Prostate Cancer Surveillance Using Data-driven Markov Decision Processes

Brian Denton, Weiyu Li

University of Michigan

July 18, 2019

Prostate Cancer (PCa)

- PCa is the most common cancer affecting men world-wide
- Many PCa's are low risk, slow growing, and **unlikely to cause harm**
- Widespread PCa screening has led to **overtreatment** – the treatment of cases that would not have caused clinical consequences if left untreated

Active surveillance for low-risk PCa patients reduces overtreatment

- **Active Surveillance** defers or avoids immediate curative treatment until there is evidence of disease progression

- Involves monitoring “low-risk” PCa with prostate specific antigen (PSA) tests and **prostate biopsies**

Biopsies can detect cancer but they cause harm to patients

- Biopsy involves sampling with hollow core needles
- Typically removes 12 cores of tissue from the different parts of the prostate
- Samples less than 1% of the entire prostate gland
- False negative results are common

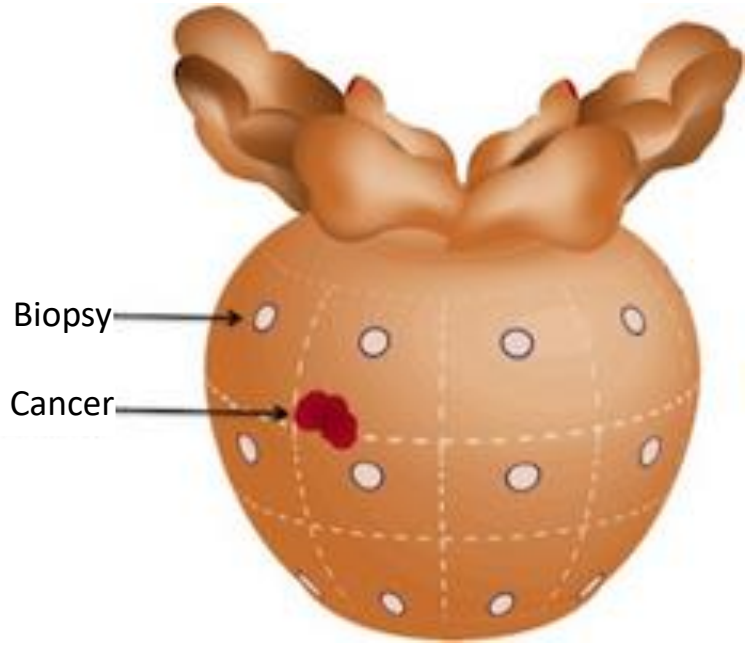

http://mdxhealth.com/confirmmdx-prostate-cancer

Doctors Don’t Agree on the Best Biopsy Strategy

Study	Recommended Biopsy Plan
Johns Hopkins (JH)	Annual biopsy
University of California, San Francisco (UCSF)	Biopsy 1 year after diagnosis, then every 1 to 2 years
PRIAS Project	Biopsy 1 year after diagnosis, then every 3 years

The PRIAS Project (Prostate Cancer Research International Active Surveillance) is a European initiative but open to all and has participation from countries all over the world.

Research Questions

- How often should patients on AS receive a prostate biopsy?
  - Under-testing → **delays detection** of progression and subsequent treatment
  - Over-testing → increases stress and **risk of infection** for patient
- Can strategies based on personalized risk improve AS?

Related Work

- Liver Transplants: Alagoz, Maillart, Schaefer, Roberts, *Management Science*, 2004
- Breast Cancer: Maillart, Ivy, Ransom, Dielhl, *Operations Research*, 2008
- HIV: Shechter, Schaefer, Roberts, *Operations Research*, 2008
- Prostate Cancer: Zhang, Denton, Balasubramanian, Shah, *M&SOM* 2012
- Adherence to Screening: Ayer, Alagoz, Stout, Burnside, Management Science, 2015
- Colorectal Cancer: Erenay, Alagoz, Said, *M&SOM* , 2014

Markov Decision Processes for Screening and Treatment of Chronic Diseases

Markov Decision Processes in Practice pp 189-222

Part of the International Series in Operations Research & Management Science book series (ISOR, volume 248)

- Lauren N. Steimle (1) Email author (steimle@umich.edu)
- Brian T. Denton (1)

1. Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, USA

A Data-driven Optimization Framework

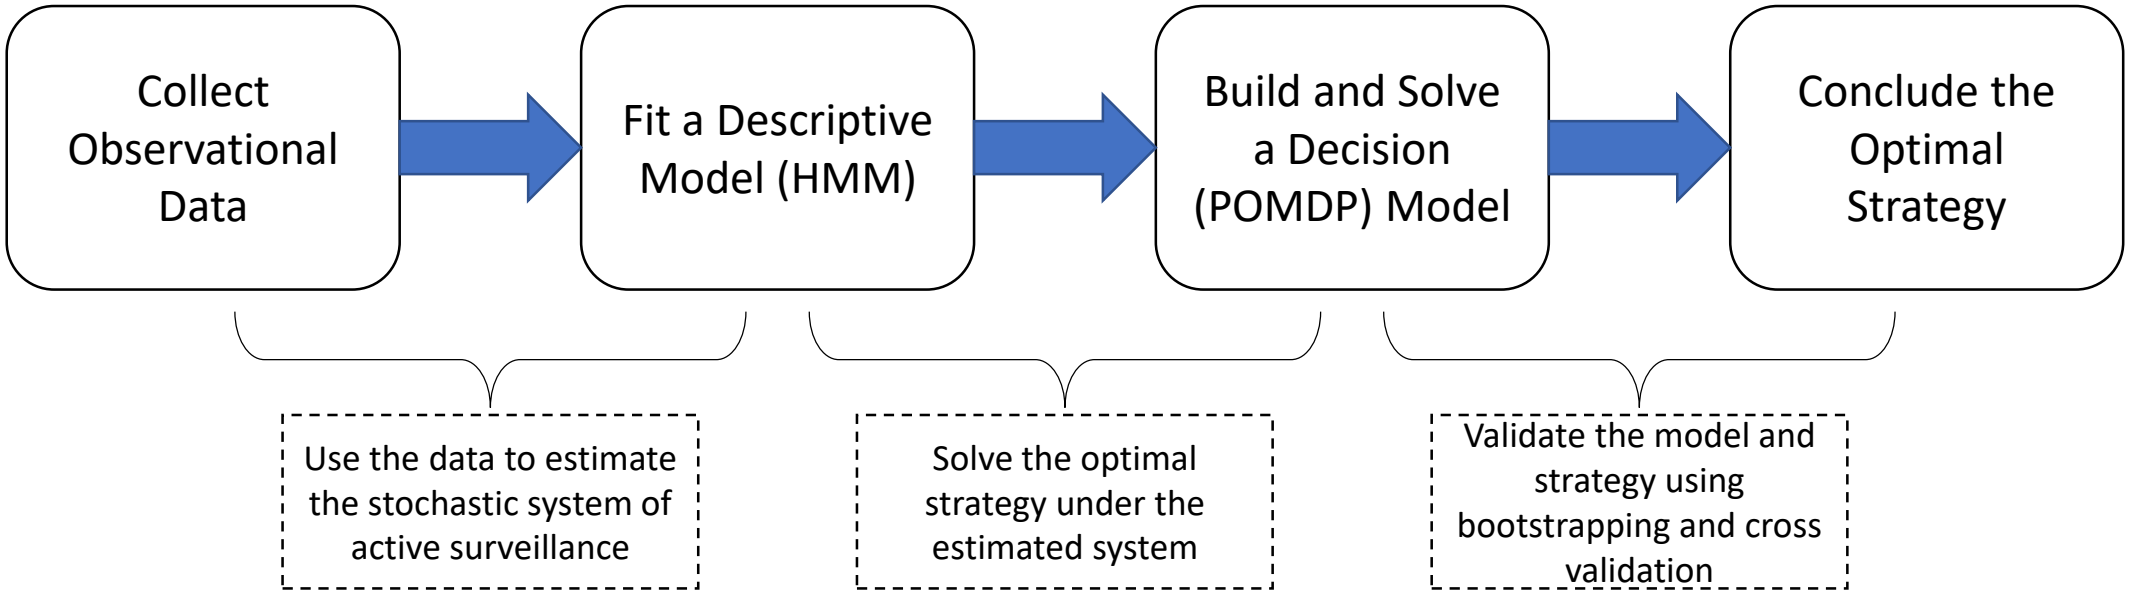

The Baum-Welch algorithm was used to estimate hidden Markov models

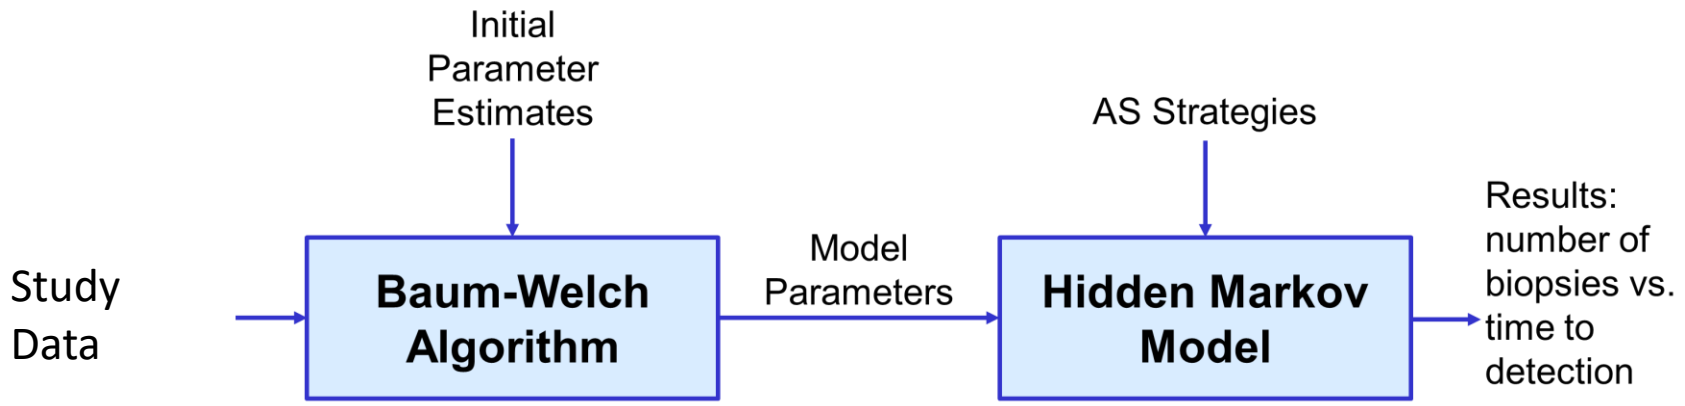

Given **observation sequences of PSA tests and biopsy results** and **initial model parameter** estimates, the Baum-Welch algorithm can find the model that locally maximizes the probability of the observation sequences

Data: The Movember Foundation's GAP3 Cohort

The Movember Foundation launched the Global Action Plan Prostate Cancer Active Surveillance (GAP3) to create a global database:

- includes 15,101 patients from 25 established AS cohorts worldwide
- records longitudinal observations of patients' clinical and demographic characteristics

Reference: Bruinsma, Sophie M., Liying Zhang, Monique J. Roobol, Chris H. Bangma, Ewout W. Steyerberg, Daan Nieboer, Mieke Van Hemelrijck et al. "The Movember Foundation's GAP3 cohort: A profile of the largest global prostate cancer active surveillance database to date." BJU international 121, no. 5 (2018): 737-744.

