Optimal Control for Tiered Remote Patient Monitoring

Siddharth Chandak

Joint work with Isha Thapa, Nicholas Bambos, David Scheinker

Outline

- Background and Motivation
- Tiered Service Architecture
- Optimal Policy
- Extensions

Remote Patient Monitoring

- Technology enabled healthcare

Continuous Glucose Monitors

Smartwatches

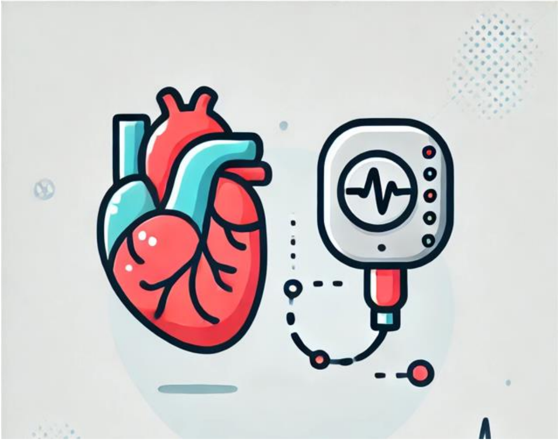

Implantable Devices

Example: RPM for Type 1 Diabetes

- Autoimmune disease requiring lifelong management
- Chronic disease management is complex
  - Cannot be automated
- Management using RPM:
  - Continuous Glucose Monitors
  - Periodic review of glucose data
- Communication with patients:
  - Text notifications
  - Virtual appointments
  - Phone calls
  - And more...

Example

Ordinary Monitoring

Intensive Monitoring

The Tradeoffs

 $

Ordinary Monitoring

 $$$

Intensive Monitoring

The Tradeoffs

Remote Monitoring

$

$$$

Critical State

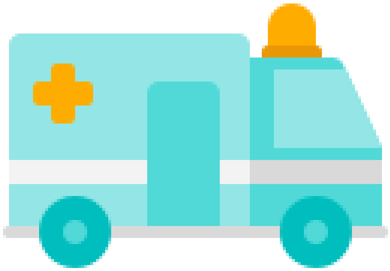

$$$$$$$$$

What is the optimal level of monitoring?

The Model

Controlled Markov Chain

The Model

Ordinary Monitoring

Intensive Monitoring

Time: $t \in \{0, 1, 2, \dots\}$

Patient State:

- Monitoring State: $m_t \in M = \{o, i\}$

The Model

Time: $t \in \{0, 1, 2, \ldots\}$

Patient State:

- Monitoring State: $m_t \in M = \{o, i\}$
- Health State: $h_t \in \{0, 1, 2, \ldots H\}$

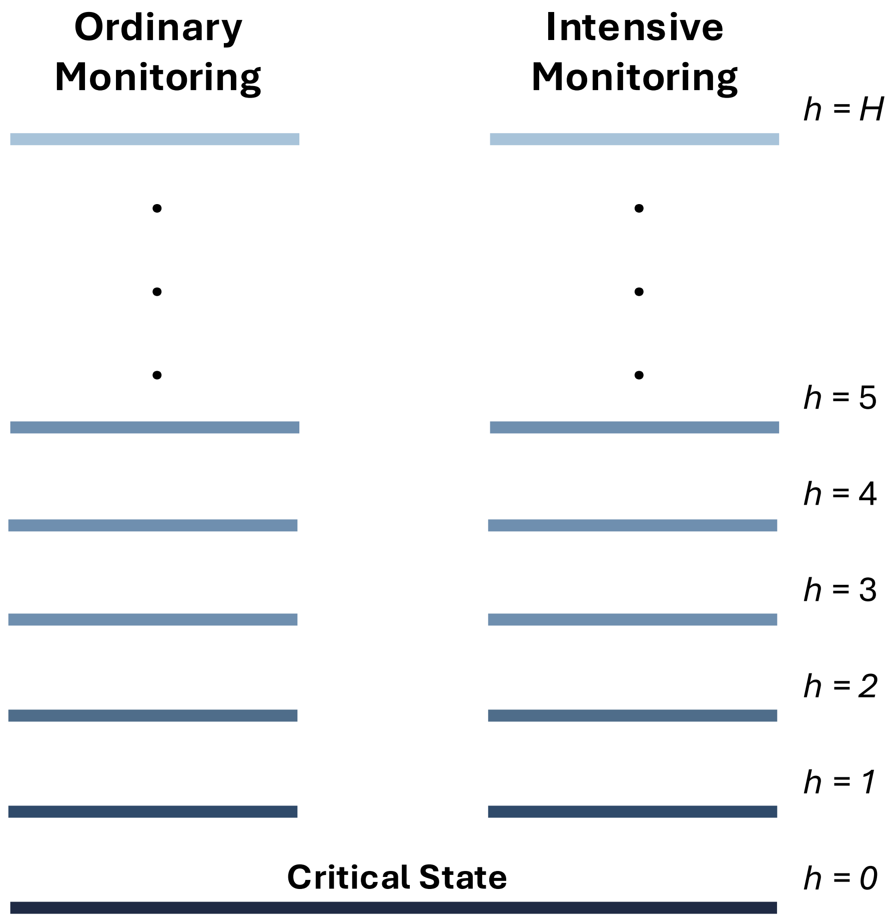

The Model

Time: $t \in \{0, 1, 2, \dots\}$

Patient State:

