On counterfactual inference in factor models with nearest neighbors


Raaz Dwivedi[1,2]

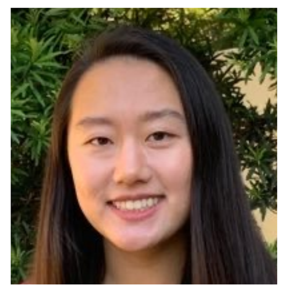
Katherine Tian[1]


Sabina Tomkins[3]


Predrag Klasnja[3]


Susan Murphy[1]


Devavrat Shah[2]

1 Harvard University
2 MIT
3 University of Michigan

raaz@mit.edu
https://raazdwivedi.github.io/
https://arxiv.org/abs/2202.06891

Supporting grants

- DMS2023528
- 1U01CA229445
- NIH NHLBI R01HL125440
- NIH/NIDA P50DA054039
- NIH/NIBIB & OD P41EB028242
- NIH/NCI U01CA229437
- NIH/NIDCR UH3DE028723

Factor models, sequential experiments, nearest neighbors

Panel data: N users with a series of treatments over T time points, where for user $i \in [N]$, at time $t \in [T]$, we observe outcome $Y_{i,t}$ that depends on the assigned treatment $A_{i,t}$, the potential outcomes $\{\theta_{i,t}^{(a)}\}$ for $a \in \mathcal{A}$, and noise $\epsilon_{i,t}$. **Our goal** is counterfactual inference with **latent factor model**, i.e., estimate $\{\theta_{i,t}^{(a)}\}$ with user x time-level guarantee

$$Y_{i,t} = \theta_{i,t}^{(A_{i,t})} + \epsilon_{i,t} \quad A_{i,t} \sim \text{policy}_{i,t}$$
$$\theta_{i,t}^{(a)} = \langle u_i^{(a)}, v_t^{(a)} \rangle$$

(choosing a=1 for simpler notation)

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

User-based nearest neighbors (User-NN):

Find NN users $\mathcal{N}_i^{(a)} = \{j : \rho_{i,j}^{(a)} \le \eta\}$ and average

$$\widehat{\theta}_{i,t,\text{user}}^{(1)} = \frac{\sum_{j\in\mathcal{N}_i^{(1)}} A_{j,t}Y_{j,t}}{\sum_{j\in\mathcal{N}_i^{(1)}} A_{j,t}}$$

For **sequentially adaptive policies**, with suitable η, and regularity conditions with high probability (whp)

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

with M distinct types of users

Simulation results with pooled ε-greedy treatment policy

- Latent factors uniform on $[-0.5, 0.5]^2$, avg. treatment effect = 0, noise-var = 1e-4
- Pooled ε-greedy treatment policy with ε = 0.5
- Methods: USVT [Chatterjee '15], User-NN, Time-NN, DR-NN, results for a single time point

