Counterfactual inference in sequential experiments via nearest neighbors


Raaz Dwivedi

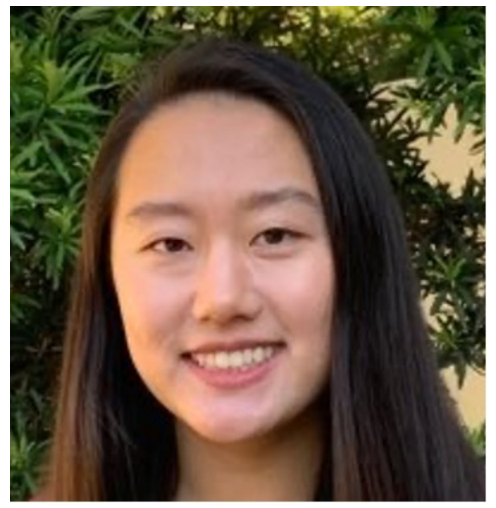
Katherine Tian


Sabina Tomkins


Predrag Klasnja


Susan Murphy


Devavrat Shah

HARVARD UNIVERSITY

Massachusetts Institute of Technology

UNIVERSITY OF MICHIGAN

Informs Annual Meeting, Oct 17, 2022

FODSI NSF NIH

DMS2023528, U101CA229445,
NIH NHLBI R01HL125440,
NIH/NIDA P50DA054039,
NIH/NIBIB & OD P41EB028242,
NIH/NCI U01CA229437,
NIH/NIDCR UH3DE028723

Sequentially adaptive experiments

- Digital treatment delivery **personalized** to user behavior via **reinforcement learning**

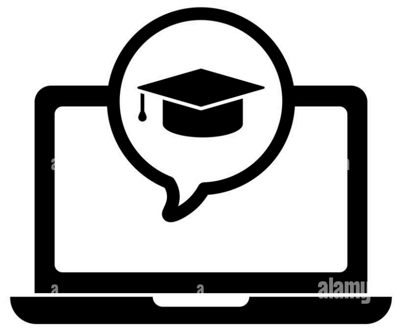

- Multiple users assigned adaptively sampled treatments over multiple time points

Personalized HeartSteps

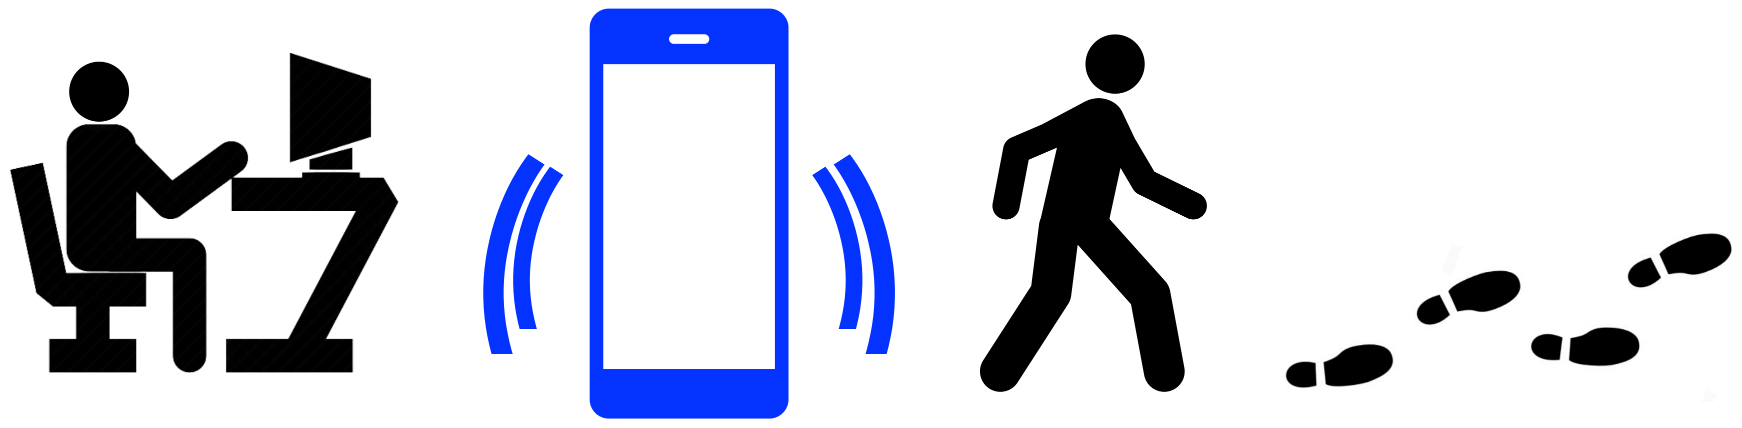

91 users sent randomized push notifications 5 times/day for 90 days; outcome = 30 min step count

[Liao et al. 2020]

