IMPERIAL
BUSINESS SCHOOL

Balancing Complexity and Personalization

Personalized Medicine vs. One-Size-Fits-All

Multi-model Sequential Decision Models for Chronic Disease Management

Reza Skandari (PhD), Narges Mohammadi, Niloofar Zamani

Imperial Business School

Alexander Pearson (MD, PhD), Oncologist

University of Chicago

Markov, Semi-Markov Models and Associated Fields

Université Paris Sorbonne, July 2025

Personalized Medicine
Intro

What is it?
- Tailors medical care to individual characteristics
- Demographics, life-style, genetics, etc.

Alternative
- One size fits all.

Benefits
- Improves patient outcomes
- Improves equality in care outcomes

Operational complexities
- Hard to implement
- Training burden
- Risk of inconsistency

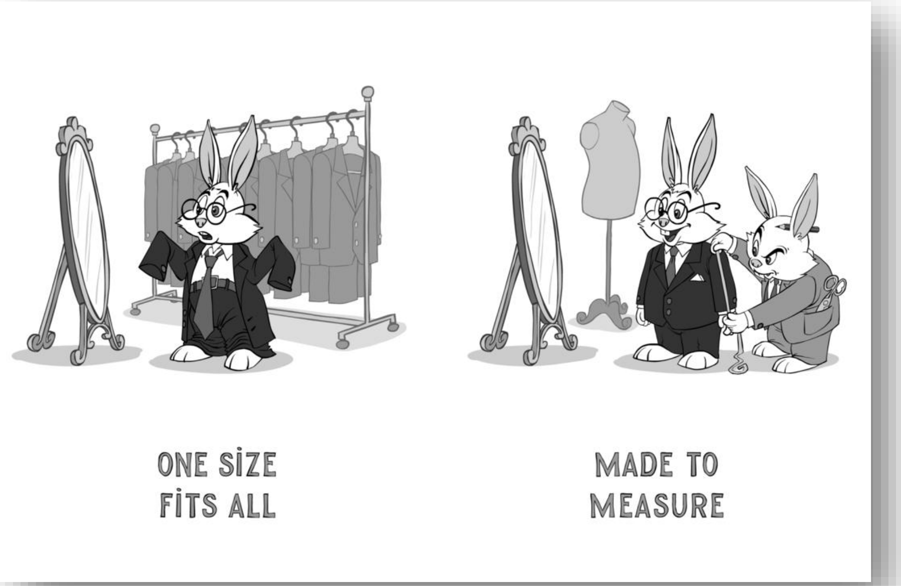

Research question
Optimal balance: trade-off

More personalization:
- ↑ effectiveness
- ↑ Fairness and equity in care outcomes
- ↑ complexity
  - ↓ efficiency

Potential solution
- Create a manageable number of patient groups
  - Handful of policies

Nested optimization
- Aggregate population into k groups
- Optimize and apply the *same* treatment strategy for a given group
  - Multi-model optimization

Benefits:
- Simplify clinical decision-making
  - Reduce complexity
- Preserves *some* level of personalization

But…
- May lead to inequality of health outcomes

Research question
Optimal balance

Current aggregation methods
- Based on clinical intuition rather than data/evidence
  - Based on disease prognosis
  - Intuition is often unaligned with operational decision
  - May overlook fairness and introduce disparities/bias

Proposed Strategy
- Create optimal clustering of patient profiles
  - Interpretable, systematic, and data-driven
- Consider fairness
  - Minimize deviations from care targets
  - timely treatment delivery

Case study
- Head and neck cancer
- Optimal post-treatment monitoring
- Care target
  - Cancer recurrence detection delay

Clinical context

Clinical context
Head and neck cancer

Incidence/prevalence
- 1 million *new* cases world-wide, annually
- 3rd most common cancer
  - HPV pandemic

Recurrence
- 50% one-year mortality in case of recurrence
  - Loco-regional or distant metastasis

Mortality and recurrence
- Demographics
- Disease history
- Lifestyle

Heterogeneity:
- 5-year recurrence rate:
  - Between 6% to 60%

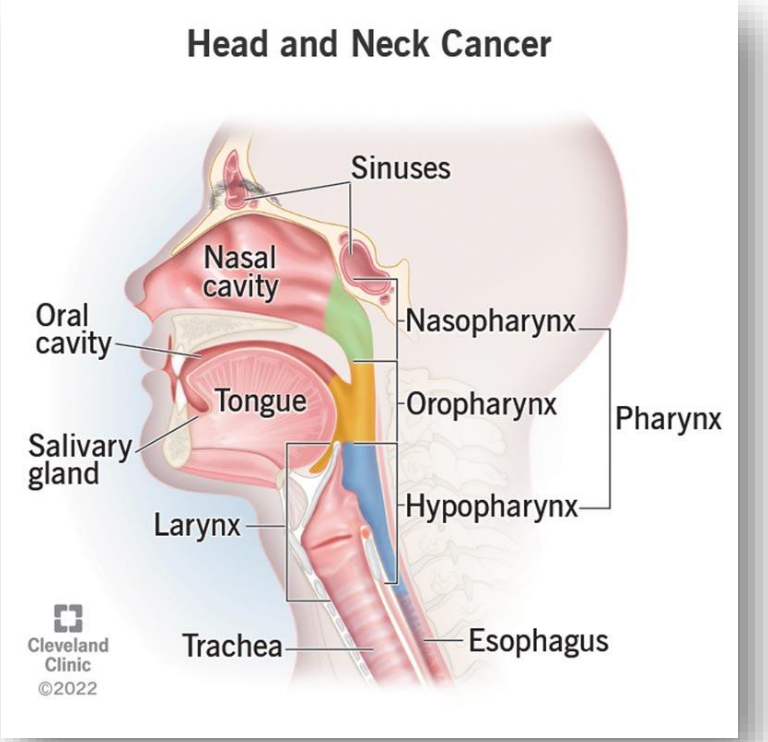

Clinical context
Pathways

Frequent monitoring
- Inform on patient's current health
- Inform on disease progression
  - Revise surveillance plan

Benefits
- Early detection
- Improved survival
  - 4 weeks reduction in delay
  - 10% less mortality
- Improved quality of life

Challenges:
- Capacity constraints
- Economic burden
- Adverse patient outcomes
  - Stress: false-positives
  - Inconvenience and pain
  - Exposure to radiation

Clinical context
Pathways

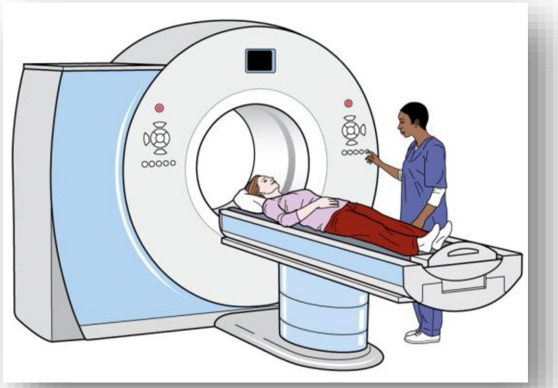

CT-Scan
- Costly
- High false-positive/negative rate
- *Pinpoints recurrence location*
  - Loco-regional or metastasis

Newer technology: ctDNA
- Simple blood test
- Cheaper
- Ease of access and implementation
- Improved adherence
- High accuracy (sensitivity, specificity)
- Detect radiologically occult (undetectable) recurrence
  - 4 Mo in advance
- *Cannot delineate loco-regional or metastasis*

Research gap
Operational question

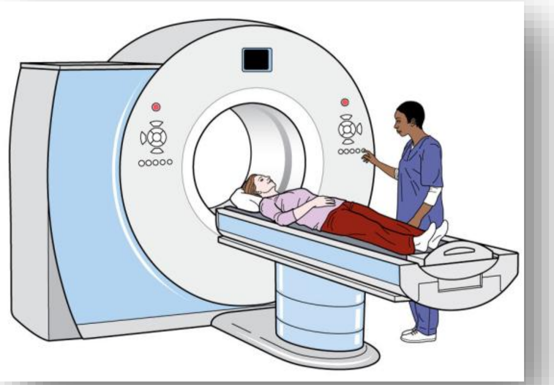

How to integrate ctDNA?
- No data-driven guideline

Operational questions:

Timing
- First test
- Subsequent tests

Modality
- ctDNA or CT scan?

Respond to test result
- Confirmatory tests
  - CT to localize
  - Biopsy to confirm
- Revise test schedule

Evaluation of status-quo.

Clinical context
Current guidelines

NCCN guideline, v1.2021:
- PET/CT scan: 3 Mo
- CT scan: 6 Mo, 9 Mo, 1 Y, 18 Mo, 2 Y, 3 Y

eviCore guideline, v2.1 :
- PET/CT scan: 3 Mo
- CT scan : 6 Mo, 1-3 Y

Kowalchuk et al. 2023:
- ctDNA: every 3 Mo, in years 1-2

Berger et al. 2022:
- ctDNA: every 3 Mo in years 1–2
- ctDNA: every 6 Mo in years ≥3

