
Scalable Causal Domain Adaptation

Mohammad Ali Javidian

School of Electrical and Computer Engineering
Purdue University
mjavidia@purdue.edu

Om Pandey

School of Computer Engineering
Kalinga Institute of Industrial Technology
opandey108@gmail.com

Pooyan Jamshidi

Computer Science & Engineering Department
University of South Carolina
pjamshid@cse.sc.edu

Abstract

One of the most important problems in transfer learning is the task of domain adaptation, where the goal is to apply an algorithm trained in one or more source domains to a different (but related) target domain. This paper deals with domain adaptation in the presence of *covariate shift* while there exist invariances across domains. One of the main limitations of existing causal inference methods for solving this problem is *scalability*. To overcome this difficulty, we propose SCTL, an algorithm that avoids an exhaustive search and identifies invariant causal features across source and target domains based on Markov blanket discovery. SCTL does not require to have prior knowledge of the causal structure, the type of interventions, or the intervention targets. There is an intrinsic locality associated with SCTL that makes SCTL practically scalable and robust, because local causal discovery increases the power of computational independence tests and makes the task of domain adaptation computationally tractable. We show the scalability and robustness of SCTL for domain adaptation using synthetic and real data sets in low-dimensional and high-dimensional settings.

1 Introduction

Standard supervised learning usually assume that both training and test data are drawn from the same distribution. However, this is a strong assumption, and often violated in practice if (1) training examples have been obtained through a biased method (i.e., *sample selection bias*), or (2) there exist a significant physical or temporal difference between training and test data sources (i.e., *non-stationary environments*) [19]. Domain adaptation approaches aim to learn domain invariant features to mitigate the problem of data shift and enhance the quality of predictions [26].

Many of domain adaptation techniques in the literature consider the *covariate shift* [5, 43, 16], where the marginal distribution of the features differs across the source and target domains, while the conditional distribution of the target given the features does not change. In causal inference, it has been noted that covariate shift corresponds to causal learning i.e., predicting effects from causes [38]. Taking into account the causal structure of a system of interest and finding causally invariant features in both source and target domains, enable us to safely transfer the predictions of the target variable based only on causally invariant features to the target domain [20, 36, 45, 46]. Following the same line of work, in this paper we formalize and study the problem of domain adaptation as a feature selection problem where our aim is to find an *optimal subset* that the conditional distribution of the target variable given this subset of predictors is invariant across domains under certain assumptions, as formally discussed in section 3. To illustrate the importance and effectiveness of this approach consider the following example.

35th Conference on Neural Information Processing Systems (NeurIPS 2021), Sydney, Australia.

Example 1 (Domain Adaptation: Prediction of Diabetes at Early Stages). *According to Diabetes Australia, early diagnosis and initiation of appropriate therapeutic treatment plays a pivotal role in: (1) helping patients to manage the disease early and (2) reducing huge economic impact of diabetes on the healthcare systems and national economies (700 million dollars each year [3]). To predict diabetes using machine learning techniques, we need its symptoms along with clinical data. The common symptoms and possible causes of diabetes Type II are weakness, obesity, delayed healing, visual blurring, partial paresis, muscle stiffness, alopecia, among others [1]. Although we do not know what causes diabetes Type II [3], one may argue that the occurrence, rate, or frequency of Delayed Healing, Blurred Vision, Partial Paresis, and Weakness increases with age, and Delayed Healing and Partial Paresis may cause some complications that result in pancreas malfunction.*




Figure 1: Prediction of diabetes at early stage: Here, age intervention leads to shift of distributions between source (Young Patients < 50) and target (Old Patients ≥ 50) domains (see also Example 1). In all cases, the orange curves indicate the distribution of tested positive patients and the blue curves indicate the distribution of tested negative patients. A standard feature selection method that does not consider the causal structure and would use all features to predict *Diabetes* would obtain biased predictions in the target domain (MSE = 0.29). Using only *Delayed Healing* and *Partial Paresis* instead yields less accurate predictions in the source domain, but much more accurate ones in the target domain (MSE = 0.221).

The causal graph of this scenario can be represented as the directed acyclic graph (DAG) in Figure 1 (a). To provide an instance of a domain adaptation problem, we divided a diabetes dataset [12] into two sub-populations: (1) source domain with patients of the age less than 50 (dubbed young) and (2) target domain with patients of the age greater than or equals to 50 (old patients). As shown in Figure 1 (b) and (c), intervention in age leads to shift distributions across domains. In our experiments, feature selection methods that do not consider the causal structure select highly relevant features to diabetes i.e., all four variables *Delayed Healing*, *Blurred Vision*, *Partial Paresis*, and *Weakness* to achieve high prediction accuracy in the source domain. This per se causes worse predictions in the target domain (MSE = 0.29 and SSE = 65.322) than the case that we only consider causally invariant features i.e., *Delayed Healing* and *Partial Paresis* for predictions in the target domain (MSE = 0.221 and SSE = 49.8927). The reason is that conditioning on the variables *Blurred Vision* and *Weakness* makes the paths between age and diabetes open, and hence the set of all features does not generalize to the target domain. However, conditioning only on *Delayed Healing* and *Partial Paresis* blocks the paths between age and diabetes, and hence these causally invariant features enable us to predict diabetes with higher accuracy in the target domain even in the presence of distribution shifts due to the age intervention.

An important limitation of existing causal inference methods [20, 36, 45, 46] is that they, currently, do not scale beyond dozens of variables due to either exponentially large number of conditional independence (CI) tests or difficulties in causal structure recovery [17]. To overcome the problem of scalability, we propose an algorithm based on Markov blanket discovery (see Appendix A for the definition) that in contrast to existing methods: (1) it takes advantage of local computation by finding only the Markov blanket of the target variable(s), and hence (2) reduces the search space for finding causally invariant features drastically because of small size of Markov blankets (≤ 10) in (many) causal models in practice e.g., see [40]. As a result, the CI tests becomes more reliable and robust when the number of variables increases, an important characteristic in high dimensional and low sample size scenarios. Our main contributions are as follows:

- We propose a new algorithm, called **Scalable Causal Transfer Learning (SCTL)**, to solve the problem of domain adaptation in the presence of covariate shift and scales to high-dimensional problems (section 3).
- For the first time we characterize Markov blankets in **Acyclic Directed Mixed Graphs (ADMGs)**

i.e., causal graphs in the presence of unmeasured confounders¹, and we prove that the standard Markov blanket discovery algorithms such as Grow-Shrink (GSMB), IAMB and its variants are still correct under the faithfulness assumption where causal sufficiency is not assumed. We prove the correctness of SCTL based on these new theoretical results (section 3).

- We demonstrate on synthesized and real-world data that our proposed algorithm improves performance over several state-of-the-art algorithms in the presence of covariate shift (section 5).

2 Related Work

Domain Adaptation. Here, we provide a brief overview of the main domain adaptation scenarios that can be found in the literature. There are three main domain adaptation problems: (1) **Covariate shift**, which is one of the most studied forms of data shift, occurs if the marginal distributions of *context variables* change across source and target domains while the posterior (conditional) distributions are the same between source and target domains [41, 47, 13]. (2) **Target shift** occurs if the marginal distributions of the *target variable* changes across source and target domains while the posterior distributions remain the same [44]. (3) **Concept shift** occurs if marginal distributions between source and target domains remain unchanged while the posteriors change across the domains [23].

We only focus on the *covariate shift* in this paper. Since assuming invariance of conditionals makes sense if the conditionals represent causal mechanisms [36], we use the relaxed version of the usual covariate shift assumption and assume that it holds true for a subset of predictor variables, as suggested and used in [36, 20].

Causal Inference in Domain Adaptation. Here, we briefly provide an overview of causal inference methods that address the problem of covariate shift: (1) **Transportability** [28, 7] expresses knowledge about differences and commonalities between the source and target domains in a formalism called *selection diagram*. Using this representation along with the *do-calculus* [27], enable us to derive a procedure for deciding whether effects in the target domain can be inferred from experiments conducted in the source domain(s). (2) **Invariant Causal Prediction** (ICP) techniques [31, 32, 33] search for finding a subset of variables to estimate individual regressions for each domain in a way that produce regression residuals with equal distribution across all domains. Theoretical identifiability guarantees for the set of direct causal predictors for ICP techniques is limited to the linear Gaussian models. In other words, when the model is non-linear, non-Gaussian, or there are latent confounders, there is no guarantee that the obtained set is the set of direct causal predictors [31, 33]. (3) **Graph surgery** [45, 46] remove variables generated by unstable mechanisms from the joint factorization to yield a distribution that is *invariant* to the differences across domains. (4) **Graph pruning** methods [20, 36] formalized as a feature selection problem in which the goal is to find the *optimal subset* that the conditional distribution of the target variable given this subset of predictors is invariant across domains. Both transportability and graph surgery methods need to know the causal model of interest in advance. However, learning causal models from data is a challenging task, especially in the presence of *unmeasured confounders* [11]. On the other hand, graph pruning methods do not rely on prior knowledge of the causal graph but they, currently, *do not scale beyond dozens of variables* [17].

3 Theory

In this section, first, we prove that the domain adaptation task under the **Causal Domain Adaptation** (CDA) assumptions can be done effectively via searching inside the set of *Markov blanket* of the target variable T rather than a brute-force search over all variables as it has been done in **Exhaustive Subset Search** (ESS) in [20], the closest work to SCTL. In the next section, we present an *efficient* and *scalable* algorithm, called the SCTL, that searches over *Markov blanket* of the target variable T to find causally invariant features that provide an accurate prediction for the target variable T given the distribution shift in terms of context variables. To apply SCTL in practice, we need an algorithm for Markov blanket discovery in the presence of unmeasured confounders. For this purpose, we prove that GSMB and IAMB and its variants are still sound under the faithfulness assumption, even when the causal sufficiency assumption does not hold. The proof of theorems can be found in Appendix B.

¹In this paper we will assume that there is no selection bias.

3.1 Problem: Causal Domain Adaptation

Here we formally state the causal domain adaptation task that we address in this work:

Task 1 (Domain Adaptation Task). *We are given data for a source and a target domain such that the marginal distribution of the context variable C changes across domains. Assume the source and target domains data are complete (i.e., no missing values), except for all values of a specific target variable T . The task is to predict these missing values of the target variable T given the available source and target domains data.*

For predicting T from a subset of features $S \subseteq V \setminus \{C, T\}$, where T is the target variable, $C = \{C_{s \in D_s}, C_{t \in D_t}\}$ is the context variable with the values $C_{s \in D_s}$ in the source domain(s) and the values $C_{t \in D_t}$ in the target domain (note that we can extend our formalism to multiple context variables), and S is a set that separates T from C in the causal model, we define the *transfer bias* as $\hat{T}_S^t - \hat{T}_S^s$, where $\hat{T}_S^t := \mathbb{E}(T|S, C = C_{t \in D_t})$ and $\hat{T}_S^s := \mathbb{E}(T|S, C = C_{s \in D_s})$. We define the *incomplete information bias* as $\hat{T}_{V \setminus \{C, T\}}^t - \hat{T}_S^t$. The *total bias* when using \hat{T}_S^s to predict T is the sum of the transfer bias and the incomplete information bias:

$$\underbrace{\hat{T}_{V \setminus \{C, T\}}^t - \hat{T}_S^s}_{\text{total bias}} = \underbrace{\hat{T}_S^t - \hat{T}_S^s}_{\text{transfer bias}} + \underbrace{\hat{T}_{V \setminus \{C, T\}}^t - \hat{T}_S^t}_{\text{incomplete information bias}}$$

For more details see [20]. Note that using invariant features only guarantees transfer bias to be zero but the incomplete information bias can be quite large for the invariant feature set. There is a trade-off between the two terms in the total bias expression which is hard to determine because the two terms are not identifiable using the source data alone.

3.2 Assumptions

We consider the same assumptions as discussed in [20], the closest work to ours.

Assumption 1 (Joint Causal Inference (JCI) assumptions). *We consider that causal graph $G = (V, E)$ is an ADMG with the variable set V . From now on, we will distinguish system variables $X_{j \in J}$ describing the system of interest, and context variables $C_{i \in I}$ describing the context in which the system has been observed:*

- (a) *Context variables are never caused by system variables: $(\forall j \in J, i \in I : X_j \rightarrow C_i \notin E)$,*
- (b) *System variables are not confounded by context variables: $(\forall j \in J, i \in I : X_j \leftrightarrow C_i \notin E)$,* (c) *All pairs of context variables are confounded: $(\forall i, k \in I : C_i \leftrightarrow C_k \in E, \text{ and } \forall i, k \in I : C_i \rightarrow C_k \notin E)$.*

The case (a) in the assumption is called *exogeneity* and captures what we mean by ‘‘context’’. Cases (b), (c) are not as important as the exogeneity and can be relaxed, depending on the application [20].

