

# Asymmetric Graph Error Control with Low Complexity in Causal Bandits

Chen Peng, Di Zhang and Urbashi Mitra

**Abstract**—In this paper, the causal bandit problem is investigated, with the objective of maximizing the long-term reward by selecting an optimal sequence of interventions on nodes in an unknown causal graph. It is assumed that both the causal topology and the distribution of interventions are unknown. First, based on the difference between the two types of graph identification errors (false positives and negatives), a causal graph learning method is proposed. Numerical results suggest that this method has a much lower sample complexity relative to the prior art by learning *sub-graphs*. However, we note that a sample complexity analysis for the new algorithm has not been undertaken, as of yet. Under the assumption of minimum-mean squared error weight estimation, a new uncertainty bound tailored to the causal bandit problem is derived. This uncertainty bound drives an upper confidence bound-based intervention selection to optimize the reward. Further, we consider a particular instance of non-stationary bandits wherein both the causal topology and interventional distributions can change. Our solution is the design of a sub-graph change detection mechanism that requires a modest number of samples. Numerical results compare the new methodology to existing schemes and show a substantial performance improvement in stationary and non-stationary settings. Averaged over 100 randomly generated causal bandits, the proposed scheme takes significantly fewer samples to learn the causal structure and achieves a reward gain of 85% compared to existing approaches.

**Index Terms**—Causal bandit, linear structural equation model, graph identification, upper confidence bound.

## I. INTRODUCTION

Investigation of cause-effect relationships is central to answering key questions in medicine, epidemiology, economics as well as the social and behavioral sciences (see *e.g.* [2]–[4]). Cause-effect relationships can be represented by Bayesian networks, in the form of directed acyclic graphs (DAGs). Identifying the causal structural model is essential to decision making and answering counterfactual questions. Herein, we will study methods for determining the underlying causal structural model in the context of causal bandits and thus will also optimize intervention design. We first discuss causal graph identification and then discuss causal bandits.

### A. Related Work

Current causal graph identification methods mainly fall into two categories: constraint-based and score-based. Constraint-

This work is funded in part by one or more of the following grants: ARO W911NF1910269, ARO W911NF2410094, DOE DE-SC0021417, Swedish Research Council 2018-04359, NSF A22-2666-S003, NSF A25-0139-S009, NSF CCF-2008927, NSF CCF-2200221, NSF CCF-2311653, NSF KR-705319, NSF RINGS-2148313, NSF PIPP phase II: Center for Pandemic Insights (CPI) 2412522, ONR 503400-78050, ONR N00014-22-1-2363.

Parts of this paper were presented in part at the 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) [1].

based approaches, such as [2], [5], search for conditional independence over all possible graphs and thus computation complexity scales exponentially. Classical score-based methods evaluate different graph structures in their ability to fit the data (see *e.g.* [6]–[8]). To lower the complexity, recent score-based methods (see *e.g.* [9]–[12]) formulate the structure learning problem as a continuous optimization task. In [10], an algorithm called GOLEM is proposed, where the likelihood score is combined with a soft DAG constraint to evaluate possible causal structures. The DAGMA scheme proposed in [11] characterizes acyclicity by a log-determinant function, with better-behaved gradients.

Decision-making or action-taking is natural within the context of causal inference. Despite this natural relationship, incorporation of causal methods within the context of multi-armed bandits (MAB) [13], [14], has only occurred recently. MABs provide a useful model for sequential decision-making problems, and as such have been employed in the design of clinical trials [15], recommendation systems [16], financial portfolio design [17], *etc.* In a MAB, an agent selects an action (arm) in each round and observes corresponding outcomes – the goal is to maximize the long-term cumulative reward. In the standard setting, the stochastic rewards generated by different arms are assumed to be statistically independent.

To model realistic scenarios with dependence, structured bandits have been considered. Within this arena, we focus on causal bandits [14], where the causal structure is exploited to improve decision-making. DAGs can encode the causal relationship among factors that contribute to the reward [14]. Herein, we interpret the arms as different interventions on the nodes of a DAG and the reward as the stochastic outcome of a certain node. In the current work, we will provide new methods for both causal discovery (graph identification) as well as the design of optimal interventions for causal MABs.

Existing literature on causal bandits can be categorized based on their assumptions about the topology of the underlying DAG and the distribution of the interventions. While many works assume prior topology or DAG knowledge [14], [18]–[22], this is often not known in practice. In this work, we assume that the underlying causal structure is unknown.

The setting without the knowledge of both topology and interventional distribution has been investigated recently [23]–[25], [27]. An algorithm based on causal tree recovery is proposed in [24], with regret scaling logarithmically with the number of nodes. In [23], an auxiliary separating set algorithm is utilized to learn the causal structure and provide improved regret over non-causal algorithms. The existing works are based on the hard (perfect) intervention model, where the

TABLE I  
KEY ASSUMPTIONS: RELATED WORK

Reference	Topology Knowledge	Intervention	Structural Model	Non-stationary Structure	Exogenous/Noise
[14], [18], [19]	✓	Hard	Discrete	✗	Bernoulli
[20]	✓	Soft	Linear	✗	Sub-Gaussian
[21]	✓	Soft	General	✗	Known distribution
[22]	✓	Soft	General	✗	Sub-Gaussian
[23]	Known Separating Sets	Hard	Discrete & Linear	✗	Bernoulli & Gaussian
[24]	✗	Hard	Discrete	✗	Sub-Gaussian
[25]	✗	Hard	Linear	✗	sub-Gaussian
[26]	✓	Soft	Linear	✓	sub-Gaussian
This work	✗	Soft	Linear	✓	Gaussian

causal relations between a node and its parents are completely cut off because a specific value is assigned to the node. Herein, we consider soft interventions, such that the node under intervention can still be causally related to other nodes.

A challenge in MAB problems is the balance between exploration and exploitation. A common approach is via derivations of upper confidence bounds (UCB) [28]. The classic UCB scheme uses the number of visits as a general measure of uncertainty [29], while in causal bandits, uncertainty can be better quantified by bounding the variance of problem-specific estimators [20], [23], [24]. We shall adapt the UCB approach to our context, exploiting properties of our graph identification scheme. Another standard assumption in causal bandits is that of a static causal model. Both [26] and [30] consider the non-stationary case, but assume prior knowledge of the initial causal graph. We shall provide a method by which to keep the causal structure up-to-date without prior graph knowledge. A comparison between prior works and our approach is provided in Table I.

### B. Approach and Contributions

We note that separating graph identification from optimized interventions for MABs is inefficient as these two goals are coupled in causal MABs. Another feature not typically considered in the causal graph identification or causal MAB literature is the distinguishing between false positive and false negative errors in graph identification – we shall actively consider this feature. A false positive means that the graph identification method outputs an edge that is not present in the true graph, while a false negative means that the graph identification method outputs no edge, where in fact, the edge is present in the true DAG. In this paper, we propose the *Causal Sub-graph Learning with Upper Confidence Bound (CSL-UCB)* scheme, without assumed knowledge of either the causal graph topology or the interventional distributions. Our algorithm seeks to minimize false negative errors in graph identification as they have a larger impact on optimal intervention design. We quantify the uncertainty in our graph identification coupled with that stemming from decision making. The learning of sub-graphs (versus the entire graph) strongly reduces complexity and provides a natural strategy for change detection in non-stationary environments. The main contributions of this paper are:

- 1) We propose a sub-graph learning approach to learn the causal structure, which has strongly improved computational complexity. Numerical results suggest that it has a

much lower sample complexity as well. A novel mutual information regularization is proposed which reduces false negative errors. It is shown that false negatives have a more critical impact on overall reward optimization in causal bandit problems.

- 2) An uncertainty bound tailored to the causal bandit framework and our graph identification method is derived and used to drive an upper confidence bound strategy for arm selection.
- 3) A sub-graph change detection mechanism with local updates is proposed enabling the consideration of non-stationary causal structure.
- 4) Numerical comparisons show that the proposed scheme identifies the optimal intervention much faster than standard MAB schemes by exploiting causal structure. Moreover, compared to strategies that only focus on causal graph identification, the proposed scheme needs much fewer samples than other methods and learns the causal structure with a limited loss with respect to the long-term reward.

The rest of the paper is organized as follows. We introduce the system model in Section II. The sub-graph learning scheme is proposed in Section III, with an analysis of the two types of error. Uncertainty quantification is provided in Section IV, which enables the balance between exploration and exploitation. The sub-graph change detection mechanism is proposed in Section V. Finally, the paper is concluded in Section VI and all the proofs are provided in the Appendix.

## II. SYSTEM MODEL

### A. The Causal Graphical Model with Soft Intervention

We assume the causal effects can be captured by a DAG with structure  $(\mathcal{V}, \mathcal{B})$ , where  $\mathcal{V} = \{1, \dots, N\}$  is the set of  $N$  nodes and  $\mathcal{B}$  is the set of directed edges. Moreover, the edge-weight matrix  $\mathbf{B} \in \mathbb{R}^{N \times N}$  captures the strength of causal effects, where the  $(i, j)$ -th entry represents the weight of the edge  $i \rightarrow j$ . To describe observations under intervention, consider node-wise intervention, defined as

$$\mathbf{a} = (a_1, \dots, a_N)^\top \in \{0, 1\}^N, \quad (1)$$

where  $a_i$  represents whether node  $i$  is intervened ( $a_i = 1$ ) or not ( $a_i = 0$ ). Specifically, instead of hard interventions, we consider soft interventions, which do not necessarily cut off causal relationships between the intervened node and its parents, but change the in-coming edges to the node. To

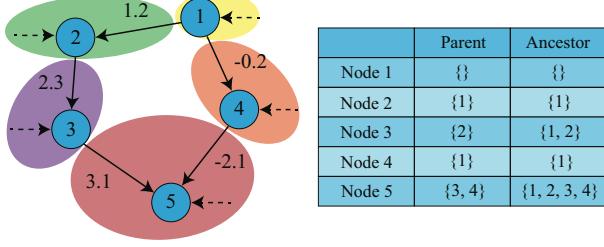


Fig. 1. A causal graph with  $N = 5$  and node 5 as the reward node. Each sub-graph is marked by a colored region. The sets of parents and ancestors are provided in the inline table.

formalize the model for soft intervention, we denote the interventional edge-weight matrix by  $\mathbf{B}' \in \mathbb{R}^{N \times N}$ , such that the post-intervention weight matrix  $\mathbf{B}_a$  can be constructed as

$$[\mathbf{B}_a]_i = \mathbb{I}(a_i = 1)\mathbf{B}'_i + \mathbb{I}(a_i = 0)\mathbf{B}_i, \quad (2)$$

where  $\mathbb{I}(\cdot)$  is the indicator function and  $[\cdot]_i$  represents the  $i$ -th column of a matrix such that  $\mathbf{B}_i, \mathbf{B}'_i \in \mathbb{R}^{N \times 1}$ . The  $i$ -th column of the post-interventional weight matrix determines the set of parents of node  $i$  and how these parents causally influence node  $i$ .

As a result of intervention, the vector of stochastic values associated with the nodes is represented by  $\mathbf{x} \in \mathbb{R}^N$ . The causal relationship among nodes is described by a linear structural equation model (LinSEM),

$$\mathbf{x} = (\mathbf{B}_a)^\top \mathbf{x} + \boldsymbol{\epsilon}, \quad (3)$$

where  $\boldsymbol{\epsilon}$  is a vector of Gaussian exogenous/noise variables. We assume  $\boldsymbol{\epsilon}$  contains independent elements with known means  $\boldsymbol{\nu}$  and unknown variances  $\boldsymbol{\sigma}^2$ . Further, we denote  $\tilde{\boldsymbol{\epsilon}} \doteq \boldsymbol{\epsilon} - \boldsymbol{\nu}$  such that  $\tilde{\boldsymbol{\epsilon}} \sim \mathcal{N}(\mathbf{0}, \text{diag}(\boldsymbol{\sigma}^2))$ .

We denote the sets of parents and ancestors of node  $i$  by  $\mathcal{P}_i(a_i)$  and  $\mathcal{A}_i(a_i)$ , the estimated set of parents by  $\widehat{\mathcal{P}}_i(a_i)$ . The set difference of the estimated and true parent sets is denoted by  $\widehat{\mathcal{P}}_i \setminus \mathcal{P}_i(a_i)$ . Taking node 3 in Fig. 1 as an example, the parent and ancestor sets are  $\mathcal{P}_3(0) = \{2\}$ ,  $\mathcal{A}_3(0) = \{1, 2\}$ . If the estimated parent set is  $\widehat{\mathcal{P}}_3(0) = \{2, 4\}$ , then we have  $\widehat{\mathcal{P}}_3 \setminus \mathcal{P}_3(0) = \{4\}$ .