- Monitoring State: $m_t \in M = \{o, i\}$
- Health State: $h_t \in \{0, 1, 2, \dots H\}$
- Overall State: $s_t \coloneqq (m_t,\, h_t)$

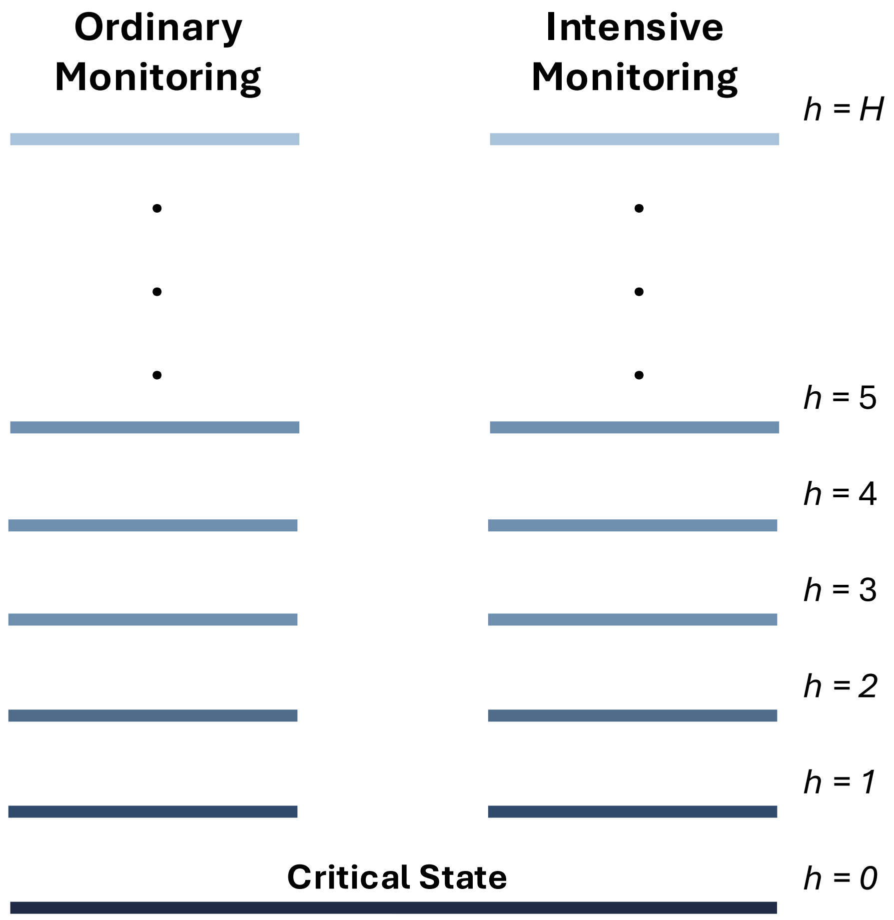

The Model

Time: $t \in \{0, 1, 2, \dots\}$

Patient State:

- Monitoring State: $m_t \in M = \{o, i\}$
- Health State: $h_t \in \{0, 1, 2, \dots H\}$
- Overall State: $s_t \coloneqq (m_t, h_t)$

