Patient panel prioritization with finite resources

Daniel F. Otero-Leon

The physician has a group of patients (panel) which consist of different backgrounds.

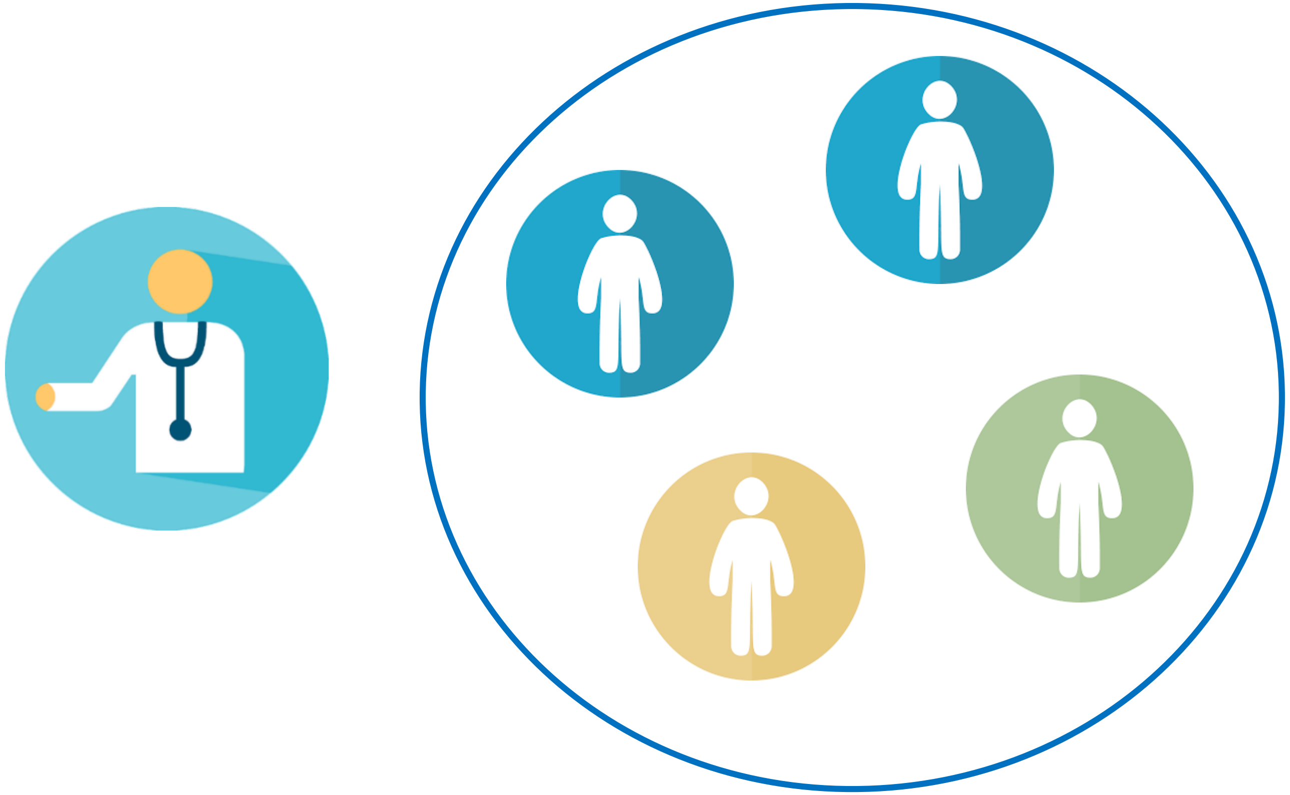

Physicians have finite resources to schedule patients, such as time and available budget.

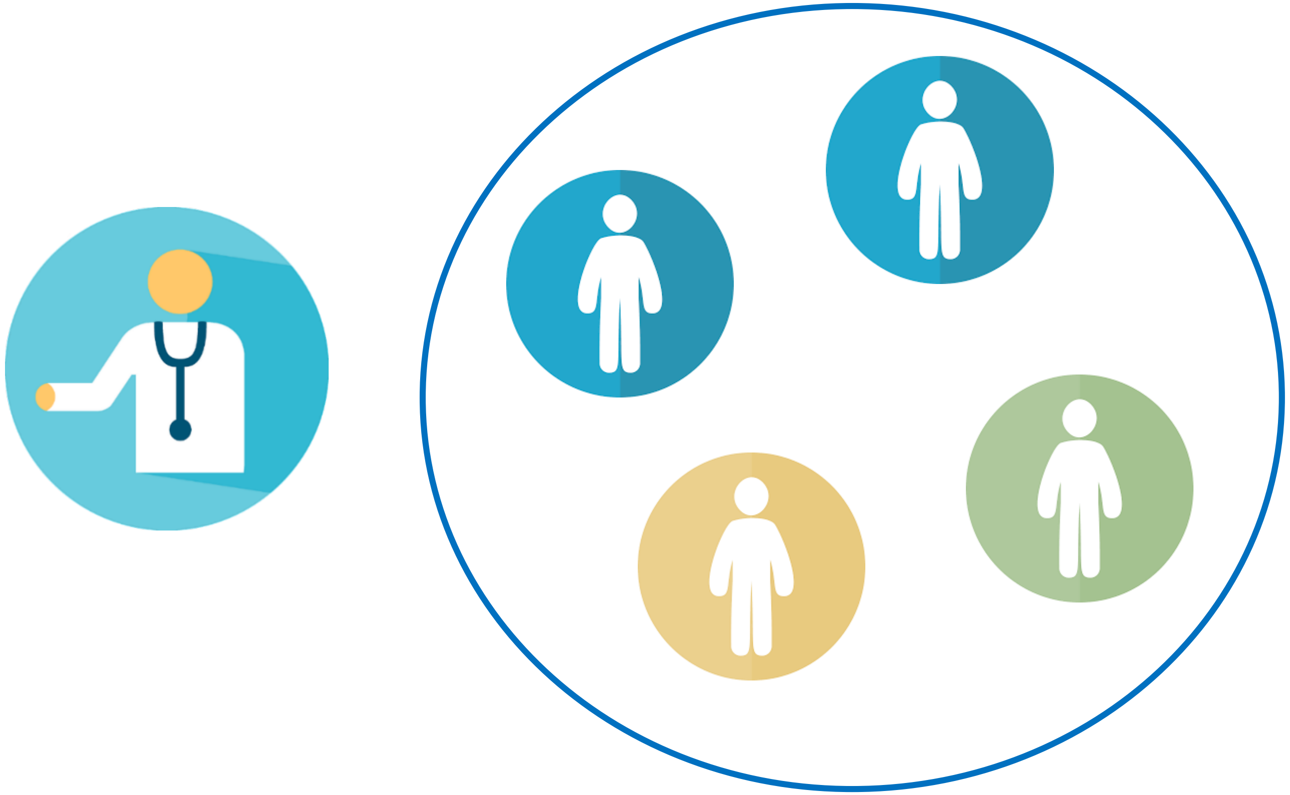

Who does the physician choose for each epoch?

Patients available information is divided into two types of variables.

