Dynamic Measurement Scheduling for Event Forecasting Using Deep RL

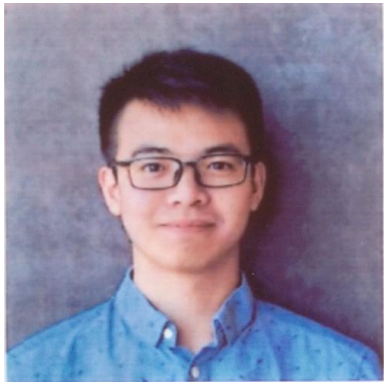

Chun-Hao Chang (Kingsley)

Mingjie Mai

Anna Goldenberg

University of Toronto, Vector Institute
The Hospital of Sickkids

Motivation

Motivation

Outcomes

Sepsis

Mortality

Treatments

⋮

Motivation

Outcomes	**Measurements**
Sepsis	Lactate
Mortality	Blood Test
Treatments	O_2
⋮	⋮

Uniform policy

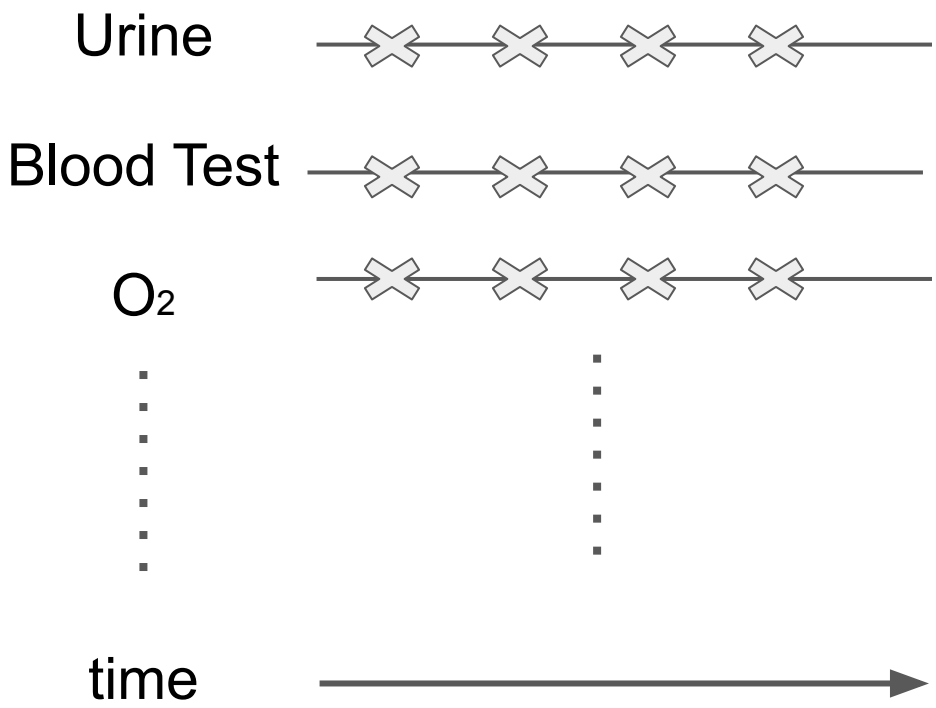

Waste lots of measurements

Dynamic policy

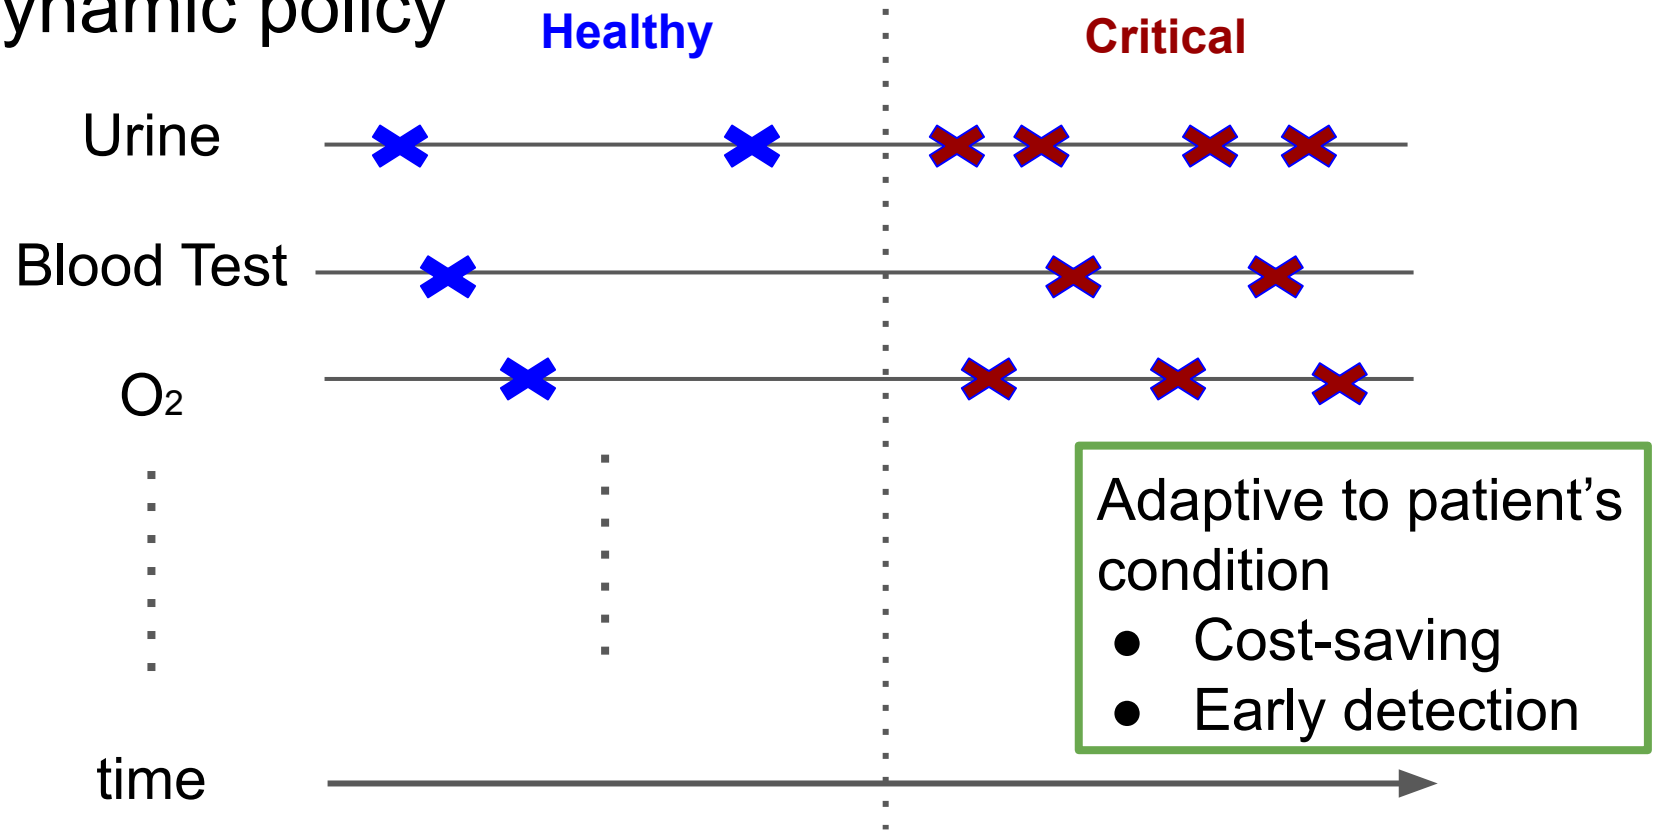

Contributions

- RL framework for **cost-sensitive scheduling** of measurements in **time-series**
- **Scalable** to large number of measurements
- Promising results in a **real-world ICU dataset** (MIMIC3)

