Aalto University
School of Science

Application: Partially observable markov decision process in health care

Emil af Björkesten
Presentation *20*
20.11.2020

MS-E2191 Graduate Seminar on Operations Research
Fall 2020

Ischemic heart disease

- Coronary arteries are blocked, which prevents oxygen from reaching parts of the heart muscle
- Diagnosis and prediction crucial
- Treatment decisions hard
- Consequences: pain, circulatory problems, myocardial infarctions

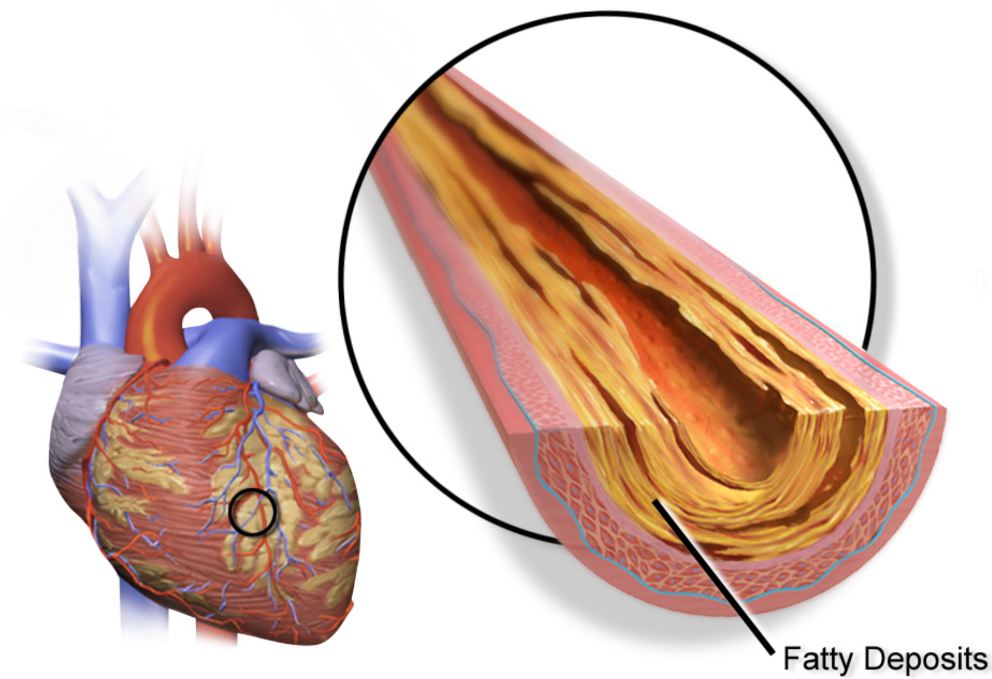

Blausen.com staff (2014). "Medical gallery of Blausen Medical 2014". *WikiJournal of Medicine* **1** (2).

Problem

- Objective:
  - Save lives
  - Minimize suffering and risks for the patient
  - Minimize costs
- Problem:
  - Uncertain outcome of actions
  - Incomplete information current state
- Approach:
  - Coronary artery disease and ischemic heart disease
  - Partially obervable Markov decision processes and solving them

