

Combinatorial Causal Bandits

Shi Feng,¹ Wei Chen²

¹ Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China

² Microsoft Research, Beijing, China

shifeng-thu@outlook.com, weic@microsoft.com

Abstract

In combinatorial causal bandits (CCB), the learning agent chooses at most K variables in each round to intervene, collects feedback from the observed variables, with the goal of minimizing expected regret on the target variable Y . We study under the context of binary generalized linear models (BGLMs) with a succinct parametric representation of the causal models. We present the algorithm BGLM-OFU for Markovian BGLMs (i.e. no hidden variables) based on the maximum likelihood estimation method, and show that it achieves $O(\sqrt{T} \log T)$ regret, where T is the time horizon. For the special case of linear models with hidden variables, we apply causal inference techniques such as the do-calculus to convert the original model into a Markovian model, and then show that our BGLM-OFU algorithm and another algorithm based on the linear regression both solve such linear models with hidden variables. Our novelty includes (a) considering the combinatorial intervention action space and the general causal models including ones with hidden variables, (b) integrating and adapting techniques from diverse studies such as generalized linear bandits and online influence maximization, and (c) avoiding unrealistic assumptions (such as knowing the joint distribution of the parents of Y under all interventions) and regret factors exponential to causal graph size in prior studies.

1 Introduction

Causal bandit problem is first introduced in (Lattimore, Lattimore, and Reid 2016). It consists of a causal graph $G = (\mathbf{X} \cup \{Y\}, E)$ indicating the causal relationship among the observed variables, where the structure of the graph is known but the underlying probability distributions governing the causal model are unknown. In each round, the learning agent selects one or a few variables in \mathbf{X} to intervene, gains the reward as the output of Y , and observes the values of all variables in $\mathbf{X} \cup \{Y\}$. The goal is to either maximize the cumulative reward over T rounds, or find the intervention closest to the optimal one after T rounds. The former setting, which is the one our study focuses on, is typically translated to minimizing cumulative regret, which is defined as the difference in reward between always playing the optimal intervention and playing interventions according to a learning algorithm. Causal bandits can be applied in many

settings that include causal relationships, such as medical drug testing, policy making, scientific experimental process, performance tuning, etc. Causal bandit is a form of multi-armed bandit (cf. (Lattimore and Szepesvári 2020)), with the main difference being that causal bandits may use the causal relationship and more observed feedback to achieve a better regret.

All the causal bandit studies so far (Lattimore, Lattimore, and Reid 2016; Sen et al. 2017; Nair, Patil, and Sinha 2021; Lu et al. 2020; Maiti, Nair, and Sinha 2021) focus on the case where the number of possible interventions is small. However, in many scenarios we need to intervene on a set of variables and the number of possible choices of sets is large. For example, in tuning performance of a large system, one often needs to tune a large number of parameters simultaneously to achieve the best system performance, and in drug testing, a combination of several drugs with a number of possible dosages needs to be tested for the best result. Intervening on a set of variables raises new challenges to the learning problem, since the number of possible interventions is exponentially large to the size of the causal graph. In this paper, we address this challenge and propose the new framework of combinatorial causal bandits (CCB) and its solutions. In each round of CCB, the learning agent selects a set of at most K observed variables in \mathbf{X} to intervene instead of one variable. Other aspects including the feedback and reward remain the same.

We use the binary generalized linear models (BGLMs) to give the causal model a succinct parametric representation where all variables are binary. Using the maximum likelihood estimation (MLE) method, we design an online learning algorithm BGLM-OFU for the causal models without hidden variables (called Markovian models), and the algorithm achieves $O(\sqrt{T} \log T)$ regret, where T is the time horizon. The algorithm is based on the earlier study on generalized linear bandits (Li, Lu, and Zhou 2017), but our BGLM model is more general and thus requires new techniques (such as a new initialization phase) to solve the problem. Our regret analysis also integrates results from online influence maximization (Li et al. 2020; Zhang et al. 2022) in order to obtain the final regret bound.

Furthermore, for the binary linear models (BLMs), we show how to transform a BLM with hidden variables into one without hidden variables by utilizing the tools in causal

inference such as do-calculus, and thus we can handle causal model even with hidden variables in linear models. Then, for BLMs, we show (a) the regret bound when applying the BGLM-OFU algorithm to the linear model, and (b) a new algorithm and its regret bound based on the linear regression method. We show a tradeoff between the MLE-based BGLM-OFU algorithm and the linear-regression-based algorithm on BLMs: the latter removes the assumption needed by the former but has an extra factor in the regret bound. We demonstrate effectiveness of our algorithms by experimental evaluations in appendix. Besides BLMs, we give similar results for linear models with continuous variables. Due to space limits, this part is put in appendix. Due to space limits, our appendix is included in the full version (Feng and Chen 2022) on arXiv.

In summary, our contribution includes (a) proposing the new combinatorial causal bandit framework, (b) considering general causal models including ones with hidden variables, (c) integrating and adapting techniques from diverse studies such as generalized linear bandits and online influence maximization, and (d) achieving competitive online learning results without unrealistic assumptions and regret factors exponential to the causal graph size as appeared in prior studies. The intrinsic connection between causal bandits and online influence maximization revealed by this study may further benefit researches in both directions.

2 Related Works

Causal Bandits. The causal bandit problem is first defined in (Lattimore, Lattimore, and Reid 2016). The authors introduce algorithms for a specific parallel model and more general causal models to minimize the simple regret, defined as the gap between the optimal reward and the reward of the action obtain after T rounds. A similar algorithm to minimize simple regret has also been proposed in (Nair, Patil, and Sinha 2021) for a graph with no backdoor. To optimally trade-off observations and interventions, they have also discussed the budgeted version of causal bandit when interventions are more expensive than observations. Sen et al. (2017) use the importance sampling method to propose an algorithm to optimize simple regret when only soft intervention on a single node is allowed. These studies all focus on simple regret for the pure exploration setting, while our work focuses on cumulative regret. There are a few studies on the cumulative regret of causal bandits (Lu et al. 2020; Nair, Patil, and Sinha 2021; Maiti, Nair, and Sinha 2021). However, the algorithms in (Lu et al. 2020) and (Nair, Patil, and Sinha 2021) are both under an unrealistic assumption that for every intervention, the joint distribution of the parents of the reward node Y is known. The algorithm in (Maiti, Nair, and Sinha 2021) does not use this assumption, but it only works on Markovian models without hidden variables. The regrets in (Lu et al. 2020; Maiti, Nair, and Sinha 2021) also have a factor exponential to the causal graph size. Yabe et al. (2018) designs an algorithm by estimating each $P(X|\mathbf{Pa}(X))$ for $X \in \mathbf{X} \cup \{Y\}$ that focuses on simple regret but works for combinatorial settings. However, it requires the round number $T \geq \sum_{X \in \mathbf{X}} 2^{|\mathbf{Pa}(X)|}$, which is still unrealistic. We

avoid this by working on the BGLM model with a linear number of parameters, and it results in completely different techniques. Lee and Bareinboim (2018, 2019, 2020) also consider the combinatorial settings, but they focus on empirical studies, while we provide theoretical regret guarantees. Furthermore, studying causal bandit problem without the full casual structure is an important future direction. One such study, (Lu, Meisami, and Tewari 2021), exists but it is only for the atomic setting and has a strong assumption that $|\mathbf{Pa}(Y)| = 1$.