Hidden Markov Model (HMM)

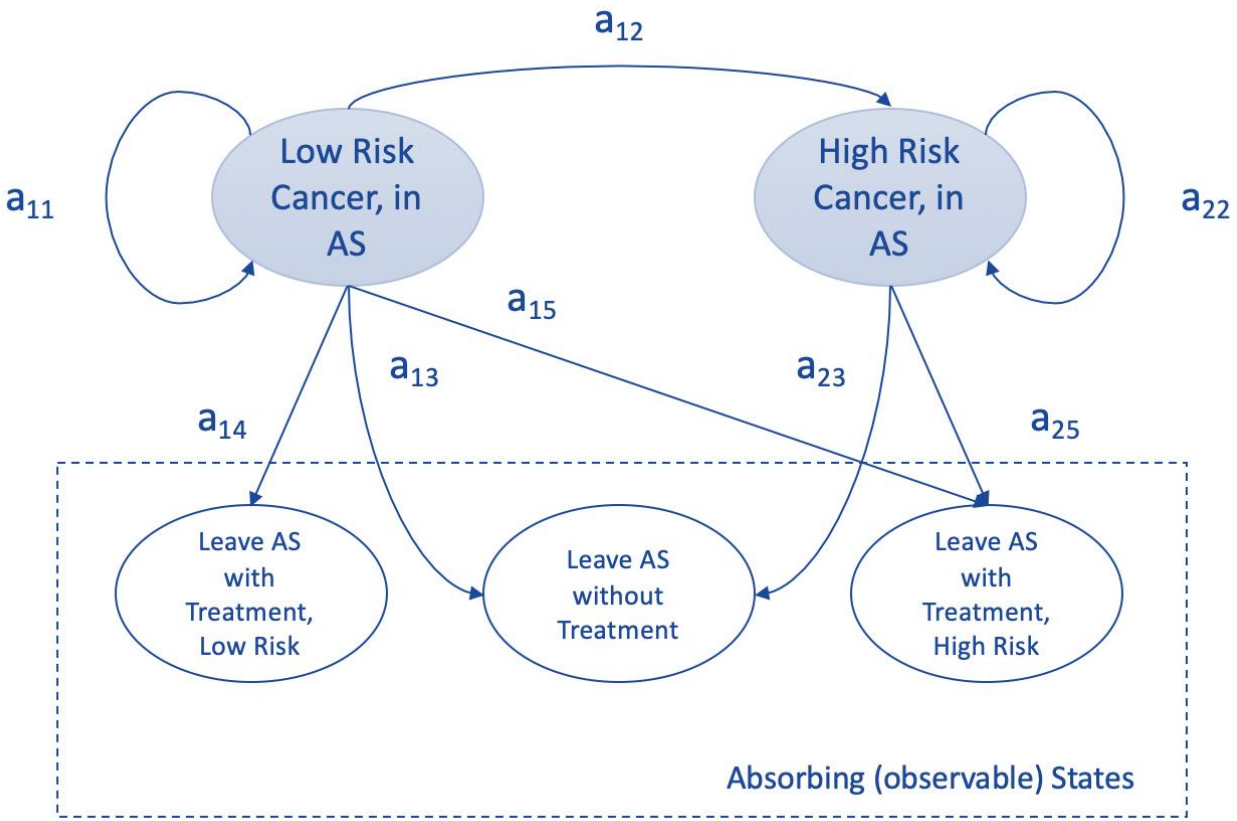

Figure: State Transition Diagram of PCa in the context of AS.

- Time periods: annual times from the start to the end of AS
- Initial distribution (at diagnosis)
  $$\phi = (\phi_1, 1-\phi_1, 0, 0, 0)$$
- Transition probability matrix:
  $$A = [\mathrm{P}(s_{t+1}|s_t)]$$
- Observation at time t:
  $$O_t = (X_t, Y_t)$$ (PSA, Biopsy)

Observation Probability Distributions of the HMM

- Biopsy observation probability matrix (discrete)

$$B = [P(y_t|s_t)]$$

- PSA observation density functions (continuous)

$$P(x_t|S_t = i) = \sum_{m=1}^{M} c_{im} \times \mathrm{Norm}(x|\mu_m, \sigma_m^2)$$

Baum-Welch Algorithm for Parameter Estimation

Given the observation sequences

$$O^{(1)} = \left(O_1^{(1)}, \dots, O_{T_1}^{(1)}\right), \dots, O^{(N)} = \left(O_1^{(N)}, \dots, O_{T_N}^{(N)}\right),$$

Baum-Welch algorithm, or equivalently *the EM (expectation-maximization) algorithm* estimates the model

$$\lambda = (A, B, \phi, c, \mu, \sigma)$$

that locally maximizes the likelihood function

$$P(O|\lambda) = \prod_{k=1}^{N} P(O^{(k)}|\lambda)$$

Reference: Rabiner, Lawrence R. "A tutorial on hidden Markov models and selected applications in speech recognition." Proceedings of the IEEE 77, no. 2 (1989): 257-286.

Results of HMMs: Estimated Parameter

Center	Number of Patients	Mis-classification Error at Diagnosis	Annual Grade Progression Rate	Biopsy Sensitivity (True Pos.)	Biopsy Specificity (True Neg.)
JH	1434	5.83% (1.74%)	6.91% (1.13%)	71.84% (0.43%)	99.72% (0.61%)
UCSF	1644	17.9% (0.95%)	14.77% (0.85%)	74.14% (0.89%)	99.06% (0.80%)
U of T	1243	14.2% (1.83%)	11.06% (2.29%)	78.62% (0.69%)	99.45% (0.75%)
PRIAS	4700	8.51% (1.44%)	8.71% (1.30%)	81.60% (0.72%)	98.96% (0.95%)

Table: Estimated Parameters (and bootstrapped standard errors) by the HMMs for Different Cohorts

Partially Observable Markov Decision Process (POMDP)

- Objective: to find the best biopsy strategy that balances the harm with the benefit of early detection
- Decision Epochs: every year
- Actions: PSA test only (No biopsy), PSA test and Biopsy

These elements define the decision process and goal

- Hidden States: Low-Risk Cancer, High-Risk Cancer
- Transition Probability Matrix: P
- Observations: PSA, Biopsy
- Biopsy Observation Probability Matrix: B
- PSA Observation Probability Matrix: Q

These elements in the POMDP model come from the HMM

POMDP Cost Function

Cost Function (negative reward): the cost function r_{iy}^{a} depends on the current state i, action a, and biopsy observation y:

$$r_{iy}^{a} = \begin{cases} 0, & a = \text{No Biopsy},\ i = \text{Low Risk} \\ \theta, & a = \text{No Biopsy},\ i = \text{High Risk} \\ \eta, & a = \text{Biopsy},\ i = \text{Low Risk},\ y = \text{Negative} \\ \eta, & a = \text{Biopsy},\ i = \text{High Risk},\ y = \text{Positive} \\ \theta + \eta, & a = \text{Biopsy},\ i = \text{High Risk},\ y = \text{Negative} \end{cases}$$

θ: cost of late detection

η: cost of biopsy burden

Both late detection and biopsy burden

Note:

- Set $\theta, \eta \in (-1, 0), \theta + \eta = -1$, i.e. give one unit of penalty for the worst case (both late detection and biopsy burden);
- A choice of (θ, η) can be interpreted as: θ years of late detection is treated as bad as the harm from η numbers of biopsies;
- Changing the values of θ and η allows trade off between two events according to patients' preference.

