Machine learning-based algorithms for dynamic patient scheduling problems with uncertainty

Tu-San Pham, Antoine Legrain,

Patrick De Causmaecker, Louis-Martin Rousseau

November 4th 2022 - Focus Period Linköping

OPTIMIZATION
IN THE HEALTHCARE DOMAIN

HANALOG

Canada Research Chair in Analytics and Logistics in Healthcare

Optimizing treatment planning and delivery in healthcare

Time to tackle health-care wait times in Canada
— *Mackenzie Moir, Bacchus Barua*

https://nationalpost.com › news › canada › why-canada...
Why Canada's hospital capacity was so easily overwhelmed ...
Jan 17, 2022 — The latest numbers from the OECD show Canada with just **one hospital bed for every 400 citizens**, a ratio that put us in the bottom tier of OECD ...

https://uwaterloo.ca › sites › files › uploads › files
SeniorS and HouSing: THe CHallenge aHead
As **Canada's** aging population grows, and as. **Canadians** live longer, there is an **overwhelming** desire among **seniors** to "age

Mental health among healthcare workers in Canada during the COVID-19 pandemic
Release date: 2021-02-02

Overwhelming evidence—it's time to fix Canadian health care
— *Appeared in the Montreal Gazette, September 9, 2022*

https://cmajnews.com › 2020/08/28 › access-qa-1095895
How can Canada improve worsening wait times? - CMAJ News
Aug 28, 2020 — Simpson: **Canadians** have accepted that some **delay** is reasonable in order to have an equitable **health care** system, and that's not necessarily a ...

https://hospitalnews.com › Topics › Health Care Po
Canada ranks last on number of hospital beds, wait times
For example, Canada ranks 26th (out of 28 countries) for the number of doctors (2.8 per 1,000 people) and 26th (out of **27**) for the number of hospital beds ...

UNCERTAINTY IN HEALTHCARE SCHEDULING

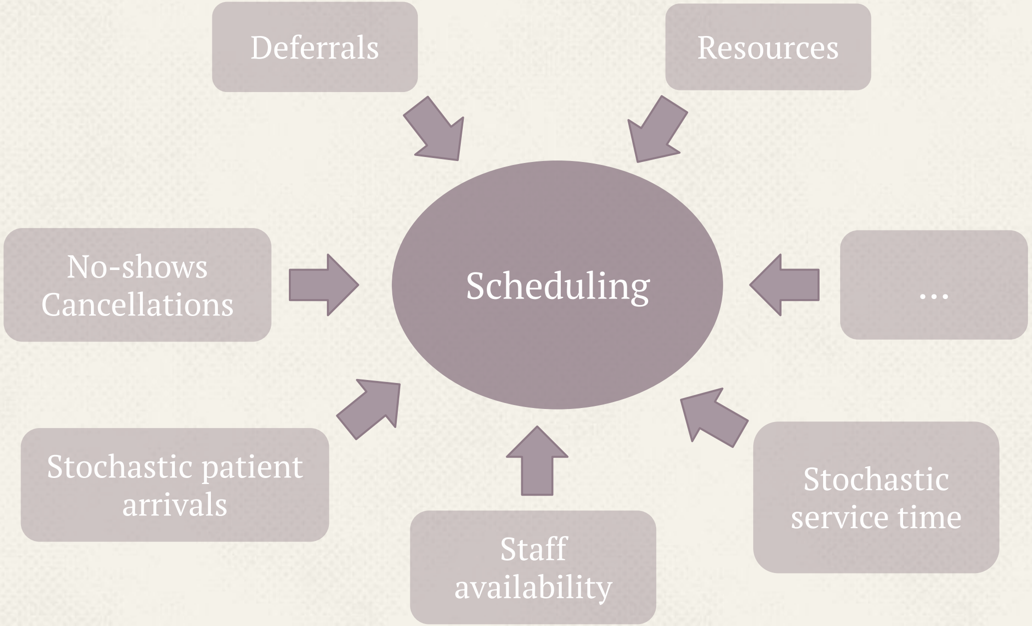

- Dynamic problems
- Online decisions

- Stochastic optimization
- MDP - (approximate) dynamic programming
- Simulation-based

Machine-learning based approaches

Radiotherapy scheduling for cancer treatments

Cancer incidence (2020)

- *10.000.000 deaths*
- *1 out of 6 deaths*

Challenges

- *Growing and aging population*

How to reduce death rates?

- *Early detection*
- *Early treatment*

Optimize treatment scheduling

https://www.who.int/news-room/fact-sheets/detail/cancer

CANCER TREATMENT

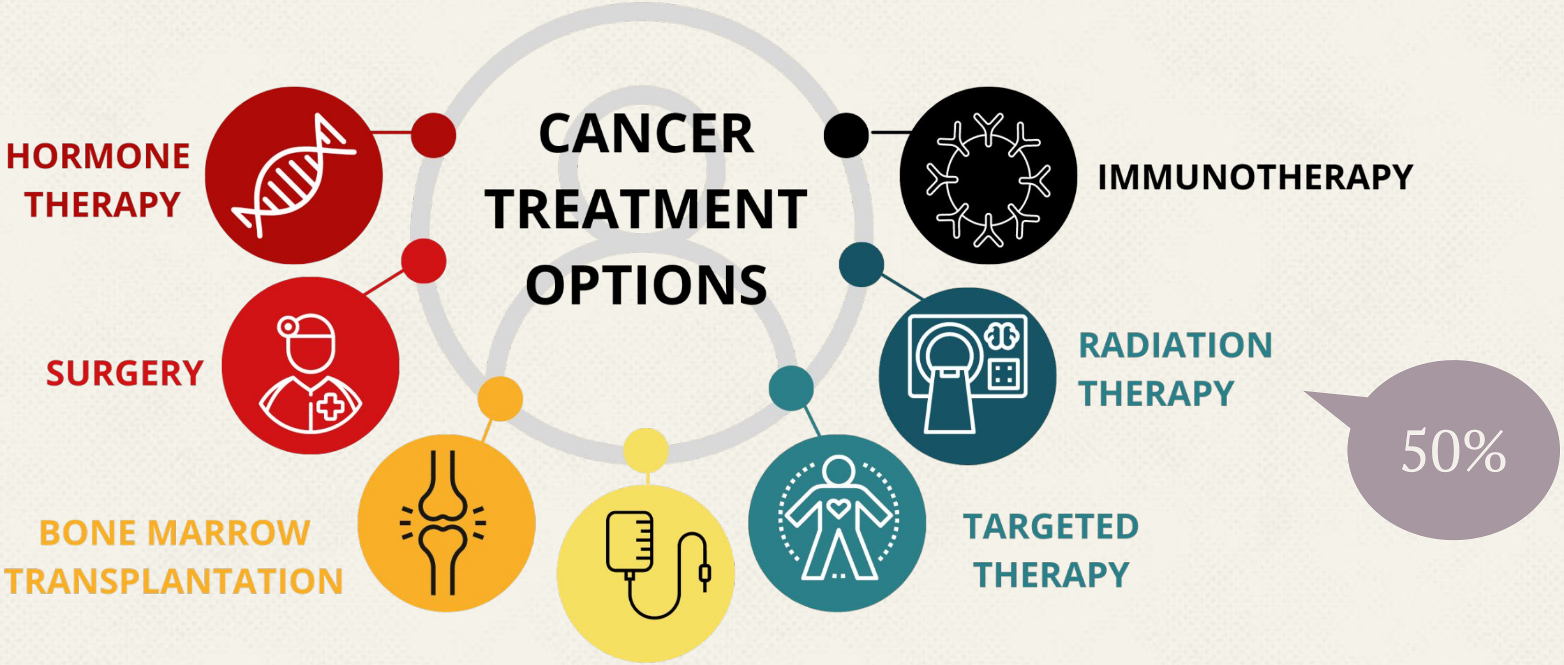

Optimizing radiotherapy treatment schedules to reduce patients’ waiting times

RADIOTHERAPY TREATMENT

- Linear accelerators (linac)
- Fraction: a small dose of radiation
- Treatment plan
  - Multiple consecutive fractions
  - Fraction duration

Patient scheduling with multi-appointments, multi-resources

LITERATURE

- Markov Decision Process & Approximate Dynamic Programming (ADP)
  - Patrick et al. 2008
  - Saure et al. 2012, 2020 (3 linacs)
  - Gocgun 2018
- Stochastic Programming
  - Legrain et al. 2015 (2 linacs)

CHUM - CENTRE HOSPITALIER DE L'UNIVERSITÉ DE MONTRÉAL

10 linacs
5 generics
4 specialized
1 cyberknife

4400 consultations

3500 new patients

40.000 fractions

(2019)

GRAY

https://www.graysuite.com/

2021 INNOVE-ACTION
BEST STARTUP AWARD

CHUM - 2019

	Category	Percentages (%)	Treatment deadline (days)	Percentage of overdue treatment (%)	Average waiting time (days)
Palliative	**P1**	0.4	1	14.29	1.09
	P2	27.2	3	79.89	6.91
Curative	**P3**	41.4	14	74.55	18.11
	P4	31.0	28	29.89	22.59

Objective: minimizing overdue treatment and waiting time

THE MAIN CHALLENGES

- Preserving linac capacity
  - Reserved capacity vs occupancy rate?
- A prediction-based approach
  - Learn to delay low-priority patients