Combinatorial Multi-Armed Bandits. CCB can be viewed as a type of combinatorial multi-armed bandits (CMAB) (Chen, Wang, and Yuan 2013; Chen et al. 2016), but the feedback model is not the semi-bandit model studied in (Chen, Wang, and Yuan 2013; Chen et al. 2016). In particular, in CCB individual random variables cannot be treated as base arms, because each intervention action changes the distribution of the remaining variables, violating the i.i.d assumption for base arm random variables across different rounds. Interestingly, it has a closer connection to recent studies on online influence maximization (OIM) with node-level feedback (Li et al. 2020; Zhang et al. 2022). These studies consider influence cascades in a social network following either the independent cascade (IC) or the linear threshold (LT) model. In each round, one selects a set of at most K seed nodes, observes the cascade process as which nodes are activated in which steps, and obtains the reward as the total number of nodes activated. The goal is to learn the cascade model and minimize the cumulative regret. Influence propagation is intrinsically a causal relationship, and thus OIM has some intrinsic connection with CCB, and our study does borrow techniques from (Li et al. 2020; Zhang et al. 2022). However, there are a few key differences between OIM and CCB, such that we need adaptation and integration of OIM techniques into our setting: (a) OIM does not consider hidden variables, while we do consider hidden variables for causal graphs; (b) OIM allows the observation of node activations at every time step, but in CCB we only observe the final result of the variables; and (c) current studies in (Li et al. 2020; Zhang et al. 2022) only consider IC and LT models, while we consider the more general BGLM model, which includes IC model (see (Zhang et al. 2022) for transformation from IC model to BGLM) and LT model as two special cases.

Linear and Generalized Linear Bandits. Our work is also based on some previous online learning studies on linear and generalized linear bandits. Abbasi-Yadkori, Pál, and Szepesvári (2011) propose an improved theoretical regret bound for the linear stochastic multi-armed bandit problem. Some of their proof techniques are used in our proofs. Li, Lu, and Zhou (2017) propose an online algorithm and its analysis based on MLE for generalized linear contextual bandits, and our MLE method is adapted from this study. However, our setting is much more complicated, in that (a) we have a combinatorial action space, and (b) we have a network of generalized linear relationships while they only consider one generalized linear relationship. As the result, our algorithm and analysis are also more sophisticated.

3 Model and Preliminaries

Following the convention in causal inference literature (e.g. Pearl 2009)), we use capital letters ($U, X, Y \dots$) to represent variables, and their corresponding lower-case letters to represent their values. We use boldface letters such as \mathbf{X} and \mathbf{x} to represent a set or a vector of variables or values.

A *causal graph* $G = (U \cup \mathbf{X} \cup \{Y\}, E)$ is a directed acyclic graph where $U \cup \mathbf{X} \cup \{Y\}$ are sets of nodes with U being the set of unobserved or hidden nodes, $\mathbf{X} \cup \{Y\}$ being the set of observed nodes, Y is a special target node with no outgoing edges, and E is the set of directed edges connecting nodes in $U \cup \mathbf{X} \cup \{Y\}$. For simplicity, in this paper we consider all variables in $U \cup \mathbf{X} \cup \{Y\}$ are $(0, 1)$ -binary random variables. For a node X in the graph G , we call its in-neighbor nodes in G the *parents* of X , denoted as $\mathbf{Pa}(X)$, and the values taken by these parent random variables are denoted $\mathbf{pa}(X)$.

Following the causal Bayesian model, the causal influence from the parents of X to X is modeled as the probability distribution $P(X|\mathbf{Pa}(X))$ for every possible value combination of $\mathbf{Pa}(X)$. Then, for each X , the full non-parametric characterization of $P(X|\mathbf{Pa}(X))$ requires $2^{|\mathbf{Pa}(X)|}$ values, which would be difficult for learning. Therefore, we will describe shortly the succinct parametric representation of $P(X|\mathbf{Pa}(X))$ as a generalized linear model to be used in this paper.

The causal graph G is *Markovian* if there are no hidden variables in G and every observed variable X has certain randomness not caused by other observed parents. To model this effect of the Markovian model, in this paper we dedicate random variable X_1 to be a special variable that always takes the value 1 and is a parent of all other observed random variables. We study the Markovian causal model first, and in Section 5 we will consider causal models with more general hidden variable forms.

An *intervention* in the causal model G is to force a subset of observed variables $\mathbf{S} \subseteq \mathbf{X}$ to take some predetermined values \mathbf{s} , to see its effect on the target variable Y . The intervention on \mathbf{S} to \mathbf{Y} is denoted as $\mathbb{E}[Y|do(\mathbf{S} = \mathbf{s})]$. In this paper, we study the selection of \mathbf{S} to maximize the expected value of Y , and our parametric model would have the monotonicity property such that setting a variable X to 1 is always better than setting it to 0 in maximizing $\mathbb{E}[Y]$, so our intervention would always be setting \mathbf{S} to all 1's, for which we simply denote as $\mathbb{E}[Y|do(\mathbf{S})]$.

In this paper, we study the online learning problem of *combinatorial causal bandit*, as defined below. A learning agent runs an algorithm π for T rounds. Initially, the agent knows the observed part of the causal graph G induced by observed variables $\mathbf{X} \cup \{Y\}$, but does not know the probability distributions $P(X|\mathbf{Pa}(X))$'s. In each round $t = 1, 2, \dots, T$, the agent selects at most K observed variables in \mathbf{X} for intervention, obtains the observed Y value as the reward, and observes all variable values in $\mathbf{X} \cup \{Y\}$ as the feedback. The agent needs to utilize the feedback from the past rounds to adjust its action in the current round, with the goal of maximizing the cumulative reward from all T rounds.