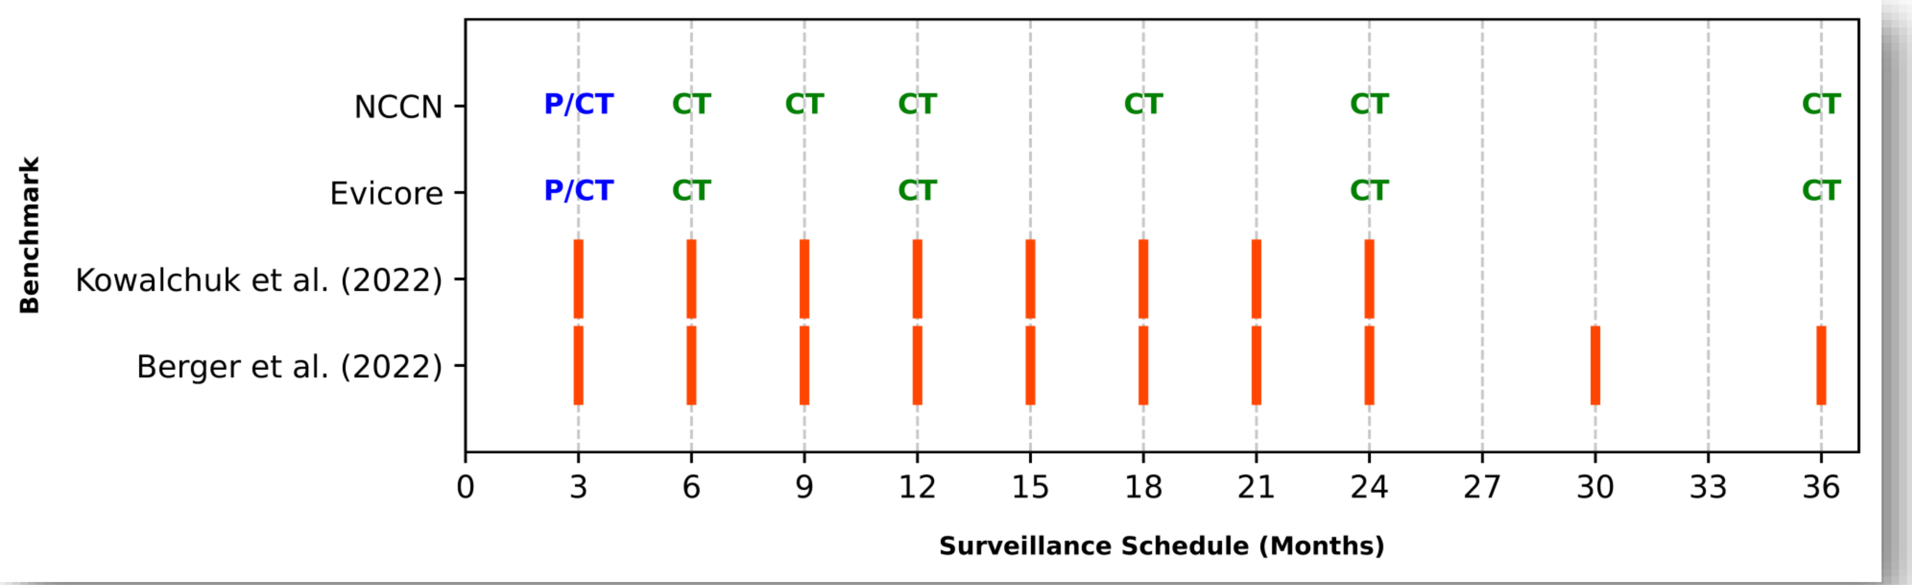

Clinical context
Evaluation

Outcomes

- System outcome
  - Testing cost
- Patient outcomes
  - Recurrence detection delay
  - False positive rate