### B. The Causal Bandit Model

In the MAB framework, an agent performs a sequence of actions in order to maximize cumulative reward over a finite horizon  $T$ . With the causal graphical model, we consider node  $N$  as the reward node, which generates a stochastic reward in each time step. An example of such a causal graph is given in Fig. 1, where the value of node 5 is considered the reward signal and the effect of exogenous variables are represented by dashed arrows.

To compute the expected reward under intervention  $\mathbf{a}$ , we recognize that in LinSEMs, there exists a causal flow between every ancestor-descendant pair. Thus each variable  $x_i$  can be decomposed as a linear combination of exogenous variables in  $\boldsymbol{\epsilon}$ , weighted by the strength of the causal flow. Define the post-intervention flow-weight matrix as

$$\mathbf{C}_a \doteq (\mathbf{I} - \mathbf{B}_a)^{-1}, \quad (4)$$

where the  $(i, j)$ -th entry represents the net flow weight from node  $i$  to  $j$ . In this way, we can rewrite (3) as

$$\mathbf{x} = (\mathbf{I} - \mathbf{B}_a)^{-1} \boldsymbol{\epsilon} = (\mathbf{C}_a)^\top \boldsymbol{\epsilon}, \quad (5)$$

where  $\mathbf{I}$  denotes the identity matrix. The expectation of  $\mathbf{x}$  under intervention  $\mathbf{a}$  is formulated as

$$\boldsymbol{\mu}_a \doteq \mathbb{E} \left[ (\mathbf{I} - \mathbf{B}_a)^{-1} \boldsymbol{\epsilon} \right] = (\mathbf{I} - \mathbf{B}_a)^{-1} \boldsymbol{\nu}. \quad (6)$$

Thus, given the post-intervention weight matrices, we are able to find the intervention that maximizes the expected reward, which is defined as the optimal intervention,

$$\mathbf{a}^* \doteq \arg \max_{\mathbf{a}} [\boldsymbol{\mu}_a]_N. \quad (7)$$

In each step, the decision-making agent selects an intervention  $\mathbf{a}^t$ , observes  $\mathbf{x}^t$  and collects a reward  $x_N^t$ . The randomness of the observations comes from the exogenous variables, which are independent of the intervention. The objective is to maximize the expected cumulative reward, or equivalently, minimize the expected cumulative regret, defined as

$$R_T \doteq \mathbb{E} \left[ \sum_{t=1}^T \left( [\boldsymbol{\mu}_{\mathbf{a}^*}]_N - [\boldsymbol{\mu}_{\mathbf{a}^t}]_N \right) \right]. \quad (8)$$

We underscore that all key derivations and theoretical results in the sequel assume the signal model in this section.

## III. CAUSAL SUB-GRAPH LEARNING

Causal bandit problems are different from ordinary MAB problems, due to the existence of an underlying causal structure, which is assumed to be unknown in our setting. Since the expected reward is a function of the post-intervention weight matrix  $\mathbf{B}_a$ , understanding the causal structure enables the agent to better select the intervention afterwards. Over a finite horizon, the agent has to balance between gaining causal knowledge (exploration) and maximizing immediate reward (exploitation), which is a common dilemma for sequential decision-making problems. In this section, we focus on identifying the causal structure, in a sub-graph learning manner. Based on that, the problem of balancing exploration and exploitation is investigated in the following section.

### A. Sub-graph Decomposition

Causal graph identification is equivalent to figuring out causal relationships between every pair of nodes in the graph, represented by directed edges. Given observation up to step  $t$ ,  $\mathbf{X}^{1:t} \doteq (\mathbf{x}^1, \dots, \mathbf{x}^t)$  of dimension  $N \times t$ , we can evaluate every possible graph structure by its ability to fit the data. However, finding the best graph structure may not be feasible in practice because the number of possible DAGs grows super-exponentially in the number of nodes. For example, the number of DAGs for 2, 4, 10 nodes are 3, 543 and  $4.2 \times 10^{18}$ , respectively [31]. The growing search space issue becomes even more serious for causal bandits. Since the effect of an intervention is characterized by the post-interventional graph  $\mathbf{B}_a$  and there are  $2^N$  possible interventions, the agent has to identify  $2^N$  different graphs with only  $t$  observations.

To improve sample efficiency, one feature of the causal bandit model is noteworthy. Even though there are  $2^N$  post-intervention distributions characterized by  $\mathbf{B}_a$ , those post-intervention weight matrices are composed of columns of  $\mathbf{B}$  and  $\mathbf{B}'$  [20]. Therefore, instead of identifying  $2^N$  causal graphs induced by  $\mathbf{B}_a$ , we can identify causal relationships induced by the columns of  $\mathbf{B}$  and  $\mathbf{B}'$ . Specifically,  $\mathbf{B}_i$  and  $\mathbf{B}'_i$  encode causal weights between node  $i$  and its parents. The following proposition ensures that causal graphs can always be uniquely decomposed into sub-graphs of our interest.<sup>1</sup>

**Proposition 1.** *Any DAG can always be uniquely decomposed into sub-graphs such that each sub-graph consists of a node and all in-coming edges to the node.*

*Proof:* First, we prove the existence of such decomposition by construction. For any complete DAG with a set of nodes  $\mathcal{V} = \{1, \dots, N\}$  and a set of edges  $\mathcal{B}$ , we can always decompose it into a group of sub-graphs  $\mathcal{G}_i, i \in \{1, \dots, N\}$ , such that

$$\mathcal{V}_i = \{i\}, \quad \mathcal{B}_i = \{(k, i) : \forall k \in \mathcal{P}_i\}. \quad (9)$$

Moreover, notice that the decomposition is unique: sub-graph  $\mathcal{G}_i$  contains and only contains node  $i$  and all of its in-coming edges, which completes the proof.  $\square$

As illustrated in Fig. 1, each sub-graph resides in one of the colored regions. For example, the sub-graph in the red region consists of node 5, directed edges (3, 5) and (4, 5). Learning sub-graphs has exponentially improved data efficiency. Consider an example where interventions are selected in turns up to step  $t$ , such that each complete graph is associated with  $t/2^N$  data samples, as there are  $2^N$  complete graphs. Each sub-graph, instead, is associated with  $t/2$  data samples, as each node can only be in one of two possible modes: under intervention, or not.

For learning causal sub-graphs, we employ the *principle of independent mechanisms*, which states that the causes and the mechanism producing the effect are independent [31]. Based on that, testing the independence of residuals has been investigated for causal discovery, in the context of additive noise models [32]–[34]. Herein, we consider minimum mean-square error (MMSE) estimation, where the estimator and residual are defined as

$$\hat{x}_i^t(a_i) \doteq \mathbb{E} \left[ x_i \middle| \mathbf{X}_{\hat{\mathcal{P}}_i(a_i)}^{1:t} \right], \quad r_i^t(a_i) \doteq \hat{x}_i^t(a_i) - x_i. \quad (10)$$

With the estimated mechanism producing  $x_i$ , the principle of independent mechanisms implies that the residual should be independent from values of the parent nodes,

$$r_i^t(a_i) \perp x_j, \quad \forall j \in \mathcal{P}_i(a_i). \quad (11)$$

Thus, the dependence between  $r_i^t$  and  $x_j$  supports the hypothesis that node  $j$  is not a parent of node  $i$  and thus the edge  $j \rightarrow i$  does not exist. By evaluating all potential in-coming edges to a node, the corresponding sub-graph can be estimated.

<sup>1</sup>The sub-graphs we consider here are not vertex-induced sub-graphs.

## B. Edge-weighted Mutual Information

As a general measure of dependence, mutual information (MI), denoted by  $I(\cdot)$ , is widely utilized for causal discovery (see e.g. [35], [36]). However, since the sub-graphs are treated separately, pure MI-based sub-graph learning with limited data suffers *significantly from estimation error*. In this subsection, we explain this issue in detail and propose the *edge-weighted mutual information* measure as a solution.

We take a closer look at the mutual information of interest. In the following derivation, we focus on node  $i$  under intervention  $a$  and a potential parent node  $j$ , with observation up to step  $t$ . The true causal relationship is represented as

$$x_i = \sum_{k \in \mathcal{P}_i(a_i)} [\mathbf{B}_a]_{ki} x_k + \epsilon_i = \sum_{k \in \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{ki} \epsilon_k + \epsilon_i. \quad (12)$$

For example, we have  $x_3 = 2.76\epsilon_1 + 2.3\epsilon_2 + \epsilon_3$  for node 3 in Fig. 1. With the estimated edge-weight matrix denoted by  $\hat{\mathbf{B}}_a^t$ , the estimated flow-weight matrix is defined as

$$\hat{\mathbf{C}}_a^t \doteq (\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-1}. \quad (13)$$

Then the estimated node value can be decomposed using flow weights as

$$\hat{x}_i^t = \sum_{k \in \mathcal{A}_i(a_i)} [\hat{\mathbf{C}}_a^t]_{ki} \epsilon_k + \sum_{l \notin \mathcal{A}_i(a_i)} [\hat{\mathbf{C}}_a^t]_{li} \epsilon_l + \nu_i. \quad (14)$$

For example, the estimate of  $x_3$  can be decomposed according to the ancestor set of node 3, as presented in Fig. 1. Since we do not assume any knowledge of the causal structure, all other nodes are considered as potential parents of node  $i$ . Based on (12) and (14), the residual can be expressed as

$$r_i^t(a_i) = -\tilde{\epsilon}_i + \underbrace{\sum_{k \in \mathcal{A}_i(a_i)} ([\hat{\mathbf{C}}_a^t]_{ki} - [\mathbf{C}_a]_{ki}) \epsilon_k}_{\text{intrinsic error}} + \underbrace{\sum_{l \notin \mathcal{A}_i(a_i)} [\hat{\mathbf{C}}_a^t]_{li} \epsilon_l}_{\text{causal error}}. \quad (15)$$

The second term in (15) is always non-zero in the finite sample regime, thus we define this error as the *intrinsic error*. The last term is defined as the *causal error* because it is non-zero only when some non-ancestor nodes are considered as direct causes (otherwise the residual would not contain  $\epsilon_l$ ). For a specific non-ancestor node  $j$ , which is a source of the causal error, we decompose its value as

$$x_j = \sum_{k \in \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{kj} \epsilon_k + \sum_{l \notin \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{lj} \epsilon_l. \quad (16)$$

Now that we have both  $r_i^t(a_i)$  and  $x_j$  for mutual information evaluation, the following proposition provides an upper bound where the impact of intrinsic and causal errors are separated.

**Proposition 2.** *With independent exogenous variables, the mutual information between the residual and a declared parent can be upper bounded as*

$$I(r_i^t(a_i); x_j) \leq I\left(\underbrace{\sum_{l \notin \mathcal{A}_i(a_i)} [\widehat{\mathbf{C}}_a^t]_{li} \epsilon_l - \tilde{\epsilon}_i; \sum_{l \notin \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{lj} \epsilon_l}_{\text{causal error induced MI}}\right) \\ + I\left(\underbrace{\sum_{k \in \mathcal{A}_i(a_i)} ([\widehat{\mathbf{C}}_a^t]_{ki} - [\mathbf{C}_a]_{ki}) \epsilon_k; \sum_{k \in \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{kj} \epsilon_k}_{\text{intrinsic error induced MI}}\right). \quad (17)$$