Mathematical set-up

For user $i \in [N]$ at time $t \in [T]$, observed outcome

$$Y_{i,t} = \theta_{i,t}^{(A_{i,t})} + \text{noise}$$

$A_{i,t}$: treatment assigned via policy$_{i,t}$

$\theta_{i,t}^{(a)}$: mean potential outcome (counterfactual) for treatment a

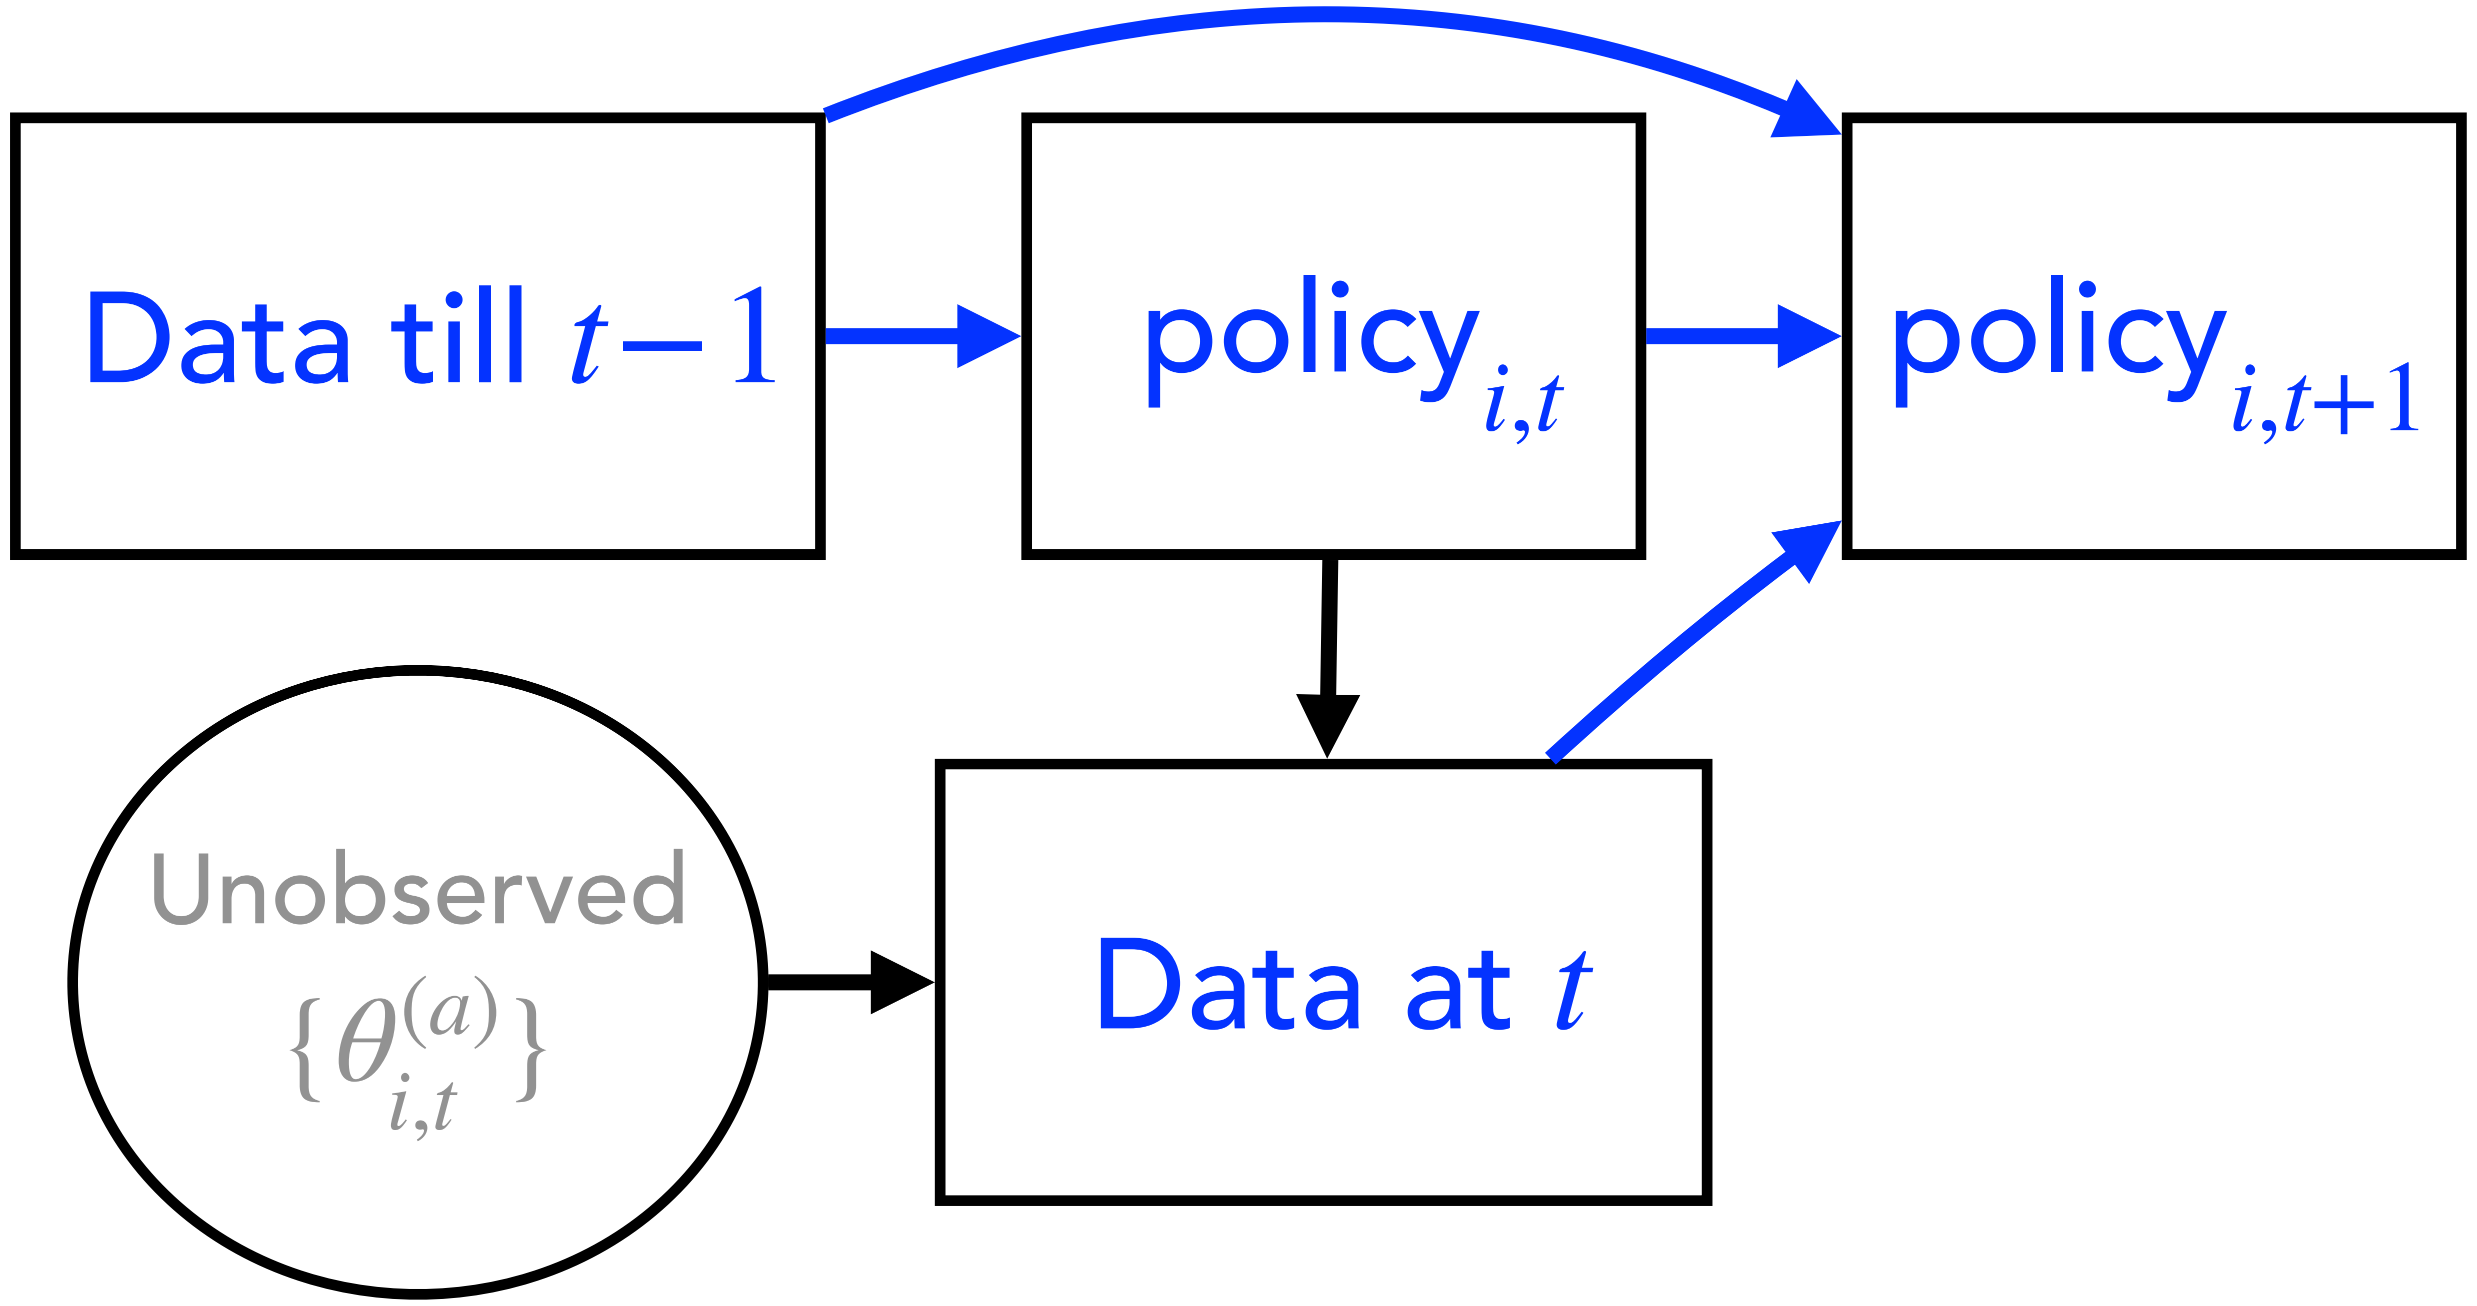

Adaptive sampling policy that can pool across useres

Key questions

- Is the **treatment effective**? [this talk]
  - **Our goal: Estimate unit x time-level counterfactual mean** $\theta_{i,t}^{(a)}$
- Is the algorithm actually personalizing?

Treatment-specific non-parametric factor model

$$\theta_{i,t}^{(a)} \triangleq f^{(a)}(u_i^{(a)}, v_t^{(a)})$$

No parametric assumptions on

- **unknown** non-linearity
- distributions of **unobserved** latent factors and noise