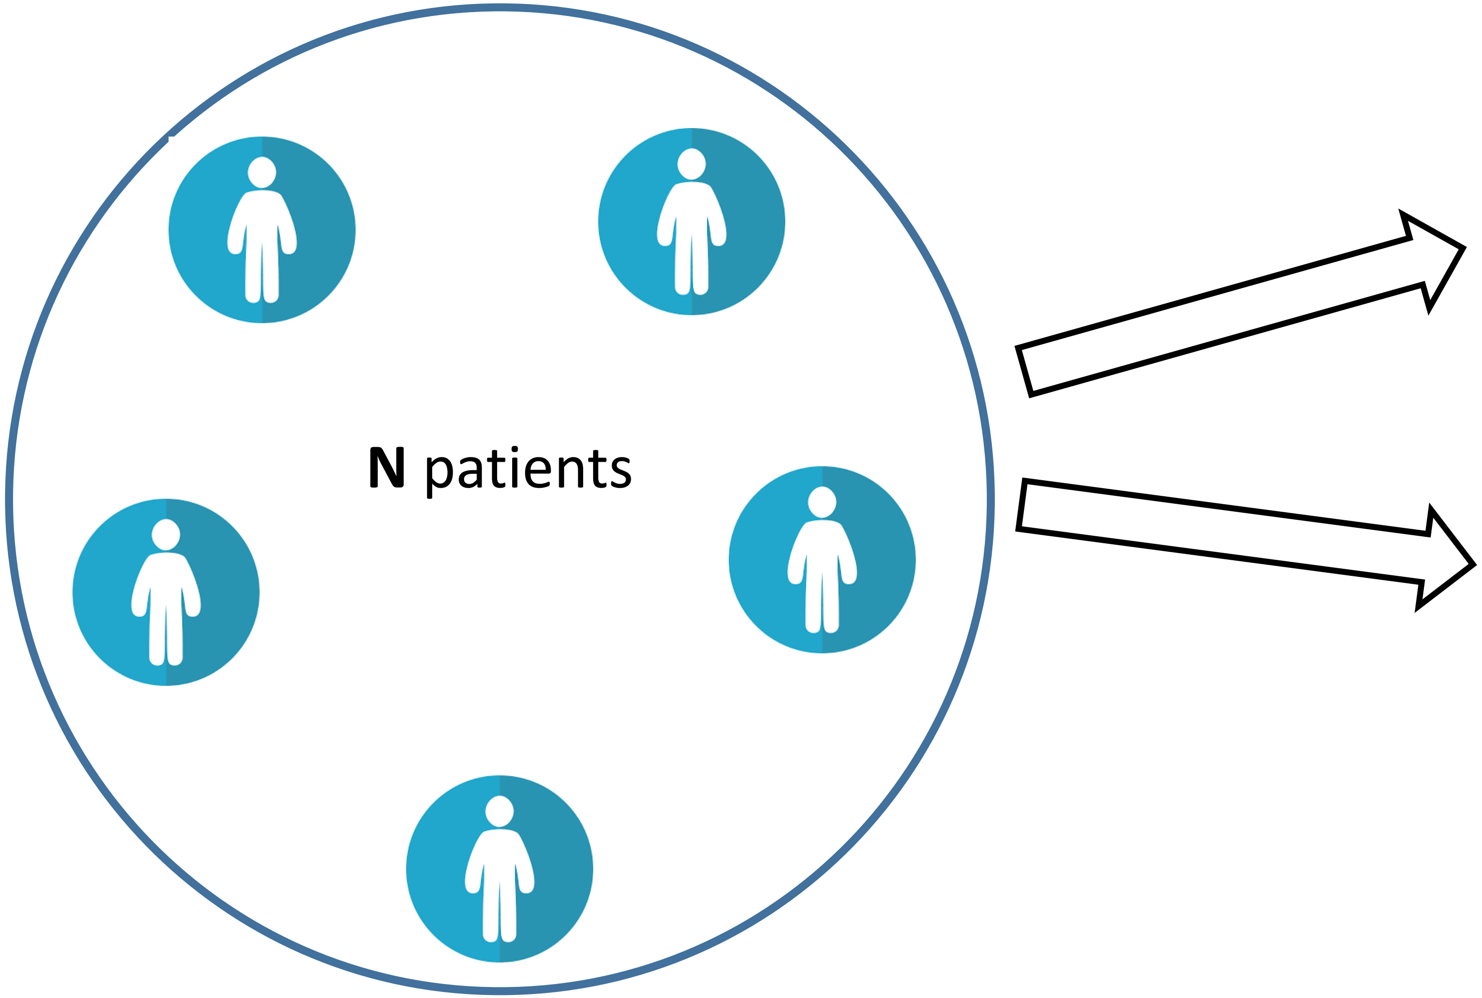

For each patient j:

Static variables X_j:

- Demographics

Time varying variables Z_{jt}:

- Health risk factors
- Age
- Treatment
- Adherence

How does each epoch works?