*Proof:* See Appendix A  $\square$

The proof relies on the data processing inequality [37] and the mutual independence of the exogenous variables. We define the first and second terms on the right-hand side of (17) as the mutual information induced by the causal and intrinsic errors, respectively. Due to the existence of the intrinsic error, we may not be able to detect the causal error and thus misjudge causal directions. To illustrate this point, consider the *causal error induced mutual information*. Denote the linear combination of exogenous variables from non-ancestors by

$$\theta(x_j, i) \doteq \sum_{l \notin \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{lj} \epsilon_l, \quad (18)$$

which is Gaussian, with mean  $\mu_{\theta x}$  and variance  $\sigma_{\theta x}^2$ . For example, since the ancestor set of node 3 in Fig. I is  $\{1, 2\}$ ,  $\theta(x_5, 3)$  is a linear combination of  $\epsilon_3, \epsilon_4$  and  $\epsilon_5$ . As we have a causal error, there is an assumed causal edge that does not exist in the true graph. This erroneous causal edge,  $j \rightarrow i$  (from non-ancestors) is used in the estimation, and as a result, induces a causal error in the residual, weighted by the estimated parameter. To focus on the effect caused by the erroneously labeled causal edge  $j \rightarrow i$ , we denote the causal error in the residual induced by other nodes by

$$\theta(r_i^t(a_i), \setminus j) \doteq \sum_{l \notin \mathcal{A}_i(a_i)} [\widehat{\mathbf{C}}_a^t]_{li} \epsilon_l - \tilde{\epsilon}_i - [\widehat{\mathbf{B}}_a^t]_{ji} \theta(x_j, i), \quad (19)$$

which is Gaussian as well, with mean  $\mu_{\theta r}$  and variance  $\sigma_{\theta r}^2$ , while  $\setminus j$  denotes the set of other non-ancestors,  $\setminus j \doteq \{l \in \{1, \dots, N\} \mid l \notin \mathcal{A}_i(a_i), l \neq j\}$ . In this way, we can rewrite and bound the causal error induced mutual information as

$$I\left(\sum_{l \notin \mathcal{A}_i(a_i)} [\widehat{\mathbf{C}}_a^t]_{li} \epsilon_l - \tilde{\epsilon}_i; \sum_{l \notin \mathcal{A}_i(a_i)} [\mathbf{C}_a]_{lj} \epsilon_l\right) \\ = I\left([\widehat{\mathbf{B}}_a^t]_{ji} \theta(x_j, i) + \theta(r_i^t(a_i), \setminus j); \theta(x_j, i)\right) \quad (20)$$

$$= h\left([\widehat{\mathbf{B}}_a^t]_{ji} \theta(x_j, i) + \theta(r_i^t(a_i), \setminus j)\right) + h\left(\theta(x_j, i)\right) \\ - h\left(\theta(r_i^t(a_i), \setminus j), \theta(x_j, i)\right) \quad (21)$$

$$\leq \log\left(\left(|[\widehat{\mathbf{B}}_a^t]_{ji}| \sigma_{\theta x} + \sigma_{\theta r}\right) \sqrt{2\pi e}\right) + \log\left(\sigma_{\theta x} \sqrt{2\pi e}\right) \\ - \left[\log(2\pi e) + \log\left(\sigma_{\theta x} \sigma_{\theta r} \sqrt{1 - \rho_\theta^2}\right)\right] \quad (22)$$

$$= \log\left(1 + \frac{\sigma_{\theta x}}{\sigma_{\theta r}} |[\widehat{\mathbf{B}}_a^t]_{ji}|\right) - \frac{1}{2} \log(1 - \rho_\theta^2), \quad (23)$$

where  $\rho_\theta$  stands for the correlation

$$\rho_\theta = \frac{\mathbb{E}[(\theta(x_j, i) - \mu_{\theta x})(\theta(r_i^t(a_i), \setminus j) - \mu_{\theta r})]}{\sigma_{\theta x} \sigma_{\theta r}}, \quad (24)$$

and  $h(\cdot)$  denotes the differential entropy. Note that the inequality (22) comes from the Gaussianity of key random variables and the fact that the variance is maximized when the variables are perfectly correlated ( $\rho_\theta = 1$ ).

We observe that the causal error induced mutual information becomes small when the estimated edge weight is small. To make matters worse, the intrinsic error induced mutual information is normally not negligible in the limited data regime. Thus, it is challenging to detect the causal error and reject a non-causal edge when the estimated weight is small. To alleviate this issue, we propose an *edge-weighted mutual information* measure, defined as

$$I_w(r_i^t(a_i); x_j) \doteq I(r_i^t(a_i); x_j) - \log |[\widehat{\mathbf{B}}_a^t]_{ji}|. \quad (25)$$

A large  $I_w$  indicates that the corresponding edge is non-causal and thus should be rejected. The remaining problem is to determine the threshold for rejecting edges.

In general, any residual based test for causal identification needs to balance between false positive and false negative errors [31], by selecting an appropriate testing threshold. However, in terms of maximizing cumulative reward, the two types of error contribute differently to the ultimate error of reward estimation. To be more specific, we argue that rejecting an actual edge (FN error) is much worse than accepting a nonexistent edge (FP error). Formally, FN errors exist in the estimated causal graph (observational or interventional) if

$$\exists (i, j) : (i, j) \in \mathcal{B}, \quad (i, j) \notin \widehat{\mathcal{B}}, \quad (26)$$

where  $\widehat{\mathcal{B}}$  denotes the estimated edge set. In the following subsections, we analyze the impact of FN and FP errors and propose an adaptive edge-rejecting strategy to minimize the FN error rate.

### C. Analysis of Graph Errors

As discussed in Section II-B, estimated rewards under different interventions are computed based on the estimated weight matrices. Thus, we can evaluate the impact of the two types of error by their effects on weight matrix estimation. Focusing on the sub-graph of node  $i$  under intervention  $a_i$ , the following proposition gives the estimated edge weights in expectation, under the two types of errors.

**Proposition 3.** *We assume MMSE estimation and determine the effect of **only** false positive or **only** false negative errors. For the presence of **only** false positive errors, the estimated edge weights satisfy*

$$\mathbb{E}\left([\widehat{\mathbf{B}}_a^t]_{i, \widehat{\mathcal{P}}_i(a_i)} \mid \mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i)\right) = [\mathbf{B}_a]_{i, \widehat{\mathcal{P}}_i(a_i)}. \quad (27)$$

*With **only** false negative errors, the estimated edge weights satisfy*

$$\mathbb{E}\left([\widehat{\mathbf{B}}_a^t]_{i, \widehat{\mathcal{P}}_i(a_i)} \mid \mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i)\right) = [\mathbf{B}_a]_{i, \widehat{\mathcal{P}}_i(a_i)} + \\ \left[\mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i) \mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i)^\top\right]^{-1} \mathbb{E}\left[\mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i) \cdot \mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i)^\top \mid \mathbf{X}_{\widehat{\mathcal{P}}_i(a_i)}^{1:t}(a_i)\right] [\mathbf{B}_a]_{i, \mathcal{P}_i \setminus \widehat{\mathcal{P}}_i(a_i)}. \quad (28)$$

*Proof:* See Appendix B  $\square$

The proof is based on the properties of MMSE estimation. Note that  $\mathbf{X}^{1:t}$  and  $\epsilon^{1:t}$  have  $a_i$  as an argument, because only the values associated with the intervention  $a_i$  are considered for learning the causal structure under intervention  $a_i$ .

Observe that with only FP errors, the estimated edge weights are asymptotically unbiased. Although the estimated weights for any non-existent parents in  $\widehat{\mathcal{P}}_i(a_i)$  is non-zero for a finite number of samples, the bias converges to zero as more samples are collected. In contrast, with the existence of FN errors, the estimated weights for both correctly identified and erroneously rejected parents are biased, resulting in inaccurate weights even with an infinite amount of data. Note that the bias is a function of the correlation between identified parents and rejected parents.

#### D. Adaptive Edge Rejection

Since estimated rewards are inaccurate when FN errors occur, we should avoid rejecting an actual edge and ensure the complete graph is a DAG. Thus, the threshold for rejecting edges should be set large enough to minimize the probability of rejecting an actual edge. On the other hand, the threshold should not be too large because redundant edges can result in cycles in the complete causal graph.

To achieve this goal, we propose to adjust the threshold adaptively such that edges are rejected only when necessary. To be more specific, the edge-weighted mutual information is calculated for each potential edge and the edge with the largest  $I_w$  is rejected at a time, until the complete graph becomes a DAG. Once the set of declared edges is determined, we construct the weight matrices by MMSE estimation. Notice that in this process, although sub-graphs are learned separately with different sets of samples, whether an edge should be rejected or not is also influenced by edges in other sub-graphs.

The overall approach for causal graph identification is denoted as the *Causal Sub-graph Learning (CSL)* scheme, and the pseudo-code is provided in Algorithm 1. Note that although the observational weight matrix  $\mathbf{B}$  is used as the target in the pseudo-code, the same procedure will be used for learning the interventional weight matrix  $\mathbf{B}'$  as well.

Lastly, we provide the time and space complexity of the proposed algorithm. Starting with a set of all possible edges of size  $N^2$ , the algorithm rejects an edge in each step until the graph becomes a DAG, and thus the number of steps is less than  $N^2$ . Each step selects an edge with the largest edge-weighted MI among the remaining edges, with time complexity  $\mathcal{O}(N^2)$ . After the rejection of an edge, estimation of the weights and residuals takes  $\mathcal{O}(N^2(N+t))$  time while estimation of the MI takes  $\mathcal{O}(Nt \log t)$  time [38]. Besides, checking whether the graph is a DAG by topological sorting takes  $\mathcal{O}(N^2)$  time [39]. Finally, if  $N \leq t$ , the time complexity of the proposed scheme is  $\mathcal{O}(N^4t + N^3t \log t)$ . The space complexity is  $\mathcal{O}(Nt)$  for storing the collected samples and a matrix of edge-weighted MI.

## IV. INTERVENTION SELECTION UNDER UNCERTAINTY

Given the ability of learning causal graphs with observed data, we can estimate the reward under each possible interven-

---

#### Algorithm 1 The Causal Sub-graph Learning Scheme

---

**Require:** The set of nodes  $\mathcal{V}$  and node values  $\mathbf{X}^{1:t}$ .

- 1: Initialize the edge set to include all possible directed edges:  $\widehat{\mathcal{B}} = \{(i, j) : \forall i, j \in \mathcal{V}\}$ .
- 2: **while**  $(\mathcal{V}, \widehat{\mathcal{B}})$  is not a DAG **do**
- 3:   Compute edge weights  $\widehat{\mathbf{B}}_{ij}$  and residuals,  $\forall (i, j) \in \widehat{\mathcal{B}}$ , by MMSE estimation, with observed values  $\mathbf{X}^{1:t}$ .
- 4:   Compute edge-weighted MI  $I_w(r_i^t(0); x_j)$ ,  $\forall (i, j) \in \widehat{\mathcal{B}}$ , with  $\mathbf{X}^{1:t}$  and empirical MI.
- 5:   Find  $(i, j) = \arg \max_{i,j} I_w(r_i^t(0); x_j)$ , remove the edge  $(i, j)$  from  $\widehat{\mathcal{B}}$ .
- 6: **end while**
- 7: Compute  $\widehat{\mathbf{B}}_{ij}$ ,  $\forall (i, j) \in \widehat{\mathcal{B}}$ , by MMSE estimation, with  $\mathbf{X}^{1:t}$ . For  $(i, j) \notin \widehat{\mathcal{B}}$ , set  $\widehat{\mathbf{B}}_{ij} = 0$ .
- 8: **return** Estimated weight matrix  $\widehat{\mathbf{B}}$ .

---

tion. Although selecting the intervention with largest estimated reward is optimal for the current step, exploring other interventions is necessary for achieving higher cumulative reward in the long run. It is because the intervention with largest estimated reward may not be the optimal intervention, due to the uncertainty of reward estimation. Thus, intervention selection should consider both how close the estimated rewards are to being maximal and the uncertainties in those estimates. In this section, we derive an uncertainty bound on the estimation error of the expected reward, so that the potential of an intervention can be taken into account.

For any specific intervention  $\mathbf{a}$ , define the matrix of edge-weight errors and the vector of expected node-value errors as

$$\Delta \mathbf{B}_\mathbf{a}^t \doteq \widehat{\mathbf{B}}_\mathbf{a}^t - \mathbf{B}_\mathbf{a}, \quad \Delta \boldsymbol{\mu}_\mathbf{a}^t \doteq \widehat{\boldsymbol{\mu}}_\mathbf{a}^t - \boldsymbol{\mu}_\mathbf{a}, \quad (29)$$

where  $\widehat{\boldsymbol{\mu}}_\mathbf{a}^t$  represents the estimated mean of  $\mathbf{x}$ . As analyzed in Section III-C, the weight errors could have non-zero means, due to missing edges (false negatives). With the centered error matrix defined as

$$\Delta \tilde{\mathbf{B}}_\mathbf{a}^t \doteq \Delta \mathbf{B}_\mathbf{a}^t - \mathbb{E}[\Delta \mathbf{B}_\mathbf{a}^t | \mathbf{X}^{1:t}(\mathbf{a})], \quad (30)$$

we bound the error in the estimated reward (see (29)) as

$$|[\Delta \boldsymbol{\mu}_\mathbf{a}^t]_N| = \left| [(\mathbf{I} - \widehat{\mathbf{B}}_\mathbf{a}^t)^{-1}]_N^T (\Delta \mathbf{B}_\mathbf{a}^t)^T \boldsymbol{\mu}_\mathbf{a} \right| \quad (31)$$

$$\leq \left| [(\mathbf{I} - \widehat{\mathbf{B}}_\mathbf{a}^t)^{-1}]_N^T \mathbb{E}[\Delta \mathbf{B}_\mathbf{a}^t | \mathbf{X}^{1:t}(\mathbf{a})]^T \boldsymbol{\mu}_\mathbf{a} \right| + \left| [(\mathbf{I} - \widehat{\mathbf{B}}_\mathbf{a}^t)^{-1}]_N^T (\Delta \tilde{\mathbf{B}}_\mathbf{a}^t)^T \boldsymbol{\mu}_\mathbf{a} \right|. \quad (32)$$

Recall that we provide an expression for the non-zero elements of  $\mathbb{E}[\Delta \mathbf{B}_\mathbf{a}^t | \mathbf{X}^{1:t}(\mathbf{a})]$  in (28), which depends on the correlation between correctly declared parents and mistakenly rejected parents. However, since we are not able to know whether and where causal edges are rejected by mistake, the uncertainty bound cannot be computed at run-time.

