Counterfactual inference in sequential experimental design

Raaz Dwivedi

https://rzrsk.github.io

Joint work with

Susan Murphy

Harvard John A. Paulson School of Engineering and Applied Sciences

Devavrat Shah

MIT EECS

Statistical Machine Learning Session, IMS Annual Meeting

June 30, 2022

Sequential decision making problems

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

Physical activity

Wearable/trackers

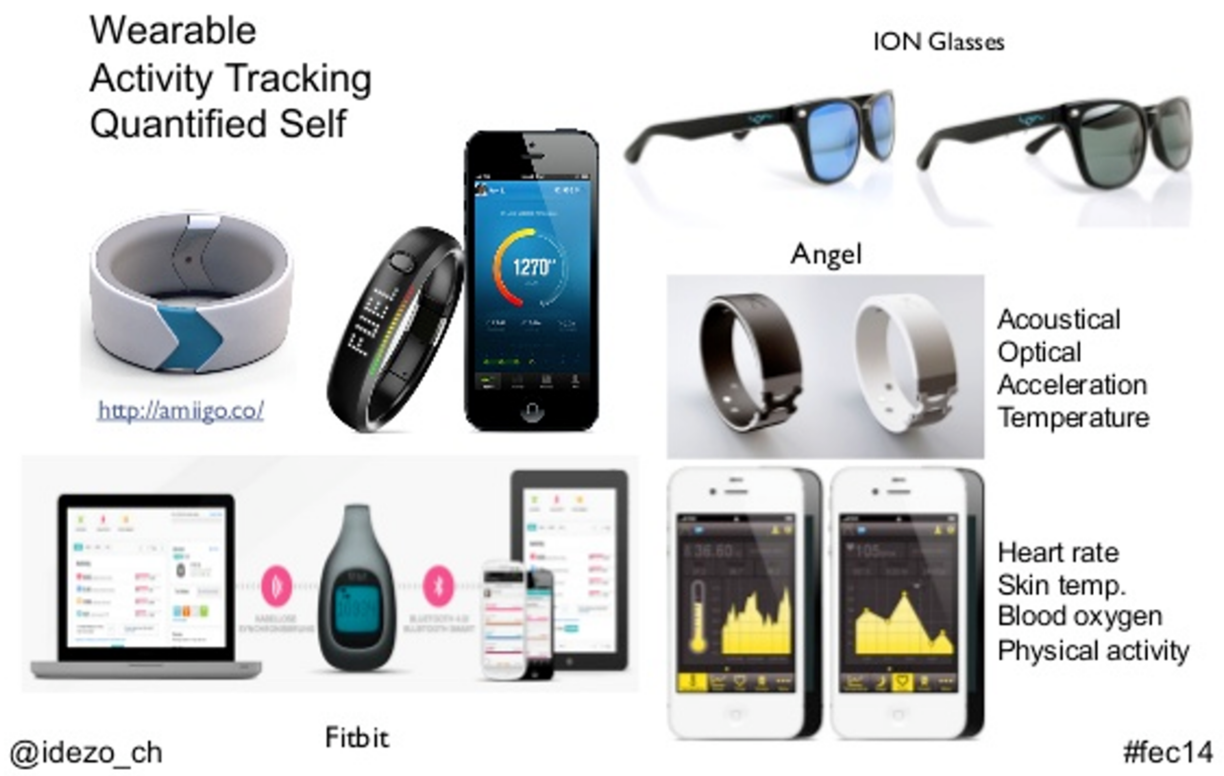

Image credits: Susan Murphy

Mobile health trial: A simplified but representative set-up

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

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

Image credits: apps.garmin.com

Mobile health trial: A simplified but representative set-up

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

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

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

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

1. **Assign** $A_{i,t} \leftarrow \begin{cases} 1 & \text{(notify) with prob. } \pi_{t,i} \\ 0 & \text{(do nothing) with prob. } 1 - \pi_{t,i} \end{cases}$

2. **Observe** $Z_{i,t} = Y_{i,t}^{(A_{i,t})} + \varepsilon_{i,t}$ (step count)

Update policy to $\pi_{t+1} \in [0,1]^N$ using history$_t$

Image credits: apps.garmin.com

Setting overview

This talk: **Sequential experiments**

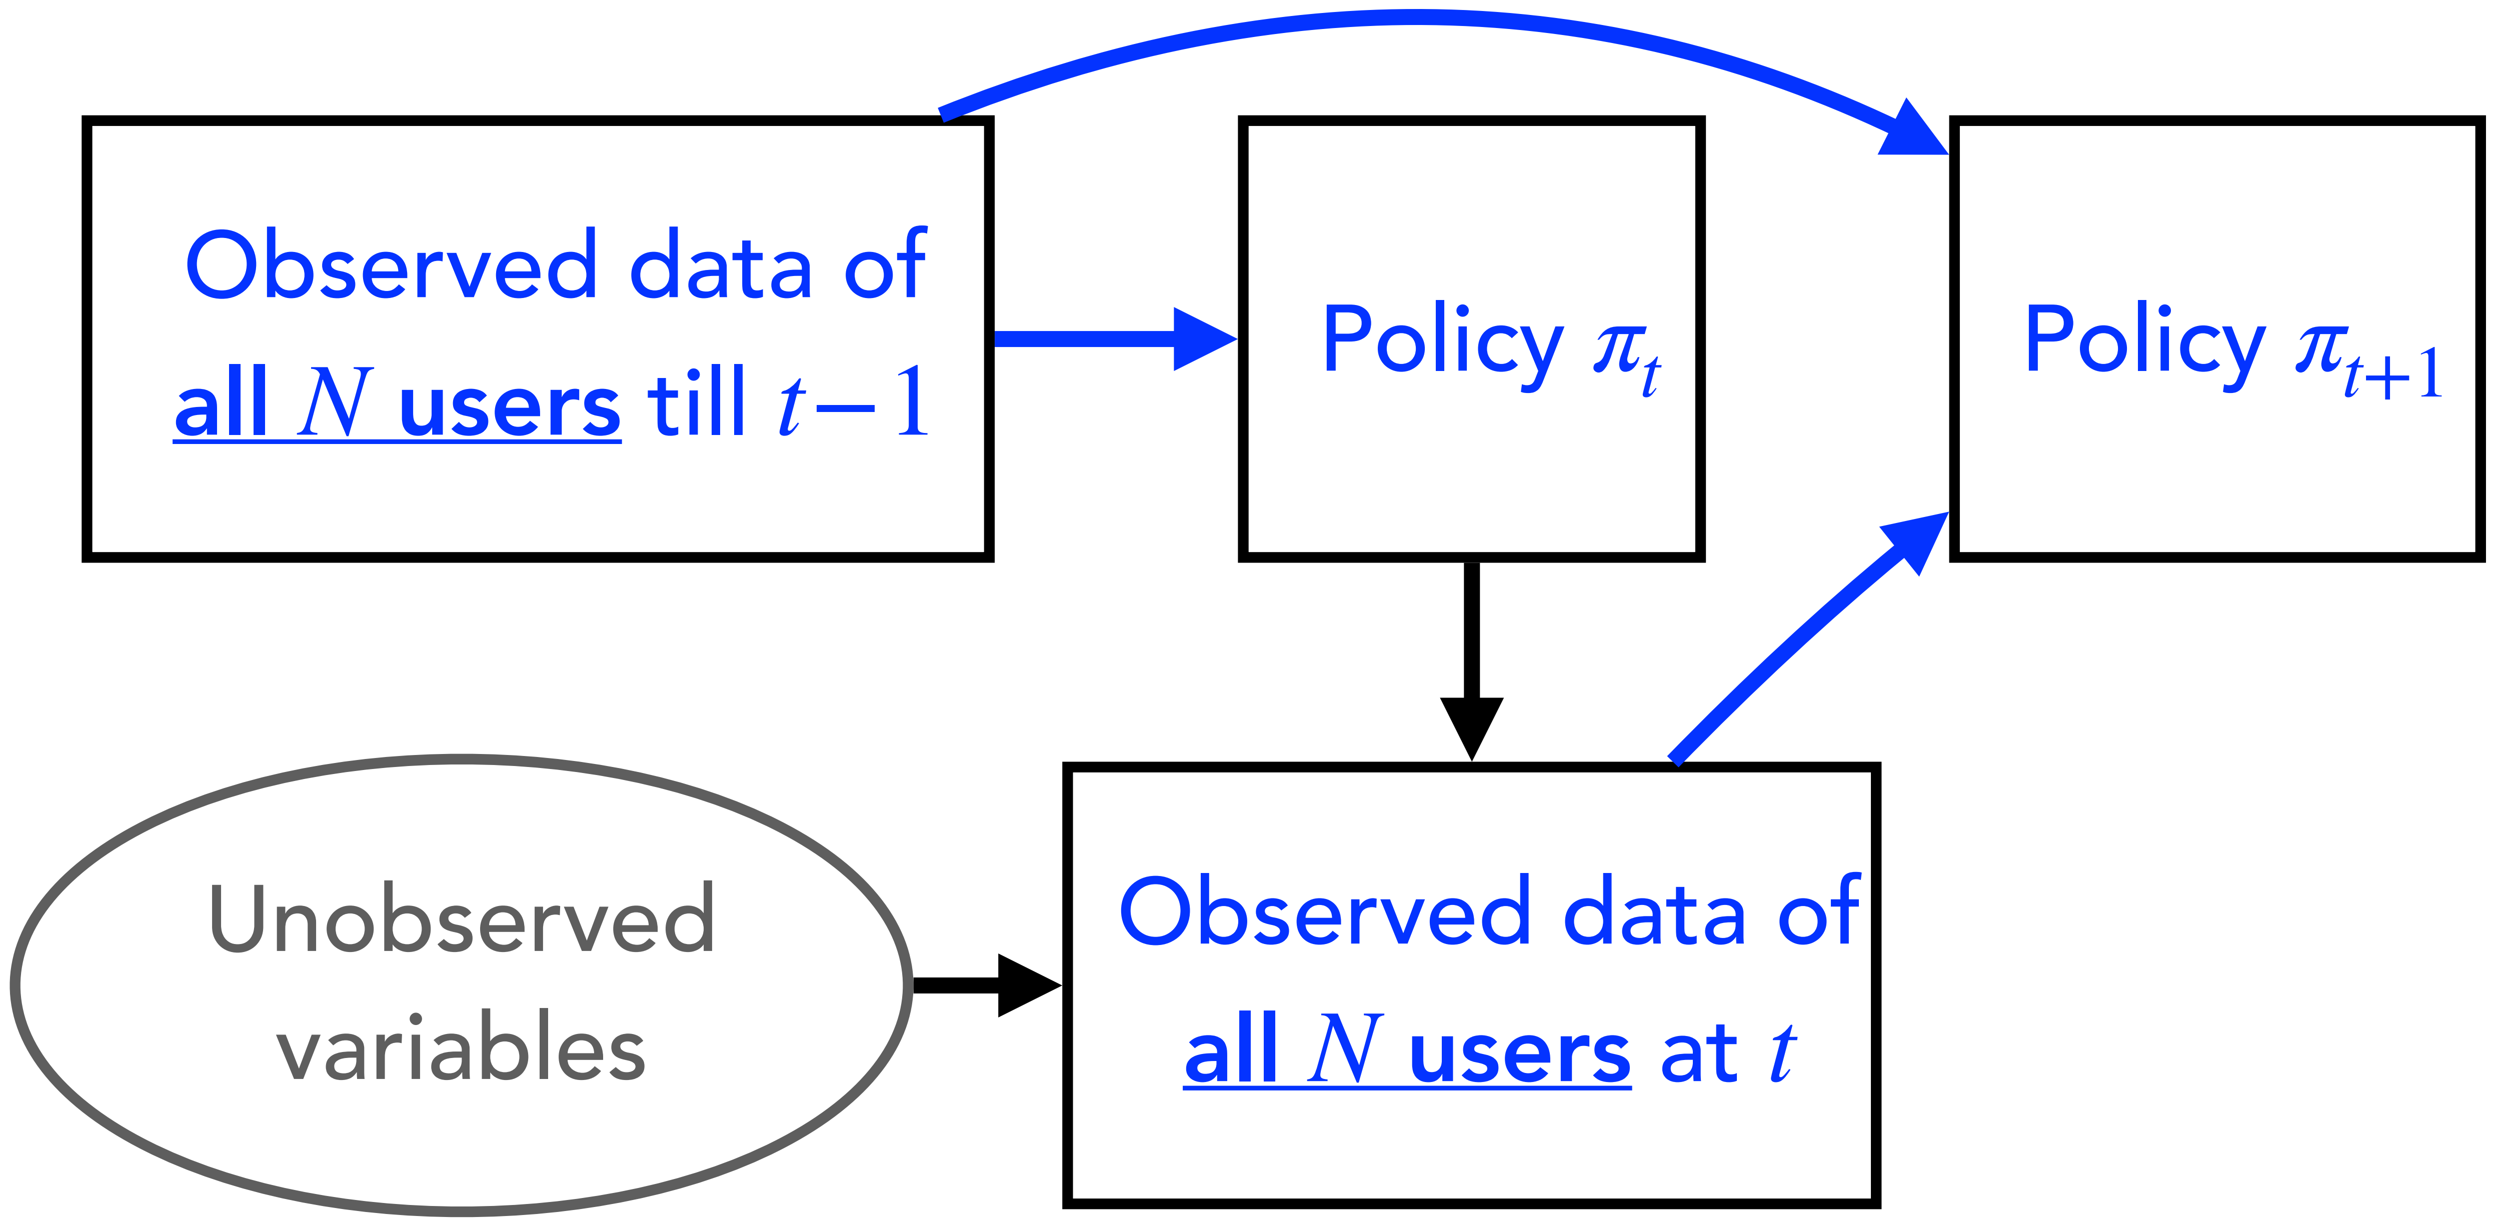

[Yom-Tov '17, Tomkins et al. '20]

Setting overview

This talk: **Sequential experiments**

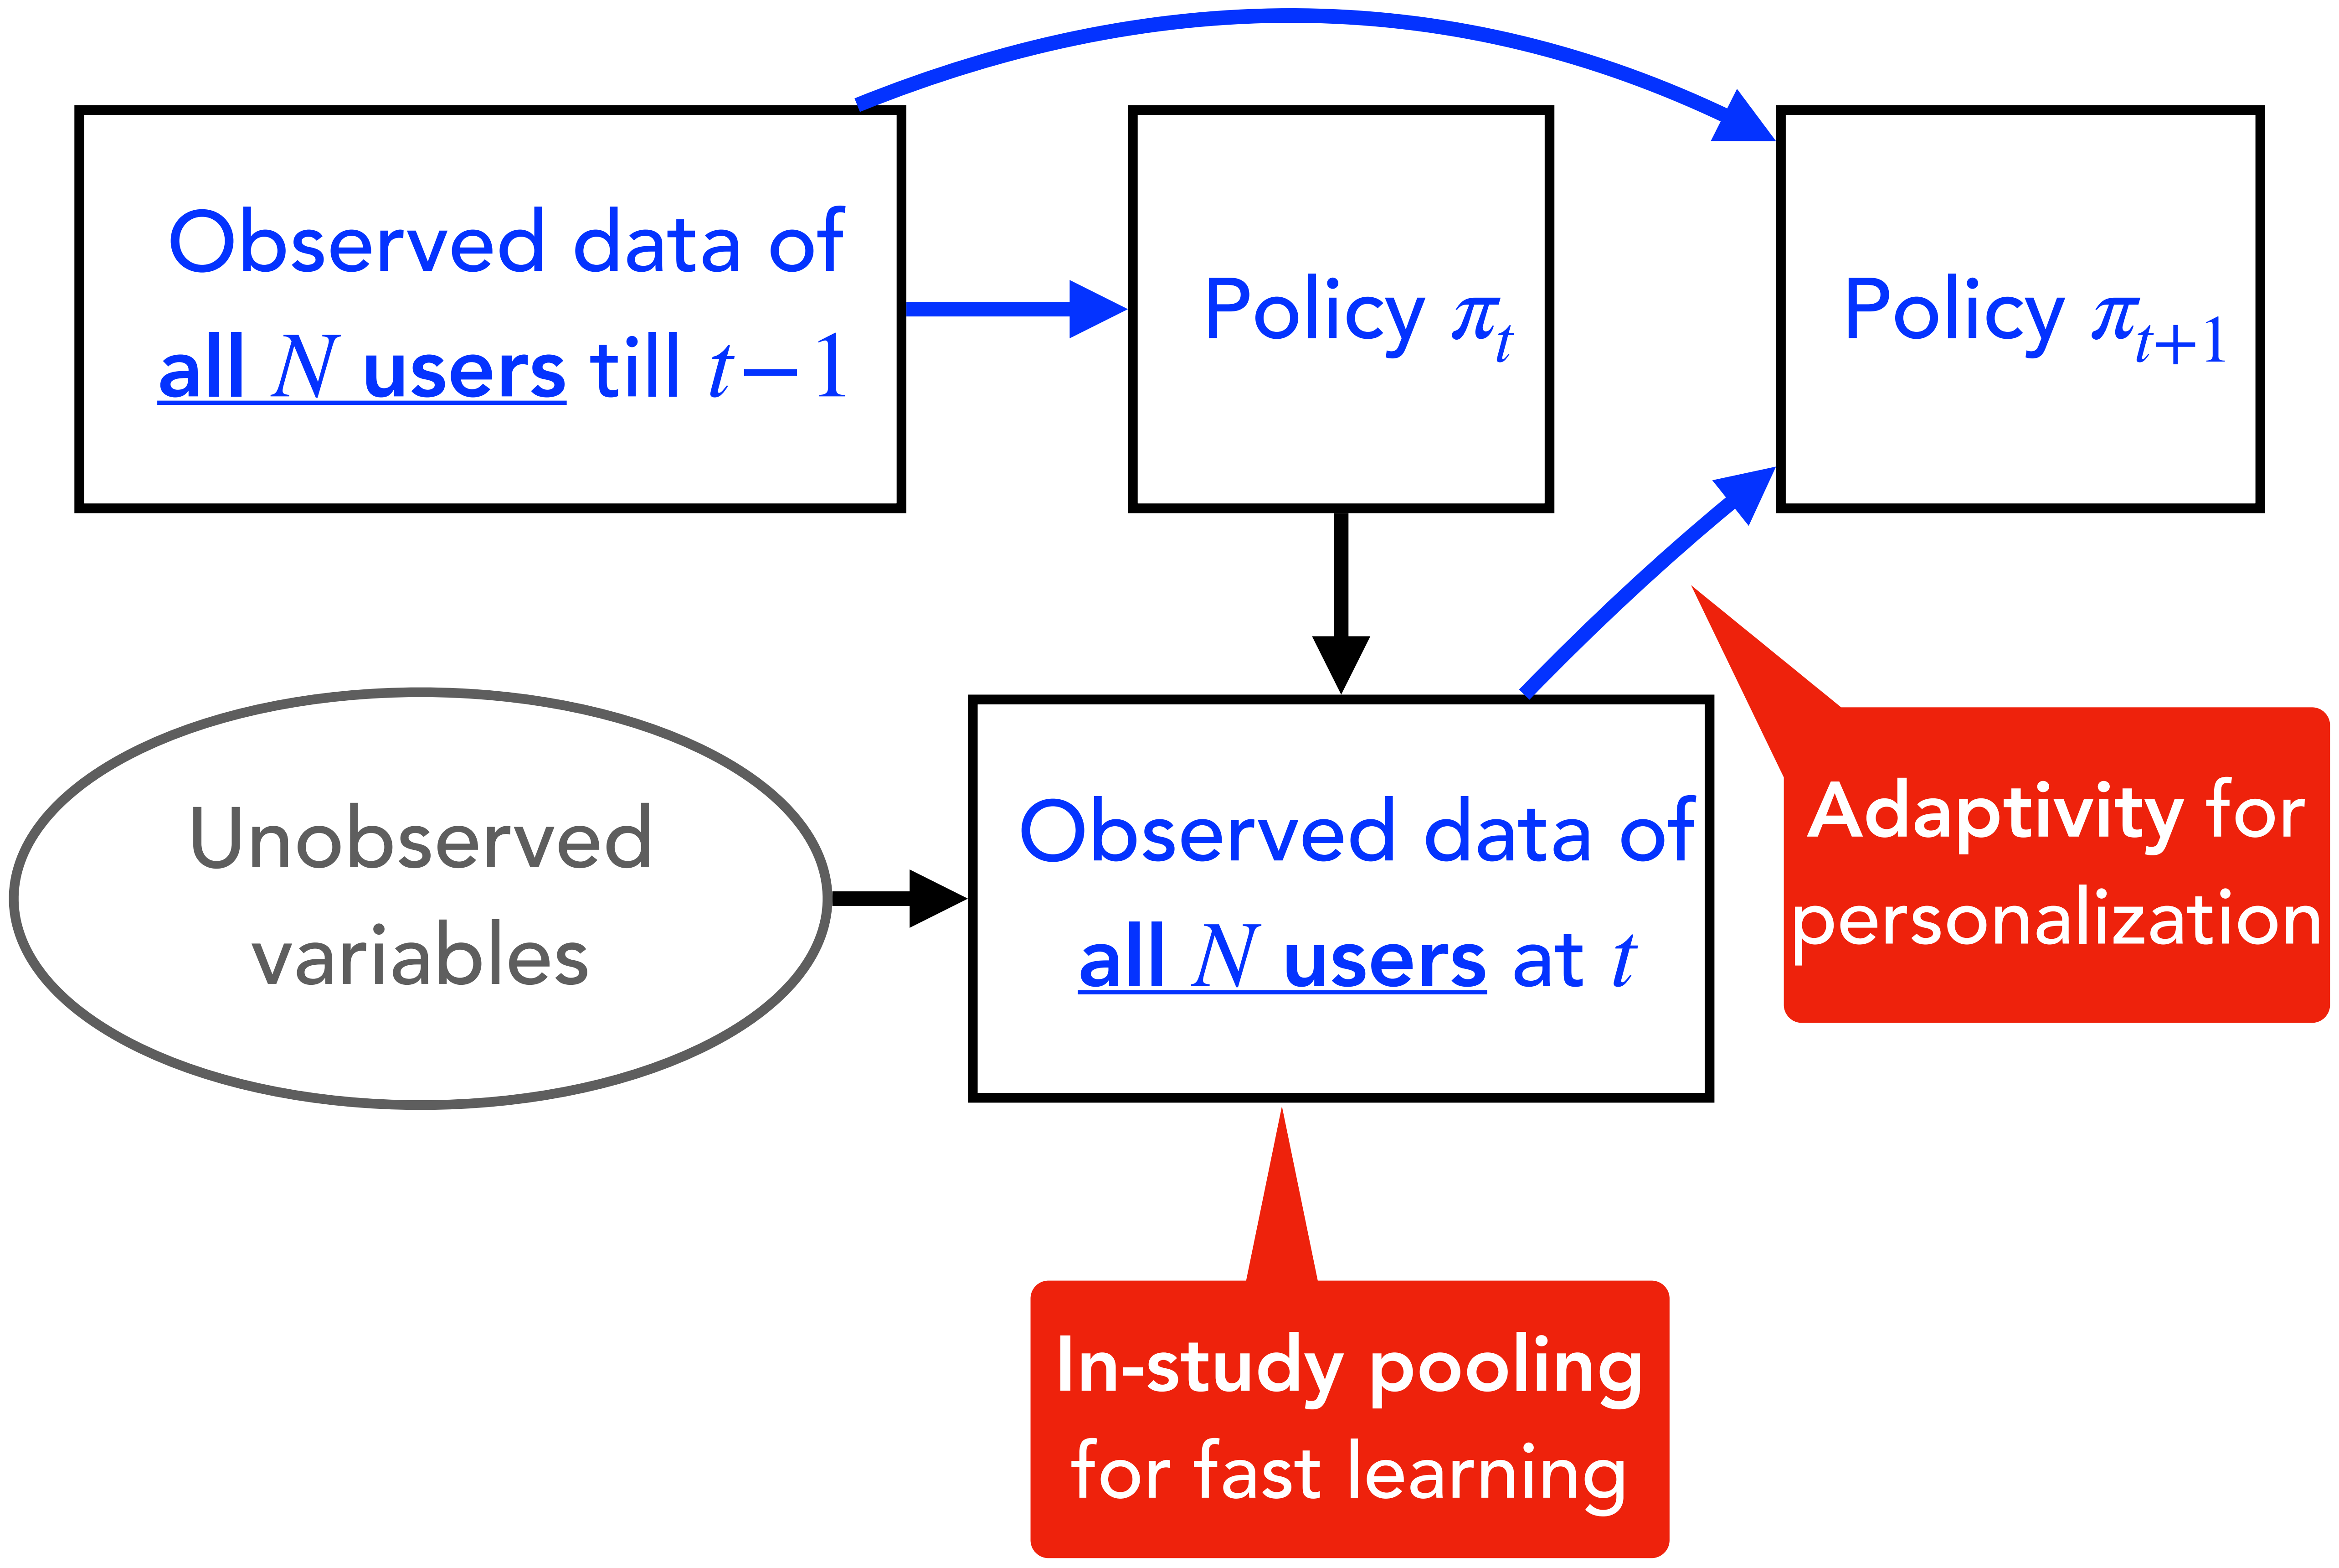

[Yom-Tov '17, Tomkins et al. '20]