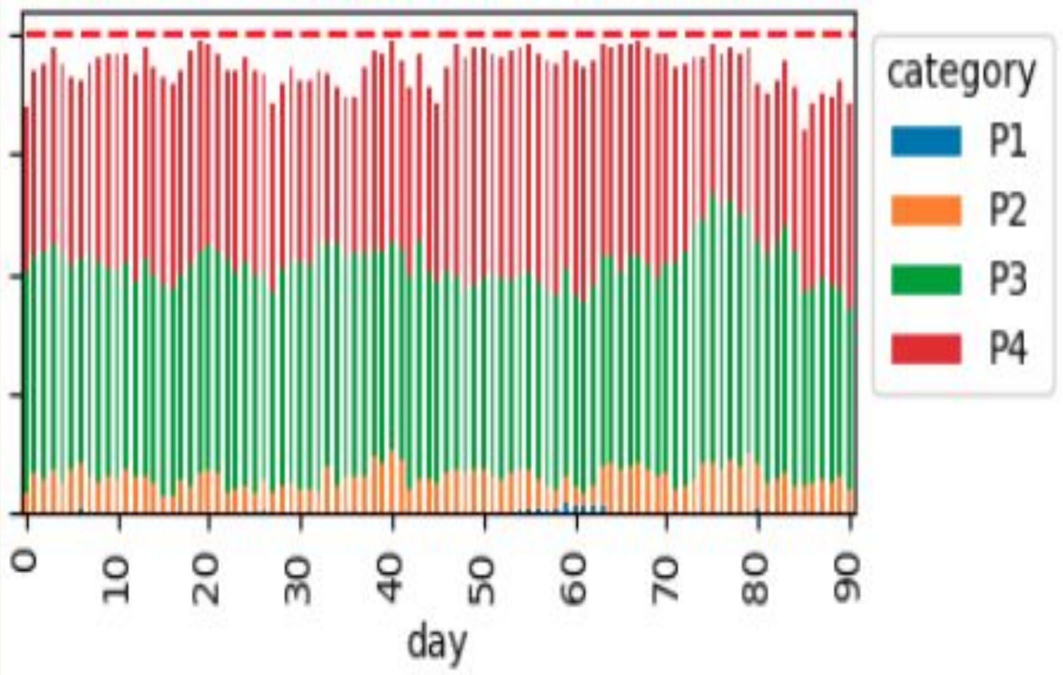

Scheduling strategies

- *Batch scheduling*
- *Offline scheduling*
- *Online scheduling*
- *Prediction-based approach*

BATCH SCHEDULING

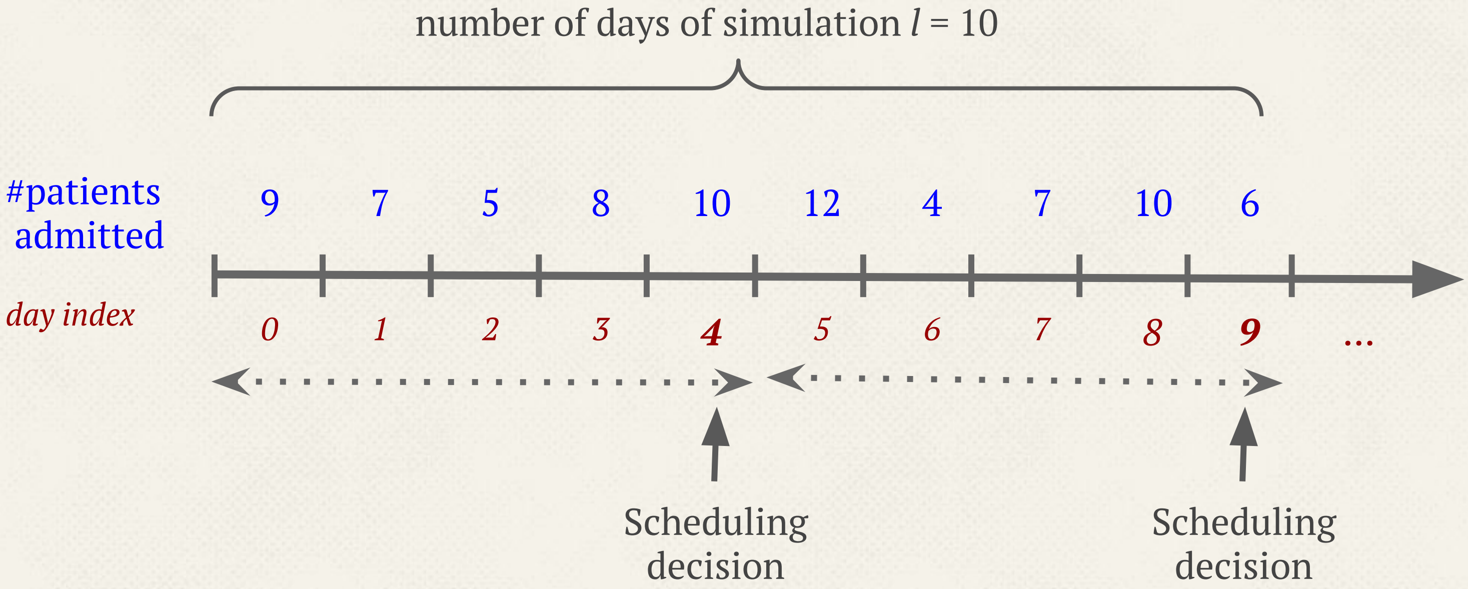

Palliative patients: schedule at arrival

AN INTEGER PROGRAMMING MODEL FOR BATCH SCHEDULING

$$x_{tl}^i = \begin{cases} 1 & \text{if patient } i \text{ receives their treatment on day } t, \text{ linac } l \\ 0 & \text{otherwise} \end{cases}$$

$$\text{minimize} \sum_{i \in \hat{\mathcal{P}}} \sum_{t \in \mathcal{T}, t > a_i} \sum_{l \in \mathcal{L}} \omega_1 \boxed{(t - a_i) log(t - a_i + 1) x_{tl}^i}$$

$$+ \sum_{i \in \hat{\mathcal{P}}} \sum_{t \in \mathcal{T}, t > d_i} \sum_{l \in \mathcal{L}} \omega_2 \boxed{(t - d_i) log(t - d_i + 1) x_{tl}^i}$$

waiting time

overdue time

IP MODEL

$$\sum_{t \in \mathcal{T}} \sum_{l \in \mathcal{L}} x_{tl}^{i} = 1 \qquad \forall i \in \hat{\mathcal{P}}$$

assignment constraint

$$x_{tl}^{i} = 0 \qquad \forall i \in \hat{\mathcal{P}}, l \in \mathcal{L}, t \in \{0, \ldots, r_i - 1\}$$

ready date

$$\sum_{i \in \hat{\mathcal{P}}} \sum_{t'=max\{0,t-I_i+1\}}^{t} p_i x_{t'l}^{i} \leq \hat{C}_l^t \qquad \forall t \in \mathcal{T}, l \in \mathcal{L}$$

capacity constraints

$$\sum_{i \in \mathcal{P}^C} \sum_{t' \in \{t-I_i+1,\ldots,t\}} p_i x_{t'l}^{i} \leq max\{0, \hat{C}_l^t - \gamma C_l^t\} \qquad \forall t \in \mathcal{T}, l \in \mathcal{L}$$

reserved capacity

$$x_{tl}^{i} \in \{0, 1\} \qquad \forall i \in \hat{\mathcal{P}}, t \in \mathcal{T}, l \in \mathcal{L}$$

OFFLINE SCHEDULING - THE PERFECT SCENARIO

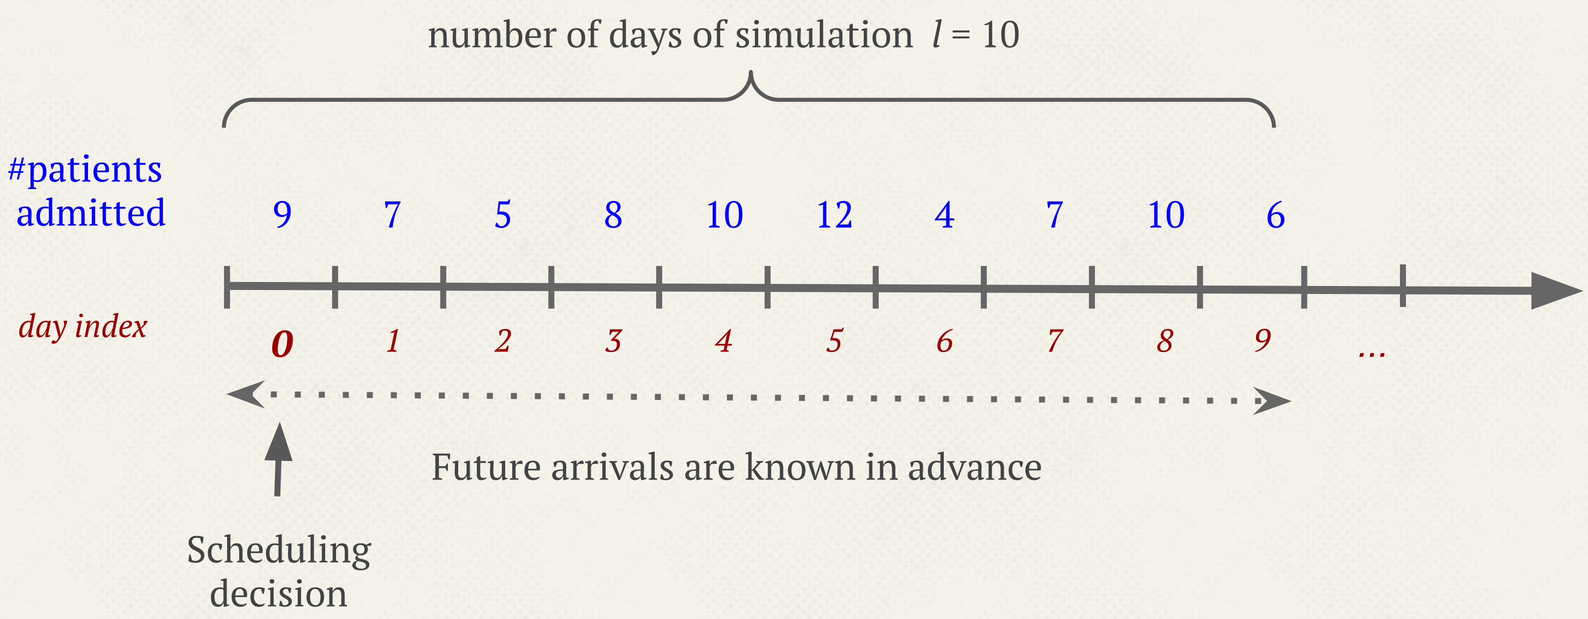

ONLINE SCHEDULING WITH A GREEDY HEURISTIC

- 1 linac, capacity 120 time slots
- a curative patient (P4) with 3 sessions, 10 time slots each

