Counterfactual Inference in Sequential Experimental Design

Preprint: https://arxiv.org/abs/2202.06891


Raaz Dwivedi
raaz@mit.edu
MIT HARVARD UNIVERSITY


Susan Murphy
samurphy@fas.harvard.edu
HARVARD UNIVERSITY


Devavrat Shah
devavrat@mit.edu
MIT

Sequential decision making problems

- Online education: Enhance teaching strategies for better learning
- Online advertising: Update ads / placements to increase revenue
- **Mobile health**: Personalized app messages to promote healthy behavior

Physical activity

Wearable/trackers

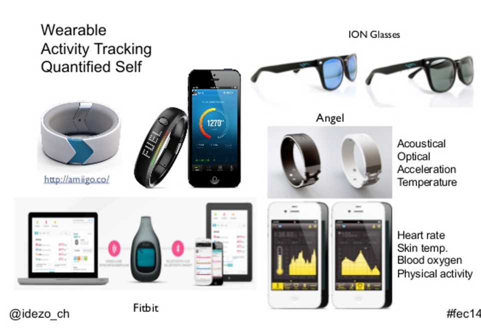

Image credits: Susan Murphy

Mobile health trial: A simplified but representative set-up

$Y_{i,t}^{(a)}$ = potential outcome of user i at time t under treatment $a \in \{0,1\}$

For time $t = 1,2,\ldots,T$

- Neyman-Rubin potential outcome framework
- **No spill-over** of treatment on future outcomes

For user $i = 1,\ldots,N$

1. Randomly send a notification ($A_{i,t} = 1$) or not ($A_{i,t} = 0$) using a sampling policy $\pi_t \in [0,1]^N$ with $\pi_t(i) = \mathbb{P}(A_{i,t} = 1 \mid \text{history}_{t-1})$
2. Observe outcome $Z_{i,t} = Y_{i,t}^{(A_{i,t})} + \varepsilon_{i,t}$ ⟵ noise (e.g., due to noisy measurements from sensors)

