the number of type j patients scanned to not exceed the number waiting?
- Waiting NCEPs, IPs and OPs were considered identical in terms of clinical urgency but this doesn't seem very realistic. What would the model look like if patients somehow had a level of urgency associated with them?
- "If NCEP's and OP's have a less worse penalty cost than an IP, why would they all have the same clinical urgency?" It would seem that making an IP in this case have a slightly higher clinical urgency in order to lower the overall penalty cost.
- Gocgun mentioned that the number of OPs not scheduled by the end of the day was relatively high. Since the MDP only considers a single day, is there a possibility of an OP never being scheduled and always pushed to the next day?
- Why aren't scheduled OPs considered in the model? Though Gocgun et al. assume that scheduled OPs are given a dedicated scanner, waiting times for scheduled OPs can often be huge. "Loss of quality" is mentioned several times throughout the paper, so if a scheduled OP is forced to wait months for their appointment couldn't this be considered loss of quality?
- why would they assume that any IPs that are not scanned by the end of the day was automatically scanned the following day. Wouldn't they still need to occupy the slot the next day to be scanned?

- What may be the worst case scenario for an IP waiting for a CT scan given the assumptions present in this model? Specifically the assumptions that IP's not seen in one work day are assumed to be seen the next work day.
-
- Would the results of their research change if they made more categories of the patient subtypes?

Penalty cost

- How did they develop their penalty costs for their MDP?
- Why incur a penalty cost for every arriving patient who is not scanned by the end of the day? They likely must still be scanned in the future, and it's obvious that not all work will always be completed at the end of each day.
- Does this model take into account the costs of being forced to wait multiple days? I would imagine a constant penalty per day would not portray the service-quality effectively if patients had to ever wait multiple days for a scan.
- The paper assumes a penalty cost of between $100-$500 for penalty costs, which was when a patient had to wait longer than a maximum recommended waiting time. However, we all know that the value of a human life is far greater than that. While the paper's calculations doesn't take into consideration the other hospital costs inherited from a patient waiting longer than they should have for a scan, and potentially negative health connotations from that. If we were to take into account that added cost from changing of the patient scheduling system from the status quo into this new optimized policy with potentially higher patient risk, would it still be as competitive as just using the simplified numbers given in the calculations in the paper?
- I wonder how the model would look like if they optimized for different outcomes, like minimizing the amount of patients who need a scan, or having more scanners available, or even finding out the actual cost involved for the different types of patients instead of using an assumed value of $500?

- Why did they not consider later in the night for in-patients as they are already at the facility, making it easier to schedule a scan for them after normal work hours?
- Their research only looked at 9-5 workday. What about the evening/overnight data? Would their results change?
-
-

Model

- How would the model be affected with patients that arrived but were late?
- Could the MDP optimize the scheduling process better with a different length horizon?
- Why does the number of patient types influence whether one should use a finite vs. infinite-horizon model?
- Does this model become problematic, and infeasible when we take into consideration the actual 24-hour work day in a hospital? How would the model treat boundary cases in such an environment? Would it then be more efficient to have a discounted infinite-stage or

indeterminate-stage MDP with a reporting function or terminating state once a certain threshold of revenue has been accrued? Thus permitting us to optimize on patient satisfaction and distance to the terminating state?

- why are the stages not a part of the states in the model?
- How large of a state space can we feasibly compute?
- Why were the five heuristic models, FCFS, the two random selection, and the two OP prioritization, chosen over other approaches? Why are they good approximations of a real-world implementation of the findings of the paper?
- If Green et al. already explicitly cover the single scanner model, why do Gocgun et al. bother considering it? Are the gains from having additional performance metrics enough to rationalize the case's consideration?
- Given the limitations in a single scanner environment, what would proposed optimizations be to the human element of deciding on scheduling,
- What if there was an additional parameter to minimize the number of patients left unscanned?
- Given that the decision rules O-1 and O-2 are next-best to the optimal policy found by the model, and that they are much simpler to implement in the hospital environment, would it have been beneficial or superfluous to include parameters in supporting experimental analysis of the model covering the ease-of-implementation?
-
- Since this is a finite horizon and the number of decisions is constrained by the time slots, would modelling this after a regular decision network be a possibility?

Applicability of model to other domains

- Can this MDP be extended to cover other patient scheduling problems? (such as surgeries, appointments with doctors, etc.)

