Intelligent Systems (AI-2)

Computer Science cpsc422, Lecture 9

Jan, 29, 2021

An MDP Approach to Multi-Category Patient Scheduling in a Diagnostic Facility

Adapted from: Matthew Dirks

Goal / Motivation

- To develop a mathematical model for multi-category patient scheduling decisions in computed tomography (CT), and to investigate associated trade-offs from economic and operational perspectives.

- Contributions to AI, OR and radiology

Types of patients:

- Emergency Patients (EP)
  - **Critical (CEP)**
  - **Non-critical (NCEP)**
- **Inpatients (IP)**
- Outpatients
  - Scheduled OP
  - **Add-on OP: Semi-urgent (OP)**
- **(Green = Types used in this model)**

Proposed Solution

- Finite-horizon MDP
- Non-stationary arrival probabilities for IPs and EPs
- Performance objective: Max $

always served in next turn

MDP Representation

\# of states 2×5^3

- State

  0/1

  - $s = (e_{CEP}, w_{OP}, w_{IP}, w_{NCEP})$ — max 4 of each type waiting
  - e_{CEP} CEP arrived
  - w_{type} Number waiting to be scanned
- Action

  $a_{OP} + a_{IP} + a_{NCEP} + e_{CEP} \le R$

  $a_j \le w_j, \quad j = \text{OP, IP, NCEP},$

  - $a = (a_{OP}, a_{IP}, a_{NCEP})$
  - a_{type} Number chosen for next slot
- State Transition
  - $s' = (d_{CEP}, w_{OP} + d_{OP} - a_{OP}, w_{IP} + d_{IP} - a_{IP}, w_{NCEP} + d_{NCEP} - a_{NCEP})$
  - d Whether a patient type has arrived since the last state

MDP Representation (cont')

- Transition Probabilities

$$P_n(s'|s, a) = p_n(d_{\mathrm{CEP}}) \times p_n(d_{\mathrm{OP}}) \times p_n(d_{\mathrm{IP}}) \times p_n(d_{\mathrm{NCEP}}),$$

Example: action

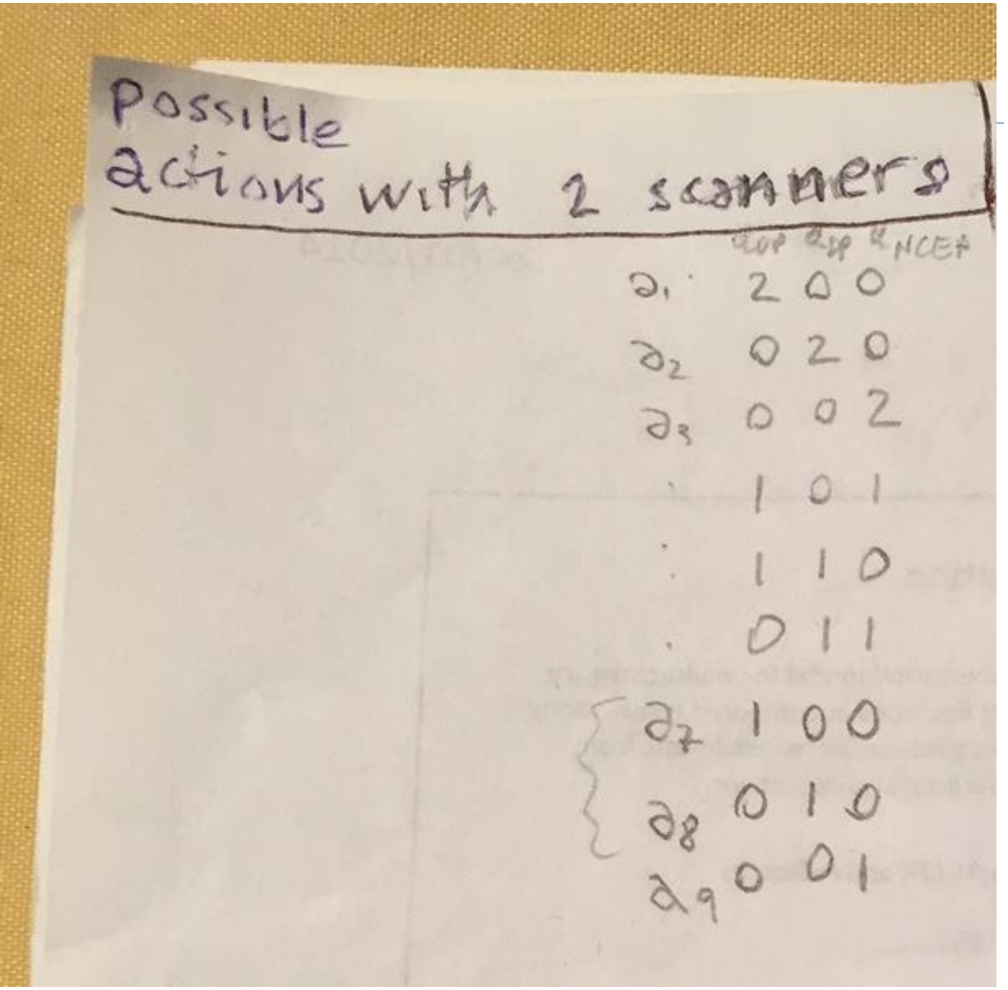

example

Possible actions with 2 scanners

	OP	IP	$NCEP$
a_1	2	0	0
a_2	0	2	0
a_3	0	0	2
	1	0	1
	1	1	0
	0	1	1
a_7	1	0	0
a_8	0	1	0
a_9	0	0	1

assume e_{CEP} so one scanner goes to this one patient

$\hat{S} = (1, \underset{OP}{3}, \underset{IP}{2}, \underset{NCP}{0})$ is the current state

applicable actions

a_7 a_8

what are the s' for a_7

after a_7 $\rightarrow (0,2,2,0) + (d_{CEP}, d_{OP}, d_{IP}, d_{NCEP})$

possible arrivals configurations

	d_{CEP}	d_{OP}	d_{IP}	d_{NCEP}	Corresponding s'
X_1	0	0	0	0	$\rightarrow s'_1 = (0,2,2,0)$
	1	0	0	0	$\rightarrow s'_2 = (1\ 2\ 2\ 0)$
	0	1	0	0	$\rightarrow s'_3 = (0\ 3\ 2\ 0)$
⋮	—	—	—	—	⋮
X_{16}	1	1	1	1	$\rightarrow s'_{16} = (1,3,3,1)$

$P(s'_i \mid \hat{S}, a_7)$?