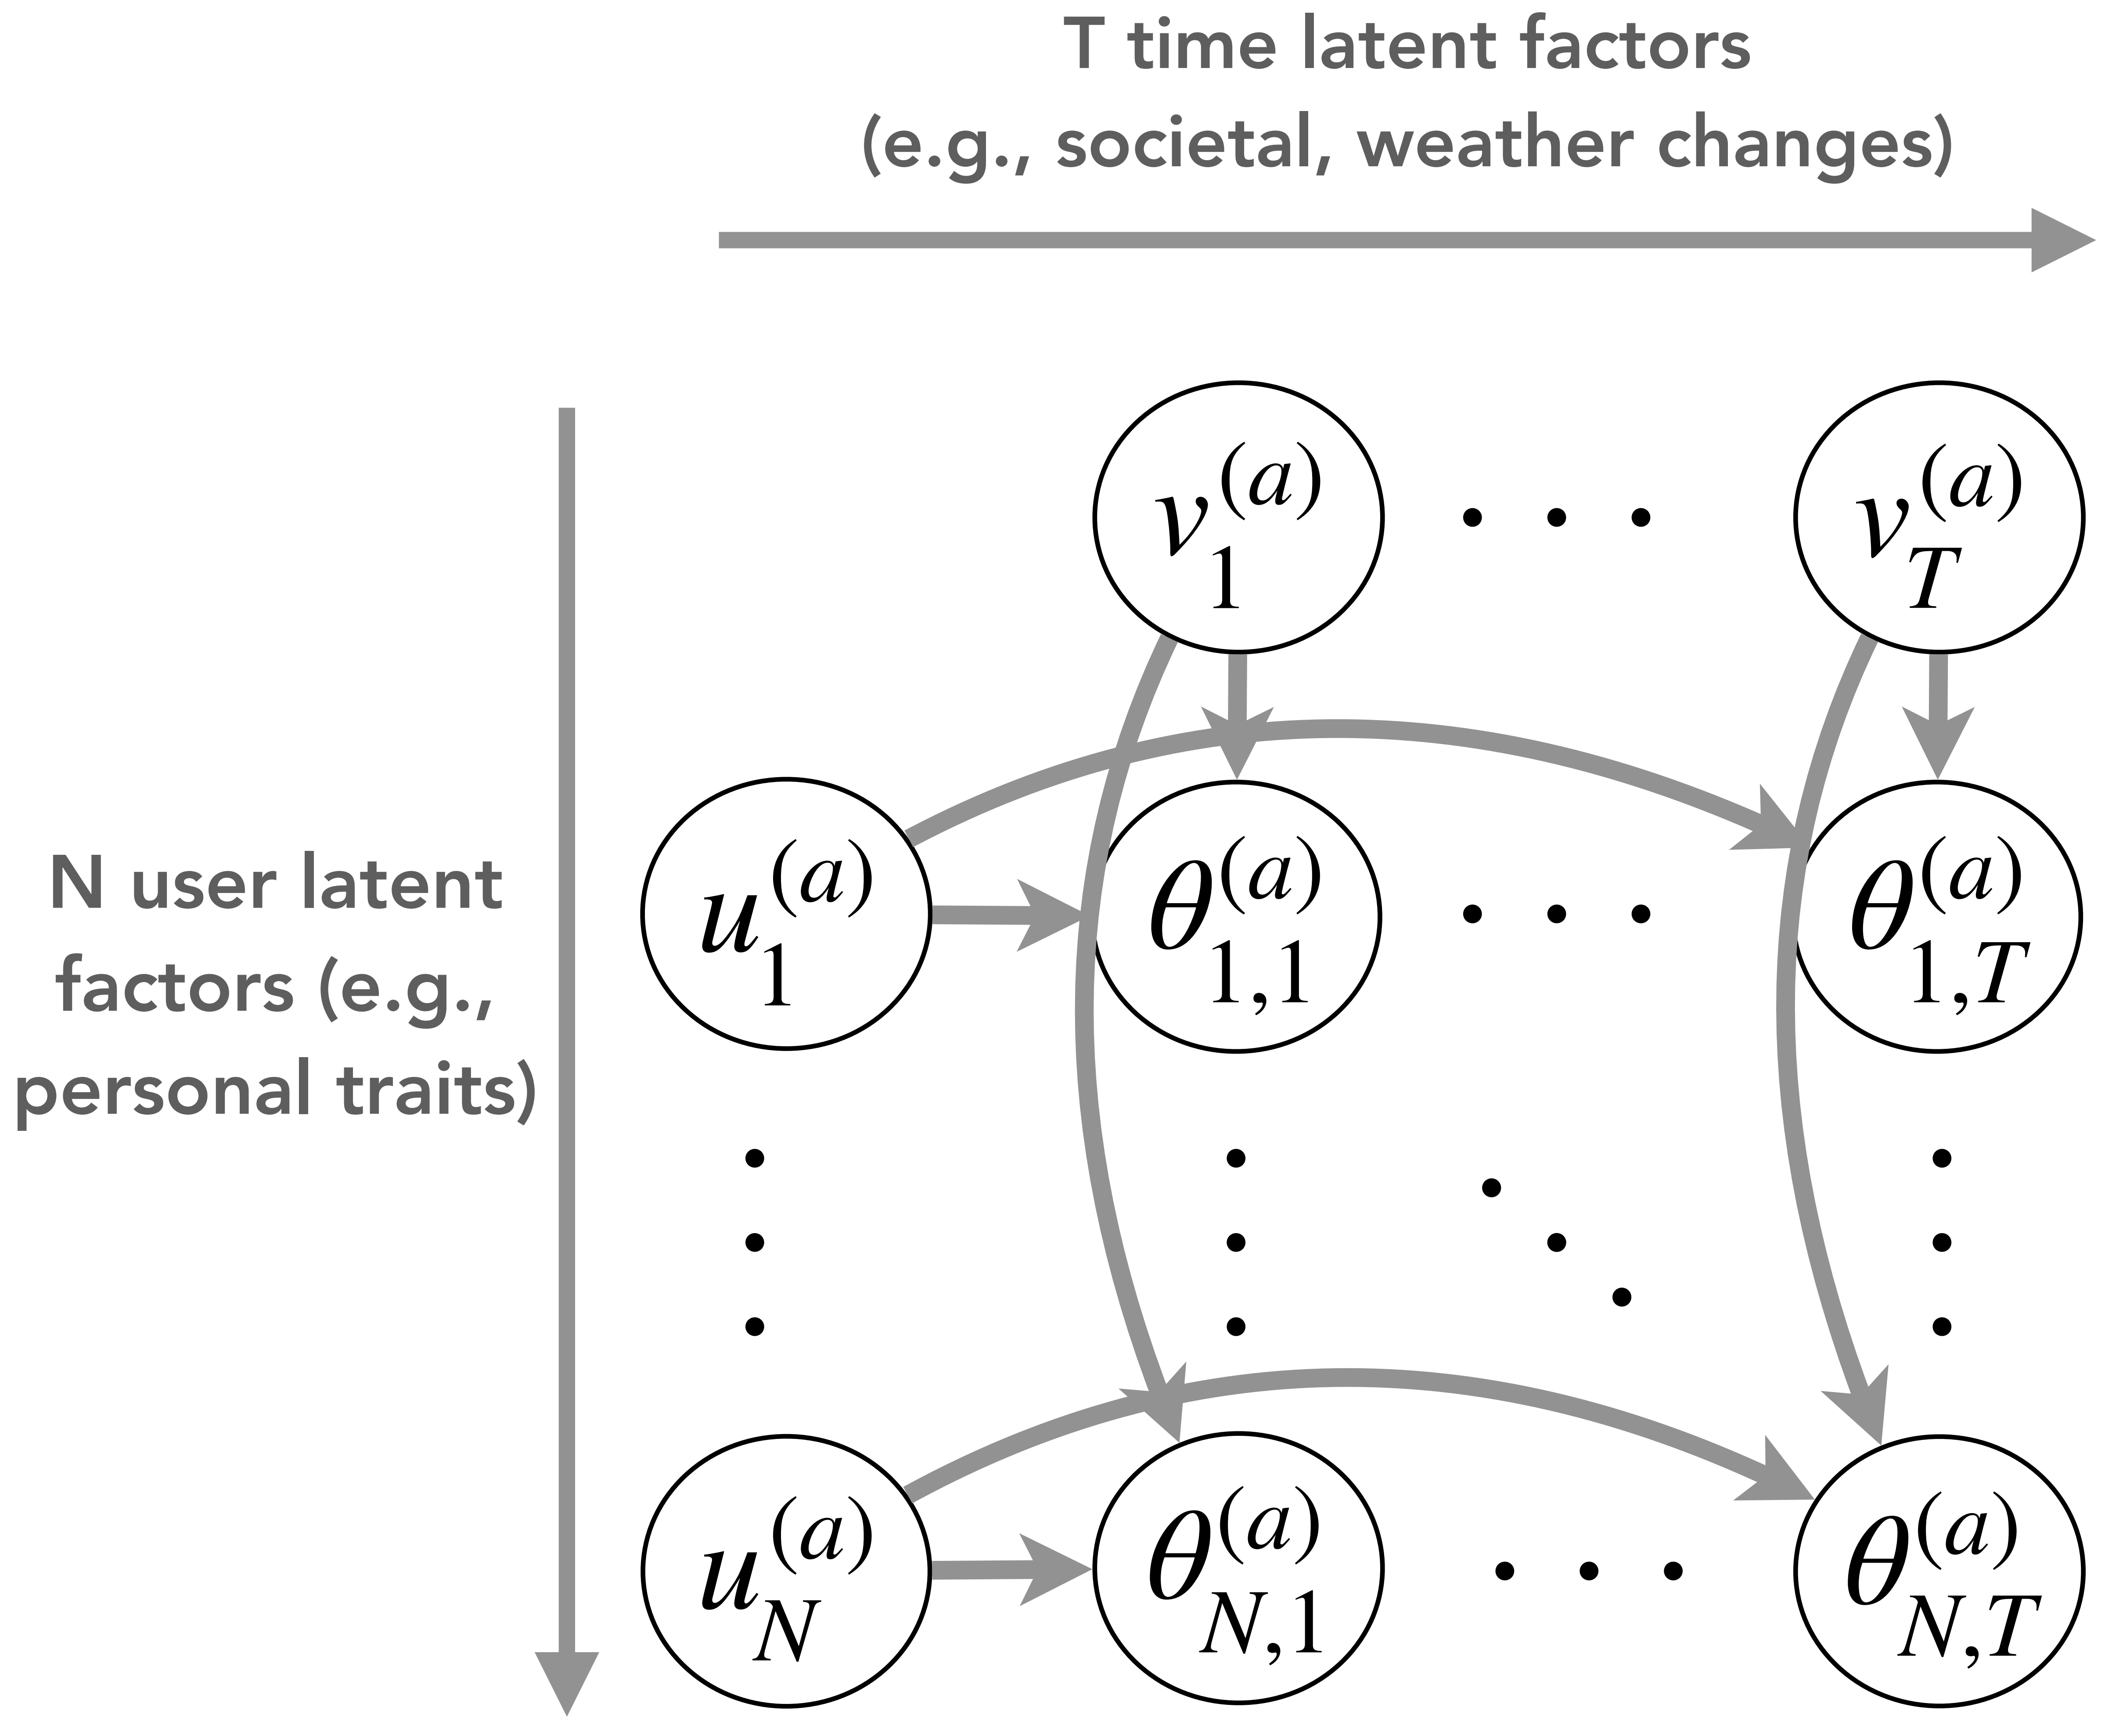

User nearest neighbors estimator for $\theta_{i,t}^{(a)}$

Compute user distance:

$$\rho_{i,j}^{(a)} = \frac{\sum_{t'} \mathbf{1}(A_{j,t'}=A_{i,t'}=a)(Y_{j,t'}-Y_{i,t'})^2}{\sum_{t'} \mathbf{1}(A_{j,t'}=A_{i,t'}=a)}$$

User nearest neighbors estimator for $\theta_{i,t}^{(a)}$

Compute user distance: $$\rho_{i,j}^{(a)} = \frac{\sum_{t'} \mathbf{1}(A_{j,t'}=A_{i,t'}=a)(Y_{j,t'}-Y_{i,t'})^2}{\sum_{t'} \mathbf{1}(A_{j,t'}=A_{i,t'}=a)}$$

Average over user neighbors: $$\widehat{\theta}_{i,t,\mathrm{userNN}}^{(a)} = \frac{\sum_{j} \mathbf{1}(\rho_{i,j}^{(a)} \leq \eta, A_{j,t}=a) Y_{j,t}}{\sum_{j} \mathbf{1}(\rho_{i,j}^{(a)} \leq \eta, A_{j,t}=a)}$$

We prove a **high probability error for each estimate**

For suitably chosen threshold η and under regularity conditions

- Lipschitz non-linearity
- iid latent factors, sub-Gaussian noise
- generic sequentially adaptive policies that assign treatments independently conditioned on history$_t$ with probability $\geq p$

$$(\widehat{\theta}^{(a)}_{i,t,\mathrm{userNN}} - \theta^{(a)}_{i,t})^2 \lesssim \frac{1}{p\sqrt{T}} + \frac{M}{pN}$$

M = alphabet size for user factors

(We also establish results for continuous factors and p going to 0)

User-NN guarantees: Advantages and challenge

$$(\widehat{\theta}^{(a)}_{i,t,\text{userNN}} - \theta^{(a)}_{i,t})^2 \lesssim \frac{1}{\sqrt{T}} + \frac{M}{N}$$

- **First unit-time level guarantee** in sequential experiments
- Leads to asymptotic **confidence intervals**
- How do we fix the **slow rate in T** and obtain

$$(?? - \theta^{(a)}_{i,t})^2 \lesssim O\left(\frac{1}{T} + \frac{1}{N}\right)$$

A general challenge

For $\theta^\star = u^\star v^\star$ with estimates $\hat{u}$ and $\hat{v}$, the error is

$$
\begin{aligned}
|u^\star v^\star - \hat{u}\,\hat{v}| &\leq |v^\star| \cdot |\hat{u} - u^\star| + |\hat{u}| \cdot |\hat{v} - v^\star| \\
&= O(|\hat{u} - u^\star| + |\hat{v} - v^\star|)
\end{aligned}
$$

A general challenge

For $\theta^\star = u^\star v^\star$ with estimates $\hat{u}$ and $\hat{v}$, the error is

$$\begin{aligned} |u^\star v^\star - \hat{u}\,\hat{v}| &\le |v^\star| \cdot |\hat{u} - u^\star| + |\hat{u}| \cdot |\hat{v} - v^\star| \\ &= O(|\hat{u} - u^\star| + |\hat{v} - v^\star|) \end{aligned}$$

When $\theta^\star = u_i^{(a)} v_t^{(a)}$

User-NN: $\hat{u} =$ avg. user factors over user neighbors $\implies$ Error $= O(|\hat{u} - u^\star|)$

$\hat{v} = v^\star$

Time-NN: $\hat{u} = u^\star$ $\implies$ Error $= O(|\hat{v} - v^\star|)$

$\hat{v} =$ avg. time factors over time neighbors

Can we convert the $+$ to $\times$?

A general challenge: $|u^\star v^\star - \hat{u}\,\hat{v}| = O(|\hat{u} - u^\star| + |\hat{v} - v^\star|)$

A possible solution: $|u^\star v^\star - ??| = O(|\hat{u} - u^\star| \times |\hat{v} - v^\star|)$

A "doubly-robust" solution!

A general challenge: $|u^\star v^\star - \hat{u}\,\hat{v}| = O(|\hat{u} - u^\star| + |\hat{v} - v^\star|)$

A possible solution: $|u^\star v^\star - ??| = O(|\hat{u} - u^\star| \times |\hat{v} - v^\star|)$

$$u^\star v^\star - ?? = (\hat{u} - u^\star) \times (\hat{v} - v^\star)$$

$$= \hat{u}\hat{v} - u^\star \hat{v} - \hat{u} v^\star + u^\star v^\star$$

$$\Longrightarrow \qquad ?? = \hat{u} v^\star + u^\star \hat{v} - \hat{u}\hat{v}$$

A "doubly-robust" solution!

A general challenge: $|u^\star v^\star - \hat{u}\,\hat{v}| = O(|\hat{u} - u^\star| + |\hat{v} - v^\star|)$

A possible solution: $|u^\star v^\star - ??| = O(|\hat{u} - u^\star| \times |\hat{v} - v^\star|)$

$$
\begin{aligned}
u^\star v^\star - ?? &= (\hat{u} - u^\star) \times (\hat{v} - v^\star) \\
&= \hat{u}\hat{v} - u^\star \hat{v} - \hat{u} v^\star + u^\star v^\star \\
\implies \quad ?? &= \hat{u} v^\star + u^\star \hat{v} - \hat{u}\hat{v}
\end{aligned}
$$

$$Y_{j,t} + Y_{i,t'} - Y_{j,t'}$$

$$\rho^{(a)}_{i,j} \le \eta, \quad \rho^{(a)}_{t,t'} \le \eta$$

A "doubly-robust" solution!