Method

- Monte-Carlo simulation

Disease model

- Beesley et al. 2021
- Multi-state continuous-time semi-Markov model

Evaluation
Hidden Semi-Markov model

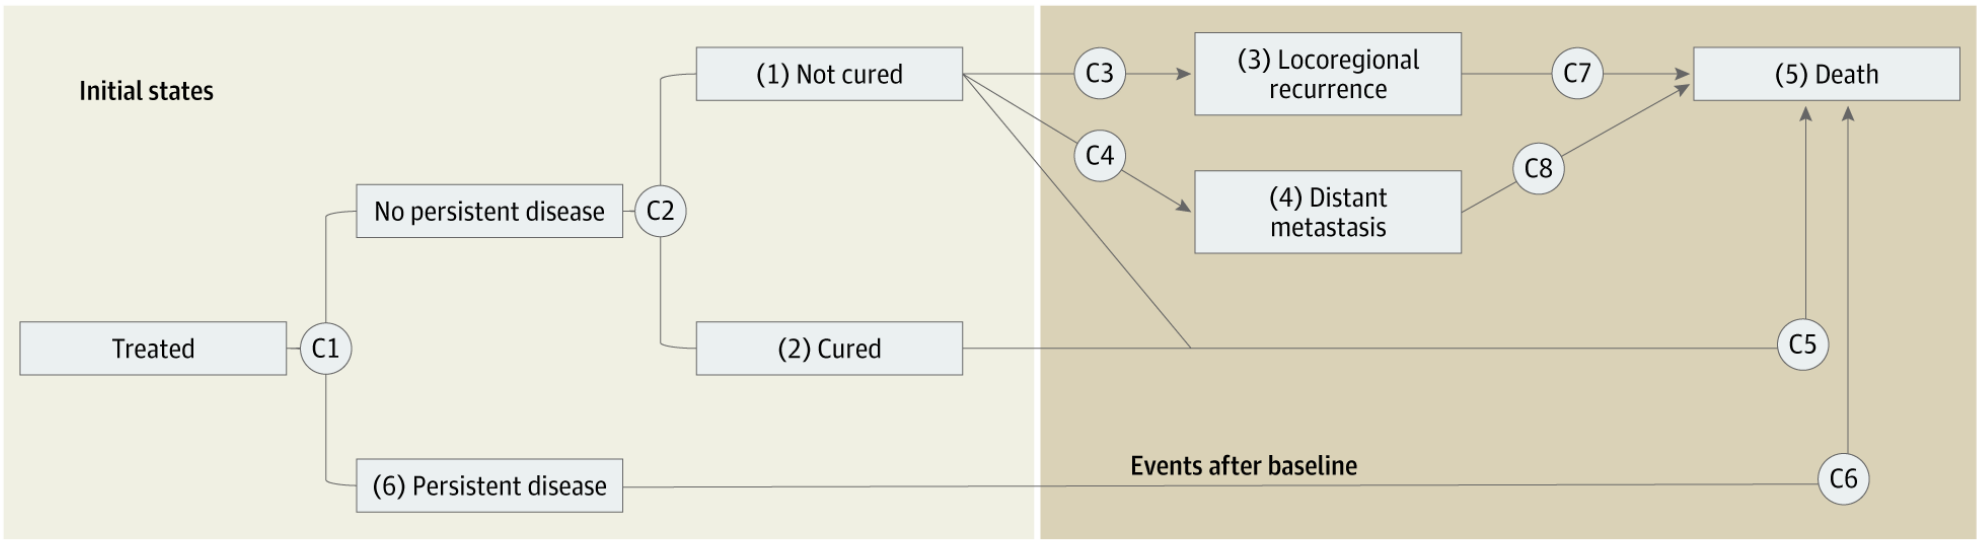

Evaluation

Hazard ratios

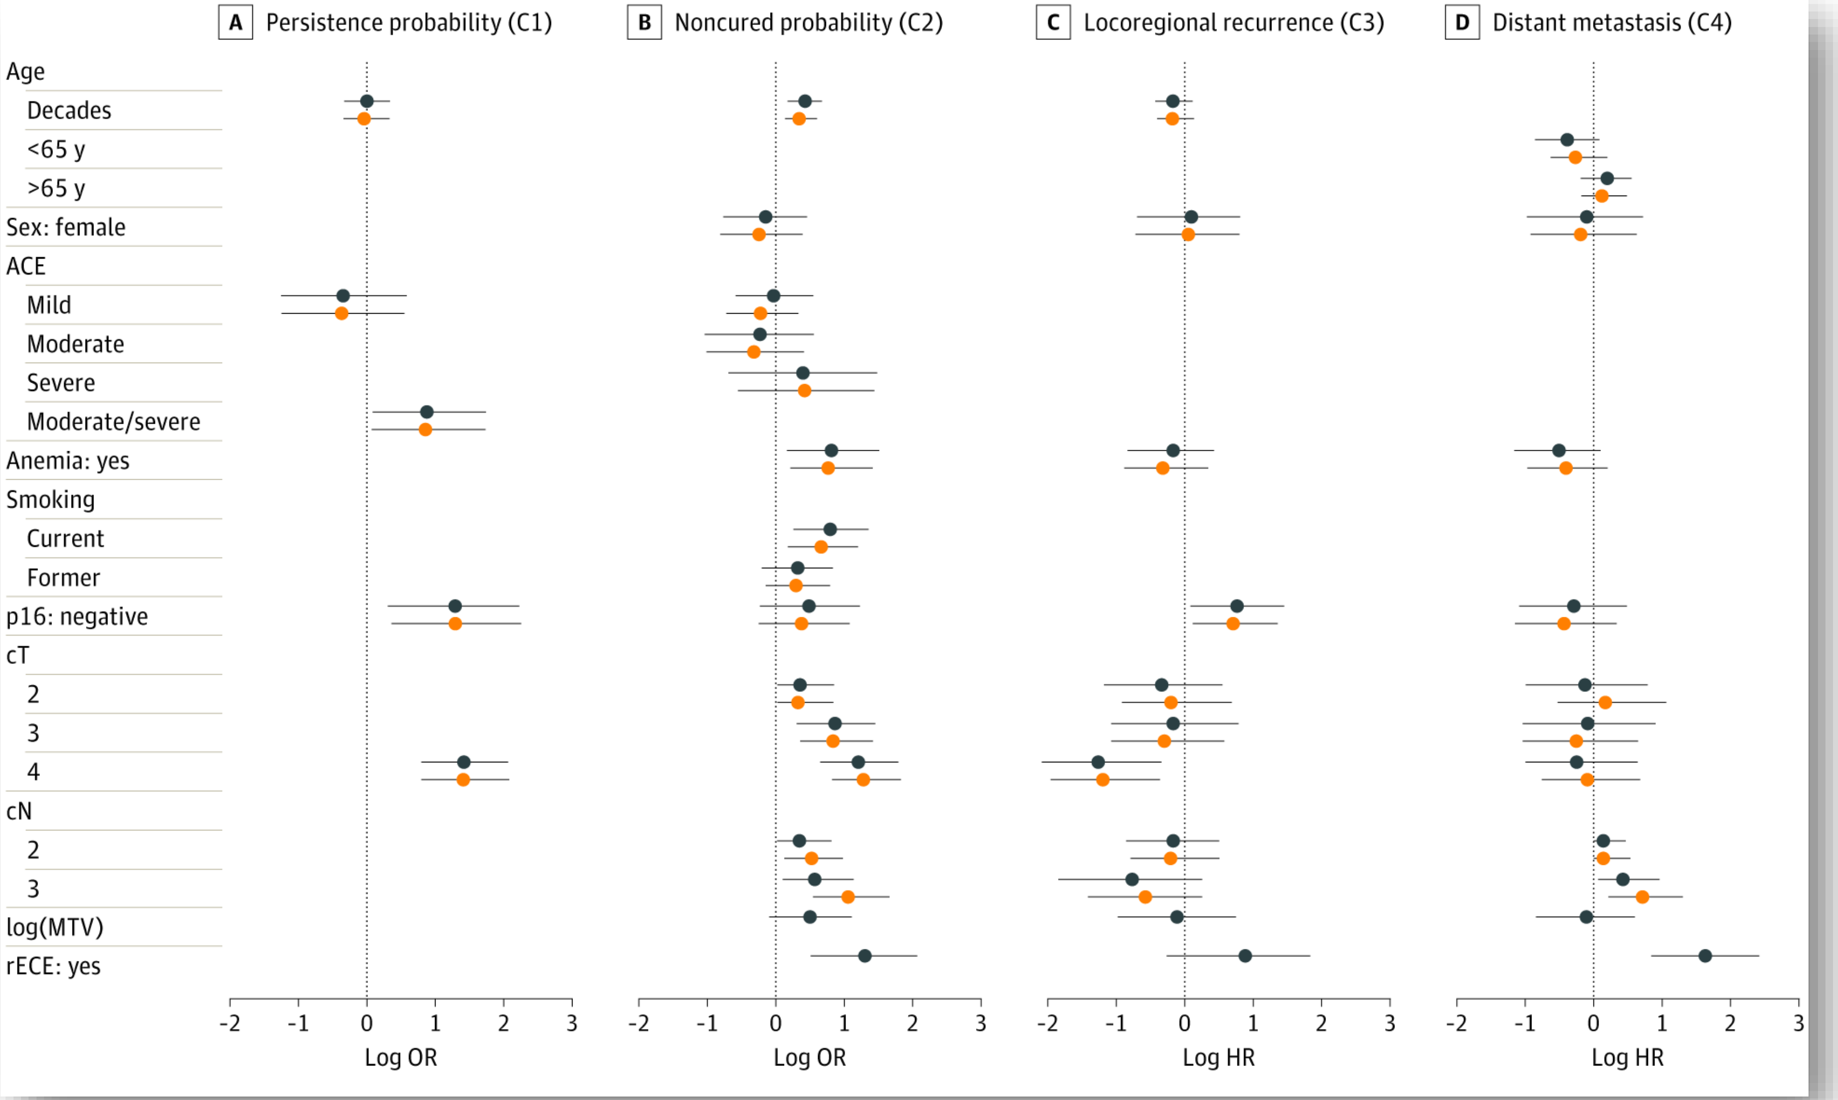

Clinical context
NCCN Guideline Evaluation

Testing cost
- $26,700 per patient

False positive rate
- 56%

Recurrence detection delay
- Approx. 6 Mo
  - Care target for personalizing

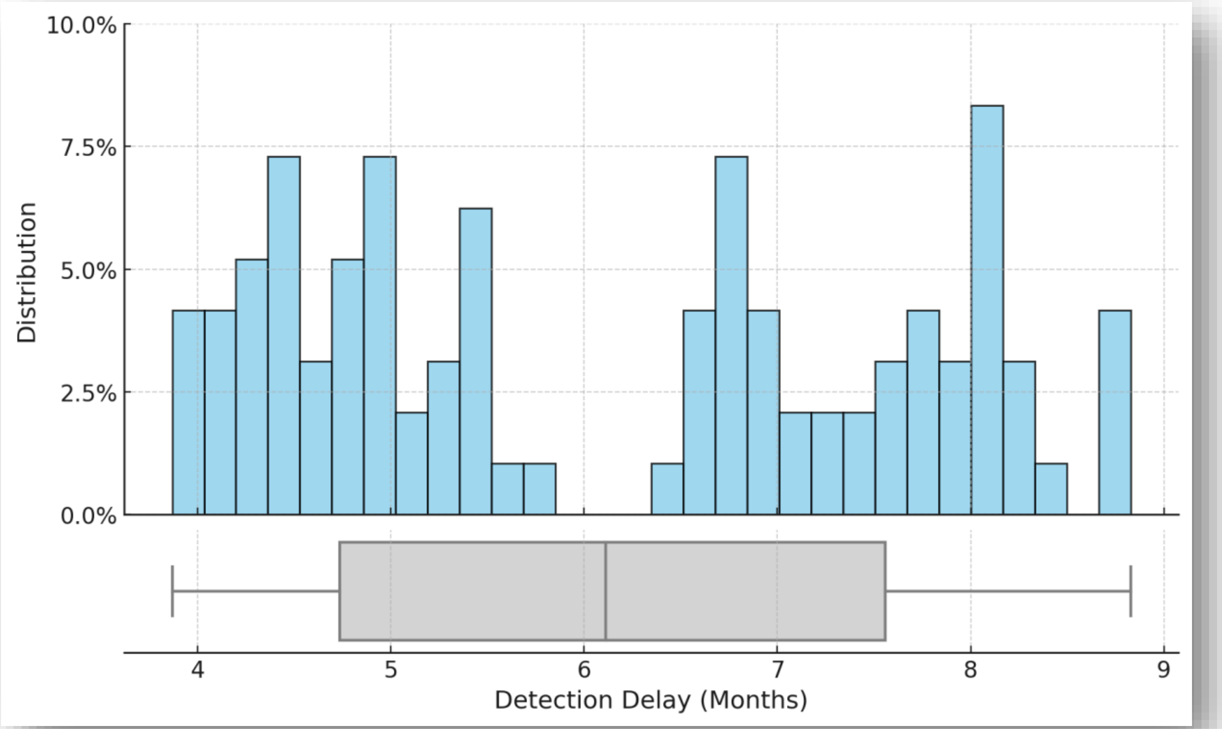

Optimize monitoring guideline

Optimization model
Personalized model

Epochs:
- Monthly

Actions:
- CT scan
- DNA test
- Wait

Horizon:
- 3-5 years

Objective:
- Minimize monitoring cost

Constraint:
- Recurrence detection delay:
  - Target: 6 Mo

Model:
- Partially observable Markov Decision model
- Constrained POMDP

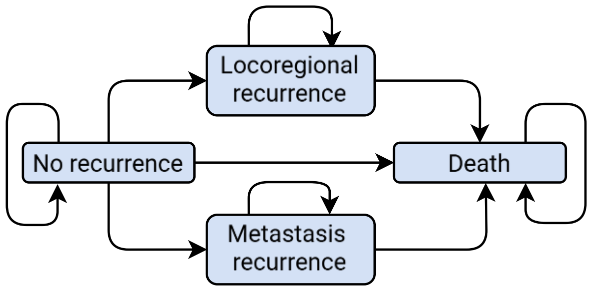

Optimization model
Personalized model

Core states
- Hidden:
  - Recurrence-free
  - Loco-regional recurrence
  - Metastatic recurrence
- Death

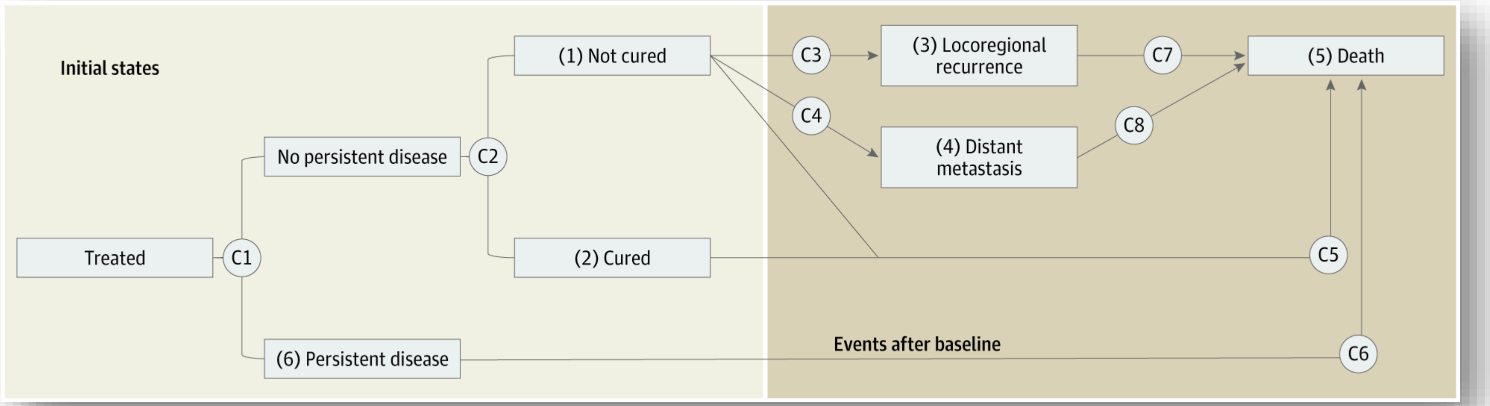

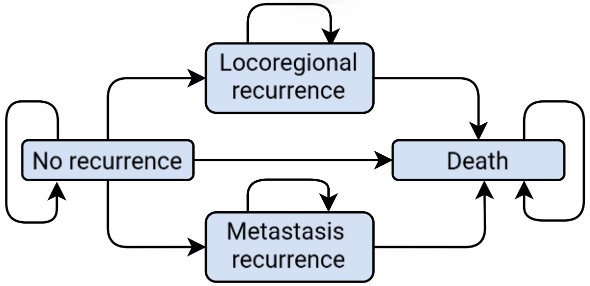

Optimization model
Personalized model

Markov model

- Converted semi-Markov model to an *equivalent* Markov model
  - Match occupancy measures
  - Non-stationary core state transitions

Multi-reward Markov model

- Monitoring costs
- Delay costs
  - One time unit for being in recurrence states, undetected

Constrained POMDP Solution

- Weighted sum of the two objectives
- Fine-tune the weight to match the target delay
  - Apply binary search