Action: $a_t \in A = \{o, i\}$

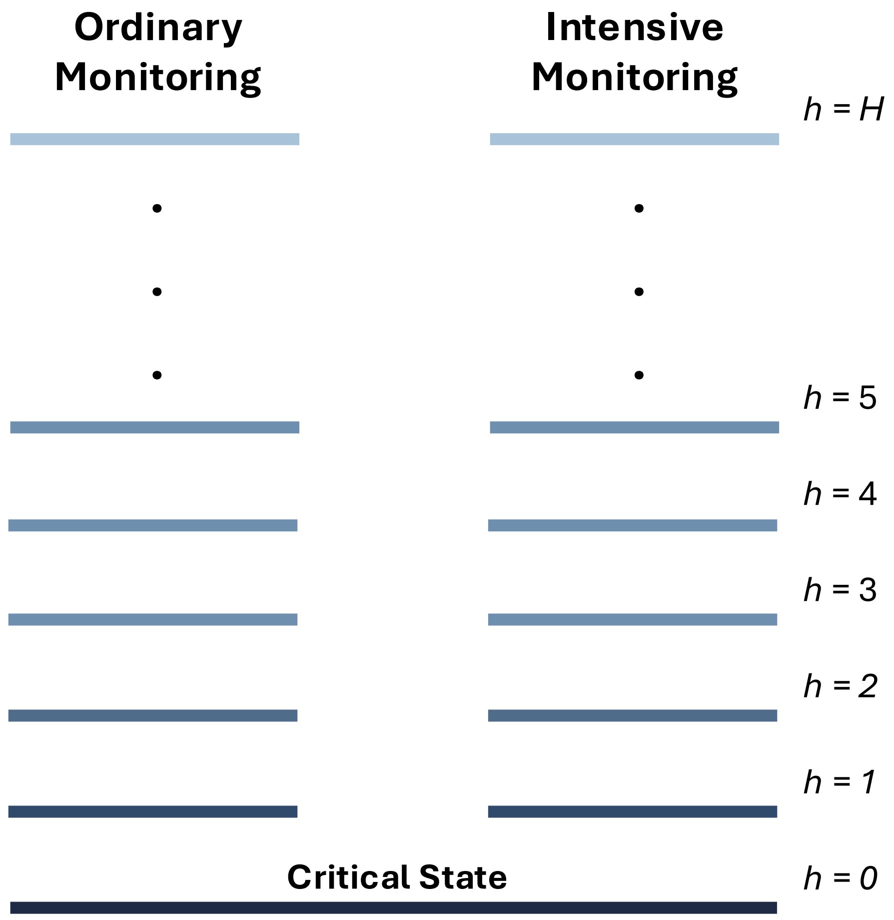

The Model

Transition Probabilities

- **Ordinary monitoring:** $\lambda_o,\ \mu_o = 1 - \lambda_o$

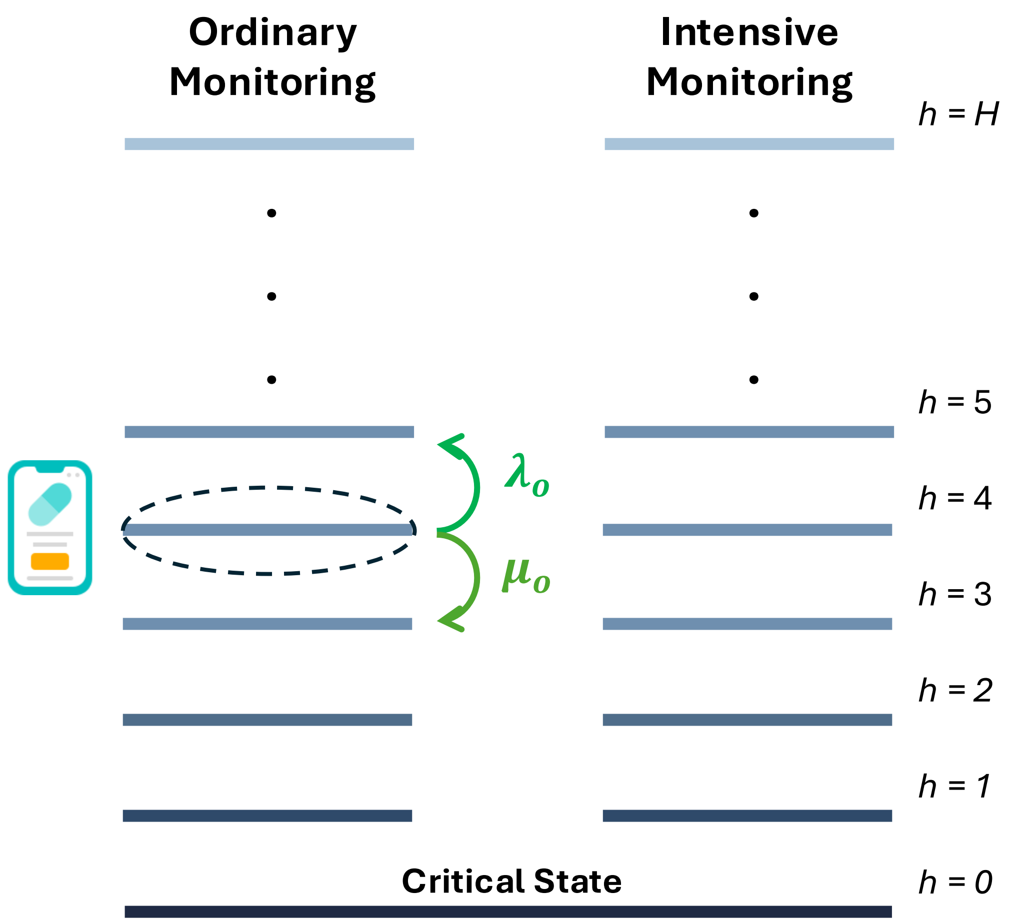

The Model

Transition Probabilities

- Ordinary monitoring: $\lambda_o,\ \mu_o = 1\ - \lambda_o$
- **Intensive monitoring: $\boldsymbol{\lambda_i,\ \ \mu_i = 1\ - \lambda_i}$**