A general challenge: $|u^\star v^\star - \hat{u}\,\hat{v}| = O(|\hat{u} - u^\star| + |\hat{v} - v^\star|)$

A possible solution: $|u^\star v^\star - ??| = O(|\hat{u} - u^\star| \times |\hat{v} - v^\star|)$

$$\begin{aligned} u^\star v^\star - ?? &= (\hat{u} - u^\star) \times (\hat{v} - v^\star) \\ &= \hat{u}\hat{v} - u^\star \hat{v} - \hat{u} v^\star + u^\star v^\star \\ \implies \qquad ?? &= \hat{u} v^\star + u^\star \hat{v} - \hat{u}\hat{v} \end{aligned}$$

$$\widehat{\theta}^{(a)}_{i,t,\mathrm{DR\text{-}NN}} = \frac{\sum_{j,t'} \mathbf{1}_{i,t,j,t'} \left(Y_{j,t} + Y_{i,t'} - Y_{j,t'}\right)}{\sum_{j,t'} \mathbf{1}_{i,t,j,t'}}$$

$$\mathbf{1}_{i,t,j,t'} = \mathbf{1}\left(A_{j,t} = A_{i,t'} = A_{j,t'} = a,\ \rho^{(a)}_{i,j} \leq \eta,\ \rho^{(a)}_{t,t'} \leq \eta \right)$$

Doubly-robust estimator fixes the slow rates!

With non-adaptive policies and $M =$ alphabet size for user and time factors

$$(\widehat{\theta}^{(a)}_{i,t,\text{DR-NN}} - \theta^{(a)}_{i,t})^2 \lesssim \frac{M}{T} + \frac{M}{N}$$

$\approx$ user-NN error $\times$ time-NN error

$\lesssim$ **min**{user-NN error, time-NN error}

$$(\widehat{\theta}^{(a)}_{i,t,\text{userNN}} - \theta^{(a)}_{i,t})^2 \lesssim \frac{1}{\sqrt{T}} + \frac{M}{N}$$

A simple view of the doubly robust principle

$$\theta^\star = u^\star v^\star$$

Error

$$\hat{u}\hat{v} \qquad O(|\hat{u} - u^\star| + |\hat{v} - v^\star|)$$

$$\downarrow$$

$$\hat{u}v^\star + u^\star\hat{v} - \hat{u}\hat{v} \qquad O(|\hat{u} - u^\star| \times |\hat{v} - v^\star|)$$

