Bilinear Bandits with Low-rank Structure

Kwang-Sung Jun

Boston University (will join the U of Arizona)

Rebecca Willett

Stephen Wright

Robert Nowak

U of Chicago

UW-Madison

Application: Drug discovery

- Choose a pair to experiment with

Application: Drug discovery

drugs

proteins

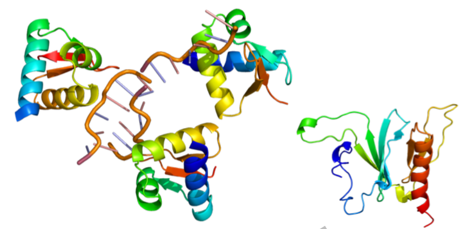

- Choose a pair to experiment with

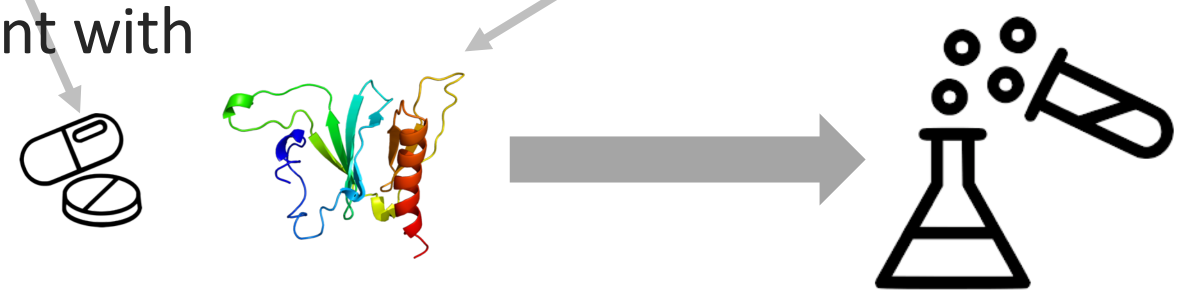

- Goal: Find as many pairs with the desired interaction as possible

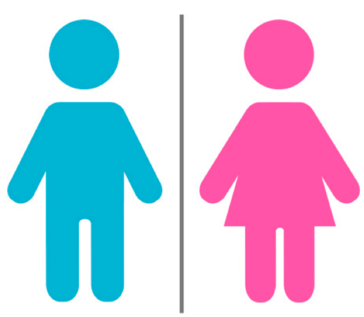

online dating

clothing recommendation

Bilinear bandits

$$y = x^\top \Theta z + \eta$$

desired interaction? (0/1) — y

drug features — x

unknown parameter (d by d) — Θ

protein features — z

- A natural model: already used for predicting drug-protein interaction.

[Luo et al., "A network integration approach for drug-target interaction prediction and computational drug repositioning from heterogeneous information", Nature Communications, 2017]

Bilinear bandits

$$y = x^{\top} \Theta z + \eta$$

- y: desired interaction? (0/1)
- x: drug features
- Θ: unknown parameter (d by d)
- z: protein features

- A natural model: already used for predicting drug-protein interaction.

[Luo et al., "A network integration approach for drug-target interaction prediction and computational drug repositioning from heterogeneous information", Nature Communications, 2017]

- Issue: d^2 number of unknowns
  - What if $\mathrm{rank}(\Theta) \ll d$?

$$\Theta = \sum_{k=1}^{r} \sigma_i u_i v_i^{\top} \quad \Longrightarrow \quad \sim dr \text{ unknowns}$$

  - Many real-world problems exhibit the low-rank structure.

Summary of the result

- A naïve method: reduction

$$\mathbb{E}[y] = x^\top \Theta z = \langle \mathrm{vec}(\Theta), \mathrm{vec}(xz^\top) \rangle$$

- Invoking linear algorithms [Abbasi-Yadkori'11], convergence rate is

$$\frac{d^2}{\sqrt{T}}$$

  - No dependence of the rank r

Summary of the result

- A naïve method: reduction

$$\mathbb{E}[y] = x^\top \Theta z = \langle \mathrm{vec}(\Theta), \mathrm{vec}(xz^\top) \rangle$$

- Invoking linear algorithms [Abbasi-Yadkori'11], convergence rate is

$$\frac{d^2}{\sqrt{T}}$$

  - No dependence of the rank r

- Can we obtain faster rates as the rank r becomes smaller?