Looking for the first eligible date that can accommodate the whole treatment

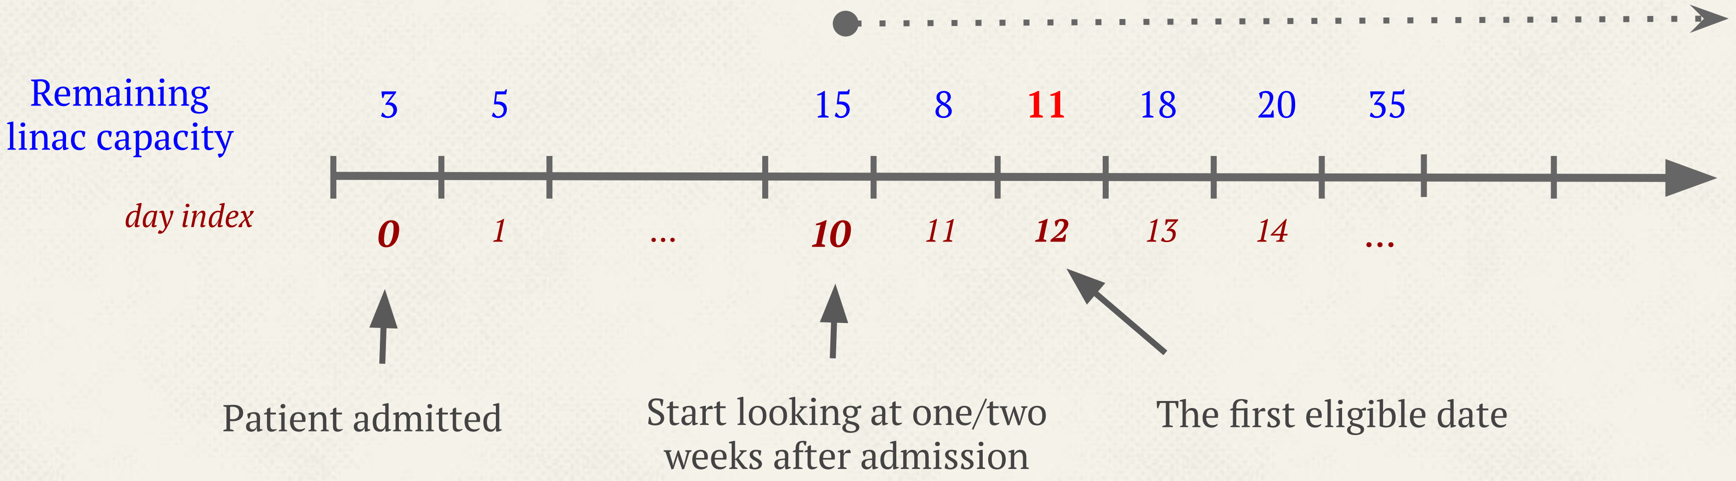

PREDICTION-BASED SCHEDULING

- 1 linac, capacity 120 time slots
- a curative patient with 3 sessions, 10 time slots each

Looking for the first eligible date that can accommodate the whole treatment

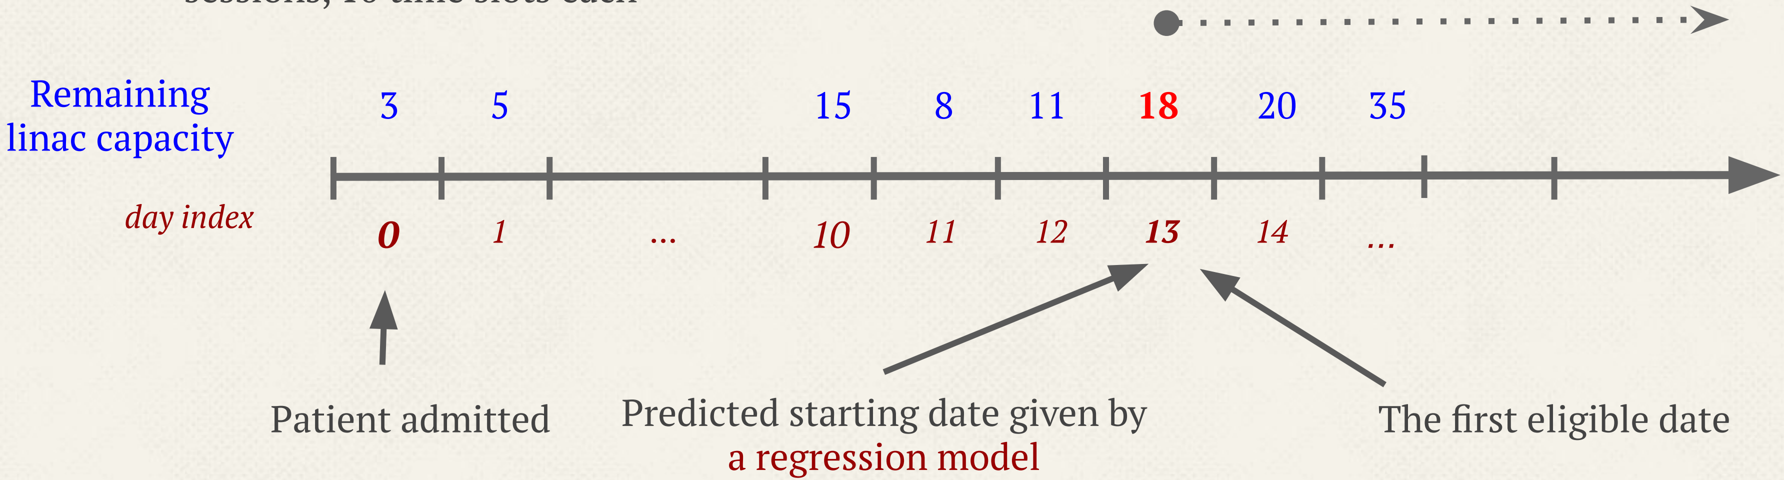

How do we predict a “good” starting date for a patient?

TRAINING THE REGRESSION MODEL

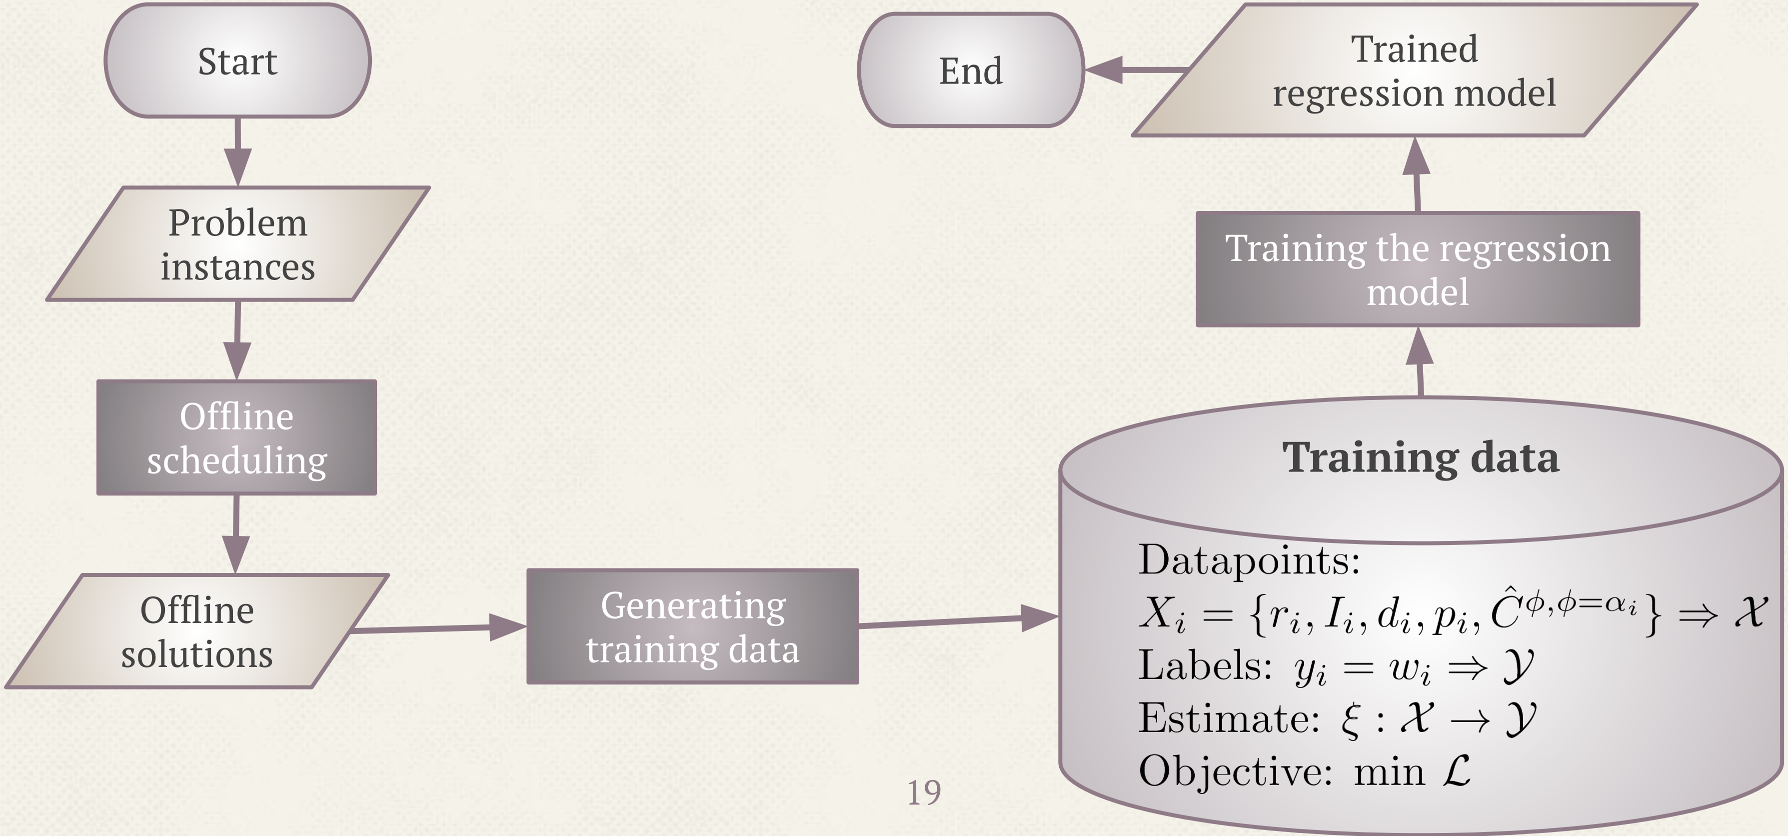

Numerical results

- *Data generation*
- *Model selection*
- *Results on simulated data*
- *Results on real data*
- *Explainability with SHAP values*