Diagnostics

- EKG (rest and stress-test)
- Angiogram

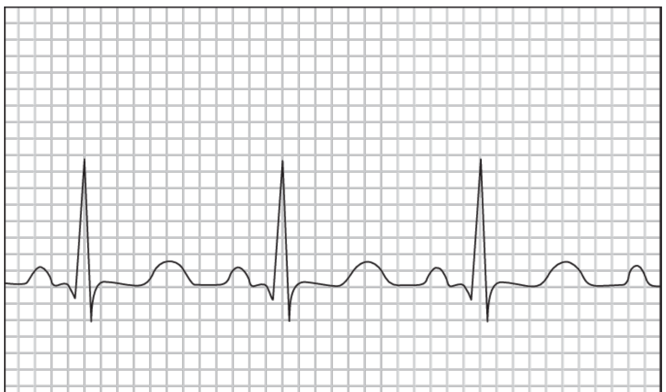

https://www.researchgate.net/figure/Example-of-ECG-curve_fig3_311319584

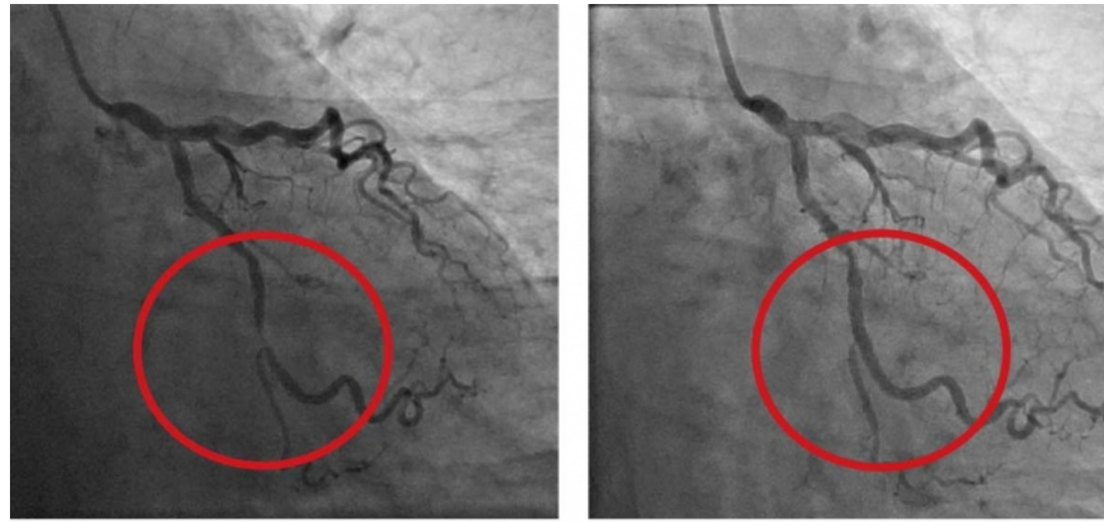

Coronary artery before angioplasty   Coronary artery after angioplasty

https://www.heartfoundation.org.nz/your-heart/heart-treatments/angioplasty-and-stents

Treatment actions

- Invasive
  - Angioplasty (PTCA)
  - Coronary artery bypass surgery (CABG)
- Non-invasive
  - No action
  - Medication

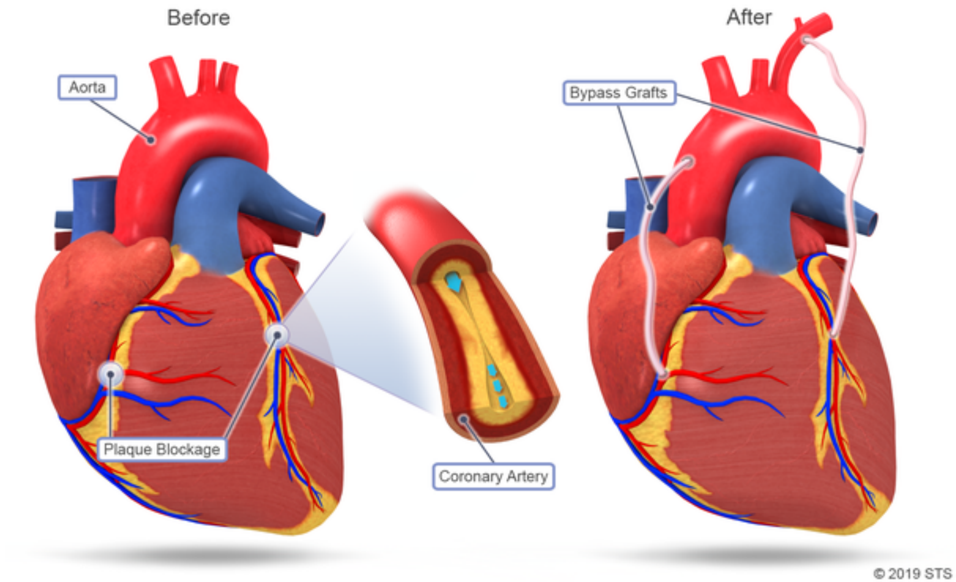

https://ctsurgerypatients.org/procedures/coronary-artery-bypass-grafting-cabg

Modelling with MDPs

- Allows finding an optimal policy with the optimal action for each state
- Solution found with dynamic techniques (value/policy iteration, linear programming)
- **Requires full observability**
- Presentations 5, 12, 13, 14