The Optimal Value Function

- Belief in high-risk cancer state

$$\pi^t \coloneqq P(S_t = \text{High Risk}), \forall t$$

- Optimal value function

$$V_t(\pi^t) \coloneqq \max_{a_t} \mathrm{E}\left[\sum_{m=t}^{\text{End}} \text{Reward}_m \,\middle|\, \pi^t, a_t\right], \forall \pi^t, \forall t$$

- Optimal equation

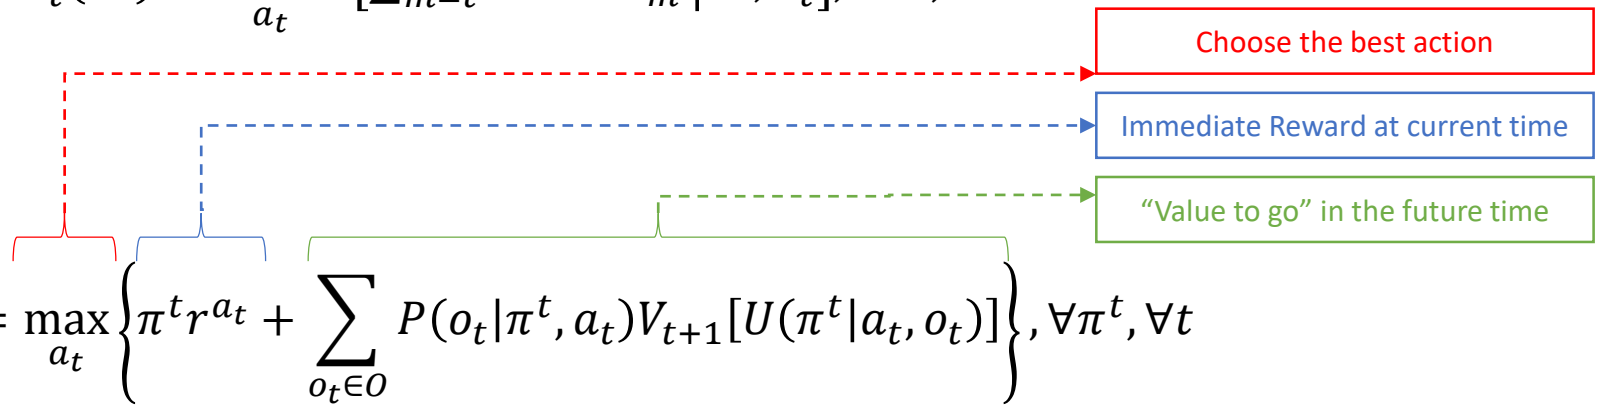

$$V_t(\pi^t) = \max_{a_t} \left\{ \pi^t r^{a_t} + \sum_{o_t \in O} P(o_t|\pi^t, a_t) V_{t+1}[U(\pi^t|a_t, o_t)] \right\}, \forall \pi^t, \forall t$$

where $U(\pi^t|a_t, o_t)$ is the updated belief calculated by the *Bayes formula*

- *Incremental pruning* algorithm (with approximation) to solve the POMDP (Cassandra et al. (1997))

Reference: Cassandra, A., Littman, M.L. and Zhang, N.L., 1997, August. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence* (pp. 54-61).

The Optimal Strategy

For a patient at time t, with the prior belief $\pi_{t-1}, \forall t$:

1. Perform a PSA test, receive the PSA observation: $\mathrm{PSA_t}$
2. Update his belief of being in high-risk cancer state:

$$\pi^t = P(S_t = \text{High Risk} | \pi^{t-1}, \mathrm{PSA}_t)$$

3. Choose the optimal action given by:

$$a_t^*(\pi^t) = \arg\max_{a_t} \left\{ \pi^t r^{a_t} + \sum_{o_t \in O} P(o_t | \pi^t, a_t) V_{t+1}[U(\pi^t | a_t, o_t)] \right\}$$

Results: Optimal Value at Age 50 in JH Study

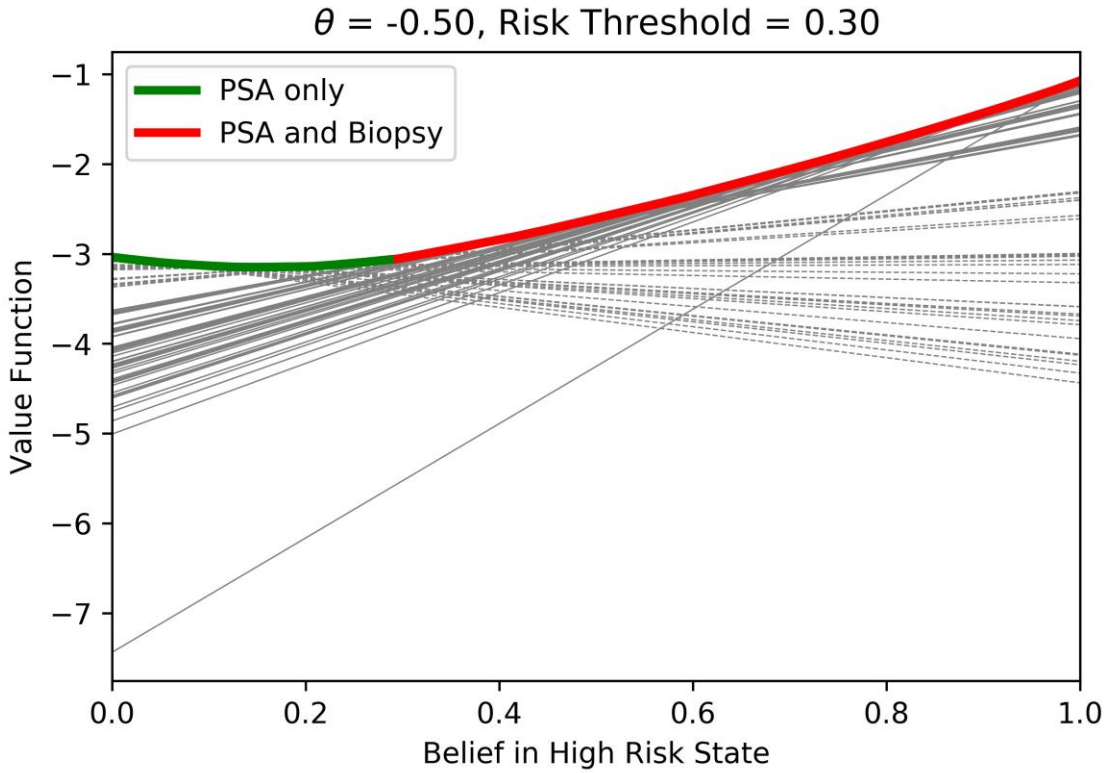

Figure: The optimal value functions at age 50 of the JH hospital's POMDP models for θ = -0.5