sample conditional prob. for s'_3 given a_7 is performed in $\hat{S}$

$$P(s'_3 \mid \hat{S}, a_7) = P(d_{CEP}=0) * P(d_{OP}=1) * \dots$$

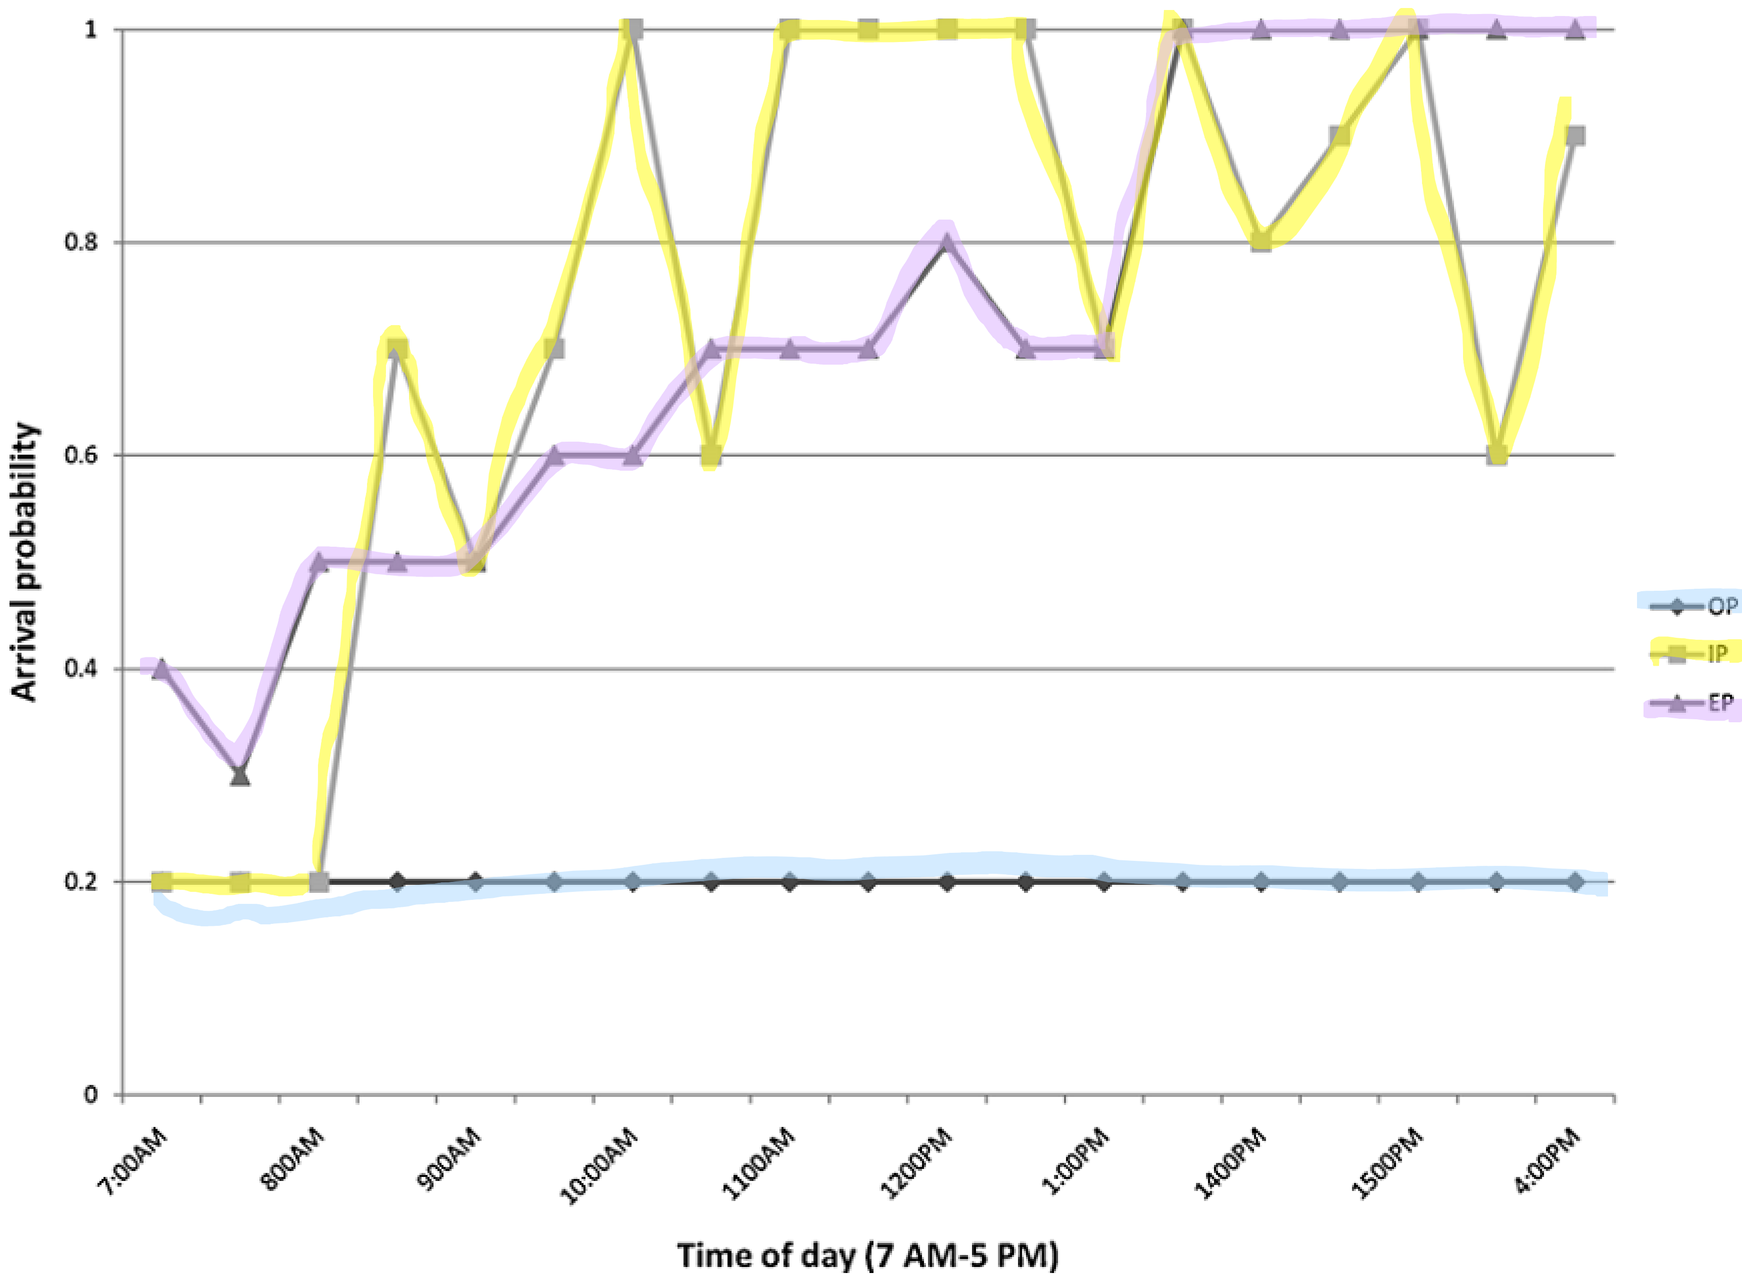

Fig. 1. Arrival probabilities for each patient-type during a work-day. EP includes both CEPs and NCEPs.