- Besides having this mathematical model contribute to health care, do you think this model can contribute to other industries like finance, agriculture, or sales?
- Can this approach be easily applied to other scheduling problems in the medical field and possibly outside the medical field?
-

Scanners

- How would the outcome change if we were to use **> two scanners**?
- Could this model be applied to a CT scheduling system with more than 2 scanners? Say 3, 4 up to N scanners, does this model still give accurate solutions?
- How might using more than 2 scanners impact the results?
- What about a three-scanner model? What does that look like? How does it compare to the one- or two-scanner?
- Will increasing the amount of scanners from 2 to 3 have the same effect as increasing the scanners from **1 to 2?**

- Could this method of finding the optimal CT scheduling policy be used in hospitals where the scanners aren't as up to date and take longer to process information?
- How could this research handle multiple patient types with multiple scanners?
- The model doesn't take into account the cost of the scanners themselves. Would it really be profitable to purchase 2 scanners with the optimal policy of this model? How long would it take to pay off a second scanner?
- Why only limit the number of scanners to 2? Why not observe a broader range of scanners to find what the optimal amount of scanners was?
-

Maximizing Revenue

- Who really stands to benefit from this research? The hospitals' revenue, or patient satisfaction? There is a delicate balance between keeping the costs of running the tests low, and making sure that patients aren't stuck waiting forever.
- What are the ethics when looking at patients in terms of revenue and cost?
- Ethically speaking, would the net revenue increase for using a scheduling system designed to maximize net revenue be, worth the potential risk of lives by not prioritizing the health of a patient over hospital profits?
- The increase in revenue as a result of the optimal policy is significant but is it economical to transition to that strategy? An small increase in revenue should not necessarily mean it is ideal to move to the new strategy because the cost in time, money and other resources can outweigh the gain in revenue.
- How does the extra revenue generated by having an additional scanner relate to the cost of buying+running the scanner? Where is the diminishing return point?
-
- Is there a way to combine the heuristic and optimal policies in order to maximize revenue as well as patients served in a day?
- It seems revenue is weighted the highest in the reward, what would happens if wait time is configured to be more important?
- would a higher number of scanners and patients give a similar revenue gap compared to the heuristics?

Methodology Criticism

- This paper seems to be building slightly on similar methods described in other papers by adding some more complexity to the problem. Will the method be able to withstand much more complex variables?
- Could we make abstractions to solve this problem using a regular scheduling algorithm (as in CPSC 320)? If so, which parts of the problem need to be abstracted out? If not, what prevents this from being possible? Would it become too simple and not indicative of what really happens?
- In my opinion, a major flaw is that they only allow one of each type of unscheduled patient to appear in each slot, so I was wondering how much different the model would look if it could take multiple emergency patients in a single time slot (like if 5 people were in an accident and all need scans)?
- The article doesn't explicitly talk about calculating the policy via value iteration, but I assume this must be the case. In comparison to heuristics, isn't it guaranteed to show improvement as it has been fit to the specific problem space that they have defined? What impact does that have on what we can say about the results?
- The equation doesn't seem to include any discount factor. Would it be beneficial to include it?

Experimental Setup criticisms

- In 3.4.1, It is stated that "...by varying revenue, cost, and arrival probability parameters we generated thirty two scenarios...". No further detail is offered. How were the scenarios actually decided upon, exactly? Why Revenue, cost, and arrival probability only?
- It says in the paper that there are many parts of the research which had to simplified due to the complexity of the scheduling problem, especially with different calculation values when taking into account multiple scanners and multiple patient types. However, instead of assuming these values, how would the results of the multi-scanner simulation compare to the current system (that is, in terms of revenue) used at the hospital if we used the closest exact data points that we had, and also took into account more realistic scenarios such as multiple arrivals of each type (as the paper said it was possible to do this calculation)?

- Why weren't more data sets tested against to compare whether the results they obtained from their model was reproducible against other medical facility data?
- Why is the study not done using more accurate cost and probability models ? as they are easy to implement in the simulation software and may lead to more optimal result

- Why is it that only testing individual parameters for high and low values is sufficient preliminary sensitivity information?

Experimental Results

- If the solution Gocgun et al. provided was similar to Green et al. and Kolisch and Sickinger, do we know how well Gocgun's solution compares or if there is a significant difference?
- What could have been the cause for the 80% decrease of average waiting time? At first thought, it seems as though increasing from one scanner to two should only decrese the wait times by 50%

Questionable assumptions

Data from only one hospital

- How can you know which assumptions can be made with respect to patients, scans and costs while still accurately reflecting the reality of CT scheduling?
- Are Gocgun et al. proposing solution to the medical scheduling problem for US hospitals only, or hospitals in general? If they are trying to solve the problem for hospitals in general, then using only data from US medical institutions might not capture cultural or regional differences on patient behaviours. This weakness might impact the assumptions that they have to make for the MDP model.
- Will there be efforts to collect data in hospitals other than Harborview?
- Since most of the data was taken from Harborview Medical Center, how sure are we that this data is representative of other hospitals. And how confident can we be of this solution when applied at a different hospital.
- Data was gathered mainly from Harborview Medical Center, is this facility an average representation of the medical centers found elsewhere? Will these results be reproducible if performed on data from other medical facilities?
- What evidence is there that the patient arrival probabilities derived from the Harborview data are representative of the conditions at most hospitals' CT scan services? That is, why can the Harborview data be generalized to all hospital CT scan services?
- What would this model have looked like if the numbers were based off a different hospital? Perhaps not an academic one? Maybe one in a developing country, or at least one that isn't as well-funded/developed?
- Is using data from one hospital enough? Should they gather data from other hospitals to make a general model?
- The data for this study was obtained from an academic hospital/Level 1 trauma center, and quote: "Our data revealed that IP and EP arrival patterns vary considerably over the course of a day". How would this variation be different for different types of hospitals. How much more useful would this model be in hospitals where traffic is more consistent?
-