- The optimal value function is the maximum of many linear functions
- Setting $\theta = \eta = -0.5$ weights equally on late detection and biopsy burden
- The optimal policy is a threshold-based policy: if the belief of high-risk cancer state exceeds the threshold, then do biopsy

Results: Thresholds vs. Age

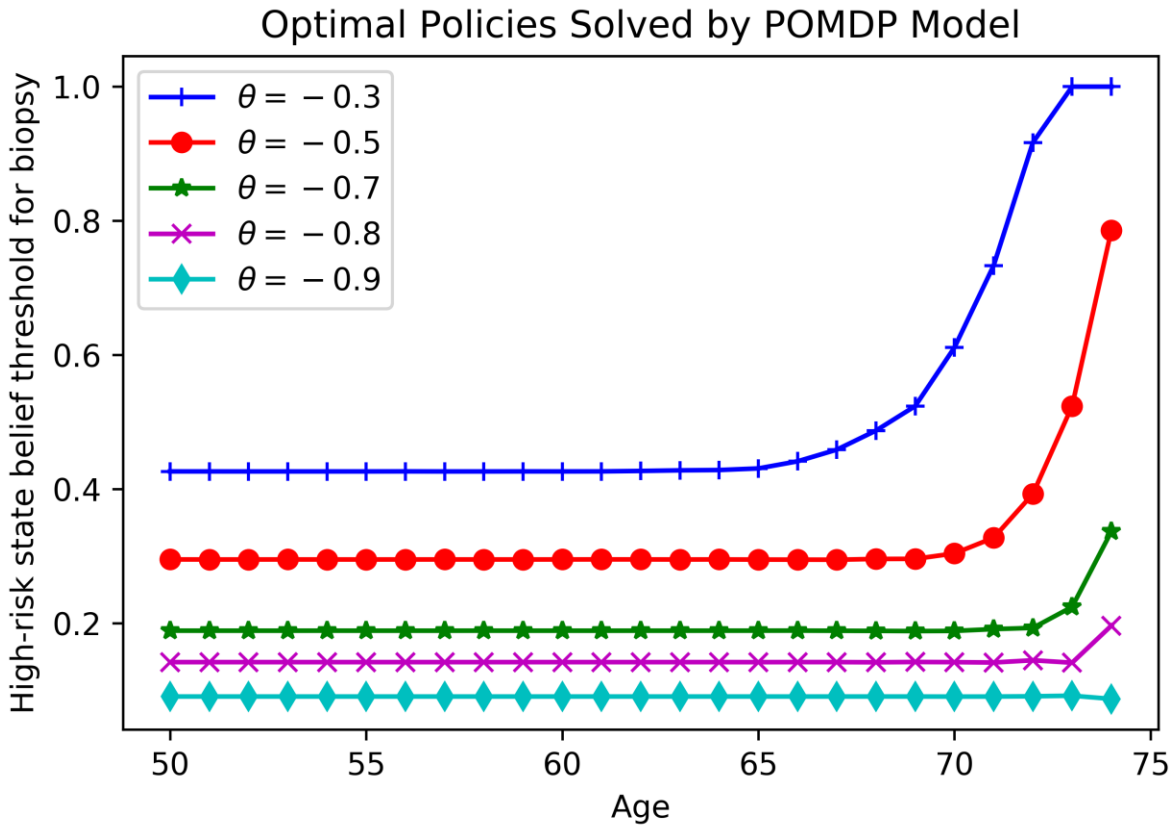

Figure 4: The relationship between the high-risk state belief thresholds for the biopsy and age for the optimal policies solved by the JH hospital's POMDP models.

- The threshold of the optimal policy depends on the value of θ and η
- The threshold varies little at early ages
- The threshold increases at older ages because years to recommended stopping time (75 year) for surveillance decrease