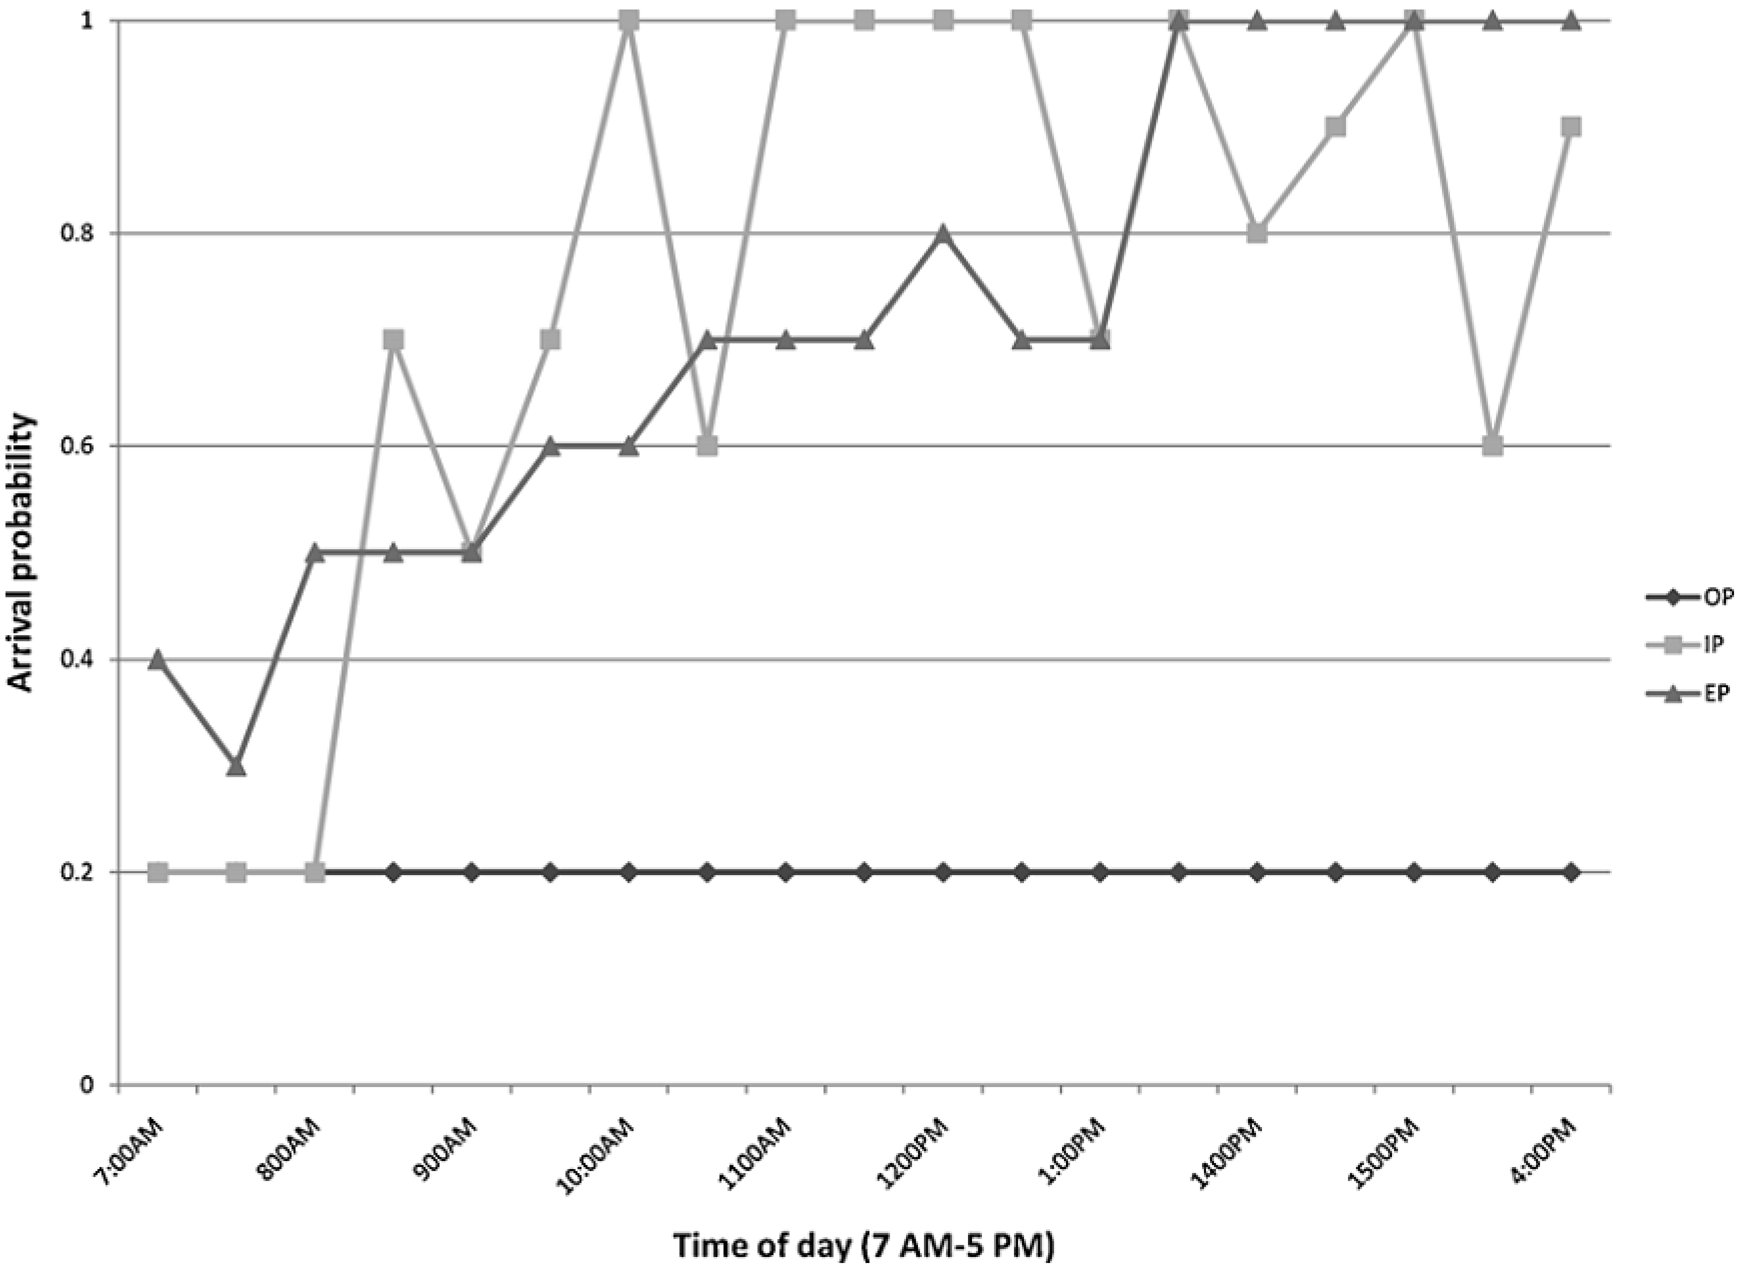

Fig. 1. Arrival probabilities for each patient-type during a work-day. EP includes both CEPs and NCEPs.