System Pipeline

Forecasting Model

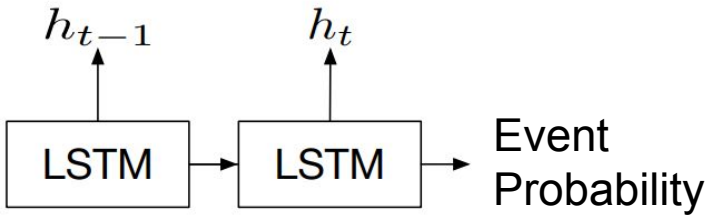

System Pipeline

Forecasting Model

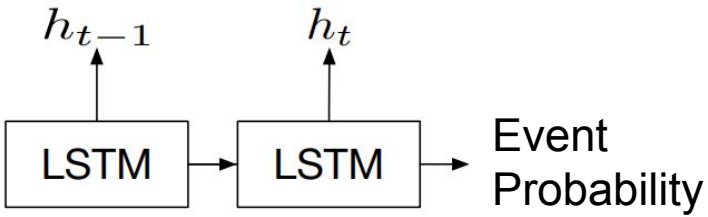

Policy Learning - Dueling Deep Q Network

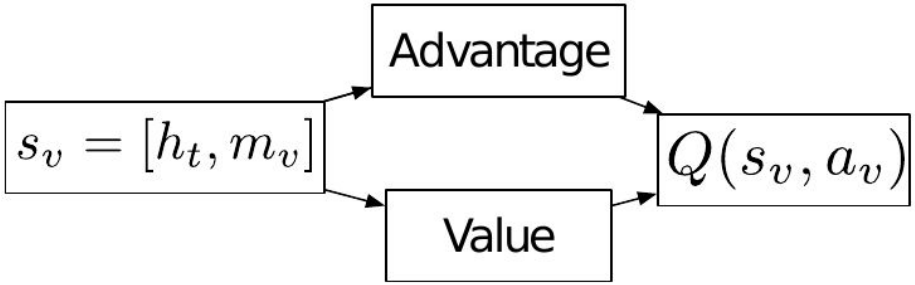

Problem of large action space

- Any combination of D measurements is a valid action
  - 2^D possible actions

Problem of large action space

- Any combination of D measurements is a valid action
  - 2^D possible actions

- Solutions:
  - Independent Policy
  - **Sequential Policy**
    - Only D+1 actions