Comparisons with Current Policies in JH

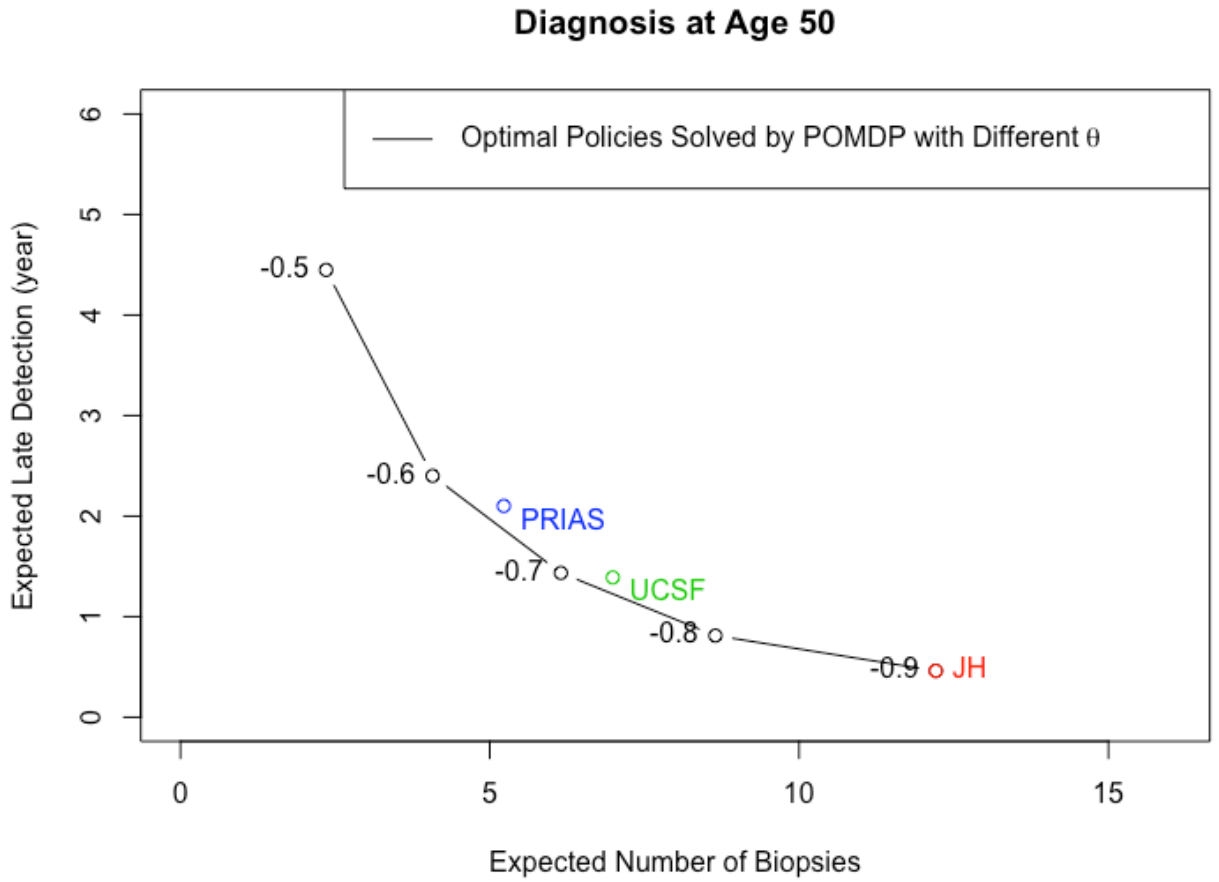

Figure: Comparisons of late detection and number of biopsies among the current policies and the optimal policies solved by the POMDP models in the JH's cohort if diagnosed at age 50.

- The current recommended policies (JH, UCSF, PRIAS) are dominated by the optimal policies from the POMDP model

- Strategies vary significantly on the basis of the tradeoff between competing goals of
  1) minimizing harm from biopsies
  2) minimizing time to detection of high risk cancer

Conclusions

- We demonstrate a data-driven optimization framework for active surveillance of prostate cancer, which can be applied to other healthcare applications with hidden states

- Current static policies are close to Pareto-optimal, but there is no single optimal strategy due to the trade off between competing criteria

- The optimal belief threshold for triggering a biopsy depends both on patient's preference and age

Opportunities for Future Research

- Robust optimal strategy given the estimated standard error of the model parameters
- The best single strategy that works well for all different cohorts
- Online learning of the natural history of cancer as well as the optimal strategy at the time of collecting new data

Brian Denton
Industrial and Operations Engineering
University of Michigan

btdenton@umich.edu

Related publications are on my website:

https://btdenton.engin.umich.edu

Simulation & Validation: PSA Density

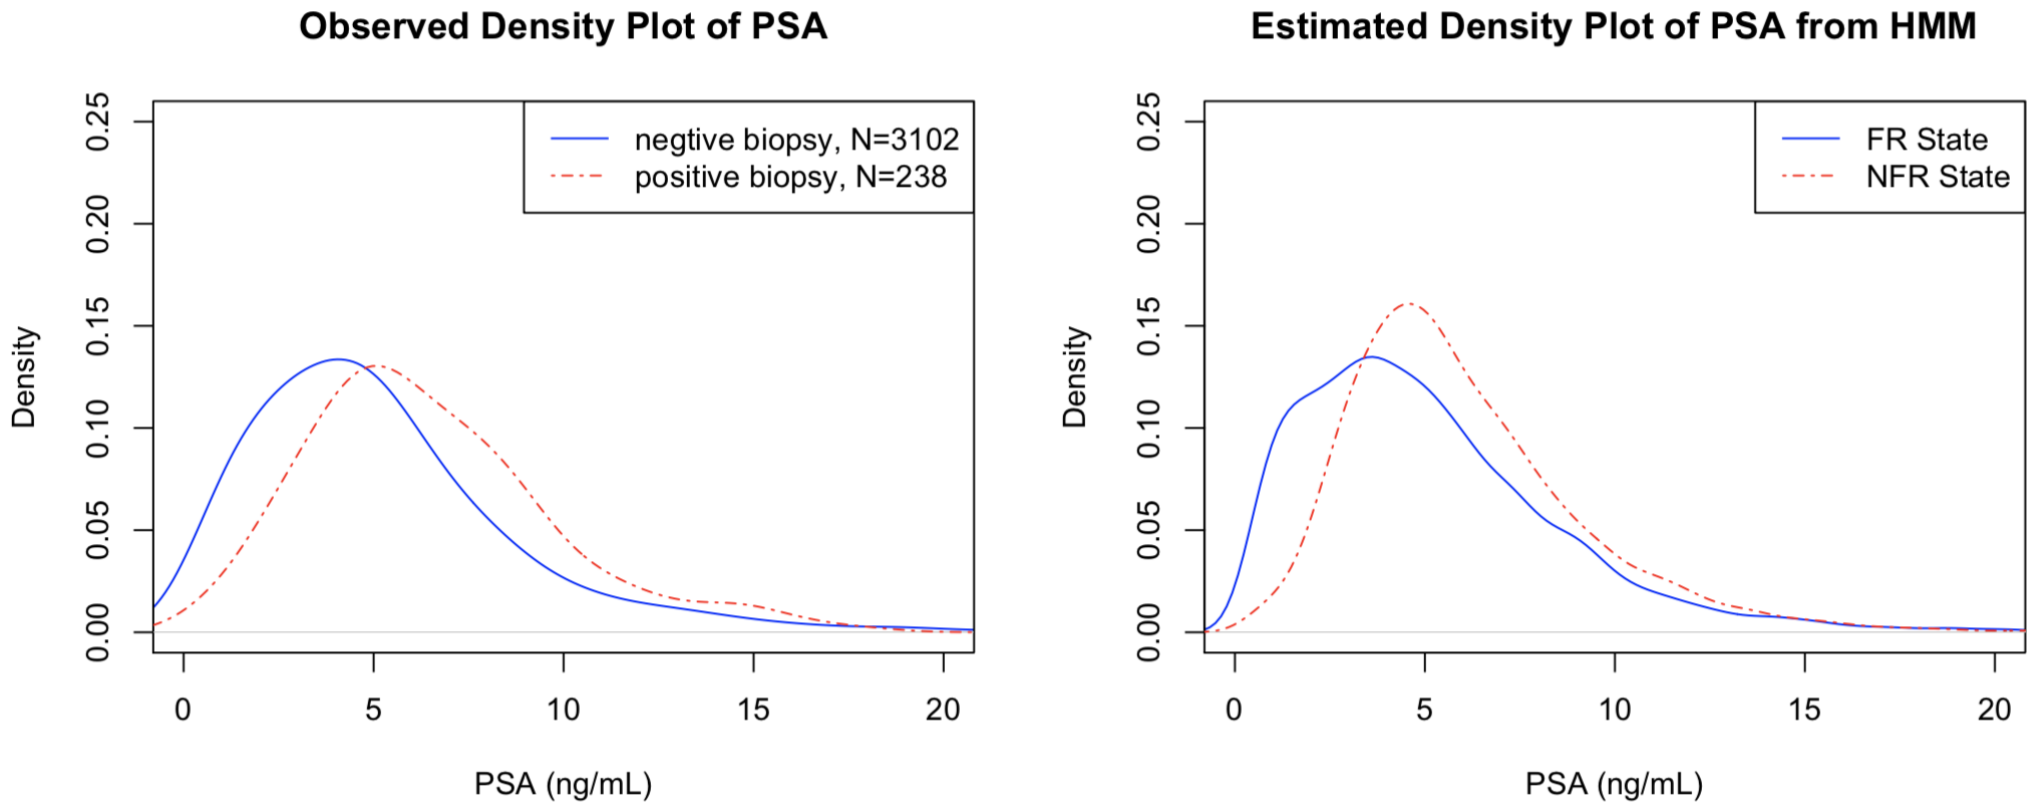

Figure 3: Comparison of the observed and estimated PSA density by HMM in JH hospital

Simulation & Validation: detection rates are consistent

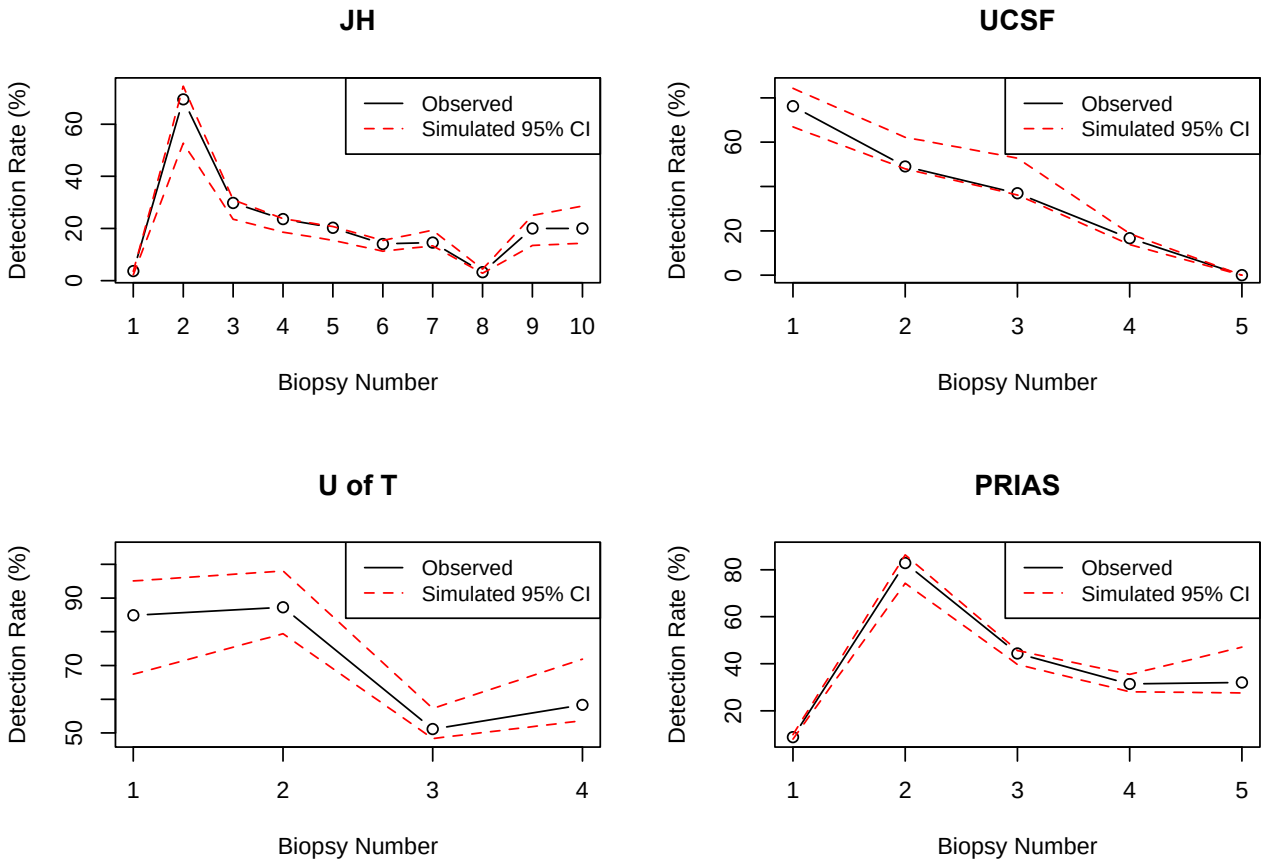

The comparisons of observed and simulated detection rates at each biopsy time for different cohorts. The detection rates at each biopsy time were calculated as the number of patients detected as high-grade cancer by the biopsy divided by the total number of patients who did this biopsy. All observed biopsy detection rates fell into the 95% CIs of the simulated detection rates.