Simulation results

DR-NN error ≈ user-NN error × time-NN error

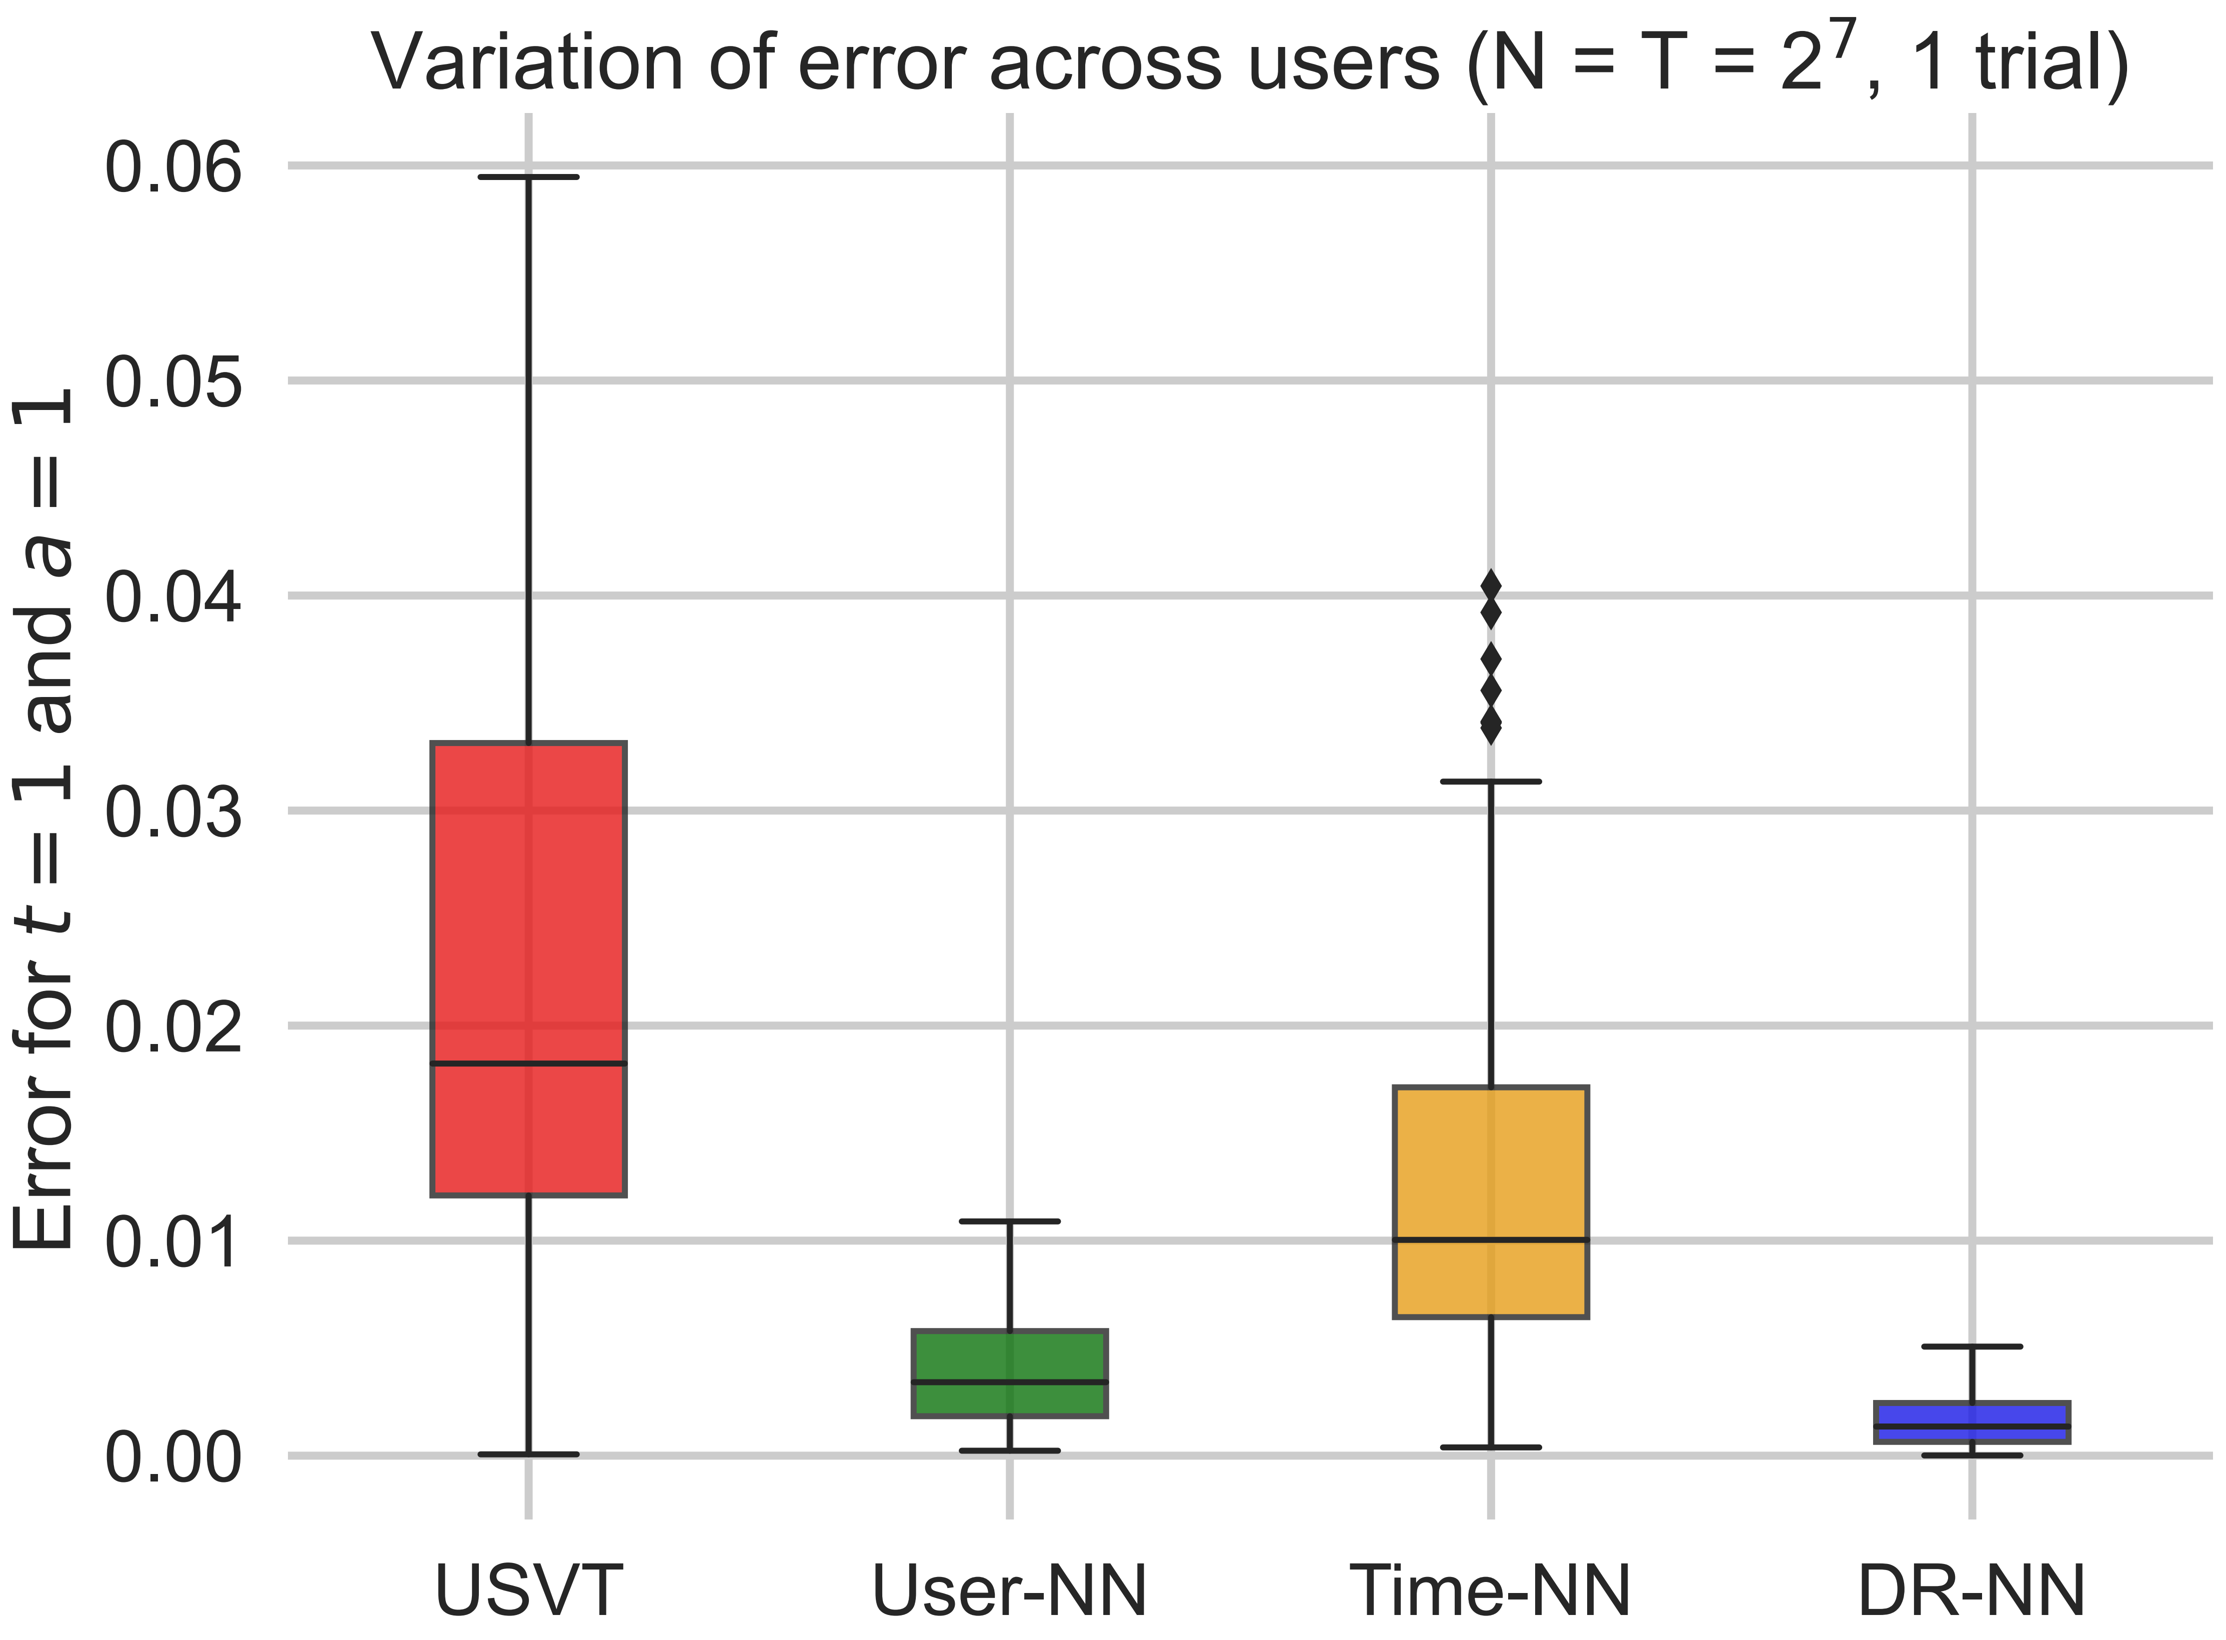

A baseline algorithm from [Chatterjee 2014]