Setting overview

This talk: Sequential experiments

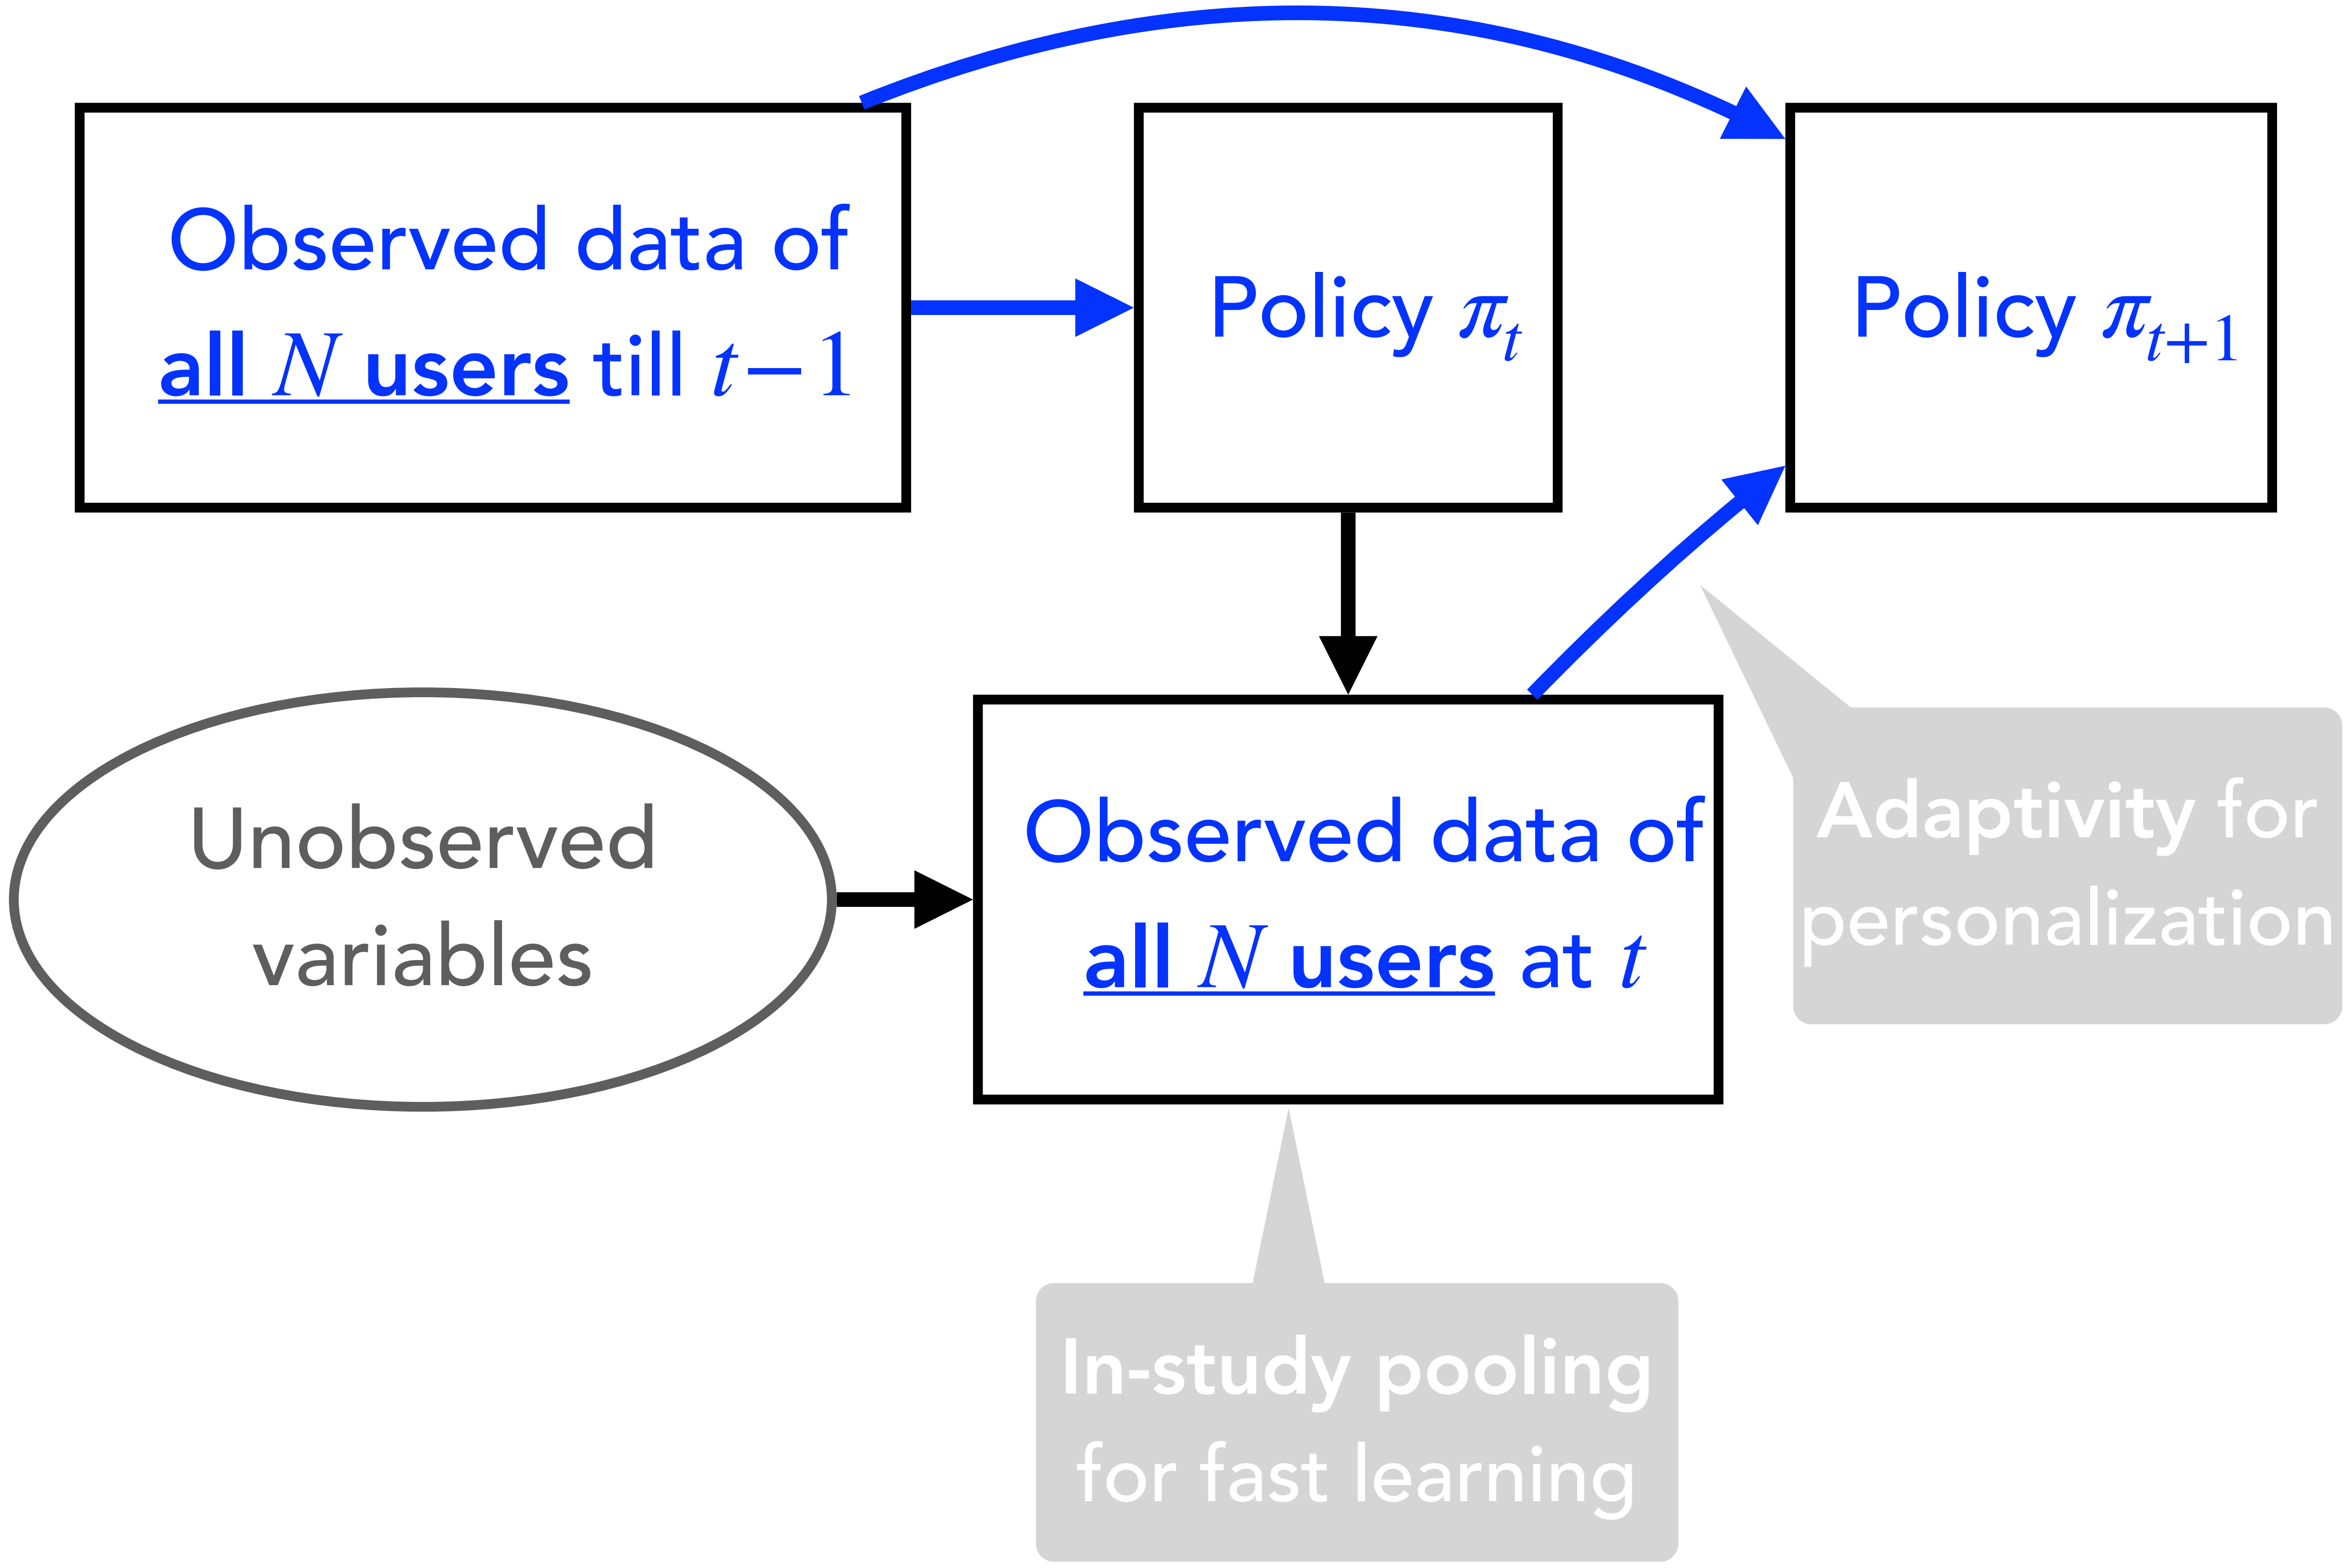

[Yom-Tov ’17, Tomkins et al. ’20]

Prior work

Dynamic treatment regime

- **i.i.d. trajectories across users**
- e.g., personalized clinical treatment

[Robins ’86, Murphy ’03, Bojinov et al. ’21]

Setting overview

This talk: Sequential experiments

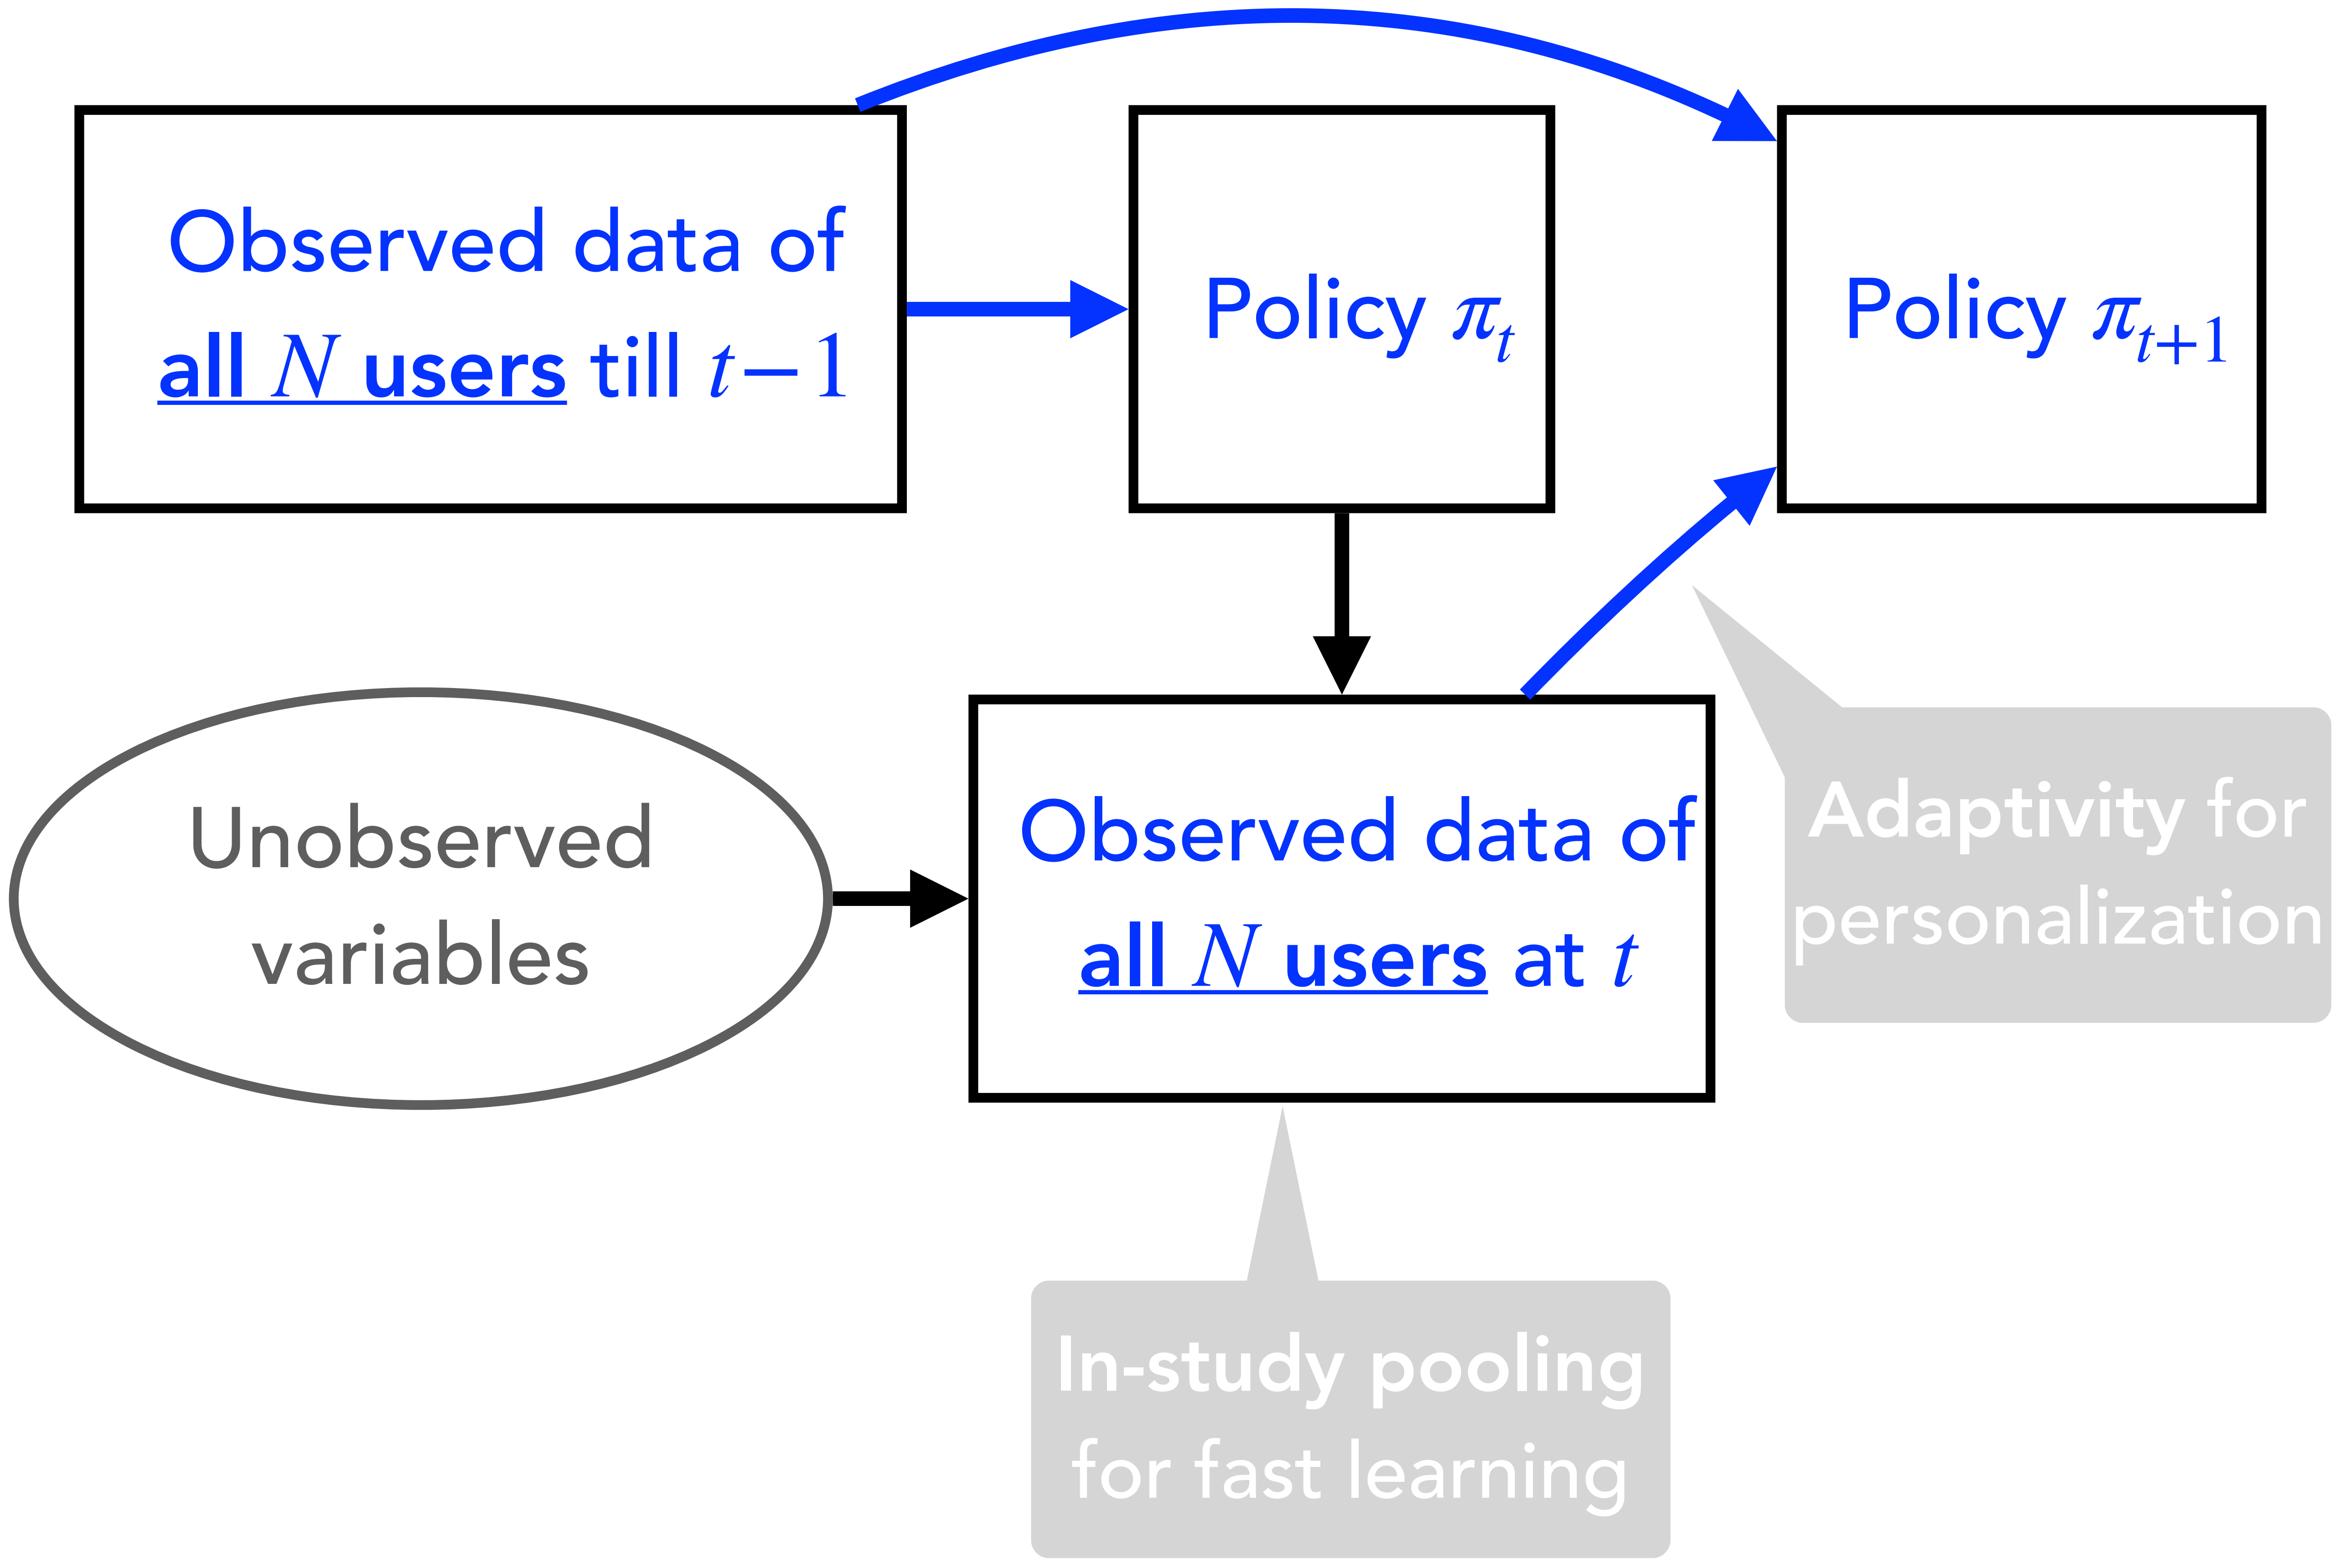

[Yom-Tov '17, Tomkins et al. '20]

Prior work

Dynamic treatment regime

- i.i.d. trajectories across users
- e.g., personalized clinical treatment

[Robins '86, Murphy '03, Bojinov et al. '21]

Policy evaluation for bandits

- **i.i.d. users at each time**
- e.g, online ads

[Zhang et al. '21, Hadad et al. '21, Bibaut et al. '21]

Setting overview

This talk: Sequential experiments

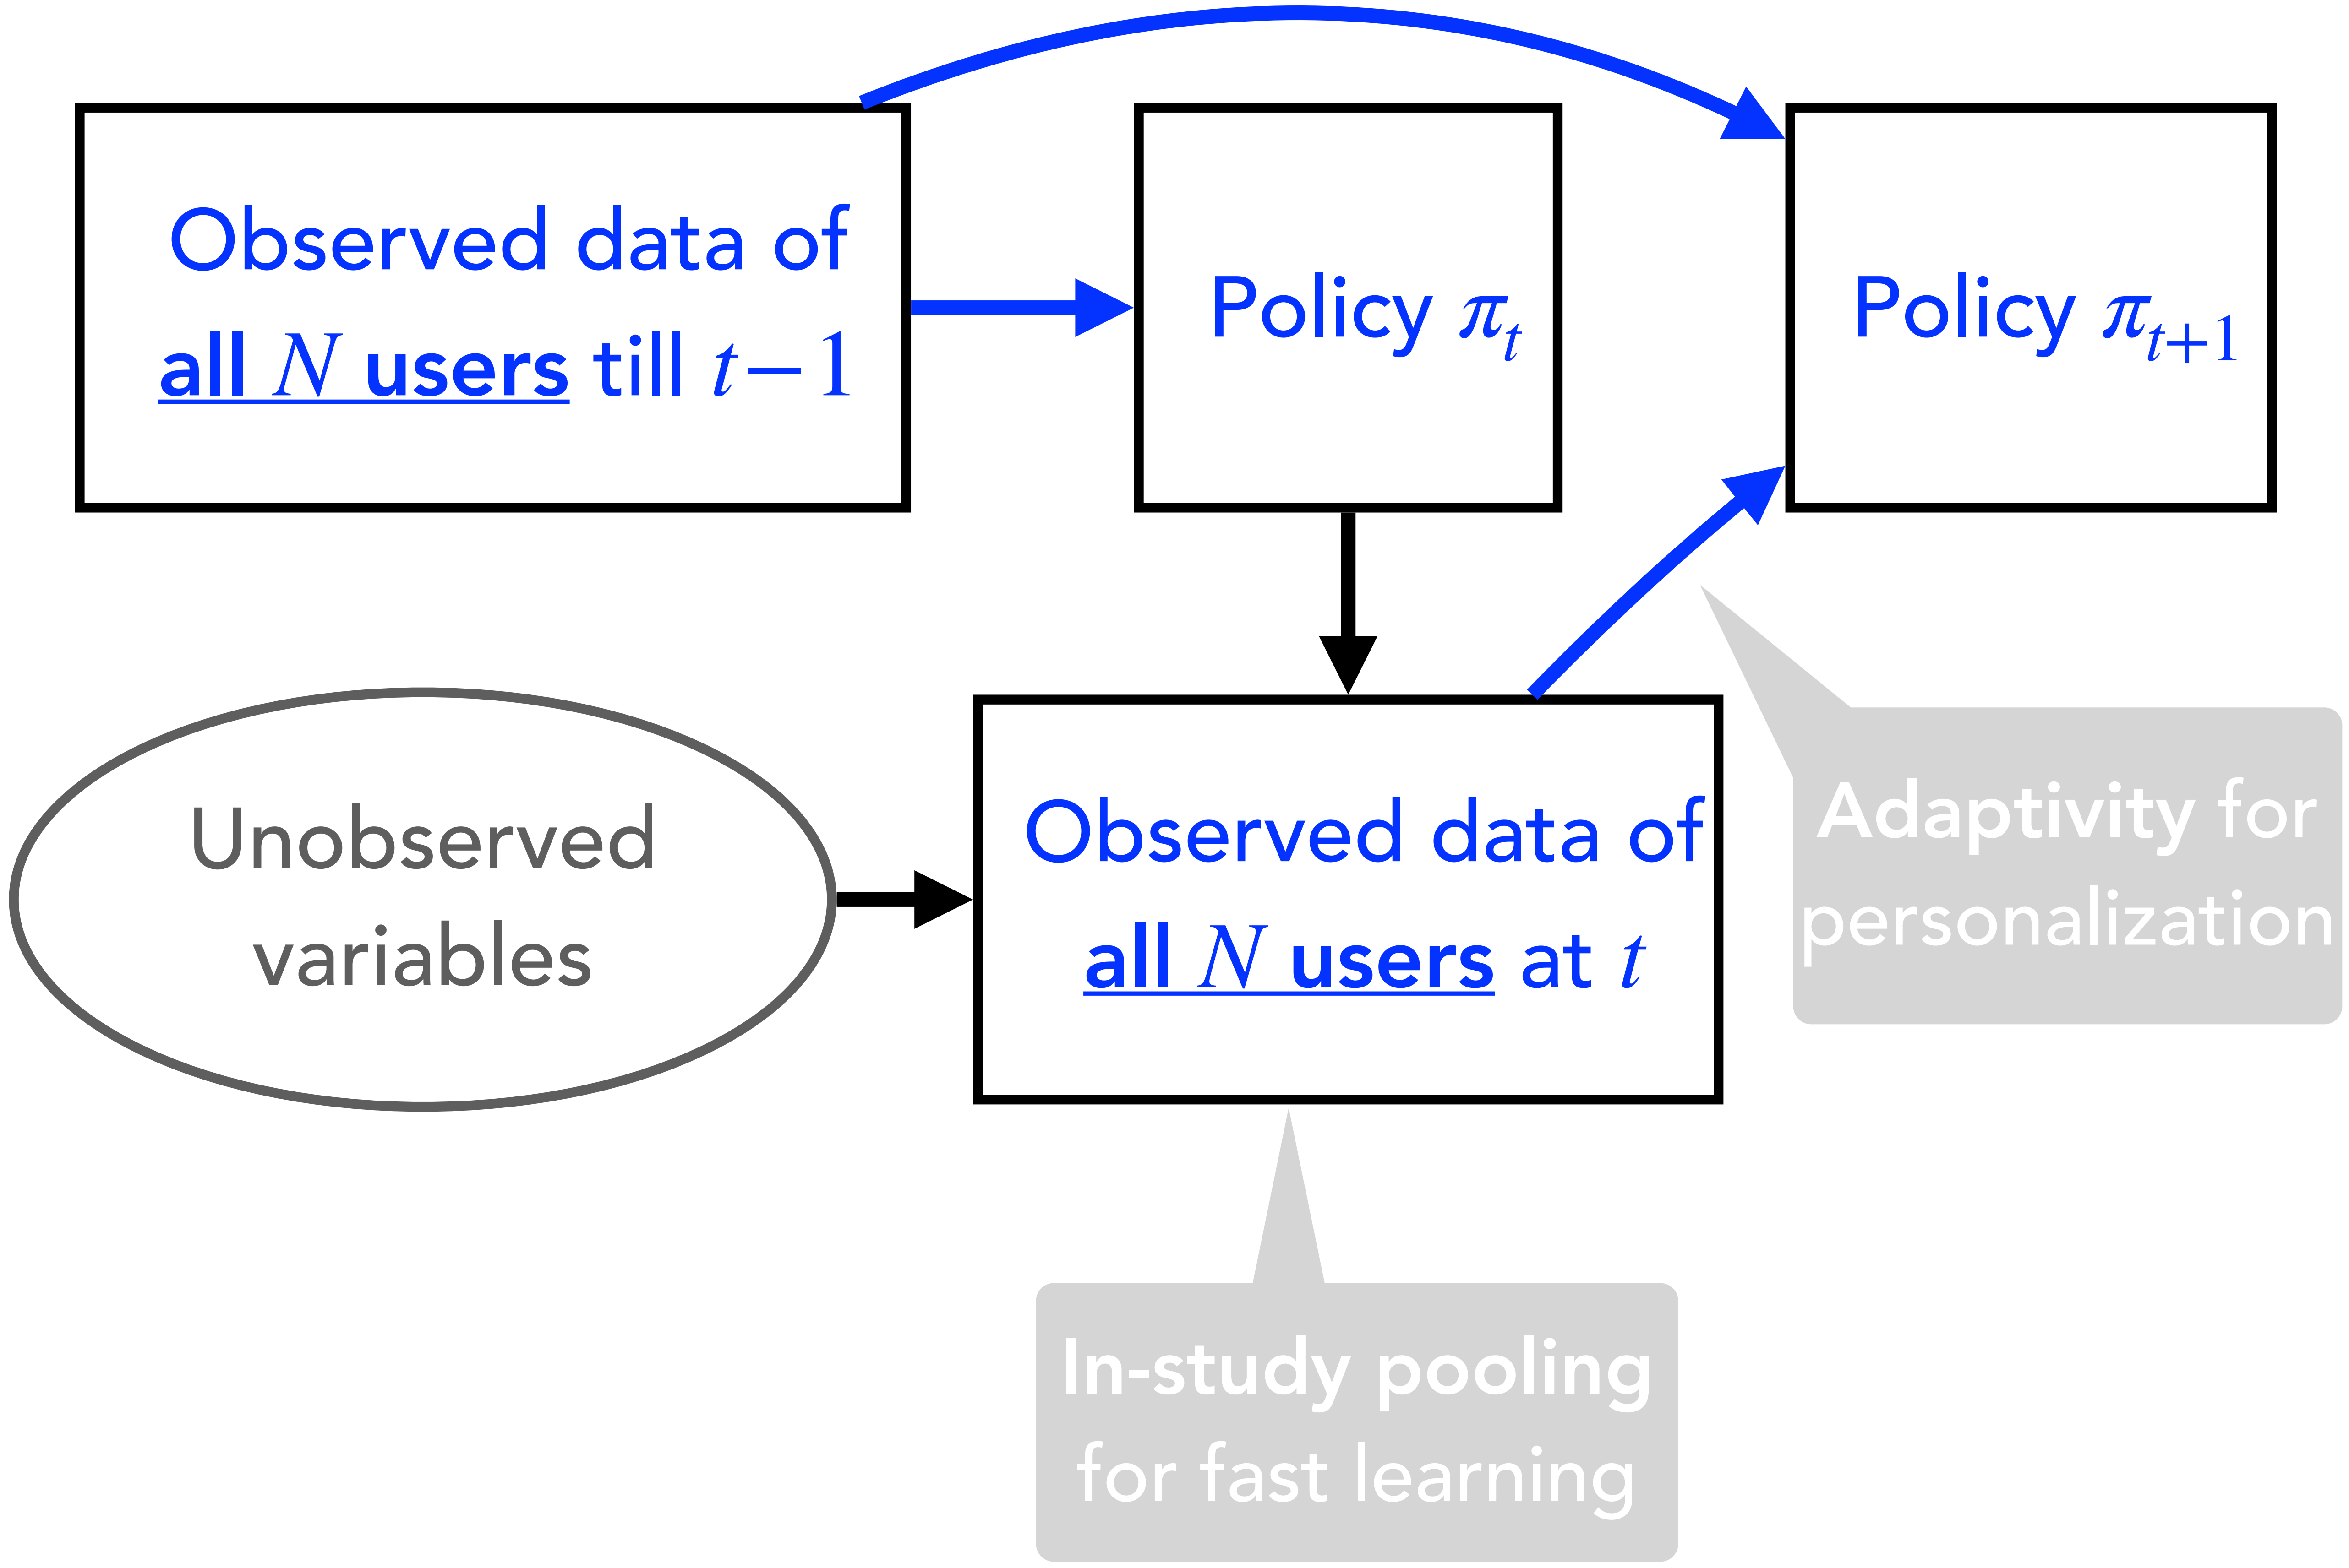

[Yom-Tov '17, Tomkins et al. '20]