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

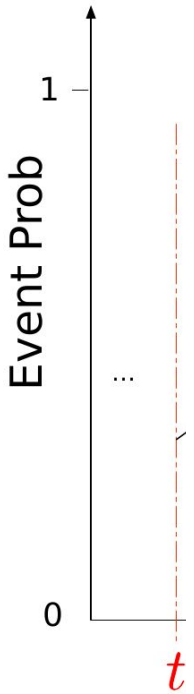

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

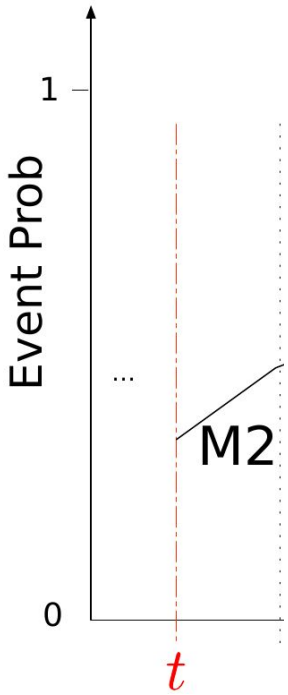

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

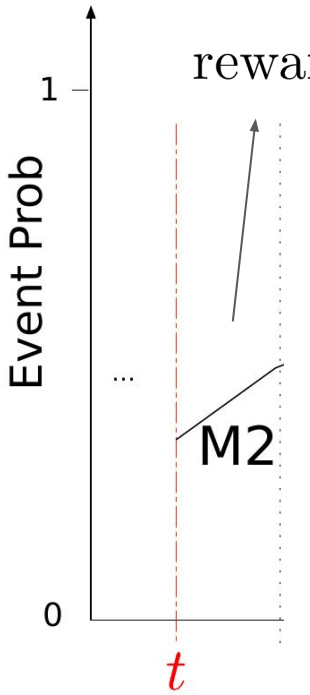