Update policy to π_{t+1} using all the data till time t

Sequential experimental design: After-study inference

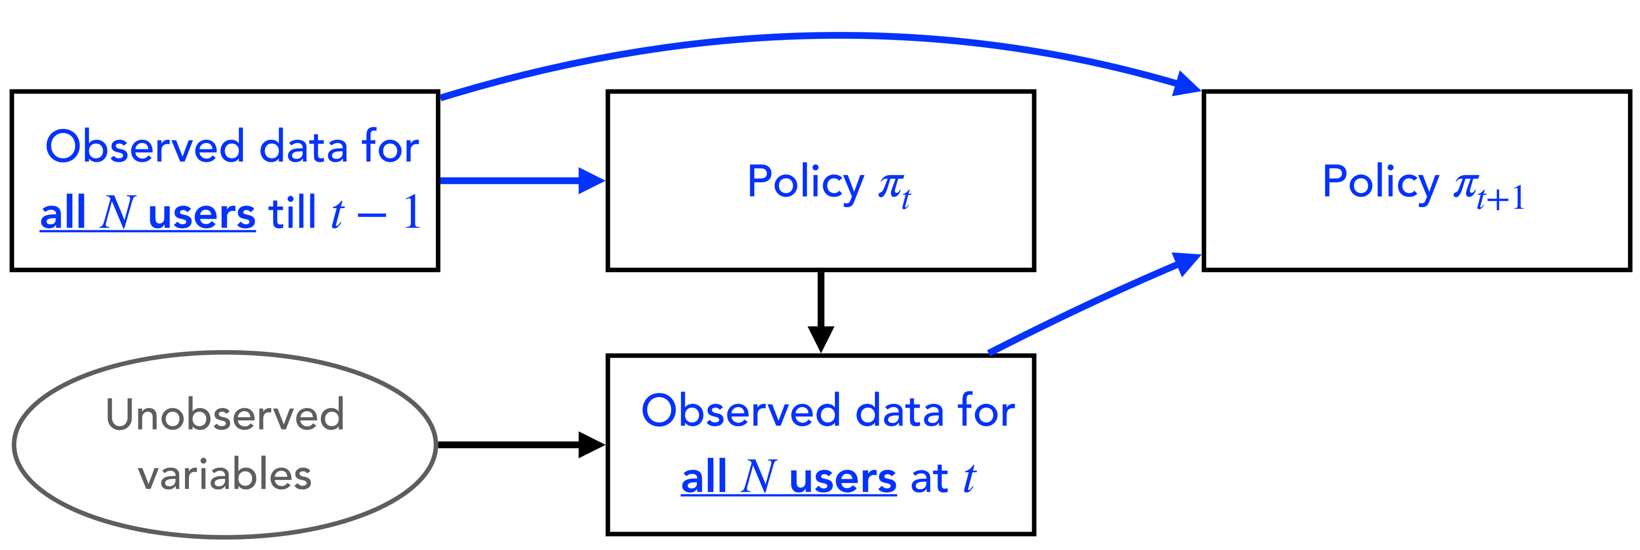

- **Adaptive** policy to personalize to users
- Data **pooled across users during study** learn a good policy quickly
- After-study questions: **Was the treatment effective? On average? Heterogeneity across users? …**

Our goal: Counterfactual inference, i.e., estimation of all missing potential outcomes—hard due to **heterogeneity across users, time and treatments**

Adaptivity and pooling of data across N users for sequential policy design makes after-study inference even more challenging

- can be used for generic after-study analyses, e.g., individual treatment effect $Y_{i,t}^{(1)} - Y_{i,t}^{(0)}$

So what do we do?

Reduce counterfactual inference to sequential matrix completion

- Fix treatment say 1, with $Y_{i,t}^{\star} \triangleq Y_{i,t}^{(1)}$
- $N \times T$ matrix of potential outcomes with missing at random entries

$$Z_{i,t} = \begin{cases} Y_{i,t}^{\star} + \varepsilon_{i,t} & \text{if} \quad A_{i,t} = 1 \\ \text{unknown} & \text{if} \quad A_{i,t} = 0 \end{cases} \quad \text{where } A_{i,t} = \text{Bernoulli}(\pi_t(i))$$

- **New goal**: Estimate missing entries $Y_{i,t}^{\star}$ (separately for each treatment)

policy can depend on observed outcomes of all treatments

Estimating $Y_{i,t}^{\star}$ via a variant of nearest neighbors

$\widehat{Y}_{i,t}$ = Average($Y_{j,t}$) over **available** users j with $\rho_t(i,j) \le \eta$ for suitable threshold η

Will such a non-parametric estimate that is policy and data agnostic work for sequential experimental design?

Yes, under a latent factor model! and suitable conditions

Latent factor model

$Y_{i,t}^{\star} \triangleq Y_{i,t}^{(1)} \triangleq f(u_i, v_t)$

u_i : latent factor for user i
v_t : latent factor for time t
f : unknown (non) linear function

Examples include:

- $Y_{i,\cdot}^{\star} \sim$ Gaussian process with covariance kernel $\mathbf{k}$

u_i = Gaussian vector
v_t = Eigenfunctions of $\mathbf{k}$
$f(u,v) = \langle u, v \rangle$