Performance Metrics (over 1 work-day)

- Expected net CT revenue
- Average waiting-time
- Average # patients not scanned by day's end

- Rewards

revenue for scanning types patient

$$r(s, a) = \sum_{j \in \{OP, IP, NCEP\}} r_j a_j - \sum_{j \in \{OP, IP, NCEP\}} (w_j - a_j) h_j.$$

waiting cost

end of day penalty

Terminal reward obtained

$$V_{N+1}(s) = -c_{OP} w_{OP} - c_{IP} w_{IP} - c_{NCEP} w_{NCEP}$$

- Discount factor? 1

Maximize total expected revenue

- Optimal Policy
  - Solving this gives the policy for each state, n, in the day

$$V_n(s) = \max_{a \in A(s)} \left\{ r(s,a) + \sum_{s'} P_n(s'|s,a) V_{n+1}(s') \right\}$$

- Finite Horizon MDP

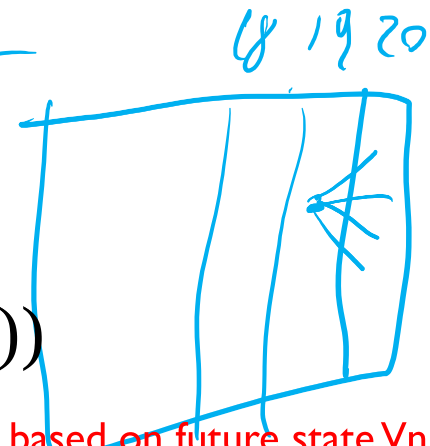

$$V^{\pi^*}(s) = R(s) + \gamma \max_a \sum_{s'} P(s'|s,a) \times V^{\pi^*}(s'))$$

- The recursive equation (3) has value of current state Vn calculated based on future state Vn+1, this contradicts with the equation given during class, where Vn+1 depends on Vn?
- The one in class was Value Iteration (the n index was for the iteration) here we have a finite horizon. We know the Vs at the end so we can compute all the Vs backward. n is an index for the time slice

Evaluation: Comparison of MDP with Heuristic Policies

- 100,000 independent day-long sample paths (one set for each of the 32 scenarios)

$\forall n$ sample $P_n(dx)$

$X = CEP, OP, IP, NCEP$

Result Metric

- Percentage Gap in avg. net revenue =

$$\frac{avg\ net\ revenue\ (optimal\ policy) - avg\ net\ revenue(heuristic\ policy)}{avg\ net\ revenue(optimal\ policy)}\ x\ 100$$

Heuristics

- **FCFS**: First come first serve
- **R-1**: One patient from randomly chosen type is scanned
- **R-2**: One patient randomly chosen from all waiting patients (favors types with more people waiting)
- **O-1**: Priority
  - OP
  - NCEP
  - IP
- **O-2**: Priority:
  - OP
  - IP
  - NCEP

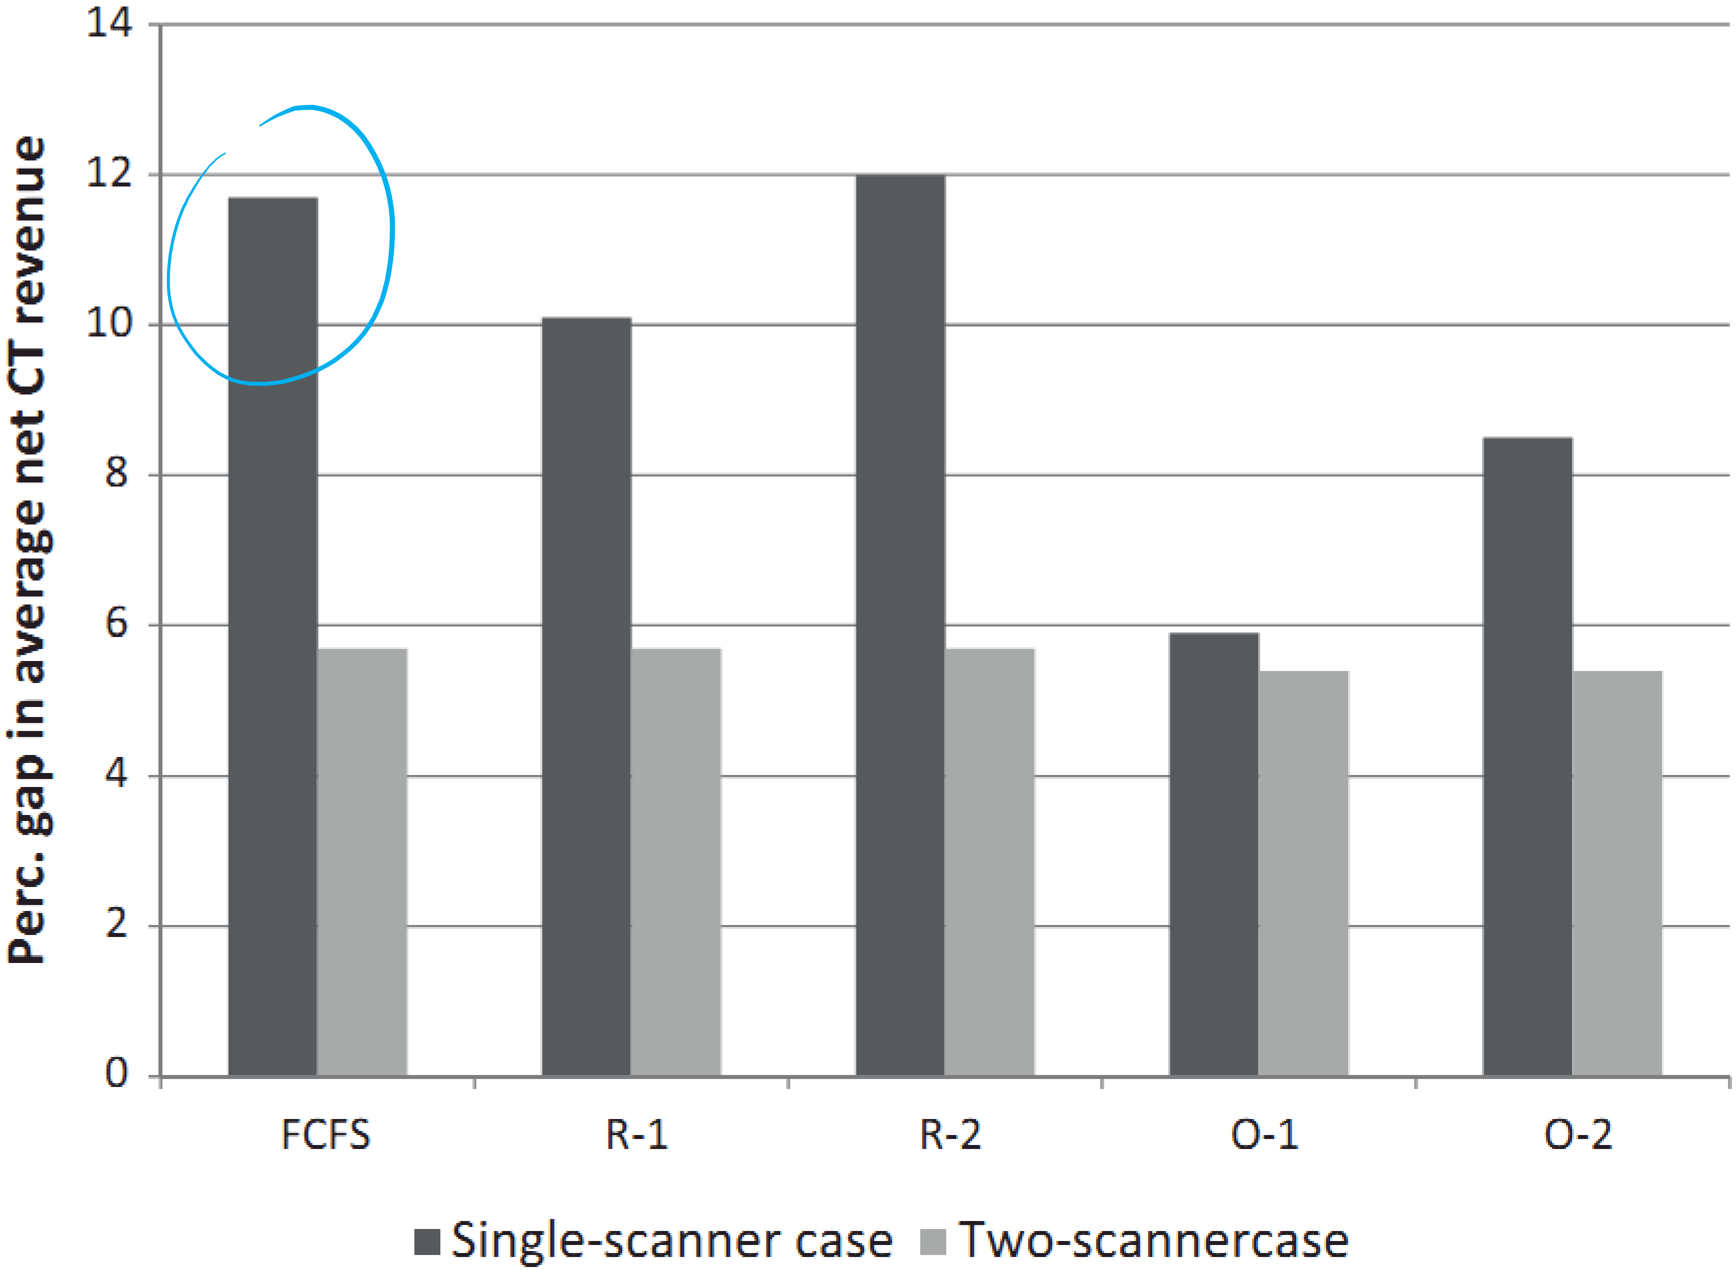
Perc. gap in average net CT revenue
14
12
10
8
6
4
2
0
FCFS
R-1
R-2
O-1
O-2
Single-scanner case
Two-scannercase