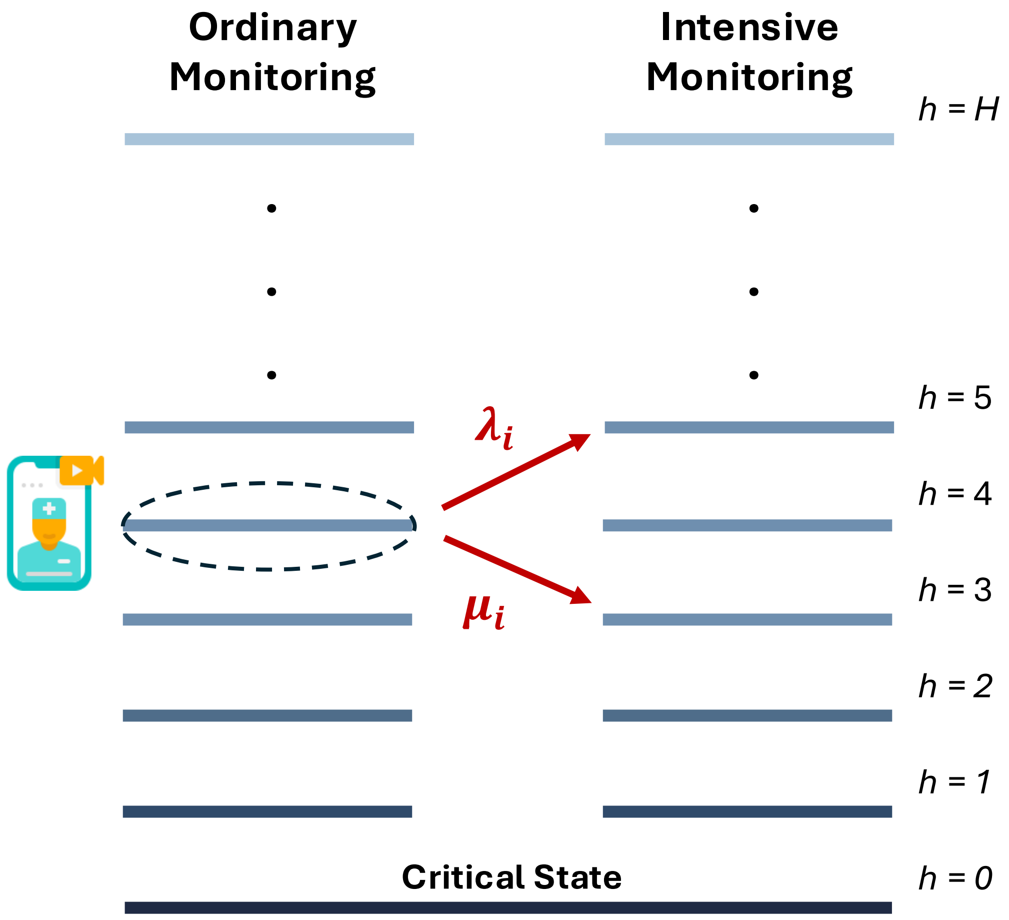

The Model

Transition Probabilities

- Ordinary monitoring: $\lambda_o,\ \mu_o = 1 - \lambda_o$
- **Intensive monitoring: $\lambda_i,\ \mu_i = 1 - \lambda_i$**