$$\text{reward } r = \Delta_P(M2) - \lambda \times c(M2)$$

Prob Gain from M2 Measurement cost

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

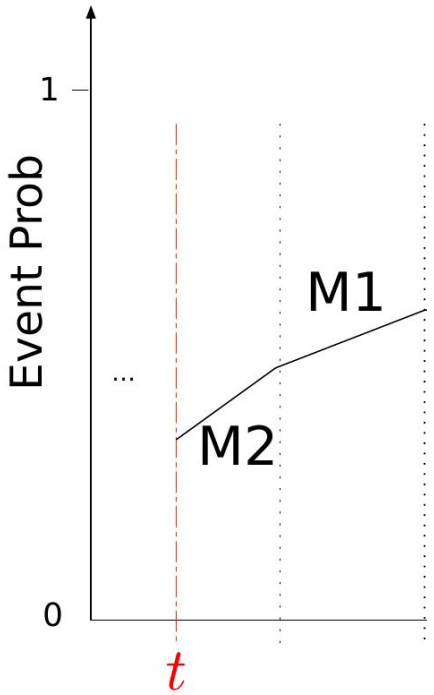

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

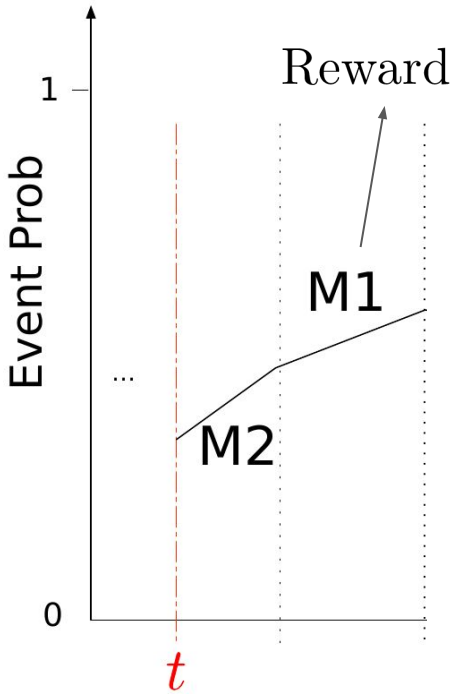

$$\text{Reward } r = \boxed{\Delta_P(M1)} - \lambda \times \boxed{c(M1)}$$