Nevertheless, when no FN error exists, we are able to bound the uncertainty of the estimated reward. With the covariance of the  $i$ -th weight error vector denoted by  $\Phi_\mathbf{a}^t(i)$ , we derive a concentration inequality for the estimation error, under the assumption of no FN errors.

**Lemma 1.** *With MMSE estimation and the assumption of no false negative errors, the norm of the weight error vectors have upper-bounded moments for  $m \geq 2$ :*

$$\mathbb{E} \left[ \left\| [\Delta \mathbf{B}_a^t]_i \right\|_2^m \right] \leq m! \sqrt{\frac{4N}{N+2}} \left[ \lambda_{\max}(\Phi_a^t(i)) \frac{N+2}{4} \right]^{m/2}, \quad (33)$$

where  $\lambda_{\max}(\cdot)$  stands for the maximum eigenvalue of a matrix.

*Proof:* See Appendix C  $\square$

The proof is based on the eigenvalue decomposition of the weight error matrix and Jensen's inequality. Lemma 1 enables us to derive a concentration inequality for the maximum singular-value of the weight error matrix.

**Lemma 2.** *With MMSE estimation and the assumption of no false negative error, the following inequality holds,*

$$\mathbb{P} \left\{ \sigma_{\max}(\Delta \mathbf{B}_a^t) \geq 2(N^2 + 2N)^{1/4} \times \sqrt{\ln \left( \frac{2N}{\delta} \right) \sum_{i=1}^N \lambda_{\max}(\Phi_a^t(i))} \right\} \leq \delta, \quad (34)$$

where  $\sigma_{\max}(\cdot)$  represents the maximum singular-value.

*Proof:* See Appendix D  $\square$

The proof relies on a key dilation transformation, Lemma 1 and the matrix Bernstein inequality [40]. Based on Lemma 2, the following theorem provides a concentration inequality for the error of the estimated reward.

**Theorem 1.** *Under the assumption MMSE estimation and no false negative errors, the following inequality holds:*

$$\mathbb{P} \left\{ \left| [\Delta \mu_a^t]_N \right| \geq U(\mathbf{X}^{1:t}, \mathbf{a}, \delta) \right\} \leq \delta, \quad (35)$$

where  $U(\mathbf{X}^{1:t}, \mathbf{a}, \delta)$  represents the error upper-bound at confidence level  $1 - \delta$ ,

$$U(\mathbf{X}^{1:t}, \mathbf{a}, \delta) = 2(N^2 + 2N)^{1/4} \left\| [(\mathbf{I} - \widehat{\mathbf{B}}_a^t)^{-1}]_N \right\|_2 \times \|\mu_a\|_2 \sqrt{\ln \left( \frac{2N}{\delta} \right) \sum_{i=1}^N \lambda_{\max}(\Phi_a^t(i))}. \quad (36)$$

*Proof:* See Appendix E  $\square$

The proof is based on the Woodbury matrix identity [41] and Lemma 2. Since the uncertainty of reward estimation cannot be measured in the presence of FN errors, we employ (35) as an approximation for the general case. The bound on uncertainty enables us to balance between exploration and exploitation, so that the cumulative regret is minimized in the long run. In the presence of FN errors, although the uncertainty bound may not hold exactly, the observed low cumulative regret in Section VI-B suggests, at least numerically, that the uncertainty bound has utility in our proposed algorithm.

Note that the uncertainty bounds are computed based on the estimation of edge-weight matrices. Initially, to ensure there exist samples for learning every sub-graph, we select interventions randomly. This phase is thus our *exploration* start. Afterwards, we select the intervention that achieves the

---

## Algorithm 2 The CSL-UCB Scheme

**Require:** The set of nodes  $\mathcal{V}$ , length of exploring start  $T_{\text{es}}$ , uncertainty level  $\delta$  and exploration parameter  $\alpha$ .

```

1: EXPLORING START PHASE
2:   for  $t = 1 : T_{\text{es}}$  do
3:     Randomly select an intervention  $\mathbf{a}^t$ .
4:     Observe node values  $\mathbf{x}^t$  and collect a reward of  $x_N^t$ .
5:   end for
6: BALANCED EXPLOITATION & EXPLORATION PHASE
7:   for  $t = T_{\text{es}} + 1 : T$  do
8:     Estimate the observational and interventional causal graphs using Algorithm 1
9:     For each intervention, construct the post-intervention weight matrix by 2 and evaluate the uncertainty bound according to (36).
10:    Select an intervention  $\mathbf{a}^t$  according to (37).
11:    Observe node values  $\mathbf{x}^t$  and collect a reward of  $x_N^t$ .
12:   end for

```

---

best balance between exploitation (large estimated reward) and exploration (high uncertainty) in each time step via:

$$\mathbf{a}^{t+1} = \arg \max_{\mathbf{a}} \left\{ [(\mathbf{I} - \widehat{\mathbf{B}}_a^t)^{-1}]_N^\top \boldsymbol{\nu} + \alpha U(\mathbf{X}^{1:t}, \mathbf{a}, \delta) \right\}, \quad (37)$$

where  $\alpha$  is a design parameter that controls the exploration level. Combining CSL with uncertainty evaluation, the overall approach is denoted as the *Causal Sub-graph Learning with Upper Confidence Bound (CSL-UCB)* scheme, whose pseudo-code is provided in Algorithm 2.

## V. CAUSAL SUB-GRAPH CHANGE DETECTION

So far, we have investigated the causal bandit problem, assuming that the underlying causal structure is invariant. However, it could be non-stationary in real-world scenarios, such that the post-intervention weight matrix  $\mathbf{B}_a$  varies over time. Although the proposed CSL-UCB algorithm keeps updating the causal structure, it should only utilize data generated by the current causal model. In this section, we develop a causal sub-graph *change detector* for both weight and topology changes, which enables the proposed scheme to maintain a high cumulative reward even when the causal structure is non-stationary.

Consider a piece-wise stationary model where the causal structure remains the same between two consecutive change points. Without loss of generality, we focus on detecting the first change-point in time, denoted by  $t_c$ . To determine whether a change occurs at step  $s$ , a straightforward approach compares the distribution of observed values during  $[1, s]$  with the distribution during  $(s, t]$ . If the distributions differ significantly,  $s$  is claimed to be a change point. However, the distribution shift of  $\mathbf{x}$  only indicates the existence of a change but not the source of the change. To understand this, recall the data generating process,

$$x_i = \sum_{k \in \mathcal{P}_i(a_i)} [\mathbf{B}_a]_{ki} x_k + \epsilon_i, \quad (38)$$

where the distribution shift of  $x_i$  could be due to the change of causal mechanism  $[\mathbf{B}_a]_i$  or/and the distribution shifts of its parents  $x_k$ . As the source of change is not identifiable, the whole graph has to be re-learned.

Since causal mechanisms are modular and independent [31], the sparse mechanism shift hypothesis [42] suggests that changes in the data distribution generally are caused by changes in only a subset of causal mechanisms. In this case, it is sufficient to only re-learn the changed part. We rely on the whitening transformation of  $\mathbf{x}$  to identify the set of changed mechanisms,

$$\hat{\mathbf{y}} \doteq (\mathbf{I} - \hat{\mathbf{B}}_a^t)^\top \mathbf{x} = (\mathbf{I} - \mathbf{B}_a)^\top \mathbf{x} - (\Delta \mathbf{B}_a^t)^\top \mathbf{x} \quad (39)$$

$$= \epsilon - (\Delta \mathbf{B}_a^t)^\top \mathbf{x}. \quad (40)$$

According to (40),  $y_i$  and  $\epsilon_i$  differ in distribution if and only if  $[\Delta \mathbf{B}_a^t]_i \neq 0$ , in which case the estimated sub-graph model fails to explain the observed data well. There are two main possible reasons for a large discrepancy between distributions: causal structure change and FN errors in estimation. Thus, besides detecting causal structure changes, checking the consistency between  $\mathbf{y}$  and  $\epsilon$  also enables us to discover FN errors.

When the sub-graph associated with node  $i$  changes, the whitened variable  $y_i^\tau$  would differ in distribution from  $\epsilon_i$ . Thus, to determine whether there is a change point, we formulate a hypothesis testing problem,

$$\mathcal{H}_0 : \forall \tau \in [1, t], y_i^\tau \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\nu_i, \sigma_i^2), \quad (41)$$

$$\mathcal{H}_1 : \exists t_c \in (1, t), y_i^\tau \stackrel{\text{i.i.d.}}{\sim} \begin{cases} \mathcal{N}(\nu_i, \sigma_i^2), & \tau \in [1, t_c] \\ \mathcal{N}(\nu'_i, \sigma'_i), & \tau \in (t_c, t] \end{cases}, \quad (42)$$

where  $\nu_i, \sigma_i^2$  are the mean and variance of  $\epsilon_i$ , while  $\nu'_i, \sigma'_i$  are the mean and variance of  $y_i^\tau, \tau \in (t_c, t]$ .

For testing the null hypothesis of no change-point versus the alternative hypothesis of a single change-point, we employ the Generalized Likelihood Ratio (GLR) statistic [43],

$$\begin{aligned} \Psi_i(s) \doteq & \sum_{\tau=1}^s \log \mathcal{L}(y_i^\tau; \nu_i, \sigma_i^2) + \sup_{\hat{\nu}, \hat{\sigma}^2} \sum_{\tau=s+1}^t \log \mathcal{L}(y_i^\tau; \hat{\nu}, \hat{\sigma}^2) \\ & - \sum_{\tau=1}^s \log \mathcal{L}(y_i^\tau; \nu_i, \sigma_i^2), \end{aligned} \quad (43)$$

where  $\mathcal{L}(\cdot; \nu, \sigma^2)$  denotes the likelihood function of a Gaussian distribution with mean  $\nu$  and variance  $\sigma^2$ . Note that  $\Psi_i(s)$  essentially measures the gain in likelihood by explaining the observed data during  $(s, t]$  with an alternative distribution. Based on the GLR statistic, the estimated change point is given as

$$\hat{t}_c(i) = \min \left\{ t \in [1, T] : \max_{s \in [1, t]} \Psi_i(s) \geq \eta \right\}, \quad (44)$$

where  $\eta$  is a positive constant. As a threshold, the value of  $\eta$  determines the false alarm and miss rates of the GLR test. Herein, we derive the functional dependence between the threshold and the false alarm rate, so that  $\eta$  can be chosen to achieve a desired false alarm probability,

$$\mathbb{P} \left( \max_{s \in [1, t]} \Psi_i(s) \geq \eta \middle| \mathcal{H}_0 \right) = \zeta. \quad (45)$$