Prior work

Dynamic treatment regime

- i.i.d. trajectories across users
- e.g., personalized clinical treatment

[Robins '86, Murphy '03, Bojinov et al. '21]

Policy evaluation for bandits

- i.i.d. users at each time
- e.g, online ads

[Zhang et al. '21, Hadad et al. '21, Bibaut et al. '21]

Causal panel data (observational studies)

- **users treated forever after t_0**
- e.g, law enforcement in california

[Abadie et al. '03, Chernozhukov et al. '17, Athey et al. '18, Agarwal et al. '21]

Goals overview

Prior work

Dynamic treatment regime

Average treatment effect (ATE)

Policy evaluation for bandits

Off-policy evaluation (OPE)

Causal panel data

Average treatment effect (ATE)

Goals overview

This talk:

Counterfactual inference in sequential experiments

user x time-level treatment effect

-allows generic after-study analyses including ATE, OPE

Prior work

Dynamic treatment regime

Average treatment effect (ATE)

Policy evaluation for bandits

Off-policy evaluation (OPE)

Causal panel data

Average treatment effect (ATE)

With stronger goals come stronger assumptions!

Equivalently, stronger responsibilities require stronger assumptions!

Or, great responsibilities requires great power!

Model: Latent factorization

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

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

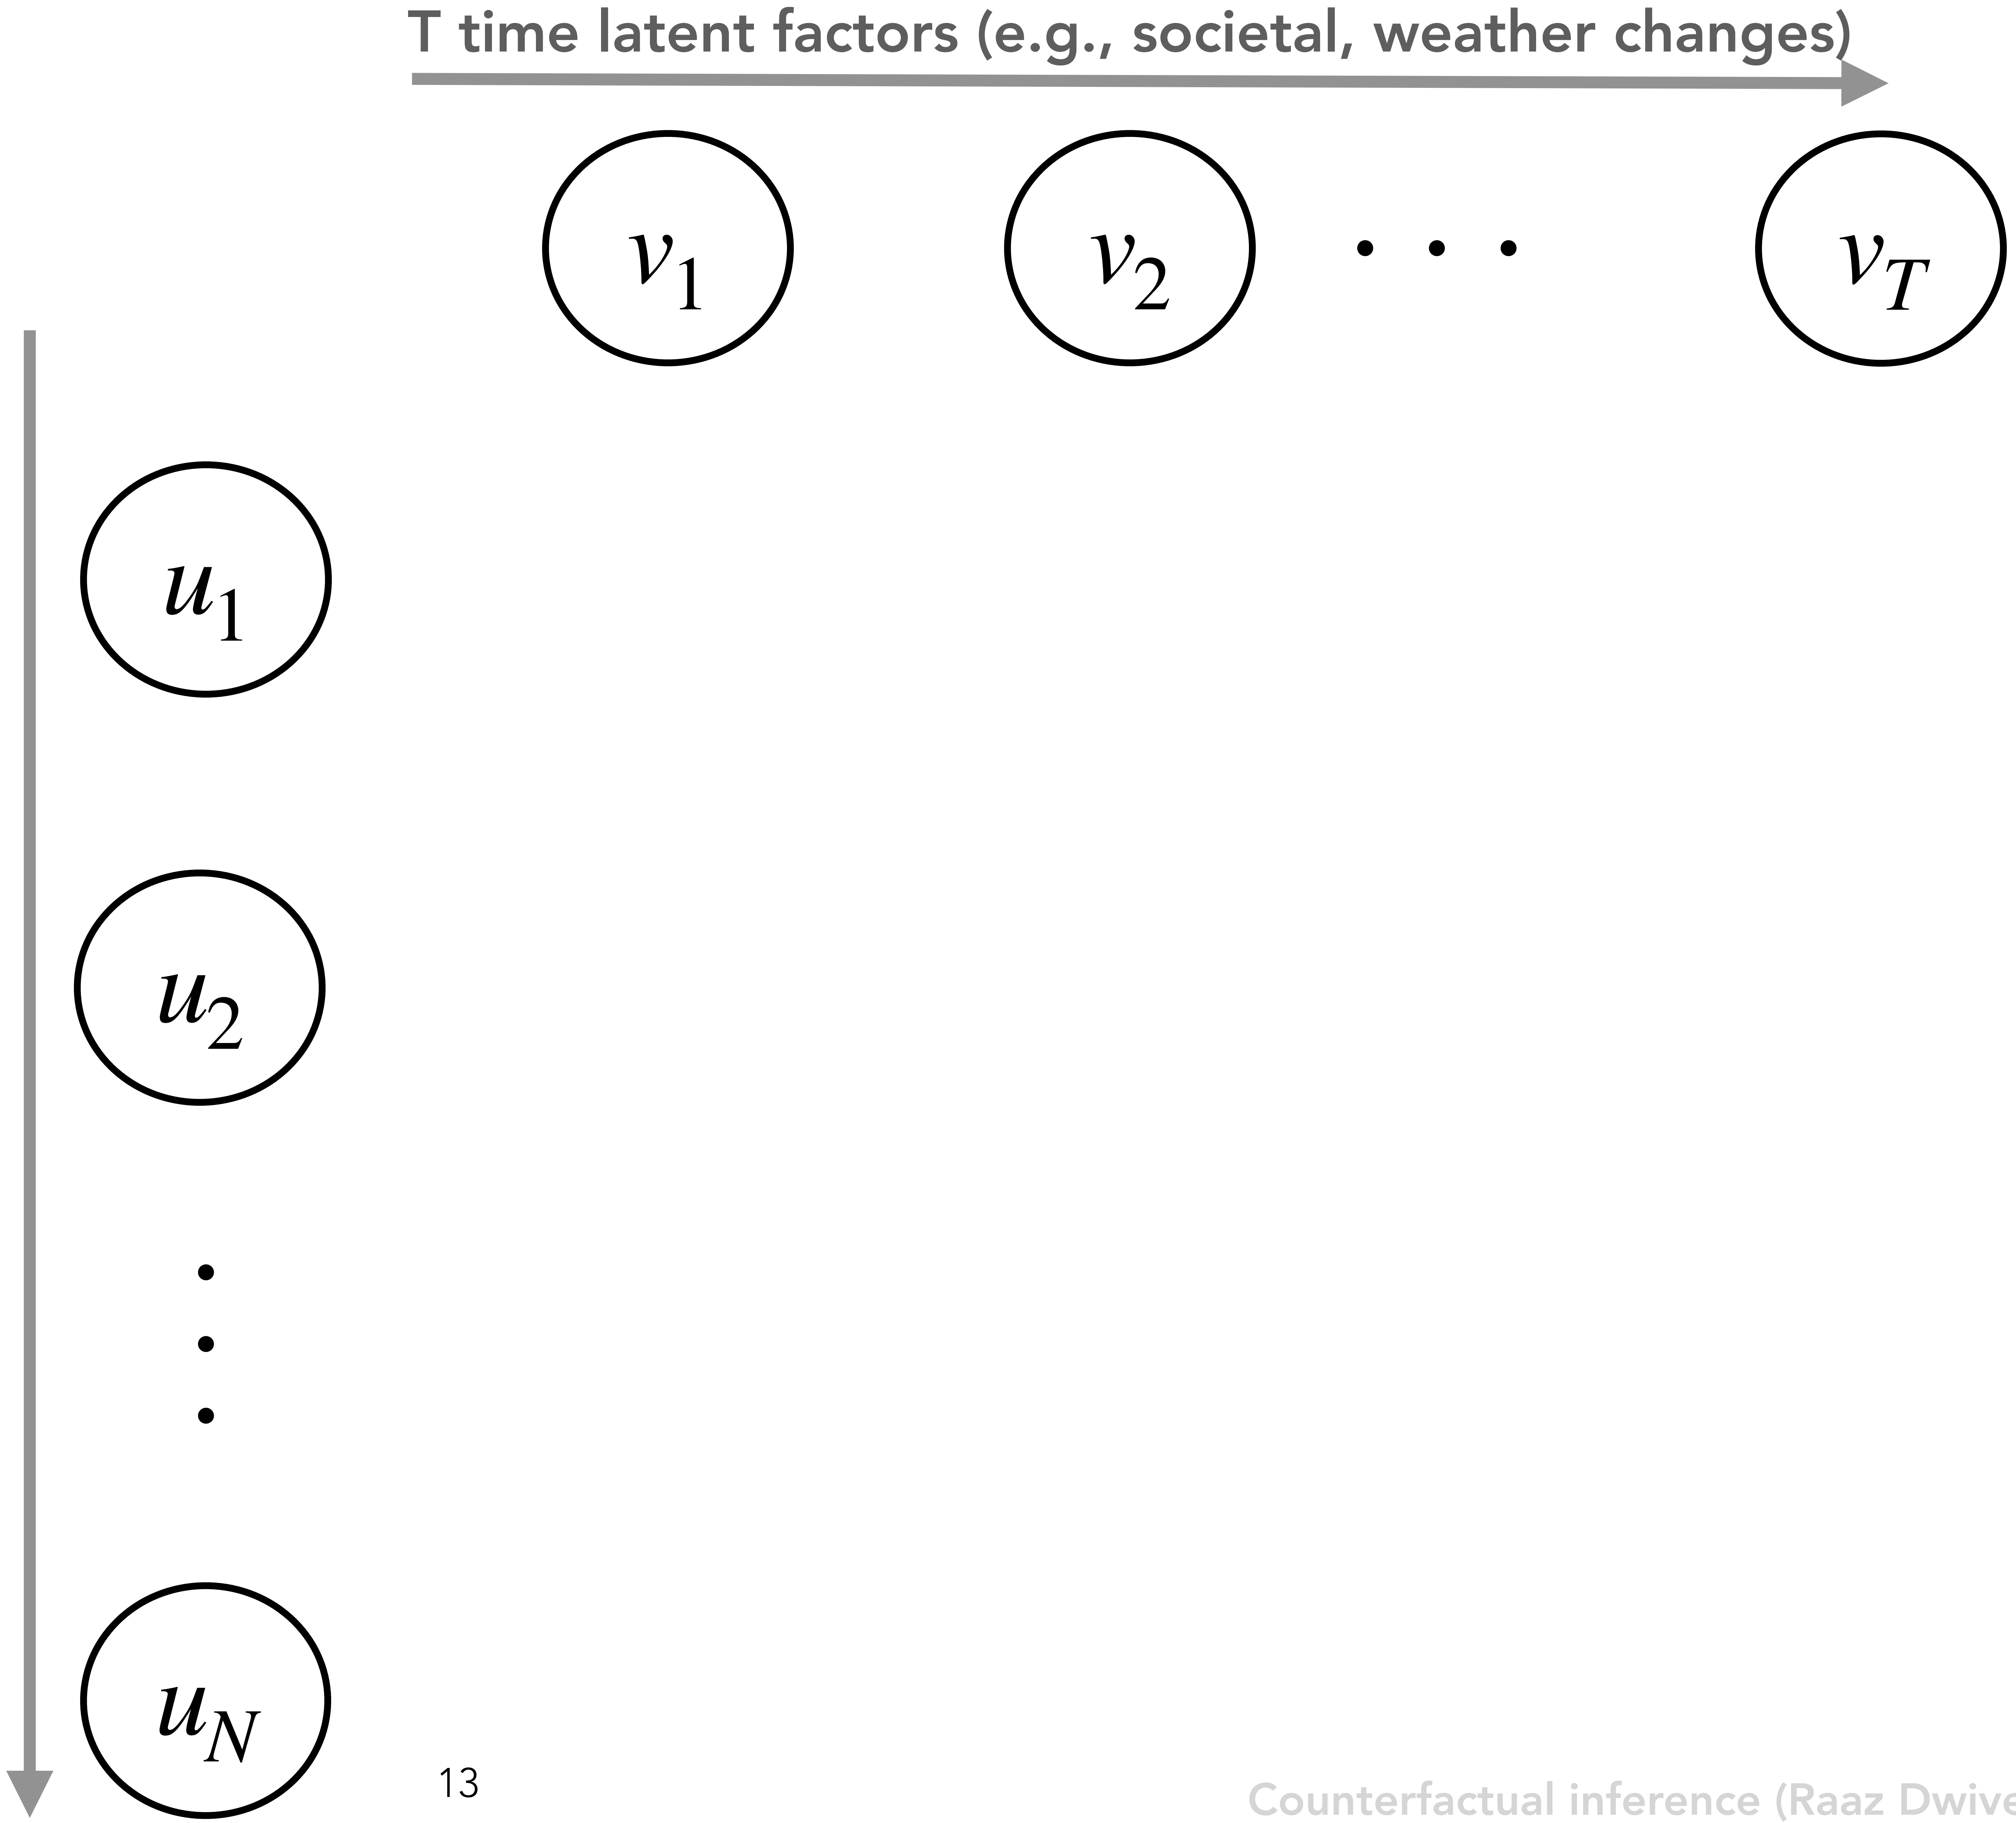

Model: Latent factorization

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

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

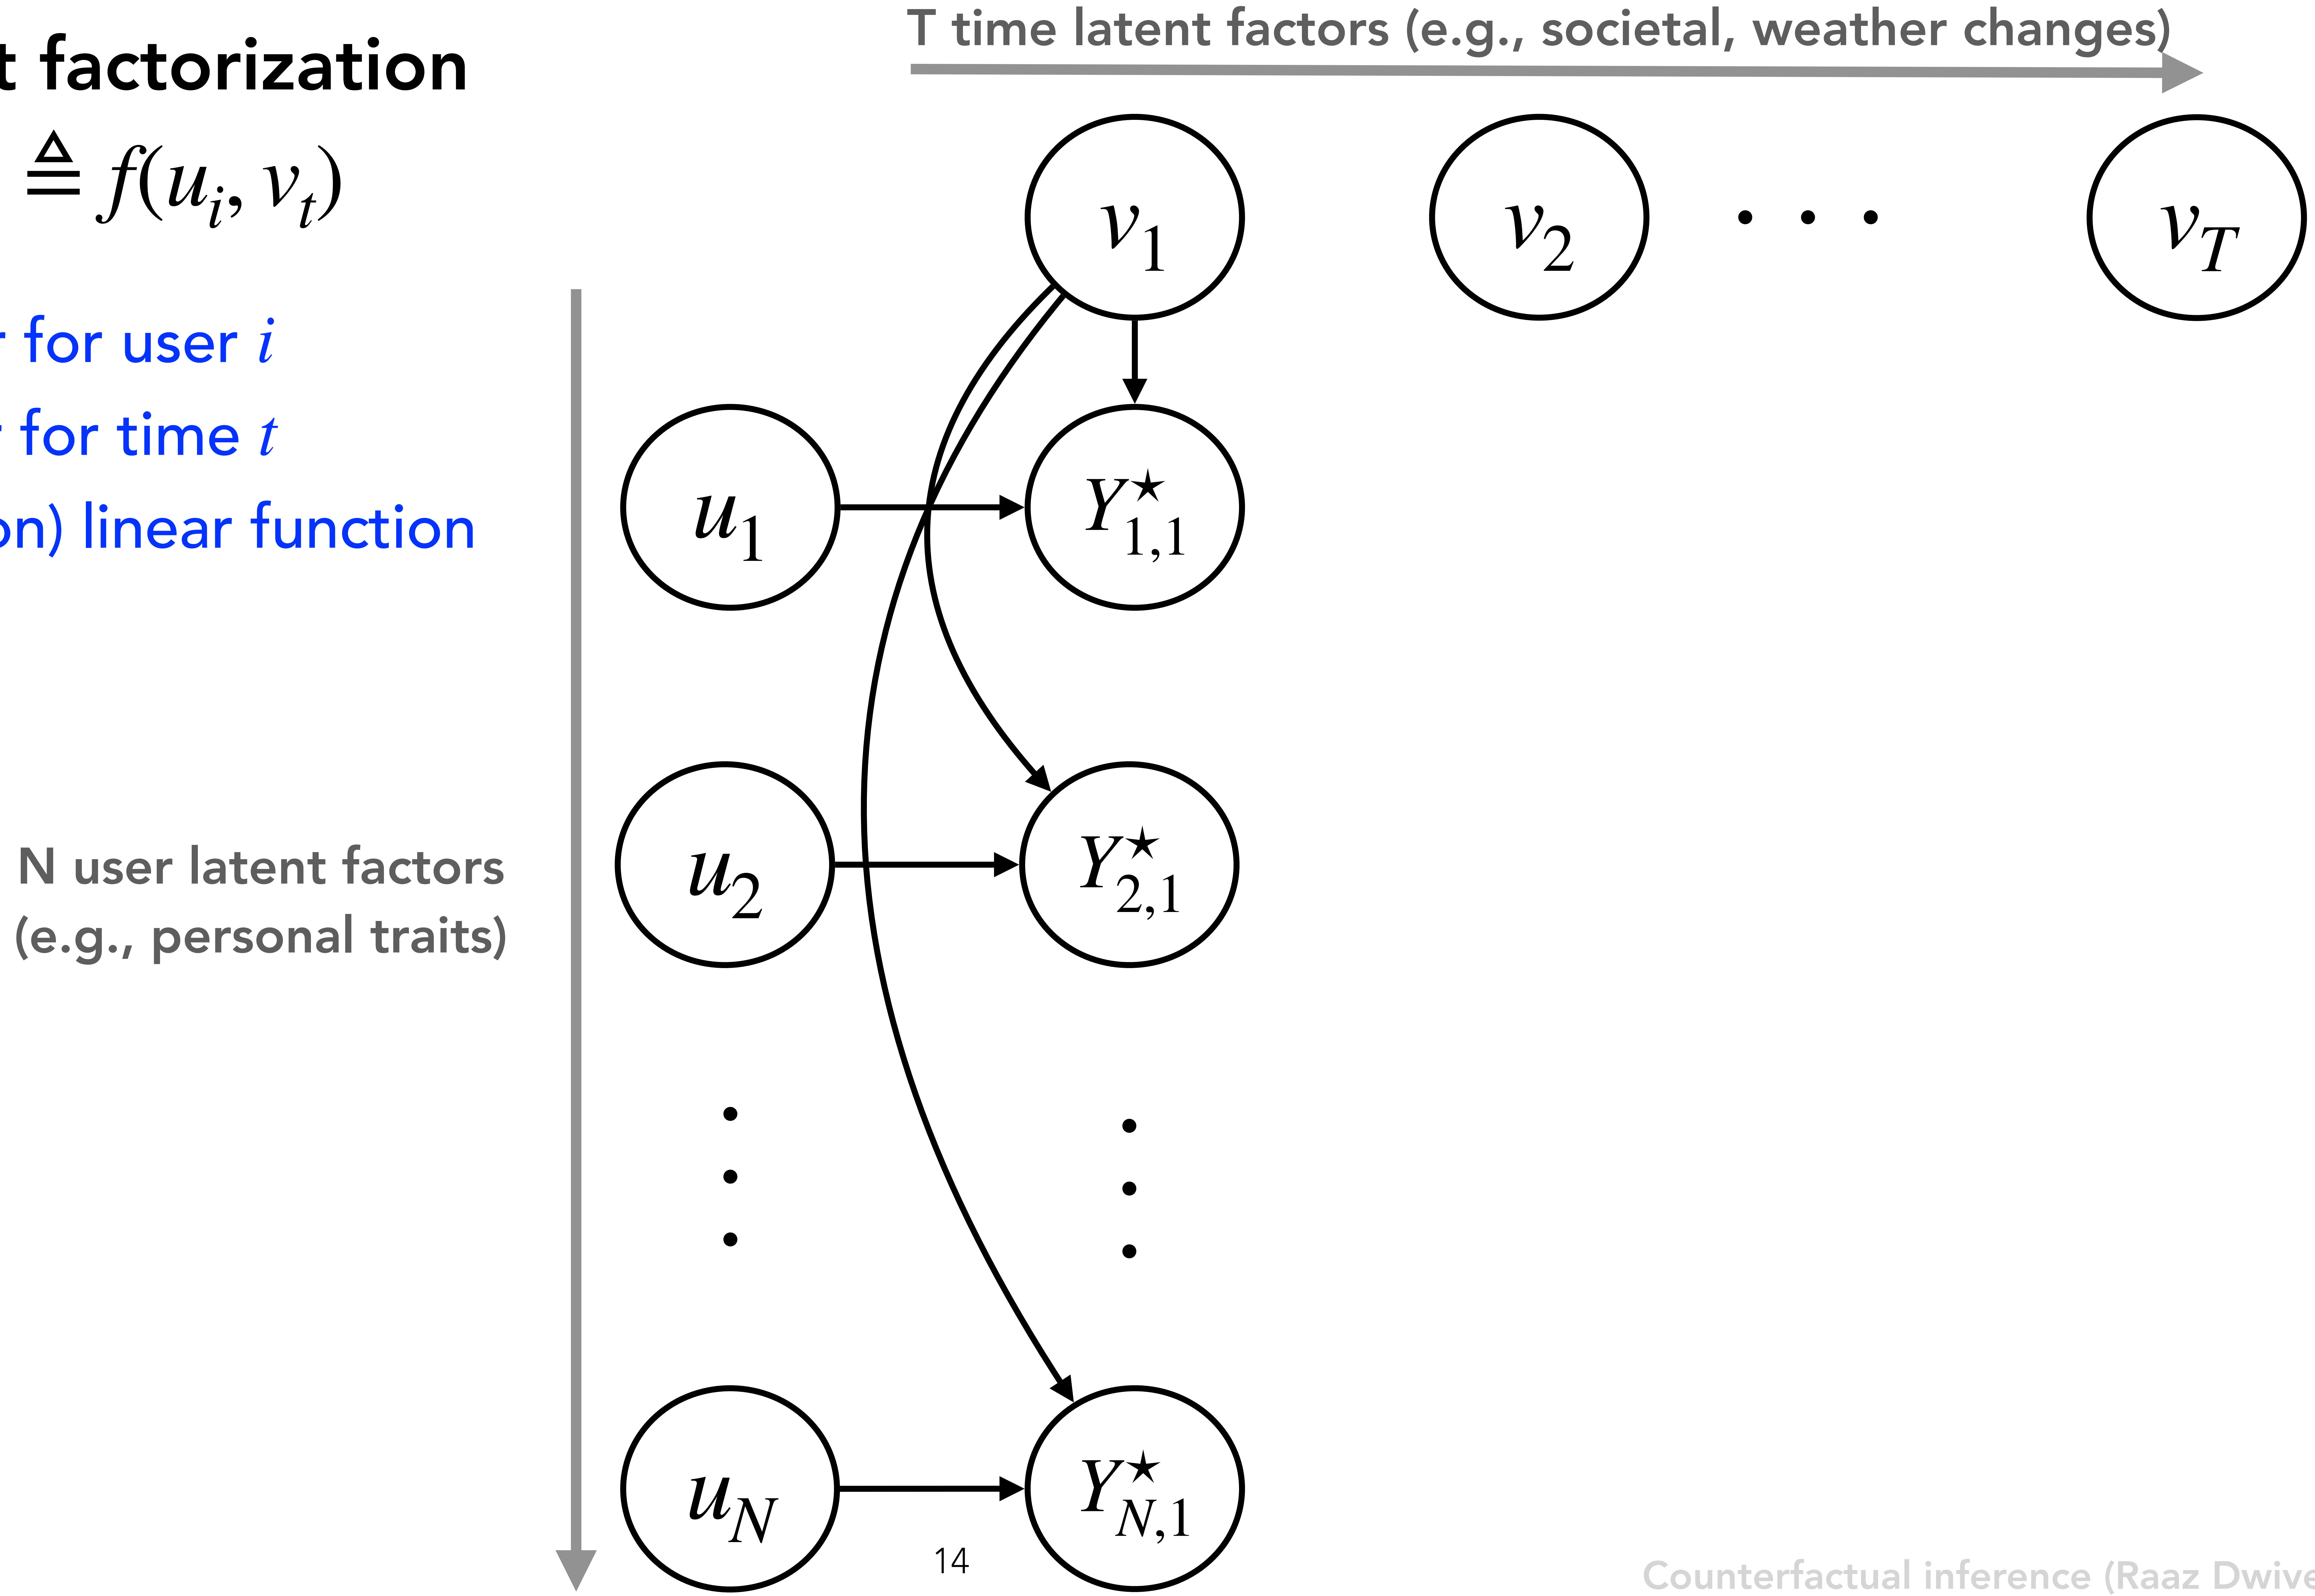

Model: Latent factorization

$$Y^\star_{i,t} \triangleq Y^{(1)}_{i,t} \triangleq f(u_i, v_t)$$

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

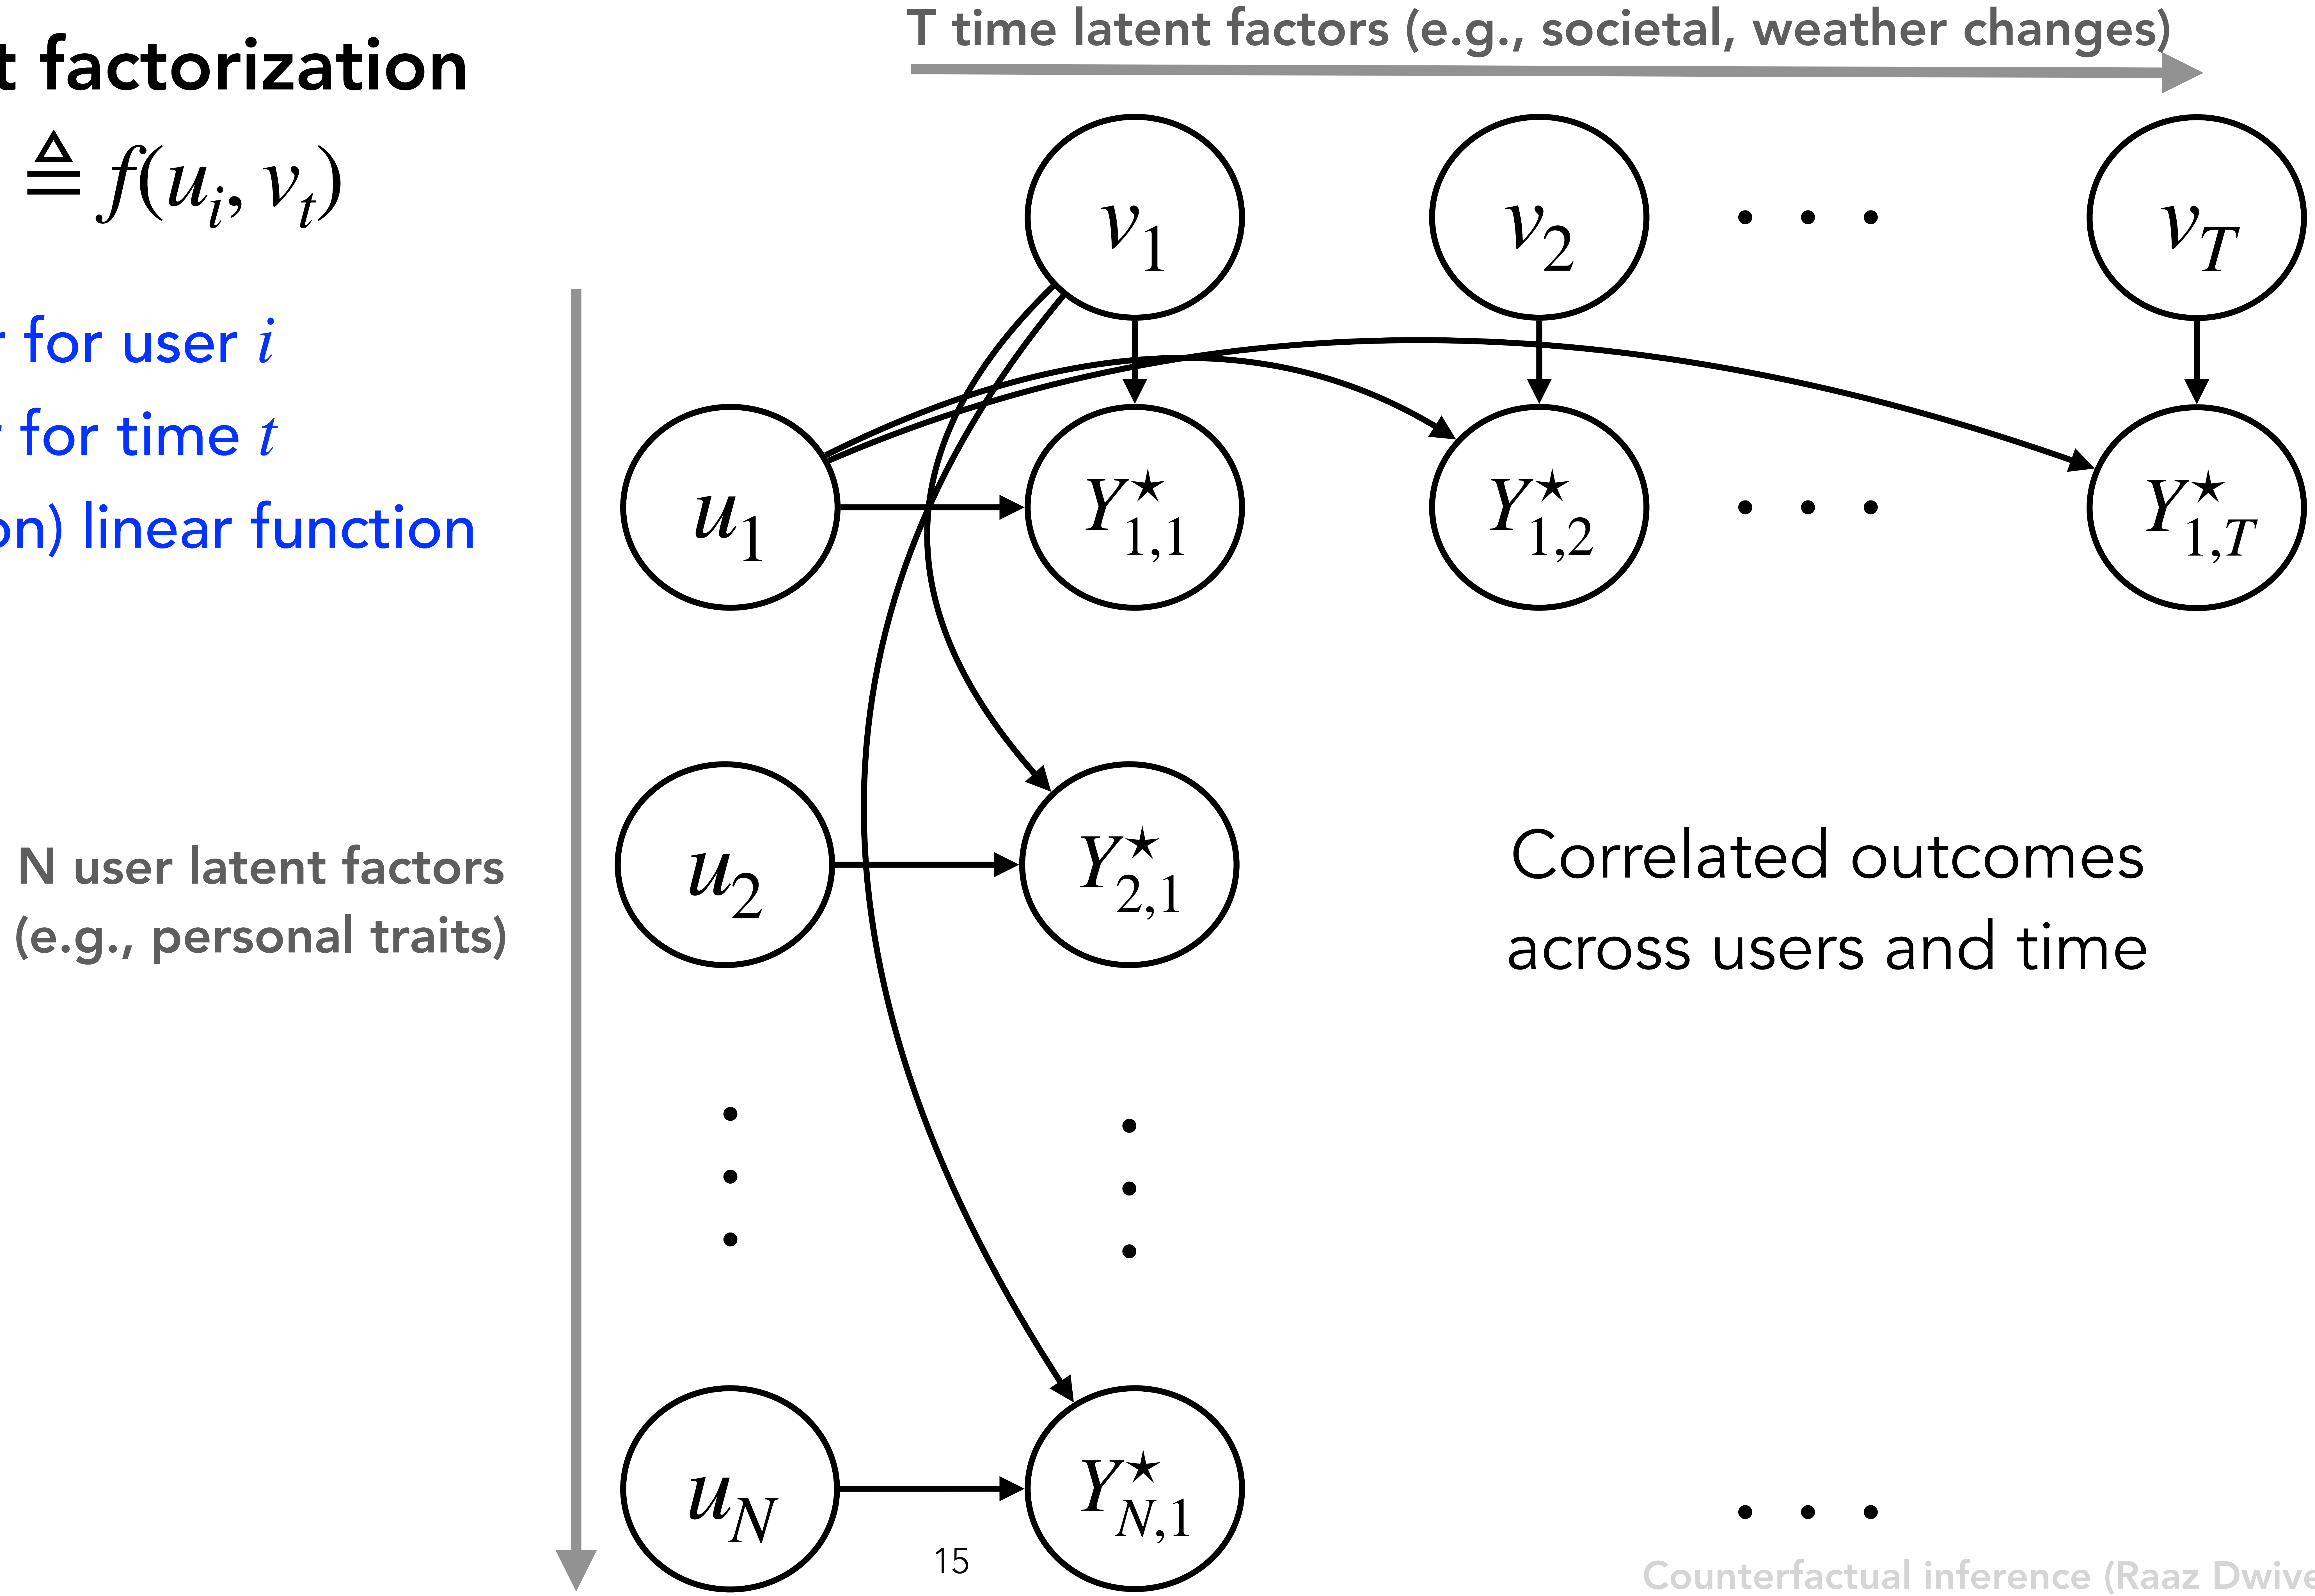

Model: Latent factorization

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

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

N user latent factors
(e.g., personal traits)

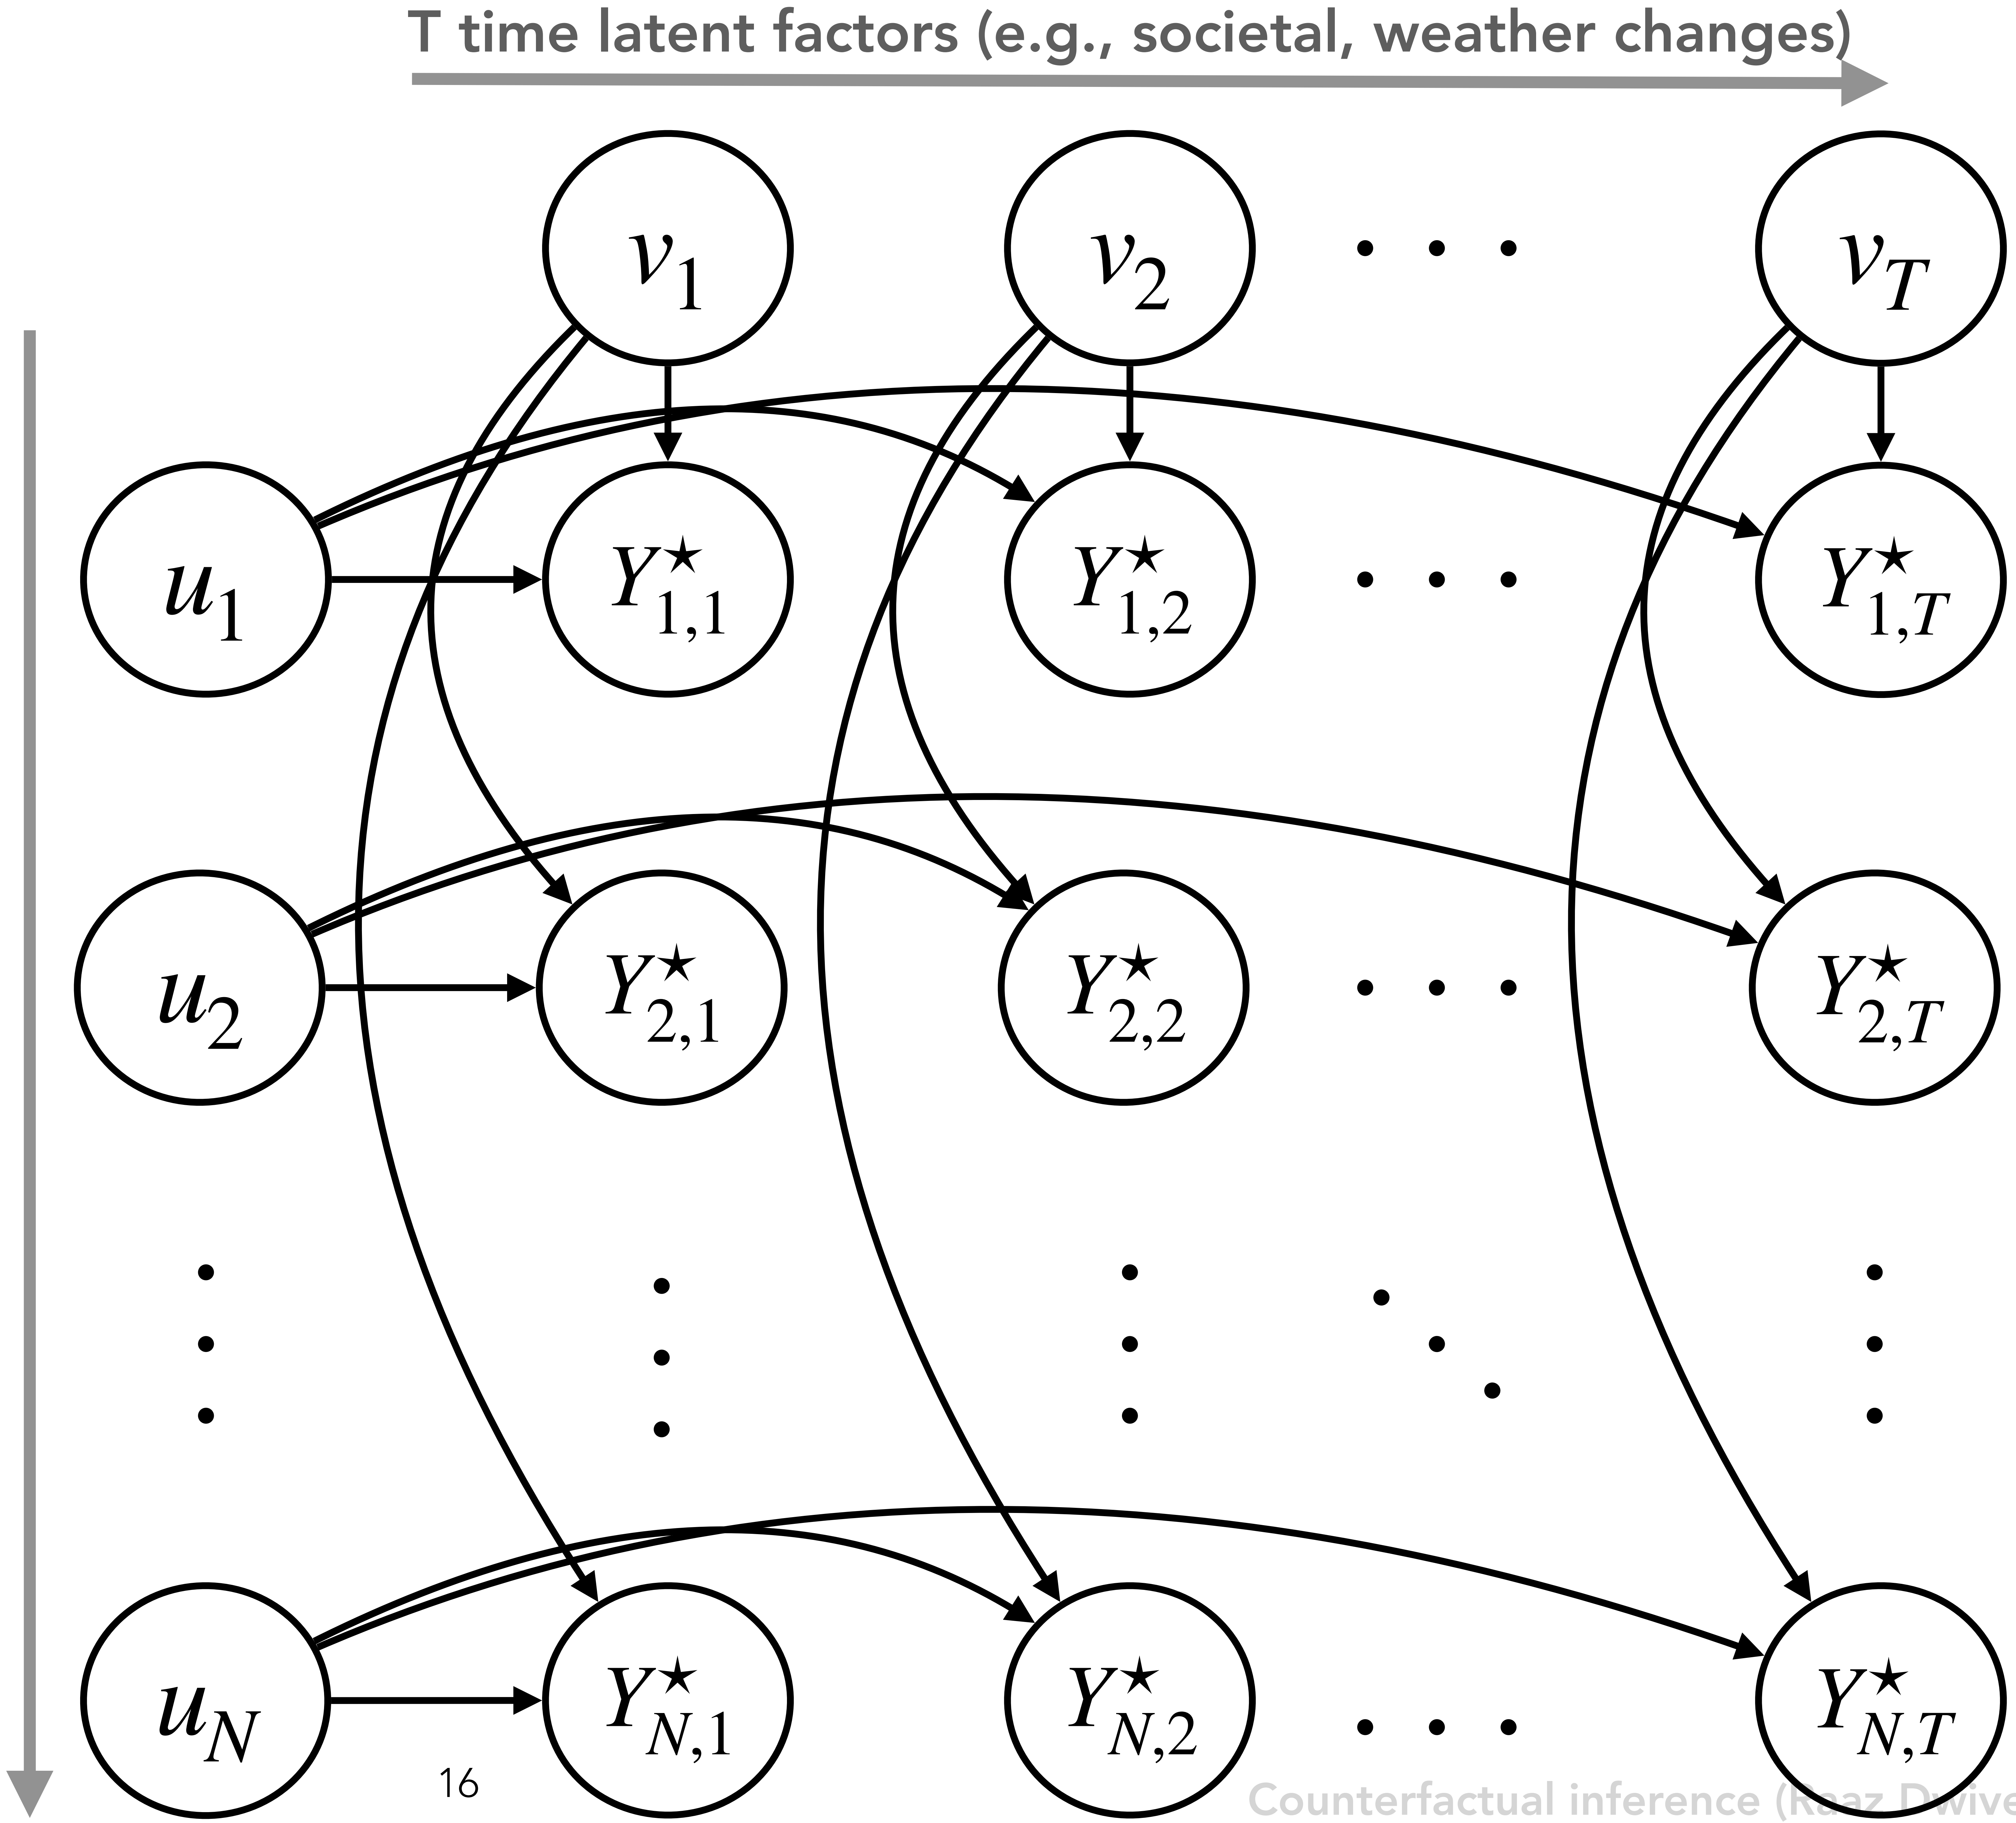

Model: Latent factorization

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

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

T time latent factors (e.g., societal, weather changes)

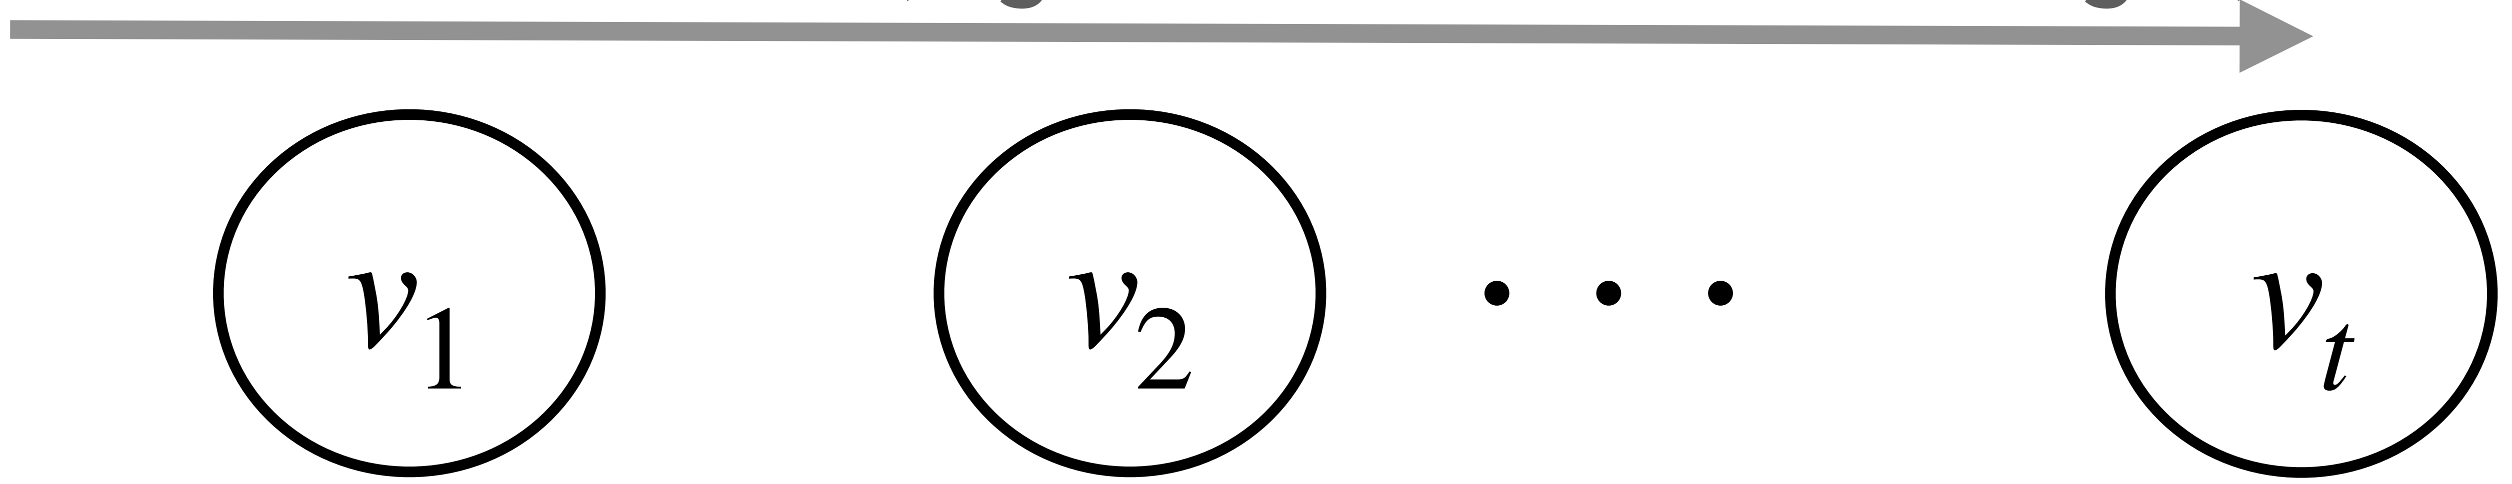

N user latent factors
(e.g., personal traits)