Assumption 2 (Causal Domain Adaptation (CDA) assumptions). *We consider causal graph G that satisfies Assumption 1. We say G satisfies CDA assumptions if*

- (a) *The probability distribution of V and the causal graph $G = (V, E)$ satisfy Markov condition and faithfulness assumption,*
- (b) *For the target variable T and a set A that $T \perp\!\!\!\perp A|S$ (i.e., T is independent of A given S) in the source domain we have the same CI in the target domain, where $\{T\} \not\subseteq A \cup S$,*
- (c) *No context variable C is the parent of the target variable T , i.e., $C \rightarrow T \notin E_G$.*

Assumption 2(b) states that the pooled source and target domains distributions are Markov and faithful to the local causal structure around the target variable of interest T . This assumption implies that the causal structure of the target variable T is invariant when going from the source to the target domain. Note that Assumption 2(b) of the CDA assumption in [20] is different from ours: both the pooled source the pooled target domains distribution are *Markov* and *faithful* to the subgraph of G which excludes the context variable C . This strong assumption forces them to do **exhaustive search**. Assumption 2(c) is strong and restrictive, which might not hold in real data sets. However, since interventions are local in their nature², if the interventions targeted precisely then Assumption 2(c) is more likely to be satisfied.

²This is called *modularity assumption* in the literature [24]. It means that intervening on a variable X_i only changes the causal mechanism for X_i , i.e., $p(X_i|\mathbf{pa}(X_i))$ and it does not change the causal mechanisms that generate any other variables.

3.3 Theoretical Results

The following theorem enables us to localize the task of finding invariant features under the CDA assumptions.

Theorem 1. *Assume that CDA assumptions hold for the context $C_{i \in I}$ and system variables $X_{j \in J}$ given data for a single or multiple source domains. To find the best separating set(s) of features that m -separate(s) $C_{i \in I}$ from the target variable T in the causal graph G , it is enough to restrict our search to the set of Markov blanket, $\mathbf{Mb}(T)$, of the target variable T .*

Theorem 1 enables us to develop a new *efficient* and *scalable* algorithm, called **Scalable Causal Transfer Learning (SCTL)**, that exploits locality for learning *invariant causal features* to be used for the task of domain adaptation. Since CDA assumptions do not require *causal sufficiency assumption*, we need sound and scalable algorithms for Markov blanket discovery in the presence of unmeasured confounders. For this purpose, we first need to provide a graphical characterization of Markov blankets in ADMGs.

Let $G = (V, E, P)$ be an ADMG model. Then, V is a set of random variables, (V, E) is an ADMG, and P is a joint probability distribution over V . Let $T \in V$, then the *Markov blanket* $\mathbf{Mb}(T)$ is the set of all variables that there is a *collider path* between them and T . We now show that the Markov blanket of the target variable T in an ADMG probabilistically shields T from the rest of the variables. Under the faithfulness assumption, the Markov blanket is the *smallest set* with this property. Formally we have:

Theorem 2. *Let $G = (V, E, P)$ be an ADMG model. Then, $T \perp\!\!\!\perp_p V \setminus \{T, \mathbf{Mb}(T)\} | \mathbf{Mb}(T)$.*

Our next theorem safely enables us to use standard Markov blanket recovery algorithms for domain adaptation task without causal sufficiency assumption:

Theorem 3. *Given the Markov assumption and the faithfulness assumption, a causal system represented by an ADMG, and i.i.d. sampling, in the large sample limit, the Markov blanket recovery algorithms GSMB [21], IAMB [48], Fast-IAMB [52], Interleaved Incremental Association (IIAMB) [48], and IAMB-FDR [30] correctly identify all Markov blankets for each variable. (Note that Causal Sufficiency is not assumed.)*

4 SCTL: Scalable Causal Transfer Learning

SCTL addresses the task of domain adaptation by finding a separating set $S \subset V$, where V is the set of context variables $C_{i \in I}$ and system variables $X_{j \in J}$ such that for the target variable T we have $C_i \perp\!\!\!\perp T | S$, for every $i \in I$ in the source domain. Since Assumption 2(b) implies that this conditional independence holds across domains, if such a separating set S can be found, S is considered as a set of causally *invariant* features for T across environments. Our SCTL algorithm, described in Algorithm 1, consists of two main steps:

Step 1. We find the Markov blanket of the target variable T , i.e., $\mathbf{Mb}(T)$ (line 3 in Algorithm 1). Using the property of Markov blankets, i.e., $T \perp\!\!\!\perp V \setminus \{T, \mathbf{Mb}(T)\} | \mathbf{Mb}(T)$, if there is no context variable in the Markov blanket of the target variable T , then $T \perp\!\!\!\perp C_{i \in I} | \mathbf{Mb}(T)$. This means that the Markov blanket of T , provides a minimal feature set required for prediction of the target variable T with maximum predictivity [2].

Step 2. If there exist a context variable C in $\mathbf{Mb}(T)$, we consider all possible subsets of the Markov blanket of the target variable T , $\mathbf{Mb}(T)$, to find those subsets S_i s that satisfy the separating condition $T \perp\!\!\!\perp C_{i \in I} | S_i$. For this purpose, the algorithm filters out those subsets for which p_{value} is below the significance level α , i.e., T is *not* conditionally independent of the context variables given those sets. At the end of line 14, we have a list S that contains all possible separating sets. If S is empty, we abstain from making prediction as the domain adaptation task is not successful. This may happen if there exist a context variable C in the adjacency of the target variable T , which means the data set does not satisfy the CDA assumptions because it violates the Assumption 1(a), (b), or Assumption 2(c). In this case, the algorithm throws a failure because there is no subset of features that is causally invariant across domains. Otherwise, we sort the list S based on the obtained prediction errors in the source domain (line 18). Then, the algorithm returns those subsets sub_i in S with the lowest prediction error in the source domain as the best possible separating sets, because according to Theorem 1, these sets m -separates T from the context variable(s) and minimize(s) the transfer bias. Computational complexity of SCTL is discussed in the supplementary material.

Input: A Dataset with variable set V is the set of context variables $C_{i \in I}$ and system variables $X_{j \in J}$, target variable T , and significance level α .

Output: A list of separating sets or null.

```

1 // Step 1: Find  $\mathbf{Mb}(T)$ , Markov blanket of  $T$ .
2 Set  $S$  as an empty list;
3 Find  $\mathbf{Mb}(T)$ ;
4 if  $\mathbf{Mb}(T) \cap C_{i, i \in I} = \emptyset$  then
5     return  $S = \mathbf{Mb}(T)$ ;
6 else
7     // Step 2: Find a subset of  $\mathbf{Mb}(T)$  that makes  $T$   $m$ -separated from the context variable(s).
8     Set Subs =  $\mathbf{Subset}(\mathbf{Mb}(T))$ ; /* All possible combinations of subsets for the obtained set  $\mathbf{Mb}(T)$  */
9     for each  $sub_i \in Subs$  do
10          $p_{val} = p_{value}(C_{i \in I} \perp\!\!\!\perp T | sub_i)$ ;
11         if  $p_{val} > \alpha$  then
12             // This means  $T$  is conditionally independent of the context variables given the set  $sub_i$ .
13             Add  $sub_i$  to  $S$ ;
14         end
15     end
16     if  $S = \emptyset$  then
17         return null;
18     end
19     Sort  $S$ ; /* Sort the list  $S$  in descending order based on prediction errors in the source domain(s) */
20     return a list of elements of  $S$  with the lowest prediction error in the source domain(s);
21 end

```

Algorithm 1: SCTL: A Scalable Transfer Learning Algorithm for Causal Domain Adaptation

Remark 1. Assume the causal graph in Figure 2 with the binary context variables C_1, C_2 , and the target variable T . Also, assume that the domain adaptation task is to predict the missing values of T in the target domains ($C_1 = 1$), based on the observed data from the source domain ($C_1 = 0$) and the target domain, without knowledge of the causal graph. In this case the true Markov blanket of the target variable T is the set of $\mathbf{Mb}(T) = \{X, Y, C_1\}$. However, since the value of the context variable C_1 is fixed in the source domain, Markov blanket discovery algorithms cannot discover the true Markov blanket in such cases. If we have access to the part of the data with no missing values for T in the target domain, which is a realistic scenario in some situations, e.g., clinical data from a target hospital, the data from the source and the target domains can be used for learning the true $\mathbf{Mb}(T)$. Otherwise, the problem is reduced to Markov blanket discovery in the presence of missing data. For this purpose, we can use the proposed method in [10].

Remark 2. It has already been suggested in [32, Theorem 3.5] that the optimal set of causally invariant features is a subset of the Markov blanket of the target variable, called stable blanket. Here, we highlight the differences between [32] and SCTL:

(1) In [32], the underlying causal structure for the system variables is considered as a DAG, which means there is no latent confounder among system variables (causal sufficiency assumption), while in our work, this restrictive assumption is relaxed. Note that causal sufficiency assumption is a strong assumption and often violated in practice. Moreover, the paper provides no theoretical guarantee that the stable blanket is a subset of the Markov blanket of the target variable T , $\mathbf{Mb}(T)$, in the presence of latent confounders. In fact, the invariance principle used in [31] shows that the stable set returned by the ICP techniques is a subset of ancestors of T . However, in our work, Theorem 1 relaxes causal sufficiency assumption and states that searching in $\mathbf{Mb}(T)$ is enough to find causally invariant features even in the presence of latent confounders.

(2) In [32], the joint probability distribution of variables is assumed to be continuous and experimental evaluations are restricted to the linear Gaussian models, while in our work, there is no restriction on the type of data and the probability distribution of variables. In fact, in section C we do a comprehensive evaluation that is not restricted to the linear Gaussian models.

5 Experimental Results

Experimental results, over various settings as discussed in Section C, have been shown in Figure 4,5,6,7,8,9, and 10, the remaining results are in Appendix E.

5.1 Synthetic Dataset

SCTL outperforms (in some cases a comparable performance) other approaches in all environment settings, considered on the synthetic dataset. In Figure 4-10, we report the most complicated scenarios involving severe domain shifts, extreme sample sizes and multiple ground truth models, results for




Figure 2: Ground truth synthetic graph: The grey node U represents the unknown confounder connecting the context variables filled in green C_1 & C_2 . The node in orange, T , represent the target variable.




Figure 3: Ground truth synthetic graph: The grey node U represents the unknown confounder connecting the context variables filled in green C_1 & C_2 . The nodes in orange like E , T , represent possible target variables. The central structure represented by solid lines is common to all graphs. The dashed lines and nodes are used to introduce variability in the structure across domains.

other scenarios can be found in the supplementary material. SCTL is, overall, as good as or even better than the state-of-the-art feature selection and domain adaptation algorithms listed in section C.

Key highlights. (1) *Scalability*: ESS, the closest work to SCTL, employs CI test based feature selection by conducting exhaustive search over the entire feature set, which increases the time required for subset generation exponentially, so cannot scale on datasets beyond 10 variables. We used ESS for computation on a 12 variable synthetic dataset for 72 hours, it crashed without generating any results. SCTL drastically reduces time required for subset search as it searches on only those variables that are in the locality of the target variable. Note that the target could be a subset of variables rather than a single variable, using parallel computation can speed up Markov blanket recovery because Markov blankets of different nodes can be learned independently [39]. This makes SCTL scalable to high-dimensional data and have potential applications in big data (as we will show in a real dataset with 400k variables in Section 5.3). (2) *Robustness to Conditional Independence Tests*: In cases where CI tests have to be estimated from data, mistakes occur in keeping or removing members from the estimated separating sets. The main source of erroneous CI tests are large condition sets in high-dimensional low sample size scenarios [6]. In such cases, the resulting changes in the CI tests can lead to different separating sets. However, SCTL that is based on Markov blanket discovery only uses a small fraction of variables in the vicinity of the target, increasing the reliance on CI tests, making the algorithm robust in practice. Our experimentation confirms that for most environment settings, where we know the ground truth causal graph, SCTL is able to find the correct invariant features (e.g., see Table 4 and 5 in the appendix).

(3) *Quantitative analysis of results on Gaussian settings*: Although we report only the most interesting scenarios in Figures 4-10, we notice that, for Gaussian settings (see Figure 9 (b)), the error rates of many of these feature selection algorithms slightly lower than SCTL. Considering total bias, as discussed in section 3, these approaches enjoy higher predictivity due to lower incomplete information bias. For example, in Figure 9 (b), Adaboost has a slightly lower MSE than SCTL for the specified setting. For all such scenarios, we perform t-tests and F-tests to confirm our observation from the results. In these tests, we found that for all such cases the difference in error rates is almost insignificant (Figures 24 - 32). (4) *Increase in Error on small sample size*: We do see a slight increase in the error on small sample size setting for SCTL (see Figure 10 (b)). Further investigation reveals that, since, we do not utilize the underlying ground truth graph structure, the lack of data is straining the Markov blanket algorithms, causing them to sometimes incorrectly learn the Markov blanket of the target variable. For example, for the setting in Figure 10 (a) as the ground truth, we found that the Markov blanket of the target variable T using IAMB is $\{P, X, C_1, G\}$. In this case the Markov blanket is wrong, but the separating set learned $\{P, X\}$ was satisfactory and causally invariant.