With the maximum likelihood estimates for Gaussian parameters,

$$\hat{\nu}_i \doteq \frac{1}{t-s} \sum_{\tau=s+1}^t y_i^\tau, \quad \hat{\sigma}_i^2 \doteq \frac{1}{t-s} \sum_{\tau=s+1}^t (y_i^\tau - \hat{\nu}_i)^2, \quad (46)$$

the test statistic under  $\mathcal{H}_0$  can be formulated as

$$\Psi_i(s) = \sum_{\tau=s+1}^t \log \mathcal{L}(y_i^\tau; \hat{\nu}_i, \hat{\sigma}_i^2) - \log \mathcal{L}(y_i^\tau; \nu_i, \sigma_i^2) \quad (47)$$

$$= \sum_{\tau=s+1}^t \log \left[ \frac{\sigma_i}{\hat{\sigma}_i} \exp \left( \frac{(y_i^\tau - \nu_i)^2}{2\sigma_i^2} - \frac{(y_i^\tau - \hat{\nu}_i)^2}{2\hat{\sigma}_i^2} \right) \right] \quad (48)$$

$$= \frac{t-s}{2} (\log \sigma_i^2 - \log \hat{\sigma}_i^2 - 1) + \sum_{\tau=s+1}^t \frac{(y_i^\tau - \nu_i)^2}{2\sigma_i^2}, \quad (49)$$

where we substitute  $\hat{\sigma}_i^2$  in (46) to arrive at (49). Applying a Taylor expansion of  $\log \hat{\sigma}_i^2$  around  $\sigma_i^2$ , the test statistic is approximated as

$$2\Psi_i(s) \approx (t-s) \left[ \frac{1}{2} \left( \frac{\hat{\sigma}_i^2 - \sigma_i^2}{\sigma_i^2} \right)^2 - \frac{\hat{\sigma}_i^2}{\sigma_i^2} \right] + \sum_{\tau=s+1}^t \frac{(y_i^\tau - \nu_i)^2}{\sigma_i^2} \quad (50)$$

$$= \frac{t-s}{2} \left( \frac{\hat{\sigma}_i^2}{\sigma_i^2} - 1 \right)^2 + \left[ \frac{\sqrt{t-s}}{\sigma_i} (\hat{\nu}_i - \nu_i) \right]^2 \quad (51)$$

$$= \left[ \frac{(t-s)\hat{\sigma}_i^2/\sigma_i^2 - (t-s)}{\sqrt{2(t-s)}} \right]^2 + \left[ \frac{\sqrt{t-s}}{\sigma_i} (\hat{\nu}_i - \nu_i) \right]^2 \quad (52)$$

$$\approx \left[ \frac{(t-s)\hat{\sigma}_i^2/\sigma_i^2 - (t-s-1)}{\sqrt{2(t-s-1)}} \right]^2 + \left[ \frac{\sqrt{t-s}}{\sigma_i} (\hat{\nu}_i - \nu_i) \right]^2. \quad (53)$$

The sample mean and variance satisfy

$$(t-s)\hat{\sigma}_i^2/\sigma_i^2 \sim \chi^2(t-s-1), \quad \hat{\nu}_i \sim \mathcal{N} \left( \nu_i, \frac{\sigma_i^2}{t-s} \right), \quad (54)$$

where  $\chi^2(d)$  refers to the chi-squared distribution with  $d$  degrees of freedom. If we approximate  $\chi^2(t-s-1)$  by a Gaussian distribution, (53) can be considered as the sum of the squares of two Gaussian random variables. Furthermore, these two variables are independent because the sample mean and sample variance of Gaussian distributions are independent [44]. Thus, we can approximate  $2\Psi_i(s)$  by a chi-squared random variable with two degrees of freedom.

Note that  $\Psi_i(s)$  is computed with samples in  $(s, t]$ , as shown in (47). As a result, a strong positive correlation exists between the statistics calculated using overlapping subsets of the samples. For example,  $\Psi_i(s_1)$  and  $\Psi_i(s_2)$  are computed with overlapped samples in the interval  $(\max(s_1, s_2), t]$ . Therefore, we approximate the false alarm probability as

$$\begin{aligned} \mathbb{P} \left( \max_{s \in [1, t]} \Psi_i(s) \geq \eta \middle| \mathcal{H}_0 \right) & \approx \\ \mathbb{P} \left( 2\Psi_i(s=1) \geq 2\eta \middle| \mathcal{H}_0 \right) & \approx 1 - F_2(2\eta), \end{aligned} \quad (55)$$

where  $F_2(\cdot)$  stands for the cumulative distribution function of  $\chi^2(2)$ . To achieve a false alarm rate of  $\zeta$ , we can set the threshold to be

$$\eta(\zeta) = F_2^{-1}(1 - \zeta)/2. \quad (56)$$

While our proposed method may not be optimal, we emphasize that the approach can detect changes at the sub-graph level, resulting in improved sample and computational efficiency.

Once a change is detected for a sub-graph at time  $\hat{t}_c$ , the edge weights associated with that sub-graph should be estimated only using observed values during  $(\hat{t}_c, t]$ . The proposed CSL-UCB algorithm combined with change detection is denoted as the *CSL-UCB-CD* scheme.

## VI. NUMERICAL RESULTS

In this section, we numerically evaluate the performance of the proposed schemes in terms of both graph identification and intervention selection. Both stationary and non-stationary causal structures are considered for evaluation. For each Monte Carlo run, the causal structure is randomly generated, with the edge weights randomly sampled from the uniform distribution  $\mathcal{U}(-2, 2)$ . The exogenous variables are independently sampled in each time step from the Gaussian distribution  $\mathcal{N}(1, 1)$ . To estimate the empirical mutual information, the  $k$ -nearest neighbor distances based approach is employed [45]. For each set of parameters, we repeat the Monte Carlo run  $M = 100$  times<sup>2</sup>

### A. Performance for Causal Graph Identification

We start by evaluating the causal graph learning ability of the proposed CSL scheme. For comparison, we consider the GOLEM algorithm from [10], the DAGMA algorithm from [11] and the CSL-nw scheme, which refers to the proposed CSL scheme without edge weight regularization. The GOLEM algorithm provides an estimate of the observational weight matrix in step  $t$  as

$$\arg \min_{\hat{\mathbf{B}} \in \mathbb{R}^{N \times N}} \left\{ -\log \mathcal{L}(\mathbf{X}^{1:t}; \hat{\mathbf{B}}) + \beta_1 \|\hat{\mathbf{B}}\|_1 + \beta_2 [\text{tr}(e^{\hat{\mathbf{B}} \circ \hat{\mathbf{B}}}) - N] \right\}, \quad (57)$$

where  $\beta_1$  and  $\beta_2$  are the penalty coefficients,  $\circ$  denotes the Hadamard product and the negative log-likelihood is

$$-\log \mathcal{L}(\mathbf{X}^{1:t}; \hat{\mathbf{B}}) = \frac{N}{2} \log \left( \sum_{i=1}^N \sum_{\tau=1}^t (x_i^\tau - \hat{\mathbf{B}}_i^\top x^\tau)^2 \right) - \log |\det(\mathbf{I} - \hat{\mathbf{B}})|. \quad (58)$$

With a different characterization of acyclicity, the DAGMA algorithm estimates the weight matrix as

$$\arg \min_{\hat{\mathbf{B}} \in \mathbb{R}^{N \times N}} \left\{ -\log \mathcal{L}(\mathbf{X}^{1:t}; \hat{\mathbf{B}}) + \beta_1 \|\hat{\mathbf{B}}\|_1 - \log \det(\beta_3 \mathbf{I} - \hat{\mathbf{B}} \circ \hat{\mathbf{B}}) + N \log \beta_3 \right\}, \quad (59)$$

where  $\beta_3$  is the log-determinant coefficient.

Fig. 2 plots the graph FN rate as a function of graph sizes, with two sample sizes,  $T = 200$  and  $400$ . The graph FN rate is defined as

$$\frac{1}{M} \sum_{m=1}^M \mathbb{I}(\exists (i, j) \in \mathcal{V} \times \mathcal{V} : \mathbf{B}_{ij}(m) \neq 0, \hat{\mathbf{B}}_{ij}(m) = 0), \quad (60)$$

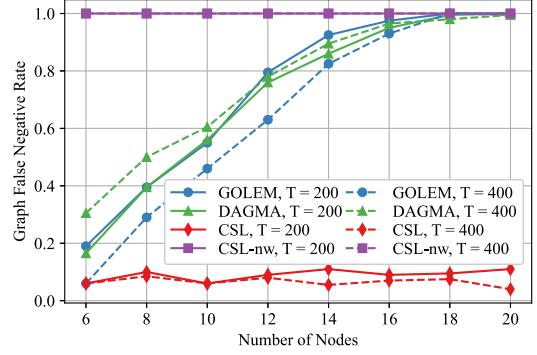


Fig. 2. Graph false negative rate as a function of graph sizes.

where  $\mathbf{B}(m)$  and  $\hat{\mathbf{B}}(m)$  represent the true and estimated edge weight matrices in the  $m$ -th Monte Carlo run. Observe that the CSL-nw scheme has a graph FN rate of 100% in all cases, which indicates that without edge weight regularization, the estimated graphs always have FN errors. The GOLEM scheme performs slightly better than the DAGMA scheme, while both of them perform much worse than the proposed CSL scheme. This is due to the fact that GOLEM and DAGMA do not distinguish between FP and FN errors, while the CSL scheme is designed to explicitly consider the FN rate. On average, the CSL scheme achieves a graph FN rate that is lower than the GOLEM scheme by 81.1% and the DAGMA scheme by 86.9%. Interestingly, given a fixed number of samples, as the graph size increases, the graph FN rates for both GOLEM and DAGMA increase rapidly. In contrast, The CSL scheme does not suffer much from the increased graph size, as learning sub-graphs has higher sample efficiency.

To better understand the balance between false positives and false negatives, we consider precision and recall for evaluation, which are defined as

$$\text{precision} \doteq \frac{\sum_{i,j} \mathbb{I}(\mathbf{B}_{ij} \neq 0, \hat{\mathbf{B}}_{ij} \neq 0)}{\sum_{i,j} \mathbb{I}(\hat{\mathbf{B}}_{ij} \neq 0)}, \quad (61)$$

$$\text{recall} \doteq \frac{\sum_{i,j} \mathbb{I}(\mathbf{B}_{ij} \neq 0, \hat{\mathbf{B}}_{ij} \neq 0)}{\sum_{i,j} \mathbb{I}(\mathbf{B}_{ij} \neq 0)}. \quad (62)$$

Note that precision and recall are inversely correlated with the numbers of FP and FN, respectively. Figure 3 and 4 plot precision and recall as functions of graph sizes. We observe that the proposed CSL scheme achieves a close to one recall by focusing on the reduction of FN errors. On average, the recall rates of GOLEM, DAGMA and CSL are 92.7%, 93.8% and 99.6%. On the other hand, the CSL scheme does not perform significantly better than GOLEM and DAGMA in terms of precision, especially for small graphs. The average precision rates achieved by GOLEM, DAGMA and CSL are 93.5%, 90.9% and 96.7%. We also observe that both precision and recall of the CSL scheme increase as the graph becomes larger, while GOLEM and DAGMA have the opposite trend. It suggests that for the CSL scheme, as  $N$  increases, the number of mistaken edges increases more slowly than the totally number of edges, which increases quadratically in  $N$ .

<sup>2</sup>The code is available at [https://github.com/CalixPeng/Causal\\_Bandit](https://github.com/CalixPeng/Causal_Bandit)

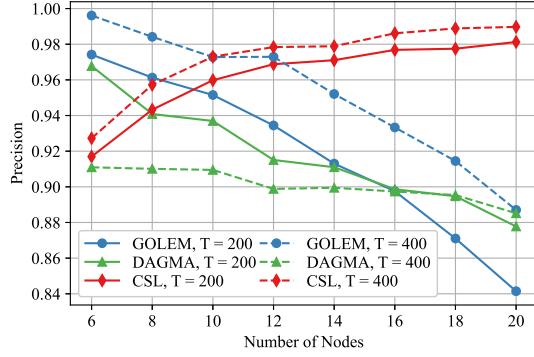


Fig. 3. Precision as a function of graph sizes.

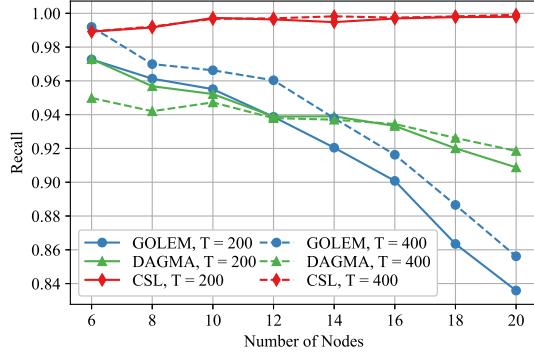


Fig. 4. Recall as a function of graph sizes.

For completeness, we also include performance for both types of graph errors, measured by the normalized structural Hamming distance (NSHD), which is defined as

$$\text{NSHD} \doteq \sum_{i,j} \left[ \mathbb{I}(B_{ij} = 0, \hat{B}_{ij} \neq 0) + \mathbb{I}(B_{ij} \neq 0, \hat{B}_{ij} = 0) - \mathbb{I}(B_{ij} \neq 0, \hat{B}_{ij} = 0, \hat{B}_{ji} \neq 0) \right] / N^2. \quad (63)$$

In Fig. 5, the NSHD is plotted as a function of graph sizes. We observe that in terms of NSHD, the CSL scheme does not perform significantly better than the GOLEM and DAGMA schemes. Particularly when the graph is small ( $N = 6$ ), DAGMA achieves the lowest normalized SHD for 200 samples, while GOLEM achieves the lowest SHD for 400 samples. Since both GOLEM and DAGMA learn the whole graph with gradient-based optimization, a sufficient amount of samples are necessary to achieve desirable performance. As the graph size increases, CSL offers strongly superior performance due to its need for fewer samples than the other methods.