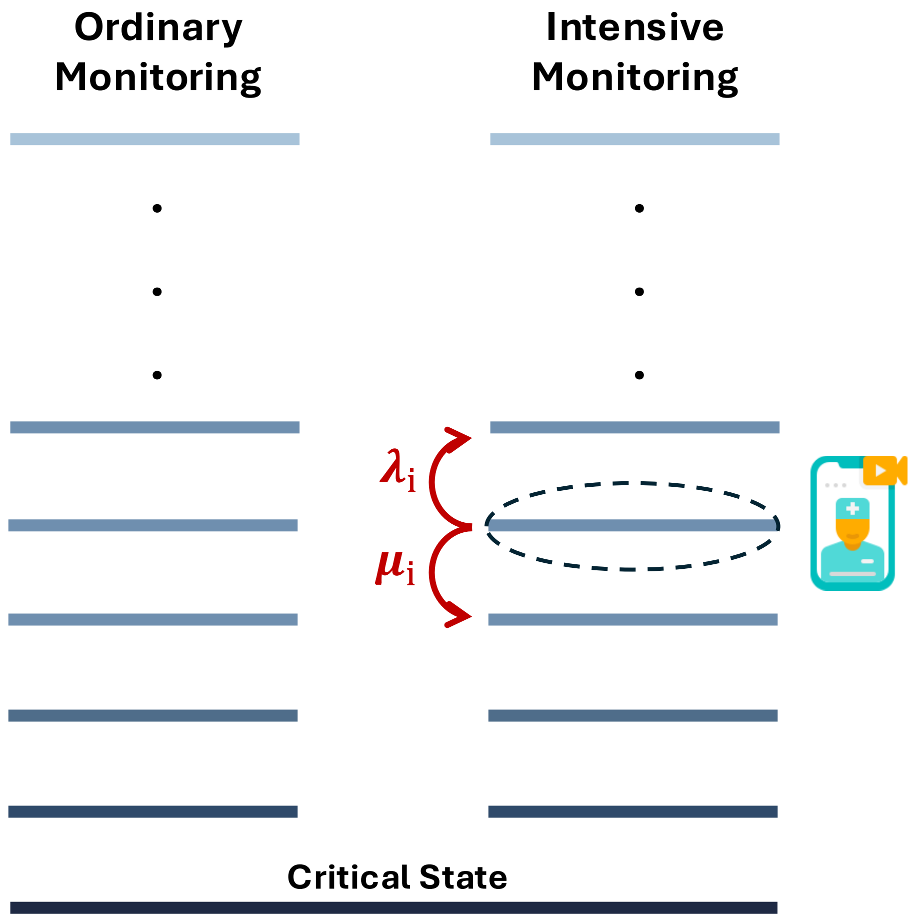

The Model

Transition Probabilities

- **Ordinary monitoring: $\boldsymbol{\lambda_o,\ \mu_o = 1 - \lambda_o}$**
- Intensive monitoring: $\lambda_i,\ \mu_i = 1 - \lambda_i$

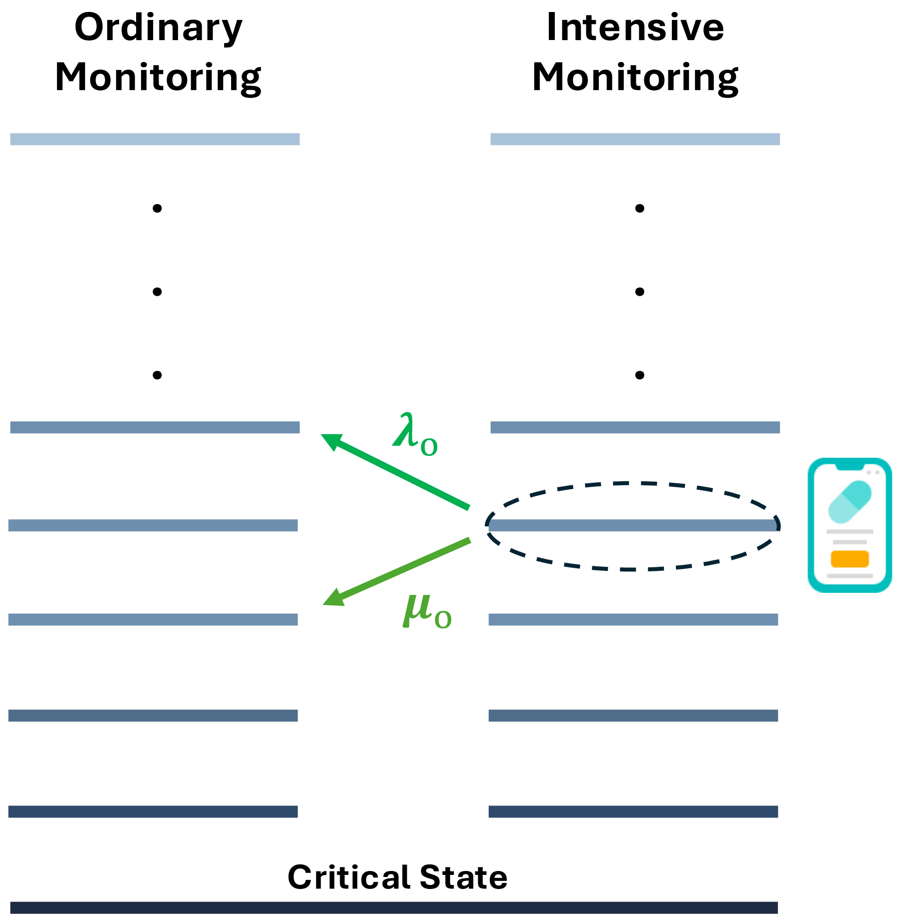

The Model

Transition Probabilities

- Ordinary monitoring: $\lambda_o, \ \mu_o = 1 - \lambda_o$
- Intensive monitoring: $\lambda_i, \ \mu_i = 1 - \lambda_i$
- Satisfies $\lambda_i \geq \lambda_o$

Costs

- Ordinary monitoring: C_o
- Intensive monitoring: C_i
- Critical State: C_c
- Satisfies $0 \leq C_o \leq C_i << C_c$

Optimal Control

We seek a policy π (mapping from states to actions) that minimizes the expected discounted cost

$$V_\pi(s) = \mathrm{E}\left[\sum_{t=0}^{T-1} \gamma^t c(s_t, \pi(s_t)) + \gamma^T C_c \mid s_0 = s\right]$$

T: hitting time for the critical state (random variable)

$\gamma \in (0,1)$: discount factor

Optimal Control

$$V_\pi(s) = \mathrm{E}\left[\sum_{t=0}^{T-1} \gamma^t c(s_t, \pi(s_t)) + \gamma^T C_c \mid s_0 = s\right]$$