Summary of the result

- A naïve method: reduction

$$\mathbb{E}[y] = x^\top \Theta z = \langle \mathrm{vec}(\Theta), \mathrm{vec}(xz^\top) \rangle$$

- Invoking linear algorithms [Abbasi-Yadkori'11], convergence rate is

$$\frac{d^2}{\sqrt{T}}$$

  - No dependence of the rank r

- Can we obtain faster rates as the rank r becomes smaller?

YES, we achieve $\frac{d^{3/2}\sqrt{r}}{\sqrt{T}}$ (factor $\sqrt{d/r}$ better)

  - Is this optimal?

Explore-Subspace-Then-Refine (ESTR)

- Stage 1: estimate the subspace

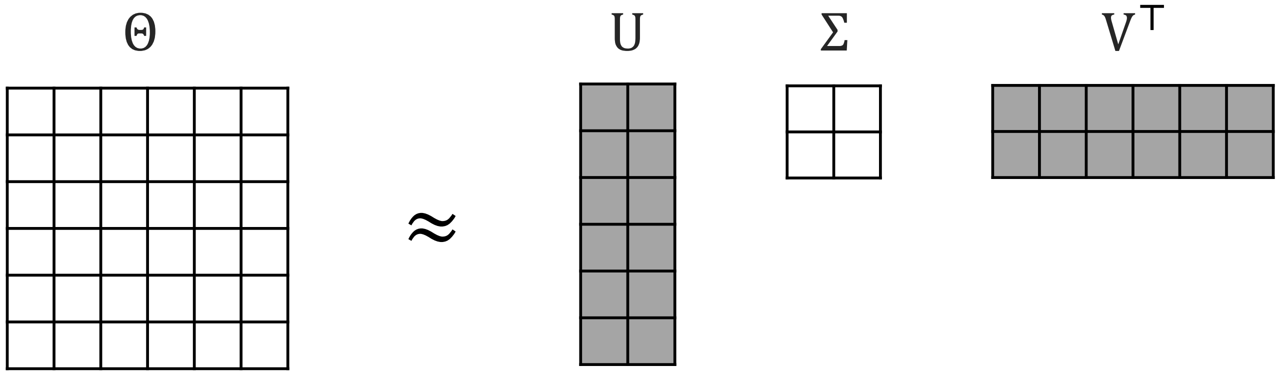

- Stage 2: linear bandit **within** the “subspace”

Explore-Subspace-Then-Refine (ESTR)

- Stage 1: estimate the subspace

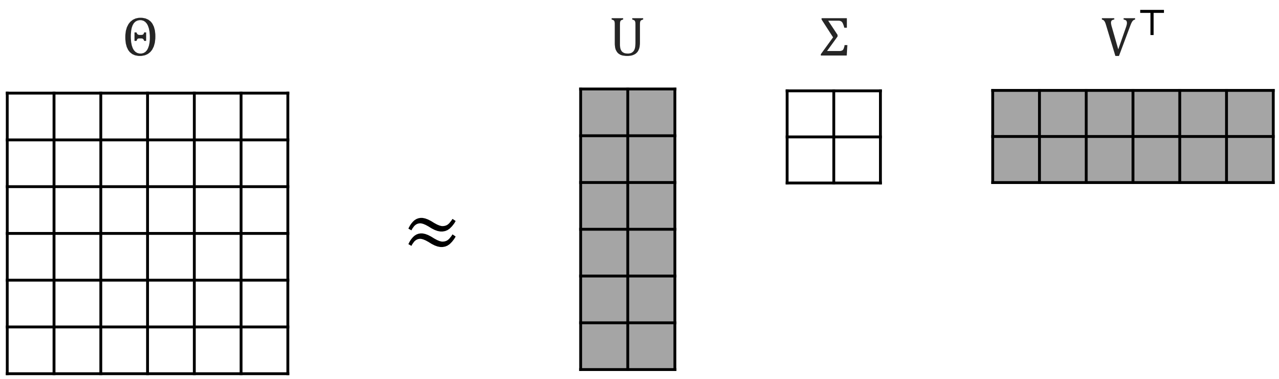

- Stage 2: linear bandit **within** the "subspace"
  - Turns out, it doesn't work.
  - Our solution: allow **"refining"** the subspace.
  - The devil is in the detail.