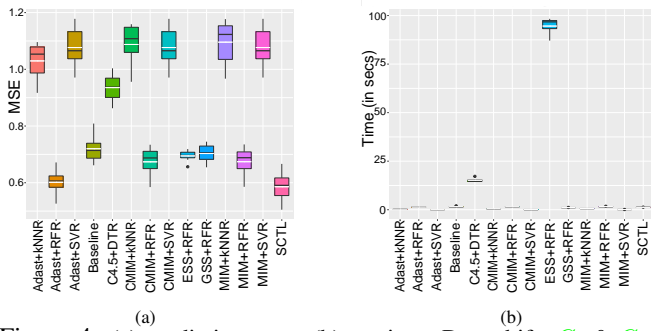


Figure 4: (a) prediction error, (b) runtime. Data shift: C_1 & C_2 , target variable: T , sample size = 1000, data type: discrete, Ground truth: Figure 3 with solid lines.




Figure 5: (a) the ground truth graph, (b) prediction error. Data shift: C_1 , target variable: T , sample size = 1000, and data type: discrete.

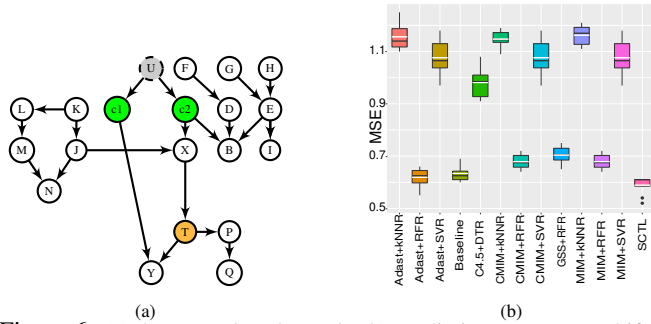


Figure 6: (a) the ground truth graph, (b) prediction error. Data shift: C_1 & C_2 , target variable: T , sample size = 1000, and data type: discrete.

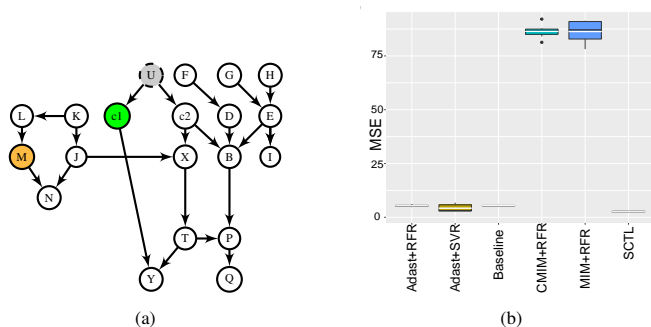


Figure 7: (a) the ground truth graph, (b) prediction error. Data shift: C_1 , target variable: M , sample size = 1000, and data type: discrete.




Figure 8: (a) the ground truth graph, (b) prediction error. Data shift: C_1 & C_2 , target variable: T , sample size = 1000, and data type: discrete.

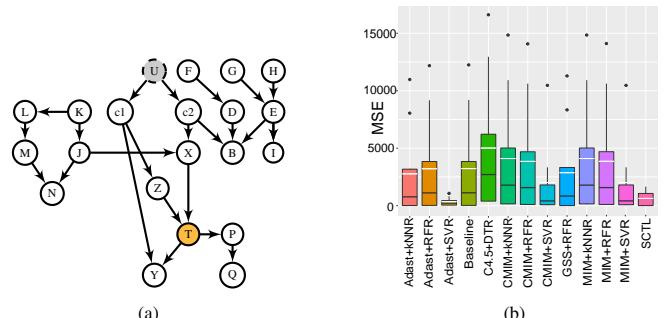


Figure 9: (a) the ground truth graph, (b) prediction error. No data shift, target variable: T , sample size = 1000, and data type: Gaussian.

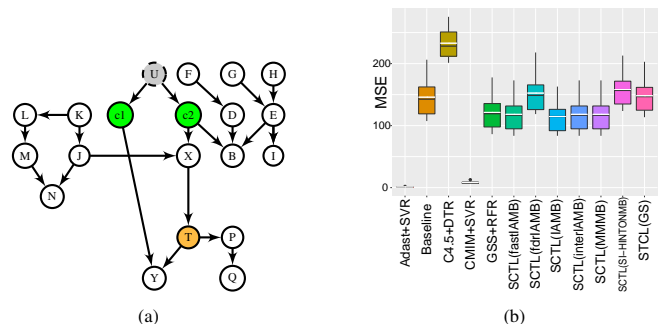


Figure 10: (a) the ground truth graph, (b) prediction error. Data shift: C_1 & C_2 , target variable: T , sample size = 50, and data type: Gaussian.




Figure 11: Error rate comparison of feature selection approaches on the Diabetes dataset with Age Shift (discussed in Section C.1.2).

(5) *Choice of Markov Blanket Algorithms*: As shown in Figure 10 (b), the Markov blanket algorithm being used also holds key significance and their practical uses may show results different from the theoretical perception. For example, GS does not consider the ordering and the strength of the association of candidate variable and the target variable T . On the other hand, IAMB orders variables based on the strength of their association with the target variable T first and then check their membership in the $\mathbf{Mb}(T)$. Since choosing different p-values, sample size of the data, and the maximum size of the conditioning sets have a big effect of the quality of learned $\mathbf{Mb}(T)$, it is necessary to choose an appropriate Markov blanket approach depending on given setting. For example, fdrIAMB [30] is particularly suitable when dealing with more nodes than samples in Gaussian models.

5.2 Real-world data: Diabetes Dataset

Although we do not have the ground truth causal graph for the real world diabetes data [12], we assume that it follows Assumption 1 & 2. We consider *Age* as context variable, because previous studies [15] have shown that older people are at higher risk for the development of Type II diabetes. As, this is a classification problem, the commonly used metrics Accuracy and F1-score were used for measuring performance of the algorithms. We compare SCTL against Adaboost and CMIM as they showed the best performance in the synthetic case. Similar to synthetic data, we ran ESS on the computer for 72 hours after which it was crashed. The results in Figure 11 show that although SCTL uses fewer number of features for prediction, it provides a higher accuracy and F1-score in average compared with other approaches. We observe that there is a greater variability for Adaboost+SVC and CMIM+SVC accuracy and F1-score compared to SCTL as well as larger outliers. This can be considered as an indication for the robustness of SCTL in real-world scenarios. On the other hand, although there is a greater variability for SCTL accuracy and F1-score compared to the baseline as well as larger outliers, SCTL enjoys higher accuracy and F1-score in general. This observations verify the importance of finding causally invariant features where data shift occurs. Notice that the total sample size for both training and testing scenarios had 200 entries for each. This signifies that even for a moderate number of samples, the Markov blankets can learn the vicinity of the target well enough for giving low error rates in the target domain. We also compared SCTL against Transfer Component Analysis (TCA), which is a state-of-the-art algorithm in unsupervised domain adaptation that identifies transfer components that remains invariant across domains. As it was shown in [25], the performance of TCA is sensitive to the choice of kernel choice and hyper-parameters of the kernel function, we therefore performed hyper-parameter tuning (details in Appendix) and reported the best result. We hypothesize that TCA cannot learn the transfer components correctly in this scenario and requires further investigations.

5.3 Real-world data: Cancer Dataset

To show the scalability of SCTL we used the Colorectal Cancer Dataset. Similar to the diabetes case, we did not have the ground truth graph. With the same reasons we compared SCTL against Adaboost and CMIM using Accuracy and F1-score. For this experiment we considered *Gender* as the context variable. Firstly, this is a dataset with 400k variables, making it impossible to scale under any circumstances using [20]. We conducted a sensitivity analysis to evaluate the robustness of algorithms in this high-dimensional and small sample size setting. For this purpose, we extracted 1000, 10k, and 100k features using an Extra-Trees classifier [29] to rank features based on their importance. In addition to the sensitivity analysis, by this intervention to the dataset, we made the setting in favor of feature selection algorithms that we compare by not only making the problem less dimensional but also sorting the features based on their importance using the complete dataset (details in

Table 1: Results for the cancer dataset (100k ranked features)

Methodology	Accuracy	F1-Score
Baseline	91.25	90.94
SCTL (GS) (167 features)	88.12	87.88
SCTL (IAMB) (80 features)	96.25	96.22
CMIM + SVC	90.05	89.79
Adaboost	96.87	97.83

Appendix E). The results for each of the datasets is shown in Table 1, 2, and 3. We found SCTL to be either as good as or better than other approaches except in compare to Adaboost that enjoys slightly better performance than SCTL in (almost) all settings. Our conjecture is that this is due to a misalignment between theory and dynamics of the domain. It also is well-known that Adaboost is a successful classifier that takes advantage of boosting [51]. Although non-causal feature selection methods such as Adaboost are often used in practice, they cannot be interpreted causally even when they achieve excellent predictivity. Note that the size of the returned feature set for SCTL(IAMB) in Table 1 is considerably smaller than for SCTL(GS), the performance of SCTL(IAMB) is better than SCTL(GS). This is an indication that IAMB is a better choice than GS in high-dimensional settings, as discussed in key highlight (5) for synthetic experiments.

Conclusion

In this paper, we proposed a new algorithm, called Scalable Causal Transfer Learning (SCTL), that identifies causal invariances in the presence of covariate shift and showed that it scales to high-dimensional and robust in low sample size setting in both synthetic and real-world scenarios. Weakening Assumption 2(b) and relaxing Assumption 2(c) are interesting directions for future work.

Acknowledgements

This work has been supported in part by NASA (Awards 80NSSC20K1720 and 521418-SC) and NSF (Awards 2007202 and 2107463). We are grateful to all who provided feedback on this work, including anonymous reviewers of NeurIPS WHY-21.

References

- [1] American Diabetes Association (ADA). Statistics about diabetes. <https://www.diabetes.org/>, 2020.
- [2] Constantin F. Aliferis, Alexander Statnikov, Ioannis Tsamardinos, Subramani Mani, and Xenofon D. Koutsoukos. Local causal and markov blanket induction for causal discovery and feature selection for classification part i: Algorithms and empirical evaluation. *JMLR*, pages 171–234, 2010.
- [3] Diabetes Australia. Diabetes in Australia. <https://www.diabetesaustralia.com.au>, 2020.
- [4] Karsten M Borgwardt, Arthur Gretton, Malte Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):49–57, 2006.
- [5] Xiangli Chen, Mathew Monfort, Anqi Liu, and Brian D. Ziebart. Robust covariate shift regression. volume 51 of *Proceedings of Machine Learning Research*, pages 1270–1279, Cadiz, Spain, 09–11 May 2016. PMLR.
- [6] Jie Cheng, David A Bell, and Weiru Liu. Learning belief networks from data: An information theory based approach. In *Proceedings of the 6th CIKM*, pages 325–331, 1997.
- [7] Juan D. Correa and Elias Bareinboim. From statistical transportability to estimating the effect of stochastic interventions. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 1661–1667. International Joint Conferences on Artificial Intelligence Organization, 7 2019.
- [8] François Fleuret. Fast binary feature selection with conditional mutual information. *Journal of Machine learning research*, 5(Nov):1531–1555, 2004.

- [9] Yoav Freund and Robert Schapire. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.
- [10] Nir Friedman. Learning belief networks in the presence of missing values and hidden variables. In *Proceedings of the Fourteenth International Conference on Machine Learning, ICML '97*, page 125–133. Morgan Kaufmann Publishers Inc., 1997.
- [11] Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10:524, 2019.
- [12] MM Faniqul Islam, Rahatara Ferdousi, Sadikur Rahman, and Humayra Yasmin Bushra. Likelihood prediction of diabetes at early stage using data mining techniques. In *Computer Vision and Machine Intelligence in Medical Image Analysis*, pages 113–125. 2020.
- [13] Fredrik D. Johansson, David Sontag, and Rajesh Ranganath. Support and invertibility in domain-invariant representations. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 527–536, 2019.
- [14] Markus Kalisch, Martin Mächler, Diego Colombo, Marloes Maathuis, and Peter Bühlmann. Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, pages 1–26, 2012.
- [15] M Sue Kirkman, Vanessa Jones Briscoe, Nathaniel Clark, Hermes Florez, Linda Haas, Jeffrey Halter, Elbert Huang, Mary Korytkowski, Medha Munshi, Peggy Soule Odegard, et al. Diabetes in older adults. *Diabetes care*, 35(12):2650–2664, 2012.
- [16] Keiichi Kisamori, Motonobu Kanagawa, and Keisuke Yamazaki. Simulator calibration under covariate shift with kernels. volume 108 of *Proceedings of Machine Learning Research*, pages 1244–1253, 2020.
- [17] Wouter Marco Kouw and Marco Loog. A review of domain adaptation without target labels. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [18] David D. Lewis. Feature selection and feature extraction for text categorization. In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*, 1992.
- [19] Yitong Li, Michael Murias, Samantha Major, Geraldine Dawson, and David Carlson. On target shift in adversarial domain adaptation. *Proceedings of Machine Learning Research*, pages 616–625, 2019.
- [20] Sara Magliacane, Thijs van Ommen, Tom Claassen, Stephan Bongers, Philip Versteeg, and Joris M. Mooij. Domain adaptation by using causal inference to predict invariant conditional distributions. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, page 10869–10879, 2018.
- [21] Dimitris Margaritis. *Learning Bayesian Network Model Structure from Data*. PhD thesis, Carnegie-Mellon University, 2003.
- [22] Dimitris Margaritis and Sebastian Thrun. Bayesian network induction via local neighborhoods. In *Proceedings of the NIPS'99*, pages 505–511, 1999.
- [23] Jose G Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V Chawla, and Francisco Herrera. A unifying view on dataset shift in classification. *Pattern recognition*, 45(1):521–530, 2012.
- [24] Brady Neal. Introduction to causal inference from a machine learning perspective. *Course Lecture Notes (draft)*, 2020.