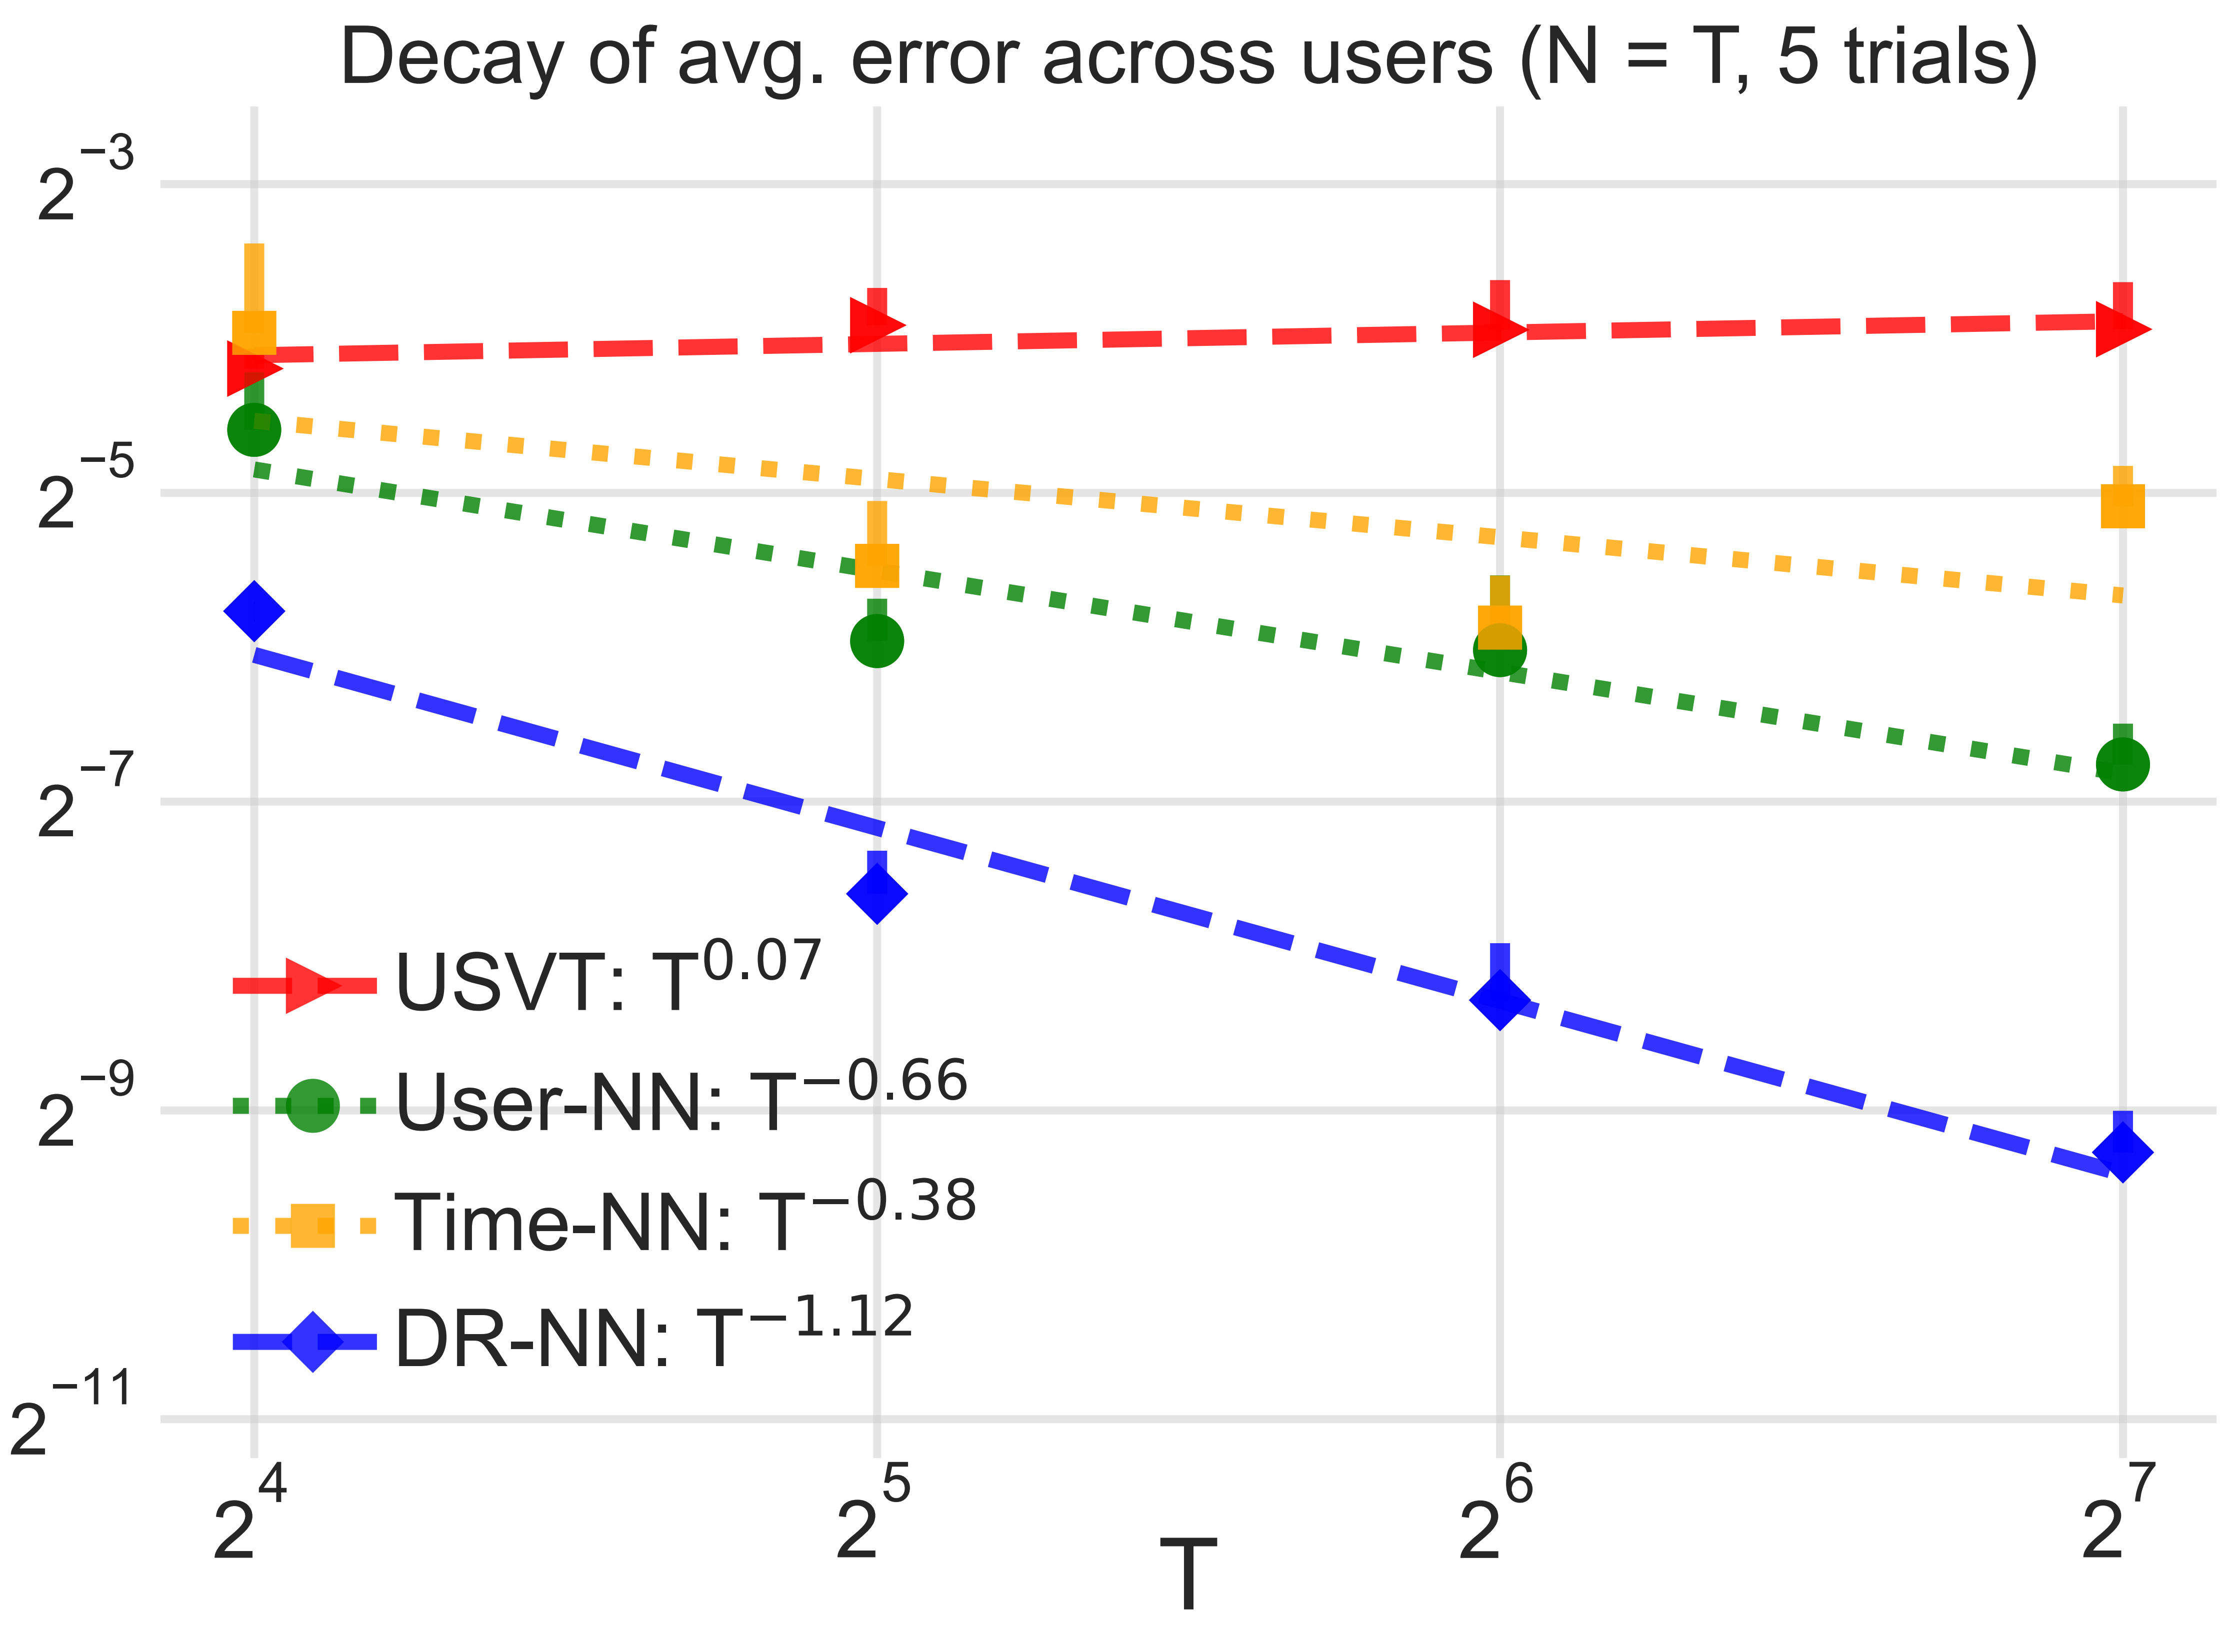

uniform latent factors on $[-0.5, 0.5]^2$,
pooled ε-greedy policy ($\varepsilon = 0.5$), noise$\sim\mathcal{N}(0, 10^{-4})$