Personalized model
Optimal policy

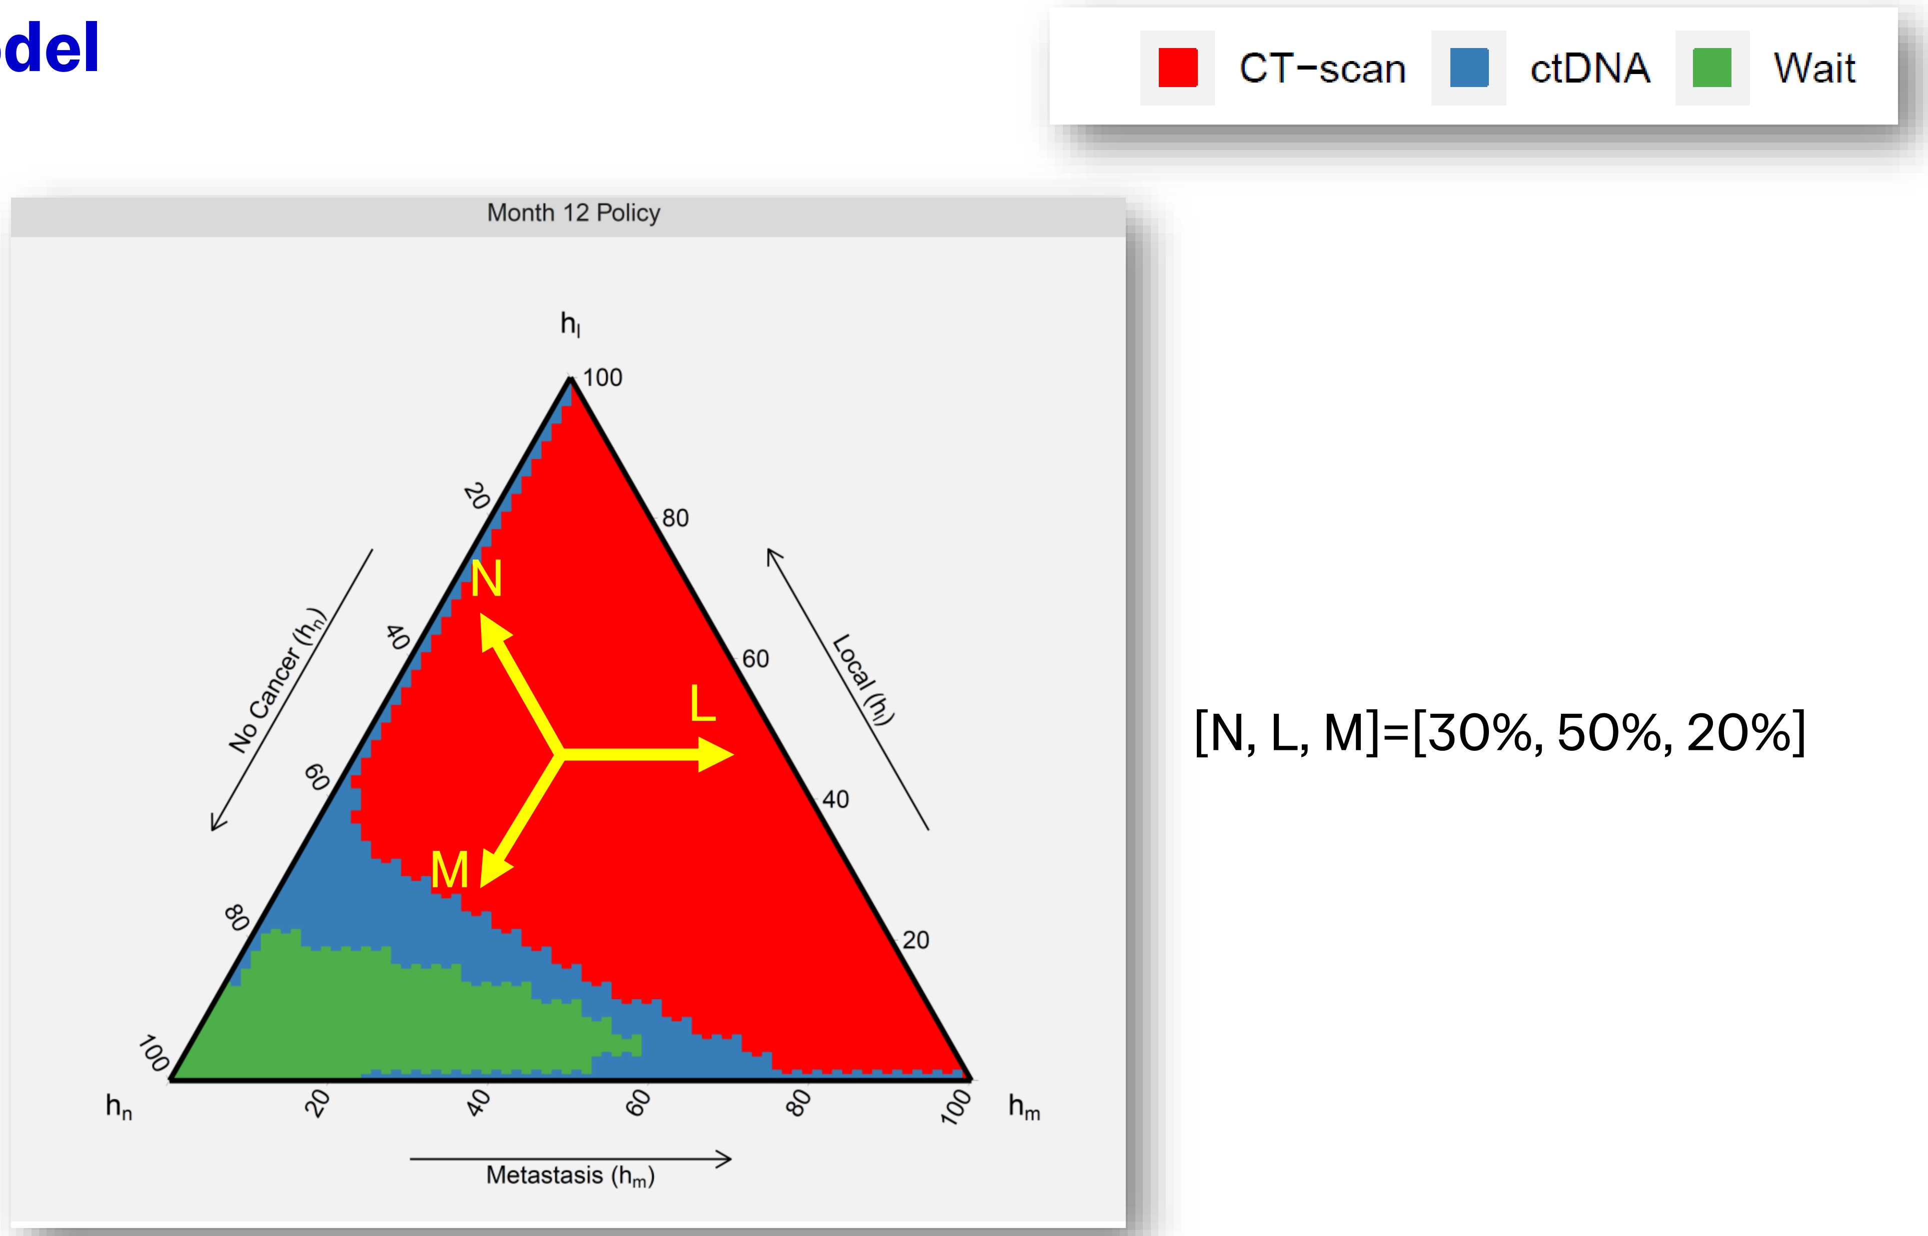

[N, L, M]=[30%, 50%, 20%]