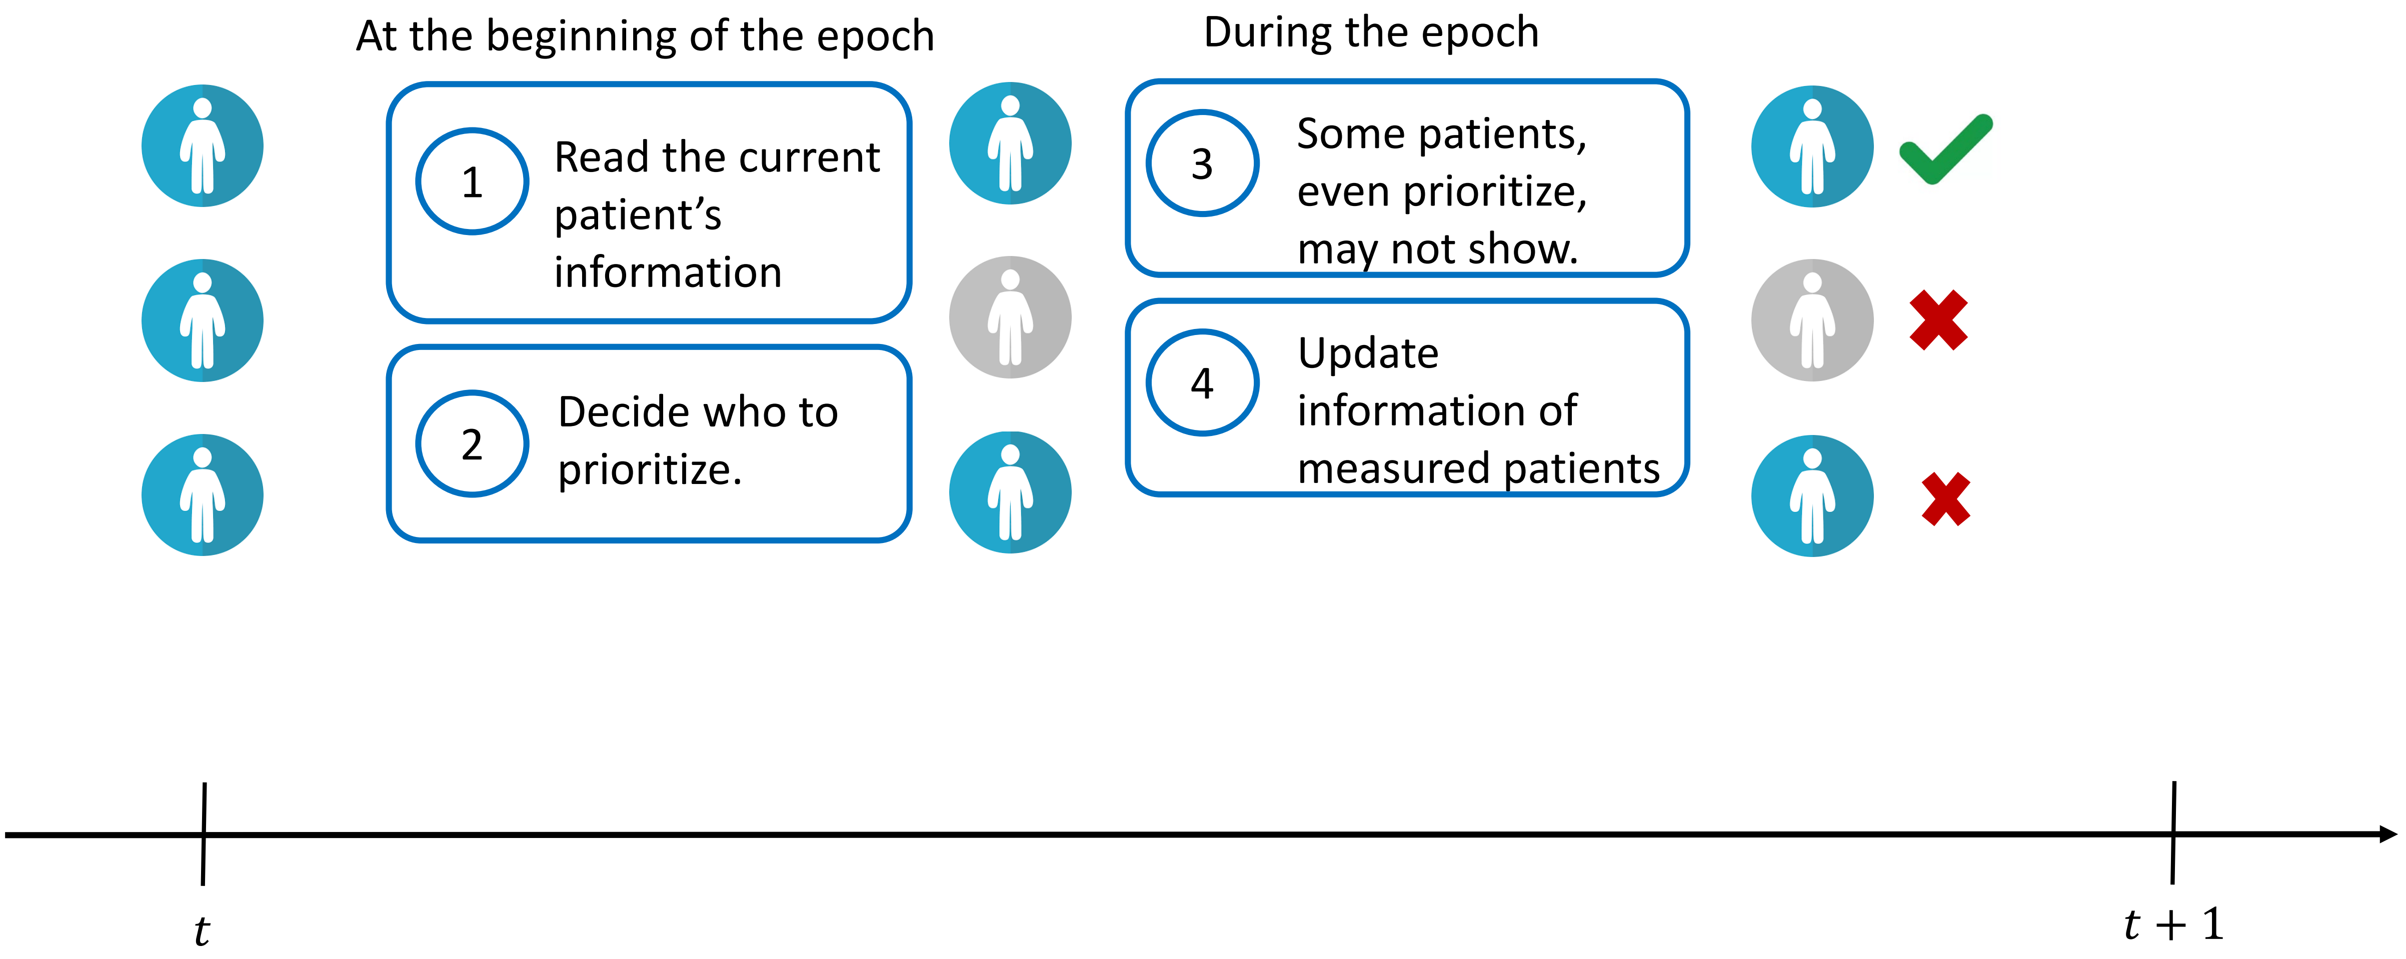

The decision to whom to prioritize is taken at the beginning of the epoch.

At the beginning of the epoch

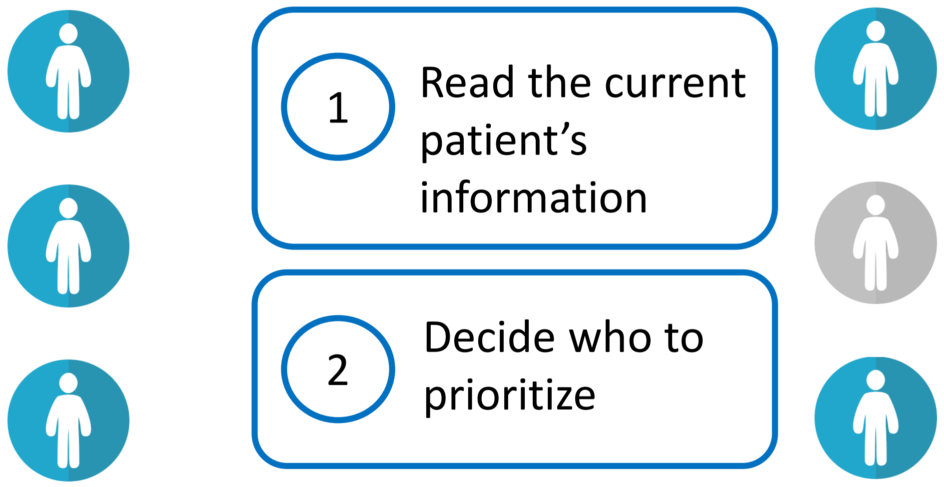

$d_{jt} \in \{0,1\} \Longrightarrow d_t := (d_{1t}, \dots, d_{Nt})$

$g_j(X_j, Z_{jt})$: cost of prioritizing patient j in epoch t

$$\sum_j d_{jt} g_j(X_j, Z_{jt}) \leq \underline{M}$$

Resource available

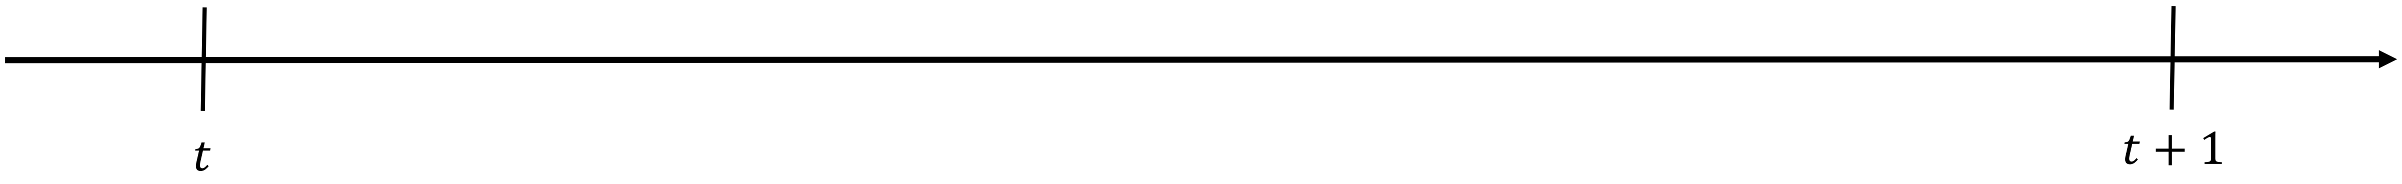

The transition probabilities are estimated assuming independence between patients.