Cost incurred at
time t by policy π

Optimal Control

$$V_\pi(s) = \mathrm{E}\left[\sum_{t=0}^{T-1} \gamma^t c(s_t, \pi(s_t)) + \gamma^T C_c \mid s_0 = s\right]$$

Time
discounting

Optimal Control

$$V_\pi(s) = \mathrm{E}\left[\sum_{t=0}^{T-1} \gamma^t c(s_t, \pi(s_t)) + \gamma^T C_c \mid s_0 = s\right]$$

Discounted cost of entering critical state

Computing the Optimal Policy

Dynamic programming can be used to numerically calculate the optimal policy for a specific set of parameters

Numerically, we observe the optimal policy exhibits a threshold structure

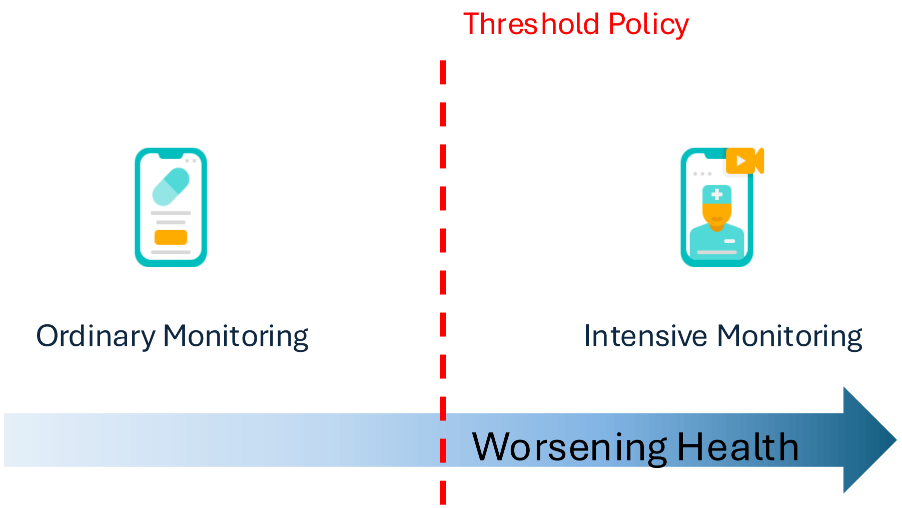

Structural Characterization of Optimal Policy

- **Threshold-based policy is optimal**
  - When $h > \bar{h}$:
    - Ordinary Monitoring
  - When $h \leq \bar{h}$:
    - Intensive Monitoring
- Characterized using:
  - Moment Generating Function (MGF) of the hitting time of critical health state