Prob Gain from M1 Measurement cost

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

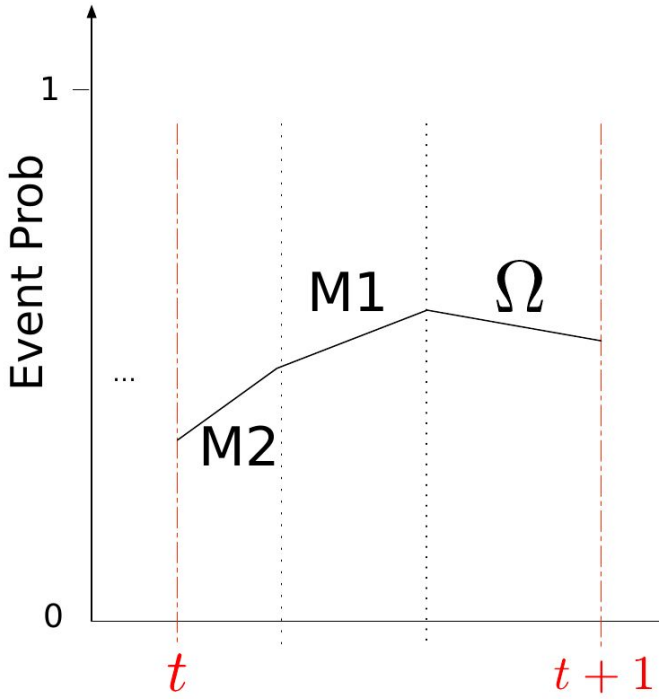

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

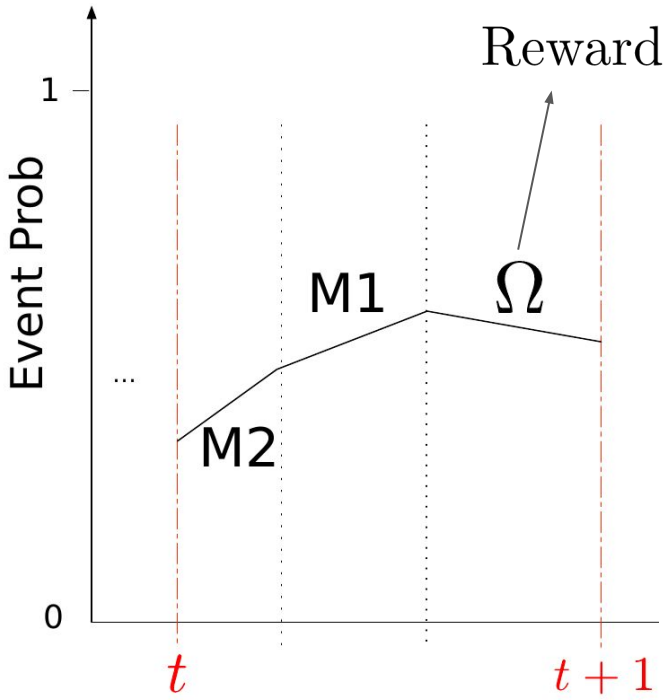

No Measurement Cost

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