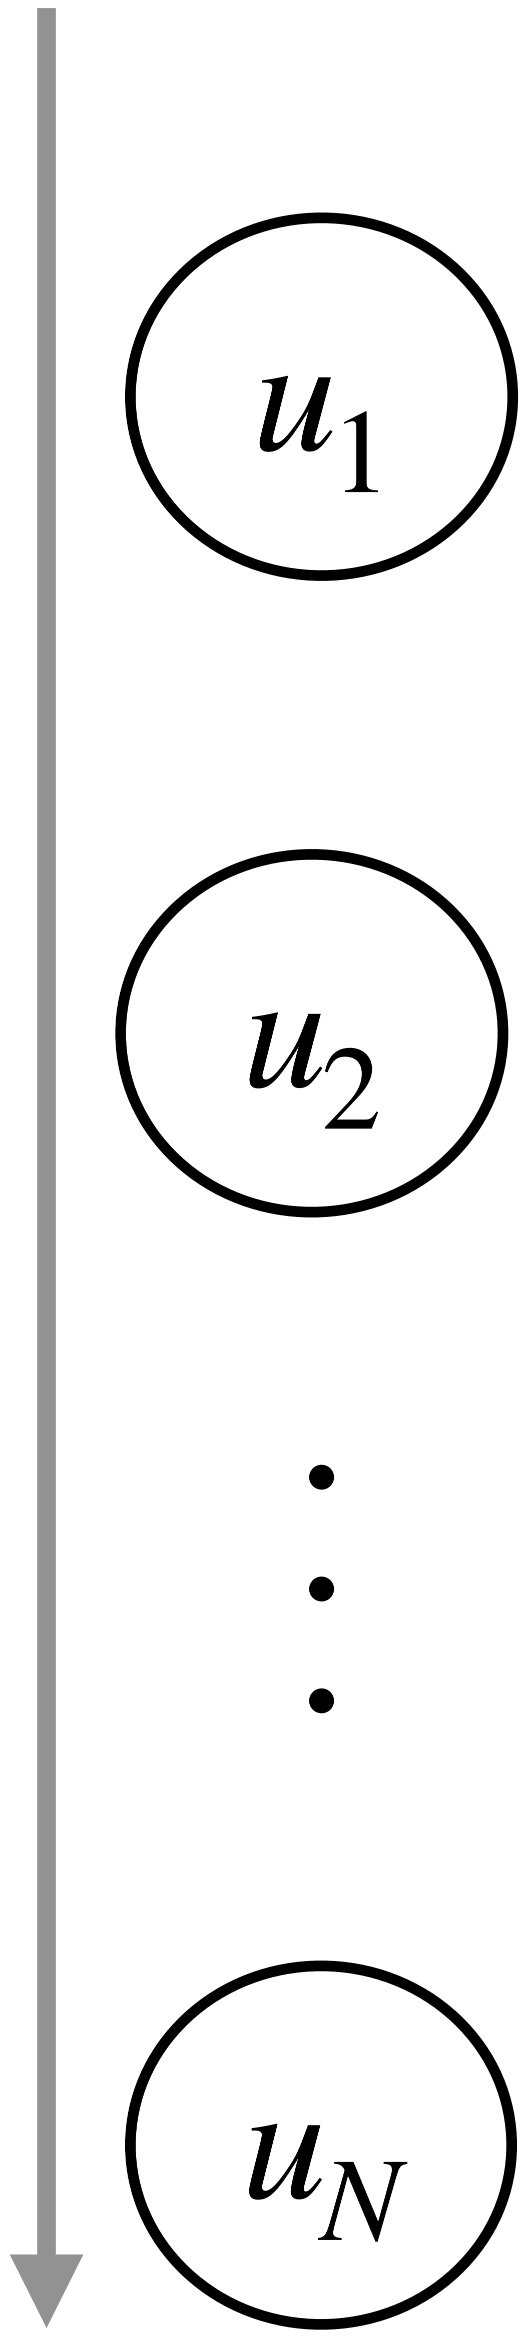

Examples include:

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

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

- Sub-class of exchangeable data

Model: Latent factorization

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

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

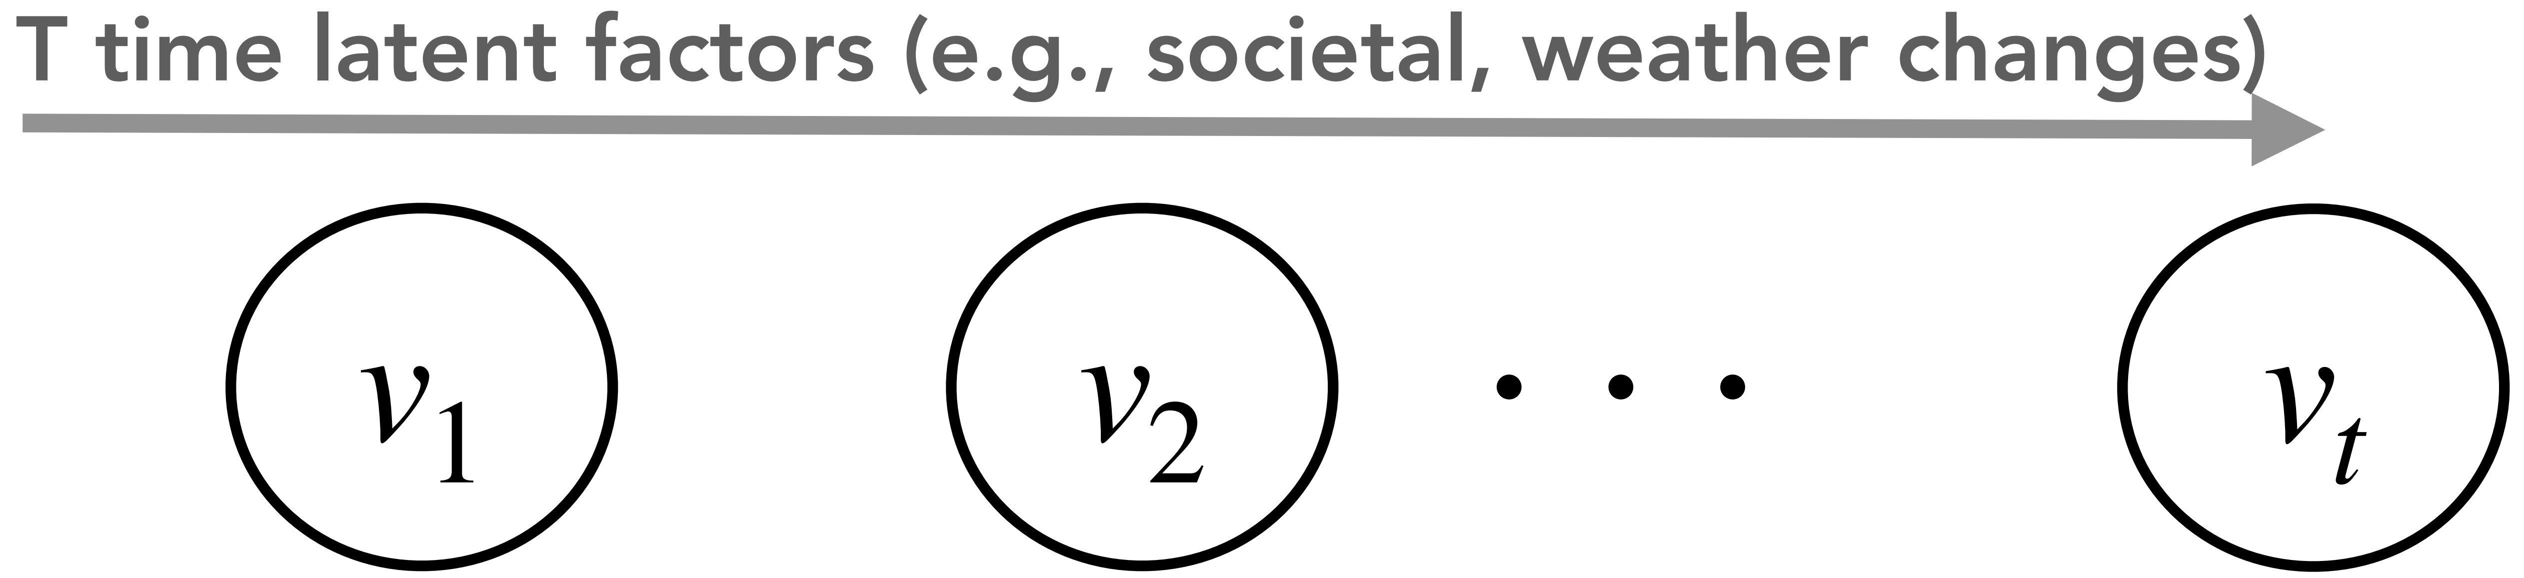

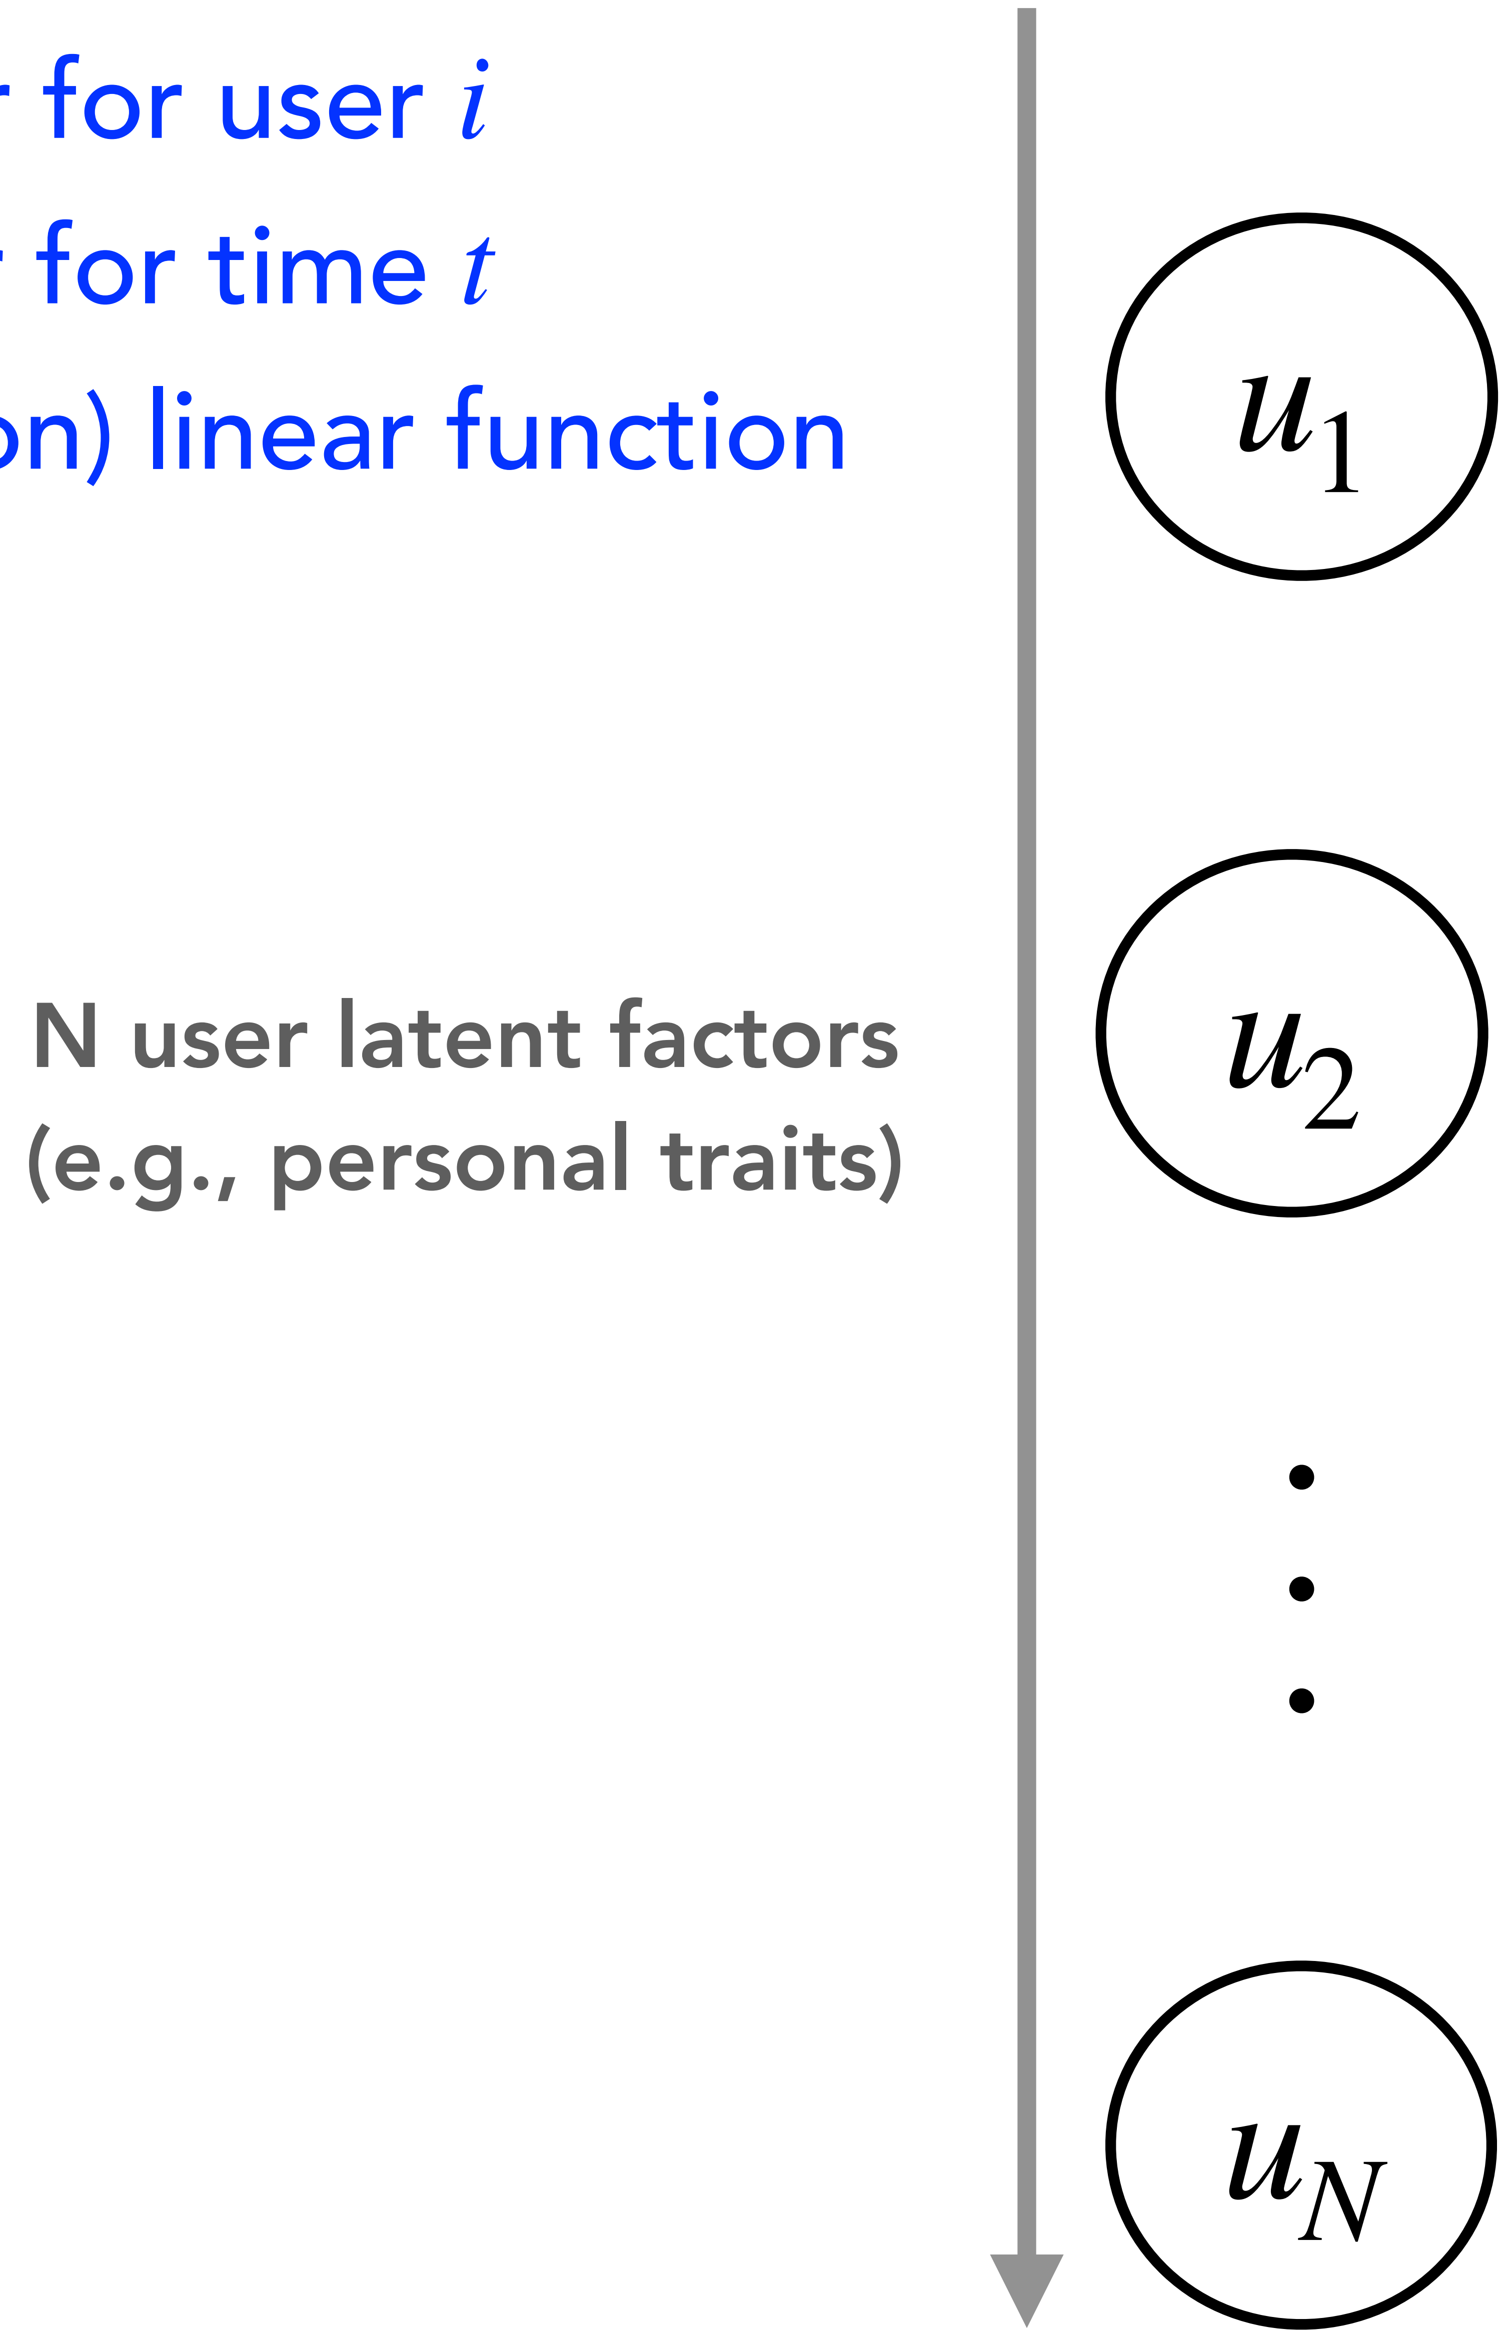

Essentially a low-rank factorization assumption—can check via **singular value decomposition!**

Algorithm:
Reduce counterfactual inference to sequential matrix completion

- Fix treatment say 1, with $Y^{\star}_{i,t} \triangleq Y^{(1)}_{i,t}$

- $N \times T$ matrix of potential outcomes with missing at random entries

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

- **New goal**: Estimate missing entries $Y^{\star}_{i,t}$ (separately for each treatment)
  policy can depend on observed outcomes of all treatments

