

# Federated Causality Learning with Explainable Adaptive Optimization

Dezhi Yang<sup>1,2</sup>, Xintong He<sup>3</sup>, Jun Wang<sup>2\*</sup>, Guoxian Yu<sup>1,2</sup>, Carlotta Domeniconi<sup>4</sup>, Jinglin Zhang<sup>5</sup>

<sup>1</sup>School of Software, Shandong University, Jinan, China

<sup>2</sup>SDU-NTU Joint Centre for AI Research, Shandong University, Jinan, China

<sup>3</sup>Department of Mathematics, National University of Singapore, Singapore

<sup>4</sup>Department of Computer Science, George Mason University, VA, USA

<sup>5</sup>School of Control Science and Engineering, Shandong University, Jinan, China

dzyang@mail.sdu.edu.cn, E0966399@u.nus.edu.sg, {kingjun, gxyu, jinglin.zhang}@sdu.edu.cn, carlotta@cs.gmu.edu

## Abstract

Discovering the causality from observational data is a crucial task in various scientific domains. With increasing awareness of privacy, data are not allowed to be exposed, and it is very hard to learn causal graphs from dispersed data, since these data may have different distributions. In this paper, we propose a federated causal discovery strategy (FedCausal) to learn the unified global causal graph from decentralized heterogeneous data. We design a global optimization formula to naturally aggregate the causal graphs from client data and constrain the acyclicity of the global graph without exposing local data. Unlike other federated causal learning algorithms, FedCausal unifies the local and global optimizations into a complete directed acyclic graph (DAG) learning process with a flexible optimization objective. We prove that this optimization objective has a high interpretability and can adaptively handle homogeneous and heterogeneous data. Experimental results on synthetic and real datasets show that FedCausal can effectively deal with non-independently and identically distributed (non-IID) data and has a superior performance.

## Introduction

Causal structure discovery is a fundamental and critical problem in many fields, such as economics (Koller and Friedman 2009) and biology (Sachs et al. 2005). Randomized controlled experiments are the golden standard for discovering causal structure, but they are often limited by cost and may even be prohibited by ethics. Causal discovery typically infers a directed acyclic graph (DAG) from observational data at a central site (Pearl 2009; Peters, Janzing, and Schölkopf 2017); this DAG encodes causal relationships between variables. However, with the increasing awareness of privacy and security, data are scattered among different clients and cannot be shared, which makes it difficult for canonical causal discovery algorithms to find reliable causal structure from limited client data.

Federated learning (FL), as a secure framework for cooperative training with multiple clients, learns a unified model from scattered data by exchanging model parameters or gradients among clients, without exposing local clients' data (McMahan et al. 2017). FL has made good progress in areas

such as image classification and recommendation systems (Chai et al. 2020). However, recent FL algorithms are based on continuous optimization and cannot be directly applied to causal discovery algorithms with a combinatorial optimization property. For example, constraint-based algorithms PC (Spirtes et al. 2000) and FCI (Spirtes, Meek, and Richardson 2013) use conditional independence between variables to judge whether there is a causal structure, while score-based algorithms GES (Chickering 2002) and the max-min hill-climbing algorithm (Tsamardinos, Brown, and Aliferis 2006) use score functions and heuristic search strategies to find causal graphs with the best scores.

Recent gradient-based causal discovery algorithms, Zheng et al. (2018) and Yu et al. (2019) make use of smooth equality constraints instead of discrete acyclic constraints to discover causal structures through continuous optimization (i.e., Augmented Lagrange method (Nemirovsky 1999)), and they provide the opportunity of learning causal graphs in a continuous manner within a federated framework. Subsequently, Lachapelle et al. (2019) and Zheng et al. (2020) introduced deep neural networks to deal with more complex causal models. However, these gradient-based algorithms need to center data, which is not feasible due to privacy protection. For dispersed data, several distributed causal discovery methods have been proposed. Some of them need to make assumptions about parameters but lack generalization (Shimizu 2012; Xiong et al. 2021), and others need to share additional learning parameters but cause privacy leakage (Na and Yang 2010; Ye, Amini, and Zhou 2022). More importantly, they are often unable to deal with non-independent identical distributed (non-IID) data among clients.

In this paper, we develop a general federated DAG learning strategy (FedCausal) to seek the global causal graph from horizontally partitioned data with different distributions. This strategy designs a global optimization process instead of traditional weighted average in the server to aggregate local causal graphs, which naturally ensures the sparsity and acyclicity of the global causal graph. The local and global optimization processes form a whole as an explainable adaptive optimization objective, which is consistent with the centralized optimization objective under statistical homogeneity. In addition, this explainable objective allows clients to flexibly learn local data distributions with statistical heterogeneity and seek a uniform global graph on

\*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

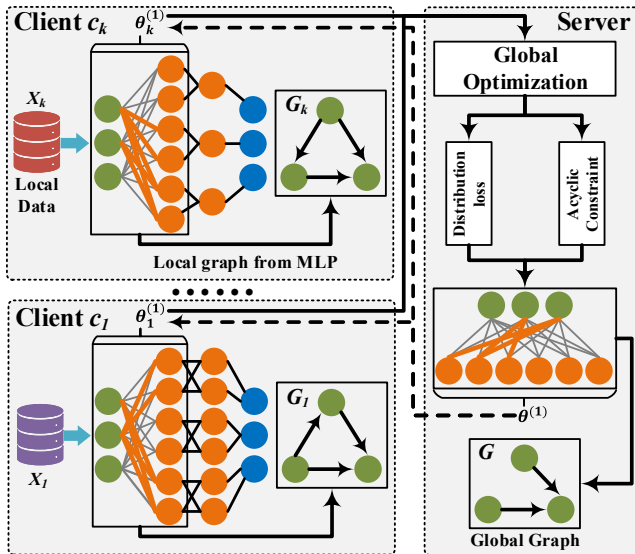


Figure 1: Schematic framework of FedCausal on the non-parametric model. In each interaction, the clients optimize local causal models based on local data and send the first layer parameters  $\theta_k^{(1)}$  to the server. The server optimizes the global model to conform to the local distributions and explicitly constrains its acyclicity to obtain a global causal graph  $G$ . FedCausal then broadcasts server parameters  $\theta^{(1)}$  to clients for the next round interaction.