The performance of the agent is measured by the *regret* of the algorithm π , which is defined as the difference between the expected cumulative reward of the best action \mathbf{S}^* and the cumulative reward of algorithm π , where $\mathbf{S}^* \in \operatorname{argmax}_{\mathbf{S} \subseteq \mathbf{X}, |\mathbf{S}|=K} \mathbb{E}[Y|do(\mathbf{S})]$, as given below:

$$R^\pi(T) = \mathbb{E} \left[\sum_{t=1}^T (\mathbb{E}[Y|do(\mathbf{S}^*)] - \mathbb{E}[Y|do(\mathbf{S}_t^\pi)]) \right], \quad (1)$$

where \mathbf{S}_t^π is the intervention set selected by algorithm π in round t , and the expectation is taking from the randomness of the causal model as well as the possible randomness of the algorithm π . In some cases our online learning algorithm uses a standard offline oracle that takes the causal graph G and the distributions $P(X|\mathbf{Pa}(X))$'s as inputs and outputs a set of nodes \mathbf{S} that achieves an α -approximation with probability β with $\alpha, \beta \in (0, 1]$. In this case, we consider the (α, β) -approximation regret, as in (Chen et al. 2016):

$$R_{\alpha, \beta}^\pi(T) = \mathbb{E} \left[\sum_{t=1}^T (\alpha\beta \mathbb{E}[Y|do(\mathbf{S}^*)] - \mathbb{E}[Y|do(\mathbf{S}_t^\pi)]) \right]. \quad (2)$$

As pointed out earlier, the non-parametric form of distributions $P(X|\mathbf{Pa}(X))$'s needs an exponential number of quantities and is difficult to learn. In this paper, we adopt the general linear model as the parametric model, which is widely used in causal inference literature (Hernán and Robins 2010; Garcia-Huidobro and Michael Oakes 2017; Han, Yu, and Friedberg 2017; Arnold et al. 2020; Vansteelandt and Dukes 2020). Since we consider binary random variables, we refer to such models as binary general linear models (BGLMs). In BGLM, the functional relationship between a node X and its parents $\mathbf{Pa}(X)$ in G is $P(X = 1|\mathbf{Pa}(X) = \mathbf{pa}(X)) = f_X(\boldsymbol{\theta}_X^* \cdot \mathbf{pa}(X)) + \varepsilon_X$, where $\boldsymbol{\theta}_X^*$ is the unknown weight vector in $[0, 1]^{|\mathbf{Pa}(X)|}$, ε_X is a zero-mean sub-Gaussian noise that ensures that the probability does not exceed 1, $\mathbf{pa}(X)$ here is the vector form of the values of parents of X , and f_X is a scalar and monotonically non-decreasing function. It is worth noticing that our BGLM here is a binary version of traditional generalized linear models (Aitkin et al. 2009; Hilbe 2011; Sakate and Kashid 2014; Hastie and Pregibon 2017). Instead of letting $X = f_X(\boldsymbol{\theta}_X^* \cdot \mathbf{pa}(X)) + \varepsilon_X$ directly, we take $f_X(\boldsymbol{\theta}_X^* \cdot \mathbf{pa}(X)) + \varepsilon_X$ as the probability for X to be 1. The vector of all the weights in a BGLM is denoted by $\boldsymbol{\theta}^*$ and the feasible domain for the weights is denoted by Θ . We use the notation $\theta_{Z, X}^*$ to denote the parameter in $\boldsymbol{\theta}^*$ that corresponds to the edge (Z, X) , or equivalently the entry in vector $\boldsymbol{\theta}_X^*$ that corresponds to node Z . We also use notation ε to represent all noise random variables $(\varepsilon_X)_{X \in \mathbf{X} \cup Y}$.

A special case of BGLM is the linear model where function f_X is the identity function for all X 's, then $P(X = 1|\mathbf{Pa}(X) = \mathbf{pa}(X)) = \boldsymbol{\theta}_X^* \cdot \mathbf{pa}(X) + \varepsilon_X$. We refer to this model as BLM. Moreover, when we remove the noise variable ε_X , BLM coincides with the *linear threshold (LT)* model for influence cascades (Kempe, Kleinberg, and Tardos 2003) in a DAG. In the LT model, each node X has a random threshold λ_X uniformly drawn from $[0, 1]$, and each

edge (Z, X) has a weight $w_{Z,X} \in [0, 1]$, such that node X is activated (equivalent to X being set to 1) when the cumulative weight of its active in-neighbors is at least λ_X . It is easy to see that when we set $\theta_{Z,X}^* = w_{Z,X}$, the activation condition is translated to the probability of $X = 1$ being exactly $\theta_X^* \cdot \mathbf{pa}(X)$. It is not surprising that a linear causal model is equivalent to an influence cascade model, since the influence relationship is intrinsically a causal relationship.

4 Algorithm for BGLM CCB

In this section, we present an algorithm that solves the on-line CCB problem for the Markovian BGLM. The algorithm requires three assumptions.

Assumption 1. For every $X \in \mathbf{X} \cup \{Y\}$, f_X is twice differentiable. Its first and second order derivatives are upper-bounded by $L_{f_X}^{(1)} > 0$ and $L_{f_X}^{(2)} > 0$.

Let $\kappa = \inf_{X \in \mathbf{X} \cup \{Y\}, \mathbf{v} \in [0,1]^{|Pa(X)|}, \|\theta - \theta_X^*\| \leq 1} \dot{f}_X(\mathbf{v} \cdot \theta)$.

Assumption 2. We have $\kappa > 0$.

Assumption 3. There exists a constant $\zeta > 0$ such that for any $X \in \mathbf{X} \cup \{Y\}$ and $X' \in Pa(X)$, for any value vector $\mathbf{v} \in \{0, 1\}^{|Pa(X) \setminus \{X', X_1\}|}$, the following inequalities hold:

$$\Pr_{\epsilon, \mathbf{X}, Y} \{X' = 1 | Pa(X) \setminus \{X', X_1\} = \mathbf{v}\} \geq \zeta, \quad (3)$$

$$\Pr_{\epsilon, \mathbf{X}, Y} \{X' = 0 | Pa(X) \setminus \{X', X_1\} = \mathbf{v}\} \geq \zeta. \quad (4)$$

The first two assumptions for our BGLM are also adopted in a previous work on GLM (Li, Lu, and Zhou 2017), which ensure that the change of $P(X = 1 | \mathbf{pa}(X))$ is not abrupt. It is worth noting that Assumption 2 only needs the lower bound of the first derivative in the neighborhood of θ_X^* , which is weaker than Assumption 1 in (Filippi et al. 2010). Finally, Assumption 3 makes sure that each parent node of X still has a constant probability of taking either 0 or 1 even when all other parents of X already fix their values. This means that each parent has some independence and is not fully determined by other parents. In appendix we give some further justification of this assumption.