priority to OPs

Number of patients not scanned

Table 5
Number of patients not receiving scans by the end of the day under different policies, averaged over all thirty two scenarios.

Different cases	Average number not scanned					
	Optimal policy	FCFS	R-1	R-2	O-1	O-2
OPs						
Single-scanner	3.38	3.50	3.27	3.62	1.73	1.73
Two-scanner	0.72	0.63	0.52	0.64	0	0
IPs						
Single-scanner	10.13	9.97	10.57	9.85	12.01	11.14
Two-scanner	1.19	1.39	1.60	1.37	2.33	1.10
NCEPs						
Single-scanner	1.94	1.99	1.62	1.99	1.71	2.58
Two-scanner	0.51	0.39	0.29	0.41	0.08	1.31

Waiting-time

Table 6
Average waiting-time in minutes of patients before service over all thirty two scenarios.

Different cases	Average waiting-time					
	Optimal policy	FCFS	R-1	R-2	O-1	O-2
OPs						
Single-scanner	28	80	74	70	45	184
Two-scanner	3	4	3	4	0	0
IPs						
Single-scanner	76	112	95	107	60	245
Two-scanner	4	3	3	3	5	3
NCEPs						
Single-scanner	24	56	56	44	36	3
Two-scanner	12	9	8	10	3	20