DATA GENERATION

- Patient arrivals: Poisson distribution
- Treatment plans: based on historical data
- Instance setting
  - Number of linacs
  - Arrival rate (average daily number of patients)
- For each instance setting: 500 instances
  - 400 for training the regression model
  - 100 for testing

PREDICTION MODELS

	Training time	Training		Testing	
		MSE	MAE	MSE	MAE
MLP	116.19	3.45	1.32	3.33	1.29
SGD	0.35	6.06	1.84	5.61	1.77
Lasso	0.44	5.97	1.81	5.52	1.74
ElasticNet	0.25	6.26	1.85	5.83	1.8
SVR	43.16	3.19	1.07	3.12	1.07
Decision Tree	0.84	2.41	0.48	6.59	1.4
Random forest	51	**0.38**	**0.39**	2.64	1.03
XGBoost	7.71	0.96	0.66	**2.44**	**0.97**

SCHEDULING STRATEGIES

	Scheduling strategy	Scheduling palliative patients	Scheduling curative patients
Batch scheduling	**Offline**	Scheduling once with all future arrivals known in advance	
	Daily	Every day	Every day
	Weekly	Every day	Every Friday
	Daily greedy	Every day	Every day
Online scheduling	**Greedy**	At admission	At admission
	Prediction-based	At admission	At admission