We first introduce some notations. Let n, m be the number of nodes and edges in G respectively. Let $D = \max_{X \in \mathbf{X} \cup \{Y\}} |Pa(X)|$, $L_{\max}^{(1)} = \max_{X \in \mathbf{X} \cup \{Y\}} L_{f_X}^{(1)}$, and c be the constant in Lecué and Mendelson's inequality (Nie 2022) (see Lemma 11 in appendix). Let (\mathbf{X}_t, Y_t) be the propagation result in the t^{th} round of intervention. It contains $\mathbf{V}_{t,X}$ and X^t for each node $X \in \mathbf{X} \cup \{Y\}$ where X^t is the propagating result of X and $\mathbf{V}_{t,X}$ is the propagating result of parents of X . Additionally, our estimation of the weight vectors is denoted by $\hat{\theta}$.

We now propose the algorithm BGLM-OFU in Algorithm 1, where OFU stands for optimism in the face of uncertainty. The algorithm contains two phases. The first phase is the initialization phase with only pure observations without doing any intervention to ensure that our maximum likelihood estimation of θ^* is accurate enough. Based on Lecué and Mendelson's inequality (Nie 2022), it is designed to ensure Eq. (5) in Lemma 1 holds. In practice, one alternative implementation of the initialization phase is doing no intervention

Algorithm 1: BGLM-OFU for BGLM CCB Problem

- 1: **Input:** Graph $G = (\mathbf{X} \cup \{Y\}, E)$, intervention budget $K \in \mathbb{N}$, parameter $L_{f_X}^{(1)}, L_{f_X}^{(2)}, \kappa, \zeta$ in Assumption 1, 2 and 3.
 - 2: Initialize $M_{0,X} \leftarrow \mathbf{0} \in \mathbb{R}^{|Pa(X)| \times |Pa(X)|}$ for all $X \in \mathbf{X} \cup \{Y\}$, $\delta \leftarrow \frac{1}{3n\sqrt{T}}$, $R \leftarrow \lceil \frac{512D(L_{f_X}^{(2)})^2}{\kappa^4} (D^2 + \ln \frac{1}{\delta}) \rceil$, $T_0 \leftarrow \max \left\{ \frac{c}{\zeta^2} \ln \frac{1}{\delta}, \frac{(8n^2 - 16n + 2)R}{\zeta} \right\}$ and $\rho \leftarrow \frac{3}{\kappa} \sqrt{\log(1/\delta)}$.
 - 3: /* Initialization Phase: */
 - 4: Do no intervention on BGLM G for T_0 rounds and observe feedback $(\mathbf{X}_t, Y_t), 1 \leq t \leq T_0$.
 - 5: /* Iterative Phase: */
 - 6: **for** $t = T_0 + 1, T_0 + 2, \dots, T$ **do**
 - 7: $\{\hat{\theta}_{t-1,X}, M_{t-1,X}\}_{X \in \mathbf{X} \cup \{Y\}} =$
BGLM-Estimate($(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_{t-1}, Y_{t-1})$)
(see Algorithm 2).
 - 8: Compute the confidence ellipsoid $\mathcal{C}_{t,X} = \{\theta'_X \in [0, 1]^{|Pa(X)|} : \|\theta'_X - \hat{\theta}_{t-1,X}\|_{M_{t-1,X}} \leq \rho\}$ for any node $X \in \mathbf{X} \cup \{Y\}$.
 - 9: $(\mathbf{S}_t, \tilde{\theta}_t) = \operatorname{argmax}_{\mathbf{S} \subseteq \mathbf{X}, |\mathbf{S}| \leq K, \theta'_{i,X} \in \mathcal{C}_{t,X}} \mathbb{E}[Y | do(\mathbf{S})]$.
 - 10: Intervene all the nodes in \mathbf{S}_t to 1 and observe the feedback (\mathbf{X}_t, Y_t) .
 - 11: **end for**
-

until Eq. (5) holds for every $X \in \mathbf{X} \cup \{Y\}$. The required number of rounds is usually much less than T_0 . Then in the second iterative phase, we use maximum likelihood estimator (MLE) method to estimate θ^* and can therefore create a confidence region that contains the real parameters θ^* with high probability around it to balance the exploration and exploitation. More specifically, in each iteration of intervention selections, we find an optimal intervention set together with a set of parameters $\tilde{\theta}$ from the region around our unbiased estimation $\hat{\theta}$ and take the found intervention set in this iteration. Intuitively, this method to select intervention sets follows the OFU spirit: the argmax operator in line 9 of Algorithm 1 selects the best (optimistic) solution in a confidence region (to address uncertainty). The empirical mean $\hat{\theta}$ calculated corresponds to exploration while the confidence region surrounding it is for exploration.

The regret analysis of Algorithm 1 requires two technical components to support the main analysis. The first component indicates that when the observations are sufficient, we can get a good estimation $\hat{\theta}$ for θ^* , while the second component shows that a small change in the parameters should not lead to a big change in the reward $\mathbb{E}[Y | do(\mathbf{S})]$.

The first component is based on the result of maximum-likelihood estimation. In the studies of (Filippi et al. 2010) and (Li, Lu, and Zhou 2017), a standard log-likelihood function used in the updating process should be $L_{i,X}^{std}(\theta_X) = \sum_{i=1}^t [X^i \ln f_X(\mathbf{V}_{i,X}^T \theta_X) + (1 - X^i) \ln(1 - f_X(\mathbf{V}_{i,X}^T \theta_X))]$. However, the analysis in their work needs