The Sequence of Operations in the POMDP

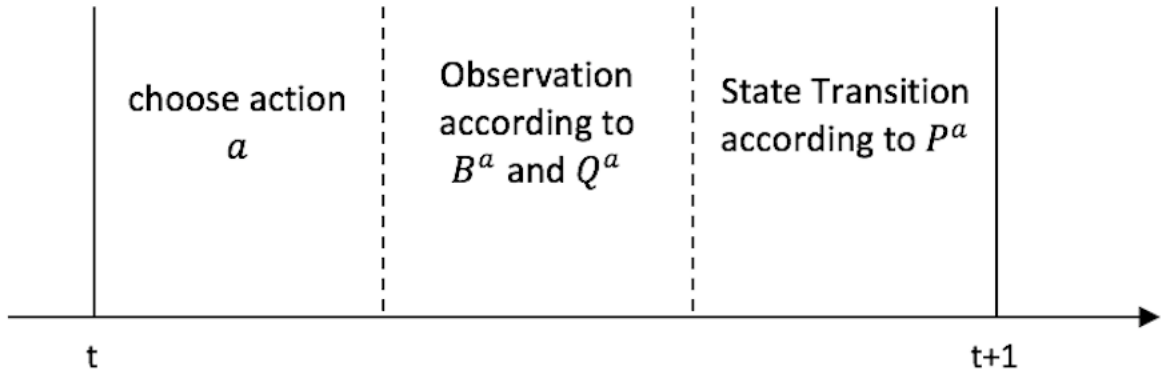

Figure 2: The sequence of operations

Hidden Markov Model (Cont.)

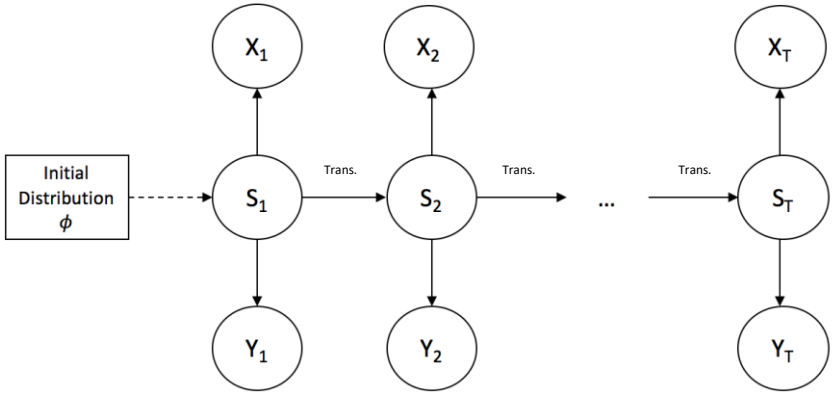

Figure 2: Bayesian Network Representation of the HMM

Multiple Types of Observations

- Under the Bayesian Network described by Figure 2, the following conditional independence holds

$$P(O_t|S_t) = P(X_t, Y_t|S_t) = P(X_t|S_t)P(Y_t|S_t), \forall t$$

- For biopsy (discrete observation of positive or negative): observation probability matrix

$$B_t = [p(y_t|s_t)]$$

- For log-PSA measurement (continuous value): assume a finite (M) mixture of Gaussian densities:

$$P(x_t|S_t = i) = \sum_{m=1}^{M} c_{im} \times \mathrm{Norm}(x|\mu_m, \sigma_m^2)$$

where c_{im} is the mixture coefficient for the mth mixture in state i, and can be interpreted as the probability that the log-PSA observation in state i was "drawn" from the mth mixture.

Results of HMMs: Estimated PSA Distributions

Range of PSA (ng/mL)		< 4	[4, 10]	>10
Johns Hopkins	LR Cancer	37.54%	54.97%	7.49%
	HR Cancer	31.16%	61.02%	7.82%
UCSF	LR Cancer	23.15%	65.54%	11.31%
	HR Cancer	22.53%	65.26%	12.21%
Toronto	LR Cancer	45.73%	34.22%	20.05%
	HR Cancer	33.12%	23.68%	43.20%
PRIAS	LR Cancer	13.79%	74.84%	11.37%
	HR Cancer	15.93%	72.05%	12.02%

Table: Estimated PSA distribution by the HMMs in different cohorts

Hidden Markov model simulation to evaluate AS strategies

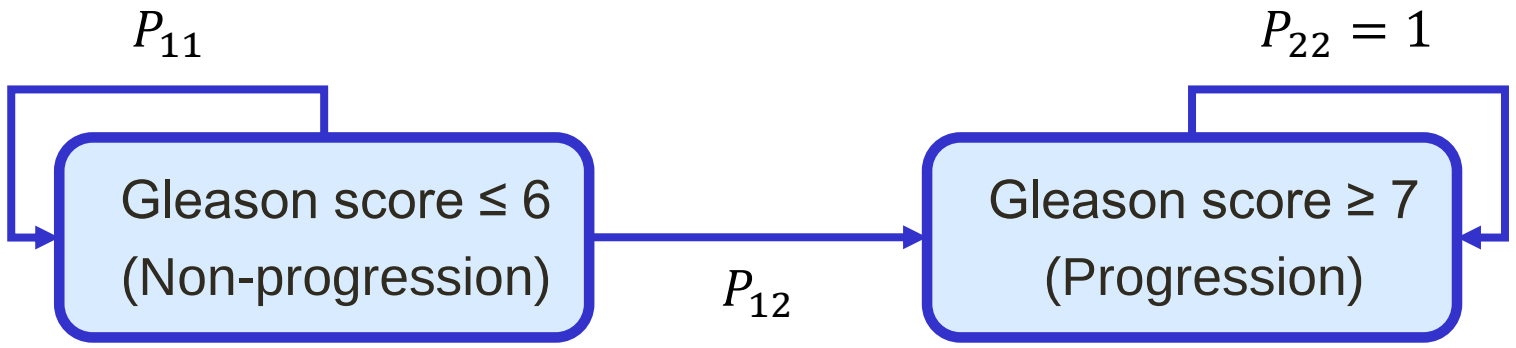

- Percentage of patients with GS ≤6 at diagnosis = 90.22%
- Annual progression probability = 3.97%
- Sensitivity of biopsy = 61.03%, Specificity of biopsy = 98.62%

Simulated all AS biopsy strategies 2^{10}= **1,024 AS strategies**

Low-risk

- Favorable Cancer state:
  - clinical stage $\leq$ T1c, PSA density $\leq$ 0.15, Gleason score $\leq$ 6, total positive core $\leq$ 2, single core positivity $\leq$ 50%
  - Recommend to stay in AS