4 LINACS

Arrival rate of 5.0

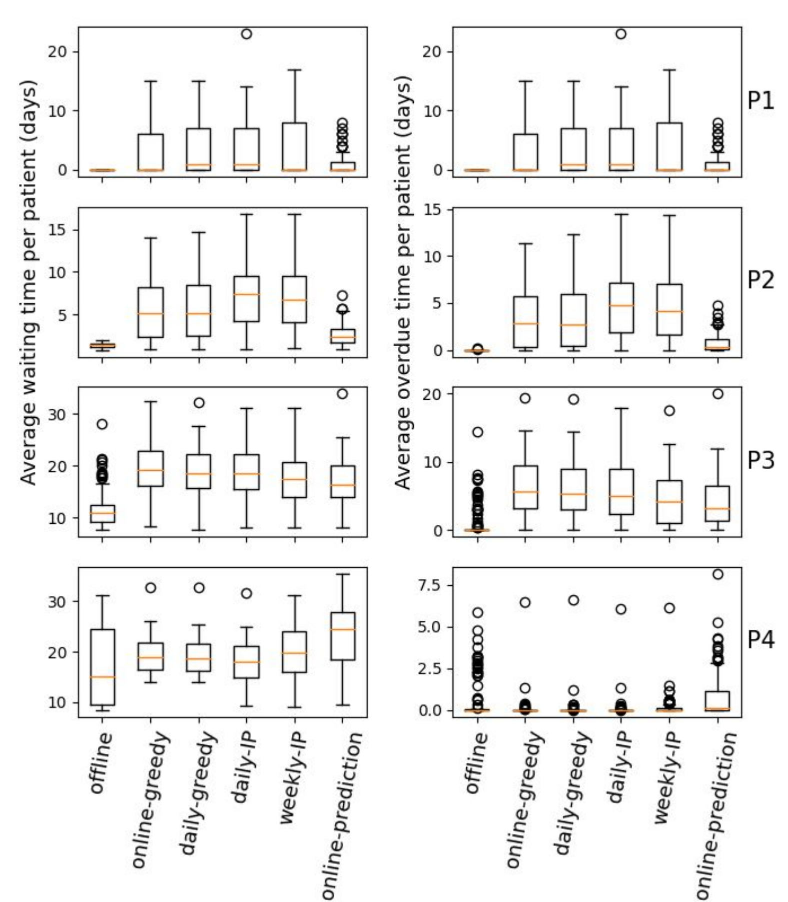

Arrival rate of 6.0

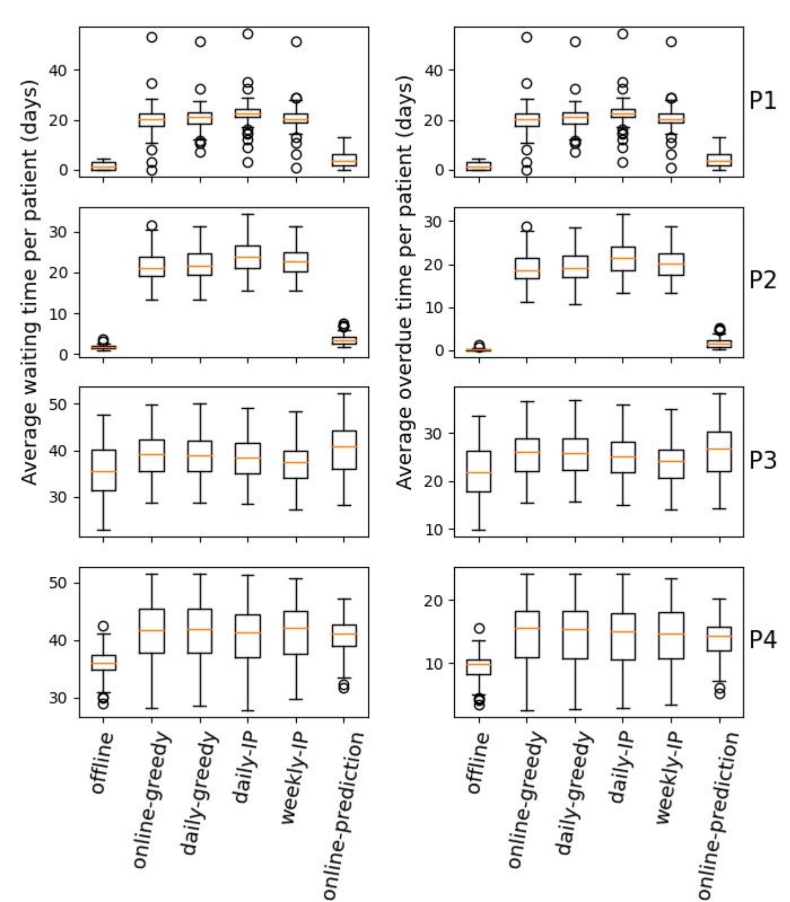

8 LINACS

Arrival rate of 10.0

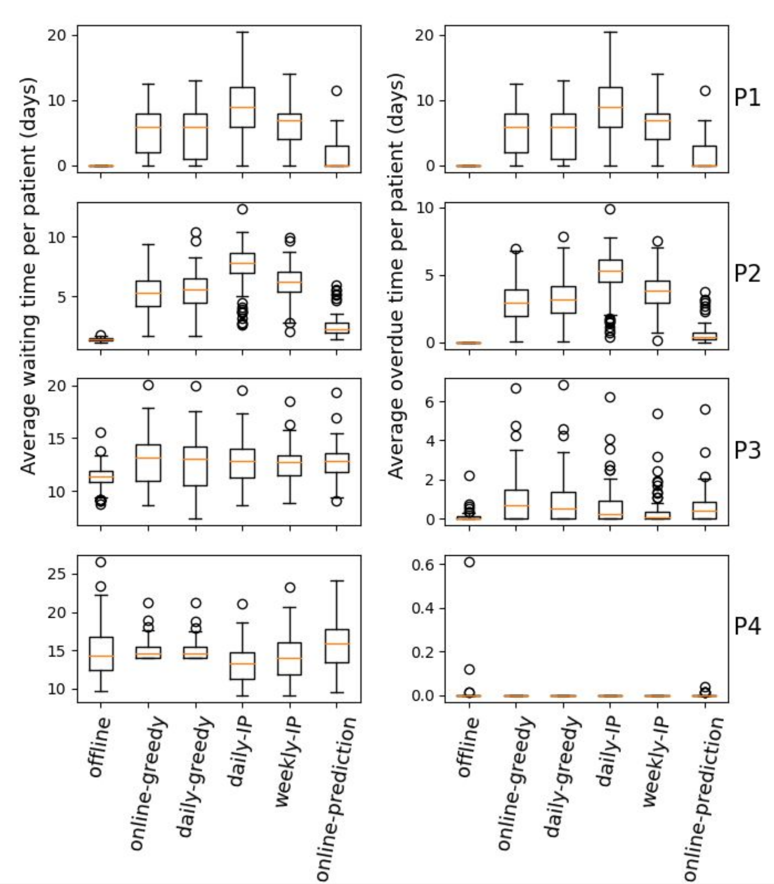

Arrival rate of 12.0

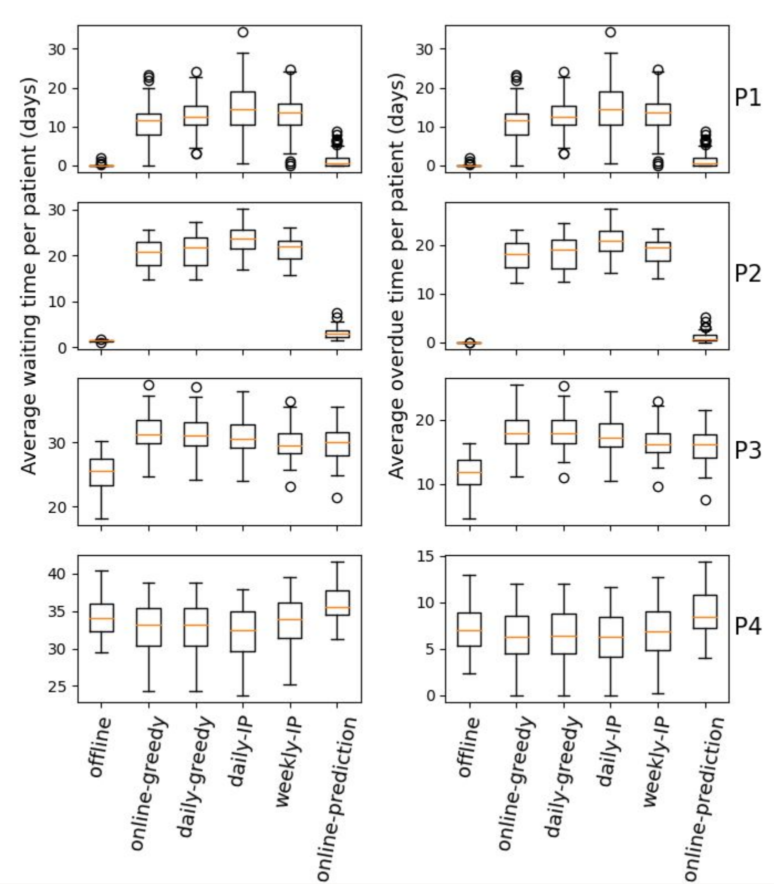

EXPERIMENT ON A REAL PATIENT FLOW

- 7 linacs operating 8 hours/day
- High fluctuation in arrival rate
  - Instance setting for training: arrival rate of 10.1
  - Simulation horizon: 30 days

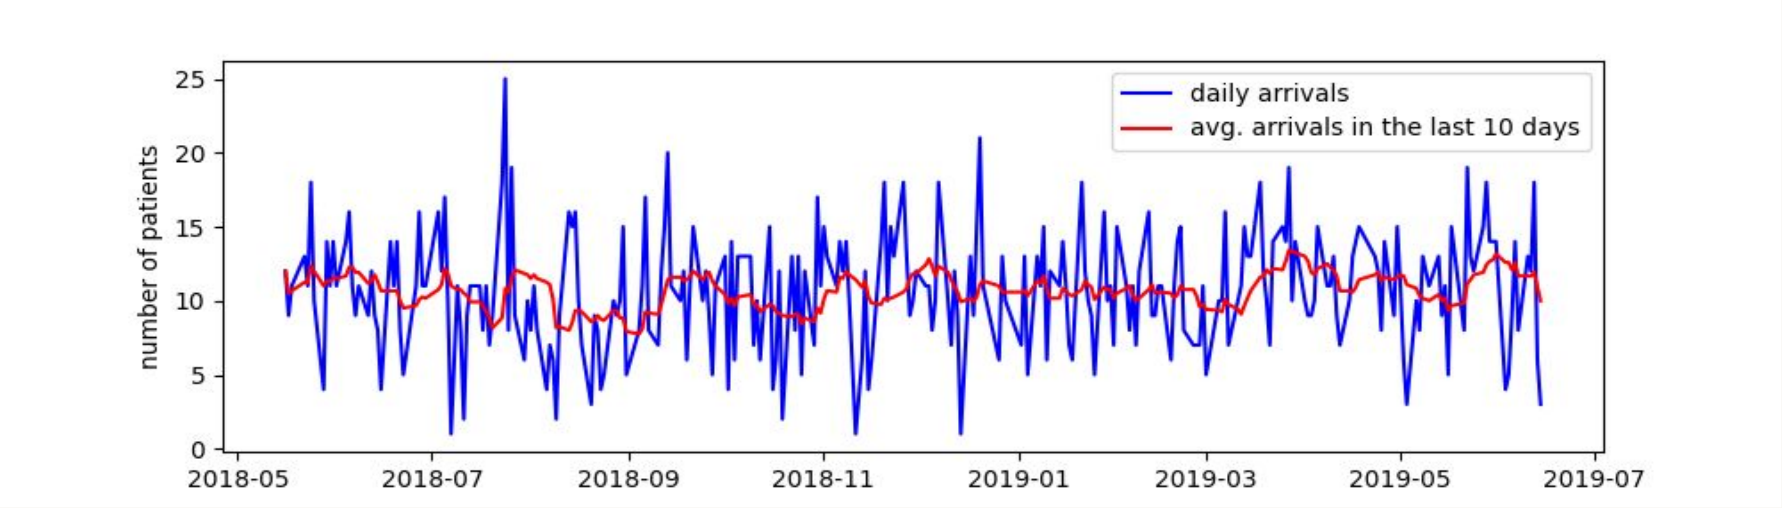

RESULTS ON THE REAL INSTANCE

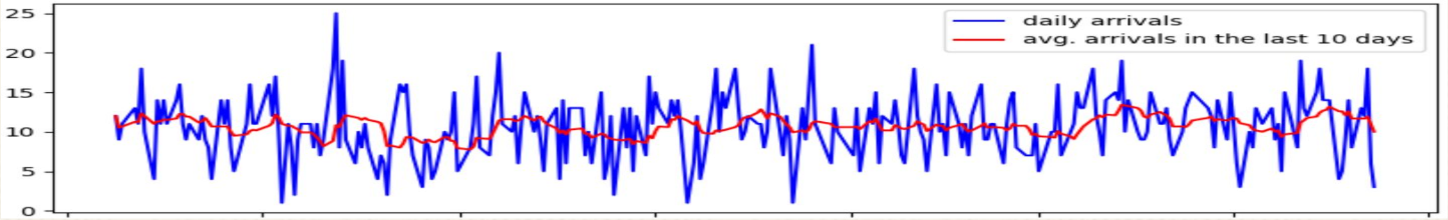

Scheduling strategy	Avg. occupancy (%)	Waiting time (days)					Overdue time (days)				
		overall	P1	P2	P3	P4	overall	P1	P2	P3	P4
online-greedy	97.45	33.02	5.14	6.13	43.67	44.02	44.02	5.14	3.91	29.74	16.18
daily-greedy	97.51	32.91	6.00	6.23	43.48	43.80	17.71	6.00	3.99	29.58	16.00
daily	97.72	33.53	9.79	9.63	42.87	**43.44**	18.25	9.79	7.15	28.93	**15.65**
weekly	97.61	33.04	7.86	7.72	**42.42**	44.10	17.76	7.86	5.37	**28.51**	16.19
prediction-based	97.14	32.93	**3.29**	**4.05**	44.21	44.94	17.69	**3.29**	**1.99**	30.22	16.96

EXPLAINABILITY WITH SHAPLEY VALUES

- SHAP (SHapley Additive exPlanations)
  - represent the relative strength of a variable on the outcome
- Widely use for explainability in machine learning
- Highly appreciated in the healthcare domain