Single-scanner

bright HIGH
dark Low

HIGH +50%
base rate
Low -50%

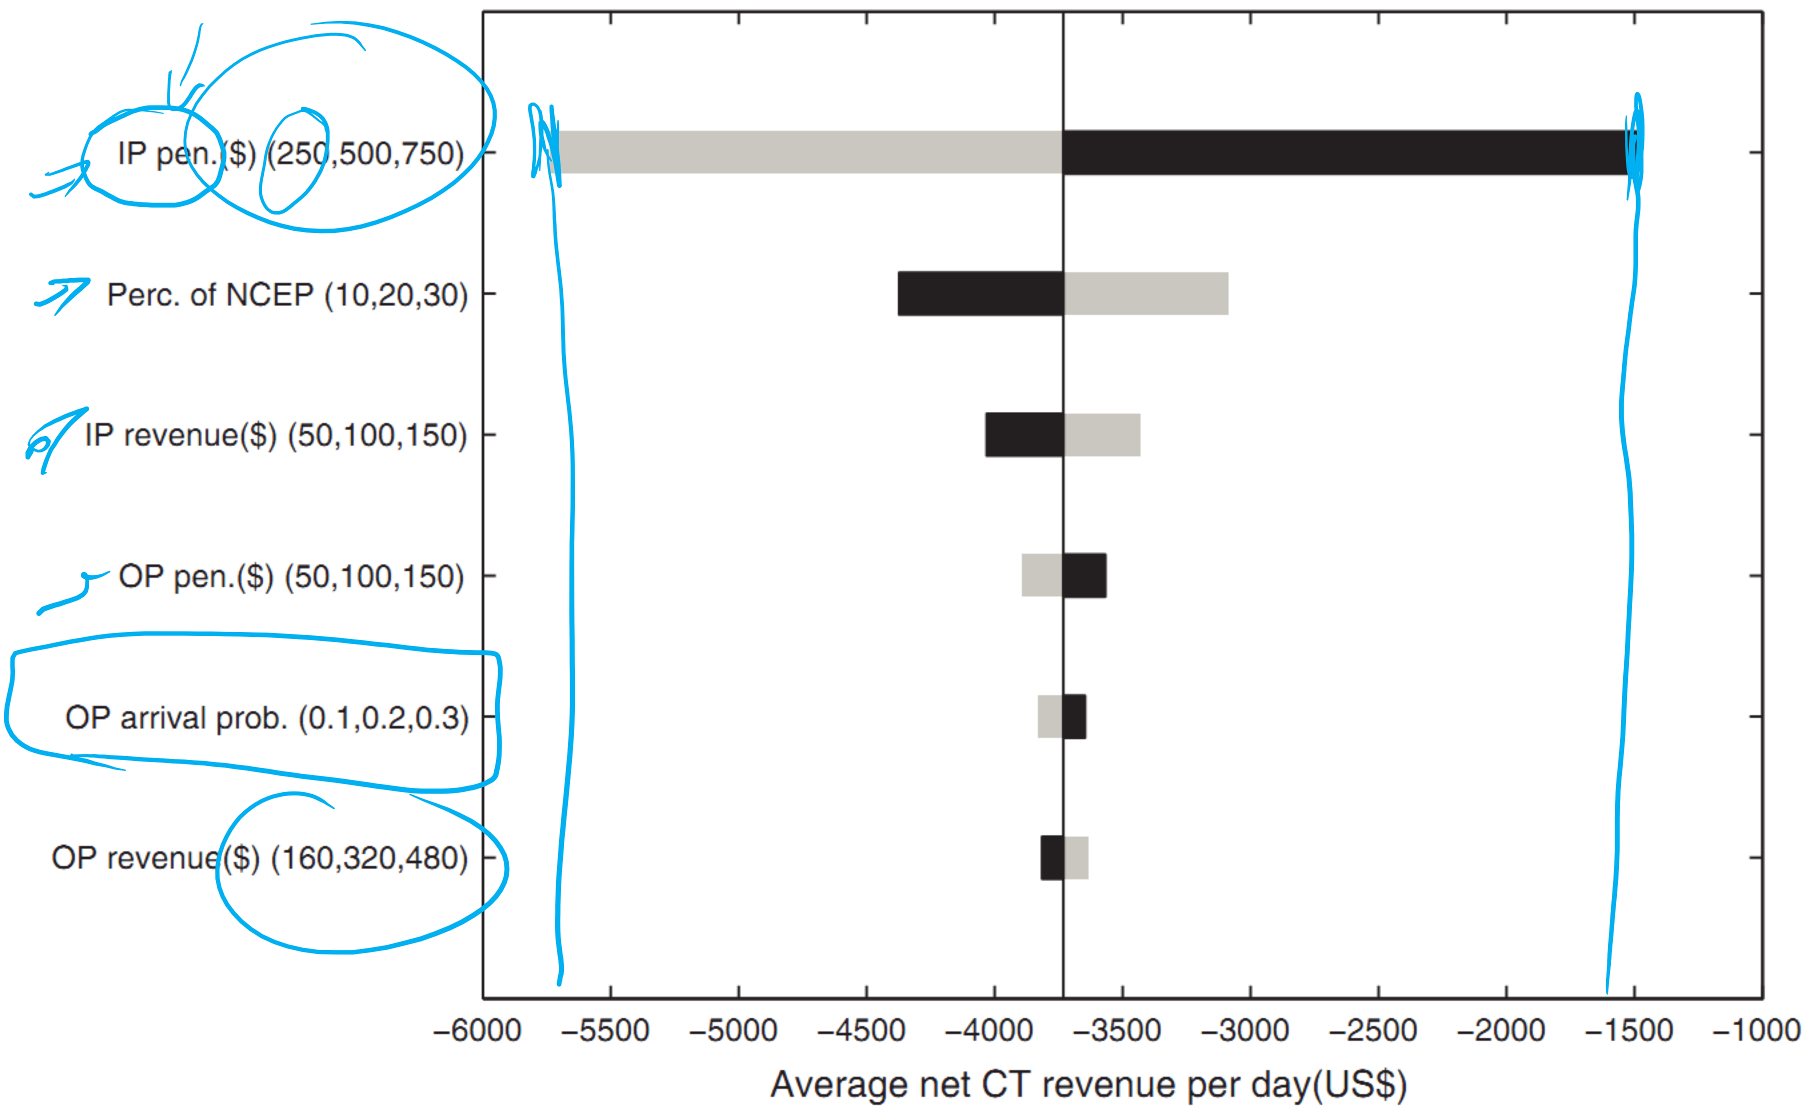

Two-scanner

bright HIGH
dark Low

HIGH +50%
base rate
Low -50%

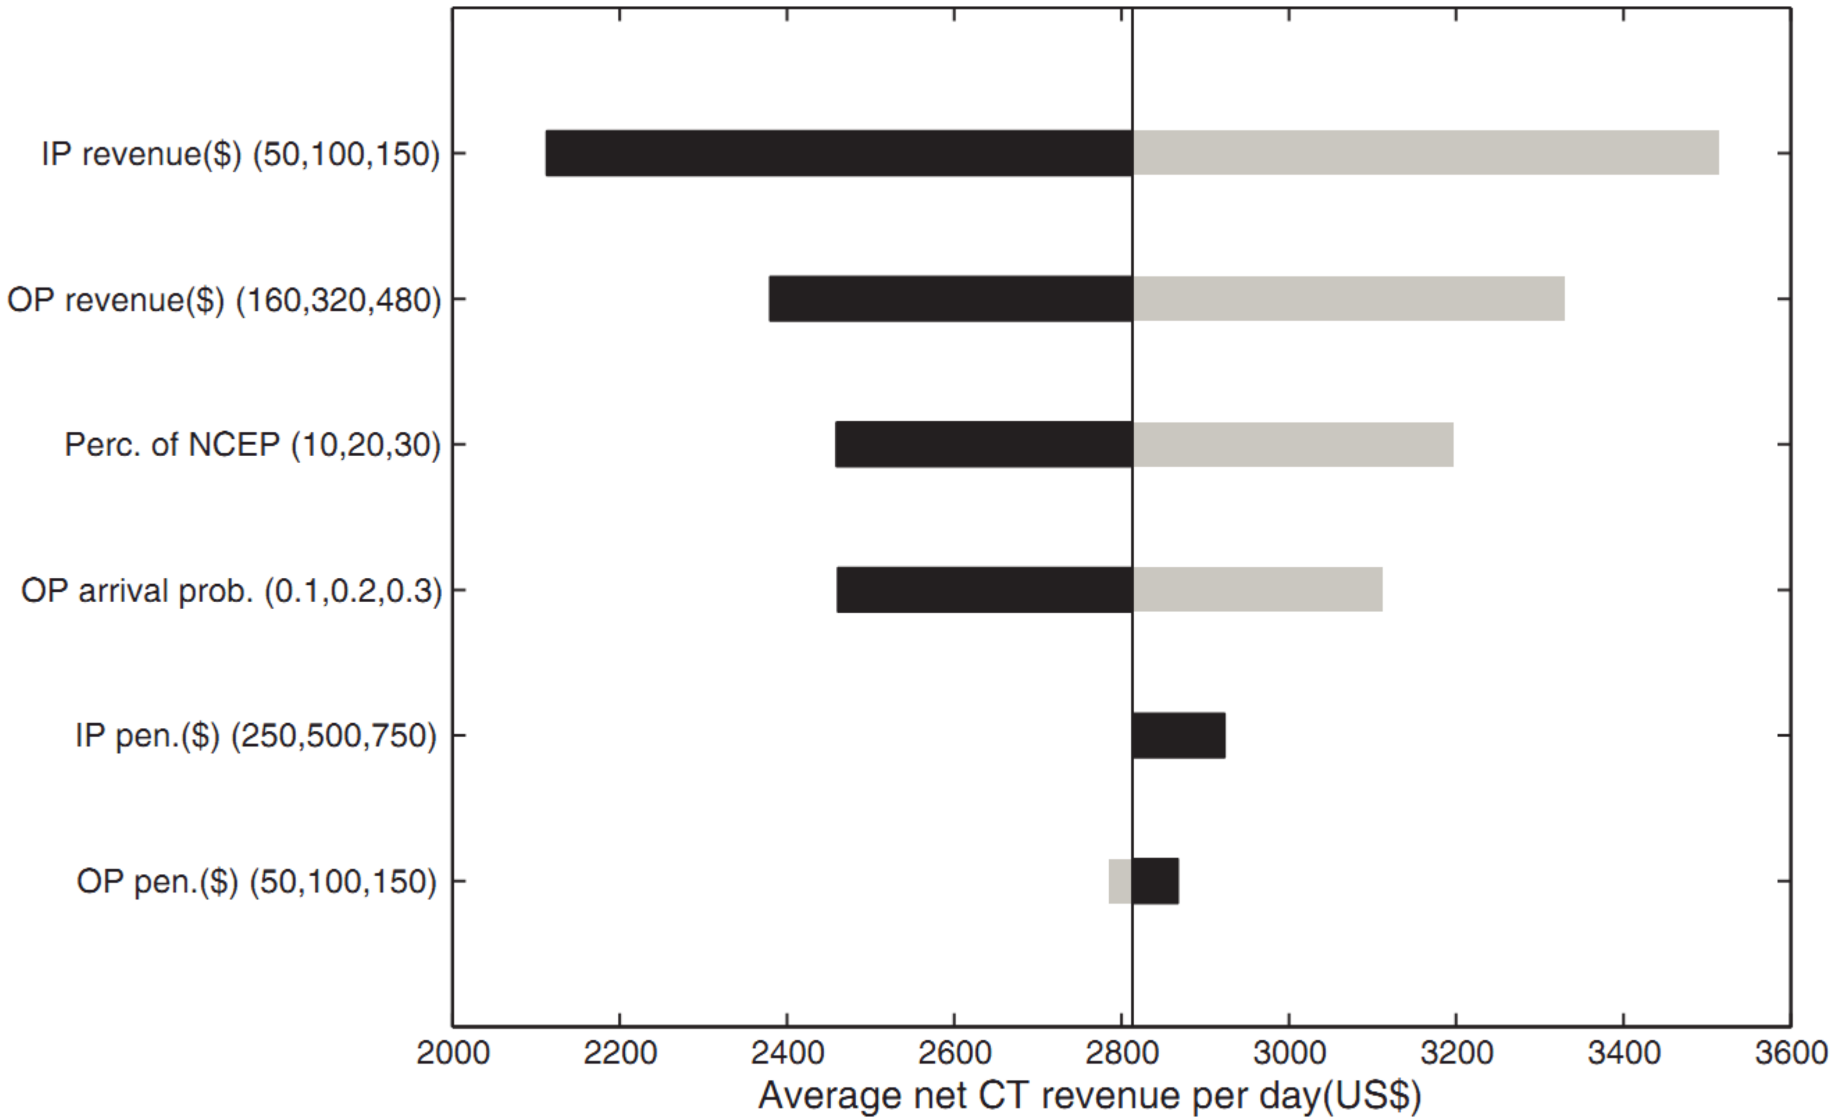

Sample Policy n=12, NCEP=5

CPE = 0

IPs →

OPs ↓

state {0, 0, 3, 5}

6	5	4	4	4	4	4	4	4	4	4	4
3	2	2	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4

i-1 OPs ; J-1 IPs

6 → scan two NCEPs

5 → scan one IP and one NCEP

4 → " two IPs

3 → one OP and one NCEP

2 → one OP and one IP

1 → " Two OPs

action to be performed in state {0, 0, 3, 5}

CPE OP IP NCEP

Student selected questions 2021

Students' (selected) questions

- Assumptions
- - why only 1 or 2 scanners. how does result change for more machines?
- - In reality, due to possible equipment failure or operator leave, the situation of numbers of CT machine can vary. If we take this into consideration, how will this effect the output optimal policy?
- - what if there are fewer or more than 20 timeslots?
- - why only 4 types of patients? Too generic.
- - Is using a fixed time for all CT scans realistic? Is there a distribution that could be used to represent this instead?
- - How did they determine that a four-month period is sufficient for deriving the arrival probabilities of IPs and Eps?

Students' (selected) questions

- Models
- why finite horizon MDP and not infinite horizon? in general how to choose between the two types?
- neural network (advantages of MDP with respect to NN?) - reinforcement learning?
- - dynamic programming (as it resembles jobs scheduling/load balancing problem)
- - Can the model take into account human suffering?
- - Patient that starts as non critical but turns to be critical
- - why did they not use value iteration here?
- - What if instead of having 0 or 1 "additional patients waiting" for each patient type and for each time slot, we have "additional patients waiting" as a random variable that can be more than 1? This more accurately reflects reality. Would this have made the methodology more complicated

Students' (selected) questions

- METRICS / EXPERIMENTS / RESULTs… IMPLEMENTATION - USE IN REAL LIFE
- evaluation metrics make sense? is the reward designed only to maximize revenue? Probably optimizing for patient care? Is it possible to measure ethics in a model?
- - Will computation costs ever be an issue with an MDP solution?
- - Is there evidence that such a theoretical analysis would work in practice?
- - Is this method used in practice nowadays?

Students' (selected) questions

- Others
- Since they already have the optimal policy, why do they develop other decision rules?
- In section 4.1, there is a paired t-test together with standard significance level. What do those terms mean? What can they tell us in this case?
- - paper is 10 years old. What is the state of the art for this problem?

More than 50 citations

About 50 results (**0.03** sec)

A Markov decision process approach to multi-category patient scheduling in a diagnostic facility

Search within citing articles

Taxonomic classification of planning decisions in health care: a structured review of the state of the art in OR/MS

PJH Hulshof, N Kortbeek, RJ Boucherie… - Health …, 2012 - orsociety.tandfonline.com

We provide a comprehensive overview of the typical decisions to be made in resource capacity planning and control in health care, and a structured review of relevant articles from the field of Operations Research and Management Sciences (OR/MS) for each planning …

Cited by 273 Related articles All 17 versions Web of Science: 103

[HTML] t
Full Vie

[HTML] Artificial intelligence framework for simulating clinical decision-making: A Markov decision process approach

CC Bennett, K Hauser - Artificial intelligence in medicine, 2013 - Elsevier