Personalized model
Optimal policy

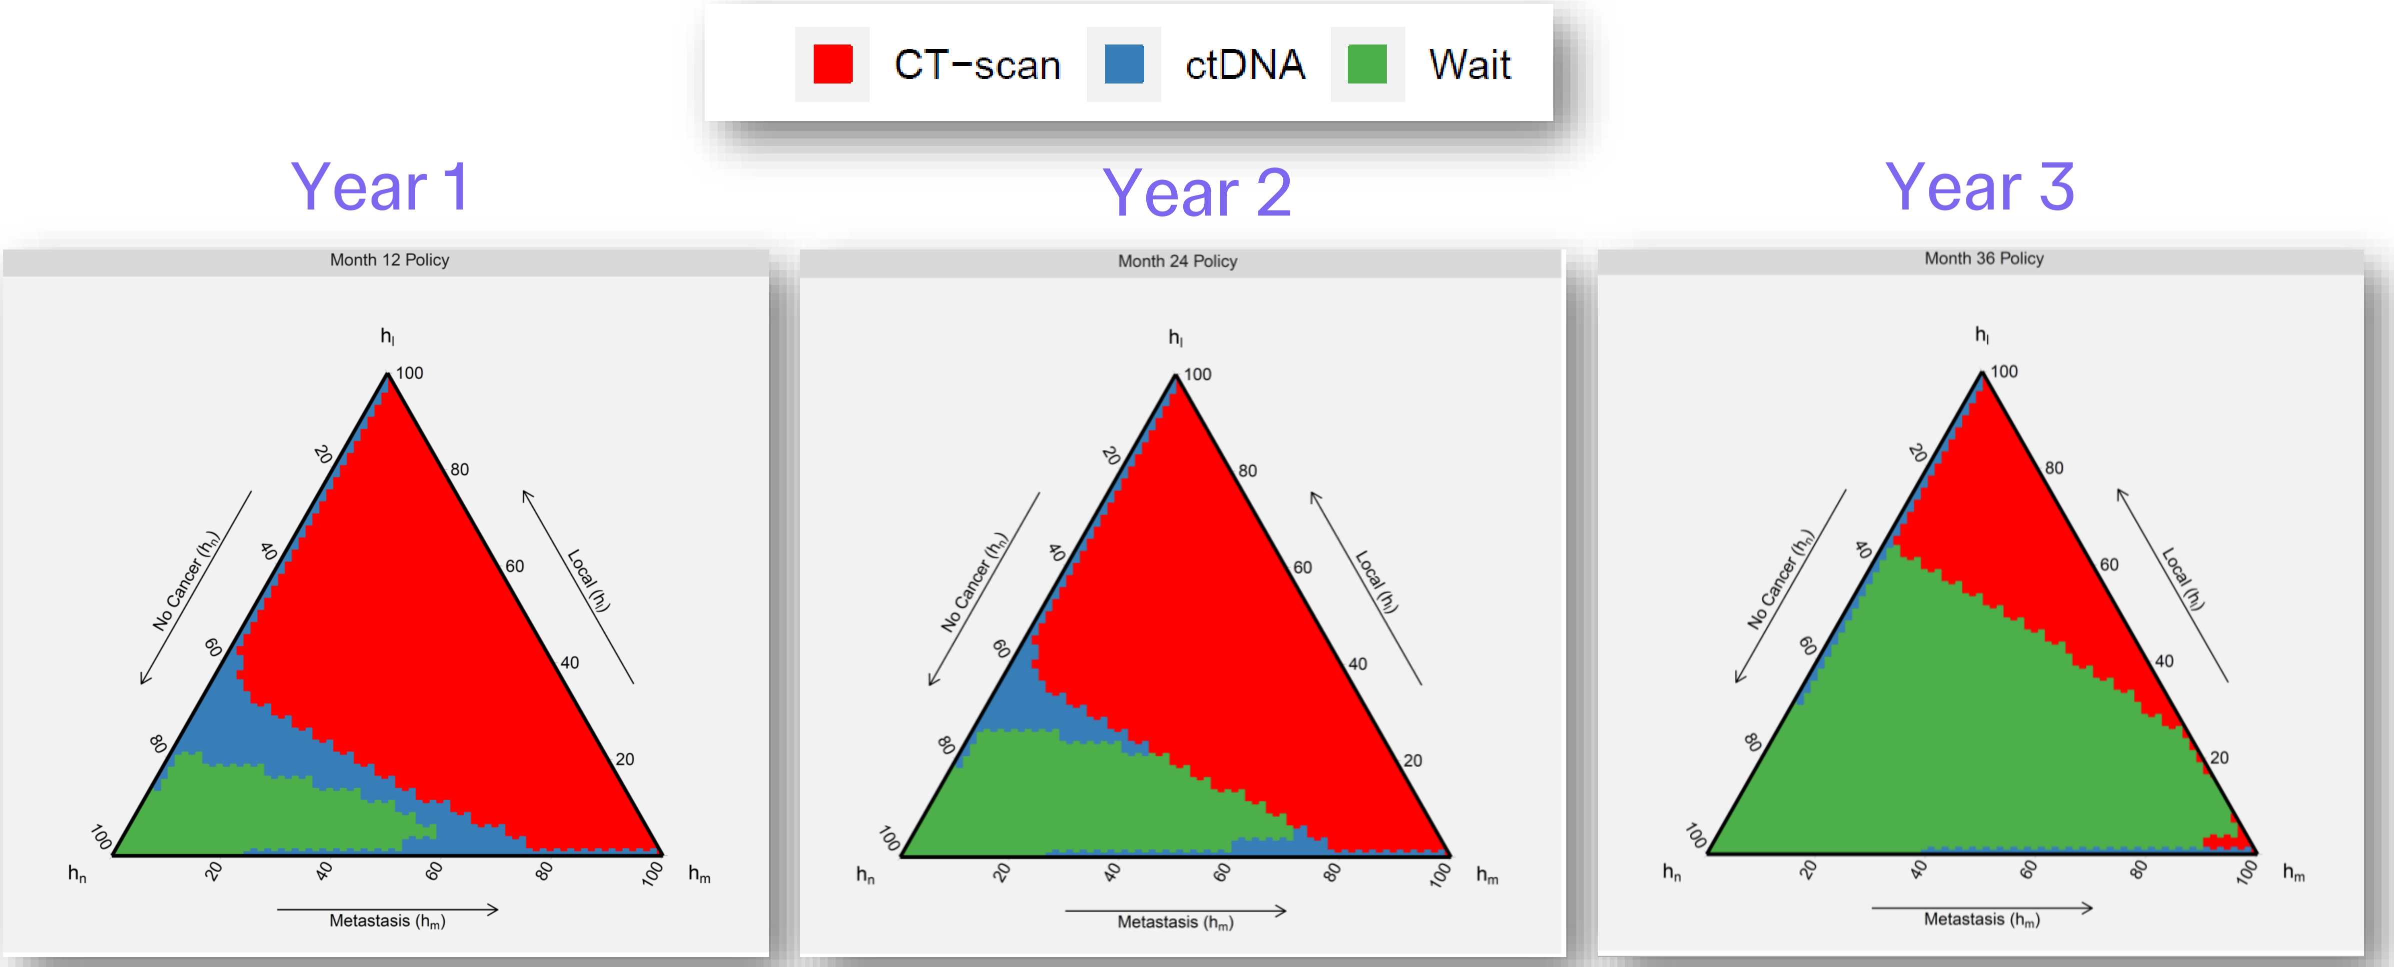

Personalized model
Policy in action

Observation
- Only DNA test until test becomes positive

Implications
  - POMDP policy can be converted to a DNA test schedule
  - Use DNA test to detect recurrence, use CT scan to pinpoint

Why?
  - Higher sensitivity/specificity
  - *Cost-effective*

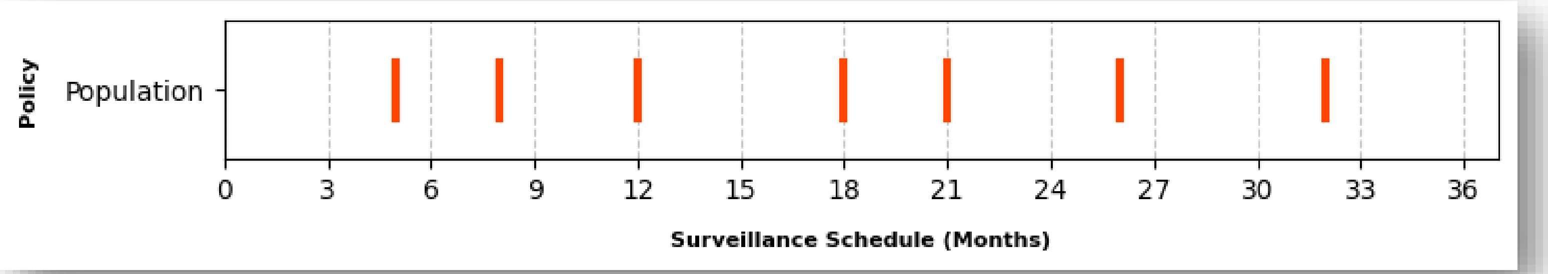

Personalization and trade-off

Optimization model
Personalized model

Age
Decades
<65 y
>65 y
Sex: female
ACE
Mild
Moderate
Severe
Moderate/severe
Anemia: yes
Smoking
Current
Former
p16: negative
cT
2
3
4
cN
2
3
log(MTV)
rECE: yes

6 Categorical variables:
- Sex: 2
- ACE: 5
- Anemia 2
- Smoking: 2
- Cancer stage: 12
  - T: 4
  - N: 3

1 Continuous variable:
- Age
- Categorize: <52 Y, ≥52 Y

Challenges:
- ≥ 960 total policies
- Hard to implement
- Diminishing return in additional stratification

Optimization model
Trade-off

Balancing objective:
- # of policies
- Inequality of patient outcomes
  - Detection delay target: 6 Mo
  - Loss: positive deviation from target
  - Weighted average positive deviation from target

Approach:
- Optimal sub-population aggregation
  - Optimal structured set partitioning
  - Interpretable

Trade-off model
Population partitioning

Unstructured partitioning
- Assignment
  - Patients to partitions

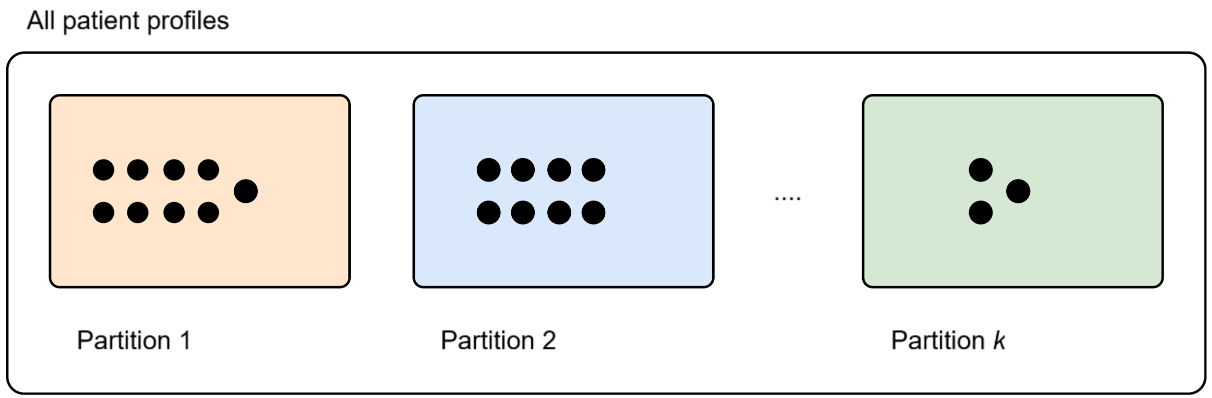

Structured partitioning:
- Based on covariates
- Follow a decision tree
- Interpretable/explainable

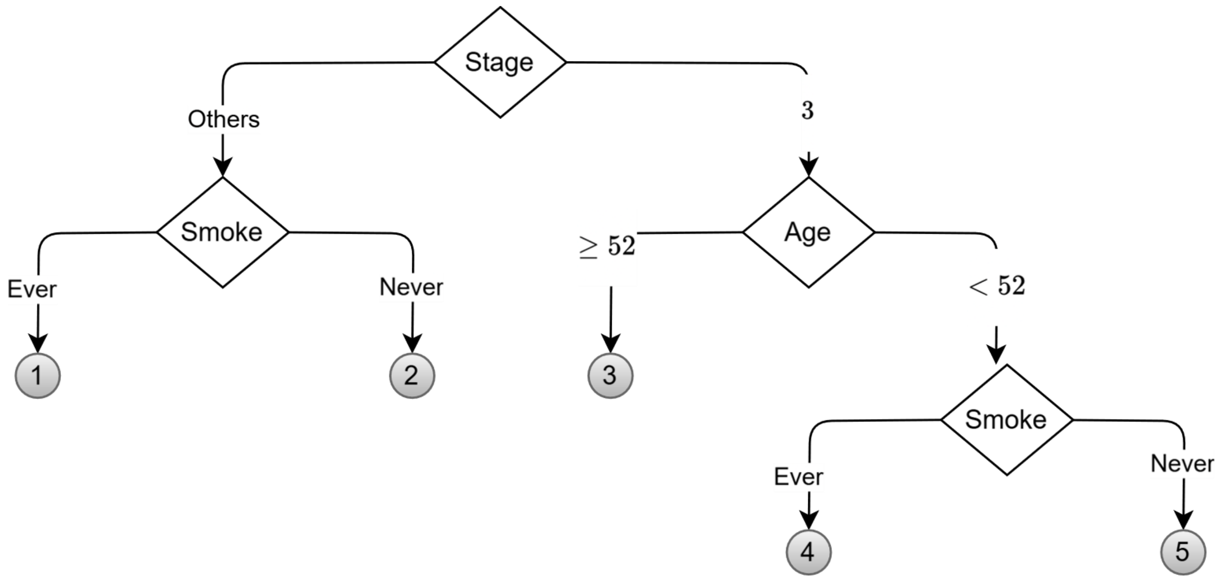

Trade-off model
Nested optimization

1. Choose partitions

2. For each partition: solve multi-model CPOMDP:

- Each patient profile, one model
  - Hence multi-model CPOMDP
- One policy across all models
  - Policy agnostic to stochastic process differences
- Minimize average partition cost
  - Constraint: *average within-partition* delay ≤ 6 Mo

Solution:

- Mixed integer LP

Trade-off model
Partitioning loss

Outcome
- Inequality of patient outcomes
  - Weighted average positive deviation
  - Patient profile weight in the population

Mathematical model:

$$\min_{\mathcal{Q},\pi^*_{\mathcal{P}}} \sum_{\mathcal{P}\in\mathcal{Q}} \sum_{x\in\mathcal{P}} w_x \,.\left(d_x^{\pi^*_{\mathcal{P}}} - \bar{d}\right)^+$$

Trade-off model
Within-partition optimization

Multi-model CPOMDP

Multi-model optimization
- Uncertain model parameters
  - Limited data, system changes
- Heterogenous dynamics/population