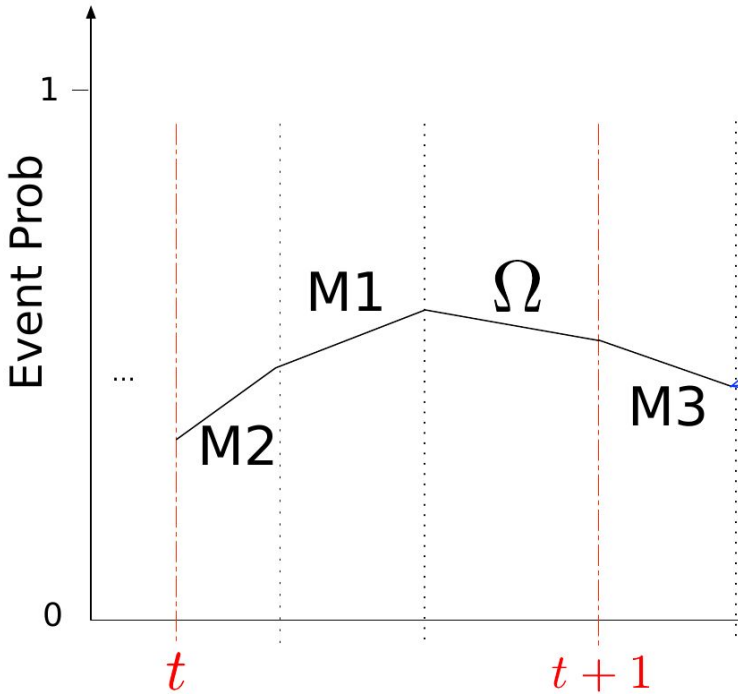

Policy Illustration

M: Measurement A = {M1, M2, M3, Ω? } Ω: Stop-Action

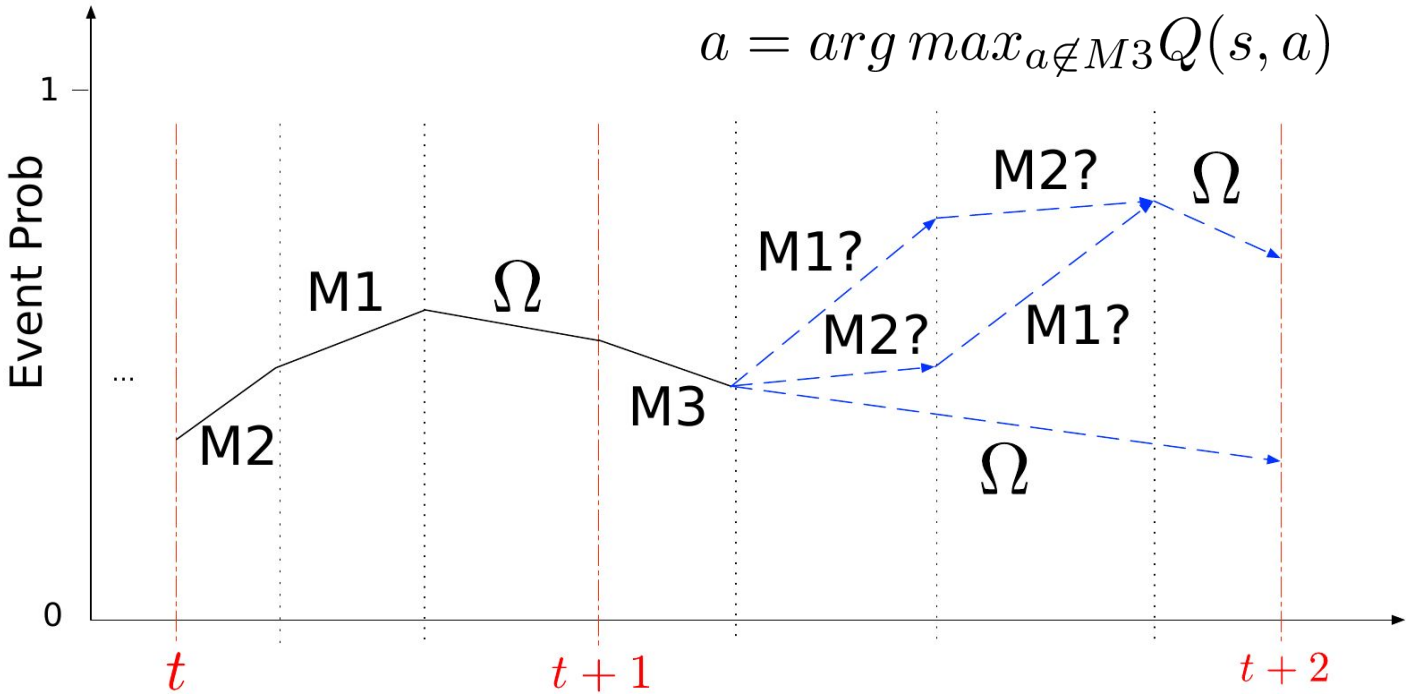

Off-Policy Policy Evaluation

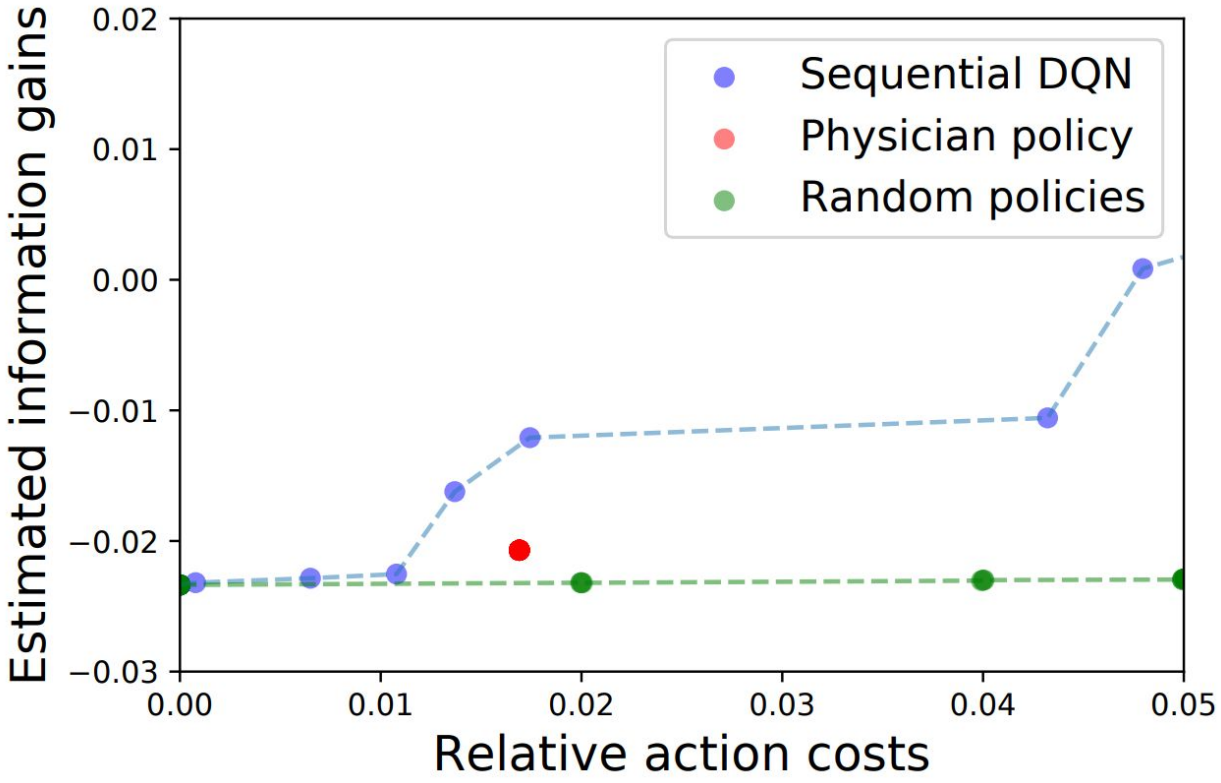

Off-Policy Policy Evaluation

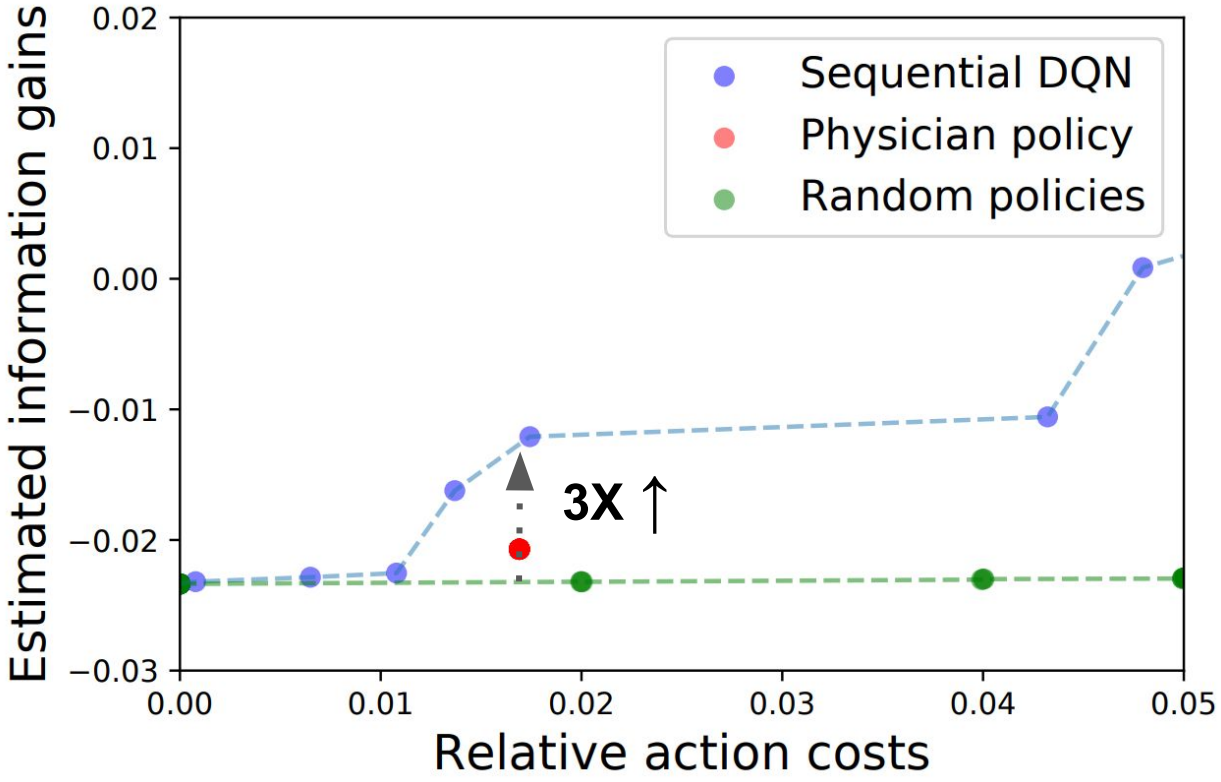