- [25] Sinno J. Pan, Ivor Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, pages 199–210, 2010.
- [26] Sangdon Park, Osbert Bastani, James Weimer, and Insup Lee. Calibrated prediction with covariate shift via unsupervised domain adaptation. volume 108 of *Proceedings of Machine Learning Research*, pages 3219–3229, Online, 26–28 Aug 2020. PMLR.
- [27] J. Pearl. *Causality. Models, reasoning, and inference*. Cambridge University Press, 2009.
- [28] Judea Pearl and Elias Bareinboim. Transportability of causal and statistical relations: A formal approach. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pages 247–254, 2011.
- [29] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [30] Jose M. Peña. Learning Gaussian graphical models of gene networks with false discovery rate control. In *Evolutionary Computation, Machine Learning and Data Mining in Bioinformatics*, pages 165–176, 2008.
- [31] Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 947–1012, 2016.
- [32] Niklas Pfister, Peter Bühlmann, and Jonas Peters. Invariant causal prediction for sequential data. *Journal of the American Statistical Association*, 114(527):1264–1276, 2019.
- [33] Niklas Pfister, Evan G Williams, Jonas Peters, Ruedi Aebersold, and Peter Bühlmann. Stabilizing variable selection and regression. *arXiv preprint arXiv:1911.01850*, 2019.
- [34] J. Ross Quinlan. Induction of decision trees. *Machine learning*, 1(1):81–106, 1986.
- [35] Thomas Richardson. Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1):145–157, 2003.
- [36] Mateo Rojas-Carulla, Bernhard Schölkopf, Richard Turner, and Jonas Peters. Invariant models for causal transfer learning. *The Journal of Machine Learning Research*, 19(1):1309–1342, 2018.
- [37] Kayvan Sadeghi. Faithfulness of probability distributions and graphs. *The Journal of Machine Learning Research*, 18(1):5429–5457, 2017.
- [38] Bernhard Schoelkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris Mooij. On causal and anticausal learning. In John Langford and Joelle Pineau, editors, *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, ICML ’12, pages 1255–1262, 2012.
- [39] Marco Scutari. Bayesian network constraint-based structure learning algorithms: Parallel and optimized implementations in the bnlearn R package. *Journal of Statistical Software, Articles*, 77(2):1–20, 2017.
- [40] Marco Scutari. bnlearn - an R package for Bayesian network learning and inference. <https://www.bnlearn.com/bnrepository/>, 2021.
- [41] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227 – 244, 2000.
- [42] Ingo Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of machine learning research*, 2(Nov):67–93, 2001.

- [43] Petar Stojanov, Mingming Gong, Jaime Carbonell, and Kun Zhang. Low-dimensional density ratio estimation for covariate shift correction. volume 89 of *Proceedings of Machine Learning Research*, pages 3449–3458. PMLR, 16–18 Apr 2019.
- [44] Amos J Storkey. *When Training and Test Sets Are Different: Characterizing Learning Transfer*, pages 3–28. MIT Press, 2009.
- [45] Adarsh Subbaswamy and Suchi Saria. Counterfactual normalization: Proactively addressing dataset shift using causal mechanisms. In Amir Globerson and Ricardo Silva, editors, *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018*, pages 947–957, 2018.
- [46] Adarsh Subbaswamy, Peter Schulam, and Suchi Saria. Preventing failures due to dataset shift: Learning predictive models that transport. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3118–3127. PMLR, 2019.
- [47] Masashi Sugiyama, Taiji Suzuki, Shinichi Nakajima, Hisashi Kashima, Paul von Bünau, and Motoaki Kawanabe. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008.
- [48] Ioannis Tsamardinos, Constantin Aliferis, Alexander Statnikov, and Er Statnikov. Algorithms for large scale Markov blanket discovery. In *In The 16th International FLAIRS Conference, St*, pages 376–380. AAAI Press, 2003.
- [49] Ting Wang, Sean K Maden, Georg E Luebeck, Christopher I Li, Polly A Newcomb, Cornelia M Ulrich, Ji-Hoon E Joo, Daniel D Buchanan, Roger L Milne, Melissa C Southey, et al. Dysfunctional epigenetic aging of the normal colon and colorectal cancer risk. *Clinical epigenetics*, 12(1):1–9, 2020.
- [50] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016.
- [51] Abraham J Wyner, Matthew Olson, Justin Bleich, and David Mease. Explaining the success of adaboost and random forests as interpolating classifiers. *JMLR*, 18(1):1558–1590, 2017.
- [52] S. Yaramakala and D. Margaritis. Speculative Markov blanket discovery for optimal feature selection. In *Proceedings of the ICDM’05*, 2005.

A Basic Definitions and Concepts

Assume that $G = (V, E)$ is a directed graph, where $V = \{v_1, v_2, v_3 \dots v_n\}$ is the set of nodes (variables), $n \geq 1$, and E is the set of directed or bidirected edges. We say v_i is a *parent* of v_j and v_j is a *child* of v_i if $v_i \rightarrow v_j$ is an edge in G . We denote the set of parents and children of a variable v by $\mathbf{pa}(v)$ and $\mathbf{ch}(v)$, respectively. Any bidirected edge $v_i \leftrightarrow v_j$ means that there exists a node $v_k \notin V$ as a hidden confounder such that $v_i \leftarrow v_k \rightarrow v_j$. If G is acyclic, then G is an *Acyclic Directed Mixed Graph* (ADMG). Formally, $\mathbf{ne}(T) = \mathbf{pa}(T) \cup \mathbf{ch}(T) \cup \{v \in V \mid v \leftrightarrow T \text{ is a bidirected edge in } G\}$ refers to the *neighbours* of T . We define *spouses* of v as $\mathbf{sp}(v) = \{u \in V \mid \exists w \in V \text{ s.t. } u \rightarrow w \leftarrow v \text{ in } G\}$. We define *Markov blanket* of node T as $\mathbf{Mb}(T) = \mathbf{pa}(T) \cup \mathbf{ch}(T) \cup \mathbf{sp}(T)$ when G is a directed acyclic graph (DAG) and the set of children, parents, and spouses of T , and vertices connected with T or children of T by a bidirected path (i.e., only with edges \leftrightarrow) and their respective parents is the Markov blanket of T when G is an ADMG.

A path of length n from x to y in an ADMG $G = (V, E)$ is a sequence $x = a_0, \dots, a_n = y$ of distinct vertices such that $(a_i, a_{i+1}) \in E$, for all $i = 1, \dots, n$. A vertex α is said to be an *ancestor* of a vertex β if either there is a directed path $\alpha \rightarrow \dots \rightarrow \beta$ from α to β , or $\alpha = \beta$. We apply this definition to sets: $\mathbf{an}(X) = \{\alpha \mid \alpha \text{ is an ancestor of } \beta \text{ for some } \beta \in X\}$.

Definition 1. A nonendpoint vertex ζ on a path is a *collider* on the path if the edges preceding and succeeding ζ on the path have an arrowhead at ζ , that is, $\rightarrow \zeta \leftarrow$, or $\leftrightarrow \zeta \leftrightarrow$, or $\leftrightarrow \zeta \leftarrow$, or $\rightarrow \zeta \leftrightarrow$. A nonendpoint vertex ζ on a path which is not a collider is a *noncollider* on the path. A path between vertices α and β in an ADMG G is said to be *m-connecting* given a set Z (possibly empty), with $\alpha, \beta \notin Z$, if:

- (i) every noncollider on the path is not in Z , and
- (ii) every collider on the path is in $\mathbf{an}_G(Z)$.

If there is no path *m-connecting* α and β given Z , then α and β are said to be *m-separated* given Z . Sets X and Y are *m-separated* given Z , if for every pair α, β , with $\alpha \in X$ and $\beta \in Y$, α and β are *m-separated* given Z (X, Y , and Z are disjoint sets; X, Y are nonempty). This criterion is referred to as a *global Markov property*. We denote the *independence model* resulting from applying the *m-separation* criterion to G , by $\mathfrak{S}_m(G)$. This is an extension of Pearl's *d-separation* criterion to mixed graphs in that in a DAG D , a path is *d-connecting* if and only if it is *m-connecting*.

Definition 2. Let G_A denote the induced subgraph of G on the vertex set A , formed by removing from G all vertices that are not in A , and all edges that do not have both endpoints in A . Two vertices x and y in an ADMG G are said to be *collider connected* if there is a path from x to y in G on which every non-endpoint vertex is a collider; such a path is called a *collider path*. (Note that a single edge trivially forms a collider path, so if x and y are adjacent in an ADMG then they are collider connected.) The augmented graph derived from G , denoted $(G)^a$, is an undirected graph with the same vertex set as G such that $c - d$ in $(G)^a \Leftrightarrow c$ and d are collider connected in G .

Definition 3. Disjoint sets $X, Y \neq \emptyset$, and Z (Z may be empty) are said to be *m*-separated* if X and Y are separated by Z in $(G_{\mathbf{an}(X \cup Y \cup Z)})^a$. Otherwise X and Y are said to be *m*-connected* given Z . The resulting independence model is denoted by $\mathfrak{S}_{m^*}(G)$.

Richardson in [35, Theorem 1] shows that for an ADMG G , $\mathfrak{S}_m(G) = \mathfrak{S}_{m^*}(G)$.

The *Markov condition* is said to hold for $G = (V, E)$ and a probability distribution $P(V)$ if $\langle G, P \rangle$ satisfies the following implication: $\forall X, Y \in V, \forall Z \subseteq V \setminus \{X, Y\} : (X \perp\!\!\!\perp_m Y \mid Z \implies X \perp\!\!\!\perp_p Y \mid Z)$. The *faithfulness condition* states that the only conditional independencies to hold are those specified by the Markov condition, formally: $\forall X, Y \in V, \forall Z \subseteq V \setminus \{X, Y\} : (X \not\perp\!\!\!\perp_m Y \mid Z \implies X \not\perp\!\!\!\perp_p Y \mid Z)$.

B Proofs of Theoretical Results

To prove Theorem 1, we need the following definitions and propositions.

Definition 4 (District and Induced Markov Blanket [35]). *Let G be an acyclic directed mixed graph. We first specify a total ordering (\prec) on the vertices of G , such that $x \prec y \Rightarrow y \notin \mathbf{an}(x)$; such an ordering is said to be consistent with G . Let $\text{pre}_{G,\prec} = \{v | v \prec x \text{ or } v = x\}$. The district of x in G is x and the set of vertices connected to x by a path on which every edge is of the form \leftrightarrow , denoted $\text{dis}_G(x)$. So, $\text{dis}_G(x) = \{v | v \leftrightarrow \dots \leftrightarrow x \in G \text{ or } v = x\}$. A set A is said to be ancestral if it is closed under the ancestor relation, i.e. if $\mathbf{an}(A) = A$. If A is an ancestral set in an ADMG G , and x is a vertex in A that has no children in A then we define the Markov blanket of a vertex x with respect to the induced subgraph on A , called induced Markov blanket, as following: $\text{imb}_A(x) = \text{pa}_{G_A}(\text{dis}_{G_A}(x)) \cup (\text{dis}_{G_A}(x) \setminus \{x\})$. It is not difficult to see that if $A = \mathbf{an}_G(T)$, then $\text{imb}_A(T) \subseteq \mathbf{Mb}(T)$ by the definition of $\mathbf{Mb}(T)$ in an ADMG.*

Proposition 1. *Given two nodes X and Y in an ADMG G and a set S of nodes not containing X and Y , there exists some subset of S which m -separates X and Y if and only if the set $S' = S \cap \mathbf{an}(X \cup Y)$ m -separates X and Y .*

Proof. (\Rightarrow) Proof by contradiction. Let $S' = S \cap \mathbf{an}(X \cup Y)$ and $\langle X, Y \not\perp_{S'} \rangle$, i.e., S' does not m -separates X and Y in G . Since $S' \subseteq \mathbf{an}(X \cup Y)$, it is obvious that $\mathbf{an}(X \cup Y \cup S') = \mathbf{an}(X \cup Y)$. So, X and Y are not separated by S' in $(G_{\mathbf{an}(X \cup Y)})^a$, hence there is an undirected path C between X and Y in $(G_{\mathbf{an}(X \cup Y)})^a$ that bypasses S' i.e., the path C is formed from nodes in $\mathbf{an}(X \cup Y)$ that are outside of S . Since $\mathbf{an}(X \cup Y) \subseteq \mathbf{an}(X \cup Y \cup S'')$, $\forall S'' \subseteq S$, then $(G_{\mathbf{an}(X \cup Y)})^a$ is a subgraph of $(G_{\mathbf{an}(X \cup Y \cup S'')})^a$. Then, the previously found path C is also a path in $(G_{\mathbf{an}(X \cup Y \cup S'')})^a$ that bypasses S'' , which means that X and Y are not separated by any $S'' \subseteq S$ in $(G_{\mathbf{an}(X \cup Y \cup S'')})^a$, which is a contradiction.