Per patient j:

Probability that patient j is measured in epoch t

Risk factors stochastic behaviors

During the epoch

3 Some patients, even prioritize, may not show.

4 Update information of measured patients

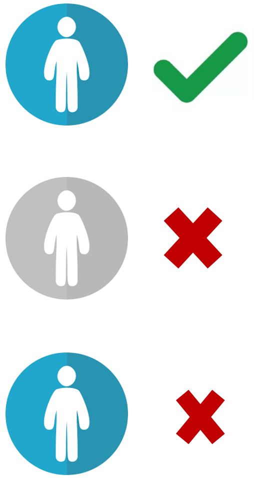

For the panel:

$p_t(\boldsymbol{X}, \boldsymbol{Z}_{t+1} \,|\boldsymbol{X}, \boldsymbol{Z}_t, d_t) \Rightarrow$ Assuming independence between patients

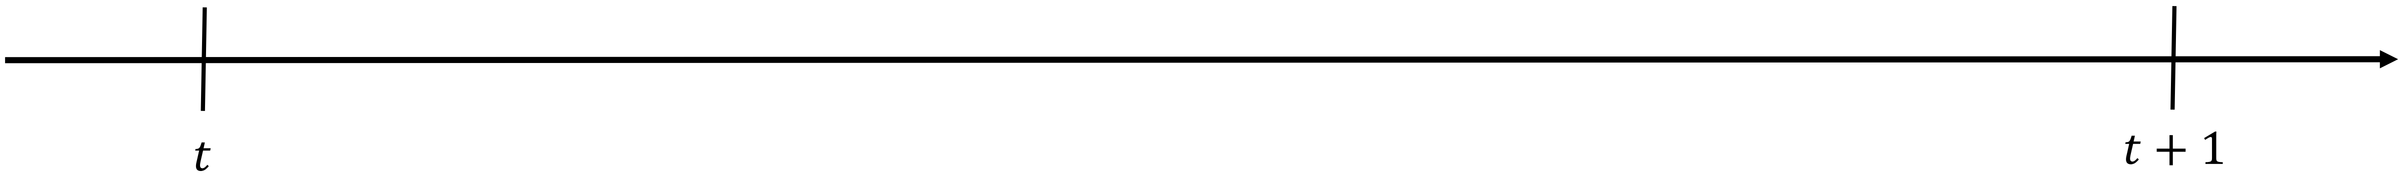

The objective is to maximize the panel's overall well-being.

Possible objectives:

- Maximize QALYs
- Minimize disease events (Cancer, Heart Attacks, etc.)

Define $r_{jt}\left(X_j, Z_{jt+1}, d_{jt}\right)$ as the rewards for patient j in epoch t.

$$r_t(X, Z_t, d_t) = \sum_j \sum_{z_j \in Z_{jt+1}} r_{jt}\left(X_j, z_j, d_{jt}\right) p_{Z_{jt}}\left(z_j \middle| X_j, Z_{jt}, d_{jt}\right)$$

Then this problem can be modeled as a Finite state and finite time MDP

$$V_t(X, Z_t) = \max_{\substack{d_t \\ s.t. \sum_j d_{jt} c_j(X_j, Z_{jt}) \leq M}} \left\{ r_t(X, Z_t, d_t) + \lambda \sum_{(z,l) \in (Z_{t+1}, L_{t+1})} p_t(X, z \mid X, Z_t, d_t) V_{t+1}(X, z) \right\}$$

Where, $V_T(X, Z_T) = 0$.