the server. Figure 1 shows the conceptual framework of FedCausal. Our main contributions can be outlined as follows: (i) We meticulously extend the centralized DAG learning framework to federated scenarios. FedCausal explicitly constrains the acyclicity of the global causal graph and optimizes it to conform to dispersed data, ensuring effective causal information interaction between local and global. (ii) FedCausal unifies the global aggregation optimization and local DAG optimization, and formulates a complete causal graph learning process. We prove that its optimization process is consistent with centralized DAG optimization on homogeneous data and is suitable for heterogeneous data. (iii) We conduct experiments on synthetic and real datasets, and prove that FedCausal is close to or even better than centralized learning on homogeneous data, and outputs more accurate and acyclic causal graphs than other methods (Ng and Zhang 2022; Gao et al. 2023) on heterogeneous data.

## Related Work

The number of DAGs is super-exponential in the number of variables. As such, learning discrete DAGs is generally NP-hard for traditional causal discovery algorithms (Spirtes et al. 2000; Chickering 2002; Bernstein et al. 2020). To avoid the difficult combinatorial optimization of the acyclic constraint, NOTEARS (Zheng et al. 2018) transforms the acyclic constraint into a smooth equality constraint, and converts the causal discovery problem into a continuous optimization one with efficient solvers. DAG-GNN (Yu et al. 2019) extends NOTEARS using graph neu-

ral networks to handle various variable types. Gran-DAG (Lachapelle et al. 2019) uses neural networks to deal with non-linear causal relationships. NOTEARS-MLP (Zheng et al. 2020) proposes a generalized function model and characterizes a non-parametric structure equation model of DAG via partial derivatives. HetDAG (Liang et al. 2023) extends NOTEARS-MLP to learn the DAGs of attributed heterogeneous network, and DARING (He et al. 2021) considers the independence of heterogeneous noises to improve the robustness. MCSL (Ng et al. 2022) adopts a Gumbel – Sigmoid function to approximate the parameter matrix of a DAG into a binary adjacency one. However, the above approaches target at centralized data, and their assumptions or strategies for identifying causal graphs cannot cope with statistically heterogeneous data. Based on the continuous optimization framework, we design a global optimization to flexibly learn the global causal graph by aggregating local causal graphs from decentralized data that may have different distributions.

Several attempts have been made to learn the causal structure from distributed datasets with different subsets of variables. (Tillman, Danks, and Glymour 2008) estimated the local partial ancestral graph from each dataset, and then found global partial ancestral graphs on the complete set of variables with the same  $d$ -connection and  $d$ -separation relationships as the local PAGs. (Triantafillou and Tsamardinos 2015) used a SAT solver in a similar process to improve the algorithm’s scalability. However, these methods cannot uniquely identify the causal graph and suffer from a large indeterminacy. For the horizontally distributed dataset, Gou, Jun, and Zhao (2007) and Na and Yang (2010) used a two-step process that independently estimates the Bayesian network by each client’s dataset, then performs conditional independent testing and voting to obtain the final causal graph. These distributed methods separately use local datasets during training and directly share local models parameters, which result in poor performance and privacy leakage. Our FedCausal aggregates the global causal graph using a very small fraction of the parameters of local models, effectively avoiding the risk of reconstructing local data from model parameters. It does not simply aggregate the local causal graphs, but adds a proximal term and global optimization to enable sufficient information to be exchanged between clients and the server.

The recent NOTEARS-ADMM (abbreviated as NO-ADMM) approach (Ng and Zhang 2022) adopts the alternating direction method of multipliers to decompose the continuous optimization objective of gradient-based methods, and learns the global causal structure by exchanging client model parameters. However, NO-ADMM requires all clients to participate in the training and makes all local models consistent, and thus lacks flexibility and does not account for statistical heterogeneity. Although FedDAG (Gao et al. 2023) takes into account heterogeneous data, it directly masks local models with a global binary matrix. As a consequence, it cannot effectively update the masked model parameters and obtain accurate causal graphs. Both NO-ADMM and FedDAG enforce too stringent local models and suffer from performance degradation. FedCausal nat-

urally combines local and global optimizations, forms an adaptive optimization process, and has a consistent optimization objective with the centralized DAG learning under IID data without additional assumptions. Compared with existing federated DAG learning methods, FedCausal is more interpretable and more suitable for non-IID data, it trains local models without strong restrictions to learn the unified global graph. FedCausal accurately identifies the unique causal graph while conforming to the heterogeneous distribution of local data.

## Background

The causal discovery problem can be defined as follows: given the observed data  $X \in \mathbb{R}^{n \times d}$  and the variables' set  $V = \{V_1, \dots, V_d\}$ , we need to learn a DAG  $\mathcal{G}$  from  $X$ , where  $n$  and  $d$  represent the sample size and the number of variables respectively. In the federated scenario, this problem becomes more complicated: let the number of clients be  $K$ ,  $X = \{X_k\}_{k=1}^K$  represent the clients' local data, and  $N = \{n_k\}_{k=1}^K$  represent the sample size of clients' data; we need to learn a uniform causal graph  $\mathcal{G}$  without exposing clients' data. The difficulty is that clients' data may be non-iid. Even though existing DAG learning methods based on continuous optimization mostly focus on centralized data and cannot handle statistical heterogeneity, they are necessary for federated causal discovery because most of the federated learning algorithms adopt continuous optimization. For this reason we introduce next some of the existing continuous DAG learning strategies.