(\Leftarrow) It is obvious. □

Now, we prove Theorem 1.

Proof of Theorem 1. Let G be a causal graph with variable set V consisting of system variables $X_{j \in J}$ and context variables $C_{i \in I}$. Assume that $C \in C_{i \in I}$, T is the target variable and $\mathbf{Mb}(T)$ is the Markov blanket of T . We have two cases:

- $C \notin \mathbf{Mb}(T)$: In this case, $C \perp\!\!\!\perp T | \mathbf{Mb}(T)$ due to the definition of Markov blanket.
- $C \in \mathbf{Mb}(T)$. In this case we show that considering the CDA assumptions, every path between C and T is blocked by $\text{imb}_A(T)$, where $A = \mathbf{an}_G(T)$, or there is no subset of variables that m -separates C from T . Proposition 1 implies that in order to find an m -separating set between C and T it is enough to restrict our search to the set $\mathbf{an}_G(C \cup T) = \{C\} \cup \mathbf{an}_G(T)$. Assume $A = \mathbf{an}_G(T)$, we have the following sub-cases:
 - (a) $C \in A$ and $C \notin \text{imb}_A(T)$. Using the ordered local Markov property for G and Theorem 2 in [35] implies that $C \perp\!\!\!\perp_m T | \text{imb}_A(T)$. Note that $\text{imb}_A(T) \subseteq \mathbf{Mb}(T)$ by the definition of $\mathbf{Mb}(T)$.
 - (b) $C \notin A$ and $C \notin \text{imb}_A(T)$. Since $\text{imb}_A(T) \subseteq A$ and A is an ancestral set, then $\mathbf{an}_G(\text{imb}_A(T)) = A$ and $\mathbf{an}_G(C \cup T \cup \text{imb}_A(T)) = \{C\} \cup \mathbf{an}_G(T)$. So, if there is an undirected path between T and C in $(G_{\mathbf{an}_G(C \cup T \cup \text{imb}_A(T))})^a$, Lemma 4 in [35] implies that this path intersect with the set $\text{imb}_A(T)$. In other words, $C \perp\!\!\!\perp_m T | \text{imb}_A(T)$.
 - (c) $C \in \text{imb}_A(T)$. Using Definition 4 implies that $C \in A$. In this case, there is no subset of variables that m -separates C from T . Proof by contradiction: Assume that there is a subset of

variables Z that m -separates C from T in G . Proposition 1 implies that $S = Z \cap \mathbf{an}_G(T \cup C) = Z \cap (\{G\} \cup \mathbf{an}_G(T)) = Z \cap \mathbf{an}_G(T)$ must m -separate T from C in G . Since $C \in \text{imb}_A(T)$ and $\text{imb}_A(T) \subseteq \mathbf{an}_G(T)$, then it is not difficult to see that in $(G_{\mathbf{an}(T \cup C \cup S)})^a$, C is directly connected to the vertex T . This means that S does not m -separate T from C in G , which is a contradiction. The following simple example illustrates such situations:



Note that in all cases the Markov condition and faithful assumptions guarantee the correctness of independence relationships. As we shown in all cases under the CDA assumptions, to find a separating set of features that m -separates C from the target variable T in the causal graph G , it is enough to restrict our search to the set of Markov blanket $\mathbf{Mb}(T)$ of the target variable T . In the case that $C_{i,i \in I}$ has more than one element, similar argument can be used to prove the theorem.

Using any subset S for prediction that satisfies the m -separating set property, implies zero transfer bias. So, the best predictions are then obtained by selecting a separating subset that also minimizes the source domains risk (i.e., minimizes the incomplete information bias). \square

Correctness of Algorithm 1. The correctness of Algorithm 1 follows from Theorem 1. \square

Proof of Theorem 2. It is enough to show that for any $A \in V \setminus \{T, \mathbf{Mb}(T)\}$, $T \perp\!\!\!\perp_m A | \mathbf{Mb}(T)$. For this purpose, we prove that any path between A and T in G is blocked by $\mathbf{Mb}(T)$. In the following cases ($A \rightarrow B$, where means $A - B$ or $A \rightarrow B$ and $A ** B$ means $A - B$, $A \rightarrow B$, or $A \leftarrow B$) we have:

1. The path ρ between A and T is of the form $A ** \dots ** B \rightarrow T$. Clearly, $B \in \mathbf{Mb}(T)$ blocks the path ρ .
2. The path ρ between A and T is of the form $A ** \dots ** C \leftarrow B \leftarrow T$. Clearly, $B \in \mathbf{Mb}(T)$ blocks the path ρ .
3. The path ρ between A and T is of the form $A ** \dots ** C \rightarrow B \leftarrow T$. $C \in \mathbf{Mb}(T)$ blocks the path ρ .
4. The path ρ between A and T is of the form $A ** \dots ** C \rightarrow D \leftrightarrow \dots \leftrightarrow B \leftarrow T$, where $\omega = C \rightarrow D \leftrightarrow \dots \leftrightarrow B \leftarrow T$ is the largest collider path between T and a node on the path ρ . Since all nodes of ω are in the Markov blanket of T , $\forall X \in \omega, X \neq A$. So, $C \in \mathbf{Mb}(T)$ blocks the path ρ .
5. The path ρ between A and T is of the form $A ** \dots ** C \rightarrow D \leftrightarrow \dots \leftrightarrow B \leftrightarrow T$, where $\omega = C \rightarrow D \leftrightarrow \dots \leftrightarrow B \leftarrow T$ is the largest collider path between T and a node on the path ρ . Since all nodes of ω are in the Markov blanket of T , $\forall X \in \omega, X \neq A$. So, $C \in \mathbf{Mb}(T)$ blocks the path ρ .

From the global Markov property it follows that every m -separation relation in G implies conditional independence in every joint probability distribution P that satisfies the global Markov property for G . Thus, we have $T \perp\!\!\!\perp_p V \setminus \{T, \mathbf{Mb}(T)\} | \mathbf{Mb}(T)$. \square

Now, we are ready to prove Theorem 3.

Sketch of proof of Theorem 3. If a variable belongs to $\mathbf{Mb}(T)$, then it will be admitted in the first step (Growing phase) at some point, since it will be dependent on T given the candidate set of $\mathbf{Mb}(T)$. This holds because of the causal faithfulness and because the set $\mathbf{Mb}(T)$ is the minimal set with that property. If a variable X is not a member of $\mathbf{Mb}(T)$, then conditioned on $\mathbf{Mb}(T) \setminus \{X\}$, it will be independent of T and thus will be removed from the candidate set of $\mathbf{Mb}(T)$ in the second phase (Shrinking phase) because the causal Markov condition entails that independencies in the distribution are represented in the graph. Since the causal faithfulness condition entails dependencies in the distribution from the graph, we never remove any variable X from the candidate set of $\mathbf{Mb}(T)$ if $X \in \mathbf{Mb}(T)$. Using this argument inductively we will end up with the $\mathbf{Mb}(T)$. \square

In order to show the details of the proof of the Theorem 3, we prove only the correctness of the Grow-Shrink Markov blanket (GSMB) algorithm without causal sufficiency assumption in details as following (for the other algorithms listed in Theorem 3, a similar argument can be used):

Proof of Correctness of the GSMB Algorithm. By ‘‘correctness’’ we mean that GSMB is able to produce the true Markov blanket of any variable in the ground truth ADMG under Markov condition and the faithfulness assumption if all conditional independence tests done during its course are assumed to be correct.

Input: a set V of nodes, a target variable T , and a probability distribution p faithful to an unknown ADMG G .

Output: The Markov blanket of T i.e., $\mathbf{Mb}(T)$.

```

1 Set  $S = \emptyset$ ;
  /* Grow Phase: */
2 while  $\exists X \in V \setminus \{T\}$  such that  $X \not\perp_p T | S$  do
3    $S \leftarrow S \cup \{X\}$ ;
  /* Shrink Phase: */
4 while  $\exists X \in S$  such that  $X \perp\!\!\!\perp_p T | S \setminus \{X\}$  do
5    $S \leftarrow S \setminus \{X\}$ ;
6 return  $(\mathbf{Mb}(T) \leftarrow S)$ ;
```

Algorithm 2: The GS Markov Blanket Algorithm [21].

We first prove that there does not exist any variable $X \in \mathbf{Mb}(T)$ at the end of the growing phase that is not in S . The proof is by induction (a semi-induction approach on finite subset of natural numbers) on the *length* of the collider path(s) between X and T . We define the length of a collider path between X and T as the number of edges between them. Let s be the length of a largest collider path between X and T .

- (*Base case*) For the base of induction, consider the set of adjacents of T i.e., $\mathbf{adj}(T) = \{X \in V | X \rightarrow T, X \leftarrow T, \text{ or } X \leftrightarrow T\}$. In this case, $\nexists S \subseteq V \setminus \{T, X\}$ such that $X \perp\!\!\!\perp_m T | S$ in G . The faithfulness assumption implies that $\nexists S \subseteq V \setminus \{T, X\}$ such that $X \perp\!\!\!\perp_p T | S$. So, at the end of the grow phase, all of adjacents of T are in the candidate set for the Markov blanket of T .
- (*Induction hypothesis*) For all $1 \leq n < s$, if there is a collider path between X and T of length n then $X \in \mathbf{Mb}(T)$ at the end of the growing phase.

- (*Induction step*) To prove the inductive step, we assume the induction hypothesis for $n = s - 1$ and then use this assumption to prove that the statement holds for s . Assume that $\rho = (v_0 = T) \ast \rightarrow v_1 \leftrightarrow \dots \leftrightarrow v_{s-1} \leftarrow \ast v_s$ is a collider path of length s . From the induction hypothesis we know that $v_i \in \mathbf{Mb}(T), \forall 1 \leq i < s$ at the end of the grow phase. Using Definition 1 implies that ρ is m -connected to T . This means that at some point of the grow phase v_s falls into the set S . The faithfulness assumption implies that $\exists S \subseteq V \setminus \{T, v_s\}$ such that $v_s \not\perp_p T | S$. So, at the end of the grow phase, all of X 's that there is a collider path between X and T fall into the candidate set for the Markov blanket of T .

For the correctness of the shrinking phase, we have to prove two things: (1) we never remove any variable X from S if $X \in \mathbf{Mb}(T)$, and (2) if $X \notin \mathbf{Mb}(T)$, X is removed in the shrink phase.

Now we prove case (1) by contradiction. Assume that $X \in \mathbf{Mb}(T)$, $X \in S$ at the end of the grow phase, and we remove X from S in the shrink phase. This means $X \perp_p T | S \setminus \{X\}$. Using the faithful assumption implies that $X \perp_m T | S \setminus \{X\}$ i.e., $S \setminus \{X\}$ m -separates X from T in G , which is a contradiction because there is a collider path between X and T and $\mathbf{Mb}(T) \setminus \{X\} \subseteq S$. In other word, the collider path between X and T is m -connected by $S \setminus \{X\}$.

To prove the case (2), assume that $X \notin \mathbf{Mb}(T)$, $\mathbf{Y} = S \setminus \{\mathbf{Mb}(T)\}$, and $X \in \mathbf{Y}$. Due to the Markov blanket property, $\mathbf{Mb}(T)$ m -separates \mathbf{Y} from T in G . Using the Markov condition implies that $\mathbf{Y} \perp_p T | \mathbf{Mb}(T)$. Since the probability distribution p satisfies the faithfulness assumption, it satisfies the *weak union* condition [37]. We recall the weak union property here: $A \perp_p B D | C \Rightarrow (A \perp_p B | D C \text{ and } A \perp_p D | B C)$. Using the weak union property for $\mathbf{Y} \perp_p T | \mathbf{Mb}(T)$ implies that $X \perp_p T | (\mathbf{Mb}(T) \cup (\mathbf{Y} \setminus \{X\}))$, which is the same as $X \perp_p T | (S \setminus \{X\})$. This means, $X \notin \mathbf{Mb}(T)$ will be removed at the end of the shrink phase. \square

C Experimental Evaluation

We empirically validated the robustness and scalability of SCTL on both synthetic (10-20 variables) and real-world high-dimensional data (400k variables) by comparing with various feature selection approaches including Conditional Mutual Information Maximization (CMIM) [8], Mutual Information Maximization (MIM) [18], Adaboost (Adast) [9], Greedy Subset Search (GSS) [36], C4.5 using Decision Trees (C4.5 + DTR) [34] and Exhaustive Subset Search (ESS) [20].