The model as a Multi-Armed Bandit

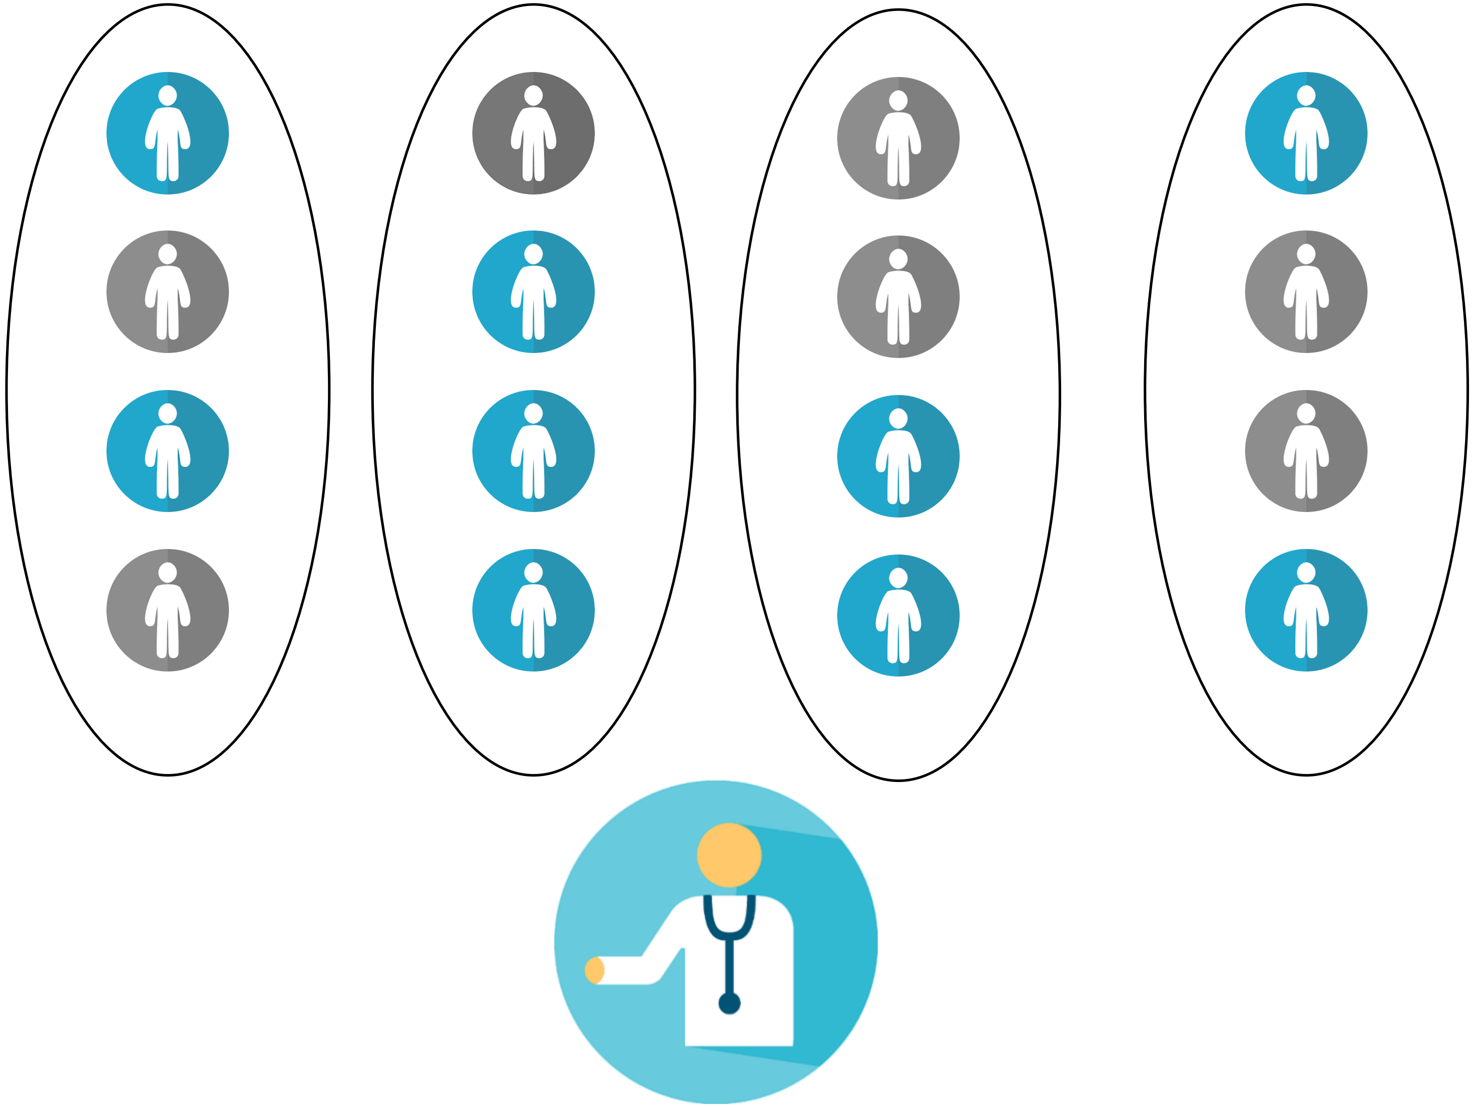

Can choose only one arm

Known as a Combinatorial MAB

This problem compared to literature.

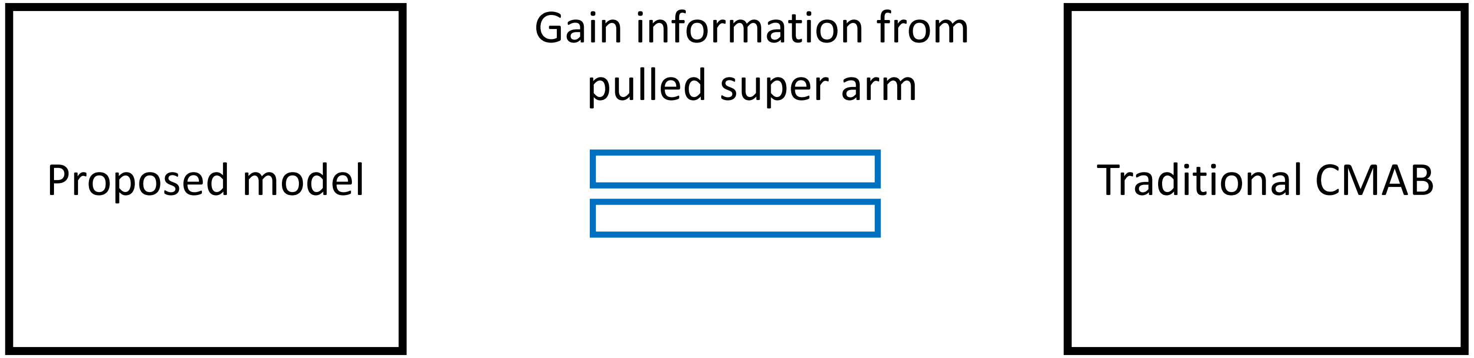

Our problem differs because information is not necessarily gain when arm is pulled.

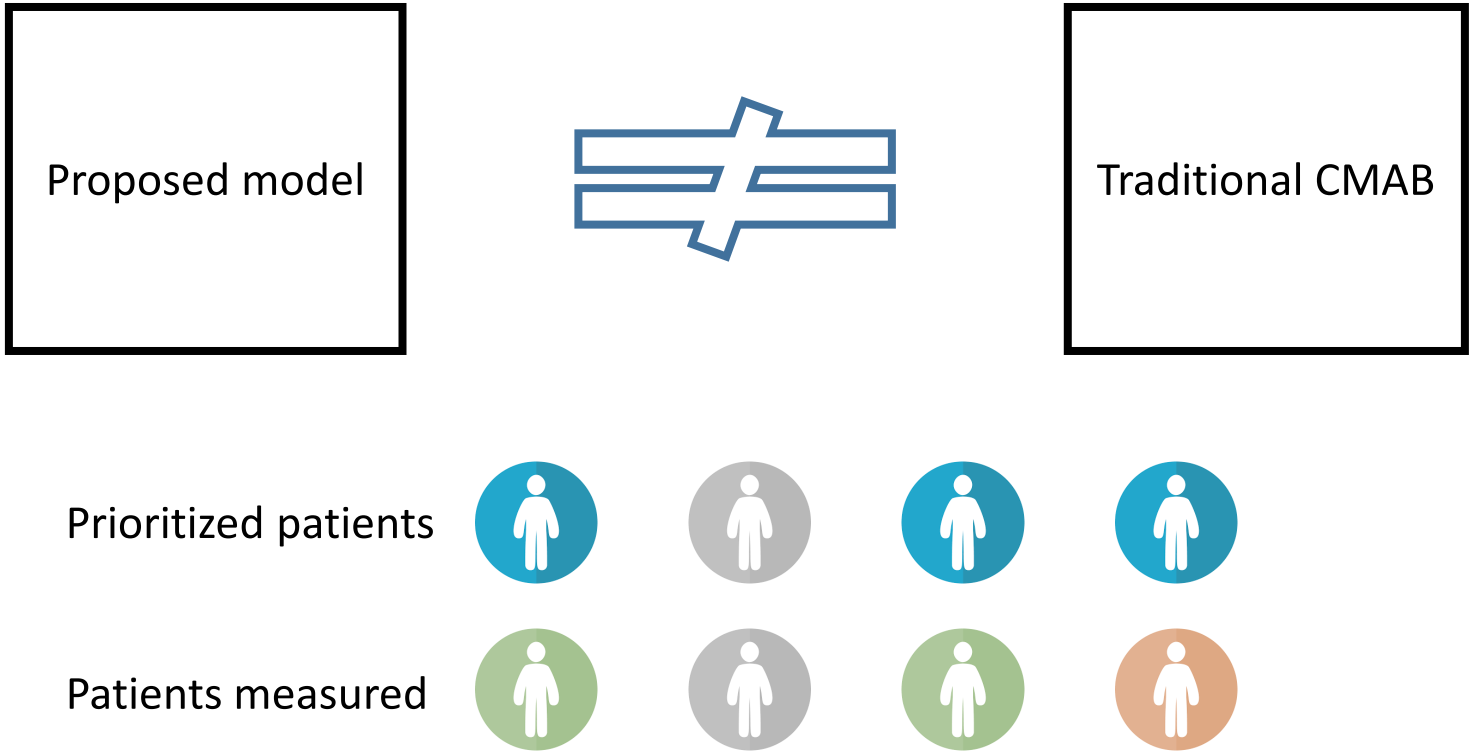

How does the literature solves this problem?