Raaz Dwivedi, Harvard & MIT

HeartSteps results

DR-NN error ≈ min { user-NN error , time-NN error }

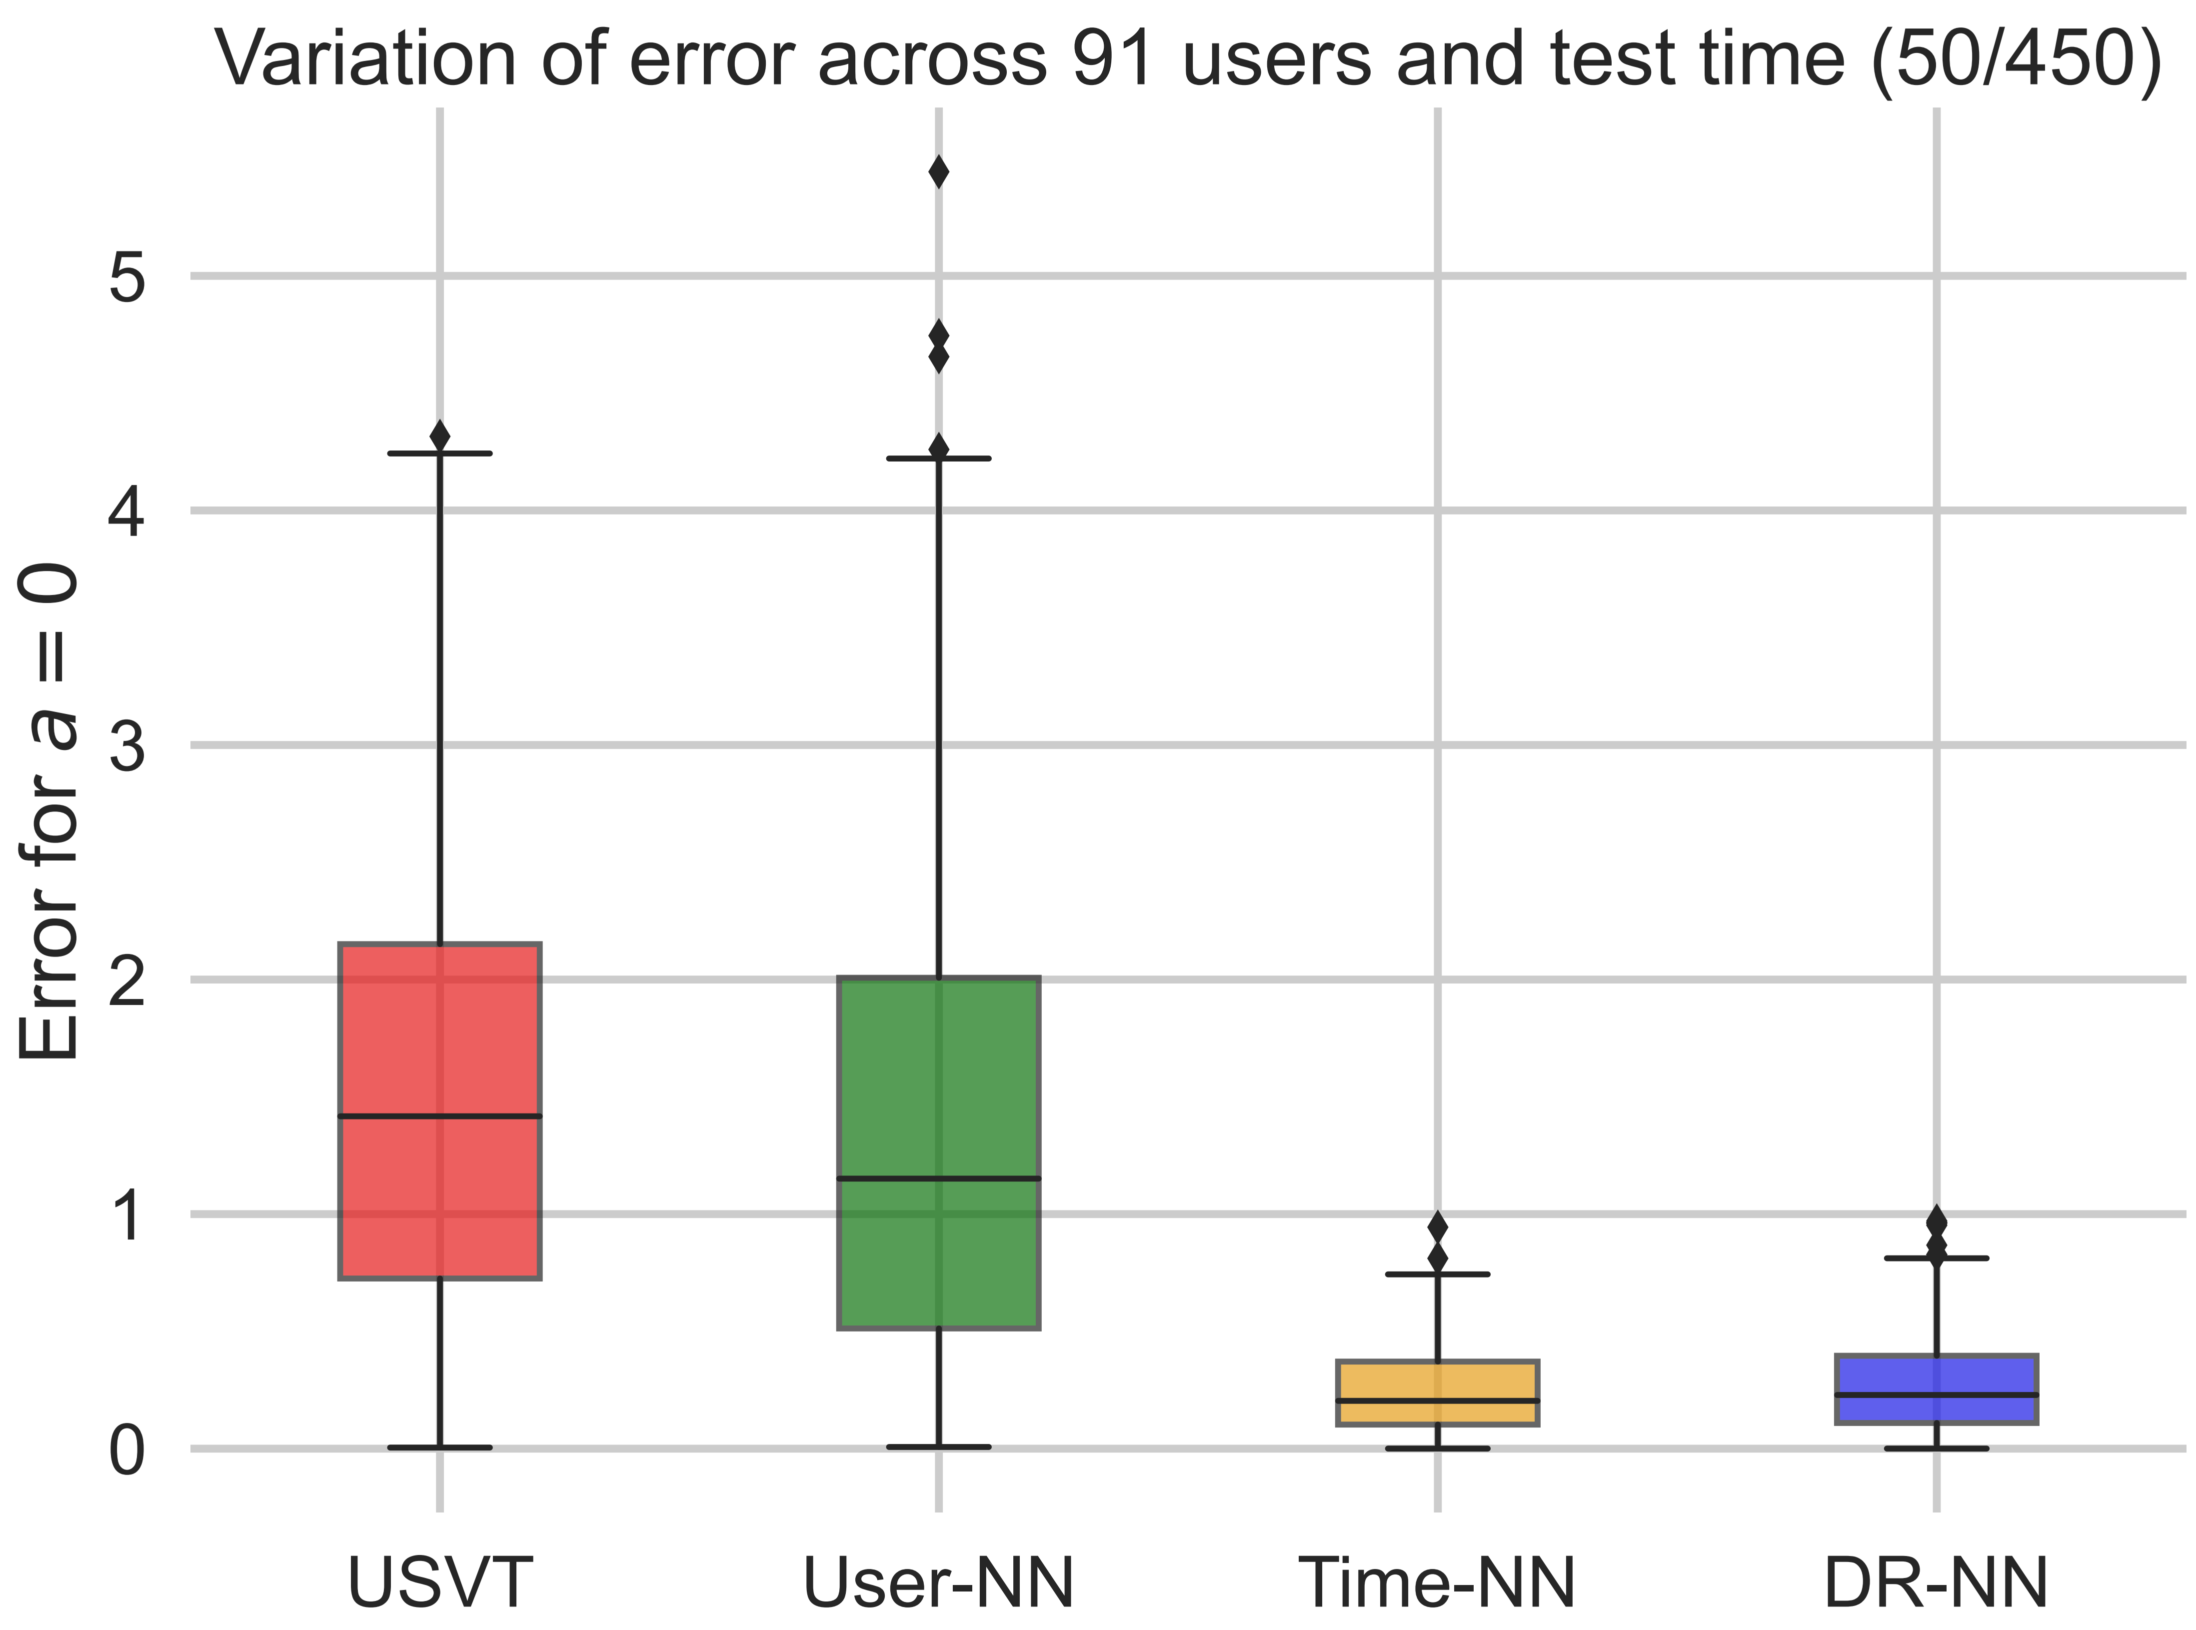

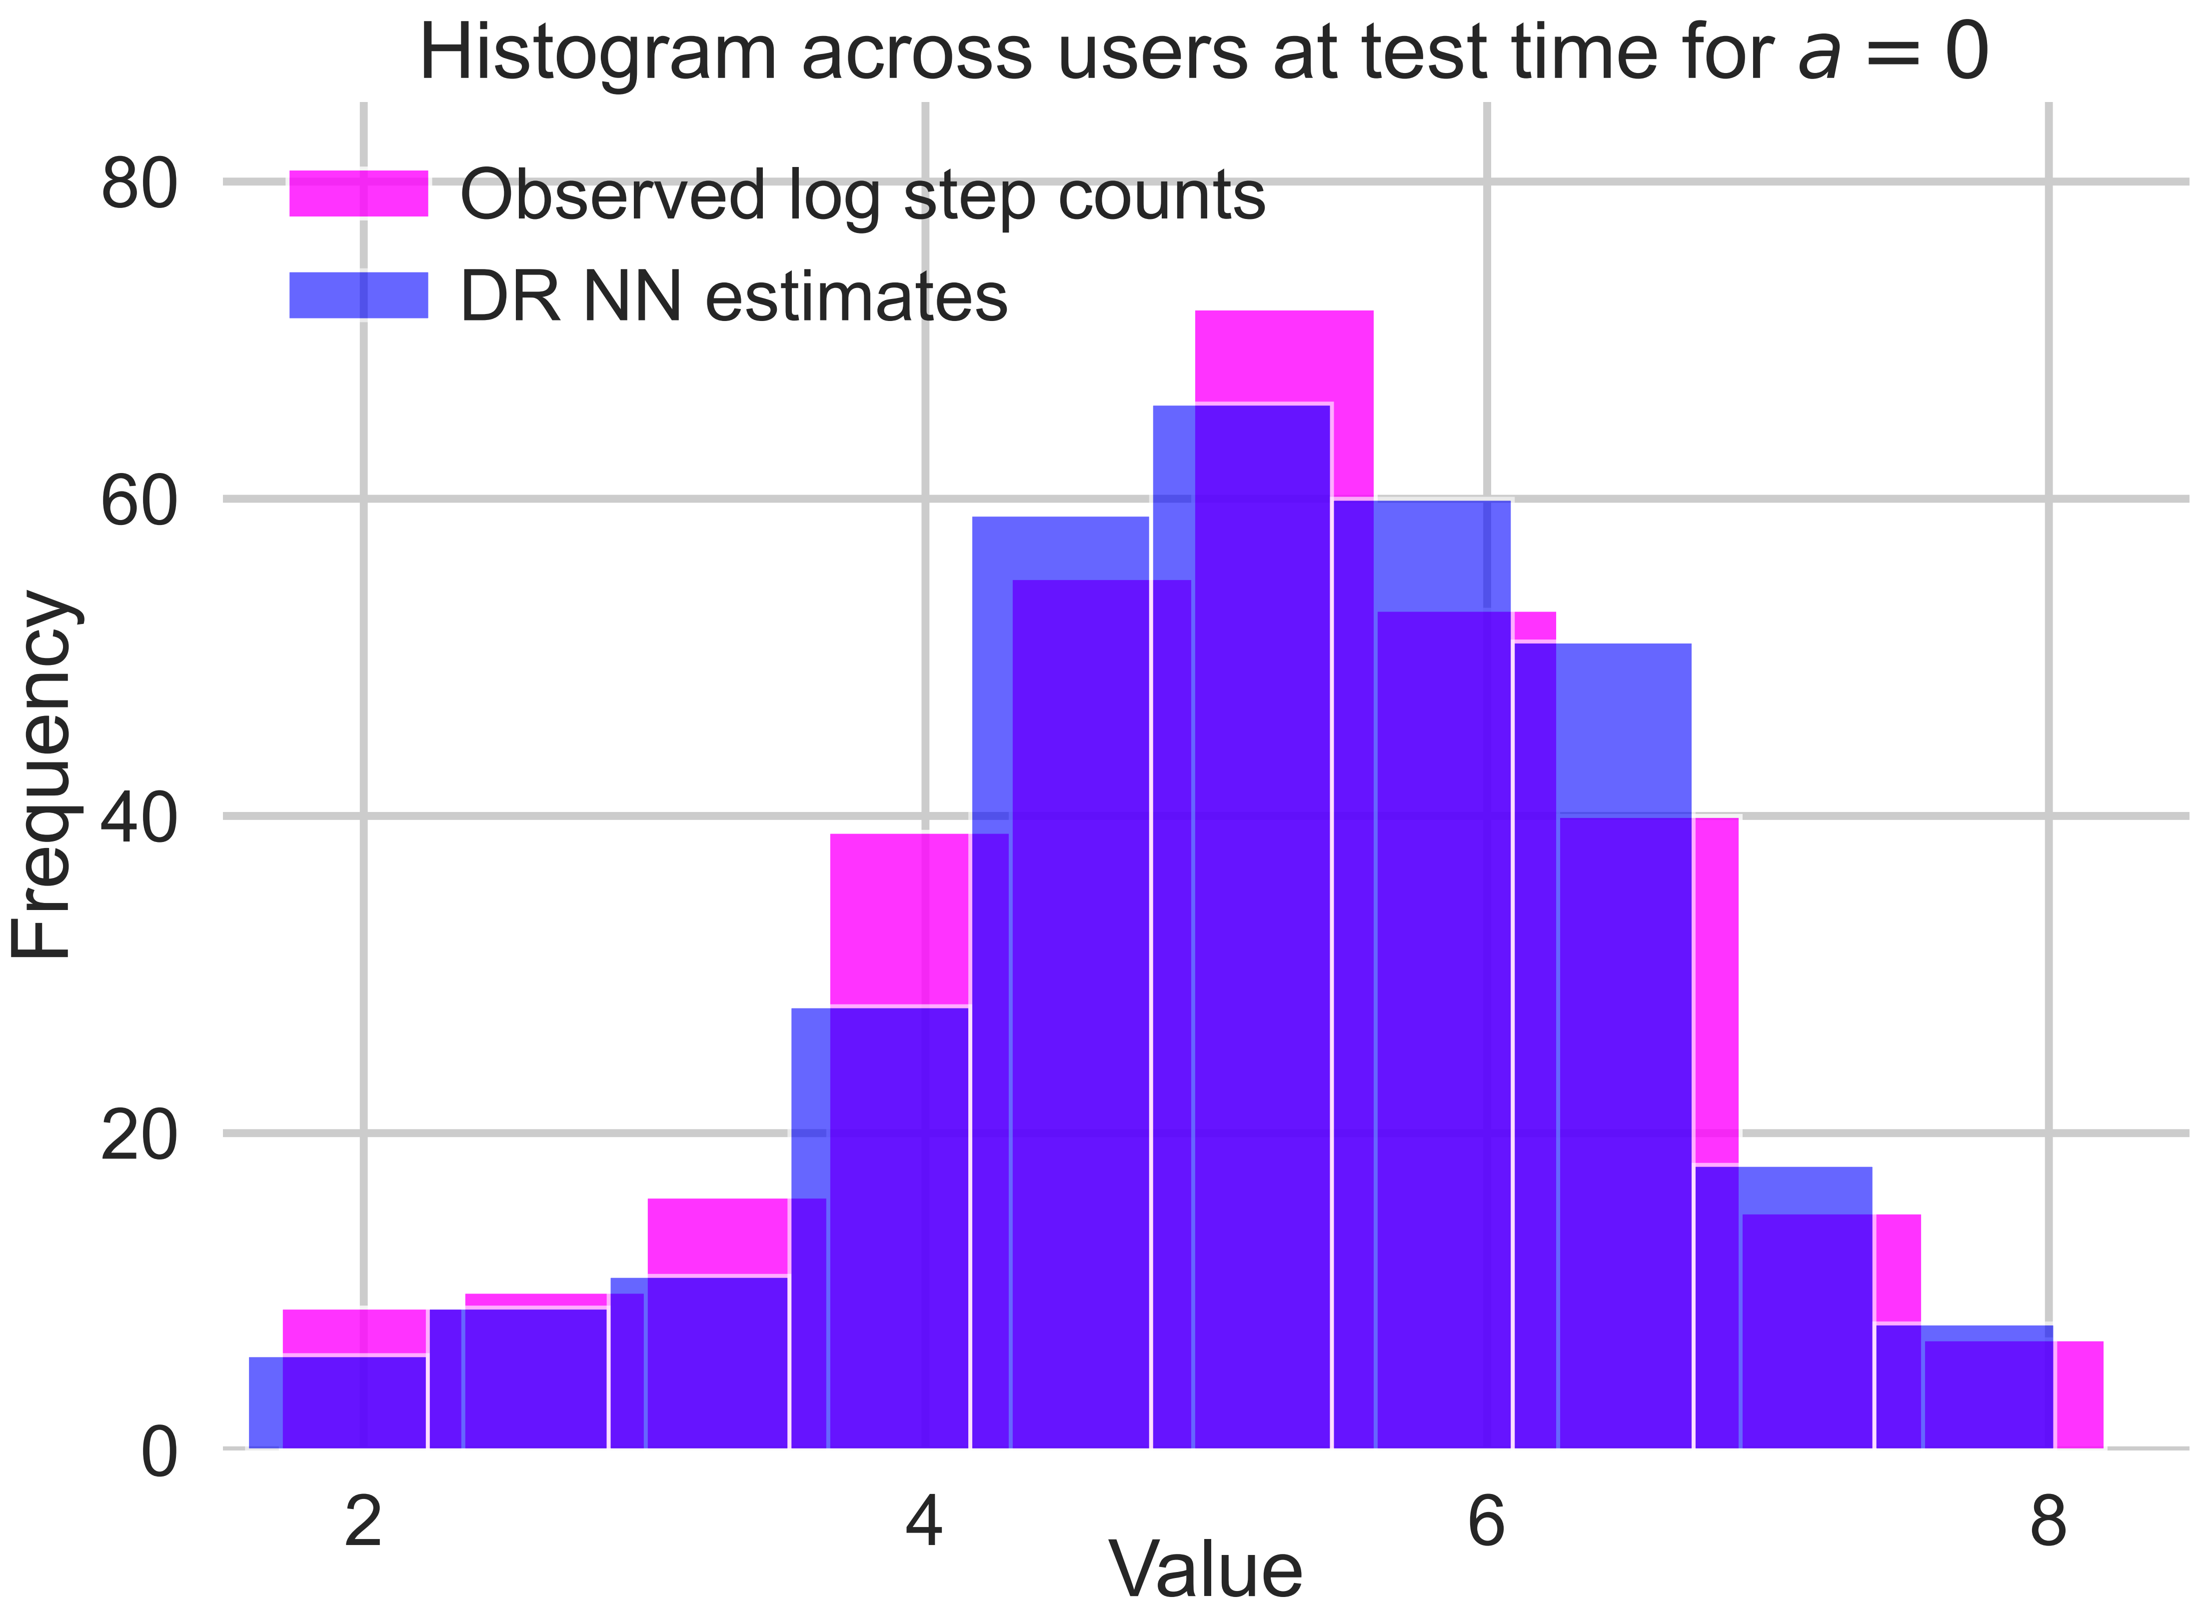

Treatments assigned with Thompson sampling independently across users

Raaz Dwivedi, Harvard & MIT

Summary

Unit-time level counterfactual inference with non-parametric factor models

A doubly robust nearest neighbor estimator satisfying

$$\textbf{DR-NN error} \approx \textbf{user-NN error} \times \textbf{time-NN error}$$
$$\lesssim \textbf{min\{user-NN error, time-NN error\}}$$