C.1 Experimental Settings

In this section, we mention the different criteria used for data generation, parameter-tuning and validation. Based on changes in parameters and user dependent variables, we divide this section into three key parts. In each sub-section we describe the combinations of hyperparameters used, why we used them and how they were implemented in practice.

C.1.1 Synthetic Data Generation

Let us consider an ADMG G as shown in Figure 3. The basic synthetic dataset based on G consists of 19 randomly generated variables, with 16 system variables, 2 context variables (C_1 & C_2) and 1 unknown confounder (U) (to be removed while using data). The data was generated in Gaussian as well as Discrete distributions. In Figure 3 the dashed lines and circles annotate the supplementary additions to the central structure which is represented by solid lines. We generate a dataset (see Appendix D for details) with desired number of samples and faithful w.r.t G . This data (not the graph structure) is used further for our experiments. We generated multiple datasets using different configurations to simulate the behaviour across different domain changes. These configurations were carefully generated to exploit characteristics of real-world scenarios:

(1) **Changes to the sample size:** We generated distributions with high disparity in the “amount” of data. The generated datasets contained 50, 1000 and 10000 cases. This would help us to see the difference in accuracy and robustness of the feature selection algorithms under extreme data-sensitive conditions (i.e. low sample size, moderate size, big data).

(2) **Changes to the network size:** We change the size of the graph, by increasing or dropping the number of nodes to monitor change in performance, ensure scalability of approach to changes in the number of environment variable. For this, we consider 3 different network sizes 20, 12 and 8 nodes, as shown in Figure 3 (further instances in Appendix). We also experimented by making changes to and on the number of context variables. This was done by using variables that are either partly, or completely unaffected by them (e.g., in Figure 3, we can use the whole or parts of graph with nodes M , N , E as target).

(3) **Changes to the complexity:** The difference in complexity of the dataset can be further divided into *structural* change i.e. changing the edge connections or relations among the nodes (e.g., in Figure 3, we can add Z to induce affect of C_1 on the target T , or connect B and P , to affect the Markov blanket of T), and *domain-distribution* change i.e. changes between the source and target domain. Since we consider two context variables, domain-distribution change can be brought about by changes in either one or both. The changes to the domain can be classified further into three sub-settings: smooth (very little change to context variable), mild (small change in context variables) and severe (drastic change in context variables between source and target) by varying the mean and variance for Gaussian distributions and changing the probabilities associated with the variables in discrete distributions. To ensure consistency in our results, for each setting, we generate 8 test datasets, by making slight changes to any one of the system variables at random. The slight change is made so that the algorithms used for prediction do not report the same error each time, instead an error range serves a better depiction of practical robustness.

C.1.2 Real-world Dataset: Low-dimensional

We used the diabetes dataset [12] that reports common diabetic symptoms of 520 person and contains 17 features, with 320 cases of diabetes Type II positive and 200 of diabetes negative patients. In the dataset we consider two context variables Age and Gender. We perform 3 sets of data shifts between training and testing domains using these context variables: (1) *Gender Shift*: We train on the Male and test on Female. (2) *Age Shift*: We split the data using an arbitrary threshold (we choose 50 as threshold for our experimentation as it gives a fair distribution of samples between training and testing data), train on young (less than 50) and test on the old (greater than 50). (3) *Double context shift*: In this case, we train on young male patients and test on old female patients. We trained by sampling 100 samples from the source domain keeping the samples in target domain constant to generates conditions for multiple source domains and single target domain. Such conditions help further test the robustness of the feature selection algorithm.

C.1.3 Real-world Dataset: High-dimensional

To assess SCTL in a high-dimension, low-sample size real world setting, we experiment on the Colorectal Cancer Dataset [49]. It contains 334 samples assigned to three cancer risk groups (low, medium, and high) based on their personal adenoma or cancer history. The entire dataset contains $\sim 400k$ features. We perform 2 sets of possible data shifts between training and testing domains using these context variables: (1) *Gender Shift*: We train on the Male and test on Female. (2) *Age Shift*: We split the data using an arbitrary threshold, in our experiments, train on young (less than 50) and test on the older (greater than 50).

C.2 Evaluation Metrics

We tested SCTL and other algorithms using different regression techniques such as RFR, SVR, and kNNR, in order to mitigate doubts about approach specific biases. We found, for SCTL, the same error more or less in (almost) all cases, but we chose Random Forest Regressor as it not only performed well, but also, gave

consistent results. Here, the *Baseline* indicates using all available features and feeding it to the Random Forest Regressor. We reported accuracy with Mean-Square-Error (MSE) and time taken in second.

D Details of Experimental Evaluation

The data generation processes, conditional independence tests and Markov blanket Algorithms were implemented in *R* by extending the *bnlearn* [39] and *pclag* [14] packages. We used Python 3.6 with *scikit-learn* [29] library for implementing the above mentioned feature selection and machine learning algorithms to compare our approach against. The library is also used for prediction over subset of features selected by SCTL using *RandomForestRegressor* function. The generated predictions from each of the algorithms are compared to the actual results for finding the error and all experiments have been reported using various comparison metrics.

Synthetic Data Generation

For Gaussian setting, a model graph as shown in Figure 3 was generated using *model2network* function from the *graphviz* package. The obtained DAG from Figure 3 is passed to the *custom.fit* function from *bnlearn* which defines the mean, variance and relationship (edges values) between the nodes. This process ensures that the generated dataset will be faithful w.r.t G . Once the graph is set and all variables have been properly defined, we use the *rbn* function from *bnlearn* to generate required number of samples.

For the Discrete setting, the same process is followed, but instead of mean, variance and edge values, we use a matrix of random probabilities. For each node, the size of the matrix would be the number of discrete values for the node, plus the number of discrete values for each of its connecting nodes multiplied by the number of nodes connected to it.

Implementation Parameters

SCTL requires the use of Markov-blanket and neighbourhood for feature selection. We use multiple Markov blanket discovery-algorithms to evaluate the effect the algorithm-choice can have on final prediction. For this purpose, we select, GS, IAMB, interIAMB, fastIAMB, fdrIAMB, MMMB, SLHINTON.MB. These algorithms were implemented in *R* using the *bnlearn* package. The variable *subs* mentioned on line 12 of Algorithm 1 contains all possible combination of subsets from the feature set S_2 . To find all possible subsets of a given feature set we use Python 3.5, one or more of these subsets will act as the separating set. To find the separating set, we sort the subsets based on p_{value} of conditional independence tests and choosing the subset/s with highest score. We use a significance level α of 0.05 as threshold for the conditional independence tests. The tests for different configurations are based on the type of data being used. For Gaussian data, we used *gaussCItest* from the *pclag* package, which uses Fisher’s z-transformation of the partial correlation, for testing correlation for sets of normally distributed random variables. For discrete data, we use mutual information (*mi*) test (an information-theoretic distance measure, which is somewhat proportional to the log-likelihood ratio) from the *bnlearn* package.

Change to the complexity As an example, for a central structure consist of 8 nodes as shown in Figure 3 having 10000 samples with severe changes in context, we train by generating one dataset on the base setting. Then make severe changes to the domain-distribution (i.e. by drastically changing the values of the context variables) and then generate a new dataset which is used for testing.

Evaluation Metrics

We used the *RandomForestRegressor* function from *scikit-learn* package, using the out-of-bag (OOB) scoring approach on the default parameters, for this process.

E More Experimental Results

In this section we include the remaining results that we got during experimentation, along with subsequent ground truth graph for representing the various settings we experiment on synthetic data. In the later part of this section we also include t-test tables as mentioned in Section 5. More experimental results regarding synthetic data can be found at the end of this appendix.

E.1 Real-world Dataset: Low-dimensional

We used the diabetes dataset [12] that contains reports of common diabetic symptoms of 520 persons. This includes data about symptoms that may cause or are potentially caused by diabetes. The dataset has been created from a direct questionnaire to people who have either recently been diagnosed as diabetic, or who are still non-diabetic but having show few or more symptoms of diabetes. The diabetes dataset contains a total of 17 features, with 320 cases of diabetes Type II positive and 200 of diabetes negative patients. The data has been collected from the patients of the Sylhet Diabetes Hospital, Sylhet, Bangladesh. In the dataset we consider two context variables Age and Gender. We perform 3 sets of data shifts between training and testing domains using these context variables: (1) *Gender Shift*: We train on the Male and test on Female. (2) *Age Shift*: We split the data using an arbitrary threshold (we choose 50 as threshold for our experimentation as it gives a fair distribution of samples between training and testing data), train on young (less than 50) and test on the older (greater than 50). (3) *Double context shift*: We do a dataset shift for both context variables, in our case, we train on young male patients and test on old female patients. In practice, while experimenting on a context shift setting we train by sampling 100 samples from the source domain keeping the samples in target domain constant. This allows us to generates conditions similar to having multiple source domains and single target domain. Such conditions will help further test the robustness of the feature selection algorithm.

E.2 Real-world Dataset: High-dimensional

To assess SCTL in a high-dimension, low-sample size real world setting, we experiment on the Colorectal Cancer Dataset [49]. Chronological age is a prominent risk factor for many types of cancers including Colorectal Cancer on which this dataset is based. The dataset is basically a genome-wide DNA methylation study on samples of normal colon mucosa, it contains 334 samples assigned to three cancer risk groups (low, medium, and high) based on their personal adenoma or cancer history. Our target variable for classification was the ‘Diagnosis’ feature, a binary variable which indicates whether or not the patient was diagnosed with cancer, the entire dataset contains ~400k features. We again perform 2 sets of data shifts between training and testing domains using these context variables: (1) *Gender Shift*: We train on the Male and test on Female. (2) *Age Shift*: We split the data using an arbitrary threshold (we choose 50 as threshold for our experimentation as it gives a fair distribution of samples between training and testing data), train on young (less than 50) and test on the older (greater than 50).

Table 2: Cancer dataset (1000 ranked features) result

Methodology	Accuracy	F1-Score
Baseline	91.87	91.70
SCTL (GS) (167 features)	87.50	87.18
SCTL (IAMB) (13 features)	96.82	97.22
CMIM + SVC	87.5	87.06
Adaboost	96.87	97.83

Table 3: Cancer dataset (10k ranked features) result

Methodology	Accuracy	F1-Score
Baseline	91.25	90.94
SCTL (GS) (167 features)	88.125	87.88
SCTL (IAMB) (79 features)	97.5	97.65
CMIM + SVC	90.15	89.79
Adaboost	96.87	97.83

Table 4: Markov blanket and separating feature sets extracted by various algorithms on the target variable T for Figure 32(a).

Methodology	Markov blanket	Separating set
GS	["K", "L", "N", "Y"]	["K", "L", "N", "Y"]
IAMB	["K", "L", "N"]	["K", "L", "N"]
interIAMB	["K", "L", "N"]	["K", "L", "N"]
fastIAMB	["K", "L", "N", "Y"]	["K", "L", "N", "Y"]
MMMB	["K", "L", "N"]	["K", "L", "N"]
fdrIAMB	["K", "L", "N"]	["K", "L", "N"]

E.3 SCTL vs Unsupervised Domain Adaptation TCA

Transfer Component Analysis (TCA) is an unsupervised technique for domain adaptation proposed by Pan et. al. [25] that discovers common latent features, called *transfer components*, that have the same marginal distribution across the source and target domains while maintaining the essential structure of the source domain data [50]. TCA learns these transfer components across domains in a reproducing kernel Hilbert space [42] using maximum mean discrepancy [4]. Once the transfer components are found, a traditional machine learning technique is used to train the final target classifier. In short, TCA is a two-step domain adaptation technique where the first step reduces the marginal distributions between the source and target domains and the second step trains a classifier with the adapted domain data. Here, we compare SCTL against TCA, which is a state-of-the-art algorithm in unsupervised domain adaptation, and report the results. Note that, in TCA the focus is to identify substructures (i.e., transfer components) that remains invariant across domains, which is different from the problem of identifying invariant features used for training a predictive model, in which we do not need to have prior knowledge about the causal structure of the model that we are interested in. For implementation of the TCA we used, the *rbf* kernel, and the *logistic* classifier to predict on the extracted components. We also tune the hyperparameters of the algorithm by using Support Vector Classifier and increasing the number of components from 1 to 5, which lead to further decrease in the accuracy of the model. We hypothesize that TCA cannot learn the transfer components correctly in this setting. Further investigation by comparing with semi-supervised versions [25] as well as strategies that also incorporate sample reweighing are left for future work.

Computational complexity of SCTL. Assume that the "learning Markov blankets" phase (Step 1) uses the grow-shrink (GS) algorithm [22], and $n = |V|$, $m = |E|$, where $G = (V, E)$ is the unknown true causal structure. Since the Markov blanket algorithm involves $O(n)$ conditional independence (CI) tests, if the cardinality of the Markov blanket of the target variable is k (in many real world cases $k \ll n$, see for example real world datasets and their properties in [40]) then the total number of needed CI tests for finding the best separating set is $O(n + 2^k)$, see [21] for more details. On the other hand, the brute force algorithm used in (Magliacane et al., 2018) needs $O(2^n)$ tests in the worst case scenario, which is infeasible in practice as we have shown in our experiments. The same argument is valid when we have multiple target variables, say c , where $c \ll n$.