Algorithm: A variant of nearest neighbors

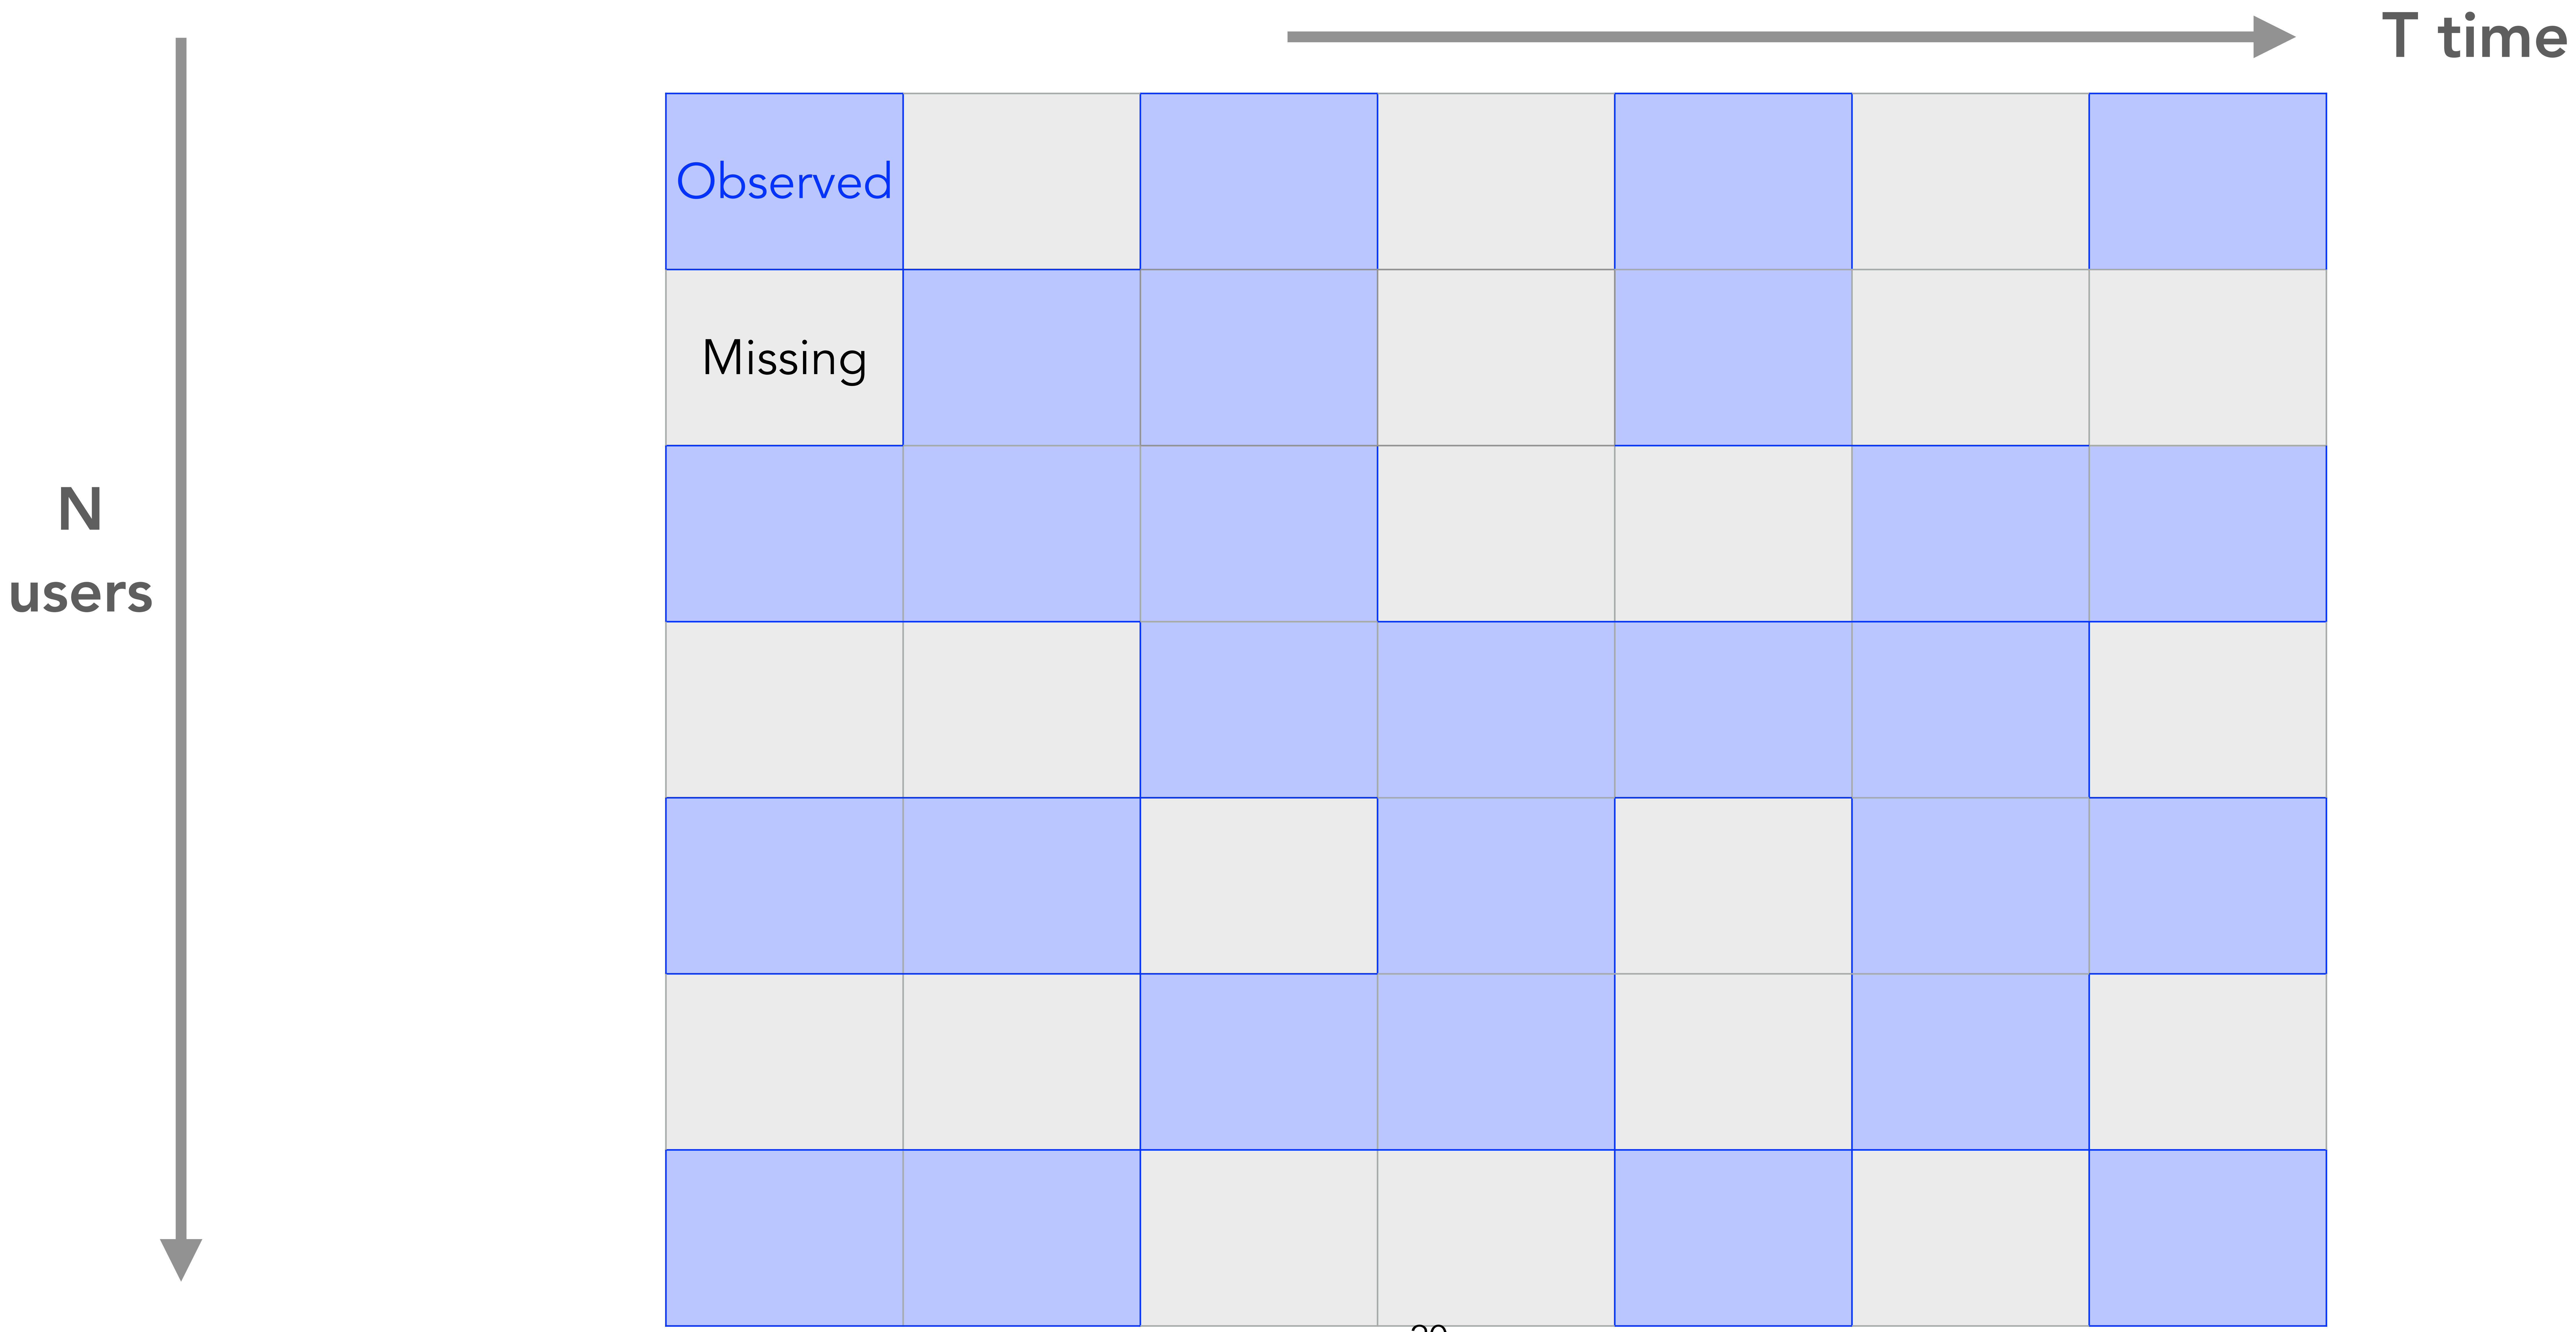

Estimating $Y^{\star}_{i,t}$ via a variant of nearest neighbors

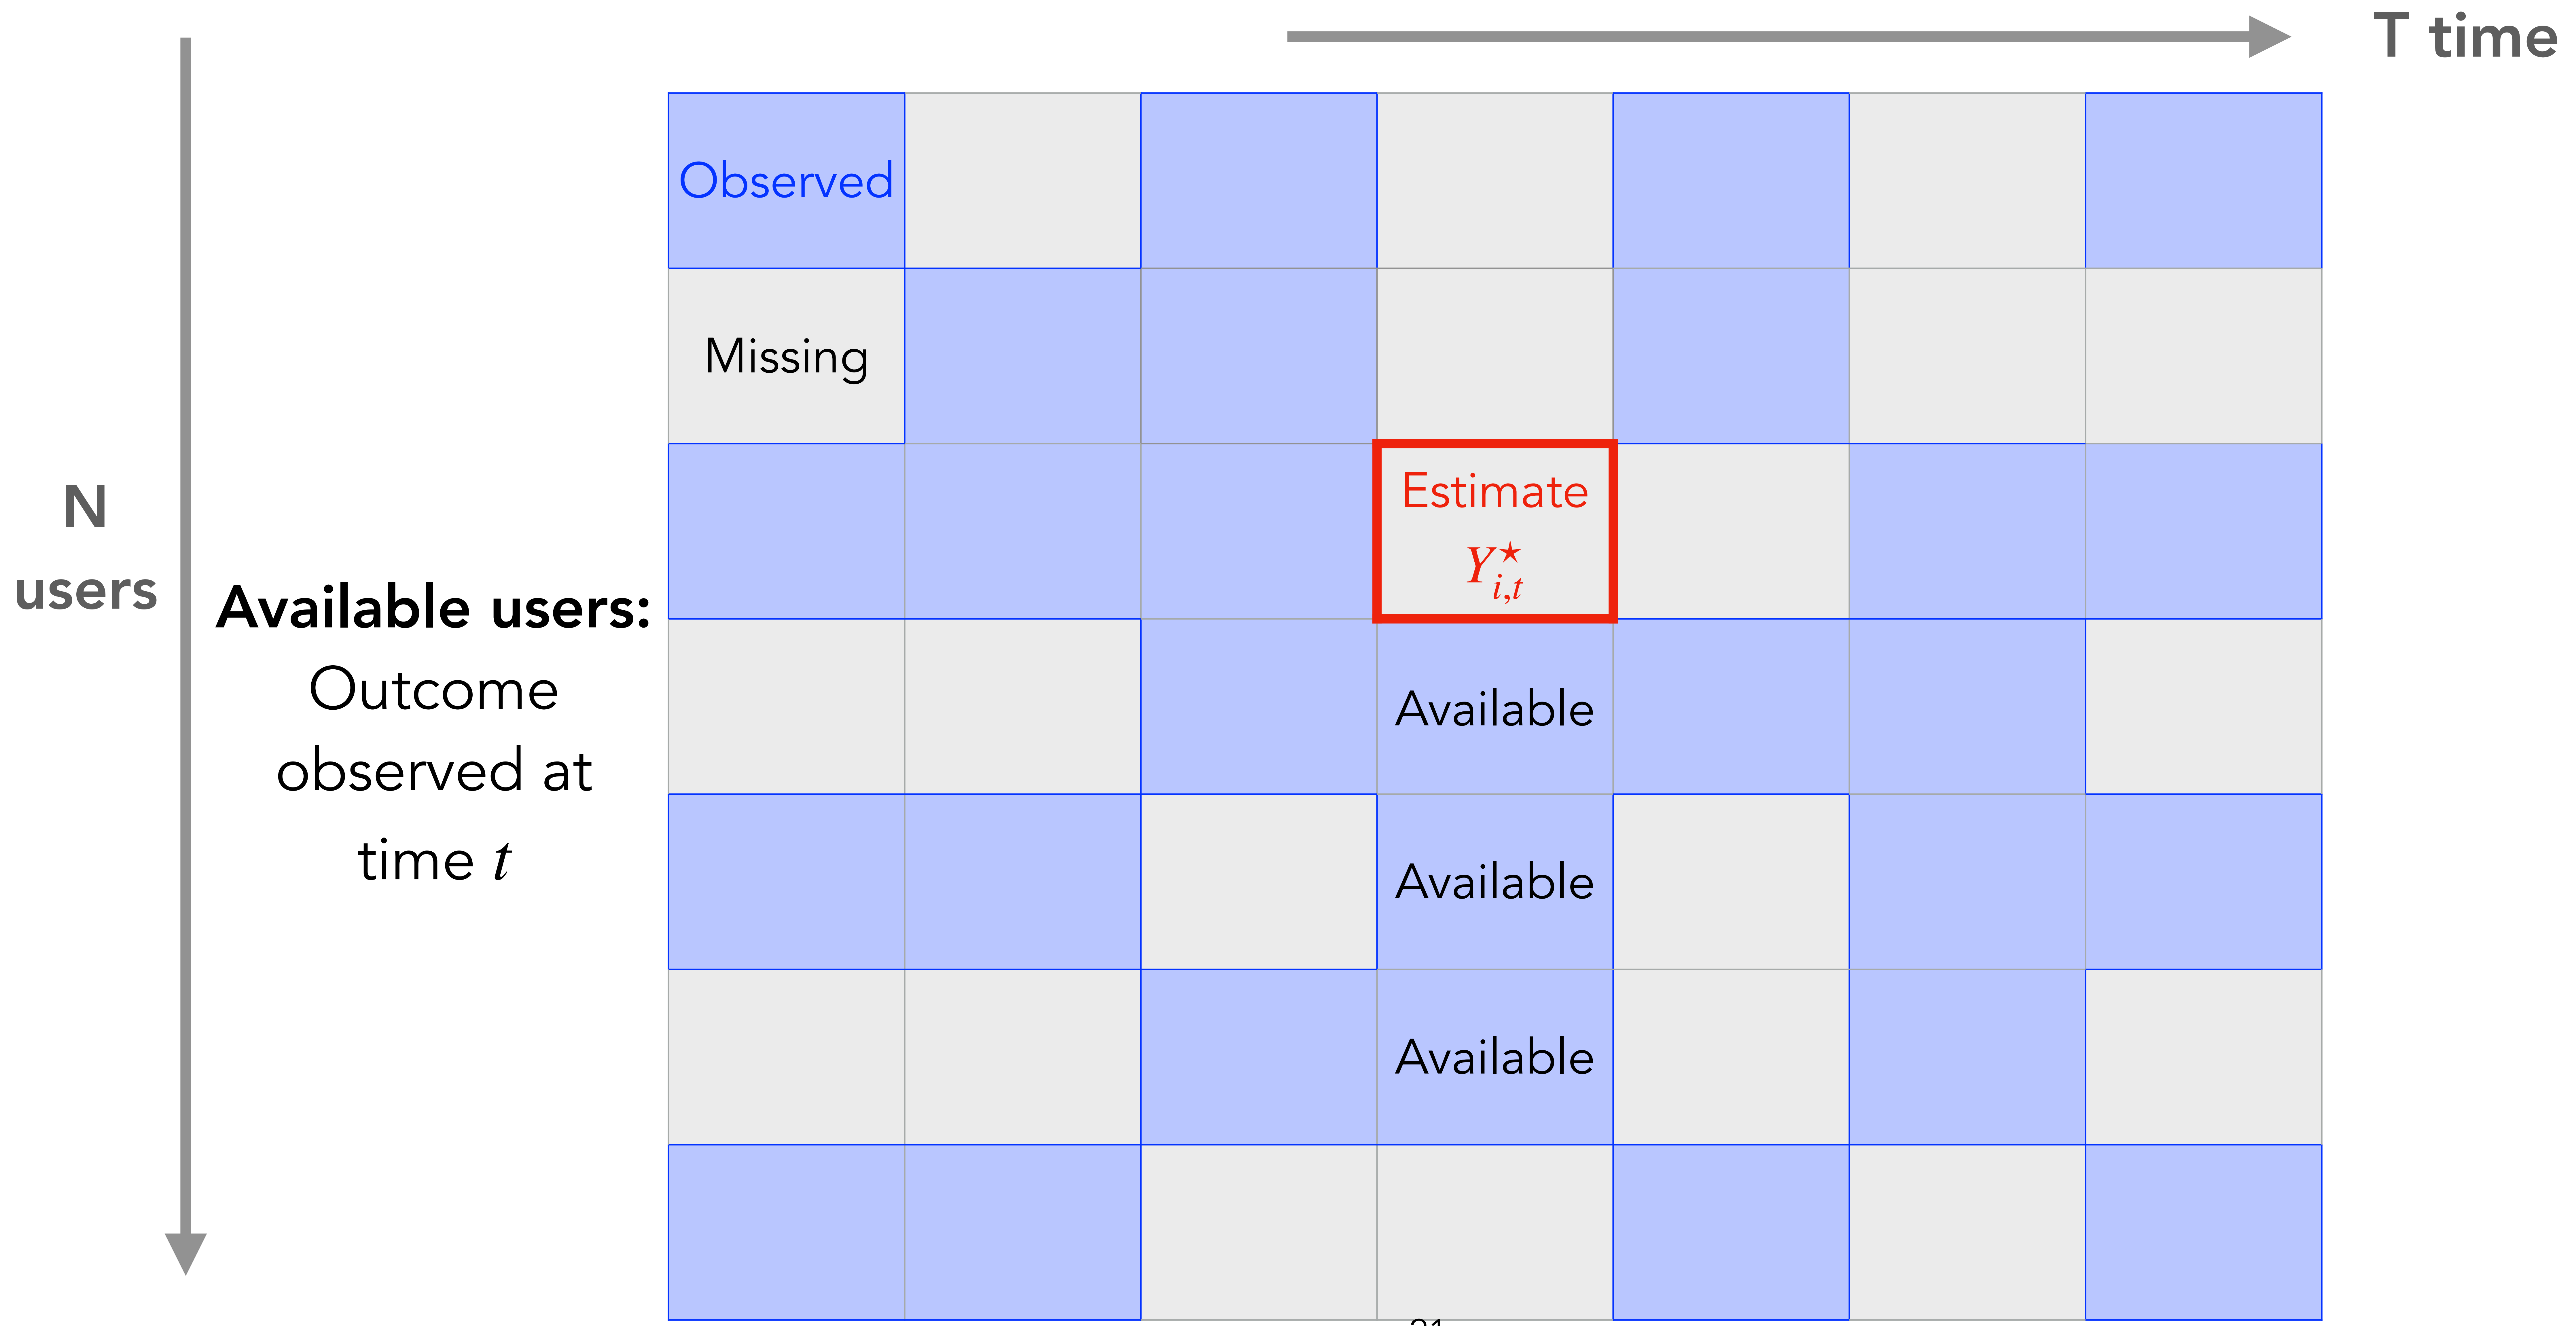

Estimating $Y^{\star}_{i,t}$ via a variant of nearest neighbors

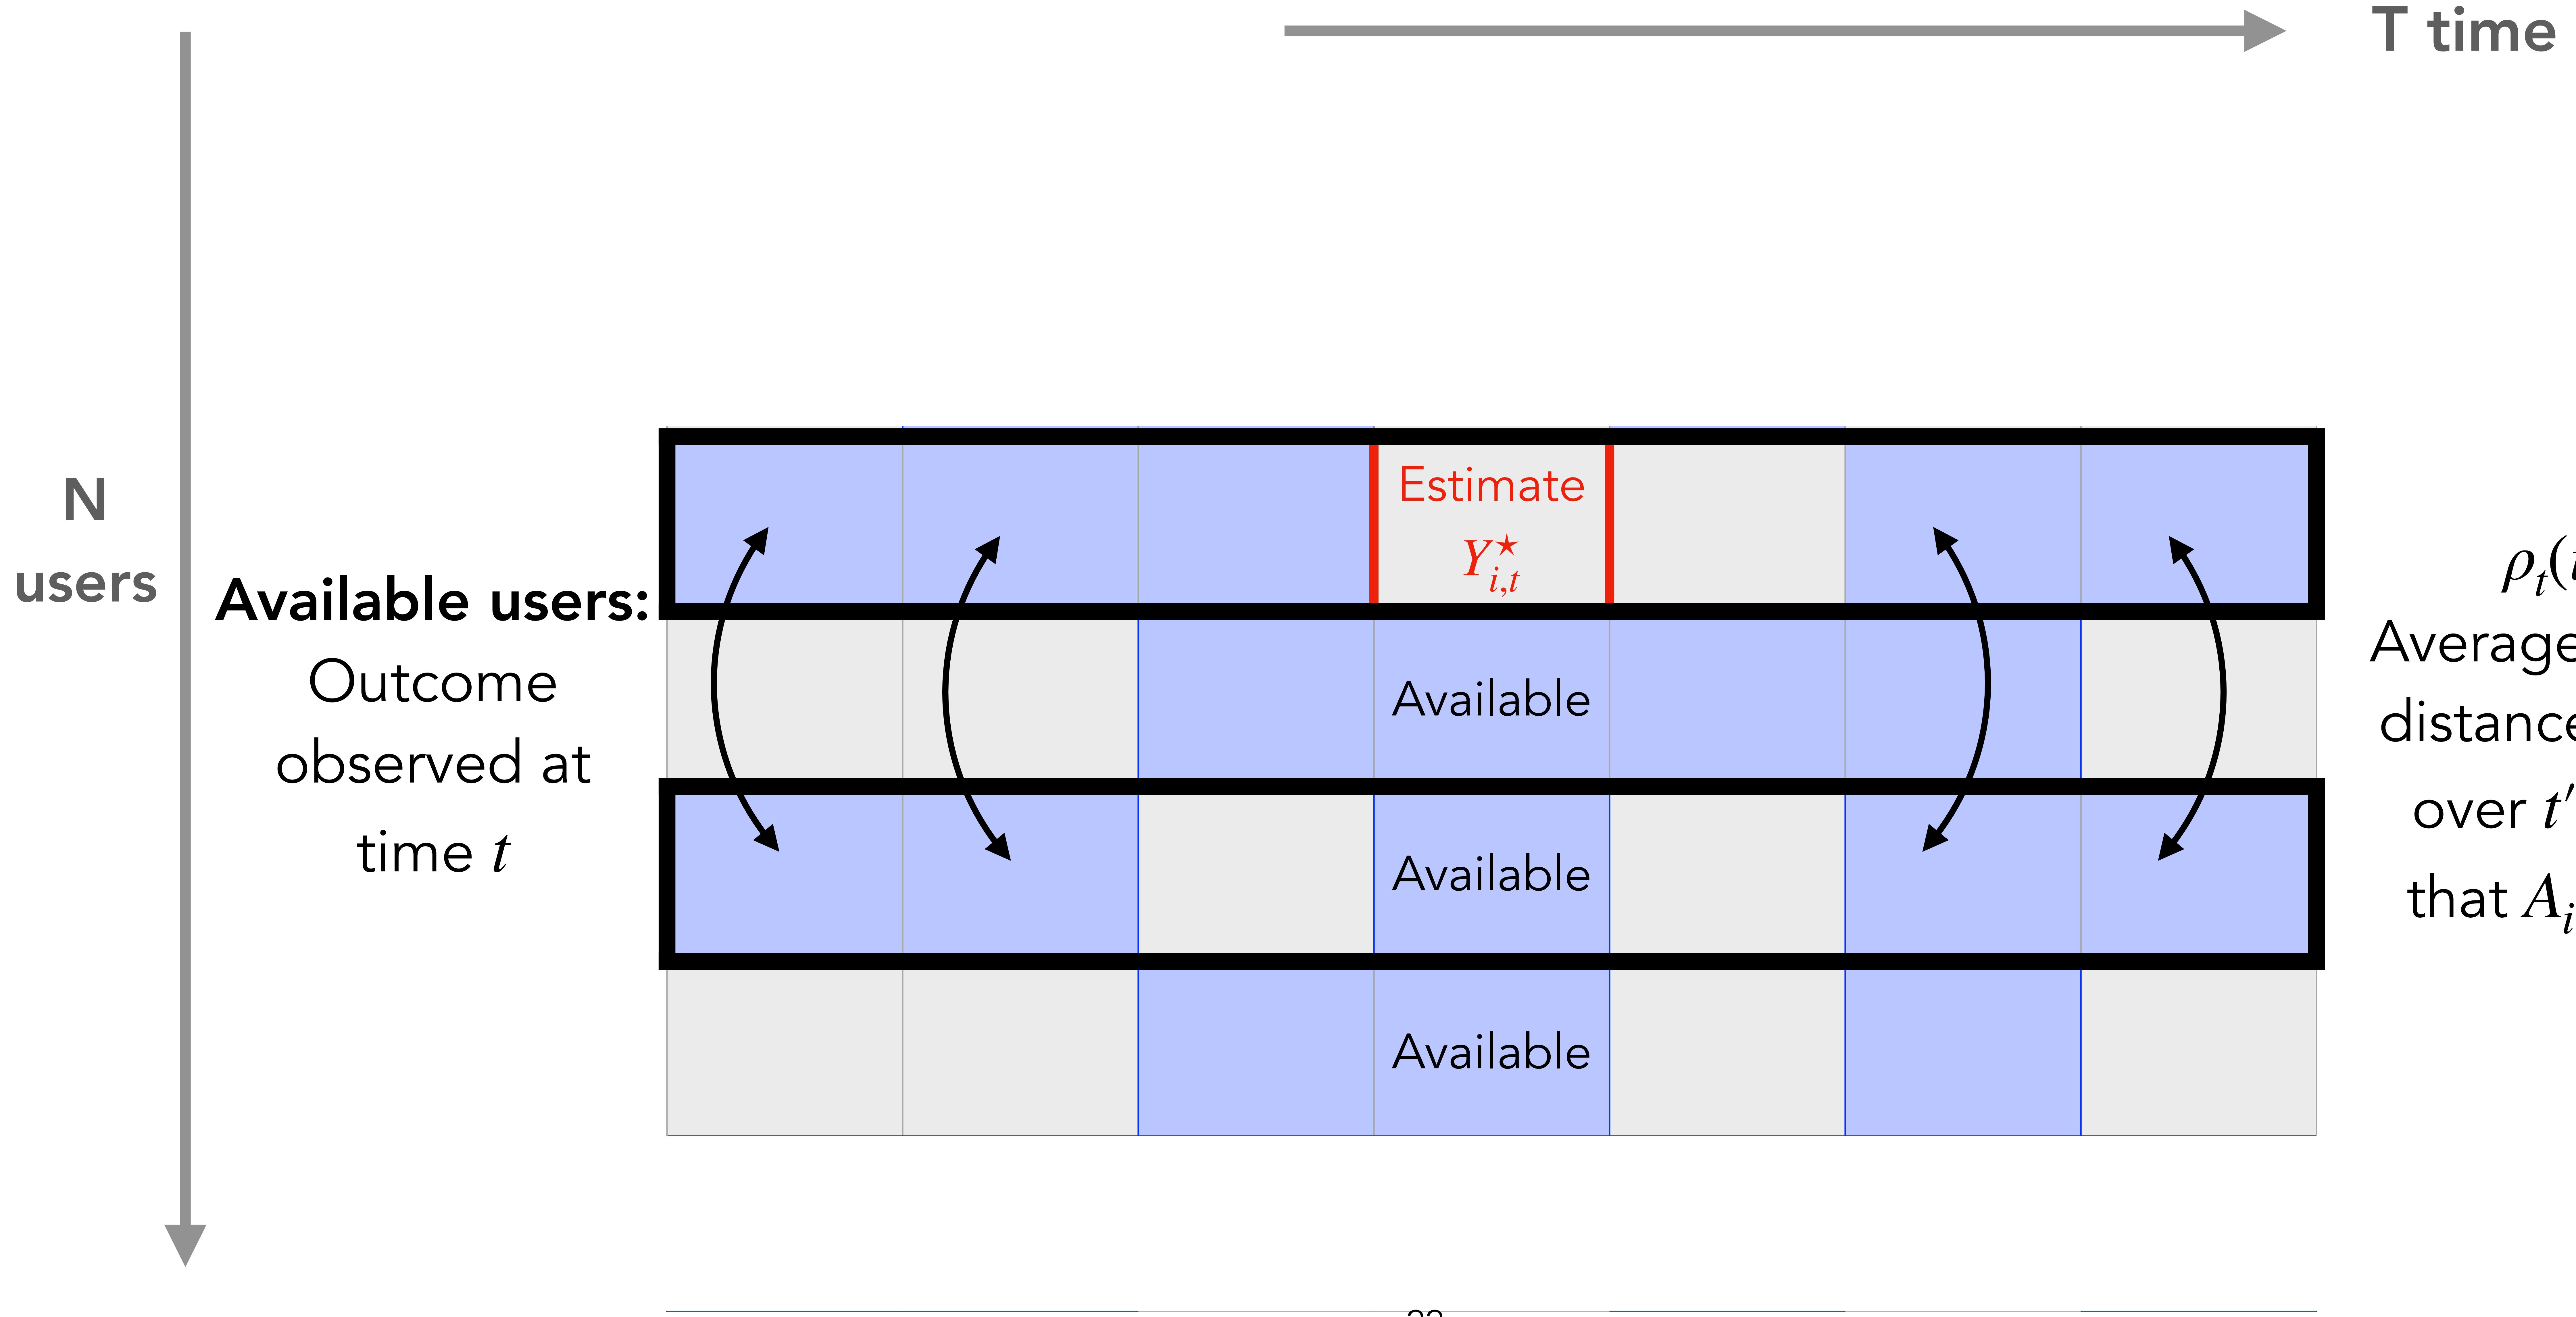

$\rho_t(i,j)$ =
Averaged squared distance between over $t' \neq t$ such that $A_{i,t'}A_{j,t'} = 1$

Non-parametric estimate: Agnostic to policy and model!

T time

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

N users

Available users: Outcome observed at time t

$\rho_t(i,j)$ = Averaged squared distance between over $t' \neq t$ such that $A_{i,t'}A_{j,t'} = 1$

Next: Theoretical guarantees

1. Sequential CLT for non-linear latent factor model
2. Anytime consistency for bilinear latent factor model

Lipschitz f

Consider a **non-linear Lipschitz** f, and suppose

- iid unit factors $\{u_j\}$, time factors $\{v_t\}$ on bounded domain
- iid bounded noise $\{\varepsilon_{j,t}\}$ with mean 0, variance σ^2
- sequential policies $\{\pi_t\}$, i.e., π_t depends on *all users'* history till $t-1$; treatments $\{A_{j,t}\}$ assigned independently given the history

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

Consider a **non-linear Lipschitz** f, and suppose

- iid unit factors $\{u_j\}$, time factors $\{v_t\}$ on bounded domain
- iid bounded noise $\{\varepsilon_{j,t}\}$ with mean 0, variance σ^2
- sequential policies $\{\pi_t\}$, i.e., π_t depends on *all users'* history till $t-1$; treatments$\{A_{j,t}\}$ assigned independently given the history

Under regularity conditions with a suitable scaling of η, for *any fixed user i at last time T* with *number of neighbors $N_{i,T}$*

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

Lipschitz f: Non-asymptotic expected squared error bound

$$\mathbb{E}\left[(\widehat{Y}_{i,T} - Y^{\star}_{i,T})^2 \mid u_i\right] \lesssim (\eta - 2\sigma^2) + \frac{D^2(1+\gamma_{i,T})}{p^2_{\min,T}\sqrt{T-1}} + \frac{\sigma^2}{p_{\min,T}\,\Phi_i N}$$

Bias due to η

Concentration of neighbor distance

Effective noise variance

Lipschitz f: Non-asymptotic expected squared error bound

$$\mathbb{E}\left[(\widehat{Y}_{i,T} - Y^\star_{i,T})^2 \,|\, u_i\right] \lesssim (\eta - 2\sigma^2) + \frac{D^2(1+\gamma_{i,T})}{p^2_{\min,T}\sqrt{T-1}} + \frac{\sigma^2}{p_{\min,T}\,\Phi_i N}$$

Bias due to η

Concentration of neighbor distance

Effective noise variance

$D =$ bound on observation

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

Lipschitz f: Non-asymptotic expected squared error bound

$$\mathbb{E}\left[(\widehat{Y}_{i,T} - Y^{\star}_{i,T})^2 \,|\, u_i\right] \lesssim (\eta - 2\sigma^2) + \frac{D^2(1+\gamma_{i,T})}{p^2_{\min,T}\sqrt{T-1}} + \frac{\sigma^2}{p_{\min,T}\,\Phi_i N}$$

Bias due to η

Concentration of neighbor distance

Effective noise variance

$D =$ bound on observation

$$p_{\min,T} = \min_{t,j} \pi_t(j)$$

min probability of sampling any entry

$$\Phi_i = \mathbb{P}_u\left(\mathbb{E}_v[f(u_i,v) - f(u,v)^2] \leq \eta/2 - \sigma^2\right)$$

Probability of sampling a nearest neighbor

Lipschitz f: Non-asymptotic expected squared error bound

$$\mathbb{E}\left[(\widehat{Y}_{i,T} - Y^{\star}_{i,T})^2 \,|\, u_i\right] \lesssim (\eta - 2\sigma^2) + \frac{D^2(1+\gamma_{i,T})}{p^2_{\min,T}\sqrt{T-1}} + \frac{\sigma^2}{p_{\min,T}\,\Phi_i N}$$

Bias due to η

Concentration of neighbor distance

Effective noise variance

$D =$ bound on observation

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

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

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

Cumulative future dependency of adaptive policies on one column

Regularity conditions needed for the CLT

$$\mathbb{E}\left[(\widehat{Y}_{i,T} - Y^{\star}_{i,T})^2 \,|\, u_i\right] \lesssim (\eta - 2\sigma^2) + \frac{D^2(1+\gamma_{i,T})}{p^2_{\min,T}\sqrt{T-1}} + \frac{\sigma^2}{p_{\min,T}\,\Phi_i N}$$

``Bias" terms go to zero after multiplying by number of neighbors $N_{i,T}$:

- $\eta - 2\sigma^2$ goes to zero fast enough
- $N_{i,T}$ can **not** grow faster than $p^2_{\min,T}\sqrt{T}$ (cap number of nearest neighbors)

The denominator (min number of nearest neighbors) goes to ∞

(See the paper for detailed examples)

$$f(u, v) = \langle u, v \rangle$$

Assuming bounded observations and

- iid latent unit $\{u_j\}$ and time factors $\{v_t\}$ with positive definite covariance $\mathbb{V}^\star$ on bounded domain
- iid bounded noise $\{\varepsilon_{j,t}\}$ with mean 0, variance σ^2
- adaptive policies $\{\pi_t\}$ that introduce ``weak correlations"