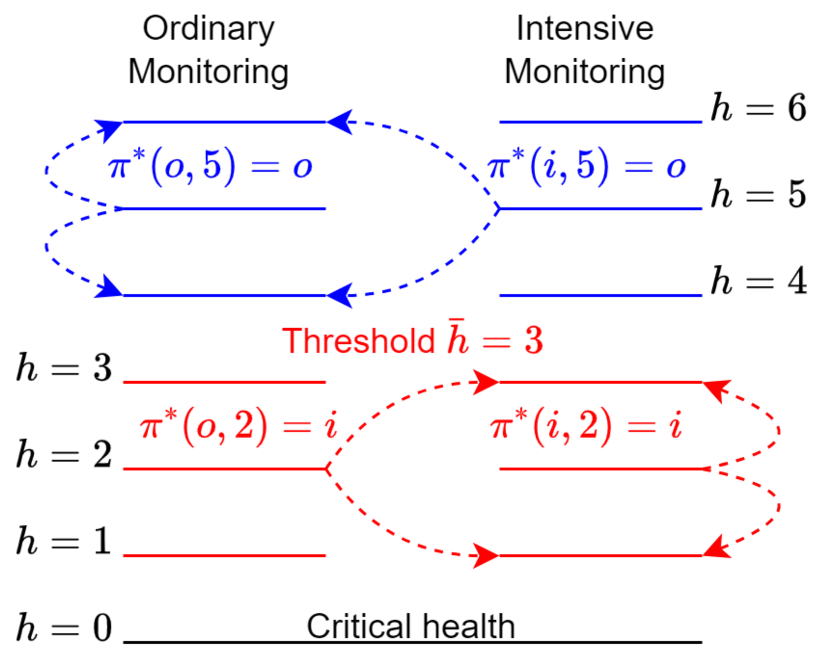

Multidimensional Health States

In practice patient health may have multiple dimensions

Multidimensional Health States

In practice patient health may have multiple dimensions

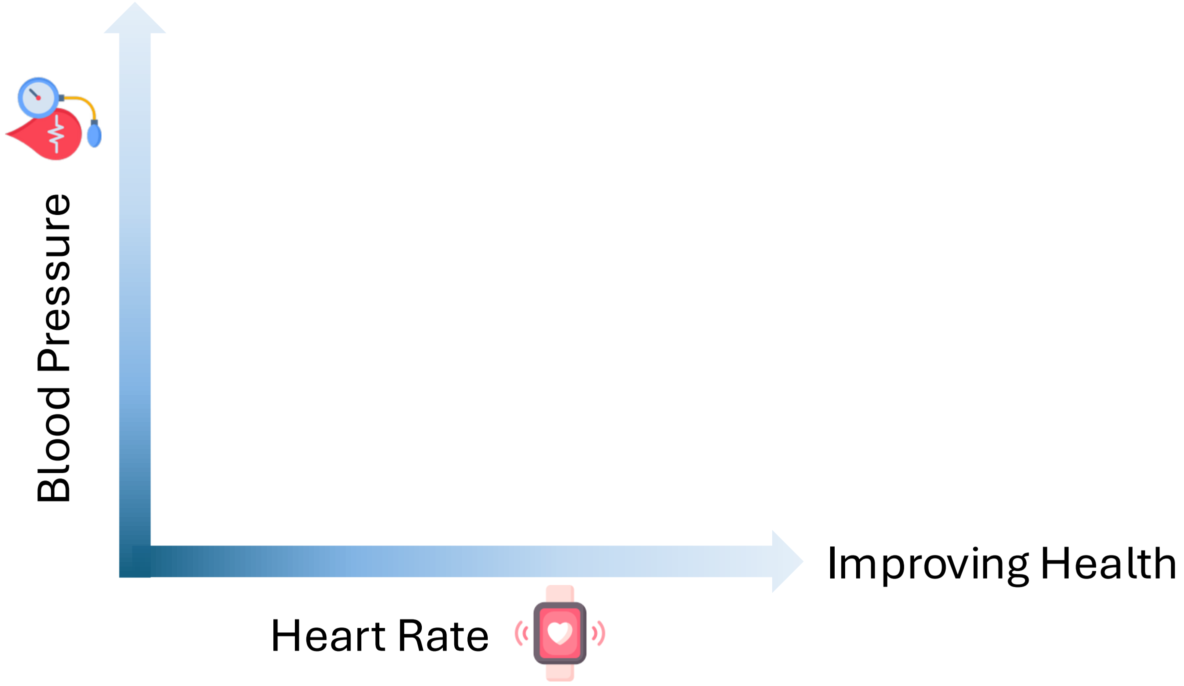

Multidimensional Health States

We prove that the optimal policy retains threshold structure

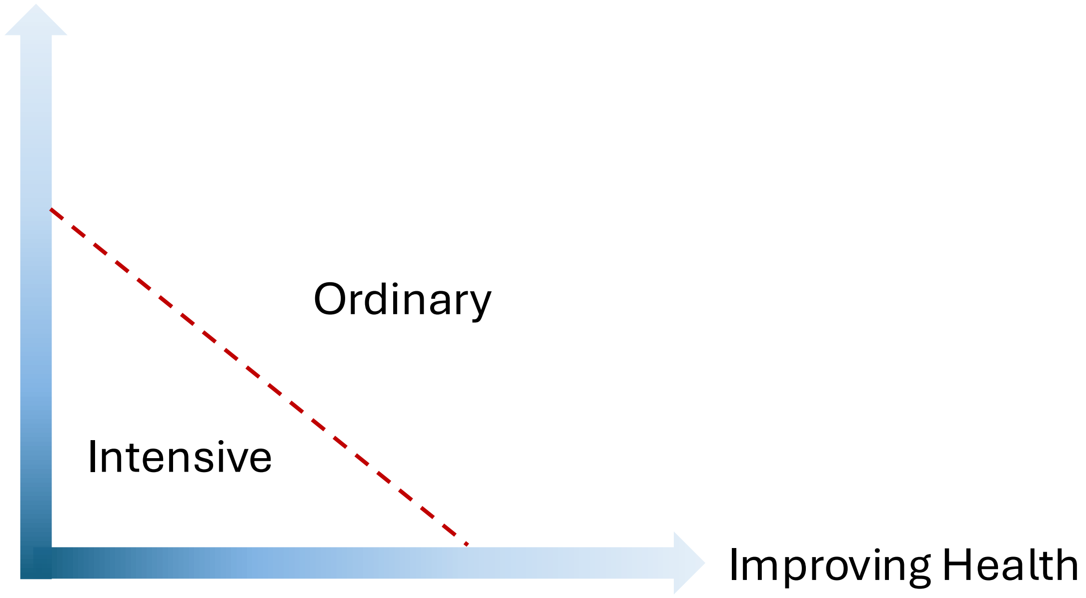

Multidimensional Health States