- Also a sub-class of exchangeable data

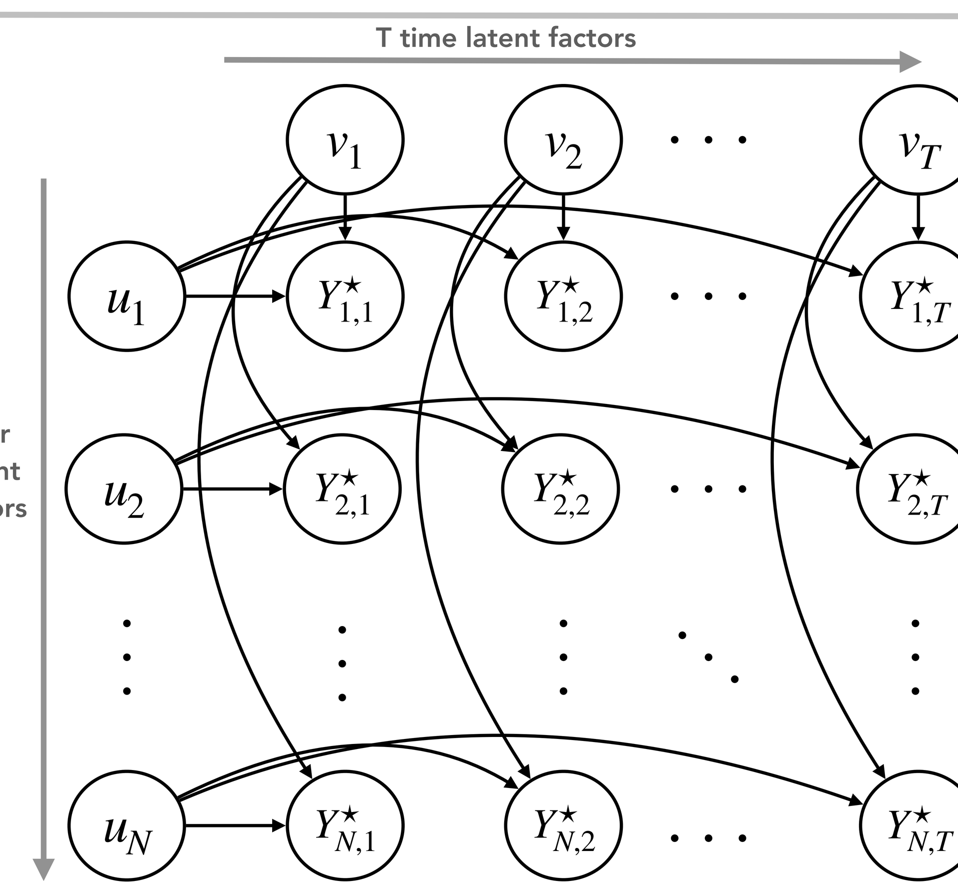

$f(u,v) = \langle u, v \rangle$: A distribution-free consistency for any (i,t)

Consider any (i,t) with enough nearest neighbors j satisfying the conditions

- **``diverse latent-time factors''**: $\frac{1}{T_{t,i,j}} \sum_{t' \in T_{t,i,j}} v_{t'} v_{t'}^{\top} \succeq \lambda I_d$ for $\lambda > 0$

$T_{t,i,j}$ = commonly observed time points other than t; used to compute distance $\rho_t(i,j)$

- **``non-adversarial noise''** across $T_{t,i,j} \cup \{t\}$: behave roughly like iid/mixing process

For suitable scaling of threshold η & mild conditions on arbitrarily dependent policy, given user i with latent factor u_i, we have

$$\widehat{Y}_{i,t} \to Y_{i,t}^{\star} \text{ for any } t \text{ as } N, T \to \infty$$

Lipschitz f: Central limit theorem for sequential estimation of $Y_{i,T}^{\star}$

Consider a **non-linear Lipschitz** f (with $\|f\|_{\infty} \le D$), and suppose

- $u_j \sim_{iid} \mathbb{P}_{user}$, $v_t \sim_{iid} \mathbb{P}_{time}$,
- $\varepsilon_{j,t} \sim_{iid} \mathbb{P}_{noise}$, $\mathbb{E}[\varepsilon_{j,t}] = 0$, $\mathbb{E}[\varepsilon_{j,t}^2] = \sigma^2$
- π_t depends on *all users'* history till $t-1$; treatments $\{A_{j,t}\}$ assigned independently given the history

Under regularity conditions, given *any user i at last time T* with *number of neighbors $N_{i,T}$*

$$\sqrt{N_{i,T}}(\widehat{Y}_{i,T} - Y_{i,T}^{\star}) \Longrightarrow \mathcal{N}(0,\sigma^2) \quad \text{as } N, T \to \infty \text{ together at suitable rates}$$

Lipschitz f: Non-asymptotic expected squared error bound

$$\mathbb{E}\left[(\widehat{Y}_{i,T} - Y_{i,T}^{\star})^2 \mid u_i\right] \precsim \underbrace{(\eta - 2\sigma^2)}_{\text{Bias due to } \eta} + \underbrace{\frac{D^2(1+\gamma_{i,T})}{p_{\min,T}^2 \sqrt{T-1}}}_{\text{Concentration of neighbor distance}} + \underbrace{\frac{\sigma^2}{p_{\min,T} \Phi_i N}}_{\text{Effective noise variance}}$$

$p_{\min,T} = \min_{t,j} \pi_t(j)$ min probability of sampling any entry

$\Phi_i = \mathbb{P}_u\left(\mathbb{E}_v[f(u_i,v) - f(u,v)^2] \le \eta/2 - \sigma^2\right)$ Probability of sampling a nearest neighbor

$$\gamma_{i,T} = \sup_{j \ne i, t < T} \left| \mathbb{E}[\Sigma_{t'=t+1}^{T-1} A_{i,t'} A_{j,t'} \mid \text{history}_t] - \mathbb{E}[\Sigma_{t'=t+1}^{T-1} A_{i,t'} A_{j,t'} \mid \text{history}_{t-1}] \right|$$

Cumulative future dependency of adaptive policies on one column