How does the literature solves this problem?

Traditional CMAB

- Myopic algorithm: Maximizes one period rewards
- Upper confidence bounds: Identify upper bounds and choose m top UCB.
- Whittle index: Identify index and choose top m.

Prioritizing Hepatitis C treatment in U.S. Prisons

Turgay Ayer, Can Zhang, Anthony Bonifonte, Anne C. Spaulding, Jagpreet Chhatwal (2019). Operations Research 67(3).

1. When a patient is chosen, this patient is measured.
2. Restless MAB: Patient not chosen changes between epochs.
3. Model satisfies the indexability property

Prioritizing Hepatitis C treatment in U.S. Prisons

Turgay Ayer, Can Zhang, Anthony Bonifonte, Anne C. Spaulding, Jagpreet Chhatwal (2019). Operations Research 67(3).

1. When a patient is choose, this patient is measure.
2. Restless MAB: Patient not chosen changes between epochs.
3. Model satisfies the indexability property

→ An index can be assign to each patient.

Background

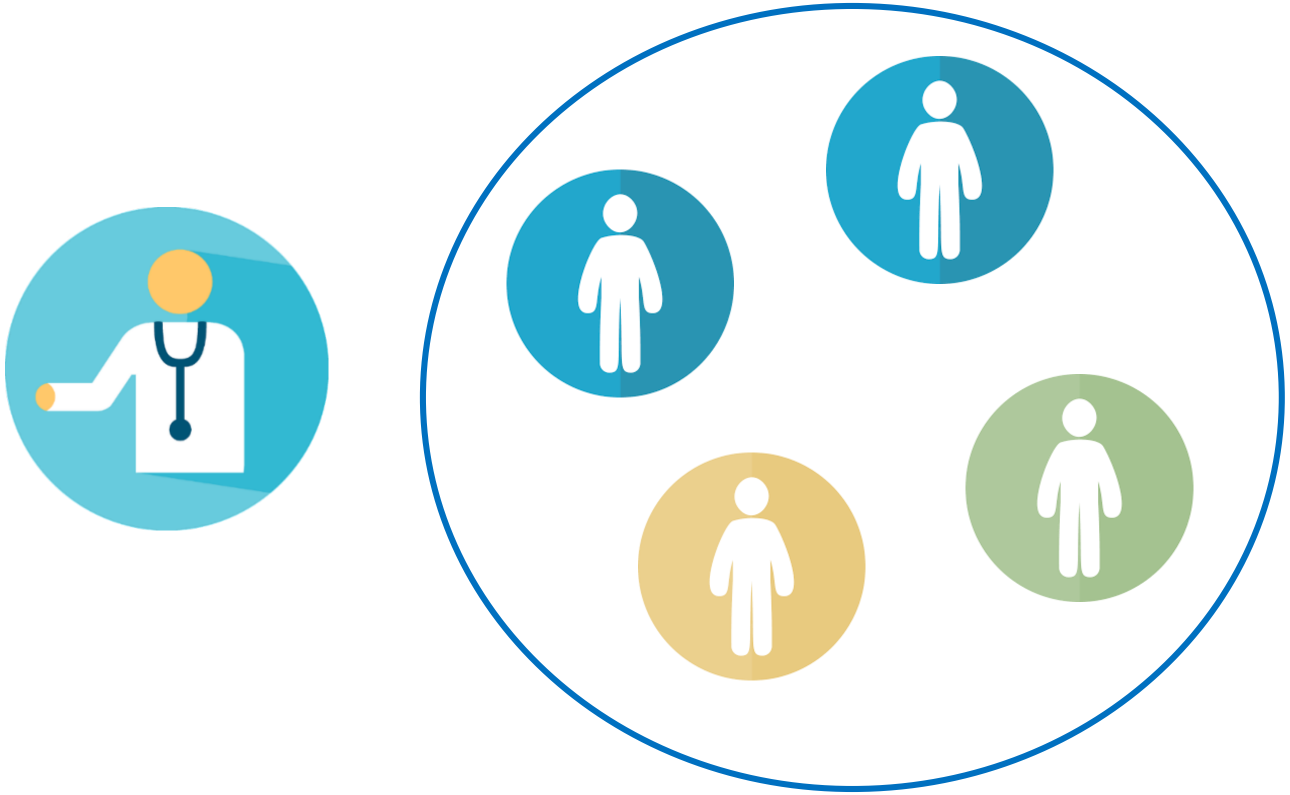

1. Treating for HCV in prisons is expensive.

2. Treating everybody is twice as high as the entire healthcare budget.

3. Patients remain in prison for 15 years on average.

The states of the patient depends if the patient is treated or not.