Objective In the modern healthcare system, rapidly expanding costs/complexity, the growing myriad of treatment options, and exploding information streams that often do not effectively reach the front lines hinder the ability to choose optimal treatment decisions over time. The …

Cited by 156 Related articles All 13 versions Web of Science: 51

[HTML] s
UBC el

[HTML] Outpatient appointment systems in healthcare: A review of optimization studies

A Ahmadi-Javid, Z Jalali, KJ Klassen - European Journal of Operational …, 2017 - Elsevier

Outpatient appointment scheduling problems have recently gained increasing attention. This paper provides a comprehensive review of recent analytical and numerical optimization

[HTML] s
UBC el

12 from 2018….

Google Scholar

Articles

12 results (**0.02** sec)

Any time
Since 2019
Since 2018
Since 2015
Custom range...

Sort by relevance
Sort by date

include citations

Create alert

A Markov decision process approach to multi-category patient scheduling in a diagnostic facility

Search within citing articles

[HTML] Literature review on multi-appointment scheduling problems in hospitals
J Marynissen, E Demeulemeester - European Journal of Operational …, 2019 - Elsevier
This paper presents a review of the literature on multi-appointment scheduling problems in hospitals. In these problems, patients need to sequentially visit multiple resource types in a hospital setting so they can receive treatment or be diagnosed. Therefore, each patient is …
Cited by 9 Related articles All 6 versions Web of Science: 3
[HTML] sciencedirect.com
UBC eLink

Coordinated Scheduling for a Multi-server Network in Outpatient Pre-operative Care
D Wang, DJ Morrice, K Muthuraman… - Production and …, 2018 - Wiley Online Library
Many parts of the healthcare system remain fragmented and outpatient surgical care is no exception. In this study, we develop a coordinated pre-operative scheduling approach between Anesthesiology and Internal Medicine to optimize patients' medical conditions prior …
Cited by 3 Related articles All 5 versions Web of Science: 2
[PDF] wiley.com
Full View

[HTML] Coordinating patient preferences through automated negotiation: A multiagent systems model for diagnostic services scheduling
J Gao, T Wong, C Wang - Advanced Engineering Informatics, 2019 - Elsevier
This paper presents a multiagent systems model for patient diagnostic services scheduling. We assume a decentralized environment in which patients are modeled as self-interested agents who behave strategically to advance their own benefits rather than the system wide …
Related articles
[HTML] sciencedirect.com
UBC eLink

[HTML] Population-level intervention and information collection in dynamic healthcare policy
LE Cipriano, TA Weber - Health care management science, 2018 - Springer
We develop a general framework for optimal health policy design in a dynamic setting. We consider a hypothetical medical intervention for a cohort of patients where one parameter varies across cohorts with imperfectly observable linear dynamics. We seek to identify the …
Cited by 5 Related articles All 12 versions Web of Science: 3
[HTML] springer.com
Full View

Optimizing Daily Service Scheduling for Medical Diagnostic Equipment Considering Patient Satisfaction and Hospital Revenue
G Du, X Li, H Hu, X Ouyang - Sustainability, 2018 - mdpi.com
Under the background of the unbalanced supply and demand of medical diagnostic equipment and rising health care costs, this study aims to optimize the service scheduling for medical diagnostic equipment so as to improve patient satisfaction by ensuring the …
Cited by 1 Related articles All 8 versions
[PDF] mdpi.com
Full View

16 documents have cited:

A Markov decision process approach to multi-category patient scheduling in a diagnostic facility
Gocgun Y., Bresnahan B.W., Ghate A., Gunn M.L.
(2011) *Artificial Intelligence in Medicine*, 53 (2), pp. 73-81.
Is cited by: Set feed

16 documents that cite: Analyze search results

Sort on: **Date** Cited by

Search within results...