$$v(s) = \max_{a \in A} \sum_{s' \in S} P(s, a, s')(R(s, a, s') + f v(s'))$$

POMDPs

- Only some states observable
  - o observation
  - b belief state
  - a action
  - s start state
  - s' end state
  - β normalization factor
- Presentation 15
- Extremely hard to solve large problems in practice
- Heuristic techniques and approximations are needed

$$P(o|b,a) = \sum_{s' \in S} P(o|s',a)P(s'|s,a)b(s)$$

$$b'(s) = \beta P(o|s,a) \sum_{s' \in S} P(s|s',a)b(s')$$

Solution for a 3-action 2-state model

$$V_i^*(b) = \max_{\alpha \in \Gamma^i} \sum_{s \in S} \alpha(s) b(s)$$

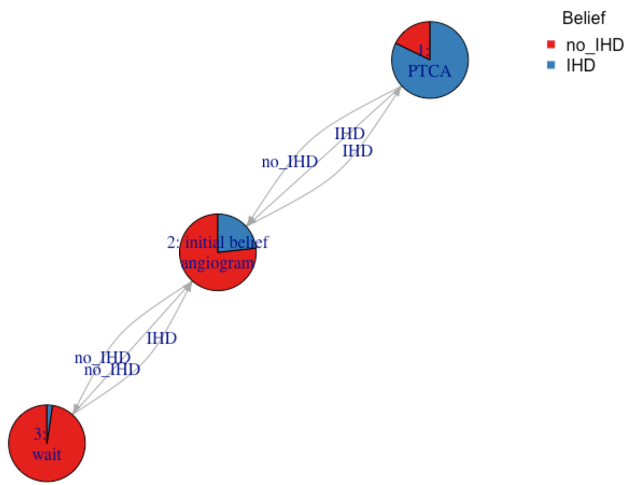

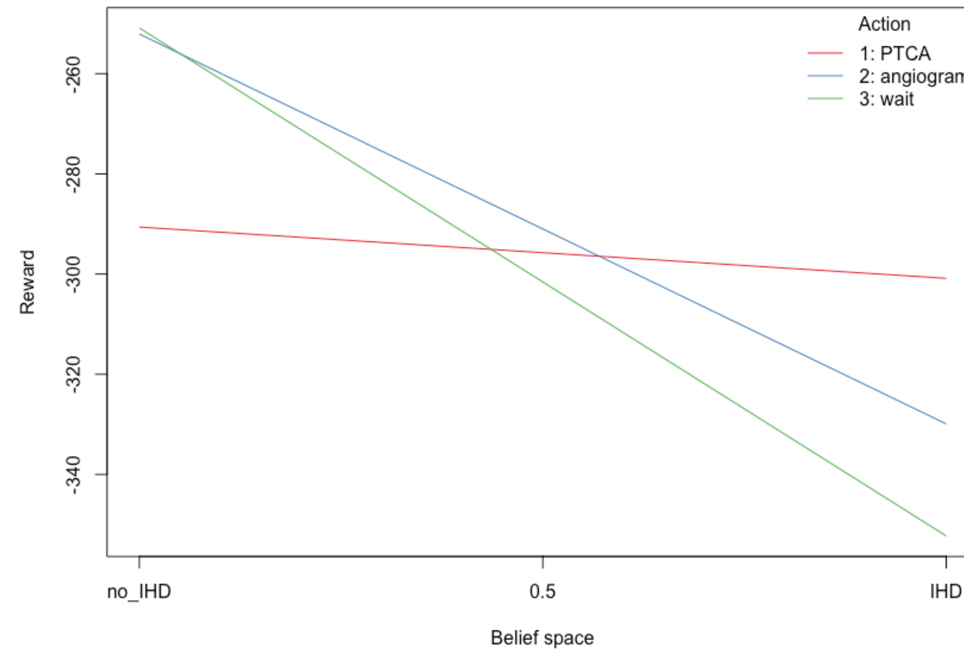

States

State variable	Possible values	Observable?
Alive	Yes, no	Yes
CAD	Normal, mild, severe	No
Ischemia level	None, mild, severe	No
Acute MI	Yes, no	Yes
Decreased ventricular function	Yes, no	Yes
History of PTCA	Yes, no	Yes
History of CABG	Yes, no	Yes
Chest pain	None, mild, severe	Yes
Resting EKG ischemia	Yes, no	Yes
Catheter coronary artery result	N/A, None, mild, severe	With action
Stress test result	N/A, unclear, negative, positive	With action