Algorithm 2: BGLM-Estimate

- 1: **Input:** All observations $((\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_t, Y_t))$ until round t .
 - 2: **Output:** $\{\hat{\theta}_{t,X}, M_{t,X}\}_{X \in \mathbf{X} \cup \{Y\}}$
 - 3: For each $X \in \mathbf{X} \cup \{Y\}$, $i \in [t]$, construct data pair $(\mathbf{V}_{i,X}, X^i)$ with $\mathbf{V}_{i,X}$ the parent value vector of X in round i , and X^i the value of X in round i if $X \notin S_i$.
 - 4: **for** $X \in \mathbf{X} \cup \{Y\}$ **do**
 - 5: Calculate the maximum-likelihood estimator $\hat{\theta}_{t,X}$ by solving the equation $\sum_{i=1}^t (X^i - f_X(\mathbf{V}_{i,X}^\top \theta_X)) \mathbf{V}_{i,X} = 0$.
 - 6: $M_{t,X} = \sum_{i=1}^t \mathbf{V}_{i,X} \mathbf{V}_{i,X}^\top$.
 - 7: **end for**
-

the gradient of the log-likelihood function to have the form $\sum_{i=1}^t [X^i - f_X(\mathbf{V}_{i,X}^\top \theta_X)] \mathbf{V}_{i,X}$, which is not true here. Therefore, using the same idea in (Zhang et al. 2022), we use the pseudo log-likelihood function $L_{t,X}(\theta_X)$ instead, which is constructed by integrating the gradient of it defined by $\nabla_{\theta_X} L_{t,X}(\theta_X) = \sum_{i=1}^t [X^i - f_X(\mathbf{V}_{i,X}^\top \theta_X)] \mathbf{V}_{i,X}$. Actually, this pseudo log-likelihood function is used in line 5 of Algorithm 2. The following lemma presents the result for the learning problem as the first technical component of the regret analysis. Let $M_{t,X}$ and $\hat{\theta}_{t,X}$ be as defined in Algorithm 2, and also note that the definition of $\hat{\theta}_{t,X}$ is equivalent to $\hat{\theta}_{t,X} = \operatorname{argmax}_{\theta_X} L_{t,X}(\theta_X)$. Let $\lambda_{\min}(M)$ denote the minimum eigenvalue of matrix M .

Lemma 1 (Learning Problem for BGLM). *Suppose that Assumptions 1 and 2 hold. Moreover, given $\delta \in (0, 1)$, assume that*

$$\lambda_{\min}(M_{t,X}) \geq \frac{512 |\mathbf{Pa}(X)| \left(L_{f_X}^{(2)}\right)^2}{\kappa^4} \left(|\mathbf{Pa}(X)|^2 + \ln \frac{1}{\delta} \right). \quad (5)$$

Then with probability at least $1 - 3\delta$, the maximum-likelihood estimator satisfies, for any $\mathbf{v} \in \mathbb{R}^{|\mathbf{Pa}(X)|}$,

$$\left| \mathbf{v}^\top (\hat{\theta}_{t,X} - \theta_X^*) \right| \leq \frac{3}{\kappa} \sqrt{\log(1/\delta)} \|\mathbf{v}\|_{M_{t,X}^{-1}},$$

where the probability is taken from the randomness of all data collected from round 1 to round t .

The proof of the above lemma is adapted from (Li, Lu, and Zhou 2017), and is included in appendix. Note that the initialization phase of the algorithm together with Assumption 3 and the Lecu e and Mendelson’s inequality would show the condition on $\lambda_{\min}(M_{t,X})$ in Eq.(5), and the design of the initialization phase, the summarization of Assumption 3 and the analysis to show Eq.(5) together form one of our key technical contributions in this section.

For the second component showing that a small change in parameters leads to a small change in the reward, we adapt the group observation modulated (GOM) bounded smoothness property for the LT model (Li et al. 2020) to show a GOM bounded smoothness property for BGLM. To do

so, we define an equivalent form of BGLM as a threshold model as follows. For each node X , we randomly sample a threshold γ_X uniformly from $[0, 1]$, i.e. $\gamma_X \sim \mathcal{U}[0, 1]$, and if $f_X(\mathbf{Pa}(X) \cdot \theta_X^*) + \varepsilon_X \geq \gamma_X$, X is activated (i.e. set to 1); if not, X is not activated (i.e. set to 0). Suppose $X_1, X_2, \dots, X_{n-1}, Y$ is a topological order of nodes in $\mathbf{X} \cup \{Y\}$, then at time step 1, only X_1 is tried to be activated; at time step 2, only X_2 is tried to be activated by X_1 ; \dots ; at time step n , only Y is tried to be activated by activated nodes in $\mathbf{Pa}(Y)$. The above view of the propagating process is equivalent to BGLM, but it shows that BGLM is a general form of the LT model (Kempe, Kleinberg, and Tardos 2003) on DAG. Thus we can show a result below similar to Theorem 1 in (Li et al. 2020) for the LT model. For completeness, we include the proof in appendix. Henceforth, we use $\sigma(\mathbf{S}, \theta)$ to represent the reward function $\mathbb{E}[Y | do(\mathbf{S})]$ under parameter θ , to make the parameters of the reward explicit. We use γ to represent the vector $(\gamma_X)_{X \in \mathbf{X} \cup Y}$.

Lemma 2 (GOM Bounded Smoothness of BGLM). *For any two weight vectors $\theta^1, \theta^2 \in \Theta$ for a BGLM G , the difference of their expected reward for any intervened set \mathbf{S} can be bounded as*

$$|\sigma(\mathbf{S}, \theta^1) - \sigma(\mathbf{S}, \theta^2)| \leq \mathbb{E}_{\varepsilon, \gamma} \left[\sum_{X \in \mathbf{X}_{\mathbf{S}, Y}} |\mathbf{V}_X^\top (\theta_X^1 - \theta_X^2)| L_{f_X}^{(1)} \right], \quad (6)$$

where $\mathbf{X}_{\mathbf{S}, Y}$ is the set of nodes in paths from \mathbf{S} to Y excluding \mathbf{S} , and \mathbf{V}_X is the propagation result of the parents of X under parameter θ^2 . The expectation is taken over the randomness of the thresholds γ and the noises ε .

We can now prove the regret bound of Algorithm 1. In the proof of Theorem 1, we use Lecu e and Mendelson’s inequality (Nie 2022) to prove that our initialization step has a very high probability to meet the needs of Lemma 1. Then we use Lemma 2 to transform the regret to the sum of $\|\mathbf{V}_{t,X}\|_{M_{t-1,X}^{-1}}$, which can be bounded using a similar lemma of Lemma 2 in (Li, Lu, and Zhou 2017). The details of the proof is put in appendix for completeness.

Theorem 1 (Regret Bound of BGLM-OFU). *Under Assumptions 1, 2 and 3, the regret of BGLM-OFU (Algorithms 1 and 2) is bounded as*

$$R(T) = O\left(\frac{1}{\kappa} n L_{\max}^{(1)} \sqrt{DT} \log T\right), \quad (7)$$

where the terms of $o(\sqrt{T})$ are omitted.