GLOBAL INTERPRETATION
BEESWARM PLOT

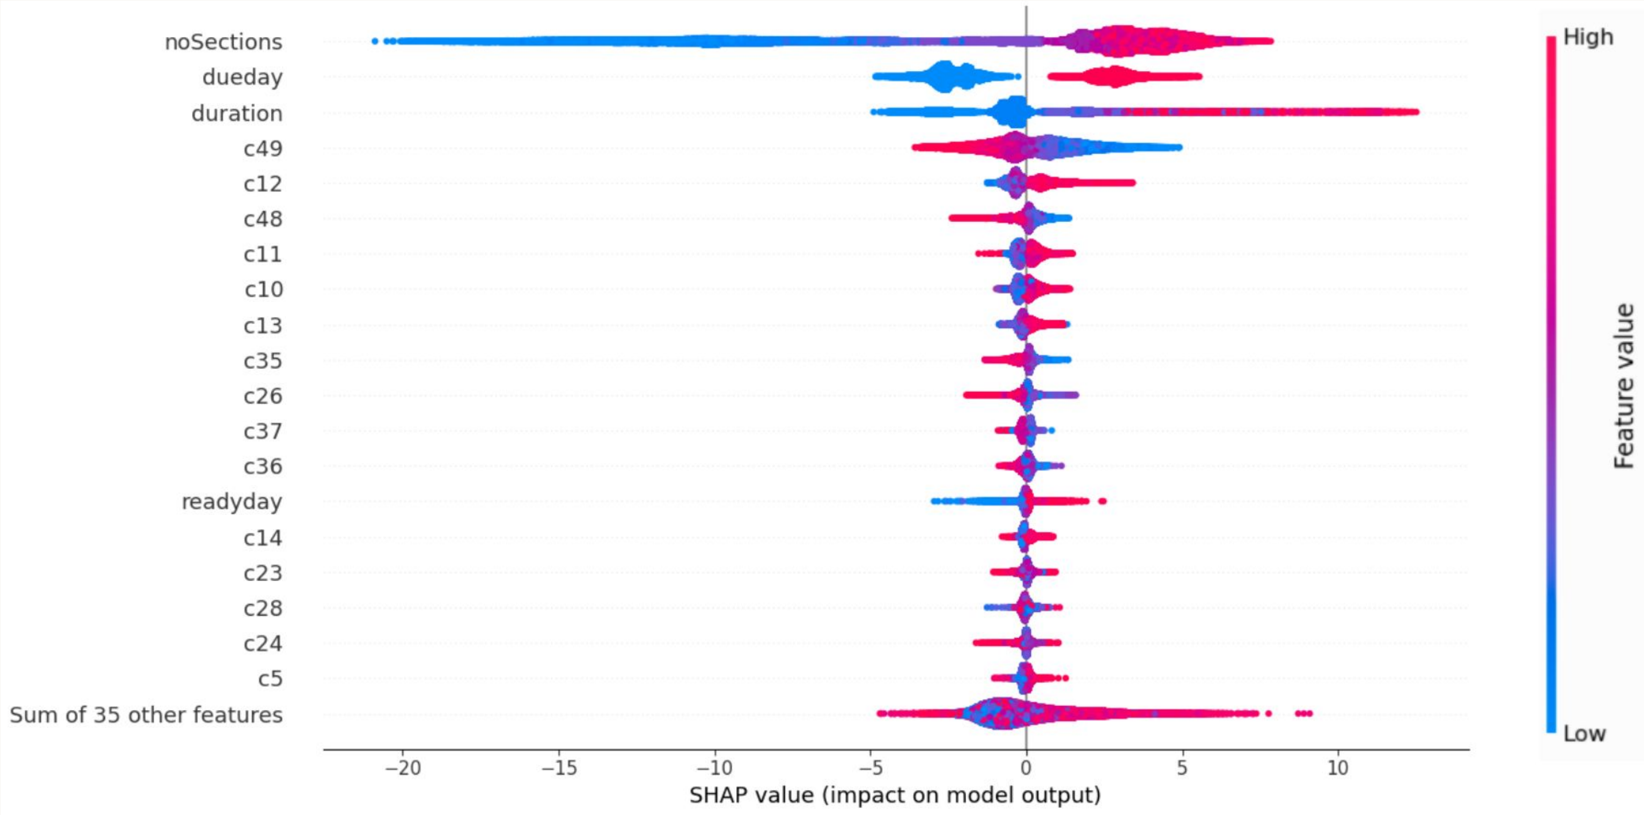

6 linacs - arrival rate 9.0

LOCAL INTERPRETATION WATERFALL PLOT

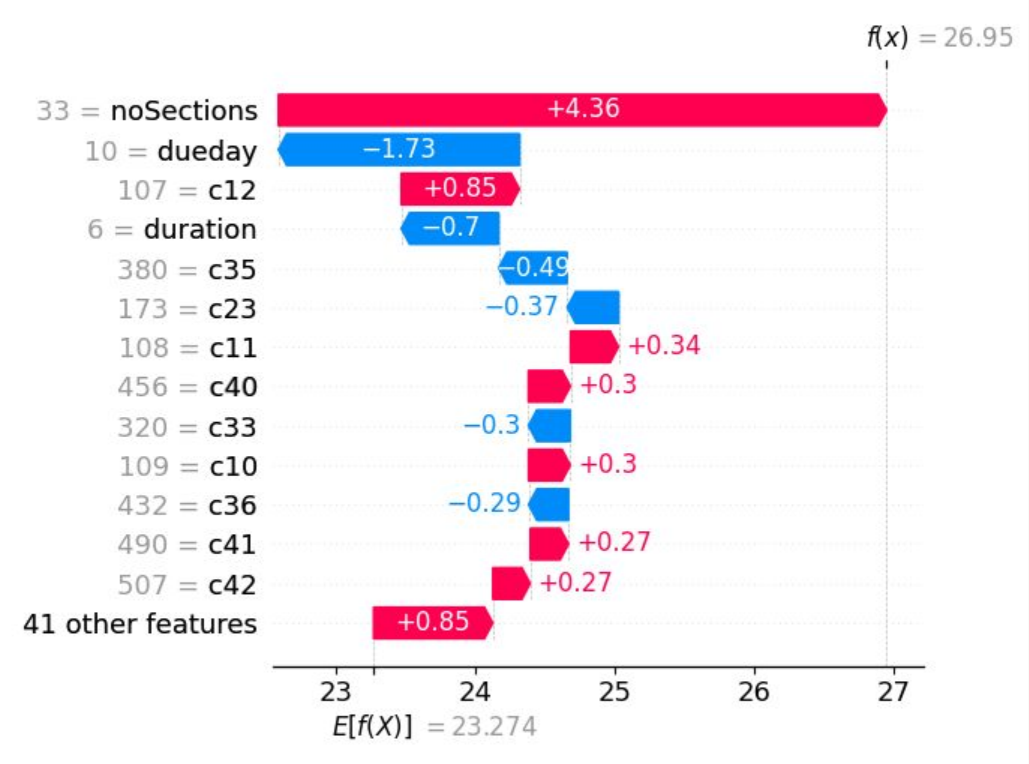

CONCLUSIONS AND CHALLENGES?

- A machine learning-based approach for online dynamic patient scheduling
- Empirical results
  - Improve overdue times of palliative patients, especially on large and crowded hospitals
  - Not too sensitive with the fluctuation on the arrival rate
- Explainability with Shapley values
- Challenges: generating offline solutions is expensive

POLYTECHNIQUE MONTRÉAL
UNIVERSITÉ D'INGÉNIERIE

CIRRELT

HANALOG

ORLAB
Better is the best

A reinforcement learning approach for the dynamic home health care scheduling and routing problem

Ta Dinh Quy, Tu-San Pham,

Minh Hoang Ha, Louis-Martin Rousseau

November 3rd 2022 - focus period Linköping

HOME HEALTHCARE (HHC)

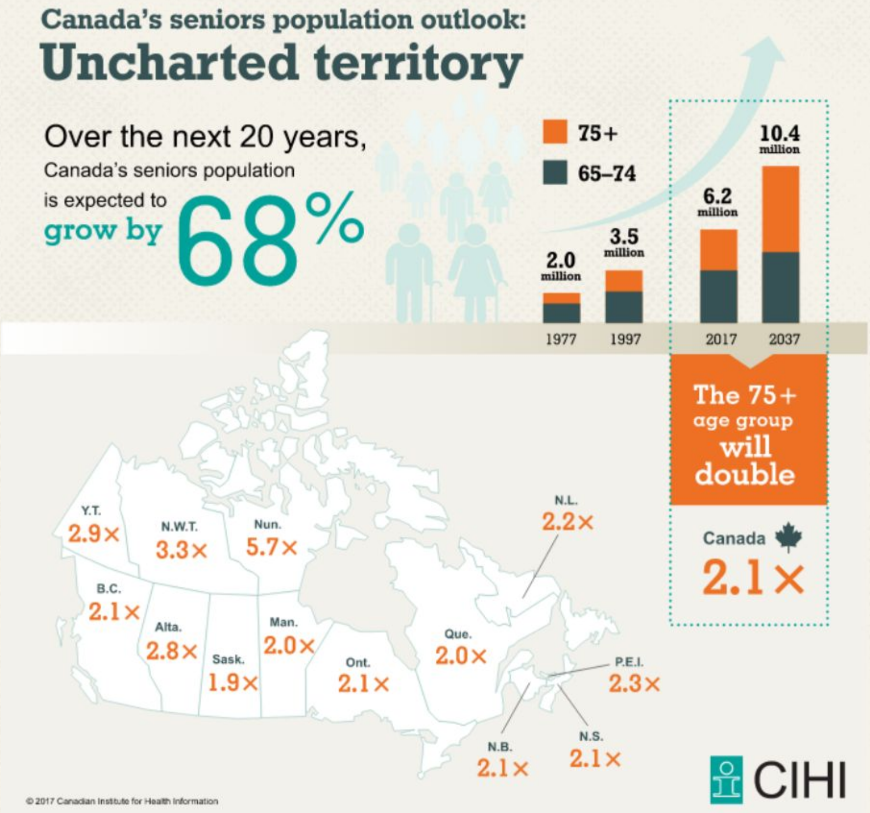

https://www.cihi.ca/en/infographic-canadas-seniors-population-outlook-uncharted-territory

- Lack of medical resources and expensive health care service costs
- People want to stay at home as long as possible
- Cost-effective and flexible
  - costs 32% less than hospital care
- HHC services is one of the fastest growing market in the US and Canada
  - In Canada, 2.2 million people relied on home care services. (2012)

Mix between an assignment problem and a VRP

Hard constraints	Soft constraints
Nurse skills	Continuity of care
Type of care	Optional requirements
Forbidden nurses	Travel time
Time windows	**Min**/Max worktime week
Available days	**Min**/Max worktime workday
Workdays	Number of visits over the week
Time-dependent travel time	

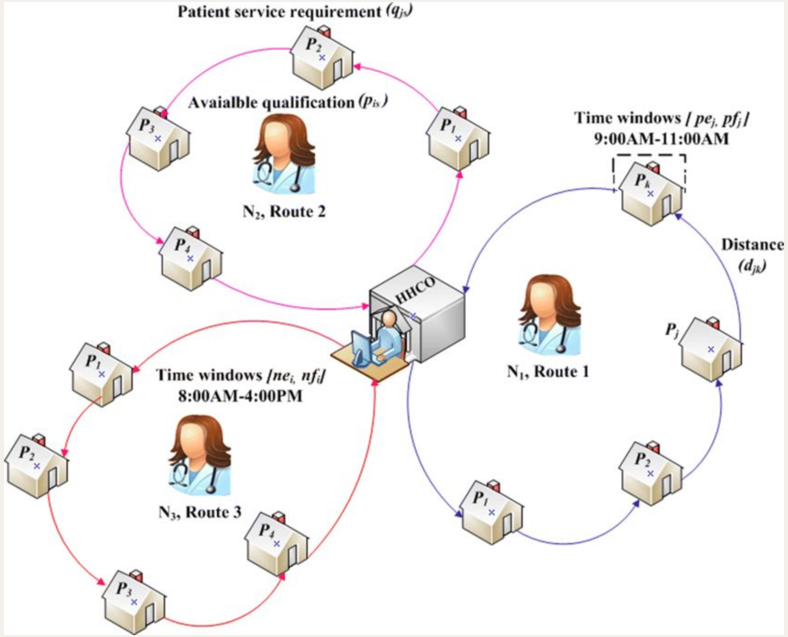

HOME HEALTHCARE (HHC)

- Delivery of professional care by workers in a client’s home.
- Objective: increase service quality and decrease costs

CHALLENGES


Travel routes
Staff availability
Continuity of care
Client needs & preferences
2.500.000 visits/year in an avg agency
Union rules
Stochastic travelling time
Stochastic request arrival
Cancellations
Sick leaves

PROBLEM DESCRIPTION

- Patient requests
  - Episode of cares (nb of weeks)
  - Frequency (nb of visits per week)
  - Visit pattern
- Nurses
  - Skill level
  - Regulation
- Constraints
  - Continuity of care
  - Time window..