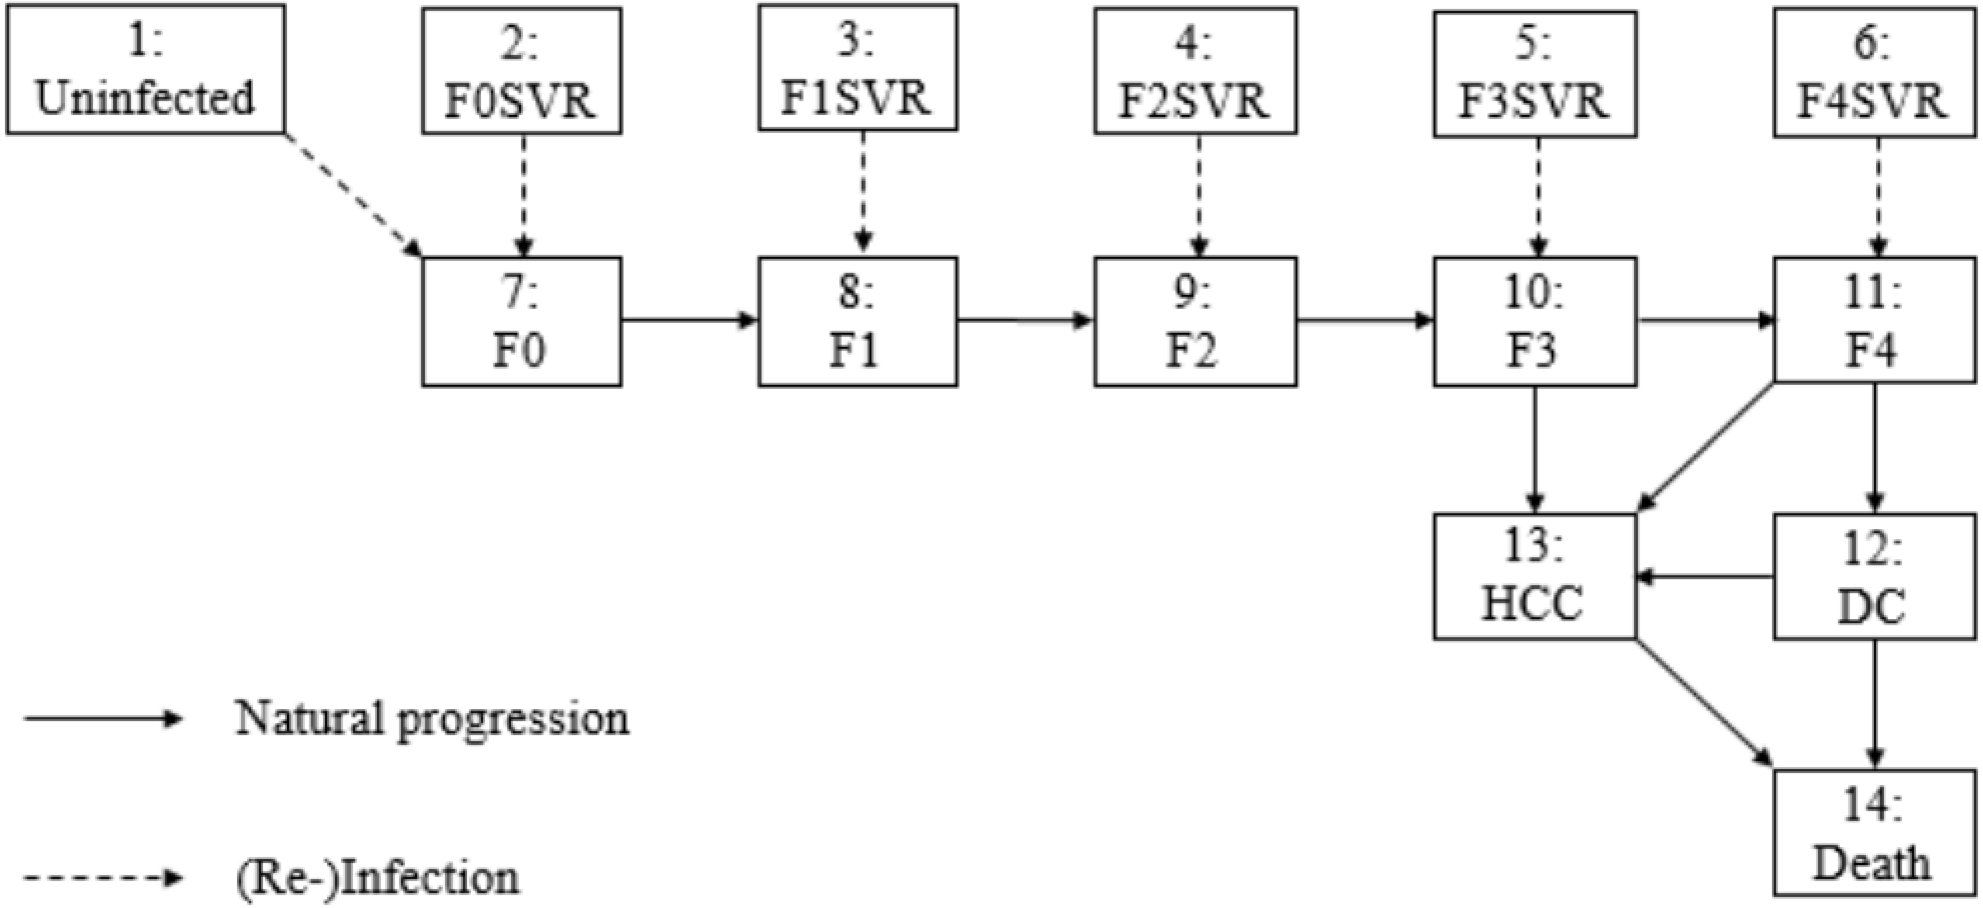

The indexability is difficult to check in general.

Notation:

- r_i^d: Reward for state i when decision d is taken.
- $d \in \{0,1\}$ where 0 is not prioritize.

Analysis assuming just one arm

The indexability is difficult to check in general.

Notation:

- r_i^d: Reward for state i when decision d is taken.
- $d \in \{0,1\}$ where 0 is not prioritize.
- W: subsidy level
- $V_T(i, W) = \max\{r_i^0 + W, r_i^1\}$

Analysis assuming just one arm

A bandit is indexable if:

Definition 1: Given subsidy level W, let $\Pi(W)$ be the set of states for which the non-prioritization action is optimal. A bandit is *indexable* if $\Pi(W)$ is increasing in W, i.e. $W_1 \leq W_2 \rightarrow \Pi(W_1) \subseteq \Pi(W_2)$.

A bandit is indexable if:

The set of states at which the non-prioritization action is optimal is increasing with the subsidy level.

Then the Whittle's Index is describe as:

Definition 2: For an indexable bandit, its Whittle's index for each state i on epoch t is defined as $W_{i,t} = \inf\{W : (i, t) \in \Pi(W)\}$

Then the Whittle’s Index is describe as:

The Whittle’s index of a given state is the smallest subsidy level at which the non-prioritization action becomes optimal.

What about limit resources?

Capacity-Adjusted Closed-Form Index Policy:

- M: Maximum number of patients per epoch
- N': Number of eligible patients that can be treated per epoch.

$$\alpha = \min\left\{\frac{M}{N'}, 1\right\}$$

Probability patient is prioritized

Results

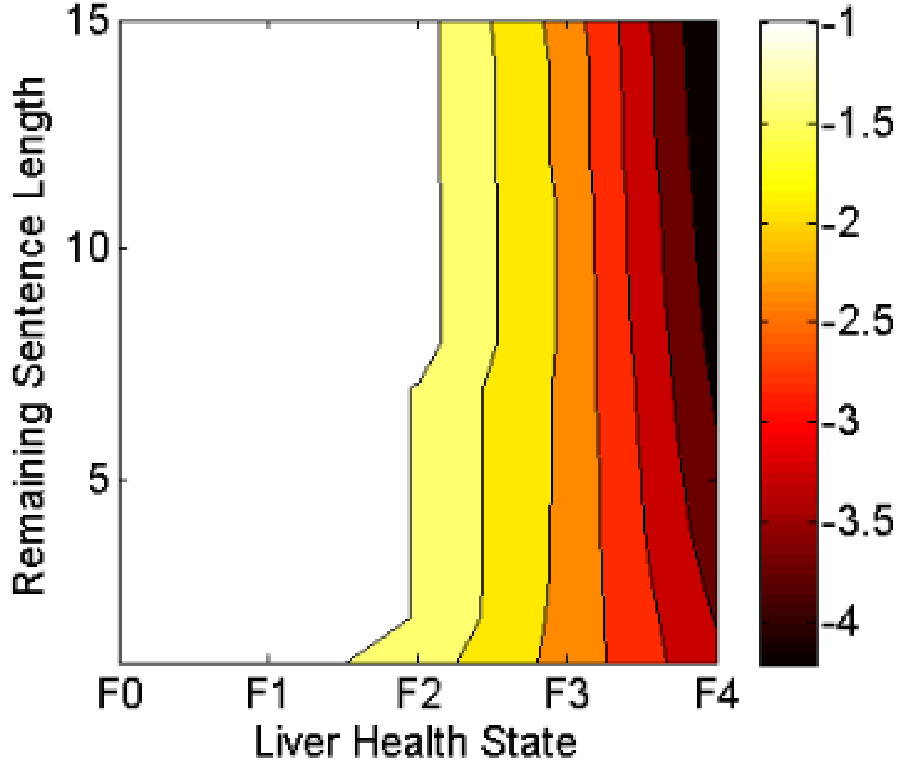

Darker the color, the patient should be treated first.