Remarks. The leading term of the regret in terms of T is in the order of $O(\sqrt{T} \log T)$, which is commonly seen in confidence ellipsoid based bandit or combinatorial bandit algorithms (e.g. (Abbasi-Yadkori, P al, and Szepesv ari 2011; Li et al. 2020; Zhang et al. 2022)). Also, it matches the regret of previous causal bandits algorithm, C-UCB in (Lu et al. 2020), which works on the atomic setting. The term $L_{\max}^{(1)}$ reflects the rate of changes in f_X ’s, and intuitively, the higher rate of changes in f_X ’s, the larger the regret since the online learning algorithm inevitably leads to some error in the parameter estimation, which will be amplified by the

rate of changes in f_X 's. Technically, $L_{\max}^{(1)}$ comes from the $L_{f_X}^{(1)}$ term in the GOM condition (Eq.(6)). Term n is to relax the sum over $\mathbf{X}_{S,Y}$ in Eq.(6), and it could be made tighter in causal graphs where $|\mathbf{X}_{S,Y}|$ is significantly smaller than n , and intuitively it means that all nodes on the path from S to Y would contribute to the regret. Term \sqrt{D} implies that the regret depends on the number of parents of nodes, and technically it is because the scale of $\|\mathbf{V}_{t,X}\|_{M_{t-1,X}^{-1}}$ is approximately $\sqrt{|\mathbf{Pa}(X)|} \leq \sqrt{D}$, and we bound the regret as the sum of $\|\mathbf{V}_{t,X}\|_{M_{t-1,X}^{-1}}$'s as we explained earlier. Term $\frac{1}{\kappa}$ comes from the learning problem (Lemma 1), which is also adopted in the regret bound of UCB-GLM of (Li, Lu, and Zhou 2017) using a similar learning problem. For the budget K , it does not appear in our regret because it is not directly related to the number of parameters we are estimating.

While our algorithm and analysis are based on several past studies (Li, Lu, and Zhou 2017; Zhang et al. 2022; Li et al. 2020), our innovation includes (a) the initialization phase and its corresponding Assumption 3 and its more involved analysis, because in our model we do not have direct observations of one-step immediate causal effect; and (b) the integration of the techniques from these separate studies, such as the maximum likelihood based analysis of (Li, Lu, and Zhou 2017), the pseudo log-likelihood function of (Zhang et al. 2022), and the GOM condition analysis of (Li et al. 2020), whereas each of these studies alone is not enough to achieve our result.

In line 9 of Algorithm 1, the argmax operator needs a simultaneous optimization over both the intervention set \mathbf{S} and parameters θ' . This could be a computationally hard problem for large-scale problems, but since we focus on the online learning aspect of the problem, we leave the computationally-efficient solution for the full problem as future work, and such treatment is often seen in other combinatorial online learning problems (e.g. (Combes et al. 2015; Li et al. 2020)).

5 Algorithms for BLM with Hidden Variables

Previous sections consider only Markovian models without hidden variables. In many causal models, hidden variables exist to model the latent confounding factors. In this section, we present results on CCB under the linear model BLM with hidden variables.

5.1 Transforming the Model with Hidden Variables to the one without Hidden Variables

To address hidden variables, we first show how to reduce the BLM with hidden variables to a corresponding one without the hidden variables. Suppose the hidden variables in $G = (\mathbf{U} \cup \mathbf{X} \cup \{Y\}, E)$ are $\mathbf{U} = \{U_0, U_1, U_2, \dots\}$, and we use X_i, X_j 's to represent observed variables. Without loss of generality, we let U_0 always be 1 and it only has out-edges pointing to other observed or unobserved nodes, to model the self-activations of nodes. We allow various types of connections involving the hidden nodes, including edges from

observed nodes to hidden nodes and edges among hidden nodes, but we disallow the situation where a hidden node U_s with $s > 0$ has two paths to X_i and X_i 's descendant X_j and the paths contain only hidden nodes except the end points X_i and X_j . Figure 1 is an example of a causal graph allowed for this section.

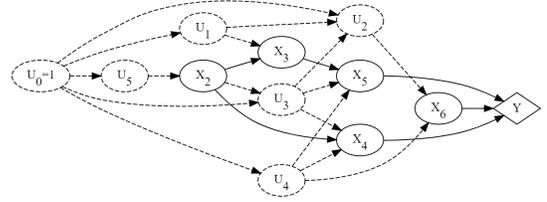


Figure 1: An Example of BLM with Hidden Variables

Our idea is to transform such a general causal model into an equivalent Markovian model $G' = (\{X_1\} \cup \mathbf{X} \cup \{Y\}, E')$. For convenience, we assume X_1 is not in the original observed variable set $\mathbf{X} \cup \{Y\}$. Henceforth, all notations with ' \prime ', such as Pr' , \mathbb{E}' , θ^* , $\mathbf{Pa}'(X)$ correspond to the new Markovian model G' , and the notations Pr , \mathbb{E} without ' \prime ' refer to the original model G . For any two observed nodes X_i, X_j , a hidden path P from X_i to X_j is a directed path from X_i to X_j where all intermediate nodes are hidden or there are no intermediate nodes. We define θ_P^* to be the multiplication of weights on path P . Let \mathcal{P}_{X_i, X_j}^u be the set of hidden paths from X_i to X_j . If $\mathcal{P}_{X_i, X_j}^u \neq \emptyset$, then we add edge (X_i, X_j) into E' , and its weight $\theta_{X_i, X_j}^* = \sum_{P \in \mathcal{P}_{X_i, X_j}^u} \theta_P^*$. As in the Markovian model, X_1 is always set to 1, and for each $X_i \in \mathbf{X} \cup \{Y\}$, we add edge (X_1, X_i) into E' , with weight $\theta_{X_1, X_i}^* = \text{Pr}\{X_i = 1 | \text{do}(\mathbf{X} \cup \{Y\} \setminus \{X_i\}) = \mathbf{0}\}$. The noise variables ε are carried over to G' without change.

The following lemma shows that the new model G' has the same parent-child conditional probability as the original model G . The proof of this lemma utilizes the do-calculus rules for causal inference (Pearl 2009).

Lemma 3. *For any $X \in \mathbf{X} \cup \{Y\}$, any $\mathbf{S} \subseteq \mathbf{X}$, any value $\mathbf{pa}'(X) \in \{0, 1\}^{|\mathbf{Pa}'(X)|}$, any value $\mathbf{s} \in \{0, 1\}^{|\mathbf{S}|}$ (\mathbf{s} is consistent with $\mathbf{pa}'(X)$ on values in $\mathbf{S} \cap \mathbf{Pa}'(X)$), we have*

$$\begin{aligned} \text{Pr}\{X = 1 | \mathbf{Pa}'(X) \setminus \{X_1\} = \mathbf{pa}'(X) \setminus \{x_1\}, \text{do}(\mathbf{S} = \mathbf{s})\} \\ = \text{Pr}'\{X = 1 | \mathbf{Pa}'(X) = \mathbf{pa}'(X), \text{do}(\mathbf{S} = \mathbf{s})\}. \end{aligned}$$

The next lemma shows that the objective function is also the same for G' and G .