Off-Policy Policy Evaluation

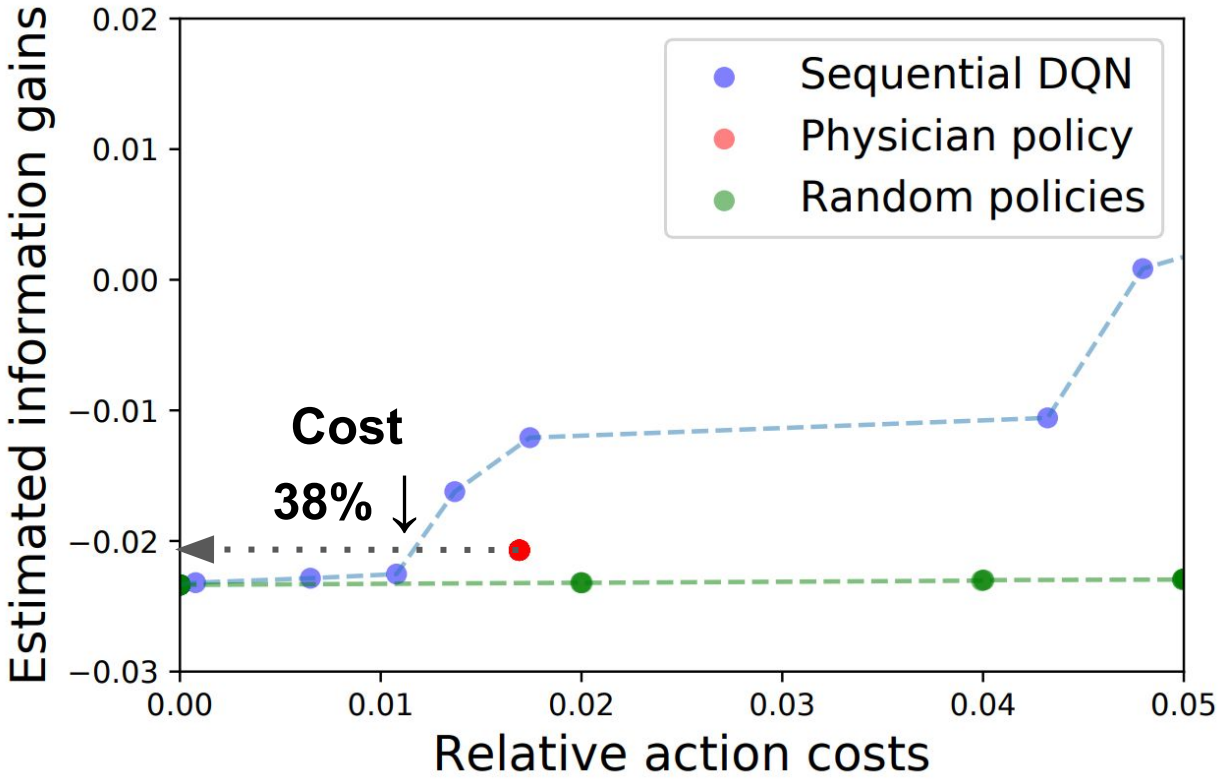

Measurements frequency

Physician’s policy

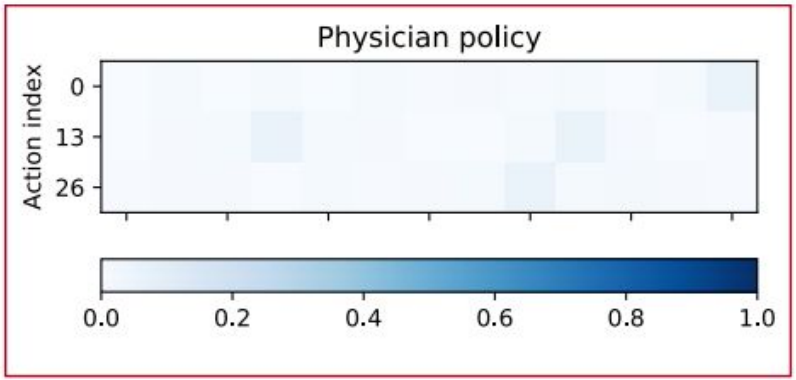

Measurements frequency

Physician's policy

Phosphate

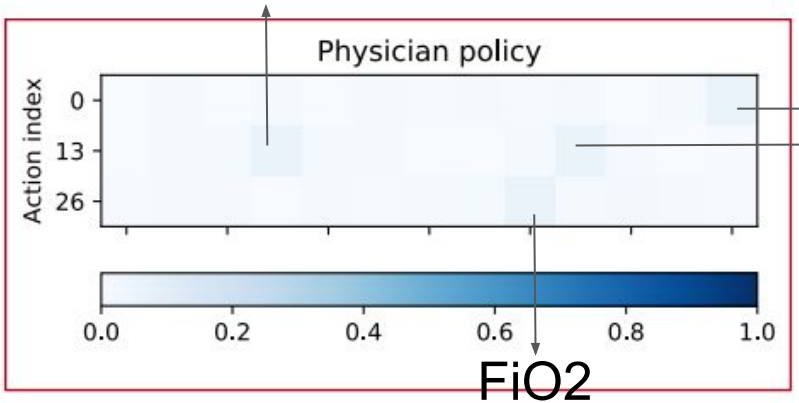

Hemoglobin

SBP

FiO2

Measurements frequency

Physician's policy

