

Multi-Granularity Causal Structure Learning

Jiaxuan Liang^{1, 2}, Jun Wang², Guoxian Yu^{1, 2, *}, Shuyin Xia^{3, 4}, Guoyin Wang^{3, 4}

¹School of Software, Shandong University, Jinan, China

²SDU-NTU Joint Centre for AI Research, Shandong University, Jinan, China

³Chongqing Key Laboratory of Computational Intelligence, Chongqing Uni. of Posts and Telecom., Chongqing, China

⁴MOE Key Laboratory of Big Data Intelligent Computing, Chongqing Uni. of Posts and Telecom., Chongqing, China

jxliang@mail.sdu.edu.cn, {kingjun, gxyu}@sdu.edu.cn, {xiasy, wanggy}@cqupt.edu.cn

Abstract

Unveiling, modeling, and comprehending the causal mechanisms underpinning natural phenomena stand as fundamental endeavors across myriad scientific disciplines. Meanwhile, new knowledge emerges when discovering causal relationships from data. Existing causal learning algorithms predominantly focus on the isolated effects of variables, overlook the intricate interplay of multiple variables and their collective behavioral patterns. Furthermore, the ubiquity of high-dimensional data exacts a substantial temporal cost for causal algorithms. In this paper, we develop a novel method called MgCSL (Multi-granularity Causal Structure Learning), which first leverages sparse auto-encoder to explore coarse-graining strategies and causal abstractions from micro-variables to macro-ones. MgCSL then takes multi-granularity variables as inputs to train multilayer perceptrons and to delve the causality between variables. To enhance the efficacy on high-dimensional data, MgCSL introduces a simplified acyclicity constraint to adeptly search the directed acyclic graph among variables. Experimental results show that MgCSL outperforms competitive baselines, and finds out explainable causal connections on fMRI datasets.

Introduction

Data science is moving from the data-centric paradigm forward the science-centric paradigm, and causal revolution is sweeping across various research fields. Causality learning endeavors to unearth causal relationships among variables from observational data and generate causal graph, that is, directed acyclic graph (DAG). Unlike correlation-based study, causality analysis reveals the causal mechanism of data generation. Identifying causality holds paramount significance for stable inference and rational decisions in many applications, such as recommendation systems (Wang et al. 2020), medical diagnostics (Richens, Lee, and Johri 2020), epidemiology (Vandenbroucke, Broadbent, and Pearce 2016) and many others (Von Kügelgen et al. 2022).

For its significance, numerous studies have been conducted toward causal structure learning. Constraint-based algorithms (Spirtes et al. 2000; Colombo, Maathuis et al. 2014; Marella and Vicard 2022) acquire a set of causal

graphs that satisfy the conditional independence (CI) inherent in the data. Nevertheless, the faithfulness assumption can be refuted, and a substantial number of CI tests are needed. Score-based algorithms (Chickering 2002; Hauser and Bühlmann 2012; Ramsey et al. 2017) define a structure scoring function in conjunction with search strategies to explore the DAG that best fits the limited data. However, these algorithms cannot differentiate DAGs that belong to the same Markov equivalence class (MEC). By virtue of additional assumptions on data distribution and functional classes, functional causal model-based methods (Hoyer et al. 2008; Zhang and Hyvärinen 2009) can distinguish DAGs within the same MEC, but the exhaustive and heuristic search of DAGs encounters the challenge of combinatorial explosion with the increasing number of variables. NOTEARS (Zheng et al. 2018) formulates the acyclicity constraint as a differentiable equation and applies efficient numerical solvers to search DAG. Subsequent efforts focus on nonlinear extension (Lachapelle et al. 2020; Zheng et al. 2020), optimization technique (Wang et al. 2021; Yang et al. 2023), robustness (He et al. 2021) and semi-structured data (Liang et al. 2023b).

However, these algorithms simply deem causal relationships stand exclusively at the level of individual variables (*micro-variable*), ignoring the collective interactions from multiple variables (*macro-variable*). For instance, the brain can be characterized at a micro granularity of neurons and their synapses, but high-order synergistic subsystems are widespread, which typically sit between canonical functional networks and may serve an integrative role (Varley et al. 2023). Actually, observational data can be regarded as knowledge in the lowest granularity level, while knowledge can be regarded as the abstraction of data at different granularity levels (Wang 2017; Wang et al. 2022). Similar viewpoints appear in the research of complex systems, which suggests that causal relationship is more pronounced at the macro-scale of a system than at its micro-scale. This phenomenon, widely known as *causal emergence*, manifests extensively across scientific domains such as physics (Loewer 2012), physiology (Noble 2012), sociology (Elder-Vass 2012), and beyond. Due to the intricacy of macro-level causality, available algorithms often overlook or misinterpret the underlying causal structure. Furthermore, the high complexity of causal discovery algorithms causes a signifi-

*Corresponding author.