More structure

- Belief state only over states that directly mediate actions
- Divide belief state into observable and unobservable parts – hybrid state $\{o_d, b_d\}$
- Hierarchic state model: disable irrelevant variables
- Use combination of decision tree and dynamic programming for calculations

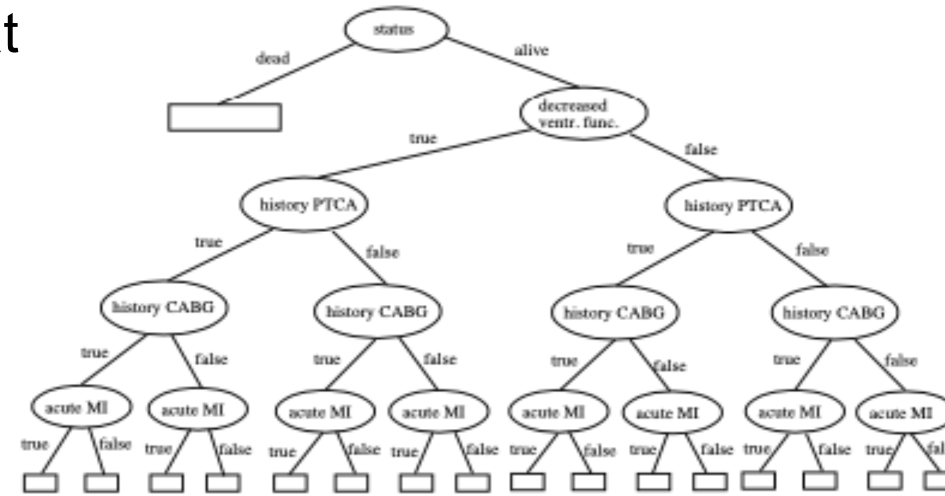

Hauskrecht & Fraser (2000)

Reward model

- Rewards (costs) from both actions and end states
- Infinite time horizon
- Discounting according to action times
  - Surgical procedures 1 day $\gamma = 1$
  - Non-invasive actions 3 months $\gamma = 0.95$
- Weights for each link
- Choose total cost and scale the other costs according to that
- AND or XOR