Mathematical modeling:
- Two stage stochastic optimization
  - Stage 1: fix action
  - Stage 2: true model unravels
- Robust and non-robust objectives

Trade-off model
Within-partition optimization

Multi-model MDP
- Admits deterministic policies
- Mixed integer linear programming
  - Stage 1: decide actions per state/time
  - Stage 2: determine value functions

Multi-model POMDPs
- Uncountable belief states

Approximation methods
- Create a *population* core state transition matrices
  - Weighted average

Bayes-adaptive approach
- Distribution weight in period 1
- Occupancy informed Bayes-adaptive weights thenceforth
- Reproduce the same occupancy measures in expectation

Tractability and Approximation Methods

Trade-off model
Problem complexity

Problem is intractable

Unstructured partitioning
- NP-complete
  - Many cases NP-hard.
- Require exponential # of costs
  - For all partitioning configurations

Partitioning costs
- Multi-model MDP is NP-Hard
  - Steimle et al. 2021
- Our case: constrained POMDP

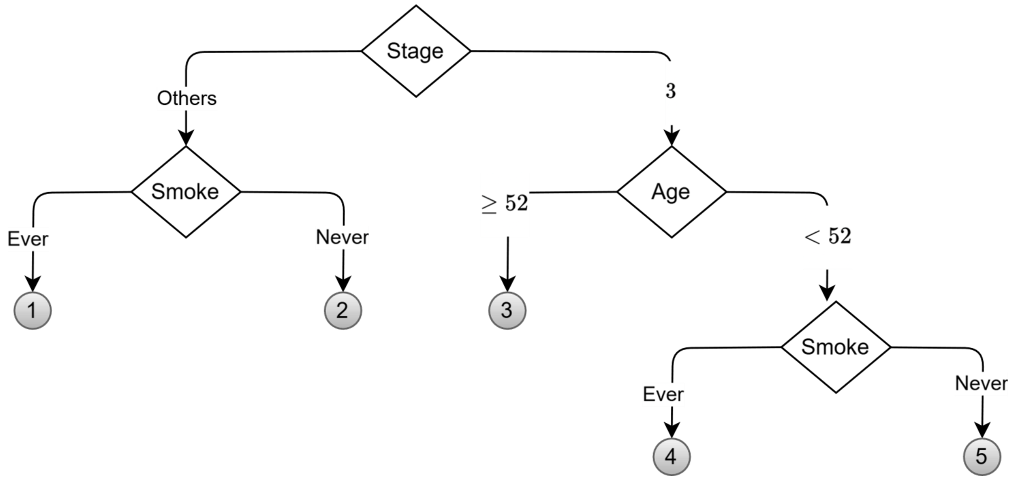

Structured partitioning
- Partitioning follows decision tree

Approximation methods

Solution approach: clustering

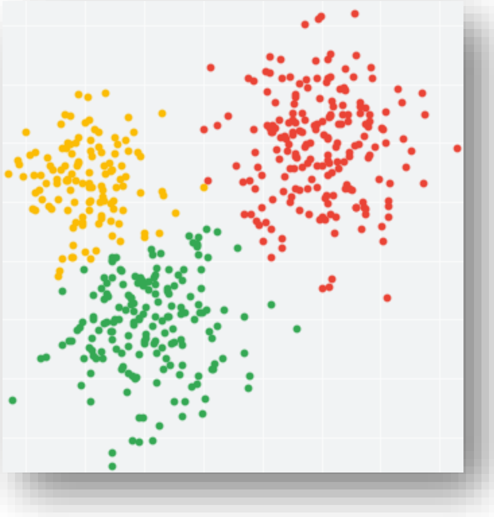

Interpretable clustering

- Partition to k clusters minimizing *average dissimilarity*

Clustering idea

- Aggregate based on *similarity* in stochastic process
  - Hazard ratio vector
  - Average core-sate transition probabilities
  - Bayesian core-state transition probabilities

Clustering loss:

  - Distance from cluster's focal point

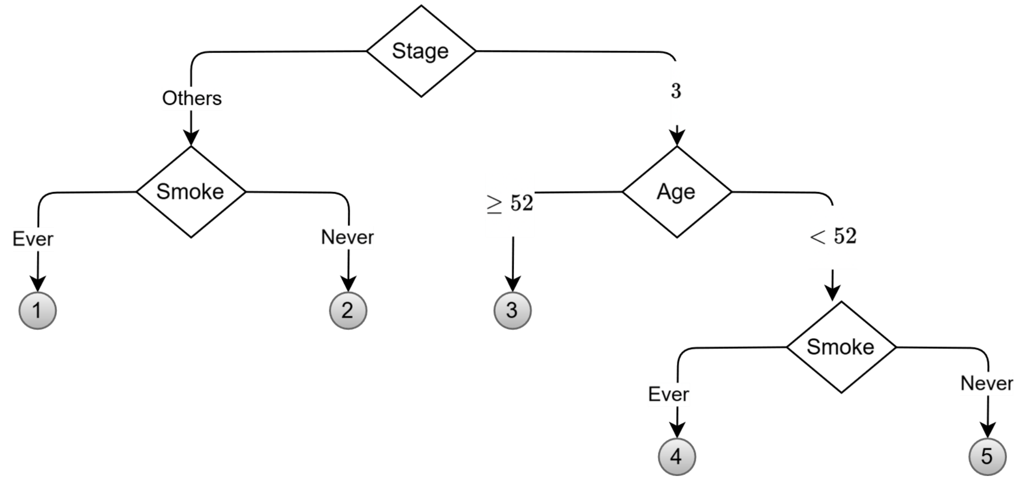

Approximation methods
Solution approach: partitioning

Partitioning idea

- Aggregate based on *consequence* of aggregation of profiles
- Loss: deviation from target when merged into a partition

Inputs:

- Profile covariate matrix, binary
- Profile weights in the population
- *Pairwise* distance matrix

Objective:

- Partition to k groups minimizing *inequality proxy*

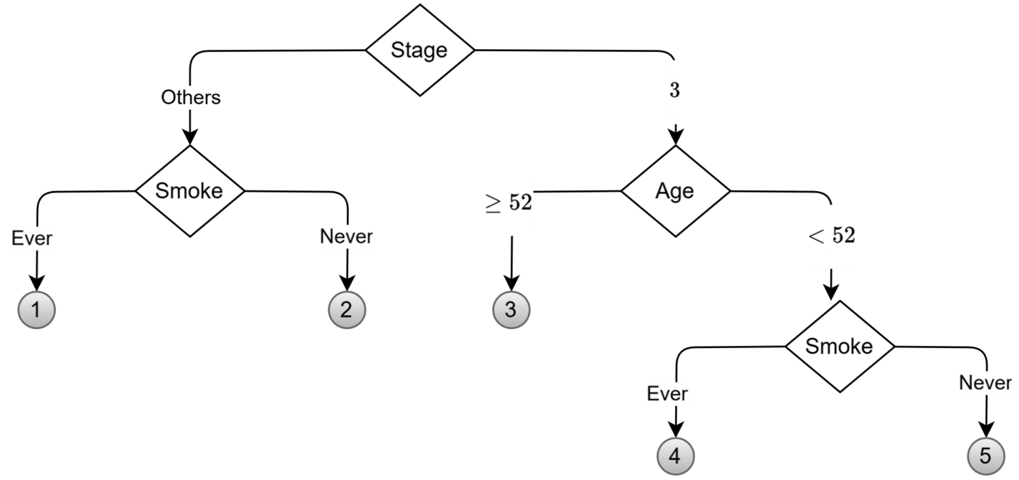

Partitioning idea
Distance matrix

Pairwise distance matrix:
- Create a 2-member partition with profiles i, j
- Solve a bi-model CPOMDP
- Evaluate optimal policy
  - Calc detection delay d_i, d_j

Cost of aggregating i, j:
- $d_{ij} = \left(d_i - \bar{d}\right)^+$
- $d_{ji} = \left(d_j - \bar{d}\right)^+$

Partitioning idea
Inequality proxy

Objective:
- Minimize weighted average of inequalities
- Across all patients

$$\min \sum_i w_i \hat{\delta}_i$$

Per partition:
- Weighted average of distance
  - From partition members
- Partition weights normalized to 1: w_j'
- Inequality proxy:

$$\hat{\delta}_i = \sum_{j \in \mathcal{P}} w_j' d_{ij}$$

Approximation methods
Optimization approach

Model:
- Adaptation of Bertsimas and Dunn (2017)
- Optimal binary tree
  - Mixed integer linear programming model
  - Local search: non-linear objectives

Model decides:
- Tree configuration
  - Branching
  - Each leaf is a partition/cluster
- Constrain or penalize # of leaves

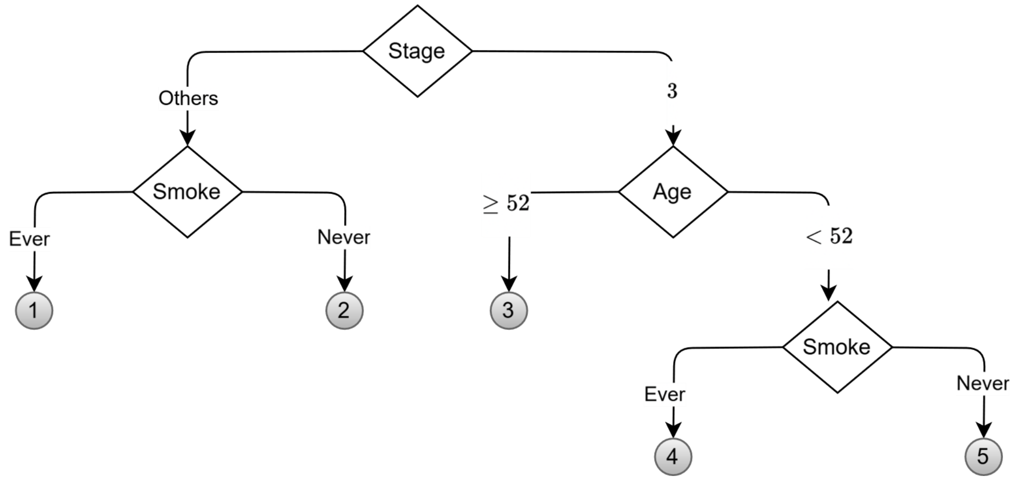

Partitioning idea
Policies

Robust objectives
- Max of max-deviations (worst case)
- Simple average of max-deviations
- Weighted-average of max-deviations

Non-robust objectives
- Average of Average deviations: simple average
- Weighted average of weighted average deviations

Clustering idea
Similarity targets

Similarity in stochastic process

- Hazard ratio vector
- Average core-sate transition probabilities
- Bayesian core-state transition probabilities