### Continuous Optimization for DAG Learning

NOTEARS uses equality to constrain the acyclicity of causal graphs and continuously optimizes a causal graph on the linear data model. NOTEARS defines a weight matrix  $W \in \mathbb{R}^{d \times d}$  to represent the causal model, thus  $W_{ij} = 0$  if and only if there is no edge from  $V_i$  to  $V_j$  in the real causal graph  $\mathcal{G}$ . Then, it uses the equation  $h(W) = \text{tr}(e^{W \circ W}) - d = 0$  to replace the discrete constraint  $\mathcal{G} \in \text{DAGs}$  (see (Zheng et al. 2018) for details). With this constraint, NOTEARS reconstructs data from  $W$  to minimize the reconstruction residual and designs the following constrained optimization formula:

$$\begin{aligned} \min_W L(W; X) &= \frac{1}{2n} \|X - XW\|_2^2 + \lambda \|W\|_1 \\ \text{s.t. } h(W) &= \text{tr}(e^{W \circ W}) - d = 0 \end{aligned} \quad (1)$$

where  $\|W\|_1$  guarantees the sparsity of the causal graph, and  $\lambda$  is the hyper-parameter. The approximate solution is sought via the L-BFGS-B algorithm (Byrd et al. 1995).

To extend NOTEARS to non-parametric models, Zheng et al. (2020) proposed NOTEARS-MLP, which learns the causal model via a multi-layer perceptron (MLP) and uses partial derivatives to express the dependence between variables. NOTEARS-MLP learns the causal generation model  $f = \{f_1, \dots, f_d\}$  for each variable. It is easy to show that there is no edge from  $V_j$  to  $V_i$ , if and only if the partial derivative  $\partial_j f_i$  of  $f_i$  with respect to  $X_j$  is equal to 0. So, NOTEARS-MLP uses the partial derivatives of  $f$  to represent the parameter matrix  $W(f)$ :  $[W(f)]_{ji} = \|\partial_j f_i\|_2$ . Let

$\theta = \{(A_1^{(1)}, \dots, A_1^{(h)}), \dots, (A_d^{(1)}, \dots, A_d^{(h)})\}$  be the parameters of all MLPs and  $A_i^{(h)}$  be the  $h$ -th layer parameter of the MLP corresponding to  $f_i$ , NOTEARS-MLP uses the first layer parameters  $\theta^{(1)} = \{A_1^{(1)}, \dots, A_d^{(1)}\}$  of all MLPs to represent the weight matrix  $[W(\theta^{(1)})]_{ji} = \|A_{i,(\cdot,j)}^{(1)}\|_2$ . The constrained optimization formula of NOTEARS-MLP is as follows:

$$\begin{aligned} \min_{\theta} L(\theta; X) &= \frac{1}{2n} [\|X - \text{mlp}(X; \theta)\|_2^2 + \lambda \|\theta\|_1] \\ \text{s.t. } h(\theta^{(1)}) &= \text{tr}(e^{W(\theta^{(1)}) \circ W(\theta^{(1)})}) - d = 0 \end{aligned} \quad (2)$$

where  $\text{mlp}(X; \theta)$  is the reconstruction result of original data  $X$  from all causal generation models parameterized by MLP and  $\|\theta\|_1$  guarantees the causal model sparsity.

Both NOTEARS and NOTEARS-MLP require centered IID data. In contrast, FedCausal targets decentralized non-IID data. It unifies local DAG learning on non-IID data into a global optimization framework, which improves the causality learning capability of client models under limited local data, and ensures unity and accuracy of global causal graphs under statistical heterogeneity. The global model of FedCausal corresponds to the unique causal graph that explains causality between variables, and the local models learn the heterogeneous distributions of clients' data.

## Proposed Methodology

### Optimization for Linear SEM

To learn causal graphs that conform to the data distribution, both centralized and federated DAG learning methods aim to minimize the reconstruction residual of observational data and guarantee the graph acyclicity. For a linear structural equation model (SEM), assume that there are  $K$  clients holding  $n = \sum_{k=1}^K n_k$  samples in total. If we can collect decentralized data  $\{X_k\}_{k=1}^K$ , the global weight matrix  $W$  can be learned as follows:

$$\begin{aligned} W &= \arg \min_W \sum_{k=1}^K \frac{1}{2n_k} [\|X_k - X_k W\|_2^2 + \lambda \|W\|_1] \\ \text{s.t. } \text{tr}(e^{W \circ W}) &- d = 0 \end{aligned} \quad (3)$$

However, clients cannot expose their own data and they can only use local data to optimize the local weight matrix  $W_k \in \mathbb{R}^{d \times d}$ . In addition, the acyclic constraint term is applied to the local matrix  $W_k$  rather than the global  $W$ . A natural idea is to aggregate the local matrices into a global one with guaranteed sparsity and acyclicity.

Without loss of generality, we assume that the probability of the observational data being distributed over  $k$ -th client is equal to the frequency  $n_k/n$  of samples appearing in it, given sufficient data. To measure local empirical risks over possible data distributions, the global matrix should be the weighted average of the local matrices based on sample sizes:  $W = \sum_{k=1}^K n_k W_k / n$ . Unfortunately, clients may learn weight matrices with opposite causal relationships, which introduces a strong numerical bias to the global matrix and violates the acyclicity. To solve this problem, we

design a global constrained optimization to aggregate local matrices and ensure the global sparsity and acyclicity:

$$\begin{aligned} \min_W L(W) &= \sum_{k=1}^K \left[ \frac{n_k}{n} \|W - W_k\|_2^2 \right] \\ \text{s.t. } h(W) &= \text{tr}(e^{W \circ W}) - d = 0 \end{aligned} \quad (4)$$

where  $n_k \|W - W_k\|_2^2 / n$  is the distribution loss term to constrain the global matrix to approach the local matrix. Obviously, Eq. (4) measures and minimizes local empirical risks, and its equality constraint also explicitly constrains the acyclicity of the global causal graph. Since the global matrix resulting from aggregating sparse local matrices is also numerically sparse, we do not add the sparse penalty term to Eq. (4). We broadcast the learned global matrix  $W$  to all clients and further optimize their local matrices as follows:

$$\begin{aligned} \min_{W_k} L_k(W_k; X_k) &= \frac{1}{2n_k} \|X_k - X_k W_k\|_2^2 + \lambda_1 \|W_k\|_1 \\ \text{s.t. } h(W_k) &= \text{tr}(e^{W_k \circ W_k}) - d = 0 \end{aligned} \quad (5)$$