- How did they determine what assumptions to make and other problems that were acknowledged to adding to the complexity of the issue?
- Why is the assumption that only one patient of each type may arrive during a single time slot reasonable for modeling a CT scan service, and is it motivated by the fact that only single patient arrival probabilities were derived from the Harborview data?
- Why assume that only one patient request for each type can come in at each slot? This seems like an arbitrary restriction. Does it simplify the MDP significantly?
- This model seems to have a lot of limitations and unrealistic assumptions such as only one request per patient type per slot is received, which seems to make this model unrealistic to

implement in practice, in general, how do we know when assumptions are good estimates of reality and can be implemented?

-
- There are many assumptions about states and actions. Would a partially observable Markov decision process be more suitable in reality?
- Would the model have been more accurate to real life if a randomness aspect was added to the potential transport time for patients to wards?
- What if the patient misses their appointment for their allotted CT scan? How does this model take into account that situation?
- CT Scans can take an indefinite amount of time. How might the MDP be altered to account for the fact that a patient may take longer than expected in the scanner (e,.g, they accidentally moved and therefore had to be rescanned). Can we account for this possibility?
- Does every CT scan take same amount of time (30 minutes)? If some scan takes longer or shorter time the optimal policy may not work anymore.
- The models ignore some key real life situations such as (e.g.. CT scans might not fit perfectly into time slots). Could these be factored in or would it be too tedious?

Heuristics:

- Do the tested heuristics actually reflect the queuing practices in hospitals? The researchers note that they had sourced the heuristics from queuing literature, but do not provide citations. Are these heuristics still used in hospitals today?
-

Optimal Policy

- How would we adjust the optimal policy in order for it to perform better for patients not scanned and waiting times?
- What is the optimal sequence/function of patients resulting from the optimal policy?
- Was the optimal policy ever tried in the hospital?
- The optimal policy obtained by solving the Markov Decision Process (MDP) model is compared to heuristic decision rules from queuing literature, but not to any policies being used at CT scan services. What form do actual, in-service policies take and how does the MDP's optimal policy compare to them?
- Why is the average number of OPs not scanned by the end of the day greater (i.e. worse) for the optimal policy than it is for the priority-based heuristics?
- How can the MDP be improved upon so that the optimal policy will not do worse than existing heuristics in any way?

MDP

- MDP has the assumption that the states are fully observable. However, I think observations of patients in a hospital waiting for medical imaging does not simply reduce to just the number of patients of each patient type. For example, there are also the moods of patients and the number of technicians and nurses that might be relevant to the problem. Some patients might be more tolerant on the issue of waiting than others. In this case, other than just the action of assigning number of patients of each class for the next time slot, action set of the MDP model should expand. Are there some explanations to choosing MDP over POMDP and restricting observation to only the numbers of patients of each patient class? Also, there might be more possible actions if the set of evidences expand.

- They mentioned that their approach is overall better than other types of scheduling techniques with priority-based techniques coming in second. At what point will it become valuable enough to implement the finite MDP approach over the easier priority-based techniques?
- How does the heuristic methods fare if they are implemented as part of the MDP (such as changing the weights to better performing heuristic, e.g. O-1)?
- Is it possible for the MDP to create a more optimized policy if it was ran for a longer duration than 4 months?
- Is it possible to use and evaluate multiple different types of rewards in MDP, in order to improve accuracy?

Clarity issues

- What is a sensitivity analysis? What are Tornado charts?
- What is the definition of a Tornado chart, and why is it a useful visualization? What are they attempting to demonstrate?
- Clarification on the one-way sensitivity analyses and the significance of three parameter levels.
- What is a switching-curve policy?
-

Future directions/ Ideas for improvement

- In the conclusion, one of the limitations states: "The fact that some of the IPs are not scanned during the day and stay overnight could be incorporated more explicitly using an infinite-horizon MDP model." I am thinking that is it possible to incorporate the findings from the finite-horizon model into running the infinite-horizon version? In this case, they might find out that if the number of IPs who do not receive scanning in the end of the day does converge. If there is a convergence, then they might try making the policy prioritizing those IPs in the following day, and see if the number of IPs who do not receive a scan in the end of day stays stable for each day. Otherwise, those IPs might accumulate.
- What happens when infinite-horizon MDP is used instead. Does policy change significantly?
-

- Would it be possible to have the model extended the horizon such that rather than any given day, have the previous day's results impact the next and simulate over a week/month? In example, if we didn't get to an inpatient today we would definitely want to get to them first thing the next day.
- Are there plans to deploy a tool which uses this decision process to hospitals?