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

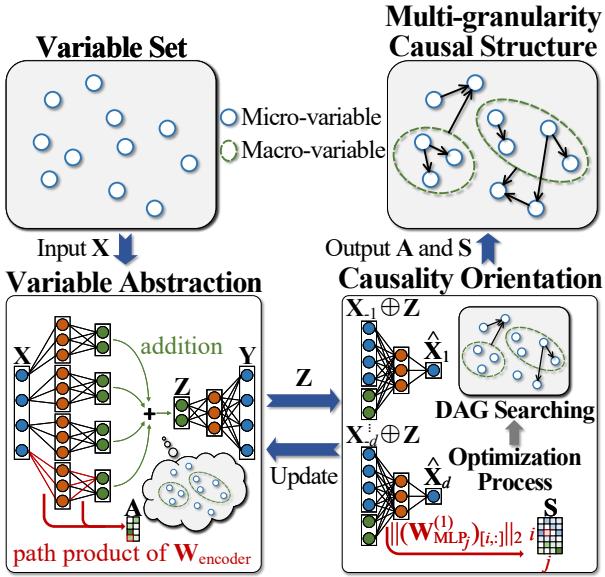


Figure 1: Framework overview of MgCSL, which takes observational data \mathbf{X} as inputs and constructs an SAE to explore coarse-grained strategies and causal abstractions. Within the encoder, each input variable is trained individually, and their encoded representations are summed to obtain the latent macro-variable representation \mathbf{Z} , which is then used by the decoder to reconstruct the observational data \mathbf{X} . The contribution matrix \mathbf{A} from micro-variables to macro-variables is extracted from the path product of the encoder parameters $\mathbf{W}_{\text{encoder}}$. Next, MgCSL feeds the concatenation of micro- and macro-variables $\mathbf{X} \oplus \mathbf{Z}$ into each MLP to explore potential causal relationships, and collects the multi-granularity weighted adjacency matrix \mathbf{S} from the first parameter matrix of MLPs $\mathbf{W}_{\text{MLP}}^{(1)}$. MgCSL formulates the DAG search as an optimization problem with a simplified acyclicity constraint, and induces the multi-granularity causal structure from \mathbf{S} and \mathbf{A} .

cant efficacy drop when dealing with high-dimensional data, hindering their practical implementations.

We aim to learn multi-granularity causal structure and propose an approach called MgCSL, as depicted in Figure 1. MgCSL firstly establishes a sparse autoencoder (SAE) (Ng et al. 2011) to automatically coarse-grain the micro-variables into latent macro-ones. Next, it constructs a multiple layer perceptron (MLP) for each micro-variable, taking both the micro- and macro-variables as inputs to explore the underlying causal mechanisms. It further introduces a simplified acyclicity constraint to efficiently orient edges. The major contributions of our work are:

(i) Diverging from existing causal learning methods that exclusively focus on causality at the micro-level of individual variables, we pioneer to learn multi-granularity causal structure across micro- and macro-variables, which is more complex and challenging but holds immense practical values in various scenarios.

(ii) Our proposed MgCSL utilizes SAE with sparsity term

on the encoder to explore potential macro-variables and extract effective coarse-graining strategies in a sensible way. MgCSL leverages MLPs to model the underlying causal mechanisms and introduces a simplified acyclicity penalty to orient edges, searching causal structure in an efficacy way. (iii) Experimental results confirm the advantages of MgCSL over competitive baselines (Spirtes et al. 2000; Chickering 2002; Yu et al. 2019; Ng et al. 2019; Lachapelle et al. 2020; Zheng et al. 2020), and it identifies multi-granularity causal connections from fMRI data.

