Causal Bandits for Linear Structural Equation Models

Burak Varıcı[1] Karthikeyan Shanmugam[2] Prasanna Sattigeri[3] Ali Tajer[4]

[1]Carnegie Mellon University

[2]Google India

[3]IBM Research

[4]Rensselaer Polytechnic Institute

NeurIPS 2024 Journal-to-conference track

Original publication: Journal of Machine Learning Research, 24(297), 1-59. 2023

Sequential Design of Interventions

designing an optimal sequence of interventions

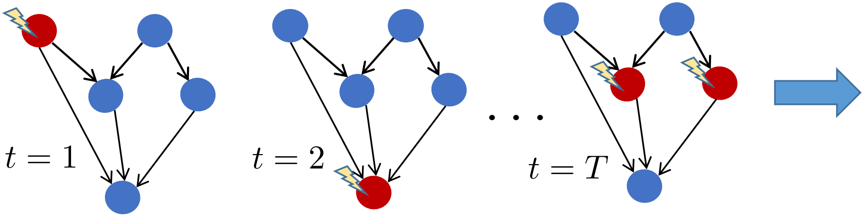

Questions:

- **structure learning:** which interventions reveal the most information?
- **identifying causal effect:** what is the smallest set of interventions needed?

Causal Bandits Problem Statement

Sequential selection of interventions: for $t \in \{1, \dots, T\}$

- **objective:** choosing $a \in \mathcal{A}$ to maximize the **expected reward** $\mu_a \triangleq \mathbb{E}_{\mathbb{P}_a}[Y]$.
- Learner selects $a_t \in \mathcal{A}$: intervene on a subset of nodes.
- Learner observes reward $Y(t) \sim \mathbb{P}_{a_t}(Y)$ and graph instance $X(t) \sim \mathbb{P}_{a_t}(X)$.

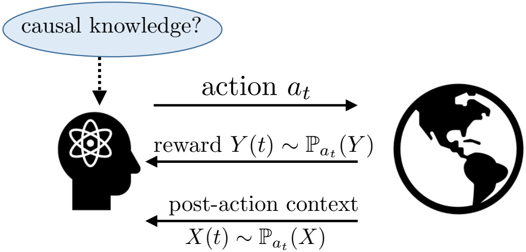

Regret minimization

Minimize expected cumulative regret

$$\mathbb{E}[R(T)] = T\mu_{a^*} - \mathbb{E}\left[\sum_{t=1}^{T} Y(t)\right]$$

Causal Bandits with Linear SEM

- Critical information for algorithm design

  (i) causal structure $\mathcal{G}$ **(ii) interventional distributions** $\{\mathbb{P}_a\}_{a\in\mathcal{A}}$

- Our focus: known structure $\mathcal{G}$, **unknown** distributions
- Prior work binary random variables and/or atomic interventions.

- **Linear Structural Equation Model:**

$$X = B^\top X + \epsilon, \quad \epsilon = (\epsilon_1, \ldots, \epsilon_N), \quad B \in \mathbb{R}^{N\times N}$$

- **Intervention model:** Soft intervention on node i changes the weights $[B]_i$ to $[B^*]_i$
  - Two mechanisms for each node
  - **Intervention space** $\mathcal{A} = 2^{\mathbf{V}}$, so $|\mathcal{A}| = 2^N$.

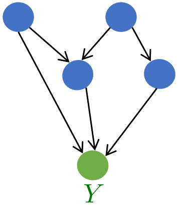

- **Graph parameters:**
  - d: maximum in-degree
  - L: length of the longest causal path.

LinSEM-UCB Algorithm

Algorithm: UCB-based + leverage causal relationships

non-causal algorithms
$\mathcal{O}(\sqrt{|\mathcal{A}|T})$

↓

unknown $\mathbb{P}_a$
$\tilde{\mathcal{O}}(d^{L+\frac{1}{2}}\sqrt{NT})$

↑

known $\mathbb{P}_a$
$\tilde{\mathcal{O}}(d\sqrt{T})$

Theorem (Upper bound)

$$\mathbb{E}[R(T)] = \tilde{\mathcal{O}}\big(d^{L+\frac{1}{2}}\sqrt{NT}\big)$$

Theorem (Minimax Lower bound)

$$\mathbb{E}[R(T)] \geq \Omega(d^{\frac{L}{2}-2}\sqrt{T}).$$

Follow-up Results

Recent developments in causal bandits with soft interventions:

- Yan et al. (JSAIT 2024): *Robust Causal Bandits for Linear Models*.
  Known graph, linear SEMs: Similar regret results for varying weights, robust
- Yan et al. (AISTATS 2024): *Causal Bandits with General Causal Models and Interventions*.
  Known graph, general SCMs: Similar regret results for general causal models
- Yan et al. (NeurIPS 2024): *Linear Causal Bandits: Unknown Graph and Soft Interventions*
  Unknown graph, linear SEMs: almost matching lower and upper bounds at $\mathcal{O}(d^{L/2})$!

At NeurIPS 2024

Causal Bandits for Linear Structural Equation Models

- Conference: Poster Session 2, Wednesday, Dec 11, 4.30-7.30pm
- Paper: https://jmlr.org/papers/v24/22-0969.html
- Code: https://github.com/bvarici/causal-bandits-linear-sem
- Contact: bvarici@andrew.cmu.edu