Weights and links

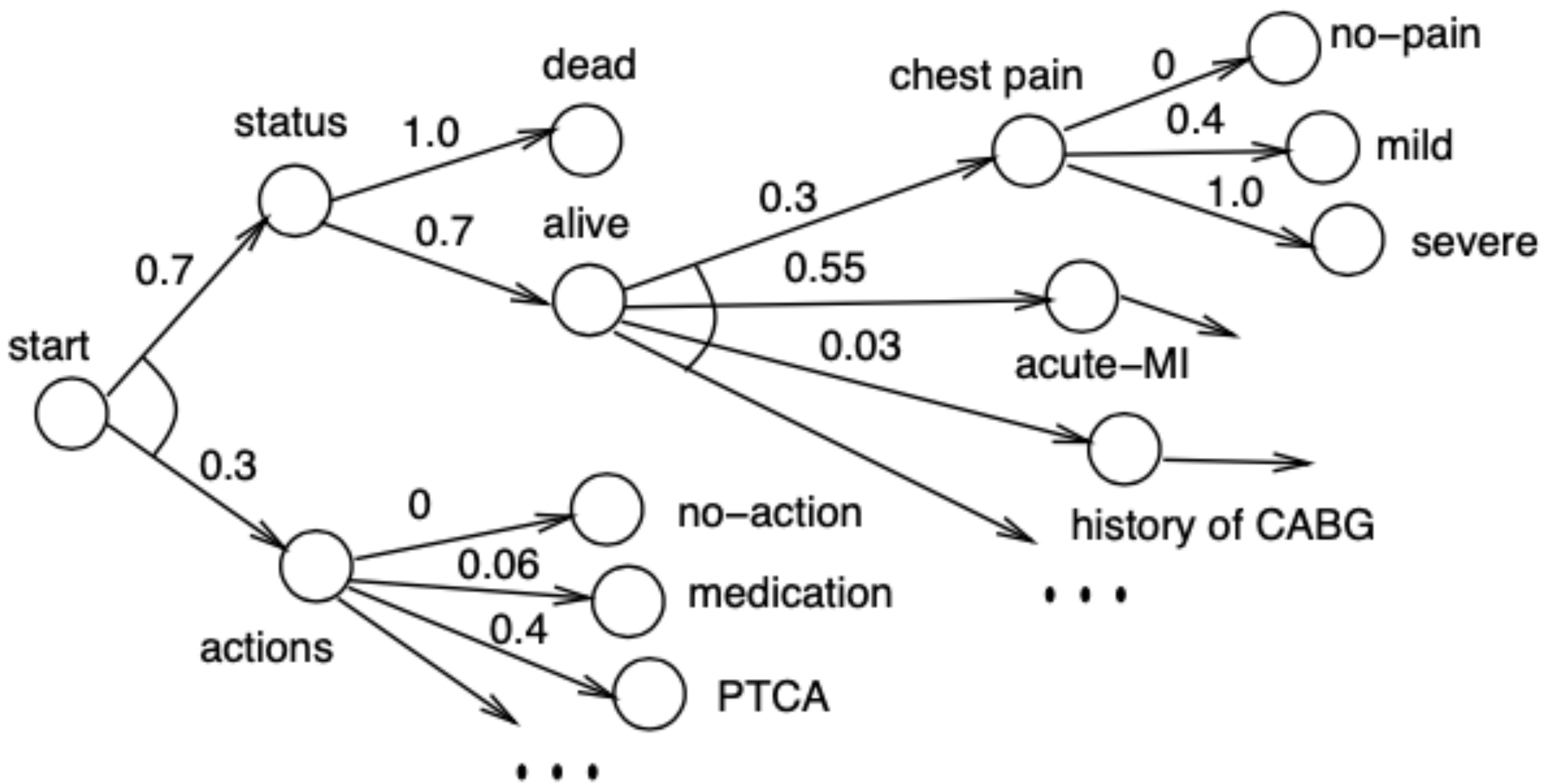

Hauskrecht & Fraser (2000)

Results

step	current patient status	actions	score (method 1)	score (method 2)
0	chest pain: mild-moderate;	**stress-test**	**285.22**	**248.53**
	acute MI: false;	no action	285.62	249.82
	rest EKG ischemia: negative;	medication	286.75	250.98
	decreased ventricular function: false;	PTCA	288.75	252.36
	catheter result: not available;	angiogram	292.92	256.68
	stress test result: not available;	CABG	491.94	427.77
	history CABG: false; history PTCA: false			

Hauskrecht & Fraser (2000)

Results

step	current patient status	actions	score (method 1)	score (method 2)
1	chest pain: mild-moderate; acute MI: false rest EKG ischemia: negative; decreased ventricular function: false; catheter result: not-available; stress test result: positive; history CABG: false; history PTCA: false	**PTCA**	**298.47**	**262.54**
		stress test	316.39	280.33
		no action	321.92	288.24
		medication	322.72	289.12
		angiogram	323.79	287.91
		CABG	503.73	440.77

Hauskrecht & Fraser (2000)

Results

step	current patient status	actions	score (method 1)	score (method 2)
2	chest pain: no chest pain; acute MI: false; rest EKG ischemia: negative; decreased ventricular function: false; catheter result: normal; stress test result: not available; history CABG: false; history PTCA: true	**no action**	**259.07**	**226.23**
		medication	260.62	227.78
		stress test	264.35	229.87
		angiogram	273.34	239.16
		PTCA	276.98	243.24
		CABG	481.36	417.28

Hauskrecht & Fraser (2000)

Summary

- Large POMDPs are hard!
- Identify the relevant state variables and assign values
- Simplify the problem by adding structure

References

Hauskrecht, M., & Fraser, H. (2000). Planning treatment of ischemic heart disease with partially observable Markov decision processes. Artificial Intelligence in Medicine, 18(3), 221-244.

Homework

ihd.R contains a simple model for treating IHD, where the patient either has or does not have ischemia. There are no other state variables.

The model contains two investigative actions (angiogram and stress test), two treatment actions (medication and PTCA), and an option to wait. Actions are chosen once a month, only one action at a time.

Add CABG as a third treatment action. The treatment is much more efficient and expensive compared to PTCA. How does the optimal policy change with your chosen values?

Send to emil.afbjorkesten@aalto.fi, DL 27.11.