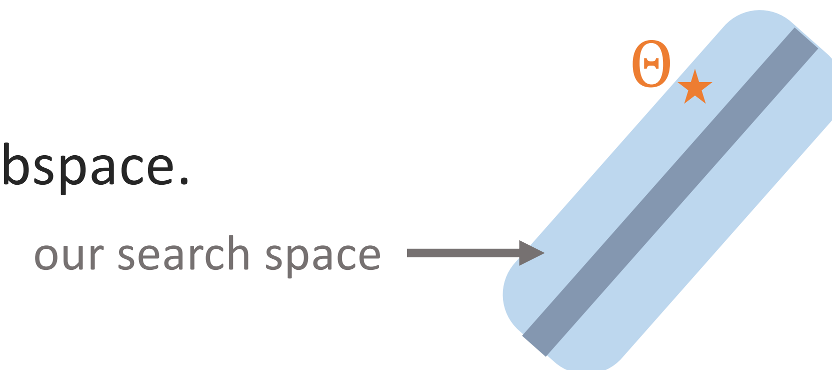

Explore-Subspace-Then-Refine (ESTR)

- Stage 1: estimate the subspace

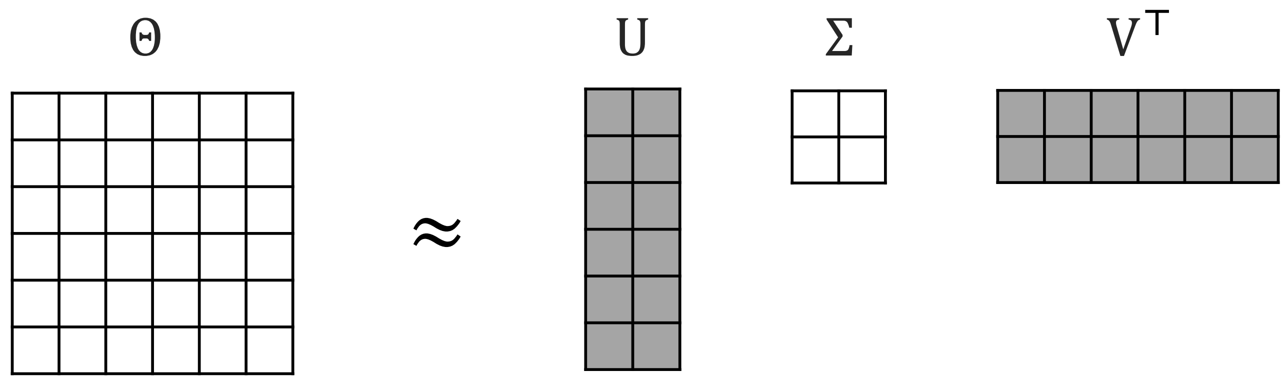

- Stage 2: linear bandit **within** the "subspace"
  - Turns out, it doesn't work.
  - Our solution: allow **"refining"** the subspace.
  - The devil is in the detail.

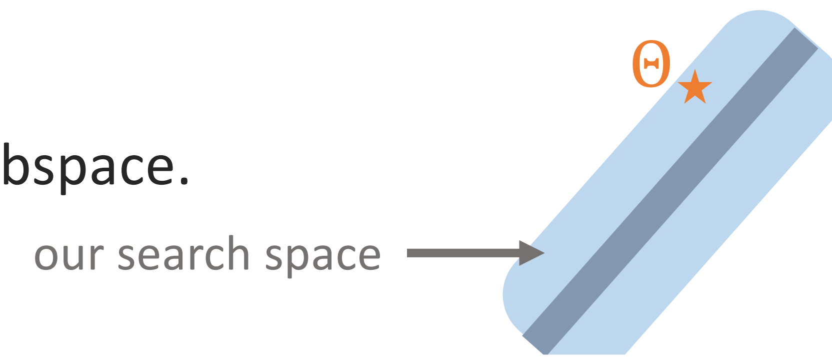

Let's chat! #127 @ Pacific Ballroom