#### B. Performance for Stationary Causal Bandits

Next, we consider the proposed CSL-UCB scheme with optimized intervention design in causal bandits. Since the goal is to maximize the cumulative reward, we start by revealing the impact of FN errors on reward estimation. To measure the impact, the estimation errors are classified into two groups

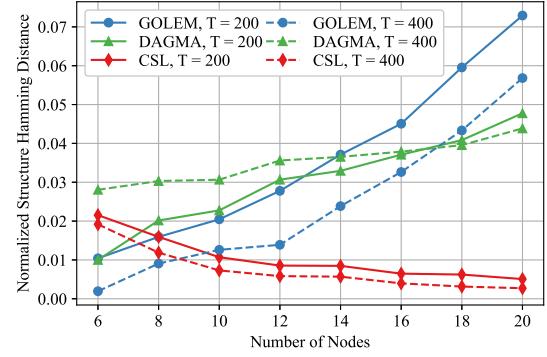


Fig. 5. Normalized Hamming distance as a function of graph sizes.

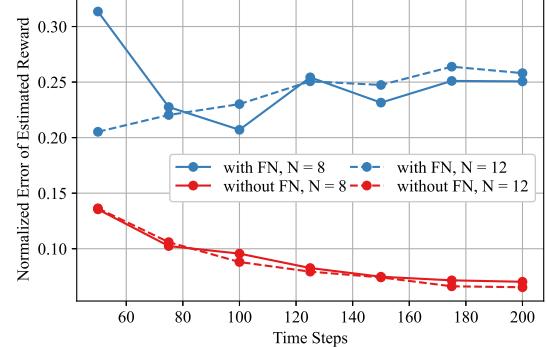


Fig. 6. Normalized error of estimated reward as a function of time steps.

based on whether FN error exists in the estimated graph. In Fig. 6, the normalized empirical reward error is plotted as a function of time for two cases: with or without FN errors. Without FN errors, the reward estimation error decreases monotonically as more samples are collected, which matches the analysis in Section III-C. With the presence of FN errors, the estimation error does not decrease as more samples are collected. On average, the estimation error is 62.8% lower in the absence of FN errors. The result is consistent with our claim that FN errors have a significant impact on reward maximization in causal bandits. The result also validates our choice to minimize the FN rate in causal graph identification.

To evaluate the CSL-UCB scheme, we focus on graphs with size  $N = 10$  and simulate causal bandits over a horizon of  $T = 1500$  time steps. To alleviate computation cost, the graphs are updated with new collected samples every 20 time slots. For comparison, we first consider the vanilla UCB algorithm [29], which selects an intervention in each step as

$$\arg \max_{\mathbf{a}} \left\{ \frac{\sum_{\tau=1}^t \mathbb{I}(\mathbf{a}^\tau = \mathbf{a}) X_N^\tau}{\sum_{\tau=1}^t \mathbb{I}(\mathbf{a}^\tau = \mathbf{a})} + \alpha' \sqrt{\frac{\ln t}{\sum_{\tau=1}^t \mathbb{I}(\mathbf{a}^\tau = \mathbf{a})}} \right\}, \quad (64)$$

where  $\alpha'$  is the parameter that controls the exploration level. Notice that the vanilla UCB algorithm does not exploit the causal structure and its sample complexity scales exponentially as  $2^N$  [46]. Another scheme we consider for comparison is

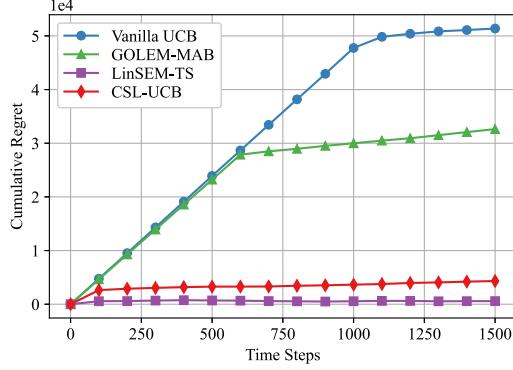


Fig. 7. Cumulative regret of stationary bandits as a function of time steps.

the GOLEM [10] algorithm based decision-making for MABs, denoted as the GOLEM-MAB scheme. The horizon is divided into two parts, such that causal structure identification is conducted in the first part, while rewards are collected in the second part by exploiting the learned causal graph. Note that since GOLEM identifies the whole graph, intervention should be fixed as  $a = 0$  for learning  $\mathcal{B}$  and  $a = 1$  for learning  $\mathcal{B}'$ . Lastly, we employ the LinSEM-TS scheme proposed in [20], which provides an upper bound on performance, as it utilizes knowledge of the causal topology.

Figure 7 plots the cumulative regret as a function of time for four different algorithms and the associated probability of selecting the optimal intervention is provided in Figure 8. The vanilla UCB algorithm gains information about the optimal intervention after exploring every possible intervention at  $t = 2^N = 1024$ . By the end of the horizon, vanilla UCB selects the optimal intervention 47.2% of the time. GOLEM-MAB tries to identify the causal graph in the first 600 steps, and exploits this knowledge to achieve an optimal intervention selection ratio of 71.0%. The value of  $t = 600$  was selected via trial and error to optimize the performance of GOLEM-MAB. The proposed CSL-UCB scheme gains and leverages causal knowledge in an alternating manner, yielding an optimal intervention selection ratio of 79.0%. With topology knowledge, LinSEM-TS is able to select the optimal intervention with probability 94.2%. In terms of cumulative regret, the CSL-UCB scheme achieves a 91.6% improvement compared with vanilla UCB and an 86.8% improvement compared with GOLEM-MAB. The LinSEM-TS algorithm can further reduce 86.6% of the cumulative regret the CSL-UCB scheme achieves when the causal structure is fully known.

An interesting observation is that, after 600 time steps for graph identification (300 for each,  $\mathcal{B}$  and  $\mathcal{B}'$ ), the graphs estimated by GOLEM are close to the ground truth with respect to the SHD. However, even rejecting one existing edge (one false negative error) could result in a significantly different expected reward. Since edges are highly possible to be part of the causal flows that influence the reward node, mistakenly rejecting an edge can cut off one or multiple relevant causal flows and thus significantly impact reward estimation. The relatively poor performance of GOLEM-MAB

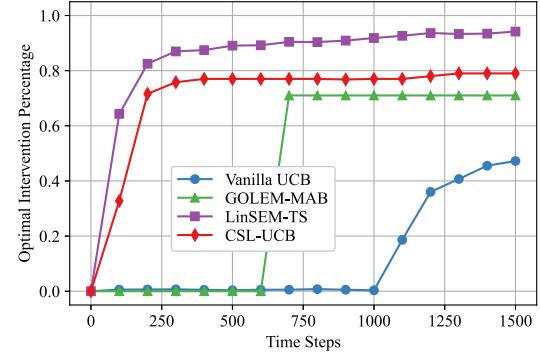


Fig. 8. Optimal percentage of stationary bandits as a function of time steps.

is not due to its generic graph identification ability, but the equal treatment of all errors and its low sample efficiency. We emphasize again that causal bandit optimization and causal graph identification are two different problems.

### C. Performance for Non-stationary Causal Bandits

Lastly, we evaluate the reward-earning ability of the proposed CSL-UCB-CD scheme in non-stationary causal bandits. Structure change occurs twice at  $t \in \{1000, 2000\}$ , where a random subset of the causal mechanisms are regenerated. For comparison, the vanilla UCB algorithm can adjust its estimate of the rewards after structure changes using newly collected samples. The GOLEM-MAB algorithm combined with a change detection mechanism is denoted as GOLEM-MAB-CD, which must re-learns the whole causal graph when a change is detected. The LinSEM-TS algorithm initially has full knowledge of the causal structure, but is unaware of the upcoming changes.

In Fig. 9, the cumulative regret is plotted as a function of time. The percentage of selecting the optimal intervention as a function of time is provided in Fig. 10. On average, the change detection mechanism is able to detect 89.3% of the sub-graph changes, with an average delay of 13.1 slots. The cumulative regret achieved by the CSL-UCB-CD scheme is 83.0% lower than vanilla UCB, 83.2% lower than GOLEM-MAB-CD and 73.5% lower than LinSEM-TS. Although LinSEM-TS obtains the lowest regret before structure change, its performance degrades dramatically when the causal knowledge becomes incorrect.

The vanilla UCB scheme performs the worst, not due to structure changes but rather, due to the lack of samples. Averaged over time, the optimal intervention ratio of the vanilla UCB scheme is only 7.0%. As illustrated in Fig. 10, the optimal percentage of GOLEM-MAB-CD drops to 0% after each change-point, which indicates that it is able to detect the changes and start to re-learn the causal graph. However, as learning the whole graph requires a large amount of samples, GOLEM-MAB-CD is only able to select the optimal intervention 18.3% of the time. Since LinSEM-TS suffers from incorrect causal knowledge caused by structure changes, it selects the optimal intervention only 32.6% of the time. The

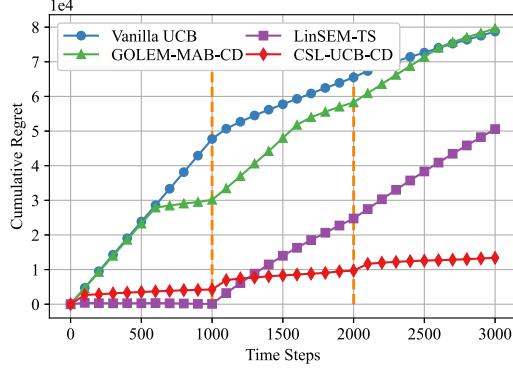


Fig. 9. Cumulative regret of non-stationary bandits as a function of time steps. Structure change occurs twice, as marked by the dashed lines.

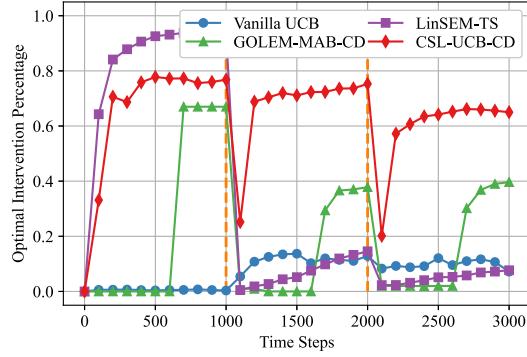


Fig. 10. Optimal percentage of non-stationary bandits as a function of time steps. Structure change occurs twice, as marked by the dashed lines.

CSL-UCB-CD scheme combines sub-graph change detection with sub-graph learning, and select the optimal intervention 63.8% of the time.

Compared with the stationary case, it is observed that the regret gap between vanilla UCB and GOLEM-MAB-CD in non-stationary bandits is smaller. The major reason is that, utilizing samples from the previous model, vanilla UCB implicitly assumes those samples contain information about the current structure. This is true, because the changes occur only for a subset of the causal mechanisms. We also observe that CSL-UCB-CD performs better than GOLEM-MAB-CD mainly due to its fast recovery from structure changes. This fast recovery property comes from the fact that re-learning sub-graphs requires fewer samples than re-learning whole graphs.

## VII. CONCLUSIONS

In this paper, we investigate the causal bandit problem with no knowledge of the causal graph topology and the interventional distribution. Distinguishing between the two types of error, a sub-graph learning-based causal identification approach is proposed, which makes efficient use of the limited data. Due to learning sub-graphs versus entire graphs, the complexity of the proposed method is significantly lower than traditional methods. Analysis reveals that optimal intervention

design is much more sensitive to false negative errors in graph identification. This feature is directly taken into consideration for the sub-graph learning algorithm. To balance exploration and exploitation, we analyze and propose an uncertainty bound. Further, sub-graph level change detection mechanism is proposed, which enables the proposed approach to adapt to non-stationary bandits. Numerical results show that the overall approach is able to identify the optimal intervention with much lower sample and computational complexity than previous methods. Furthermore, performance is improved over prior approaches due to the direct consideration of false negative errors. The proposed approach requires 67% fewer samples to learn the causal structure and achieves a lower cumulative regret by 85%, compared to the existing schemes.