Similarity in optimal policy

- WIP

Optimal Aggregation Results

Results
Inequality: clustering models

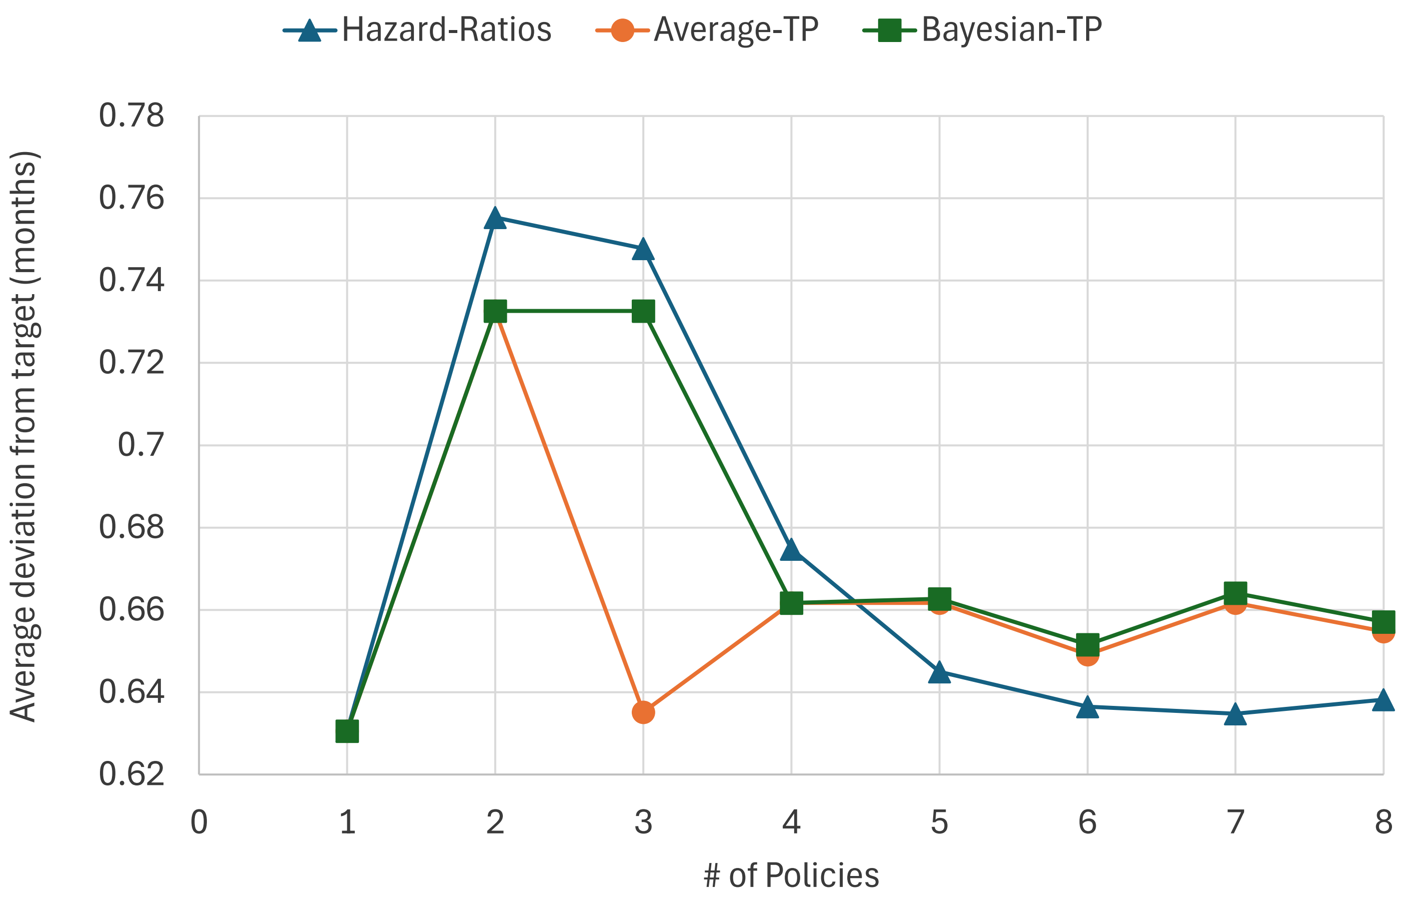

Results

Inequality: partitioning models

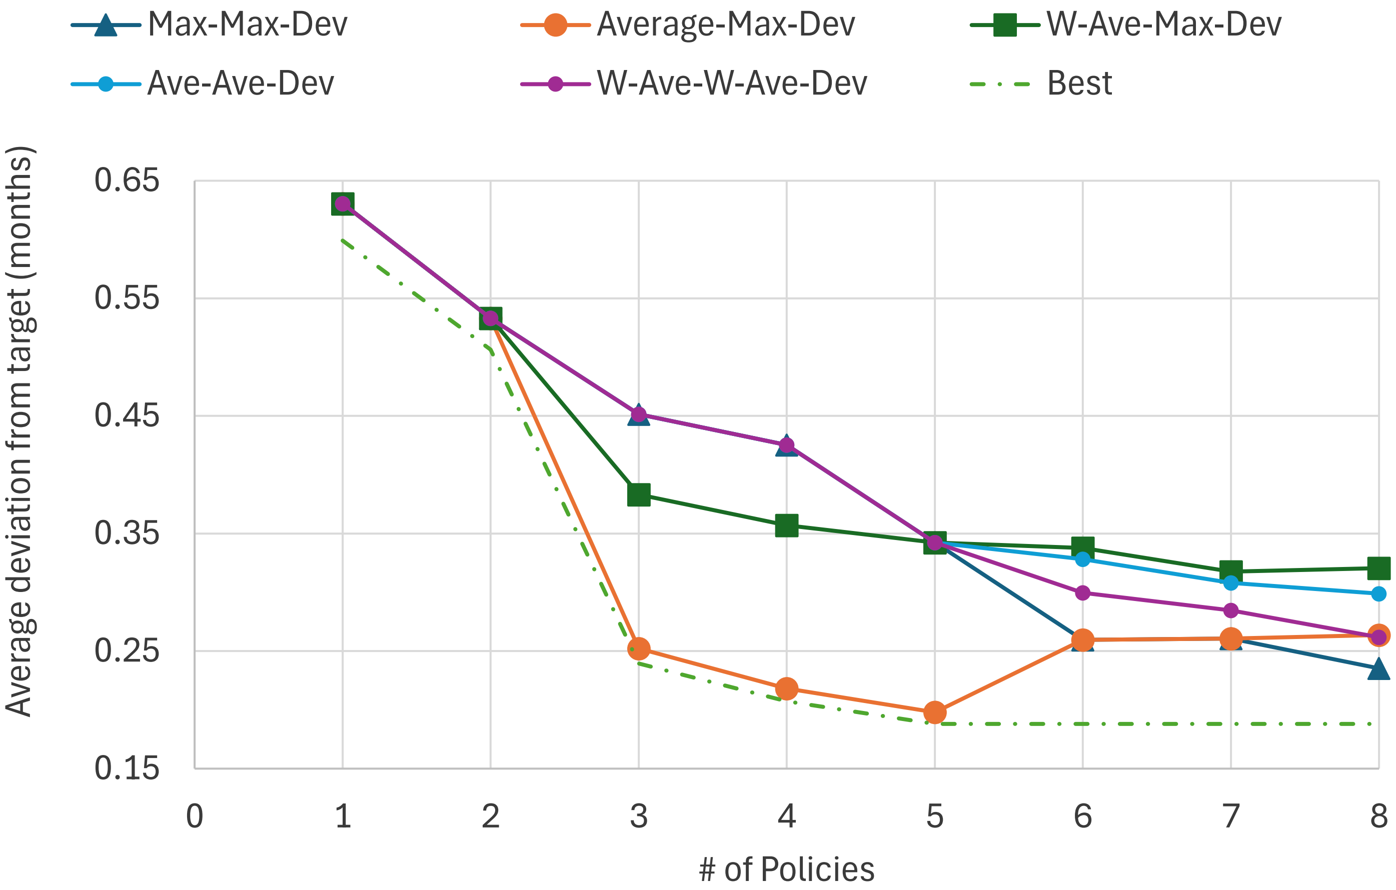

Results
Relative optimality gap

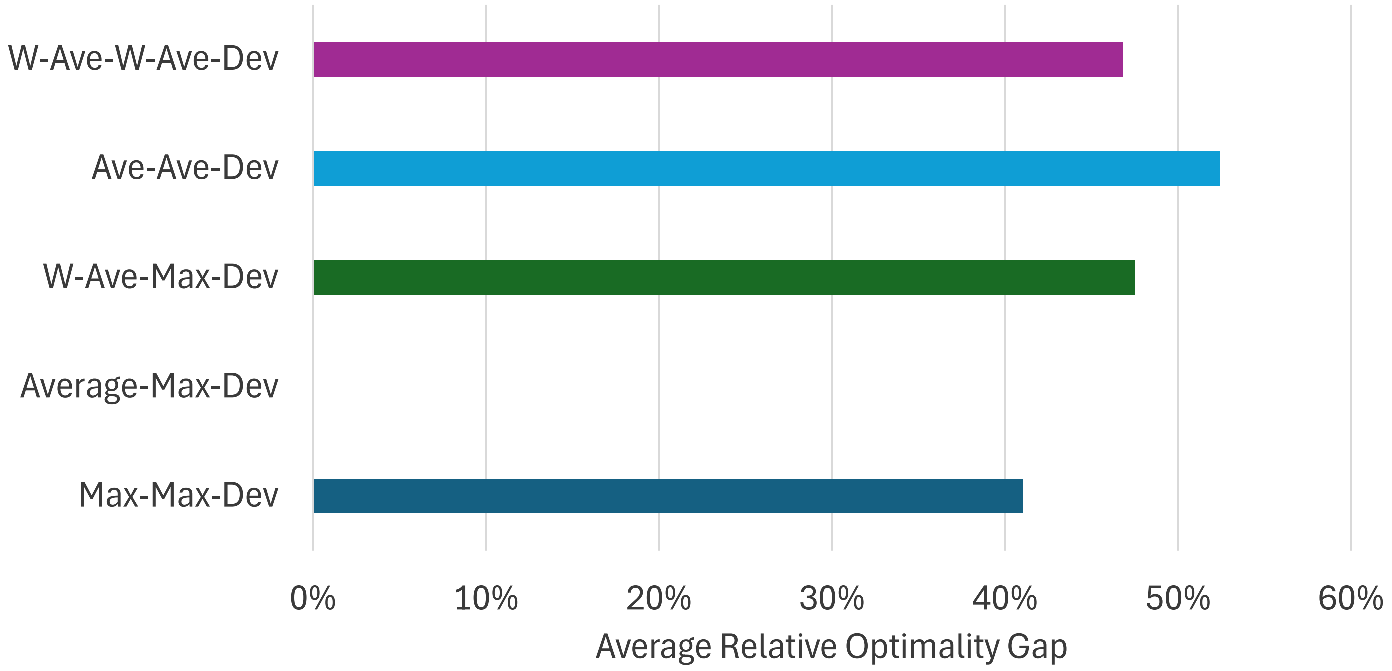

Results
Optimal partitions, k=2 to 4

k=2

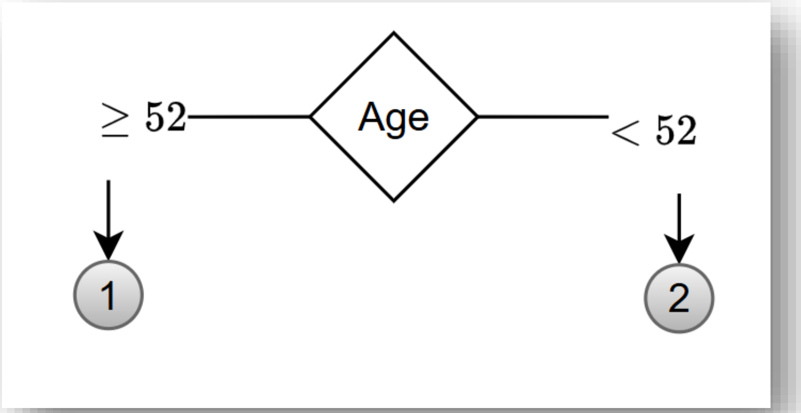

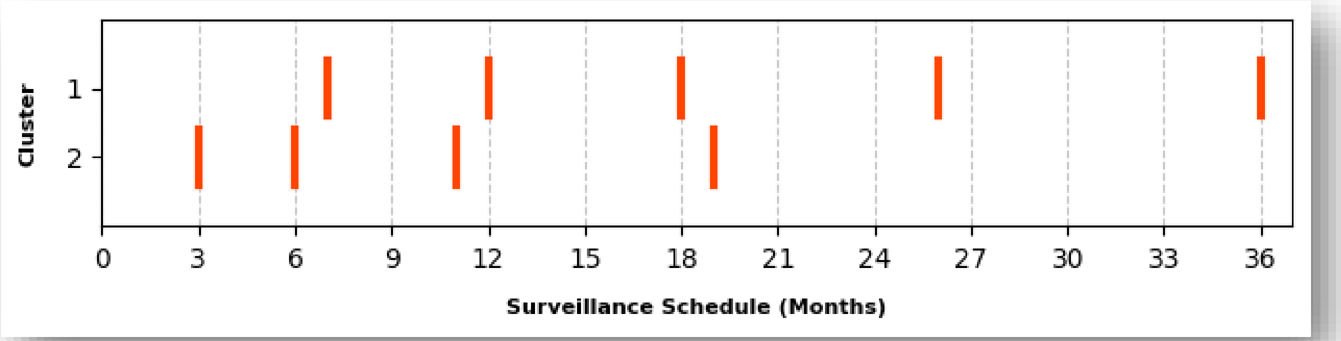

Results

Optimal partitions, k=2 to 4

k=3

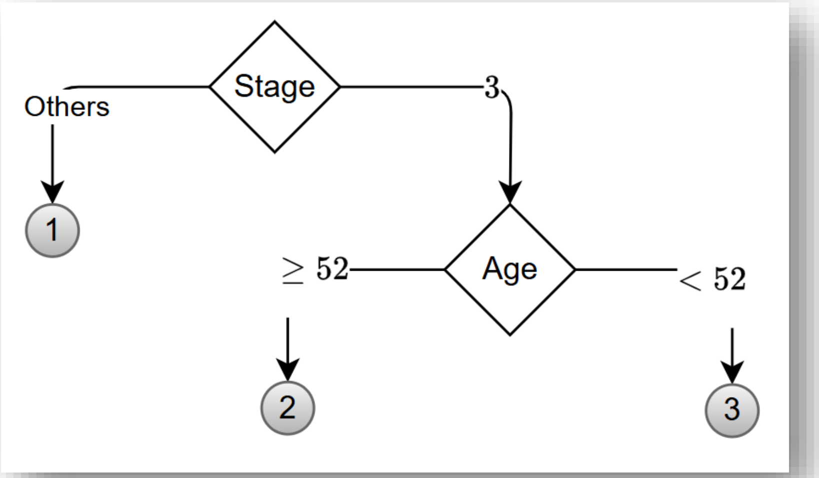

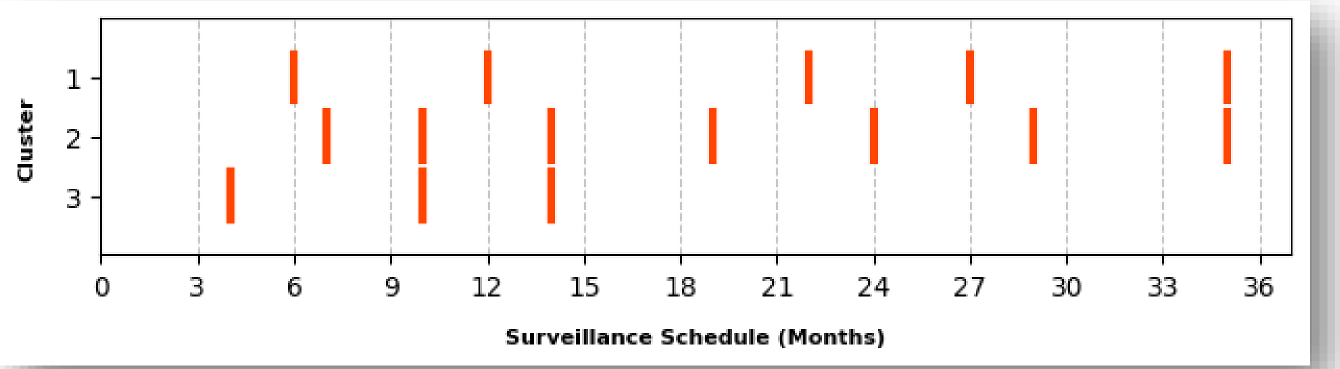

Results

Optimal partitions, k=2 to 4

k=4

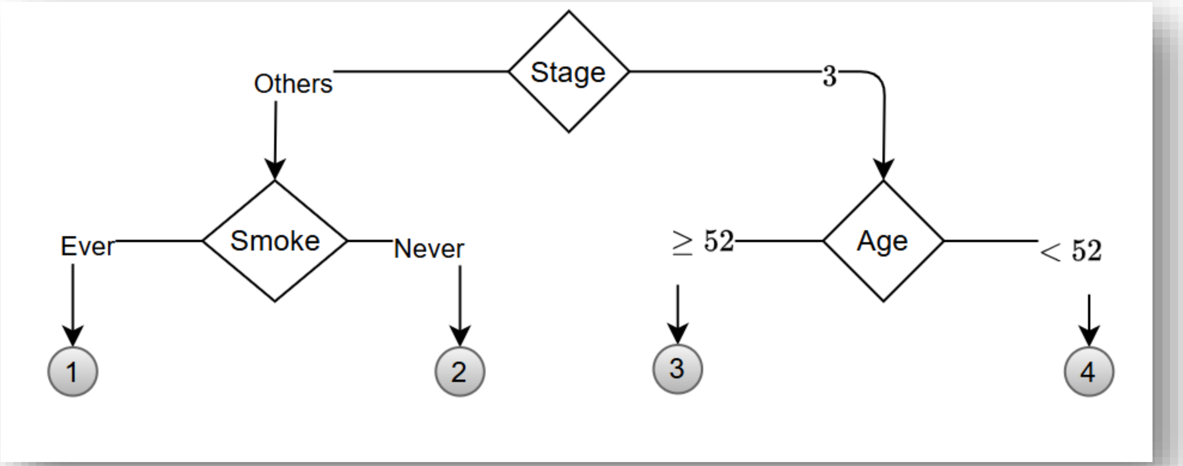

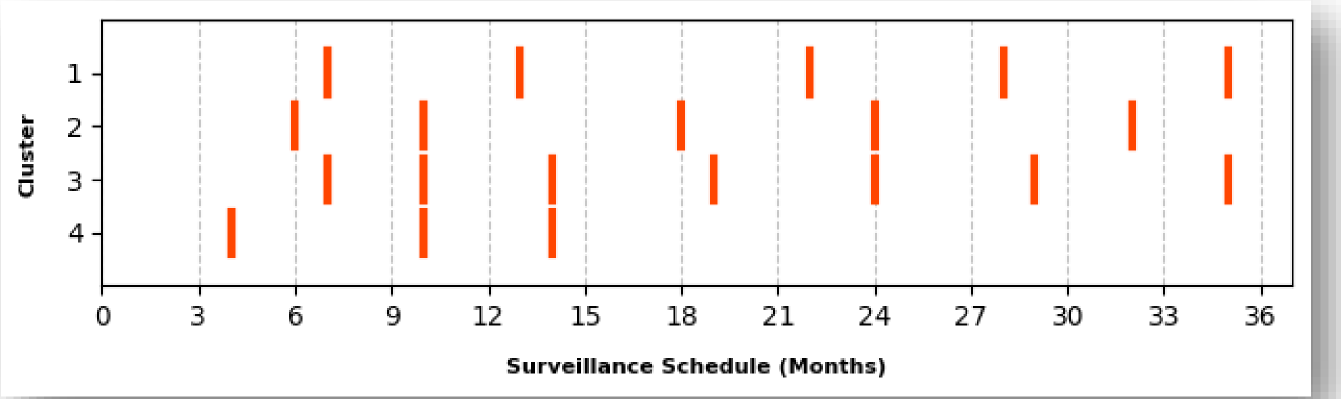

Limitations

Computations:
- Time-consuming for large tree depth

Approximation:
- Proxy vs real objective
- CPOMDP
- Multi-model CPOMDP

Data and evidence:
- Disease model
  - More stratification, less data, less confidence
- Confidence in test performance

Conclusion

Challenge:
- Trade-off
  - personalized care vs one-size-fits-all
  - Complexity vs fairness

Solution:
  - Strategic population stratification/aggregation
  - Data-driven and systematic
  - Interpretable
  - optimal tree-based partitioning

Case study:
- post-cancer surveillance
- generalizable beyond the case study

Takeaway

Smart design of population aggregation/stratification enhances healthcare delivery by balancing personalization, interpretability, and equality.


Thank you for your
Attention!