1. Prioritize patients with advance fibrosis
2. Prioritize patients with longer sentences (F4)

If there is a capacity constrained, the result is adjusted by α.

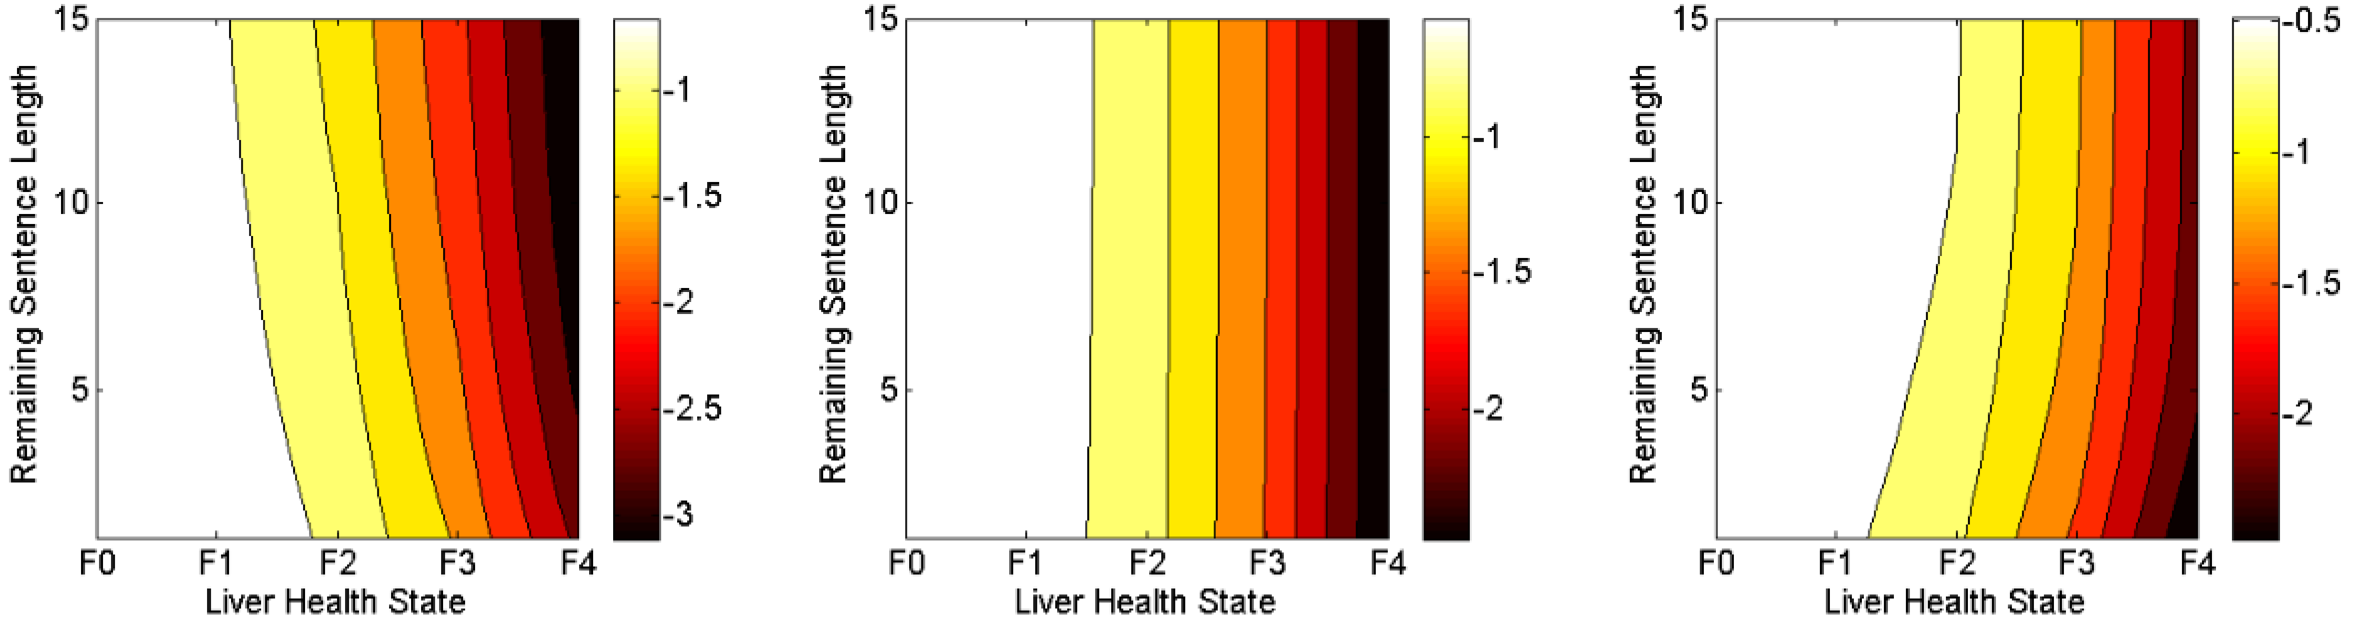

(a) $\alpha = 0.05$ (b) $\alpha = 0.10$ (c) $\alpha = 0.15$

Comparison with other algorithms

1. Health state policy: Prioritize patients with more advances fibrosis stages and randomized within the same stage.

2. Myopic: For each period, maximizes without considering future periods.

3. Whittle's Index without the capacity adjustment.

4. Primal-dual Index: Presented in Bertsimas and Niño-Mora (2000)

5. Capacity-adjusted index

Comparison with other algorithms

Policies	$M = 1$	$M = 5$	$M = 10$	$M = 15$	$M = 20$
Health state policy	37.6 (±0.5, –)	167.0 (±1.2, –)	282.1 (±1.6, –)	359.3 (±1.9, –)	415.0 (±2.1, –)
Myopic policy	42.2 (±0.6, 12.1%)	172.4 (±1.3, 3.3%)	280.3 (±1.6, -0.6%)	355.6 (±1.8, -1.0%)	408.0 (±2.0, -1.7%)
Whittle's index	43.8 (±0.6, 16.5%)	176.3 (±1.3, 5.6%)	285.3 (±1.6, 1.1%)	361.3 (±1.9, 0.5%)	412.7 (±2.0, -0.6%)
Primal-dual index	43.7 (±0.6, 16.2%)	175.3 (±1.3, 5.0%)	282.6 (±1.7, 0.2%)	358.4 (±1.9, -0.3%)	416.9 (±1.9, 0.5%)
Capacity-adjusted index	44.0 (±0.6, 17.1%)	177.7 (±1.3, 6.4%)	290.4 (±1.6, 2.9%)	368.3 (±1.9, 2.5%)	423.2 (±2.1, 2.0%)

Conclusions and Remarks

- Healthcare providers prefer policies such as the Whittle’s index where a value is assign to patients.

- Fairness issues: For example prioritizing younger patients.

Thank you

dfotero@umich.edu