Figure 12: Error rate comparison of multiple feature selection approaches (supervised and unsupervised) on a simulation of the Diabetes dataset with *Age Shift* (discussed in Section C.1.2).

Table 5: Markov blanket and separating feature sets extracted by various algorithms on the target variable T for Figure 19(a).

Methodology	Markov blanket	Separating set
GS	["C1", "P", "X", "Y"]	["P", "X"]
IAMB	["C1", "P", "X", "Y"]	["P", "X"]
interIAMB	["C1", "P", "X", "Y"]	["P", "X"]
fastIAMB	["C1", "P", "X", "Y"]	["P", "X"]
MMMB	["C1", "P", "X", "Y"]	["P", "X"]
fdrIAMB	["C1", "P", "X", "Y"]	["P", "X"]
ESS	NA	["Q", "X", "Y"]

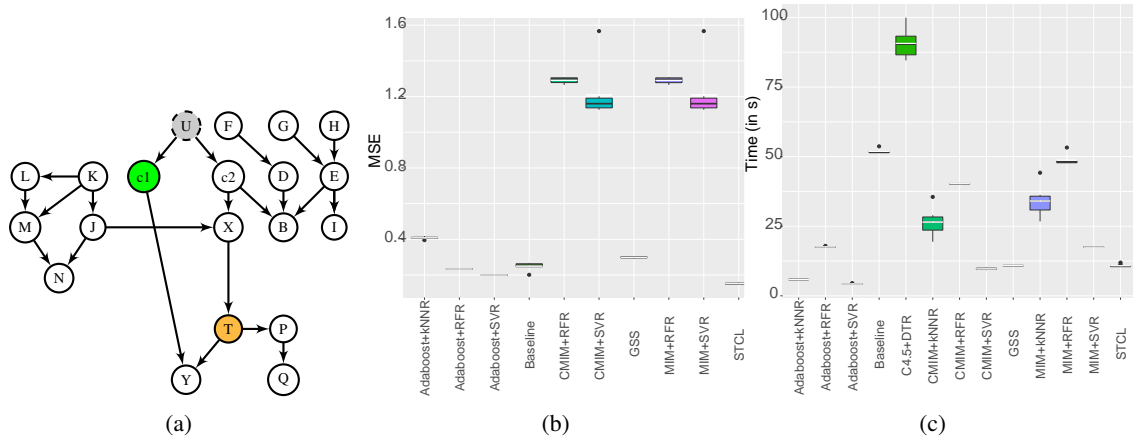


Figure 13: Results on (b), (c) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the probability distribution of the context variable C_1 , where sample size = 10000 for a Gaussian distribution.

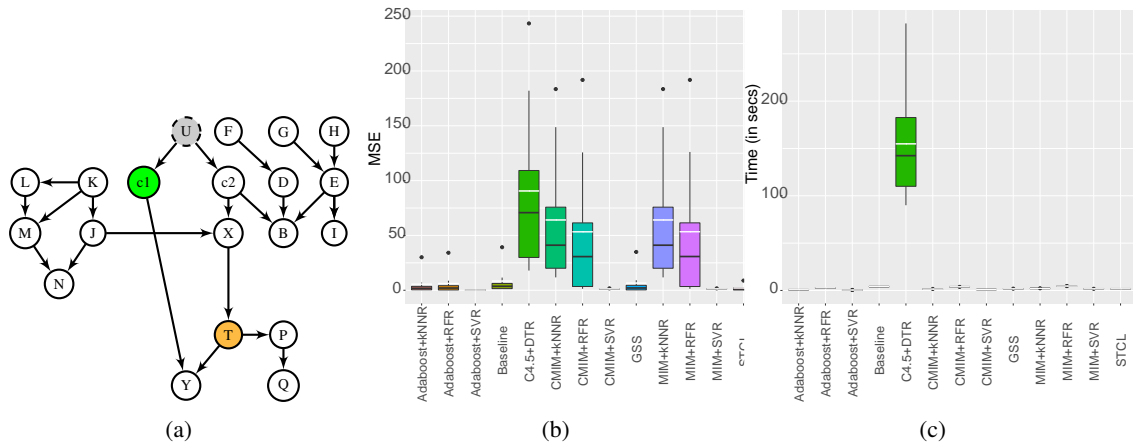


Figure 14: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the probability distribution of the context variable C_1 , where sample size = 1000 for a Gaussian distribution.

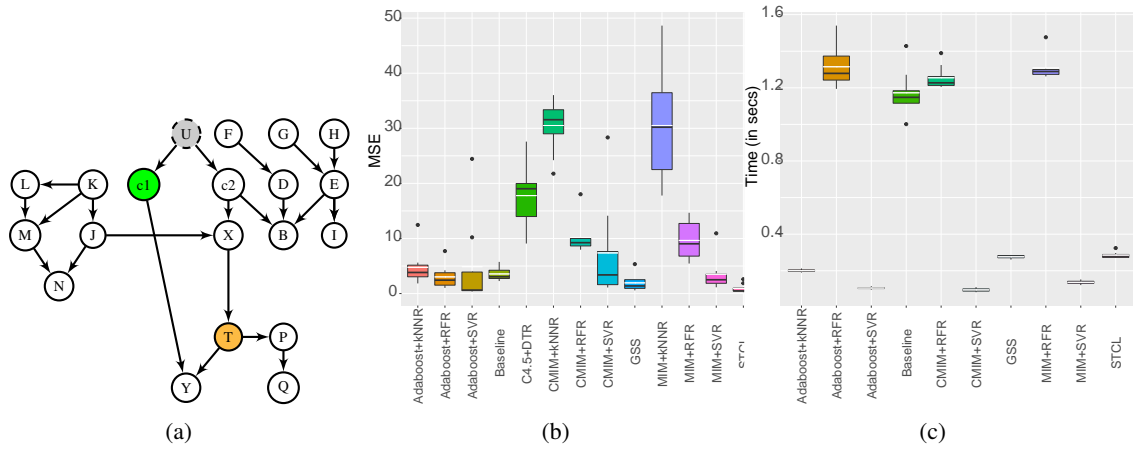


Figure 15: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the probability distribution of the context variable C_1 , where sample size = 50 for a Gaussian distribution.

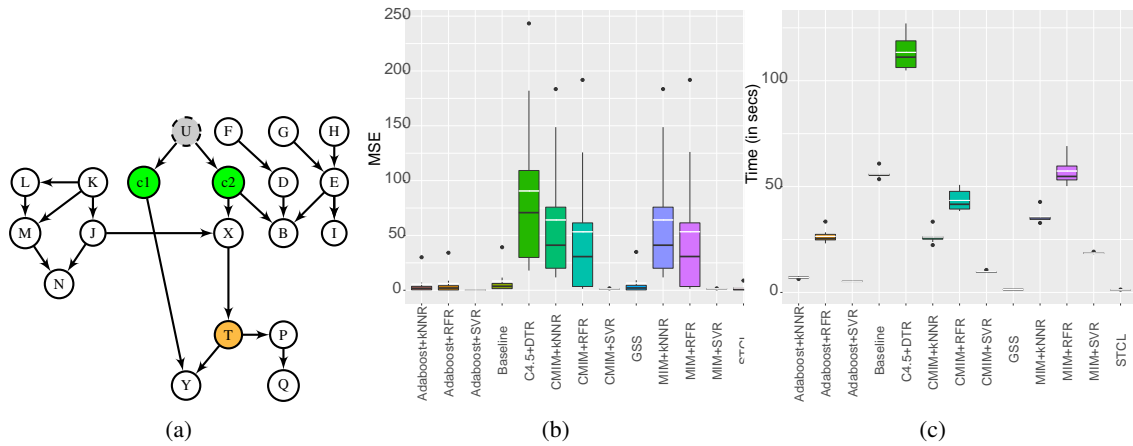


Figure 16: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the probability distribution of the context variable C_1 & C_2 , where sample size = 10000 for a Gaussian distribution.

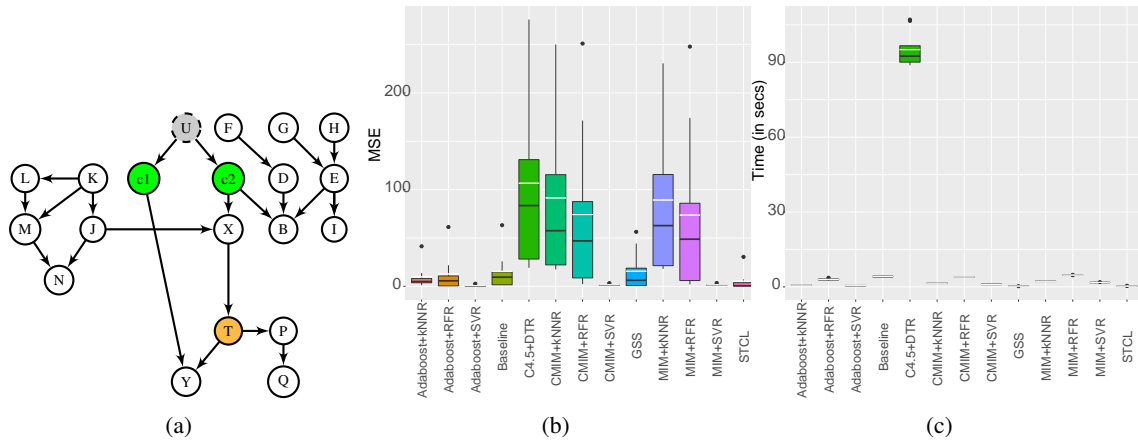


Figure 17: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the probability distribution of the context variable C_1 & C_2 , where sample size = 1000 for a Gaussian distribution.

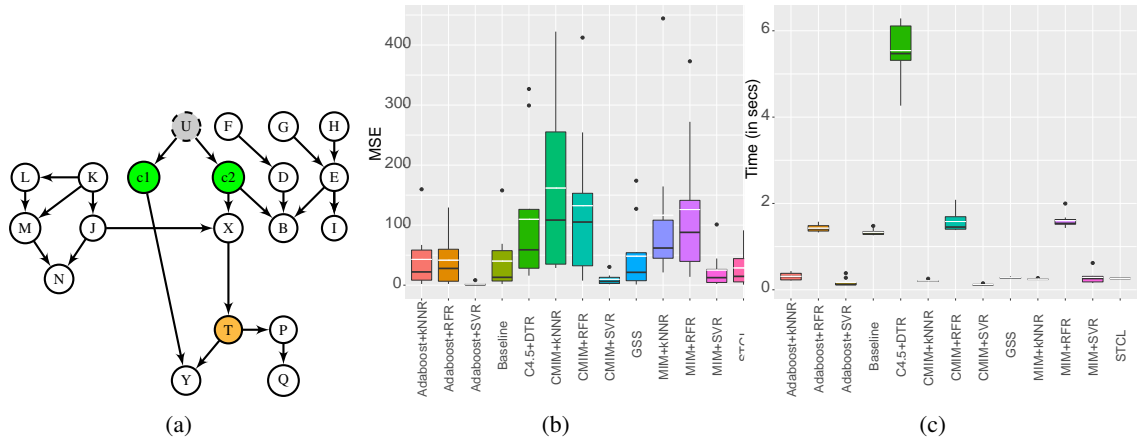


Figure 18: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 & C_2 , where sample size = 50 for a Gaussian distribution.

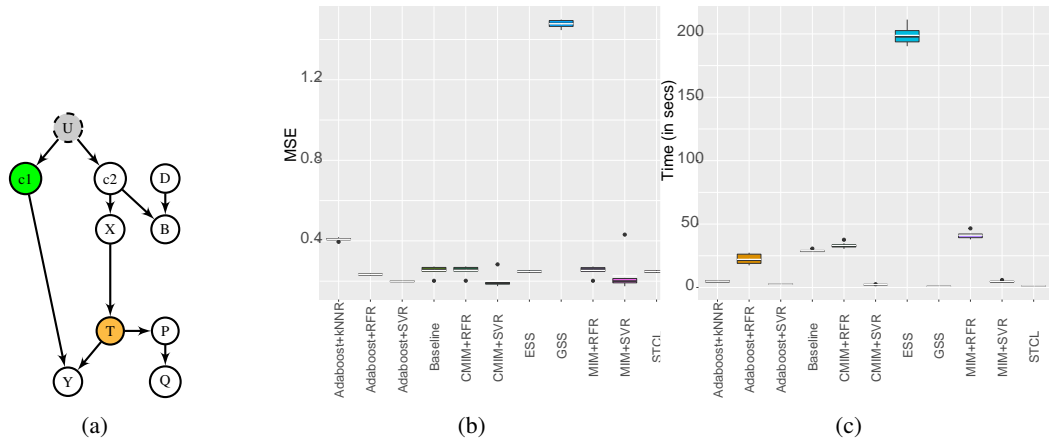


Figure 19: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 , where sample size = 10000 for a Gaussian distribution.