$f(u, v) = \langle u, v \rangle$: A consistency result for <u>any time t</u>

Assuming bounded observations and

- iid latent unit $\{u_j\}$and time factors $\{v_t\}$ with positive definite covariance $\mathbb{V}^\star$ on bounded domain
- iid bounded noise $\{\varepsilon_{j,t}\}$ with mean 0, variance σ^2
- adaptive policies $\{\pi_t\}$ that introduce ``weak correlations"

then with suitable scaling of (η, N, T)

$$\widehat{Y}_{i,t} \to Y^\star_{i,t} \text{ in probability for any fixed } (i, t).$$

Proof sketch: Advantage of bilinearity

Under the regularity conditions

$$\rho_t(i,j) \le \eta \quad \overset{\textbf{weak correlations}}{\Longrightarrow} \quad (u_i - u_j)^\top \left(\frac{\sum_{t' \neq t} A_{j,t'} A_{i,t'} v_{t'} v_{t'}^\top}{\sum_{t' \neq t} A_{j,t'} A_{i,t'}} \right) (u_i - u_j) \lesssim \eta - c\sigma^2$$

Proof sketch: Advantage of bilinearity

Under the regularity conditions

$$\rho_t(i,j) \leq \eta \quad \overset{\textbf{weak correlations}}{\Longrightarrow} \quad (u_i - u_j)^\top \left(\frac{\sum_{t' \neq t} A_{j,t'} A_{i,t'} v_{t'} v_{t'}^\top}{\sum_{t' \neq t} A_{j,t'} A_{i,t'}} \right) (u_i - u_j) \lesssim \eta - c\sigma^2$$

$$\overset{\textbf{pos. def. covariance + weak correlations}}{\Longrightarrow} \quad \|u_i - u_j\|_2 \lesssim \sqrt{\frac{\eta - c\sigma^2}{\lambda}}$$

$$\overset{\textbf{Cauchy-Schwarz}}{\Longrightarrow} \quad |Y^\star_{i,t} - Y^\star_{j,t}| \lesssim \|v_t\|_2 \cdot \sqrt{\frac{\eta - c\sigma^2}{\lambda}}$$

where $\lambda = \lambda_{\min}(\mathbb{V}^\star)$

Summary:
Counterfactual inference in sequential experiments

https://rzrsk.github.io

Sequential experimental design

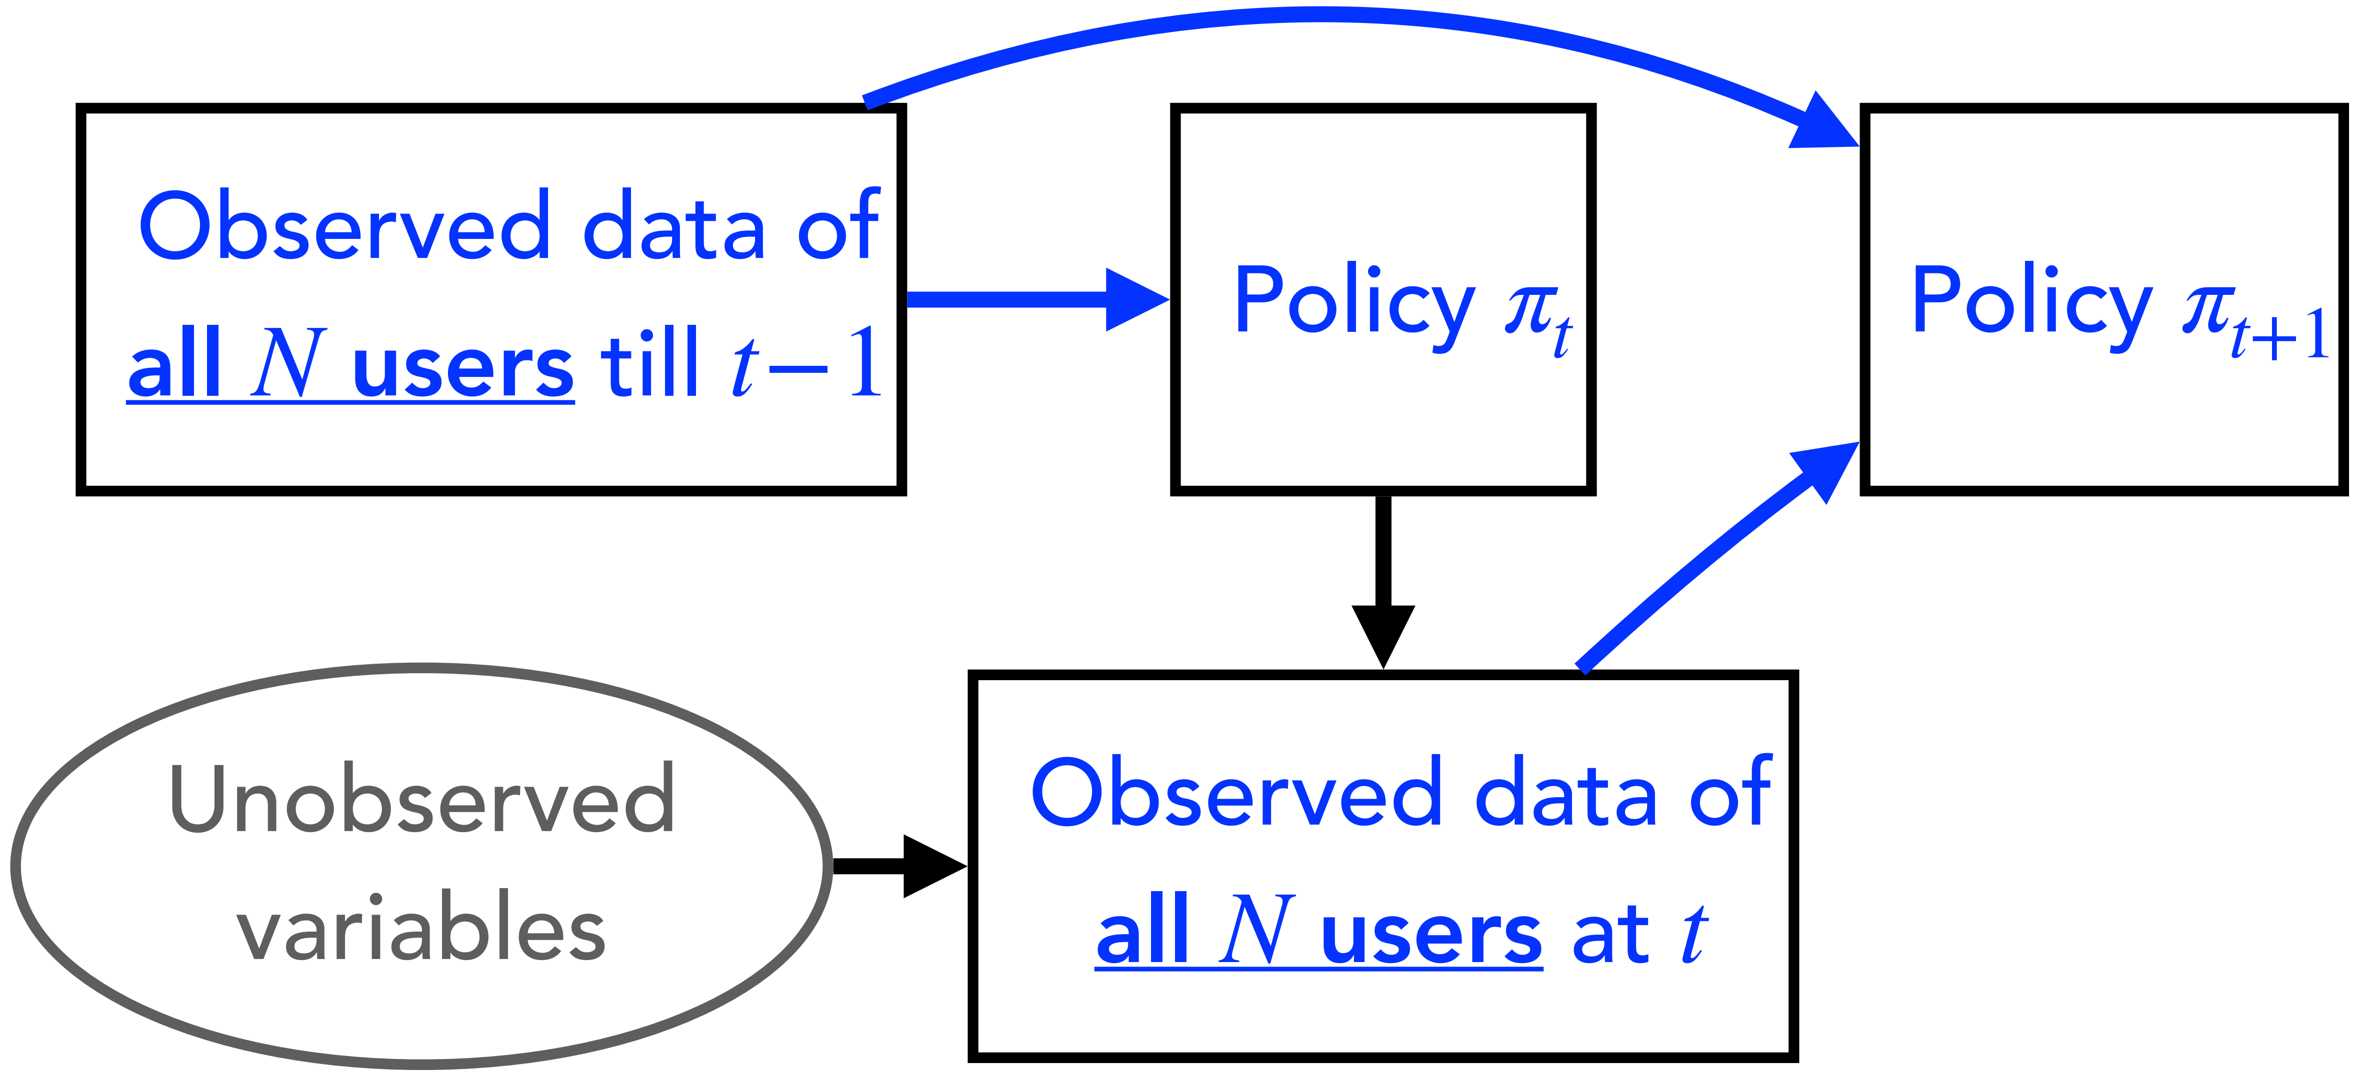

Summary:
Counterfactual inference in sequential experiments

https://rzrsk.github.io

Sequential experimental design

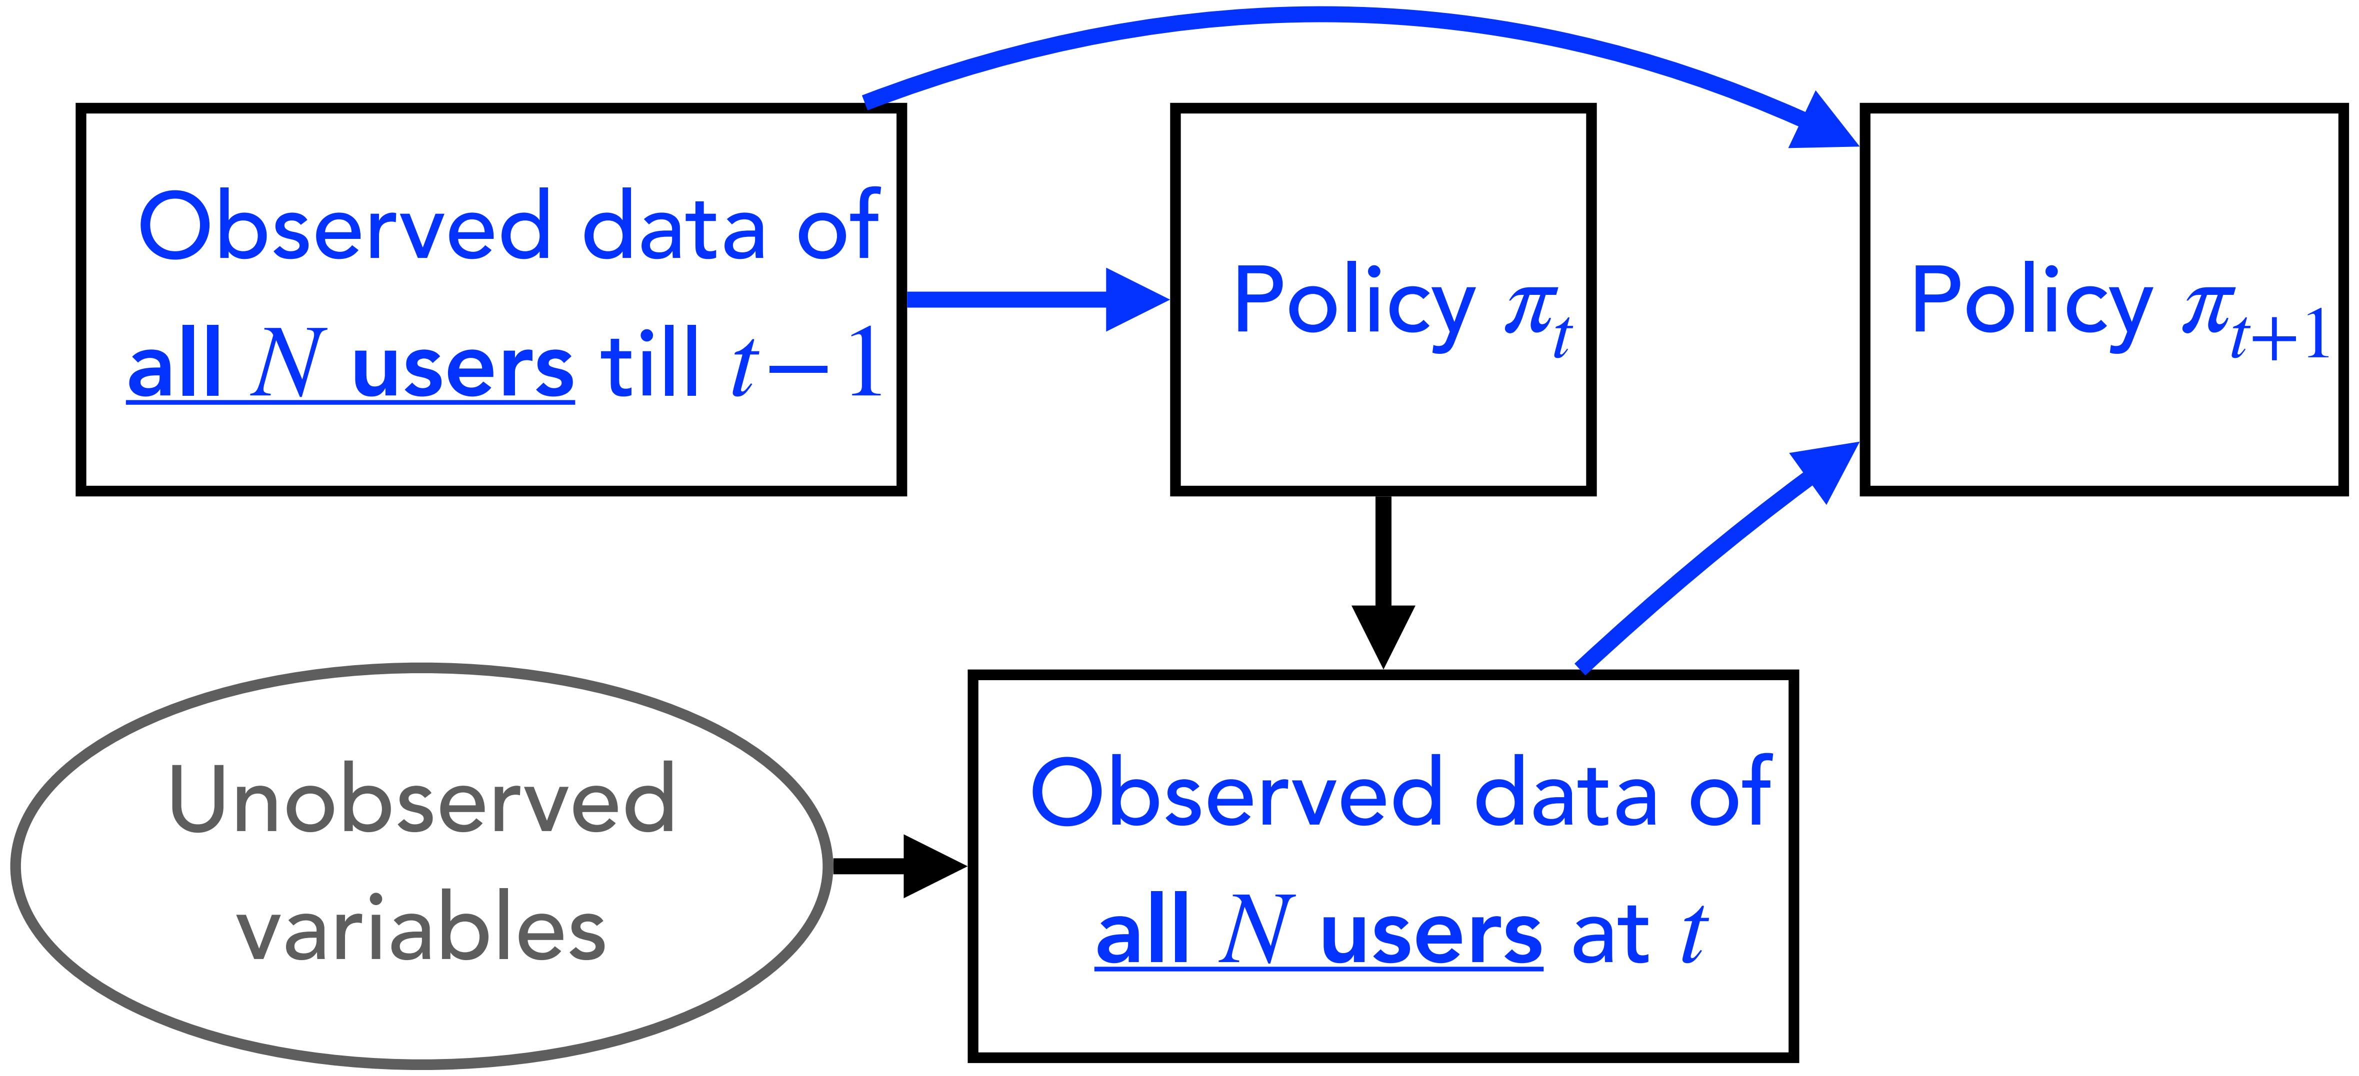

Modeling assumptions

(For each treatment separately)

$$Z_{i,t} = \begin{cases} Y^{\star}_{i,t} + \varepsilon_{i,t} & \text{if} \quad A_{i,t} = 1 \\ \text{unknown} & \text{if} \quad A_{i,t} = 0 \end{cases}$$

Latent factor model:

$$Y^{\star}_{i,t} = f(u_i, v_t)$$

Algorithm

Nearest neighbor estimate $\widehat{Y}_{i,t}$ for $Y^{\star}_{i,t}$ by measuring distance over time

Summary:
Counterfactual inference in sequential experiments

https://rzrsk.github.io

Sequential experimental design

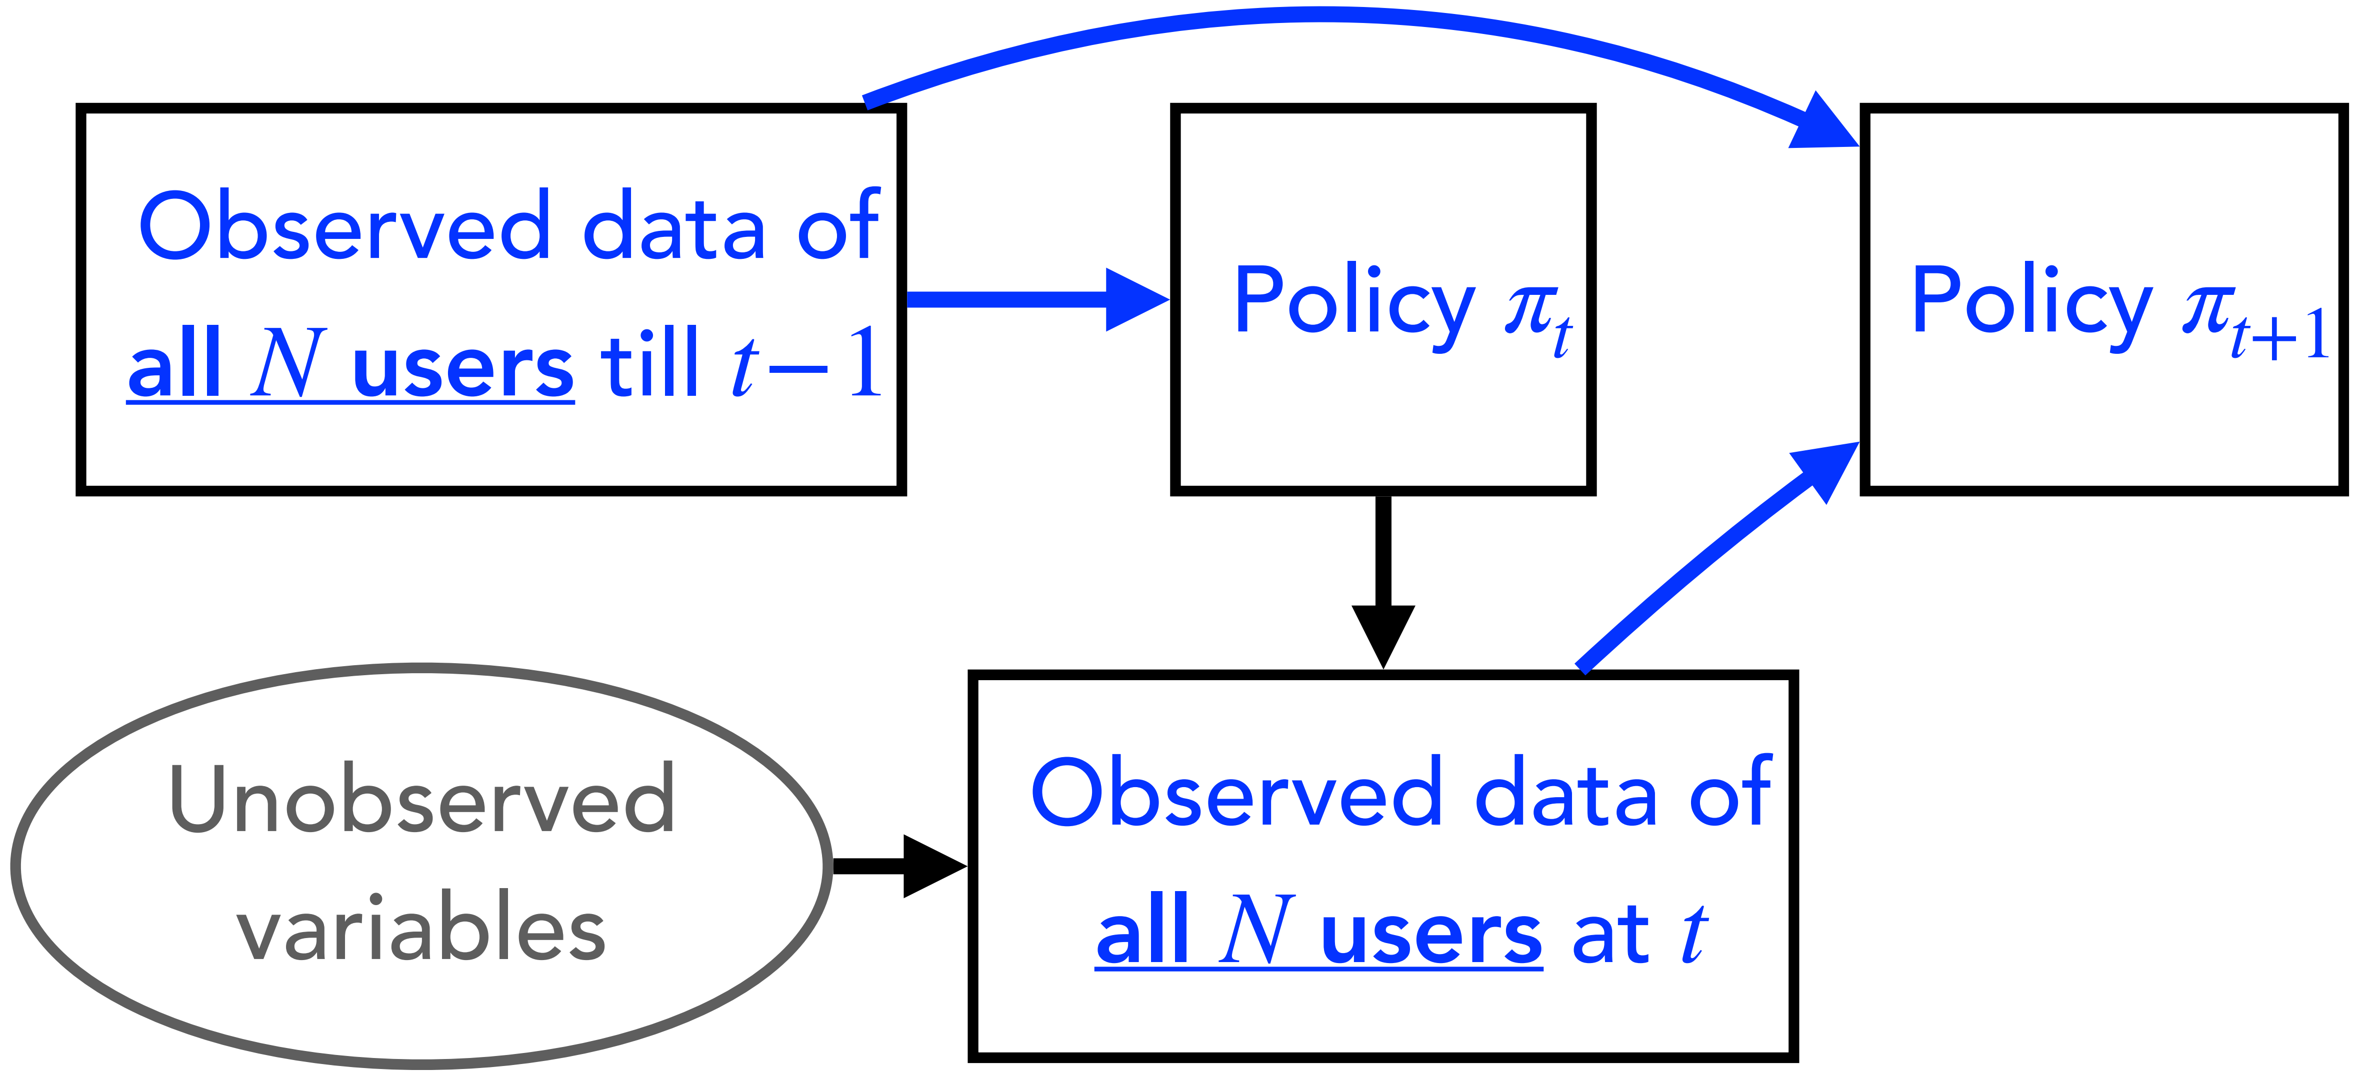

Modeling assumptions

(For each treatment separately)

$$Z_{i,t} = \begin{cases} Y^\star_{i,t} + \varepsilon_{i,t} & \text{if} \quad A_{i,t} = 1 \\ \text{unknown} & \text{if} \quad A_{i,t} = 0 \end{cases}$$

Latent factor model:

$$Y^\star_{i,t} = f(u_i, v_t)$$

Algorithm

Nearest neighbor estimate $\widehat{Y}_{i,t}$ for $Y^\star_{i,t}$ by measuring distance over time

Guarantees

Bilinear f → Any-time consistency $\widehat{Y}_{i,t} \to Y^\star_{i,t}$

Lipschitz non-linear f → Sequential CLT $\sqrt{N_{i,T}}(\widehat{Y}_{i,T} - Y^\star_{i,T}) \Rightarrow \mathcal{N}(0, \sigma^2)$

Coming this fall…

- **Confidence intervals for average treatment effects**
- **An improved ``doubly-robust'' variant of nearest neighbors**
- Several future directions:
  - Use of covariates/contexts
  - Temporal structure (dynamical system)
  - Leveraging information across treatments

Thank you!