Local matrix optimization still requires the sparsity penalty and acyclic constraint to prevent its overfitting. In practice, we find that adding a penalty term  $\lambda_2 \|W_k - W\|_2^2$  to bring local matrices close to the global one is beneficial for the convergence of local optimization. This penalty term is weaker than the equality constraint of NO-ADMM (Ng and Zhang 2022), which forces the local matrix to be equal to the global matrix, but it is more flexible and suitable for addressing the statistical heterogeneity. FedCausal first optimizes local matrices  $\{W_k\}_{k=1}^K$  at clients using Eq. (5) and uploads them to the server. The server then optimizes the global matrix  $W$  using Eq. (4) and distributes it to the local models. FedCausal repeats the above local and global optimizations for a sufficiently large number of iterations, or stops when the global acyclic constraint term falls below a pre-defined threshold.

### Optimization for Non-Parametric SEM

The non-parametric model is more complex than the linear model, since it has unknown parameters and clients' data distributions. Exchanging all models' parameters between the clients and the server not only causes excessive communication overheads, but also leads to the same local models and to the violation of statistical heterogeneity. To address these issues, we adjust the global optimization and pass only a few parameters between the clients and the server. We denote the parameters of the local causal generation model as  $\theta_k$ , the parameters of the first layer as  $\theta_k^{(1)}$  and those of other layers as  $\theta_k^{(-1)}$ . The global optimization formula of FedCausal in the non-parametric model is:

$$\begin{aligned} \min_{\theta^{(1)}} L(\theta^{(1)}) &= \sum_{k=1}^K \left[ \frac{n_k}{n} \|\theta^{(1)} - \theta_k^{(1)}\|_2^2 \right] \\ \text{s.t. } h(\theta^{(1)}) &= \text{tr}(e^{W(\theta^{(1)}) \circ W(\theta^{(1)})}) - d = 0 \end{aligned} \quad (6)$$

FedCausal receives only the first layer parameters  $\theta_k^{(1)}$  of the local models to optimize the global model. The unified

---

### Algorithm 1: FedCausal: Federated Causal Learning

---

**Input:** Observed data  $X = \{X_1, \dots, X_K\}$

**Output:** Causal weight matrix  $W(\theta^{(1)})$

---

```

1: Initial global parameters  $\theta^{(1)}$ , local parameters
    $\{\theta_k = (\theta_k^{(1)}, \theta_k^{(-1)})\}_{k=1}^K$  and hyperparameter-list
    $\{\alpha, \rho, h_{tol}, \rho_{max}, \gamma, \beta\}$ 
2: for  $t = 0, 1, 2, \dots$  do
3:   while  $\rho < \rho_{max}$  do
4:     Broadcast  $\theta^{(1)}, \alpha, \rho: \theta_k^{(1)} = \theta^{(1)}, \alpha_k = \alpha, \rho_k = \rho$ 
5:     Update local parameters  $\{\theta_k\}_{k=1}^K$  using Eq. (7)
6:     Upload  $\{\theta_k^{(1)}\}_{k=1}^K$  to the server
7:     Update global parameters  $\theta^{(1)}$  using Eq. (6)
8:     if  $h(\theta^{(1)}) > \gamma H$  then
9:        $\rho = \beta \rho$ 
10:    else
11:      Break
12:    end if
13:  end while
14:  Set  $H = h(\theta^{(1)})$ ,  $\alpha = \alpha + \rho h(\theta^{(1)})$ 
15:  if  $h(\theta^{(1)}) \leq h_{tol}$  or  $\rho \geq \rho_{max}$  then
16:    Break
17:  end if
18: end for
19: return  $W(\theta^{(1)})$ 

```

---

global causal graph is derived from the partial derivative of the first-layer global model parameters  $\theta^{(1)}$ , so there is no need to send  $\theta_k$ . This design not only reduces the communication overhead compared to passing all parameters as in NO-ADMM (Ng and Zhang 2022), but also enforces privacy protection, because it is difficult for attackers to reconstruct data from the first layer parameters of a complex model. We then broadcast  $\theta^{(1)}$  to all clients, replace their  $\theta_k^{(1)}$ , and optimize their  $\theta_k$  with the following formula:

$$\begin{aligned} \min_{\theta_k} L_k(\theta_k; X_k) &= \frac{1}{2n_k} \|X_k - \text{mlp}(X_k; \theta_k)\|_2^2 + \lambda_1 \|\theta_k^{(1)}\|_1 \\ \text{s.t. } h(\theta_k^{(1)}) &= \text{tr}(e^{W(\theta_k^{(1)}) \circ W(\theta_k^{(1)})}) - d = 0 \end{aligned} \quad (7)$$

FedCausal learns the global causal graph using Eq. (6), allows clients to learn different causal models conforming to non-IID local data using Eq. (7), and thus provides high flexibility for causal discovery in non-IID data. Algorithm 1 gives the main process of FedCausal on non-parametric data, where  $\alpha$  and  $\rho$  are the Lagrangian multiplier and the penalty parameter,  $h_{tol}$  and  $\rho_{max}$  are the thresholds that control the loop termination,  $\gamma$  and  $\beta$  are the condition and the step size that control the update of  $\rho$  and  $\alpha$ . After obtaining the causal weight matrix  $W(\theta^{(1)})$ , we prune the edges with small weight ( $< 0.3$ ) to get the DAG.

Let's consider the case of  $K$  client models with one hidden layer of  $m$  units. FedCausal's clients take  $O(n_k d^2 m + d^2 m + d^3)$  to compute the objective and the gradient. The computation complexity of FedDAG's clients is  $O(n_k d^2 m + n_k d^2 + d^2 m + d^3)$ , due to the additional mask operation.