## APPENDIX A PROOF OF PROPOSITION 2

*Proof:* We start by proving a general statement for continuous random variables. That is, the following inequality holds,

$$I(Z_1 + Z_2; Z_3 + Z_4) \leq I(Z_1; Z_3) + I(Z_2; Z_4), \quad (65)$$

if  $Z_1 \perp Z_2, Z_4$  and  $Z_3 \perp Z_2, Z_4$ . To prove that, we first employ the data processing inequality [37] to bound the mutual information as

$$I(Z_1 + Z_2; Z_3 + Z_4) \leq I(Z_1, Z_2; Z_3 + Z_4) \quad (66)$$

$$\leq I(Z_1, Z_2; Z_3, Z_4). \quad (67)$$

Then, based on the definition of mutual information, we have

$$I(Z_1, Z_2; Z_3, Z_4) = \int f(z_1, z_2, z_3, z_4) \log \frac{f(z_1, z_2, z_3, z_4)}{f(z_1, z_2)f(z_3, z_4)} dz \quad (68)$$

$$= \int f(z_1, z_3)f(z_2, z_4) \log \frac{f(z_1, z_3)}{f(z_1)f(z_3)} dz + \int f(z_1, z_3)f(z_2, z_4) \log \frac{f(z_2, z_4)}{f(z_2)f(z_4)} dz \quad (69)$$

$$= \int f(z_1, z_3) \log \frac{f(z_1, z_3)}{f(z_1)f(z_3)} dz_1 dz_3 \cdot 1 + \int f(z_2, z_4) \log \frac{f(z_2, z_4)}{f(z_2)f(z_4)} dz_2 dz_4 \cdot 1 \quad (70)$$

$$= I(Z_1; Z_3) + I(Z_2; Z_4). \quad (71)$$

Given this inequality and the fact that the exogenous variables are independent, the substitution

$$\begin{aligned} Z_1 &= \sum_{l \notin \mathcal{A}_i(a_i)} [\hat{C}_a^t]_{li} \epsilon_l - \tilde{\epsilon}_i, \\ Z_2 &= \sum_{k \in \mathcal{A}_i(a_i)} ([\hat{C}_a^t]_{ki} - [C_a]_{ki}) \epsilon_k, \\ Z_3 &= \sum_{l \notin \mathcal{A}_i(a_i)} [C_a]_{lj} \epsilon_l, \quad Z_4 = \sum_{k \in \mathcal{A}_i(a_i)} [C_a]_{kj} \epsilon_k \end{aligned} \quad (72)$$

completes the proof.  $\square$

## APPENDIX B PROOF OF PROPOSITION 3

*Proof:* Consider two sets  $\mathcal{Q}$ ,  $\mathcal{S}$  such that  $\mathcal{Q} \subset \mathcal{S}$  and the observed values of nodes up to step  $t$  satisfy

$$\mathbf{X}_i^{1:t}(a_i) = [\mathbf{B}_a]_{i,\mathcal{Q}}^\top \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i) + \mathbf{0}^\top \mathbf{X}_{\mathcal{S} \setminus \mathcal{Q}}^{1:t}(a_i) + \epsilon_i^{1:t}(a_i) \quad (73)$$

$$= [\mathbf{B}_a]_{i,\mathcal{S}}^\top \mathbf{X}_{\mathcal{S}}^{1:t}(a_i) + \epsilon_i^{1:t}(a_i) \quad (74)$$

where  $\mathbf{0}$  is a vector of zeros and  $\mathcal{S} \setminus \mathcal{Q}$  denotes the set difference of  $\mathcal{S}$  and  $\mathcal{Q}$ . The expected vector of weights associated with nodes in  $\mathcal{S}$ , is given by the MMSE estimation as

$$\begin{aligned} \mathbb{E}[[\hat{\mathbf{B}}_a^t]_{i,\mathcal{S}} | \mathbf{X}_{\mathcal{S}}^{1:t}(a_i)] &= \mathbb{E}[[\mathbf{X}_{\mathcal{S}}^{1:t}(a_i) \mathbf{X}_{\mathcal{S}}^{1:t}(a_i)^\top]^{-1} \mathbf{X}_{\mathcal{S}}^{1:t}(a_i) \\ &\quad \cdot (\mathbf{X}_{\mathcal{S}}^{1:t}(a_i)^\top [\mathbf{B}_a]_{i,\mathcal{S}} + \tilde{\epsilon}_i^{1:t}(a_i)^\top) | \mathbf{X}_{\mathcal{S}}^{1:t}(a_i)] \quad (75) \end{aligned}$$

$$= [\mathbf{B}_a]_{i,\mathcal{S}} + \mathbf{0} = [\mathbf{B}_a]_{i,\mathcal{S}}. \quad (76)$$

With the substitution  $\mathcal{Q} = \mathcal{P}_i(a_i)$ ,  $\mathcal{S} = \widehat{\mathcal{P}}_i(a_i)$ , we obtain the effect of including non-existent parents (false positive error). Equation (76) holds because the exogenous variable  $\tilde{\epsilon}_i$  is independent of the variables associated with the nodes in  $\mathcal{S} = \widehat{\mathcal{P}}_i(a_i)$ . Without a false negative error,  $\widehat{\mathcal{P}}_i(a_i)$  should not contain node  $i$  and its descendants.

Next, consider the case that the values are generated as

$$\mathbf{X}_i^{1:t}(a_i) = [\mathbf{B}_a]_{i,\mathcal{Q}}^\top \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i) + [\mathbf{B}_a]_{i,\mathcal{S} \setminus \mathcal{Q}}^\top \mathbf{X}_{\mathcal{S} \setminus \mathcal{Q}}^{1:t}(a_i) + \epsilon_i^{1:t}(a_i). \quad (77)$$

The vector of expected weights associated with  $\mathcal{Q}$  is given by the MMSE estimation as

$$\begin{aligned} \mathbb{E}[[\hat{\mathbf{B}}_a^t]_{i,\mathcal{Q}} | \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i)] &= \mathbb{E}[[\mathbf{X}_{\mathcal{Q}}^{1:t}(a_i) \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i)^\top]^{-1} \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i) \\ &\quad \cdot (\mathbf{X}_{\mathcal{Q}}^{1:t}(a_i)^\top [\mathbf{B}_a]_{i,\mathcal{Q}} + \mathbf{X}_{\mathcal{S} \setminus \mathcal{Q}}^{1:t}(a_i)^\top [\mathbf{B}_a]_{i,\mathcal{S} \setminus \mathcal{Q}} + \\ &\quad \tilde{\epsilon}_i^{1:t}(a_i)^\top) | \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i)] \quad (78) \end{aligned}$$

$$\begin{aligned} &= [\mathbf{B}_a]_{i,\mathcal{Q}} + [\mathbf{X}_{\mathcal{Q}}^{1:t}(a_i) \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i)^\top]^{-1} \mathbb{E}[\mathbf{X}_{\mathcal{Q}}^{1:t}(a_i) \cdot \\ &\quad \mathbf{X}_{\mathcal{S} \setminus \mathcal{Q}}^{1:t}(a_i)^\top [\mathbf{B}_a]_{i,\mathcal{S} \setminus \mathcal{Q}} | \mathbf{X}_{\mathcal{Q}}^{1:t}(a_i)]. \quad (79) \end{aligned}$$

Note that having FN errors is equivalent to have  $\widehat{\mathcal{P}}_i(a_i) \subset \mathcal{P}_i(a_i)$ . Thus with the substitution  $\mathcal{Q} = \widehat{\mathcal{P}}_i(a_i)$ ,  $\mathcal{S} = \mathcal{P}_i(a_i)$ , we obtain the effect of FN errors (rejecting actual edges).  $\square$

## APPENDIX C PROOF OF LEMMA 1

*Proof:* When there is no false negative error, each weight error vector follows a multivariate normal distribution according to the theory of linear least-square estimation,

$$[\Delta \mathbf{B}_a^t]_i \sim \mathcal{N}(\mathbf{0}, \Phi_a^t(i)), \quad (80)$$

where the covariance matrix can be decomposed as  $\Phi_a^t(i) = \mathbf{Q}_i \Lambda_i \mathbf{Q}_i^\top$ . With the orthonormal matrix  $\mathbf{Q}_i$ , we have

$$\|[\Delta \mathbf{B}_a^t]_i\|_2 = \|\mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i\|_2, \quad \mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i \sim \mathcal{N}(\mathbf{0}, \Lambda_i). \quad (81)$$

Since the squared norm of  $\mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i$  is a weighted sum of independent squared Gaussian variables, it can be bounded from above as

$$\|\mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i\|_2^2 \leq \lambda_{\max}(\Lambda_i) \|\Lambda_i^{-1/2} \mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i\|_2^2 \quad (82)$$

$$= \lambda_{\max}(\Phi_a^t(i)) \|\Lambda_i^{-1/2} \mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i\|_2^2. \quad (83)$$

Notice that  $\|\Lambda_i^{-1/2} \mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i\|_2^2$  follows a Chi-squared distribution with  $N$  degrees of freedom, with the  $m$ -th moment defined as

$$\mathbb{E} \left[ \left( \|\Lambda_i^{-1/2} \mathbf{Q}_i^\top [\Delta \mathbf{B}_a^t]_i\|_2^2 \right)^m \right] = 2^m \frac{\Gamma(m + N/2)}{\Gamma(N/2)}, \quad (84)$$

where  $\Gamma(\cdot)$  represents the Gamma function. Thus we can upper bound the  $m$ -th moment of the error norm as

$$\mathbb{E} \left[ \left\| [\Delta \mathbf{B}_a^t]_i \right\|_2^m \right] \stackrel{(a)}{\leq} \sqrt{\mathbb{E} \left[ \left( \left\| [\Delta \mathbf{B}_a^t]_i \right\|_2^2 \right)^m \right]} \quad (85)$$

$$\stackrel{(b)}{\leq} \sqrt{\lambda_{\max}(\Phi_a^t(i))^m 2^m \frac{\Gamma(m + N/2)}{\Gamma(N/2)}} \quad (86)$$

$$\stackrel{(c)}{\leq} m! \sqrt{\frac{4N}{N+2}} \left[ \lambda_{\max}(\Phi_a^t(i)) \frac{N+2}{4} \right]^{m/2}, \quad (87)$$

where (a) comes from Jensen's inequality while (b) is a direct consequence of (81), (83) and (84). To show (c), we first rewrite the quotient of the Gamma function as

$$2^m \frac{\Gamma(m + N/2)}{\Gamma(N/2)} = N(N+2) \cdots (N+2m-2) \quad (88)$$

$$= (m!)^2 \cdot \prod_{k=1}^m \frac{N+2k-2}{k^2}. \quad (89)$$

When  $m \geq 2$ , we have the following inequalities

$$\frac{N+2k-2}{k^2} = \frac{N-2}{k^2} + \frac{2}{k} \leq \frac{N+2}{4}, \quad \forall k \geq 2, \quad (90)$$

$$2^m \frac{\Gamma(m + N/2)}{\Gamma(N/2)} \leq (m!)^2 \cdot N \cdot \left( \frac{N+2}{4} \right)^{m-1}, \quad (91)$$

which shows (c) and thus completes the proof.  $\square$

## APPENDIX D PROOF OF LEMMA 2

*Proof:* For a specific intervention  $\mathbf{a}$ , consider the dilation [47] of the weight error matrix,

$$\mathbf{H}(\Delta \mathbf{B}_a^t) \doteq \begin{bmatrix} \mathbf{0} & \Delta \mathbf{B}_a^t \\ (\Delta \mathbf{B}_a^t)^\top & \mathbf{0} \end{bmatrix}. \quad (92)$$

An important property is that the dilation preserves spectral information:

$$\lambda_{\max}(\mathbf{H}(\Delta \mathbf{B}_a^t)) = \sigma_{\max}(\mathbf{H}(\Delta \mathbf{B}_a^t)) = \sigma_{\max}(\Delta \mathbf{B}_a^t), \quad (93)$$