Counterfactual inference in sequential experiments

https://rzrsk.github.io

Sequential experimental design

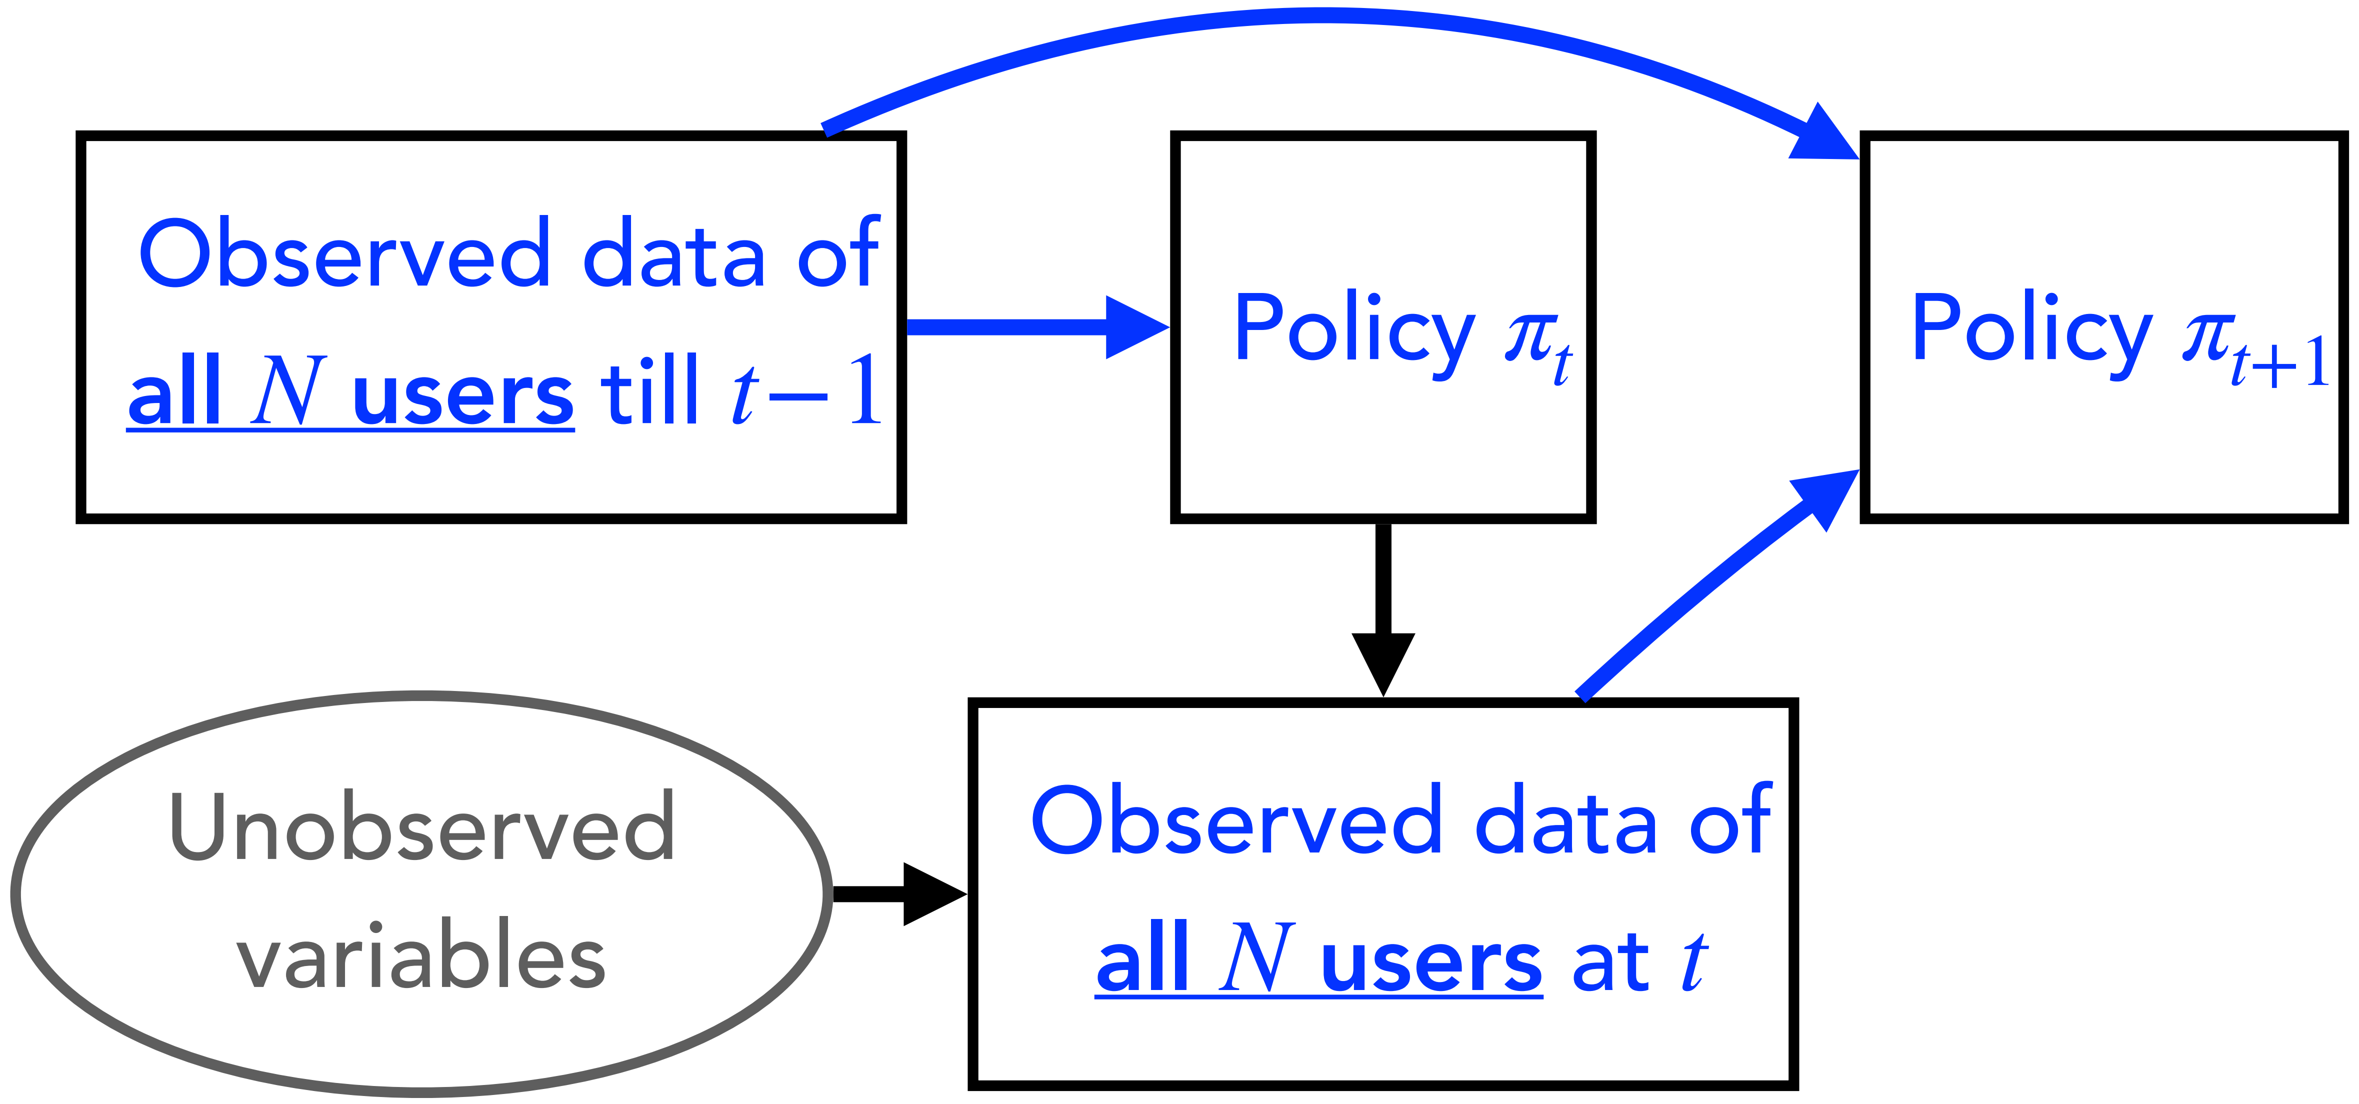

Modeling assumptions

(For each treatment separately)

$$Z_{i,t} = \begin{cases} Y^\star_{i,t} + \varepsilon_{i,t} & \text{if} \quad A_{i,t} = 1 \\ \text{unknown} & \text{if} \quad A_{i,t} = 0 \end{cases}$$

Latent factor model:

$$Y^\star_{i,t} = f(u_i, v_t)$$

Algorithm

Nearest neighbor estimate $\widehat{Y}_{i,t}$ for $Y^\star_{i,t}$ by measuring distance over time

Guarantees

Bilinear f → Any-time consistency $\widehat{Y}_{i,t} \rightarrow Y^\star_{i,t}$

Lipschitz non-linear f → Sequential CLT $\sqrt{N_{i,T}}(\widehat{Y}_{i,T} - Y^\star_{i,T}) \Rightarrow \mathcal{N}(0, \sigma^2)$

Additional slides

Explicit non-asymptotic bound for the bilinear case

$f(u, v) = \langle u, v \rangle$: A deterministic error bound for any (i, t)

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

- $N_{i,t}$ = nearest neighbors for (i, t)

- Deterministic error bound:

$$(\widehat{Y}_{i,t} - Y^{\star}_{i,t})^2 \lesssim \frac{\|v_t\|_2^2}{\lambda} \left(\eta - c\sigma^2 + \frac{\langle \text{noise}, \{v_t\} \rangle}{\min_{j \in N_{i,t}} T_{t,i,j}} \right) + \left(\frac{\sum_{j \in N_{i,t}} \text{noise}_{j,t}}{|N_{i,t}|} \right)^2$$

Gaussian process as a bilinear latent factor model

Latent factor model

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

u_i : latent factor for user i

v_t : latent factor for time t

f : unknown (non) linear function

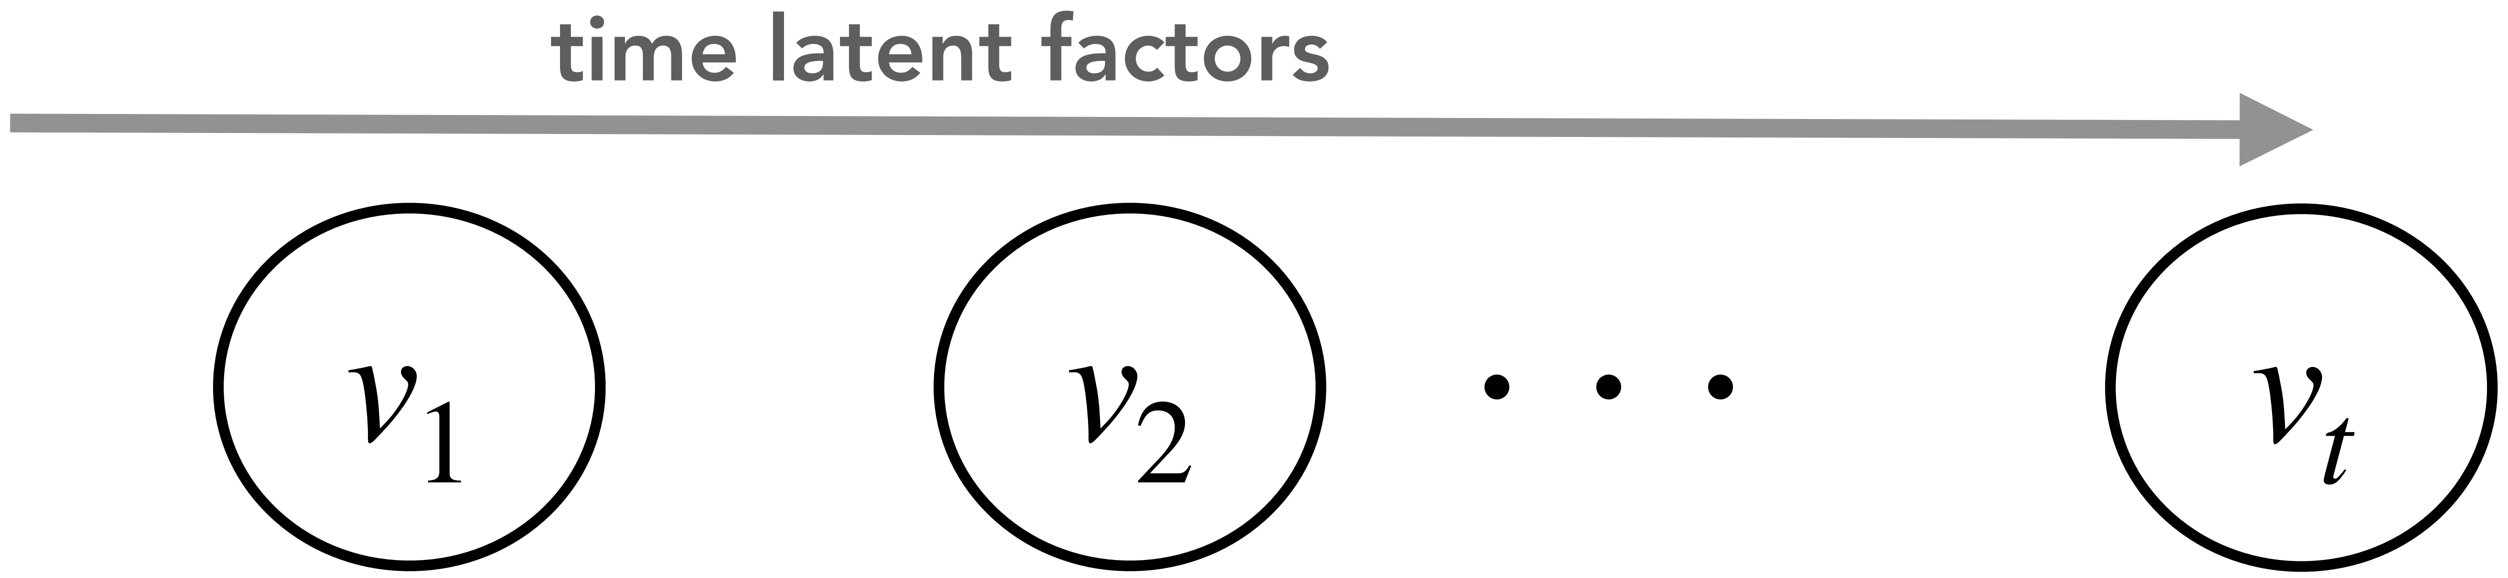

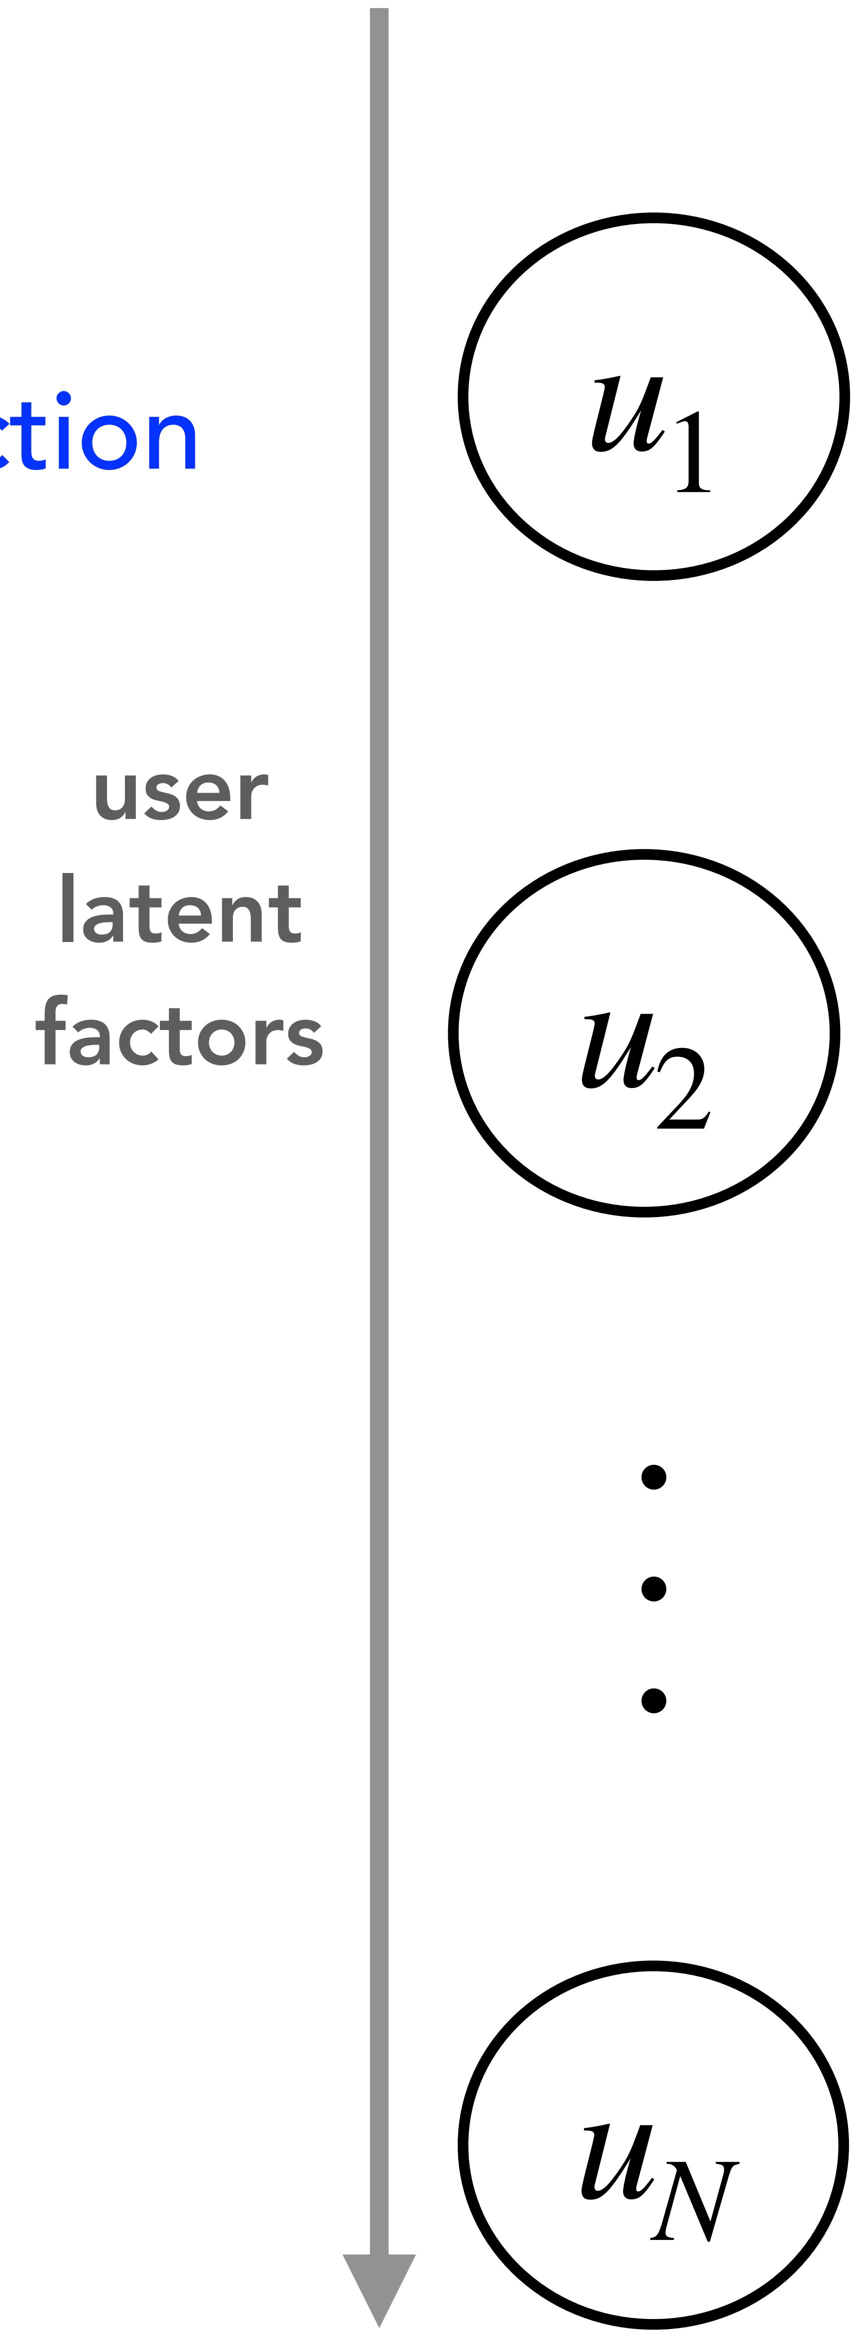

- $Y^{\star}_{i,\cdot} \sim$ Gaussian process with mean m in the reproducing kernel Hilbert space of the covariance kernel $\mathbf{k}$, which has eigenfunctions ϕ_j then

$$f(u, v) = \langle u, v \rangle,$$

$$u_i = (\langle m, \phi_j \rangle)_{j=1}^{r} + \mathcal{N}(0, I_r), \text{ and}$$

$$v_t = (\phi_j(x_t))_{j=1}^{r}.$$

Gaussian process as a latent factor model

- If each user's data is a sample from $\mathcal{GP}(0,\mathbf{k})$ where $\mathbf{k}$ is Mercer's kernel such that
$$\mathbf{k}(t_1, t_2) = \sum_{\ell=1}^{\infty} \lambda_\ell \phi_\ell(t_1)\phi_\ell(t_2),$$
where λ_ℓ, ϕ_ℓ denote eigenvalue-eigenfunctions with $\{\phi_\ell\}$ orthonormal

- Then for $\xi_{i,\ell} \sim_{iid} \mathcal{N}(0,1)$, we have
$$Y_{i,t} = \sum_{\ell=1}^{\infty} \xi_{i,\ell}\sqrt{\lambda_\ell}\phi_\ell(t) \text{ almost surely } \implies Y_{i,t} = f(u_i, v_t) = \langle u_i, v_t \rangle$$
for $u_i = (a_1, a_2, \ldots) + (\xi_{i,1}, \xi_{i,2}, \ldots)$, and $v_t = (\sqrt{\lambda_1}\phi_1(t), \sqrt{\lambda_2}\phi_2(t), \ldots)$

Example: Exchangeable data

- The latent factor model also holds if the matrix $\{Y^{\star}_{i,t} + \varepsilon_{i,t}, i = 1,\ldots.N, t = 1,\ldots,T\}$ for a sub-class of exchangeable data i.e., exchangeable under row and column permutations

- See Sec II.C Li et al. 2017

Mathematical description of nearest neighbors

Nearest neighbors estimate for $Y^{\star}_{i,t}$

- **Input**: Partially observed matrix; **Output**: Estimate of noiseless entry (i, t)
- **Algorithm**: Compute
  - Available neighbors at time $t = \{j : A_{j,t} = 1\}$
  - Good neighbors for user i at time $t = \{j : \rho_t(i, j) \leq \eta\}$ where

$$\rho_t(i, j) = \frac{1}{\sum_{t' \neq t} A_{i,t'} A_{j,t'}} \sum_{t' \neq t} (Z_{i,t'} - Z_{j,t'})^2 A_{i,t'} A_{j,t'}$$

  - $\widehat{Y}_{i,t}$ = Simple average of $Y_{j,t}$ over $\{j : A_{j,t} = 1 \text{ and } \rho_t(i, j) \leq \eta\}$.

Further details about prior work

Data collected with a <u>fixed</u> policy: Off policy evaluation

- Set-up considers either
  - i.i.d. users, e.g., multi-armed bandits
  - or, one user over time, e.g., Markov decision proces
  - but ***not*** multiple users over multiple time
- **Quantities of interest:** Average reward under alternative policy, estimated using IPW-based estimates, switch estimators etc.

[…, Li et al 2015, Wang et al. 2021, Ma et al. 2021,…]

Diverse mobile health applications

Smoking addiction

Well-being

Wearable/trackers

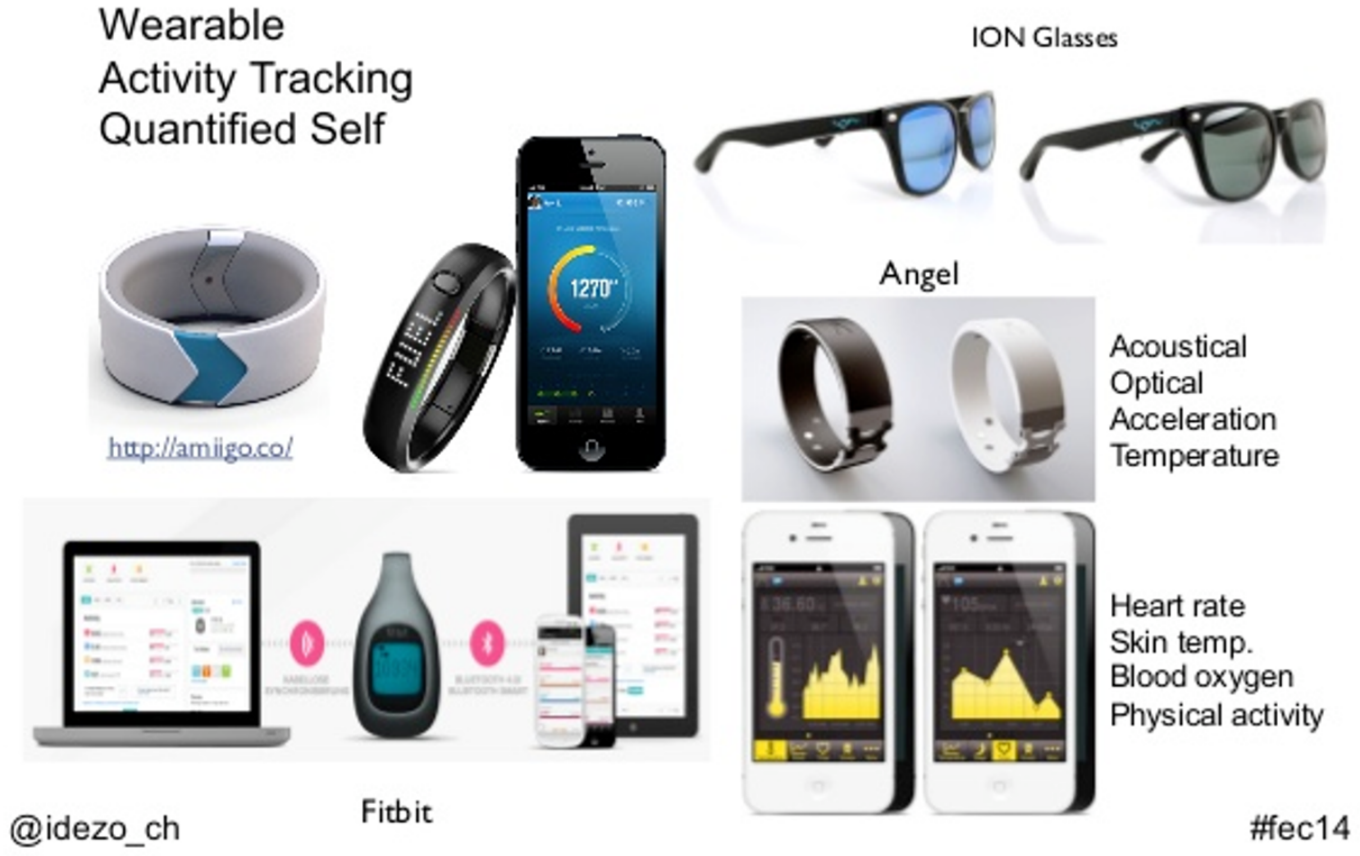

Binge drinking

Recovery support

Physical activity

Image credits: Susan Murphy