Lemma 4. *For any $\mathbf{S} \subseteq \mathbf{X}$, any value $\mathbf{s} \in \{0, 1\}^{|\mathbf{S}|}$, we have $\mathbb{E}[Y | \text{do}(\mathbf{S} = \mathbf{s})] = \mathbb{E}'[Y | \text{do}(\mathbf{S} = \mathbf{s})]$.*

The above two lemmas show that from G to G' the parent-child conditional probability and the reward function are all remain unchanged.

5.2 Applying BGLM-OFU on BLM

With the transformation described in Section 5.1 and its properties summarized in Lemmas 3 and 4, we can apply

Algorithm 1 on G' to achieve the learning result for G . More precisely, we can use the observed data of X and $\mathbf{Pa}'(X)$ to estimate parameters θ_X^* and minimize the regret on the reward $\mathbb{E}[Y|do(\mathbf{S})]$, which is the same as $\mathbb{E}[Y|do(\mathbf{S})]$. Furthermore, under the linear model BLM, Assumptions 1 and 2 hold with $L_{f_X}^{(1)} = \kappa = 1$ and $L_{f_X}^{(2)}$ could be any constant greater than 0. We still need Assumption 3, but we change the $\mathbf{Pa}(X)$'s in the assumption to $\mathbf{Pa}'(X)$'s. Then we can have the following regret bound.

Theorem 2 (Regret Bound of Algorithm 1 for BLM). *Under Assumption 3, Algorithm 1 has the following regret bound when running on BLM with hidden variables:*

$$R(T) = O\left(n\sqrt{DT} \log T\right), \quad (8)$$

where n is the number of nodes in G' , and $D = \max_{X \in \mathbf{X} \cup \{Y\}} |\mathbf{Pa}'(X)|$ is the maximum in-degree in G' .

Remarks. We compare our regret bound to the one in (Li et al. 2020) for the online influence maximization problem under the LT model, which is a special case of our BLM model with $\varepsilon_X = 0$ for all X 's. Our regret bound is $O(n^{\frac{3}{2}}\sqrt{T} \ln T)$ while theirs is $O(n^{\frac{3}{2}}\sqrt{T} \ln T)$. The n^3 factor saving comes from three sources: (a) our reward is of scale $[0, 1]$ while theirs is $[0, n]$, and this saves one n factor; (b) our BLM is a DAG, and thus we can equivalently fix the activation steps as described before Lemma 2, causing our GOM condition (Lemma 2) to save another factor of n ; and (c) we use MLE, which requires an initialization phase and Assumption 3 to give an accurate estimate of θ^* , while their algorithm uses linear regression without an initialization phase, which means we tradeoff a factor of n with an additional Assumption 3.

5.3 Algorithm for BLM based on Linear Regression

Now we have already introduced how to use Algorithm 1 and 2 to solve the online BLM CCB problem with hidden variables. However, in order to meet the needs of Lemma 1, we have to process an initialization phase. Assumption 3 needs to hold for the Markovian model G' after the transformation. We can remove the initialization phase and the dependency on Assumption 3, by noticing that our MLE based on pseudo log-likelihood maximization is equivalent to linear regression when adopted on BLMs. In particular, we use Lemma 1 in (Li et al. 2020) to replace Lemma 1 for MLE. We rewrite it as Lemma 11 in appendix.

Based on the above result, we introduce Algorithm 3 using the linear regression. Algorithm 3 is designed for Markovian BLMs. For a BLM G with more general hidden variables, we can apply the transformation described in Section 5.1 to transform the model into a Markovian model G' first. The following theorem shows the regret bound of Algorithm 3.

Theorem 3 (Regret Bound of Algorithm 3). *The regret of BLM-LR (Algorithm 3) running on BLM with hidden variables is bounded as*

$$R(T) = O\left(n^2\sqrt{DT} \log T\right).$$

Algorithm 3: BLM-LR for BLM CCB Problem

- 1: **Input:** Graph $G = (\mathbf{X} \cup \{Y\}, E)$, intervention budget $K \in \mathbb{N}$.
 - 2: Initialize $M_{0,X} \leftarrow \mathbf{I} \in \mathbb{R}^{|\mathbf{Pa}(X)| \times |\mathbf{Pa}(X)|}$, $\mathbf{b}_{0,X} \leftarrow \mathbf{0}^{|\mathbf{Pa}(X)|}$ for all $X \in \mathbf{X} \cup \{Y\}$, $\hat{\theta}_{0,X} \leftarrow 0 \in \mathbb{R}^{|\mathbf{Pa}(X)|}$ for all $X \in \mathbf{X} \cup \{Y\}$, $\delta \leftarrow \frac{1}{n\sqrt{T}}$ and $\rho_t \leftarrow \sqrt{n \log(1 + tn) + 2 \log \frac{1}{\delta} + \sqrt{n}}$ for $t = 0, 1, 2, \dots, T$.
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: Compute the confidence ellipsoid $\mathcal{C}_{t,X} = \{\theta'_X \in [0, 1]^{|\mathbf{Pa}(X)|} : \|\theta'_X - \hat{\theta}_{t-1,X}\|_{M_{t-1,X}} \leq \rho_{t-1}\}$ for any node $X \in \mathbf{X} \cup \{Y\}$.
 - 5: $(\mathbf{S}_t, \tilde{\theta}_t) = \operatorname{argmax}_{\mathbf{S} \subseteq \mathbf{X}, |\mathbf{S}| \leq K, \theta'_{t,X} \in \mathcal{C}_{t,X}} \mathbb{E}[Y|do(\mathbf{S})]$.
 - 6: Intervene all the nodes in \mathbf{S}_t to 1 and observe the feedback (\mathbf{X}_t, Y_t) .
 - 7: **for** $X \in \mathbf{X} \cup \{Y\}$ **do**
 - 8: Construct data pair $(\mathbf{V}_{t,X}, X^t)$ with $\mathbf{V}_{t,X}$ the parent value vector of X in round t , and X^t the value of X in round t if $X \notin \mathbf{S}_t$.
 - 9: $M_{t,X} = M_{t-1,X} + \mathbf{V}_{t,X} \mathbf{V}_{t,X}^\top$, $\mathbf{b}_{t,X} = \mathbf{b}_{t-1,X} + X^t \mathbf{V}_{t,X}$, $\hat{\theta}_{t,X} = M_{t,X}^{-1} \mathbf{b}_{t,X}$.
 - 10: **end for**
 - 11: **end for**
-

The proof of this theorem is adapted from the proof of Theorem 2 in (Li et al. 2020). For completeness, the proof is put in appendix. Comparing the algorithm and the regret bound of BLM-LR with those of BGLM-OFU, we can see a tradeoff between using the MLE method and the linear regression method: When we use the linear regression method with the BLM-LR algorithm, we do not need to have an initialization phase so Assumption 3 is not required anymore. However, the regret bound of BLM-LR (Theorem 3) has an extra factor of n in regret bound, comparing to the regret bound of BGLM-OFU on BLM (Theorem 2).