NO-ADMM does not compute the acyclic term on clients, so its clients' computation complexity is  $O(n_k d^2 m + d^2 m)$ , but it suffers from overfitting local models. Both FedCausal and NO-ADMM optimize the global graph on the server, so their complexity on the server side is  $O(K d^2 m + d^3)$ . However, FedCausal only needs to compute the first layer parameters of the model, so it is usually faster than NO-ADMM, especially when the model has at least two layers. The communication overheads of FedCausal, NO-ADMM, and FedDAG for a single round interaction are  $d^2 m$ ,  $d^2 m + d m$  and  $d^2$ , respectively. As the model becomes complex with more layers, the overhead of NO-ADMM sharply increases. FedDAG averages only local matrices and has a lower server computation than NO-ADMM and FedCausal, but its global matrix may not uncover an accurate global causal graph, as we show in our experiments.

### Analysis of Explainable Adaptive Optimization

**Lemma 1.** *Under the condition that local data are homogeneous and obey linear SEM, the optimization objective of FedCausal is consistent with that of centralized DAG learning, i.e. the global matrix  $W \rightarrow \arg \min_W \|X - XW\|_2^2/2n + \lambda\|W\|_1 + \rho h^2(W)/2 + \alpha h(W)$ .*

For the linear data model, the optimization objective of the penalty term  $\|W - W_k\|_2^2$  in Eq. (4) is  $W \rightarrow W_k$ , and that of Eq. (5) is  $W_k \rightarrow \arg \min_{W_k} L_k(W_k; X_k) + \rho_k h^2(W_k)/2 + \alpha_k h(W_k)$ , so the actual optimization objective of  $\|W - W_k\|_2^2$  is  $W \rightarrow \arg \min_W L_k(W; X_k) + \rho_k h^2(W)/2 + \alpha_k h(W)$ . We can then replace the penalty term in the global optimization (Eq. (4)) with the actual objective, and obtain the following:

$$W \rightarrow \arg \min_W \sum_{k=1}^K \frac{n_k}{n} \left[ \frac{1}{2n_k} \|X_k - X_k W\|_2^2 + \lambda \|W\|_1 \right] + \frac{\rho}{2} h^2(W) + \alpha h(W) \quad (8)$$

Eq. (8) only preserves the global acyclic constraint term, because the local acyclic constraint acts on the local matrices and does not constrain the global matrix. By merging the terms in Eq. (8), we find that the optimization objective of the global matrix is:  $W \rightarrow \arg \min_W \|X - XW\|_2^2/2n + \lambda\|W\|_1 + \rho h^2(W)/2 + \alpha h(W)$ . Obviously, this optimization objective is consistent with the centralized optimization objective of NOTEARS, which means that FedCausal may achieve the same results as NOTEARS, even though the data are scattered.

**Lemma 2.** *Under the condition that local data are homogeneous and obey non-parametric SEM, the optimization objective of FedCausal is consistent with that of centralized DAG learning, i.e. the global parameters  $\theta^{(1)} \rightarrow \arg \min_{\theta^{(1)}} \|X - \text{mlp}(X; \theta^{(1)}, \theta^{(-1)})\|_2^2/2n + \lambda\|\theta^{(1)}\|_1 + \rho h^2(\theta^{(1)})/2 + \alpha h(\theta^{(1)})$ .*

For the non-parametric data model, the optimization objective of the penalty term  $\|\theta^{(1)} - \theta_k^{(1)}\|_2^2$  in Eq. (7) is  $\theta^{(1)} \rightarrow \theta_k^{(1)}$ . By expressing the local parameters  $\theta_k$  optimized by Eq. (6) as  $(\theta_k^{(1)}, \theta_k^{(-1)})$ , we obtain

the actual optimization objective of  $\|\theta^{(1)} - \theta_k^{(1)}\|_2^2$  as  $\theta^{(1)} \rightarrow \arg \min_{\theta^{(1)}} L_k(\theta^{(1)}, \theta_k^{(-1)}; X_k) + \rho_k h^2(\theta^{(1)})/2 + \alpha_k h(\theta^{(1)})$ . Similarly, we can deduce the global optimization objective in the non-parametric data model as:

$$\theta^{(1)} \rightarrow \arg \min_{\theta^{(1)}} \sum_{k=1}^K \frac{n_k}{n} \left[ \frac{1}{2n_k} \|X_k - \text{mlp}(X_k; \theta^{(1)}, \theta_k^{(-1)})\|_2^2 + \lambda \|\theta^{(1)}\|_1 \right] + \frac{\rho}{2} h^2(\theta^{(1)}) + \alpha h(\theta^{(1)}) \quad (9)$$

We consider two cases. On the one hand, if clients' data are *homogeneous*, the local parameter  $\theta_k^{(-1)}$  can be the same or very similar, so we can merge the terms in Eq. (9) to obtain the final optimization objective:  $\theta^{(1)} \rightarrow \arg \min_{\theta^{(1)}} \|X - \text{mlp}(X; \theta^{(1)}, \theta^{(-1)})\|_2^2/2n + \lambda\|\theta^{(1)}\|_1 + \rho h^2(\theta^{(1)})/2 + \alpha h(\theta^{(1)})$ . This optimization objective is also consistent with the centralized optimization objective of NOTEARS-MLP. On the other hand, if clients have decentralized *heterogeneous* data, Eq. (9) allows clients to learn different  $\theta_k^{(-1)}$ , and still optimizes a unified  $\theta^{(1)}$ . In other words, regardless whether the local data fit linear or non-linear models, FedCausal learns different data joint distributions on the clients and optimizes a uniform causal graph on the server. In other words, FedCausal adapts to the case where clients have different causal generation models but the same causal graph. In summary, our FedCausal can flexibly adapt to statistical heterogeneous/homogeneous data with an explainable proof.