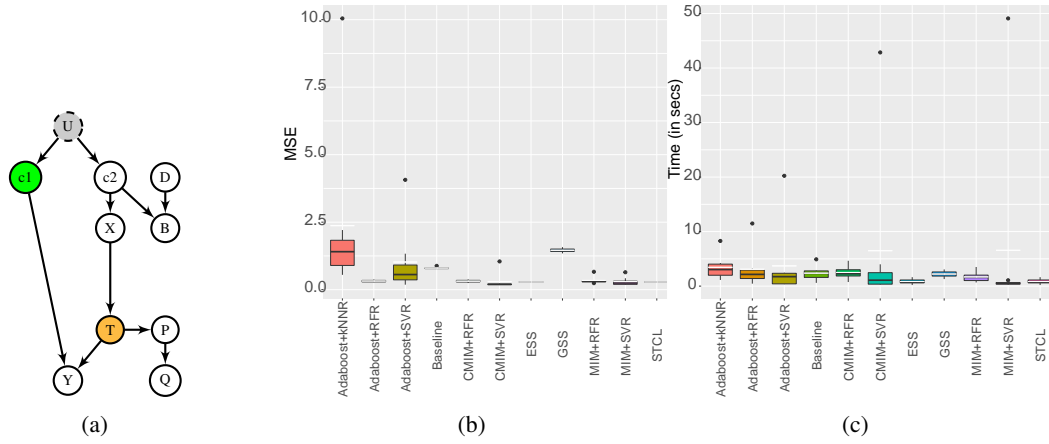


Figure 20: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 , where sample size = 1000 for a Gaussian distribution.

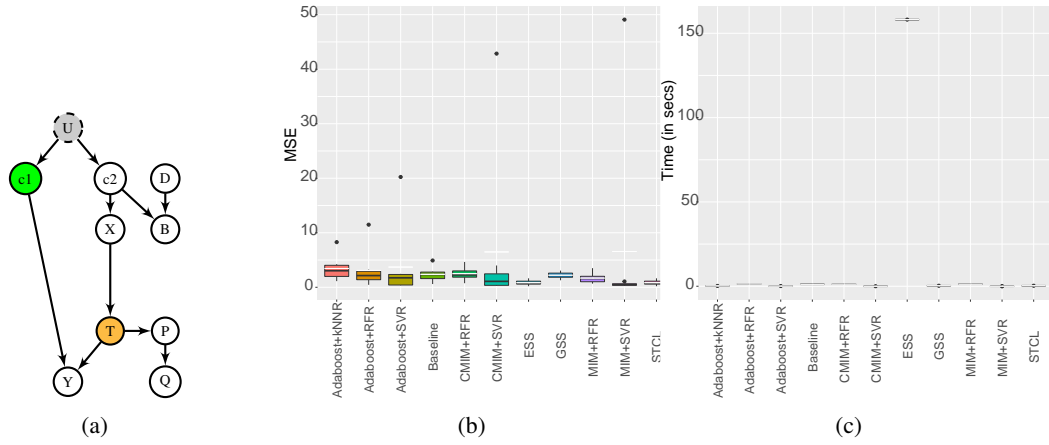


Figure 21: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 , where sample size = 50 for a Gaussian distribution.

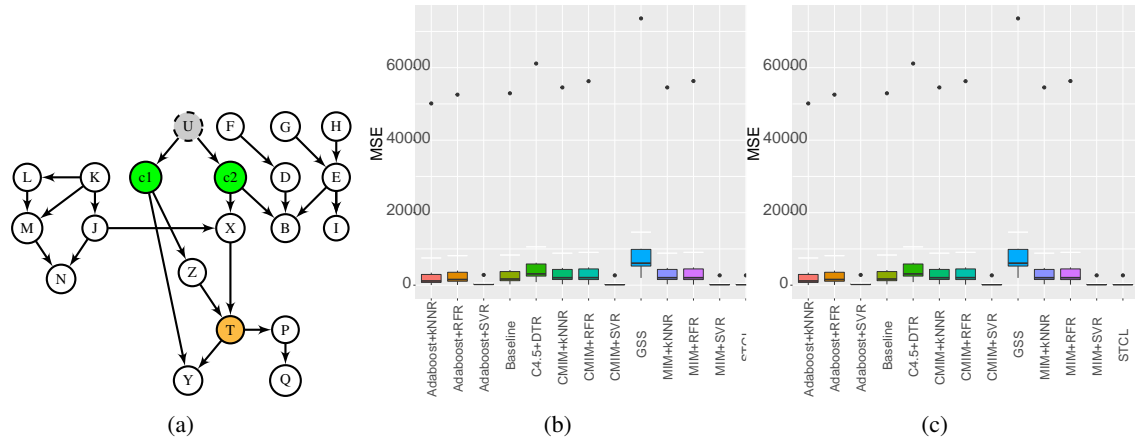


Figure 22: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 , where sample size = 1000 for a Gaussian distribution.

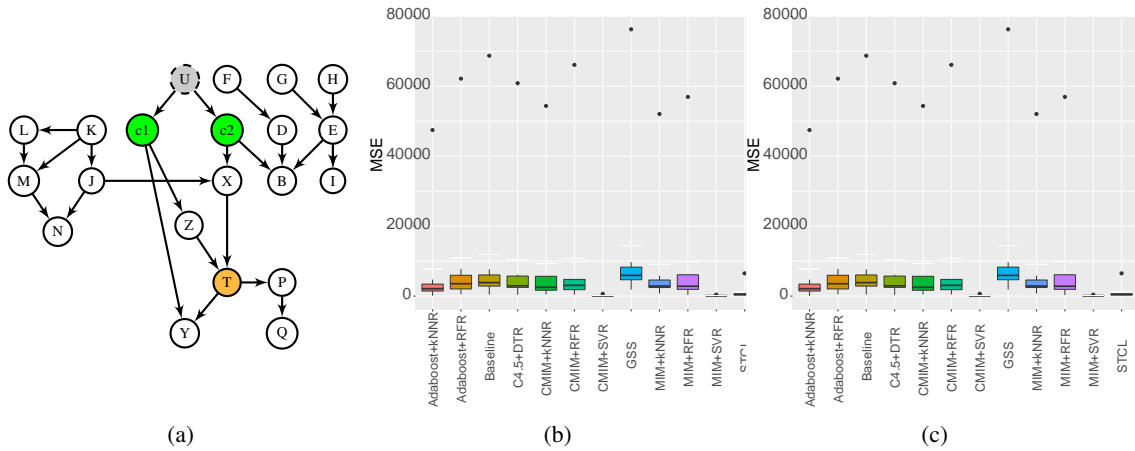




Figure 23: Results on (b), (c), (d) are generated over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 , where sample size = 50 for a Gaussian distribution.



Methodology	p-value
Baseline	0.1694
GSS	0.19617
CMIM+SVR	0.3187
CMIM+kNNR	0.1258
CMIM+RFR	0.1376
MIM+SVR	0.318705
MIM+kNNR	0.1258
MIM+RFR	0.13759
Adaboost+SVR	0.279
Adaboost+kNNR	0.20346
Adaboost+RFR	0.169
C4.5+ DTR	0.0889

(b)

Figure 24: Results in the table are p-values generated by performing T-test on SSE over the ground truth graph (a) for the target variable T . Data shift between source and target occurs by change in the distribution of the context variable C_1 , over Gaussian distribution with sample size = 1000. The bold results indicate that the average performance of SCTL is not significantly different from the average performance of other approaches for the p-value 0.05.




(a)

Methodology	p-value
Baseline	0.1535
GSS	0.1681
CMIM+SVR	0.2712
CMIM+kNNR	0.1224
CMIM+RFR	0.1885
MIM+SVR	0.65673
MIM+kNNR	0.04307
MIM+RFR	0.16345
Adaboost+SVR	0.2660
Adaboost+kNNR	0.0526
Adaboost+RFR	0.1185
C4.5+ DTR	0.0295
ESS	0.0222

(b)

Figure 25: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Datashift between source and target occurs by change in the distribution of the context variable C_1 & C_2 , over Gaussian distribution with sample size = 1000. The bold results indicate that the average performance of SCTL is not significantly different from the average performance of other approaches for the p-value 0.05.




(a)

Methodology	p-value
Baseline	0.01747
GSS	6.14934e-08
CMIM+SVR	2.9747e-08
CMIM+kNNR	0.00445
CMIM+RFR	1.7186e-17
MIM+SVR	5.415e-07
MIM+kNNR	0.0057
MIM+RFR	8.48536e-17
Adaboost+SVR	0.0001
Adaboost+kNNR	1.0681e-07
Adaboost+RFR	0.16304
C4.5+ DTR	4.2182e-08

(b)

Figure 26: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Datashift between source and target occurs by change in the distribution of the context variable C_1 , over Gaussian distribution with sample size = 1000. The bold results indicate that the average performance of SCTL is significantly different from the average performance of other approaches for the p-value 0.05.




(a)

Methodology	p-value
Baseline	2.5901e-09
GSS	9.1341e-06
CMIM+SVR	0.7350
CMIM+kNNR	0.1154
CMIM+RFR	0.0265
MIM+SVR	0.94294
MIM+kNNR	0.1687
MIM+RFR	0.0190
Adaboost+SVR	0.87851
Adaboost+kNNR	9.9208e-09
Adaboost+RFR	0.1205
C4.5+ DTR	2.6341e-05
ESS	0.6341

(b)

Figure 27: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Data shift between source and target occurs by change in the distribution of the context variable C_1 , over Gaussian data with sample size = 1000. The bold results indicate that the average performance of SCTL is significantly better than the average performance of other approaches for the p-value 0.05.




(a)

Methodology	p-value
Baseline	0.0225
GSS	1.2961e-05
CMIM+SVR	1.212e-10
CMIM+kNNR	1.6860e-14
CMIM+RFR	0.0001
MIM+SVR	1.212e-10
MIM+kNNR	3.344e-14
MIM+RFR	0.0001
Adaboost+SVR	1.212e-10
Adaboost+kNNR	5.148e-13
Adaboost+RFR	0.1290
C4.5+ DTR	2.562e-10

(b)

Figure 28: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Data shift between source and target occurs by change in the distribution of the context variables C_1 & C_2 , over Discrete data with sample size = 1000. The bold results indicate that the average performance of SCTL is significantly better than the average performance of other approaches for the p-value 0.05.




(a)

Methodology	p-value
Baseline	0.00013
GSS	0.0001
CMIM+SVR	4.785e-10
CMIM+kNNR	5.4748e-10
CMIM+RFR	0.0059
MIM+SVR	4.7853e-10
MIM+kNNR	4.7034e-10
MIM+RFR	0.0056
Adaboost+SVR	4.785e-10
Adaboost+kNNR	7.088e-10
Adaboost+RFR	0.5531
C4.5+ DTR	3.136e-09
ESS	0.0006

(b)

Figure 29: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Data shift between source and target occurs by change in the distribution of the context variables C_1 & C_2 , over Discrete data. The bold results indicate that the average performance of SCTL is significantly better than the average performance of other approaches for the p-value 0.05.




(a)

Methodology	p-value
Baseline	0.0001
GSS	1.585e-06
CMIM+SVR	7.227e-15
CMIM+kNNR	4.352e-09
CMIM+RFR	5.212e-06
MIM+SVR	7.227e-15
MIM+kNNR	7.327e-09
MIM+RFR	0.0001
Adaboost+SVR	7.229e-15
Adaboost+kNNR	3.760e-08
Adaboost+RFR	0.054
C4.5+ DTR	8.380e-10


(b)

Figure 30: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Data shift between source and target occurs by change in the distribution of the context variable C_1 , over Discrete data. The bold results indicate that the average performance of SCTL is significantly better than the average performance of other approaches for the p-value 0.05.



Methodology	p-value
Baseline	0.1651
GSS	0.1687
CMIM+SVR	3.5196e-07
CMIM+kNNR	0.0277
CMIM+RFR	0.066
MIM+SVR	3.519e-07
MIM+kNNR	0.0277
MIM+RFR	0.0664
Adaboost+SVR	0.3315
Adaboost+kNNR	0.1858
Adaboost+RFR	0.1865
C4.5+ DTR	0.0228

Figure 31: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **T**. Data shift between source and target occurs by change in the distribution of the context variable C_1 & C_2 , over Discrete data with sample size = 10000. The bold results indicate that the average performance of SCTL is comparable to the average performance of other approaches for the p-value 0.05.



Methodology	p-value
Baseline	6.7513e-12
GSS	0.0009
CMIM+SVR	0.0056
CMIM+kNNR	1.03614e-06
CMIM+RFR	1.2254e-11
MIM+SVR	1.0319e-05
MIM+kNNR	1.1642e-07
MIM+RFR	4.537e-10
Adaboost+SVR	0.0305
Adaboost+kNNR	0.0138
Adaboost+RFR	3.993e-10
C4.5+ DTR	2.0194e-07

Figure 32: Results in the table are p-values generated by performing T-test on MSE over the ground truth graph (a) for the target variable **M**. Data shift between source and target occurs by change in the distribution of the context variable C_1 , over Discrete data with sample size = 1000. The bold results indicate that the average performance of SCTL is significantly better than the average performance of other approaches for the p-value 0.05.