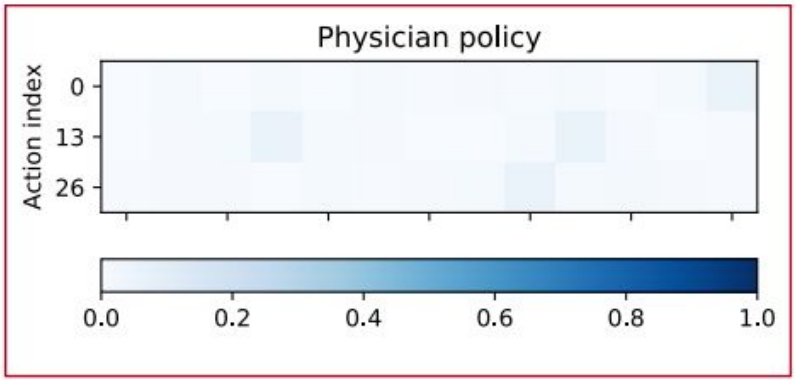

RL policy

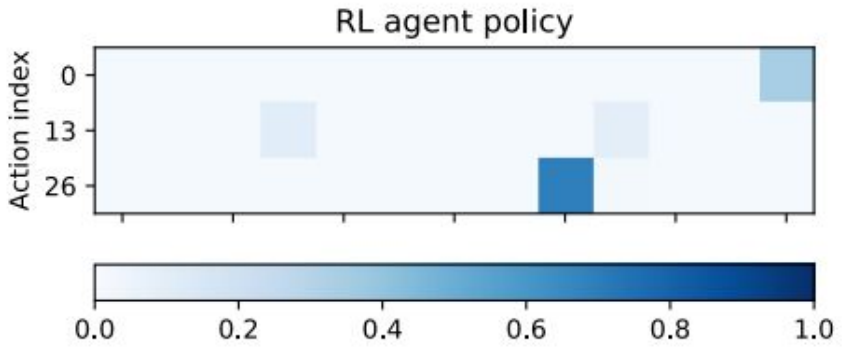

Dynamic Measurement Scheduling for Event Forecasting Using Deep RL

- Code and data preprocessing are released at https://github.com/zzzace2000/autodiagnosis

- Poster **# 247** @ Pacific Ballroom
- Wed 06:30 -- 09:00 PM