Since the weight matrix consists of independently learned columns, we decompose the dilation as

$$\mathbf{H}(\Delta \mathbf{B}_a^t) = \sum_{i=1}^N \mathbf{H}_i, \quad \mathbf{H}_i \doteq \begin{bmatrix} \mathbf{0} & (\Delta \mathbf{B}_a^t)_i \\ (\Delta \mathbf{B}_a^t)_i^\top & \mathbf{0} \end{bmatrix}, \quad (94)$$

where the  $N \times N$  matrix  $(\Delta \mathbf{B}_a^t)_i$  is composed of the  $i$ -th column of  $\Delta \mathbf{B}_a^t$  and zero everywhere else. Applying block matrix multiplication repeatedly, we have

$$\mathbf{H}_i^2 = \begin{bmatrix} (\Delta \mathbf{B}_a^t)_i (\Delta \mathbf{B}_a^t)_i^\top & \mathbf{0} \\ \mathbf{0} & (\Delta \mathbf{B}_a^t)_i^\top (\Delta \mathbf{B}_a^t)_i \end{bmatrix}, \quad (95)$$

$$\mathbf{H}_i^m = \begin{cases} \|[\Delta \mathbf{B}_a^t]_i\|_2^{m-1} \cdot \mathbf{H}_i, & m \text{ is odd} \\ \|[\Delta \mathbf{B}_a^t]_i\|_2^{m-2} \cdot \mathbf{H}_i^2, & m \text{ is even} \end{cases}. \quad (96)$$

We further observe that

$$\lambda_{\max}(\mathbf{H}_i) = \|[\Delta \mathbf{B}_a^t]_i\|_2, \quad \lambda_{\max}(\mathbf{H}_i^2) = \|[\Delta \mathbf{B}_a^t]_i\|_2^2, \quad (97)$$

which enables us to draw a general conclusion that

$$\mathbf{H}_i^m \preceq \|[\Delta \mathbf{B}_a^t]_i\|_2^m \cdot \mathbf{I}, \quad (98)$$

where  $\preceq$  represents the matrix semi-definite ordering. Further, taking expectation and applying Lemma 1, we have

$$\mathbb{E}[\mathbf{H}_i] = \mathbf{0} \quad (99)$$

$$\mathbb{E}[\mathbf{H}_i^m] \preceq m! \sqrt{\frac{4N}{N+2}} \left[ \lambda_{\max}(\Phi_a^t(i)) \frac{N+2}{4} \right]^{\frac{m}{2}} \mathbf{I}, \quad m \geq 2. \quad (100)$$

These properties enable us to apply the matrix Bernstein inequality for the sub-exponential case [40] and obtain

$$\mathbb{P}\{\sigma_{\max}(\Delta \mathbf{B}_a^t) \geq \xi\} = \mathbb{P}\{\lambda_{\max}(\mathbf{H}(\Delta \mathbf{B}_a^t)) \geq \xi\} \quad (101)$$

$$= \mathbb{P}\left\{\lambda_{\max}\left(\sum_{i=1}^N \mathbf{H}_i\right) \geq \xi\right\} \quad (102)$$

$$\leq 2N \exp \left[ -\frac{\xi^2}{4\sqrt{N(N+2)} \cdot \sum_{i=1}^N \lambda_{\max}(\Phi_a^t(i))} \right], \quad (103)$$

which is valid when

$$\xi \leq \sqrt{\frac{4N}{\max_i \lambda_{\max}(\Phi_a^t(i))}} \cdot \sum_{i=1}^N \lambda_{\max}(\Phi_a^t(i)). \quad (104)$$

Lastly, we denote the probability bound by  $\delta$  and solve  $\xi$  from

$$\delta = 2N \exp \left[ -\frac{\xi^2}{4\sqrt{N(N+2)} \cdot \sum_{i=1}^N \lambda_{\max}(\Phi_a^t(i))} \right], \quad (105)$$

to arrive at the desired inequality.  $\square$

## APPENDIX E PROOF OF THEOREM 1

*Proof:* First, recognize that the error of estimated reward is the last element of  $\Delta \mu_a^t$ . To understand the behavior of  $\Delta \mu_a^t$ , consider the following two linear systems,

$$(\mathbf{I} - \mathbf{B}_a)^\top \mu_a = \nu, \quad (106)$$

$$(\mathbf{I} - \hat{\mathbf{B}}_a^t)^\top \hat{\mu}_a^t = \nu, \quad (107)$$

where (106) and (107) represent the true and estimated systems, respectively. Based on the Woodbury matrix identity [41], we have

$$(\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-1} = (\mathbf{I} - \mathbf{B}_a)^{-1} + (\mathbf{I} - \mathbf{B}_a)^{-1} \Delta \mathbf{B}_a^t (\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-1}. \quad (108)$$

Combining (107) with (108), we can express  $\Delta \mu_a$  as

$$\Delta \mu_a^t = -\mu_a + (\mathbf{I} - \mathbf{B}_a)^{-\top} \nu + (\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-\top} (\Delta \mathbf{B}_a^t)^\top (\mathbf{I} - \mathbf{B}_a)^{-\top} \nu \quad (109)$$

$$= (\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-\top} (\Delta \mathbf{B}_a^t)^\top \mu_a, \quad (110)$$

where (110) is a direct consequence of (106). Now, we can bound the error of estimated reward as

$$|[\Delta \mu_a^t]_N| = \left| [(\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-1}]_N^\top (\Delta \mathbf{B}_a^t)^\top \mu_a \right| \quad (111)$$

$$\leq \left\| [(\mathbf{I} - \hat{\mathbf{B}}_a^t)^{-1}]_N \right\|_2 \cdot \sigma_{\max}(\Delta \mathbf{B}_a^t) \cdot \|\mu_a\|_2. \quad (112)$$

Based on that, we employ Lemma 2 to provide a high probability bound on the largest singular value, which completes the proof.  $\square$

## REFERENCES

- [1] C. Peng, D. Zhang, and U. Mitra, “Graph identification and upper confidence evaluation for causal bandits with linear models,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 7165–7169.
- [2] J. Pearl, *Causality*. Cambridge university press, 2009.
- [3] S. C. Wu and A. L. Swindlehurst, “Algorithms and bounds for dynamic causal modeling of brain connectivity,” *IEEE transactions on signal processing*, vol. 61, no. 11, pp. 2990–3001, 2013.
- [4] E. L. Ogburn, O. Sofrygin, I. Diaz, and M. J. Van der Laan, “Causal inference for social network data,” *Journal of the American Statistical Association*, vol. 119, no. 545, pp. 597–611, 2024.
- [5] P. Spirtes, C. Glymour, and R. Scheines, *Causation, prediction, and search*. MIT press, 2001.
- [6] D. M. Chickering, “Optimal structure identification with greedy search,” *Journal of machine learning research*, vol. 3, no. Nov, pp. 507–554, 2002.
- [7] N. Meinshausen and P. Bühlmann, “High-dimensional graphs and variable selection with the lasso,” *The Annals of Statistics*, pp. 1436–1462, 2006.
- [8] M. Koivisto, “Advances in exact bayesian structure discovery in bayesian networks,” in *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, 2006, pp. 241–248.
- [9] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing, “Dags with no tears: Continuous optimization for structure learning,” *Advances in neural information processing systems*, vol. 31, 2018.
- [10] I. Ng, A. Ghassami, and K. Zhang, “On the role of sparsity and dag constraints for learning linear dags,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 17943–17954, 2020.
- [11] K. Bello, B. Aragam, and P. Ravikumar, “Dagma: Learning dags via matrices and a log-determinant acyclicity characterization,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 8226–8239, 2022.
- [12] C. Cui, P. Banelli, and P. M. Djurić, “Topology inference of directed graphs by gaussian processes with sparsity constraints,” *IEEE Transactions on Signal Processing*, 2024.
- [13] E. Bareinboim, A. Forney, and J. Pearl, “Bandits with unobserved confounders: A causal approach,” *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [14] F. Lattimore, T. Lattimore, and M. D. Reid, “Causal bandits: Learning good interventions via causal inference,” *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [15] S. Liu, K. C. See, K. Y. Ngiam, L. A. Celi, X. Sun, and M. Feng, “Reinforcement learning for clinical decision support in critical care: comprehensive review,” *Journal of medical Internet research*, vol. 22, no. 7, p. e18477, 2020.
- [16] J. Xu, T. Xing, and M. Van Der Schaar, “Personalized course sequence recommendations,” *IEEE Transactions on Signal Processing*, vol. 64, no. 20, pp. 5340–5352, 2016.
- [17] X. Cai and Z. Ye, “Gaussian weighting reversion strategy for accurate online portfolio selection,” *IEEE Transactions on Signal Processing*, vol. 67, no. 21, pp. 5558–5570, 2019.
- [18] V. Nair, V. Patil, and G. Sinha, “Budgeted and non-budgeted causal bandits,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 2017–2025.
- [19] A. Maiti, V. Nair, and G. Sinha, “A causal bandit approach to learning good atomic interventions in presence of unobserved confounders,” in *The 38th Conference on Uncertainty in Artificial Intelligence*, 2022.
- [20] B. Varici, K. Shanmugam, P. Sattigeri, and A. Tajer, “Causal bandits for linear structural equation models,” *Journal of Machine Learning Research*, vol. 24, no. 297, pp. 1–59, 2023.
- [21] S. Sussex, A. Makarova, and A. Krause, “Model-based causal bayesian optimization,” in *The Eleventh International Conference on Learning Representations*, 2022.

- [22] Z. Yan, D. Wei, D. A. Katz, P. Sattigeri, and A. Tajer, “Causal bandits with general causal models and interventions,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2024, pp. 4609–4617.
- [23] A. De Kroon, J. Mooij, and D. Belgrave, “Causal bandits without prior knowledge using separating sets,” in *Conference on Causal Learning and Reasoning*. PMLR, 2022, pp. 407–427.
- [24] Y. Lu, A. Meisami, and A. Tewari, “Causal bandits with unknown graph structure,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 24 817–24 828, 2021.
- [25] A. Malek, V. Aglietti, and S. Chiappa, “Additive causal bandits with unknown graph,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 23 574–23 589.
- [26] Z. Yan, A. Mukherjee, B. Varıcı, and A. Tajer, “Robust causal bandits for linear models,” *IEEE Journal on Selected Areas in Information Theory*, 2024.
- [27] B. Bilodeau, L. Wang, and D. Roy, “Adaptively exploiting d-separators with causal bandits,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 20 381–20 392, 2022.
- [28] C. Tekin and E. Turğay, “Multi-objective contextual multi-armed bandit with a dominant objective,” *IEEE Transactions on Signal Processing*, vol. 66, no. 14, pp. 3799–3813, 2018.
- [29] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [30] R. Alami, “Non-stationary causal bandits,” in *NeurIPS 2022 Workshop on Causality for Real-world Impact*, 2022.
- [31] J. Peters, D. Janzing, and B. Schölkopf, *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- [32] J. M. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf, “Distinguishing cause from effect using observational data: methods and benchmarks,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1103–1204, 2016.
- [33] A. Ghassami, N. Kiyavash, B. Huang, and K. Zhang, “Multi-domain causal structure learning in linear systems,” *Advances in neural information processing systems*, vol. 31, 2018.
- [34] H. Wu and M. D. Wang, “An information theoretic learning for causal direction identification,” in *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2020, pp. 287–294.
- [35] D. Janzing and B. Schölkopf, “Causal inference using the algorithmic markov condition,” *IEEE Transactions on Information Theory*, vol. 56, no. 10, pp. 5168–5194, 2010.
- [36] L. Sun and A. G. Nikolaev, “Mutual information based matching for causal inference with observational data,” *Journal of Machine Learning Research*, vol. 17, no. 199, pp. 1–31, 2016.
- [37] T. M. Cover, *Elements of information theory*. John Wiley & Sons, 1999.
- [38] M. Vollmer, I. Rutter, and K. Böhm, “On complexity and efficiency of mutual information estimation on static and dynamic data,” in *21st International Conference on Extending Database Technology, EDBT 2018*. OpenProceedings.org, 2018, pp. 49–60.
- [39] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*. MIT press, 2022.
- [40] J. A. Tropp, “User-friendly tail bounds for sums of random matrices,” *Foundations of computational mathematics*, vol. 12, pp. 389–434, 2012.
- [41] N. J. Higham, *Accuracy and stability of numerical algorithms*. SIAM, 2002.
- [42] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio, “Toward causal representation learning,” *Proceedings of the IEEE*, vol. 109, no. 5, pp. 612–634, 2021.
- [43] T. L. Lai and H. Xing, “Sequential change-point detection when the pre- and post-change parameters are unknown,” *Sequential analysis*, vol. 29, no. 2, pp. 162–175, 2010.
- [44] C. R. Rao, C. R. Rao, M. Statistiker, C. R. Rao, and C. R. Rao, *Linear statistical inference and its applications*. Wiley New York, 1973, vol. 2.
- [45] A. Kraskov, H. Stögbauer, and P. Grassberger, “Estimating mutual information,” *Physical review E*, vol. 69, no. 6, p. 066138, 2004.
- [46] T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 2020.
- [47] V. Paulsen, *Completely bounded maps and operator algebras*. Cambridge University Press, 2002, no. 78.