## Experiments

We compare FedCausal against recent federated causal discovery algorithms, including FedDAG (Gao et al. 2023) and NO-ADMM (Ng and Zhang 2022). We also set three baselines based on NOTEARS (Zheng et al. 2018). The first baseline (NO-ALL) learns DAGs from all data by centering local data; the second (NO-Avg) independently learns local causal graphs and combines these graphs by weighted average; the third one (NO-w/oAcy) is similar to our strategy but disregards the acyclic constraint on the global causal graph. For nonlinear or heterogeneous data, we adapt the nonlinear or non-parametric version of FedCausal, FedDAG and NO-ADMM, and replace NOTEARS with NOTEARS-MLP (Zheng et al. 2020) for NO-ALL, NO-Avg and NO-w/oAcy. For real datasets, we add two classic baselines PC (Spirtes et al. 2000) and GES (Chickering 2002), which use all data but identify complete partially DAG (CPDAG) with undirected edges. We provide the hyperparameter settings of FedCausal and other baselines in the supplementary file (Yang et al. 2023).

We study the performance of the compared methods on synthetic IID and non-IID datasets, as well as real datasets. We use the Erdős-Rényi (ER) and Scale-free (SF) models to generate the ground truth DAGs. We use the linear Gaussian (LG) model and the additive noise model with MLP (ANM-MLP) (Zheng et al. 2020) respectively to synthesize linear and nonlinear data. For non-IID data, we fix the generated DAG and randomly sample models from the linear Gaussian model (LG), the additive noise model with

MLP (ANM-MLP), the additive model with Gaussian processes (ADD-GP) (Bühlmann, Peters, and Ernest 2014) and the additive index model (MIM) (Yuan 2011) to generate data for different clients. We guarantee that, even if different clients choose the same model, their model parameters will be different. We use the same server (Ubuntu 18.04.5, Intel Xeon Gold 6248R and Nvidia RTX 3090) to perform experiments and report the structural hamming distance (SHD), true positive rate (TPR) and false discovery rate (FDR) of the estimated DAGs, averaged over 10 random runs. A larger TPR indicates a better performance, while the opposite trend holds for SHD and FDR. The code of FedCausal is shared at <https://www.sdu-idea.cn/pubDetail?pubId=279>.

## Results on IID Data

To test FedCausal on IID data, we set up 10 clients and generate 200 samples for each client with linear (LG) and nonlinear (ANM-MLP) IID data models. We conduct the experiment to estimate randomly generated DAGs with {10, 20, 40, 80} variables. Based on the results of FedCausal and other baselines in Figure 2, we have some observations: (i) FedCausal obtains the highest TPR on the linear model across different numbers of variables. NO-ALL learns the causal graph using all data, so its results can be taken as an upper bound. FedCausal even outperforms NO-ALL, probably because there are sufficient data for clients to learn local causal graphs on the simple linear model. In addition, the collaboration between clients works alike an ensemble system that helps to reinforce correct edges and remove erroneous ones. NO-ADMM achieves the lowest TPR and highest FDR and SHD, since it does not consider the sparsity and acyclicity on clients, which cause overfitted local graphs and fail to form a good global causal graph. FedDAG performs well in the linear model, approaching NO-ALL with only a few variables; it does not with more variables. NO-Avg per-

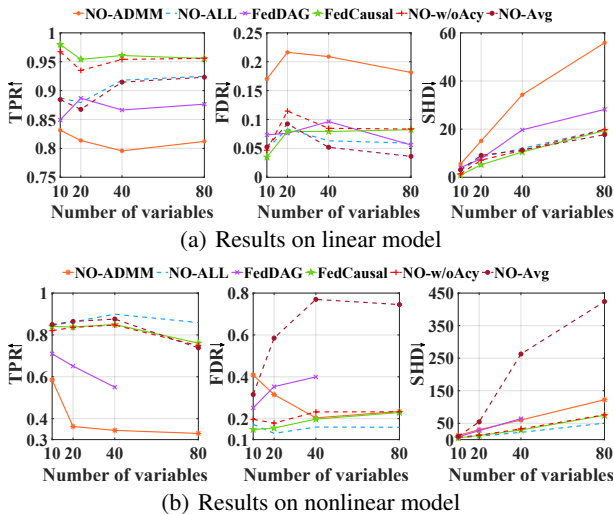


Figure 2: Results on linear and non-linear IID data. Solid lines correspond to federated methods and dotted lines are NOTEARS-based baselines.

forms almost as well as NO-ALL on the linear model, since the linear model is simple and clients can confidently learn local graphs without any collaboration.

(ii) FedCausal also achieves the lowest FDR and SHD values on the nonlinear model, being clearly superior to NO-ALL, with very close TPR values. This is because our global optimization aligns with the centralized optimization objective of NO-ALL that uses all data, and federated learning allows clients to collaborate with each other to reduce erroneous edges. For the same reason as for the linear model, NO-ADMM still performs poorly on the nonlinear model. FedDAG does not perform as well on the nonlinear model, as it did on the linear one, because the Gumbel-Sigmoid function it uses is not suitable for dealing with nonlinear causality, which leads to biased results. We do not report the results of FedDAG on 80 nodes, since Gao et al. (2023) only provided the hyperparameters of FedDAG on 10, 20, and 40 nodes, and FedDAG is very sensitive to the hyperparameters and fails on 80 nodes. NO-Avg yields unreliable results with very high FDR and SHD values, even though its TPR is also high. This indicates that simply averaging local causal graphs may result in more erroneous edges.

(iii) Our aggregation optimization effectively improves the performance of FedCausal. Compared to NO-w/oAcy, FedCausal performs better on the linear model and obtains significantly lower FDR and SHD values on the nonlinear model. This is because our adaptive optimization not only measures the local empirical risk, but also constrains the acyclicity of the global graph, while NO-w/oAcy disregards this constraint. FedCausal trains the global graph to conform to the local data distribution and eliminates the numerical bias caused by aggregation that may lead to a cyclic graph, thus effectively improves the global graph.