Personalized decision-making that accounts for user-specific behavior involves **two fundamental tasks**

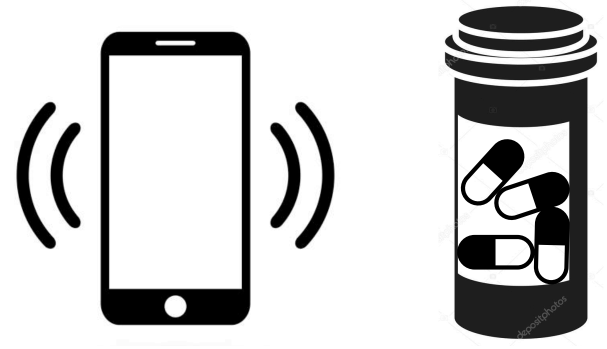

Mobile health, medicine

1. Infer decision's effect from data when the model is unknown

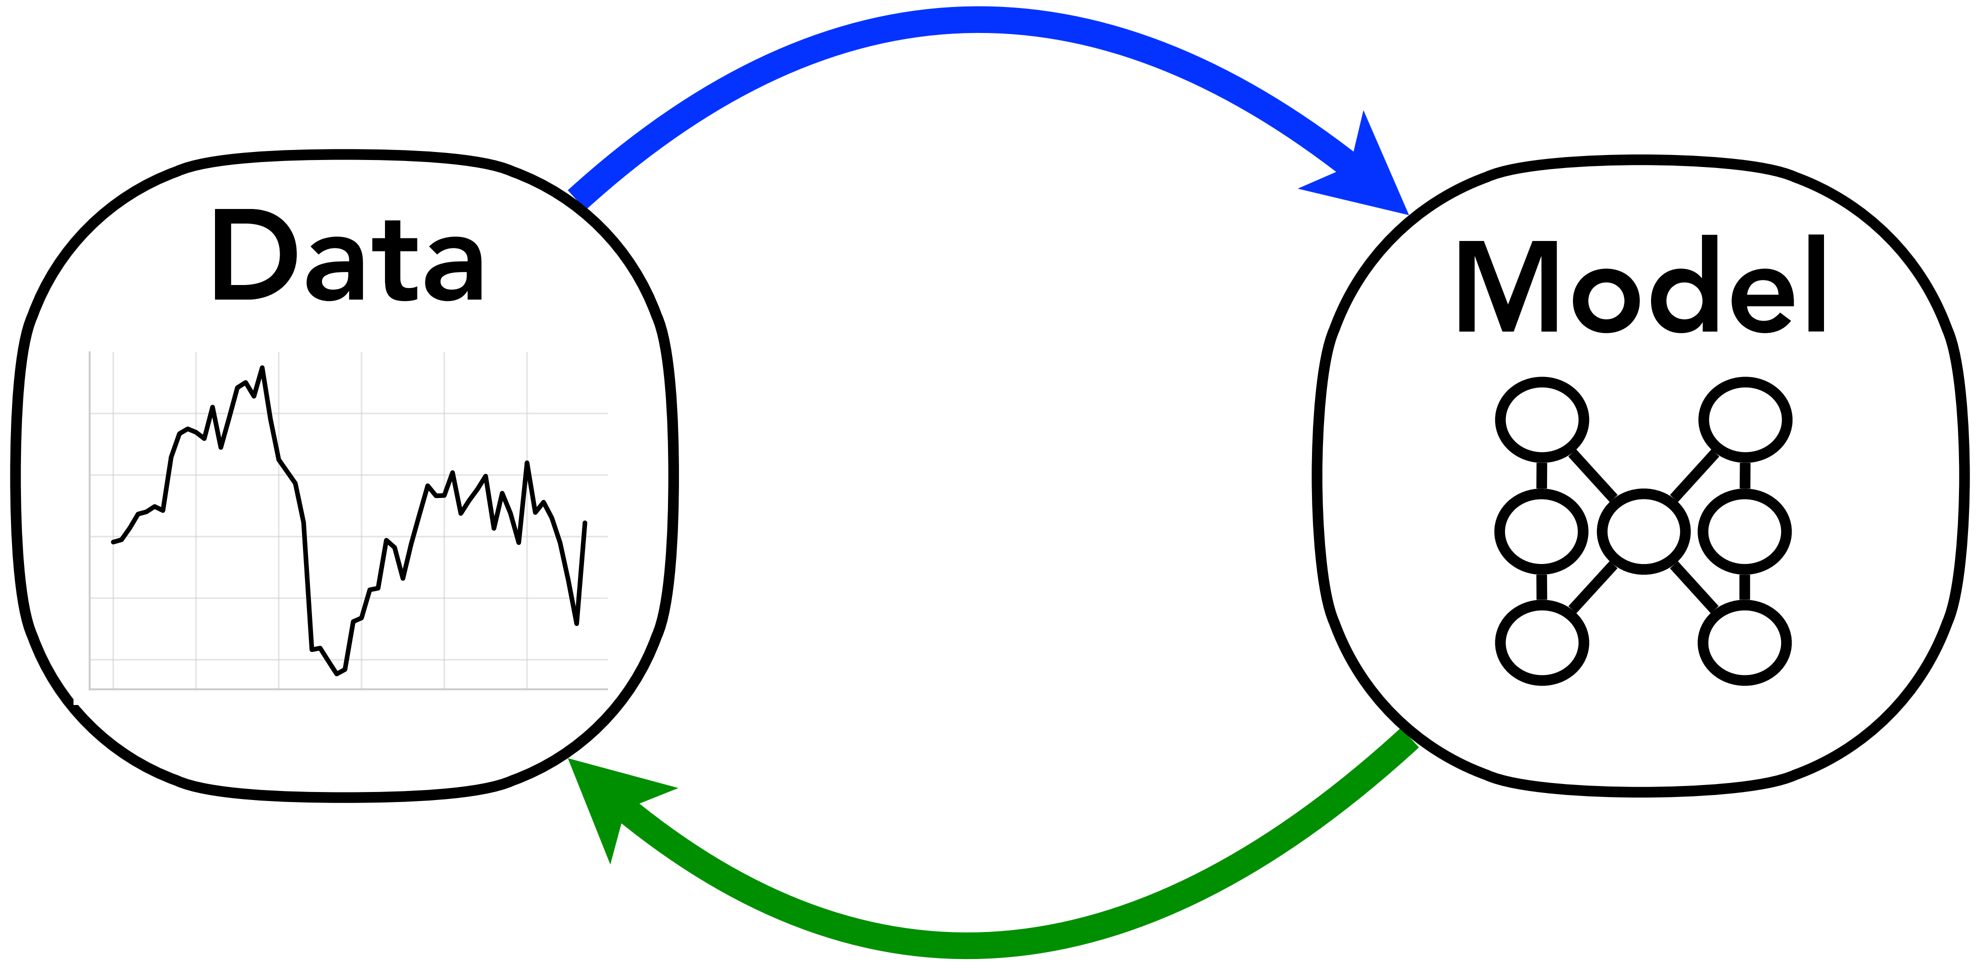

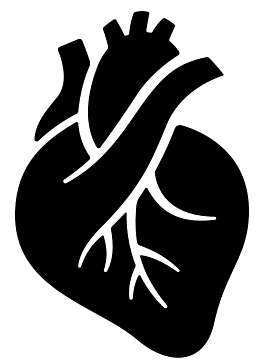

Computational cardiology

2. Simulate decision's effect with a known model

I am on the job market!

raaz@mit.edu
raazdwivedi.github.io
HARVARD UNIVERSITY
Massachusetts Institute of Technology