Export | Download | View citation overview | View cited by | More...

Refine

Limit to   Exclude

Year
- 2017 (4)
- 2016 (3)
- 2015 (2)
- 2014 (3)
- 2013 (3)

Author name
- Gocgun, Y. (2)
- Abe, T.K. (1)
- Agus, J. (1)
- Ahmadi-Javid, A. (1)
- Barz, C. (1)

Subject area
- Computer Science (6)
- Decision Sciences (6)
- Engineering (6)
- Medicine (6)
- Mathematics (5)

Strategic level proton therapy patient admission planning: a Markov decision process modeling approach View abstract | Related documents	Gedik, R., Zhang, S., Rainwater, C.	2017	Health Care Management Science	1
Outpatient appointment systems in healthcare: A review of optimization studies View abstract | Related documents	Ahmadi-Javid, A., Jalali, Z., Klassen, K.J.	2017	European Journal of Operational Research	1
A goal programming and simulation based study for overall process improvement in an Indian hospital View abstract | Related documents	Sah, B., Titiyal, R., Sonia	2017	International Journal of Services and Operations Management	0
A markov decision process model of allocating emergency medical resource among multi-priority injuries View abstract | Related documents	Ni, Y., Wang, K., Zhao, L.	2017	International Journal of Mathematics in Operational Research	0
Optimal Dynamic Outpatient Scheduling for a Diagnostic Facility with Two Waiting Time Targets View abstract | Related documents	Geng, N., Xie, X.	2016	IEEE Transactions on Automatic Control	1
Operations research applications in hospital operations: Part III View abstract | Related documents	Abe, T.K., Beamon, B.M., Storch, R.L., Agus, J.	2016	IIE Transactions on Healthcare Systems Engineering	0
Stochastic modeling of multi-state disease dynamics under random environments	Manoharan, M., Xavier, T.D.	2016	Proceedings - 2nd International Symposium on Stochastic Models in Reliability Engineering, Life Science, and Operations Management, SMRLO	0

ocess approach to multi-category patient scheduling in a diagnostic facility
B.W., Ghate A., Gunn M.L.
ence in Medicine, 53 (2), pp. 73-81.
d

nat cite: Analyze search results

Sort on: **Date** Cited by ...

ults...

Export | Download | View citation overview | View cited by | More...

Exclude

Strategic level proton therapy patient admission planning: a Markov decision process modeling approach View abstract | Related documents	Gedik, R., Zhang, S., Rainwater, C.	2017	Health Care Management Science	1
Outpatient appointment systems in healthcare: A review of optimization studies View abstract | Related documents	Ahmadi-Javid, A., Jalali, Z., Klassen, K.J.	2017	European Journal of Operational Research	1
A goal programming and simulation based study for overall process improvement in an Indian hospital View abstract | Related documents	Sah, B., Titiyal, R., Sonia	2017	International Journal of Services and Operations Management	0
A markov decision process model of allocating emergency medical resource among multi-priority injuries View abstract | Related documents	Ni, Y., Wang, K., Zhao, L.	2017	International Journal of Mathematics in Operational Research	0
Optimal Dynamic Outpatient Scheduling for a Diagnostic Facility with Two Waiting Time Targets View abstract | Related documents	Geng, N., Xie, X.	2016	IEEE Transactions on Automatic Control	1
Operations research applications in hospital operations: Part III View abstract | Related documents	Abe, T.K., Beamon, B.M., Storch, R.L., Agus, J.	2016	IIE Transactions on Healthcare Systems Engineering	0
Stochastic modeling of multi-state disease dynamics under random environments	Manoharan, M., Xavier, T.D.	2016	Proceedings - 2nd International Symposium on Stochastic Models in Reliability Engineering, Life Science, and Operations Management,	0

5)
5)

Question from students (2017)

- Would the model cause ethical problems in hospitals? Is revenue a good metric of performance if we put life and death situations into consideration?

- Finite vs. infinite
  - Simplicity. Lots of uncertainty about what can happen overnight
  - Non stationary process – best action depends on time
- Use machine learning / reinforcement learning?
- Arrival Probabilities (seasonal trends? More than one patient of each type?)
- Only comparison with simple heuristics
- More scanners Why only 1 and 2 scanners?
- Modeling more patient types (urgency) / different hospital….. can easily extend the model
- Only data from one Hospital (general?)
- Uniform slot length (realistic?) Finer granularity of the time slots
- Modeling even more uncertainty "Accidents happen randomly without any pattern." "Scanner not working"
- 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?

- Benefits classifying more patient types? Could same solution be applied to scheduling other functions of the hospital?
- How would this model handle two CEPs that came in at the same time? *Randomly Push one to the next slot ☹*
- What happens if you add a sudden influx of patients? Example, due to a nearby accident. Will it still perform better than the heuristics?
- Transfer model to other facilities? *Yes…*
- Discount factor 1? *Yes*
- This work failed to take into account human suffering, or the urgency of scans for in and out patients. Could the reward function to tailored to include such concepts or is it beyond the capabilities of the model?
- This model is specific to the target hospital
- Operational Cost of Implementing the policy (take into account): *compute the policy vs. apply the policy*

Question Types from students

- Finite vs. infinite
  - Simplicity. Lots of uncertainty about what can happen overnight
  - Non stationary process – best action depends on time
- Arrival Probabilities
- More scanners
- Modeling more patient types (urgency) / different hospital….. can easily extend the model, Only data from one Hospital (general?)
- Uniform slot length (realistic?)
- the probability distribution of the time for CT scans to be completed rather than to make the assumption that they are all of fixed duration? Finer granularity of the time slots
- Operational Cost of Implementing the policy (take into account): *compute the policy vs. apply the policy*
- Modeling even more uncertainty "Accidents happen randomly without any pattern." "Scanner not working"
- 2 patients at once (need to collect all the prob and consider those in the transition prob)
- P-value
- Why no VI?
- Used in practice ?

- Other models: Is it better to use continuous Markov Chain and queuing theory in analyzing this scheduling problem?
- How would this model handle two CEPs that came in at the same time? Randomly Push one to the next slot ☹
- How does approximate dynamic programming compare to value iteration? *(approximate method, can deal with bigger models but not optimal)*
- Transfer model to other facilities? *Yes…*
- Discount factor 1? *Yes*
- This work failed to take into account human suffering, or the urgency of of scans for in and out patients. Could the reward function to tailored to include such nebulous concepts or is it beyond the capabilities of the model?
- This model is specific to the target hospital
- I think outperforming other MDP-based models can better illustrate the effectiveness of this model's features, so are the choices of comparison methods good in this paper?
- *First step showing that sound probabilistic models can be build and outperform heuristics then you can do the above*