## Results on Non-IID Data

In this experiment, we also set 10 clients, 200 samples for each client and graphs with {10, 20, 40, 80} variables. We randomly select models from LG, ANM-MLP, ADD-GP or MiM to generate non-IID data for different clients. Figure 3 shows the results of FedCausal and other baselines. We have some important observations:

(i) FedCausal manifests the best performance among the compared methods on decentralized heterogeneous data. NO-ADMM pursues the consistency of all local models and thus does not identify the global causal graph in this practical and challenging non-IID setting. NO-ALL gets extremely high FDR values because the training data are mixed with multiple causal models, which prevent NO-ALL to accurately identify causal relationships. This observation suggests that causal algorithms with the prerequisite of centralized data cannot identify the causal graphs from non-IID data. Federated causal discovery algorithms (FedCausal, FedDAG, and NO-w/oAcy) perform better on heterogeneous data than on homogeneous nonlinear data. This is because clients discover the same false causal relationships on IID data and reinforce these false edges in aggregation, while clients on non-IID data can offset the false edges. FedCausal has a much lower FDR than FedDAG and a higher TPR than NO-w/oAcy. The graph found by NO-

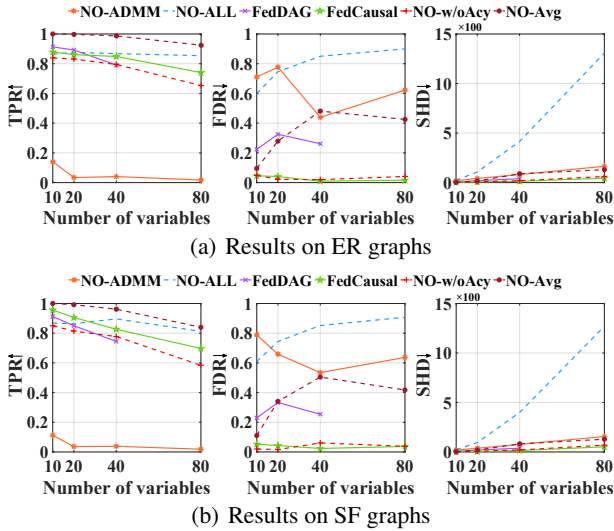


Figure 3: Results on linear and non-linear non-IID data.

Avg has a very high TPR and FDR values without guaranteed acyclicity.

(ii) FedCausal effectively improves the local causal graphs and guarantees the acyclicity. NO-Avg cannot obtain a reliable and acyclic causal graph on non-IID data, due to the poor local results and high  $h(\theta^{(1)})$  (a smaller  $h(\theta^{(1)})$  indicates a better acyclicity) in Table 1. This is because the clients of NO-Avg are trained independently, but the limited local data can't support them to accurately learn the nonlinear and heterogeneous data models, which are more complex than linear ones. In contrast, by aggregating local results at each interaction, FedCausal, NO-w/oAcy and FedDAG significantly improve the quality of local graphs. FedDAG learns local graphs masked by an approximate binary global matrix, which allows clients to get very similar results with global  $h(\theta^{(1)})$  equal to 0. However, the masked model parameters cannot be effectively updated, which causes the model to be unable to correct error edges and limits its accuracy. In terms of the client performance and global acyclicity, FedCausal is superior to NO-w/oAcy. This proves the effectiveness of our adaptive optimization that guarantees the acyclic global causal graph for clients.

In addition, to study the communication efficiency of FedCausal and other federated baselines, we show the performance trend of the algorithm as a function of the number

40 nodes	Metrics of local graphs		$h(\theta^{(1)}) \downarrow$
	TPR $\uparrow$	FDR $\downarrow$	
FedDAG	82.4 $\pm$ 0.0	10.8 $\pm$ 0.3	$\approx 0$
FedCausal	<b>92.4<math>\pm</math>1.9</b>	<b>4.4<math>\pm</math>2.8</b>	$\approx 4.2 \times 10^{-12}$
NO-w/oAcy	84.7 $\pm$ 3.4	<b>8.3<math>\pm</math>5.0</b>	$\approx 7.8 \times 10^{-7}$
NO-Avg	77.0 $\pm$ 11.2	71.9 $\pm$ 6.0	$\approx 2.2 \times 10^{-1}$

Table 1: Evaluation of global acyclic constraint term  $h(\theta^{(1)})$  and local causal graph

		TPR $\uparrow$	FDR $\downarrow$	SHD $\downarrow$	NNZ
CP	PC	<b>50.00</b>	61.54	24	26
	GES	40.00	78.95	31	38
DAG	FedCausal	<b>45.00</b>	<b>52.63</b>	<b>15</b>	19
	NO-ALL	<b>45.00</b>	57.14	18	21
	NO-w/oAcy	40.00	69.23	24	26
	NO-Avg	40.00	70.37	25	27
	FedDAG	40.00	71.43	26	28
	NO-ADMM	20.00	80.00	29	20

Table 2: Results on the Sachs dataset. DAG and CP refer to algorithms that identify directed acyclic graph and complete partially directed acyclic graph, respectively.

of communications in the supplementary file (Yang et al. 2023). The results show that FedCausal has the most superior communication efficiency. We also explore the robustness of FedCausal against different numbers of clients and local sample sizes, and present the results in the supplementary file (Yang et al. 2023). We observe that FedCausal is stable and in all cases outperforms the rivals.

### Results on a Real Dataset

We evaluate FedCausal on a protein signaling network based on expression levels of proteins and phospholipid given by (Sachs et al. 2005), which is generally accepted by the biology community. There are a total of  $n = 7466$  samples,  $d = 11$  variables, and 20 edges in the ground truth DAG of this dataset (Sachs). We randomly select  $n = 7460$  samples and evenly distribute them to 10 clients to mimic scattered data. The results are shown in Table 2.

FedCausal achieves the best and most reliable results. PC and GES, as centralized and typical causal discovery algorithms, both show good TPR but high SHD, because they can only identify CPDAGs. NO-ALL also uses all data to identify DAGs, so its results are better than most federated approaches. Federated methods NO-w/oAcy, NO-Avg, FedDAG and NO-ADMM are limited by decentralized data and are outperformed by centralized PC and NO-ALL. The performance of FedCausal is not only better than other federated methods, but even better than NO-ALL, which may be owed to the effective interaction between the clients and the server powered by our adaptive optimization.