	Time slot 1	Time slot 2	Time slot 3	Time slot 4	Time slot 5	Time slot 6	Time slot 7	Time slot 8	Time slot 9	Time slot 10	Time slot 11	Time slot 12	Time slot 13	Time slot 14	Time slot 15
Monday		Patient 1			→	Patient 2									
Tuesday						Patient 2				→	→	Patient 3			
Wednesday		Patient 1			→	Patient 2									
Thursday		Patient 1			→	Patient 2				→	→	Patient 3			
Friday						Patient 2									

→ *Travel time*

THE DYNAMIC HHCSP (DHHCSP)

- Patient's request arrives dynamically.
- Decision:
  - Accept or reject?
  - Online decision
- Obj: maximize the number of patients served.

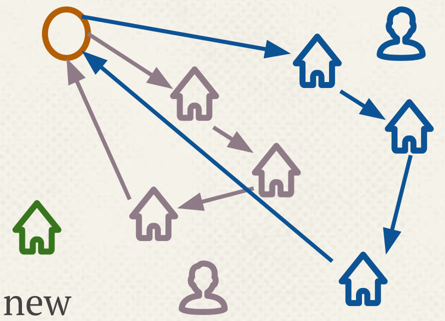

new patient

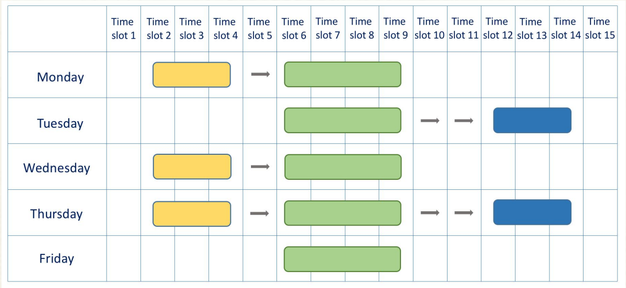

New patient

?

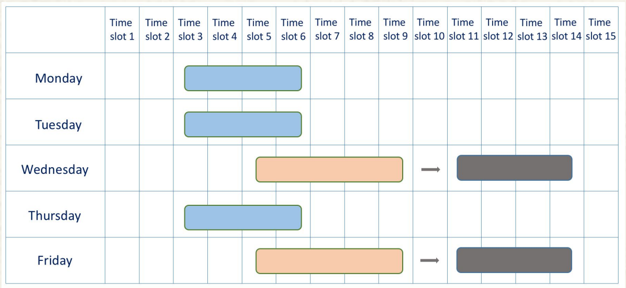

Nurse 1

Nurse 2

SOLUTION APPROACH

- *Greedy approach*
- *Scenario-based*
- *Reinforcement learning*

DISTANCE-BASED INSERTION HEURISTIC

Bennett & Erera (2011)

Greedy heuristic:

- ➔ Each nurse
- ➔ Each visit pattern
- ➔ Find the cheapest insertion
- ➔ Accept if there's an eligible slot
- ➔ Reject otherwise

Nurse

N1	-> 600
N2	-> 400
N3	-> 300

Visit pattern

(M, W)	-> 100
(M, Th)	-> 60
(Tu, Th)	-> 70

Route

Cost: total travelling time

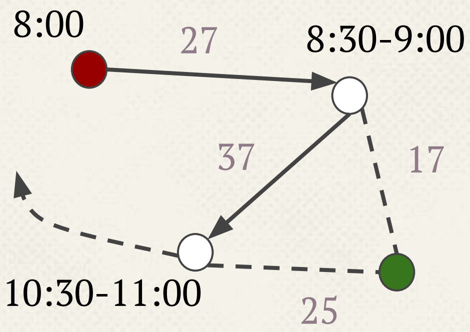

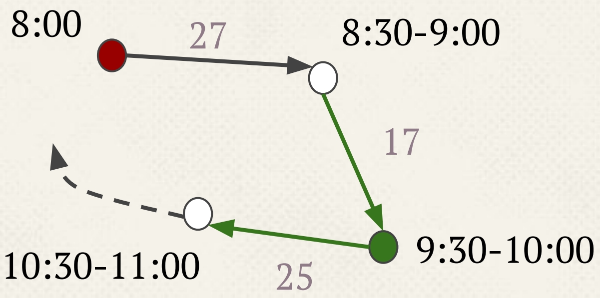

SCENARIO BASED APPROACH (SBA)

Demirbilek et al. (2019a)

1. Generate a set of future scenarios
2. For each scenario
3. Run greedy heuristic
4. If the patient is not accepted in any scenario
5. Reject
6. Else
7. Schedule the patient to the most frequently assigned

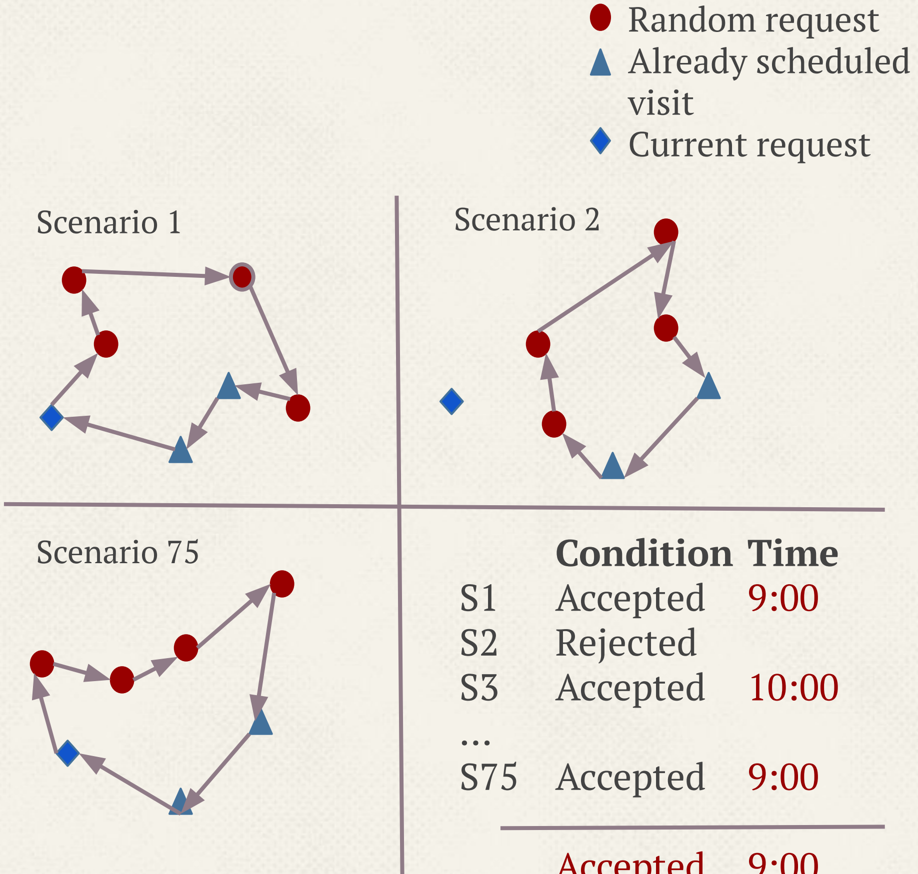

	Condition	Time
S1	Accepted	9:00
S2	Rejected	
S3	Accepted	10:00
…		
S75	Accepted	9:00
	Accepted	9:00

REINFORCEMENT LEARNING

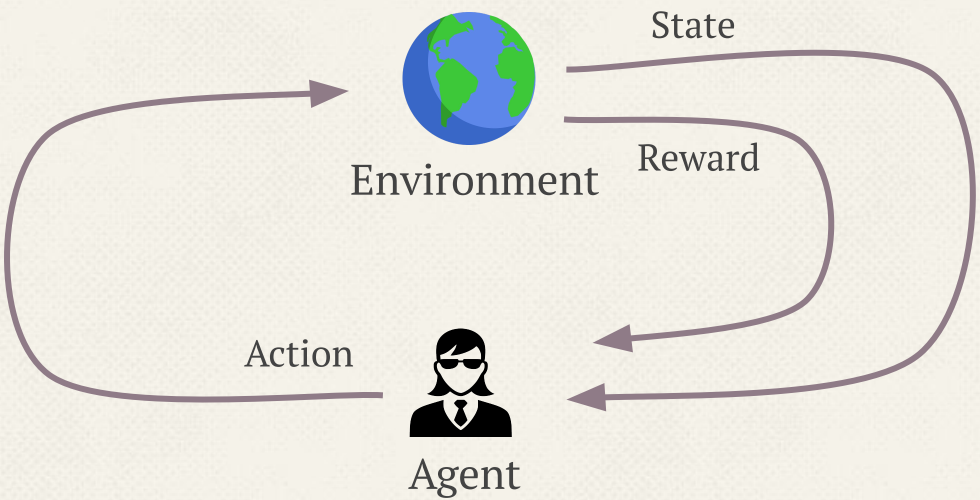

Interacts in a
Markov Decision Process

REINFORCEMENT LEARNING FOR DYNAMIC REQUEST ACCEPTANCE

Decision: when a new request arrives

STATE	ACTION	REWARD	POST-DECISION STATE
Request's information and current schedules	Accept, schedule patient using a heuristic	1	New schedule
	Reject	0	Schedule doesn't change