6 Conclusion and Future Work

In this paper, we propose the combinatorial causal bandit (CCB) framework, and provide a solution for CCB under the BGLM. We further study a special model, the linear model. We show that our algorithm would work for models with many types of hidden variables. We further provide an algorithm for linear model not relying on Assumption 3 based on the linear regression.

There are many open problems and future directions to extend this work. For the BGLM, one could study how to remove some assumptions (e.g. Assumption 3), how to include hidden variables, or how to make the computation more efficient. For the linear model, one could consider how to remove the constraint on the hidden variable structure that does not allow a hidden variable to connect to an observed variable and its observed descendant via hidden variables. More generally, one could consider classes of causal models other than the BGLM.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments and constructive feedback.

References

- Abbasi-Yadkori, Y.; Pál, D.; and Szepesvári, C. 2011. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24.
- Aitkin, M.; Francis, B.; Hinde, J.; and Darnell, R. 2009. *Statistical Modelling in R*. Oxford University Press Oxford.
- Arnold, K. F.; Davies, V.; de Kamps, M.; Tennant, P. W.; Mbotwa, J.; and Gilthorpe, M. S. 2020. Reflection on modern methods: generalized linear models for prognosis and intervention—theory, practice and implications for machine learning. *International Journal of Epidemiology*, 49(6): 2074–2082.
- Chen, W.; Wang, Y.; and Yuan, Y. 2013. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, 151–159. PMLR.
- Chen, W.; Wang, Y.; Yuan, Y.; and Wang, Q. 2016. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *The Journal of Machine Learning Research*, 17(1): 1746–1778.
- Combes, R.; Shahi, M. S. T. M.; Proutiere, A.; et al. 2015. Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, 2116–2124.
- Feng, S.; and Chen, W. 2022. Combinatorial Causal Bandits. *arXiv preprint arXiv:2206.01995*.
- Filippi, S.; Cappe, O.; Garivier, A.; and Szepesvári, C. 2010. Parametric bandits: The generalized linear case. *Advances in Neural Information Processing Systems*, 23.
- Garcia-Huidobro, D.; and Michael Oakes, J. 2017. Squeezing observational data for better causal inference: Methods and examples for prevention research. *International Journal of Psychology*, 52(2): 96–105.
- Han, B.; Yu, H.; and Friedberg, M. W. 2017. Evaluating the impact of parent-reported medical home status on children’s health care utilization, expenditures, and quality: a difference-in-differences analysis with causal inference methods. *Health Services Research*, 52(2): 786–806.
- Hastie, T. J.; and Pregibon, D. 2017. Generalized linear models. In *Statistical Models in S*, 195–247. Routledge.
- Hernán, M. A.; and Robins, J. M. 2010. Causal inference.
- Hilbe, J. M. 2011. Logistic regression. *International Encyclopedia of Statistical Science*, 1: 15–32.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146.
- Lattimore, F.; Lattimore, T.; and Reid, M. D. 2016. Causal bandits: learning good interventions via causal inference. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 1189–1197.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit Algorithms*. Cambridge University Press.
- Lee, S.; and Bareinboim, E. 2018. Structural causal bandits: where to intervene? *Advances in Neural Information Processing Systems*, 31.
- Lee, S.; and Bareinboim, E. 2019. Structural causal bandits with non-manipulable variables. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, 4146–4172.
- Lee, S.; and Bareinboim, E. 2020. Characterizing optimal mixed policies: Where to intervene and what to observe. *Advances in Neural Information Processing Systems*, 33: 8565–8576.
- Li, L.; Lu, Y.; and Zhou, D. 2017. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, 2071–2080. PMLR.
- Li, S.; Kong, F.; Tang, K.; Li, Q.; and Chen, W. 2020. On-line influence maximization under linear threshold model. *Advances in Neural Information Processing Systems*, 33: 1192–1204.
- Lu, Y.; Meisami, A.; and Tewari, A. 2021. Causal bandits with unknown graph structure. *Advances in Neural Information Processing Systems*, 34.
- Lu, Y.; Meisami, A.; Tewari, A.; and Yan, W. 2020. Regret analysis of bandit problems with causal background knowledge. In *Conference on Uncertainty in Artificial Intelligence*, 141–150. PMLR.
- Maiti, A.; Nair, V.; and Sinha, G. 2021. Causal Bandits on General Graphs. *arXiv preprint arXiv:2107.02772*.
- Nair, V.; Patil, V.; and Sinha, G. 2021. Budgeted and non-budgeted causal bandits. In *International Conference on Artificial Intelligence and Statistics*, 2017–2025. PMLR.
- Nie, Z. 2022. Matrix anti-concentration inequalities with applications. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, 568–581.
- Pearl, J. 2009. *Causality*. Cambridge University Press. 2nd Edition.
- Sakate, D.; and Kashid, D. 2014. Comparison of Estimators in GLM with Binary Data. *Journal of Modern Applied Statistical Methods*, 13(2): 10.
- Sen, R.; Shanmugam, K.; Dimakis, A. G.; and Shakkottai, S. 2017. Identifying best interventions through online importance sampling. In *International Conference on Machine Learning*, 3057–3066. PMLR.
- Vansteelandt, S.; and Dukes, O. 2020. Assumption-lean inference for generalised linear model parameters. *arXiv preprint arXiv:2006.08402*.
- Yabe, A.; Hatano, D.; Sumita, H.; Ito, S.; Kakimura, N.; Fukunaga, T.; and Kawarabayashi, K.-i. 2018. Causal bandits with propagating inference. In *International Conference on Machine Learning*, 5512–5520. PMLR.

Zhang, Z.; Chen, W.; Sun, X.; and Zhang, J. 2022. Online influence maximization with node-level feedback using standard offline oracles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 9153–9161.