### Conclusion

This paper introduces a federated approach (FedCausal) to learn a unified global causal graph from decentralized non-IID data. FedCausal uses an explainable and adaptive optimization process to coordinate clients to optimize the local causal graphs based on clients' data and to learn the global graph with ensured acyclicity. Our analysis shows that the optimization objective of FedCausal under statistically homogeneous data is consistent with that of causal discovery algorithms for centralized data, and FedCausal can flexibly learn DAGs from decentralized heterogeneous data. Experimental results confirm its effectiveness, generality and reliability on IID, and non-IID data.

## Acknowledgements

This work is supported by NSFC (No. 62072380, 62031003 and 62272276), National Key Research and Development Program of China (No. 2022YFC3502101), Taishan Scholars Project Special Funding, Xiaomi Young Talents Program, CAAI-Huawei MindSpore Open Fund.

## References

- Bernstein, D.; Saeed, B.; Squires, C.; and Uhler, C. 2020. Ordering-based causal structure learning in the presence of latent variables. In *AISTATS*, 4098–4108.
- Bühlmann, P.; Peters, J.; and Ernest, J. 2014. CAM: Causal additive models, high-dimensional order search and penalized regression. *The Annals of Statistics*, 42(6): 2526–2556.
- Byrd, R. H.; Lu, P.; Nocedal, J.; and Zhu, C. 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5): 1190–1208.
- Chai, D.; Wang, L.; Chen, K.; and Yang, Q. 2020. Secure federated matrix factorization. *IEEE Intelligent Systems*, 36(5): 11–20.
- Chickering, D. M. 2002. Optimal structure identification with greedy search. *JMLR*, 3(11): 507–554.
- Gao, E.; Chen, J.; Shen, L.; Liu, T.; Gong, M.; and Bondell, H. 2023. FedDAG: Federated DAG Structure Learning. *TMLR*.
- Gou, K. X.; Jun, G. X.; and Zhao, Z. 2007. Learning Bayesian network structure from distributed homogeneous data. In *In ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing*, volume 3, 250–254. IEEE.
- He, Y.; Cui, P.; Shen, Z.; Xu, R.; Liu, F.; and Jiang, Y. 2021. Daring: Differentiable causal discovery with residual independence. In *KDD*, 596–605.
- Koller, D.; and Friedman, N. 2009. *Probabilistic graphical models: principles and techniques*. MIT press.
- Lachapelle, S.; Brouillard, P.; Deleu, T.; and Lacoste-Julien, S. 2019. Gradient-Based Neural DAG Learning. In *ICLR*.
- Liang, J.; Wang, J.; Yu, G.; Guo, W.; Domeniconi, C.; and Guo, M. 2023. Directed Acyclic Graph Learning on Attributed Heterogeneous Network. *IEEE Transactions on Knowledge and Data Engineering*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 1273–1282.
- Na, Y.; and Yang, J. 2010. Distributed Bayesian network structure learning. In *ISIE*, 1607–1611.
- Nemirovsky, A. 1999. Optimization II. Numerical methods for nonlinear continuous optimization.
- Ng, I.; and Zhang, K. 2022. Towards federated bayesian network structure learning with continuous optimization. In *AISTATS*, 8095–8111.
- Ng, I.; Zhu, S.; Fang, Z.; Li, H.; Chen, Z.; and Wang, J. 2022. Masked gradient-based causal structure learning. In *SDM*, 424–432.
- Pearl, J. 2009. *Causality*. Cambridge university press.
- Peters, J.; Janzing, D.; and Schölkopf, B. 2017. *Elements of causal inference: foundations and learning algorithms*. MIT Press.
- Sachs, K.; Perez, O.; Pe’er, D.; Lauffenburger, D. A.; and Nolan, G. P. 2005. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721): 523–529.
- Shimizu, S. 2012. Joint estimation of linear non-Gaussian acyclic models. *Neurocomputing*, 81: 104–107.
- Spirites, P.; Glymour, C. N.; Scheines, R.; and Heckerman, D. 2000. *Causation, prediction, and search*. MIT Press.
- Spirites, P. L.; Meek, C.; and Richardson, T. S. 2013. Causal inference in the presence of latent variables and selection bias. *arXiv preprint arXiv:1302.4983*.
- Tillman, R. E.; Danks, D.; and Glymour, C. 2008. Integrating locally learned causal structures with overlapping variables. In *NeurIPS*, 1665–1672.
- Triantafillou, S.; and Tsamardinos, I. 2015. Constraint-based causal discovery from multiple interventions over overlapping variable sets. *JMLR*, 16(1): 2147–2205.
- Tsamardinos, I.; Brown, L. E.; and Aliferis, C. F. 2006. The max-min hill-climbing Bayesian network structure learning algorithm. *Mach. Learn.*, 65(1): 31–78.
- Xiong, R.; Koenecke, A.; Powell, M.; Shen, Z.; Vogelstein, J. T.; and Athey, S. 2021. Federated causal inference in heterogeneous observational data. *arXiv preprint arXiv:2107.11732*.
- Yang, D.; He, X.; Wang, J.; Yu, G.; Domeniconi, C.; and Zhang, J. 2023. Federated Causality Learning with Explainable Adaptive Optimization. *arXiv:2312.05540*.
- Ye, Q.; Amini, A. A.; and Zhou, Q. 2022. Distributed Learning of Generalized Linear Causal Networks. *arXiv preprint arXiv:2201.09194*.
- Yu, Y.; Chen, J.; Gao, T.; and Yu, M. 2019. DAG-GNN: DAG structure learning with graph neural networks. In *ICML*, 7154–7163.
- Yuan, M. 2011. On the identifiability of additive index models. *Statistica Sinica*, 1901–1911.
- Zheng, X.; Aragam, B.; Ravikumar, P.; and Xing, E. P. 2018. DAGs with NO TEARS: continuous optimization for structure learning. In *NeurIPS*, 9492–9503.
- Zheng, X.; Dan, C.; Aragam, B.; Ravikumar, P.; and Xing, E. 2020. Learning sparse nonparametric dags. In *AISTATS*, 3414–3425.