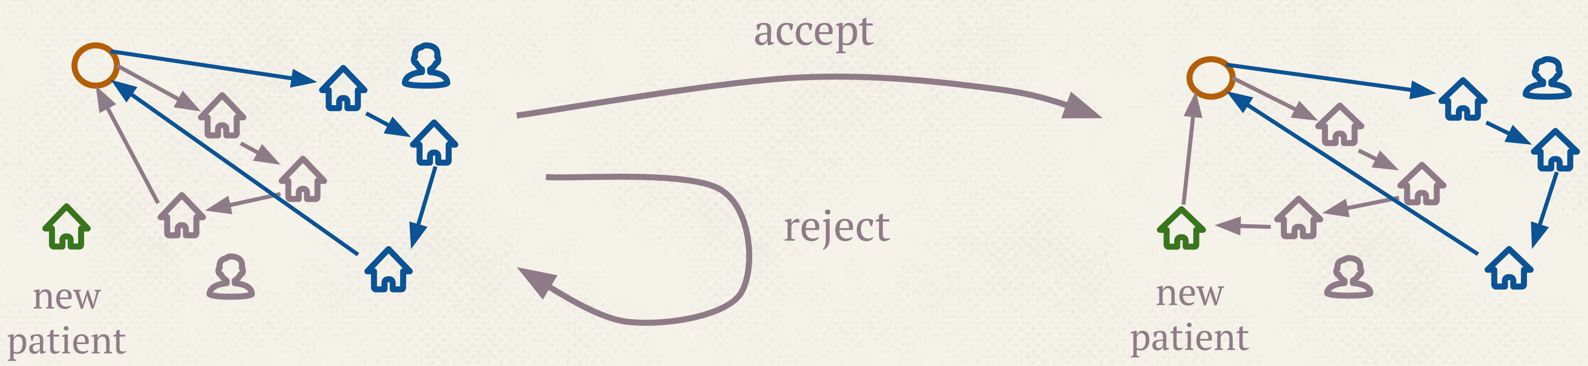

STATE PRESENTATION

Patient's characteristic	Nurse j's information	…	Cheapest insertion cost

- *Episode of care*
- *Frequency*
- *Duration of visits*

- *Number of assigned visits*
- *Total idle time*
- *Total traveling time*

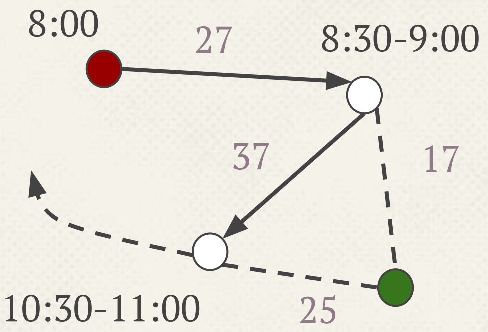

DOUBLE DEEP Q-LEARNING WITH EXPERIENCE REPLAY

Initialize replay memory $\mathcal{D}$
Initialize network θ, target network $\hat{\theta}$
for *episode* $= 1, M$ **do**
    $s_t \leftarrow s_0$
    for $t = 0, T$ **do**
        With probability ϵ select a random action a_t
        Execute a_t, observe r_t, s_{t+1}
        Store (s_t, a_t, r_t, s_{t+1}) in $\mathcal{D}$
        $s_t \leftarrow s_{t+1}$
        Sample random *minibatch* of $\mathcal{N}$ transitions from $\mathcal{D}$
        for every transition (s_j, a_j, r_j, s_{j+1}) in *minibatch* **do**
            $y_j \leftarrow r_j + \gamma Q_\theta(s_{j+1}, argmax_{a'} Q_{\hat{\theta}}(s_{j+1}, a'))$
        end for
        Calculate the loss $\mathcal{L} = 1/N \sum_{j=0}^{N-1} (y_j - Q(s_j, a_j; \theta))^2$
        Update θ by gradient descent by minimizing the loss $\mathcal{L}$
        Update target network $\hat{\theta} \leftarrow \theta$ every K steps
    end for
end for

Calculate target

Back-tracking

Update target network every K steps

Experimental results

EXPERIMENTAL DESIGN

- Nurses,
  - Working time 8:00 to 16:30
  - 3 skill levels
- Patients
  - Location: uniform distribution in a square of 80*80
  - Travelling time: 1’ per unit
  - Time between requests: exponential distribution with expected value
    - 150 (~10 requests/day)
    - 240 (~6 requests/day)
    - 360 (4 requests/day)
- Baselines
  - Greedy
  - SBA

1 nurse

Avg. decision time (s)	
SBA	10.17
RL	0.033

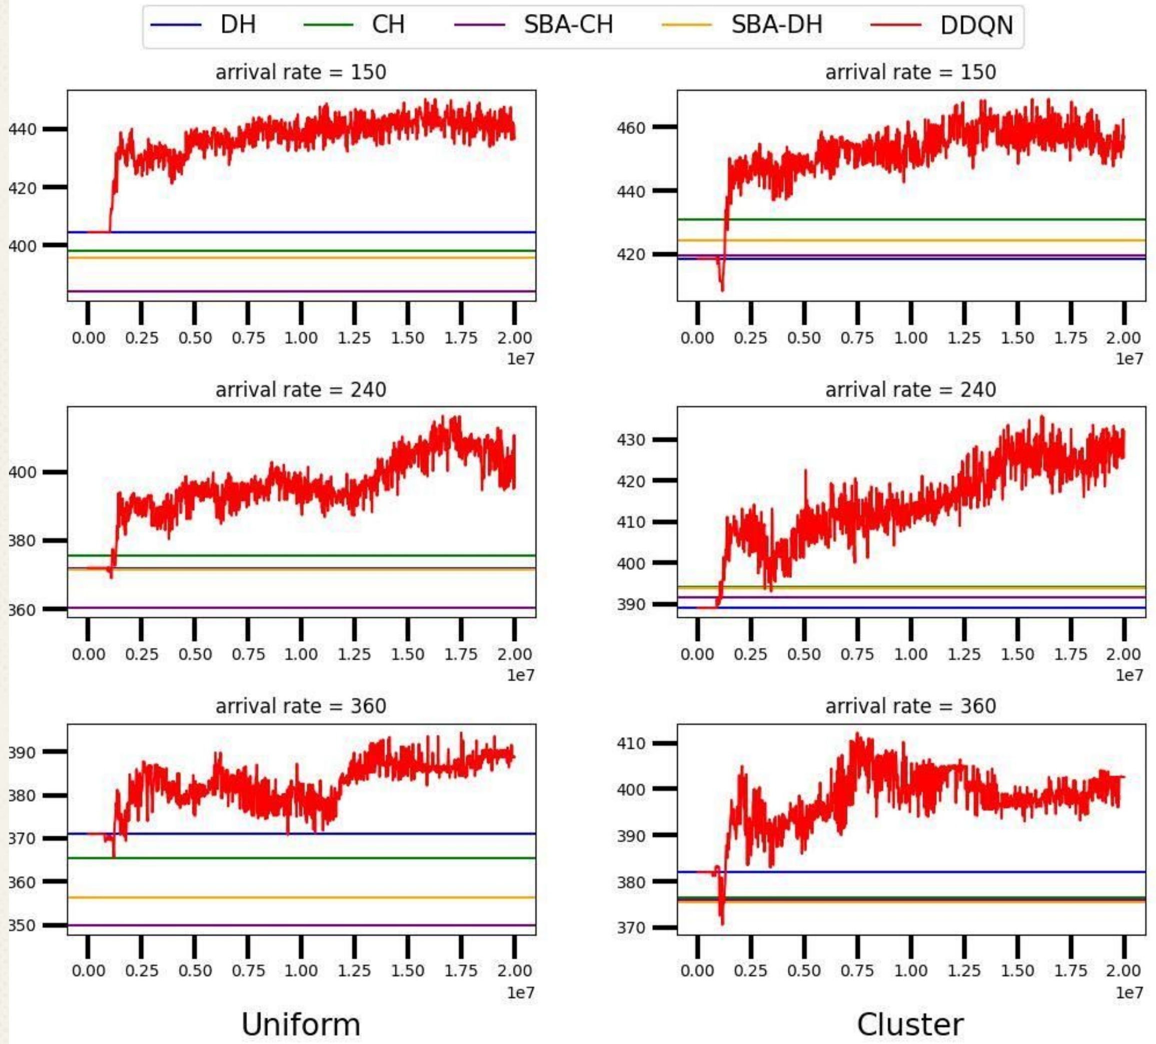

CONCLUSION AND FUTURE WORK

- RL approach for a dynamic HHC scheduling problem
- Future work
  - Multiple objectives (soft constraints, travelling time…)
  - More sources of uncertainty

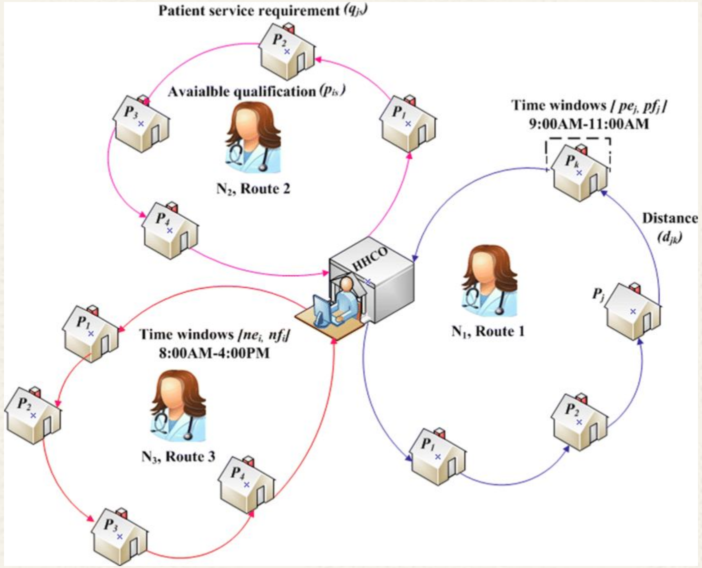

THANKS!

Any questions?

tu-san.pham@polymtl.ca