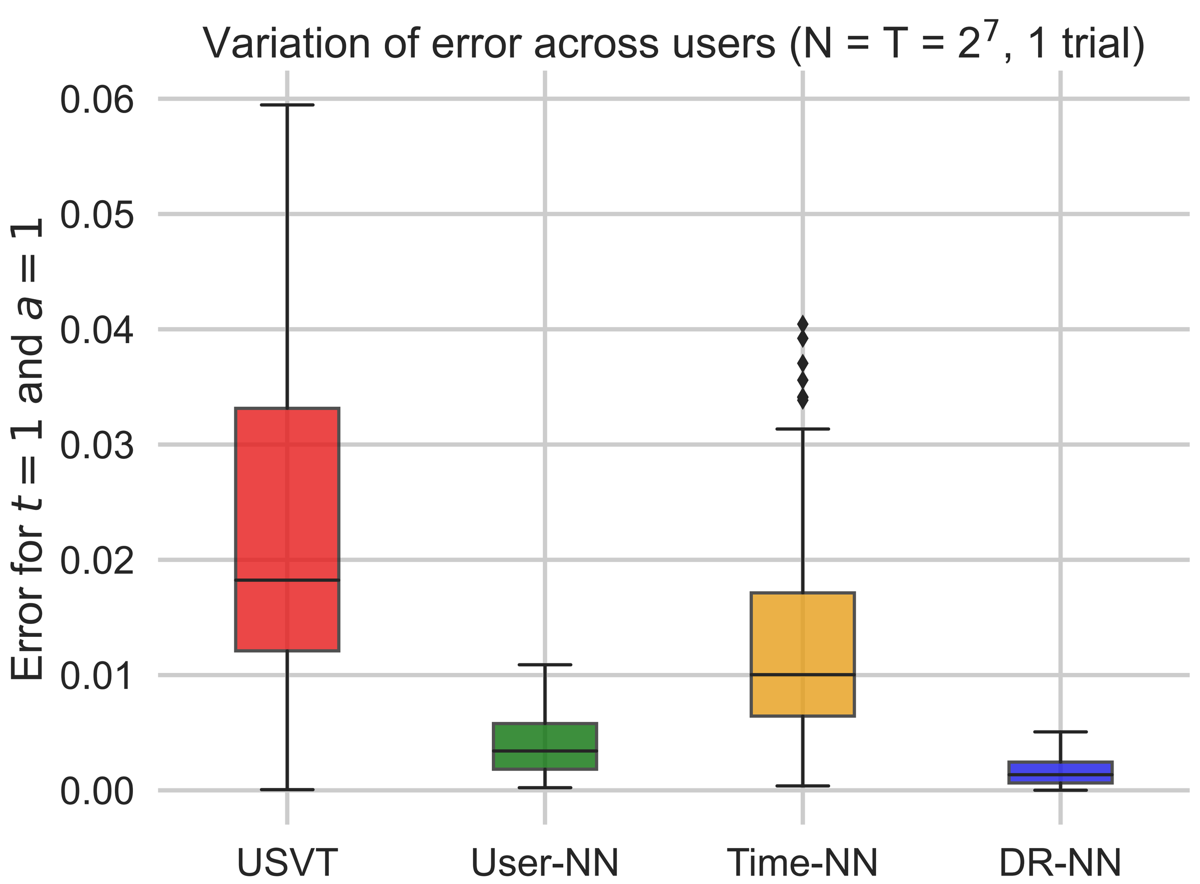

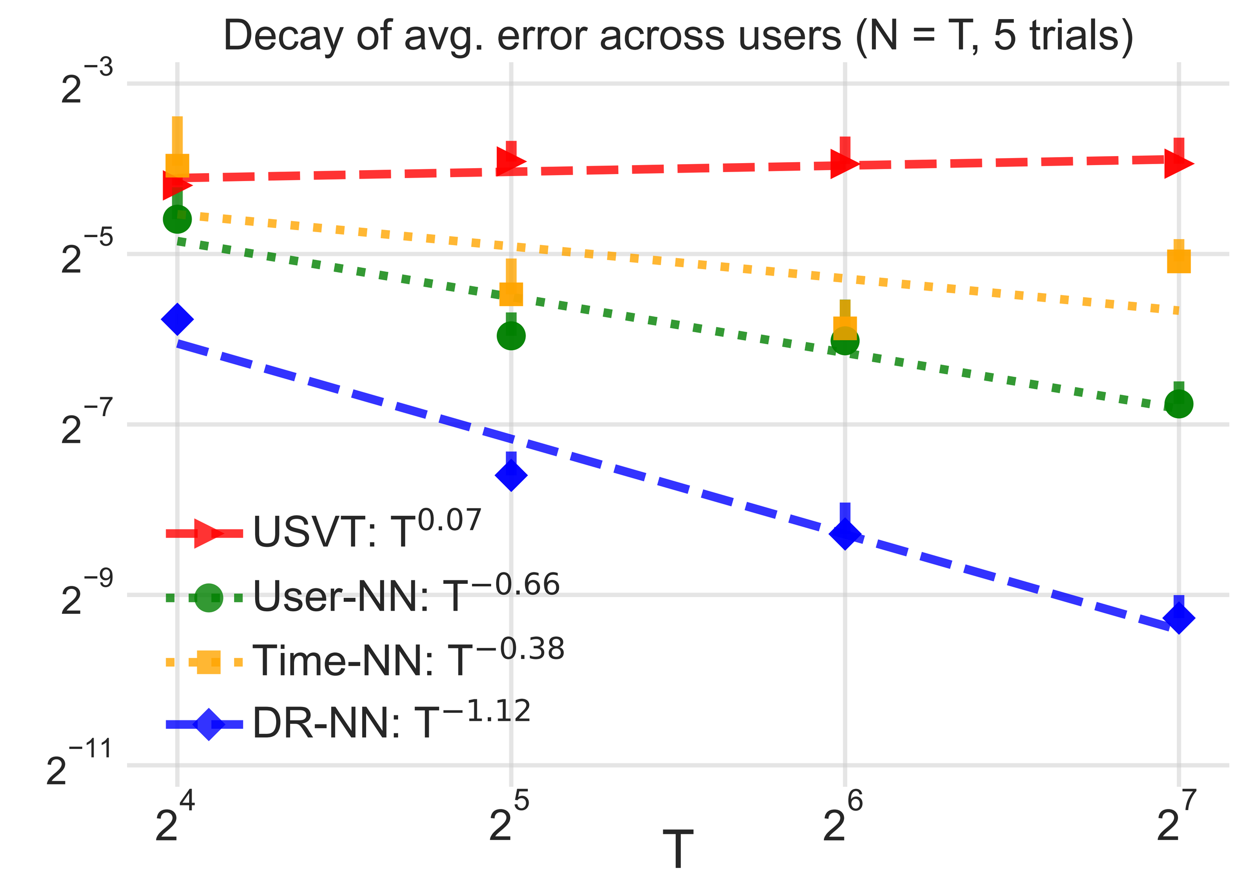

A "doubly robust" approach for nearest neighbors

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

Panel data: $u^\star = u_i, v^\star = v_t$
(dropping 'a' in notation)

User NN: $\hat{u} = \frac{\sum_{j\in\mathcal{N}_i} u_j}{|\mathcal{N}_i|}, \ \hat{v} = v_t$

Time NN: $\hat{u} = u_i, \ \hat{v} = \frac{\sum_{t'\in\mathcal{T}_t} v_{t'}}{|\mathcal{T}_t|}$

Motivating question:

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

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

Doubly robust (DR) NN

$$\widehat{\theta}_{i,t,\text{DR}}^{(1)} = \frac{\sum_{j\in\mathcal{N}_i^{(1)}, t'\in\mathcal{T}_t} A_{j,t}A_{i,t'}A_{j,t'}(Y_{j,t} + Y_{i,t'} - Y_{j,t'})}{\sum_{j\in\mathcal{N}_i^{(1)}, t'\in\mathcal{T}_t} A_{j,t}A_{i,t'}A_{j,t'}}$$

Guarantee: For *non-adaptive policies*, with suitable hyper-parameters and regularity conditions, w.h.p.,

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

(M distinct types of user & time latent factors)

–**a quadratic improvement.** More generally,

Err(DR-NN) ≲ min{ Err(User-NN), Err(Time-NN) }

HeartSteps: A mobile health study with Thompson sampling

- 3 months sequentially adaptive trial with 91 users: Mobile notifications sent (a=1) or not (a=0), 5 times a day (T=450) using Thompson sampling independently for each user, outcome= = log of step-count in the 30-minute window after the decision time
- Results for a = 0 at 50 held out decision times