Numerical Example

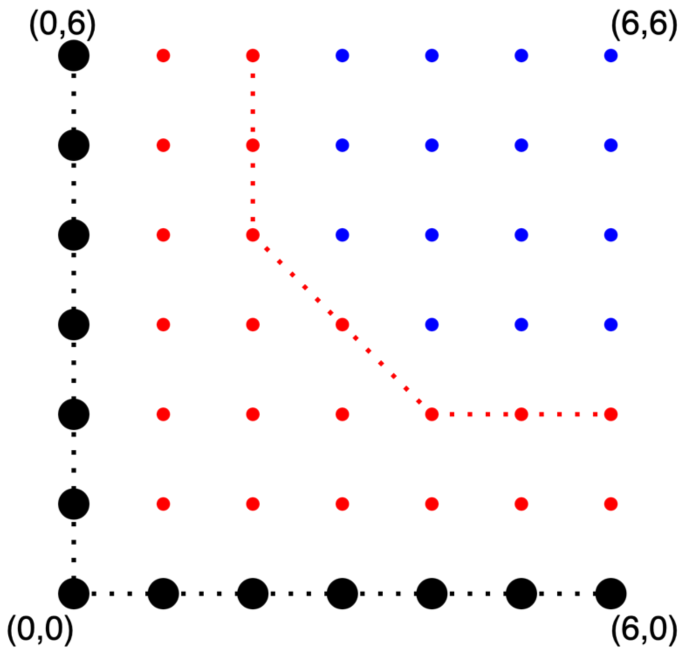

Multidimensional Health States

This framework allows for different definitions of the critical state

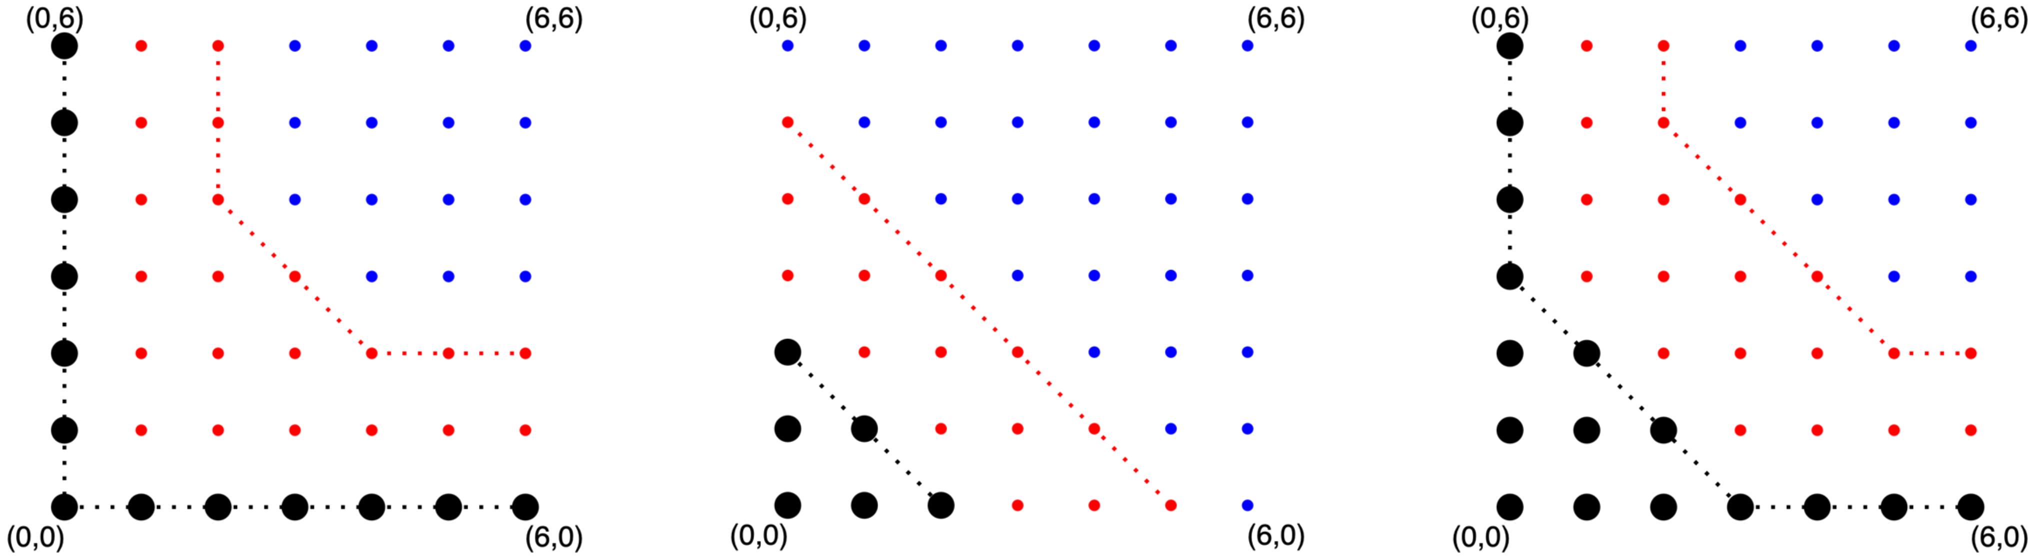

Key Takeaways

- Modeled patient monitoring as a dynamic decision problem
- Developed a theoretically grounded service architecture for remote monitoring

Provides simple, interpretable policies to support clinical decision making

Thank You!