Related Work

Causal structure learning is an indispensable and intricate task pervading in various scientific fields. Existing causal learning algorithms can be grouped into two types: constraint-based and score-based. Constraint-based methods (Spirtes et al. 2000; Bird and Burgess 2008; Marella and Vicard 2022) perform CI tests to obtain causal skeleton, then orient edges via elaborate rules to meet the requirements of DAG. Score-based solutions (Chickering 2002; Hauser and Bühlmann 2012; Ramsey et al. 2017) leverage the score function and search strategies to find the graph yielding the highest score. NOTEARS (Zheng et al. 2018) recasts the combinatoric graph search problem as a continuous optimization problem, stimulating a proliferation of literature. DAG-GNN (Yu et al. 2019) and GAE (Ng et al. 2019) extend NOTEARS to nonlinear cases using autoencoder. GraN-DAG (Lachapelle et al. 2020) and NOTEARS-MLP (Zheng et al. 2020) employ MLPs to approximate the underlying data generation functions. DARING (He et al. 2021) imposes residual independence constraint in an adversarial way to facilitate DAG learning. RCL-OG (Yang et al. 2023) uses order graph instead of Markov chain Monte Carlo to approximate the posterior distribution of DAGs. However, these causal solutions only focus on the causality among micro-variables, can not cope with multiplex causality at different granularity levels. Our MgCSL bridges this gap by leveraging SAE to explore multi-granularity causality, and gradient-based search with simplified acyclicity constraint to seek DAG in an efficacy way.

Actually, the behavior and properties of a system are regulated by causal relations at different granularity levels. The micro-level causality cannot reveal the complexity of a system, especially for the cross-regulations among different levels. Moreover, the extracted macro-variables represent a form of knowledge that unveils higher-level characteristics and collective behavioral patterns within a system. Inspired by the human cognitive principle ‘from coarser to finer’, granular cognitive computing aims to process information at various granularity levels (Wang 2017; Wang et al. 2022). Causal learning from the perspective of granular computing has received much less attention, compared with that from individual micro-variables. MaCa (Yang et al. 2022) models the text with multi-granularity way and uses a matrix capsule network to cluster the emotion-cause pairs. MAGN (Qiang et al. 2023) performs backdoor adjustment based on structural causal model to achieve cross-granularity few-shot learning. EMGCE (Wu et al. 2023) represents sentences at different granulation layers to extract explicit and implicit

causal triplets of words or phrases. While MgCSL focuses on exploring the multi-granularity causal graph among micro- and macro-variables.

The Proposed MgCSL Algorithm

Problem Definition

In this section, we introduce main concepts used in this paper, followed by a formal definition of our problem.

Definition 1 (Structural Causal Model (SCM)). *The SCM is defined on a set of variables $\mathcal{V}=(v_1, v_2, \dots, v_d)$, and consists of a causal graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ along with structural equations. Each edge $(i, j) \in \mathcal{E}$ represents a direct causal relation from v_i to v_j . v_i represents a micro-variable, with its observational data denoted as \mathbf{x}_i . Multiple micro-variables can be abstracted into a macro-variable u_i , with its representative data denoted as \mathbf{z}_i . The distribution P is said to be Markov with respect to \mathcal{G} , allowing the joint probability P to be decomposed into the product of conditional probabilities as $P(\mathbf{X})=\prod_{i=1}^d P(\mathbf{x}_i|\mathbf{x}_{pa(i)})$, where $\mathbf{x}_{pa(i)}$ acts as the set of parents of \mathbf{x}_i . We assume that the distribution P is entailed by Additive Noise Models (ANMs) (Peters et al. 2014) of the form:*

$$\mathbf{x}_i = f_i(\mathbf{x}_{pa(i)}) + N_i \quad (1)$$

where f_i is a nonlinear function that denotes the generative process of \mathbf{x}_i , N_i is the external noise of \mathbf{x}_i .

Definition 2 (Directed acyclic graph (DAG)). *A DAG consists of variables and edges, with each edge directed from one variable to another. For instance, $v_i \rightarrow v_j$ means there is a directed edge from v_i to v_j and hence v_i is a direct cause of v_j . A path between v_i and v_j in a DAG is a sequence of edges. The ending variable of each edge acts as the starting variable of the next edge in the sequence. If there is no path from any variable to itself, then the graph is acyclic. In this paper, we further extend the concept of DAG to encompass multi-granularity causal structures, wherein if there exists an edge $u_i \rightarrow v_j$, there should be no path from v_j to any micro-variable composing u_i .*

Problem 1. *Given observational data $\mathbf{X}=(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_d)$, our task is to uncover multi-granularity causal DAG. This DAG can reflect the collective behavior of variable clusters and reveal complex causal interactions.*

Variable Abstraction

Abstract descriptions provide the foundation for system interventions and explanations for observed phenomena at a coarser granularity level than the most fundamental account of the system. While the emerging causal representation learning aim to reconstruct disentangled causal variable representations from unstructured data (e.g. images) (Schölkopf et al. 2021), our intention is to learn from observational data and generate low-dimensional representations of macro-variables, which typically encapsulate high-order interactions and collective behaviors among multiple micro-ones. SAE is a unsupervised learning neural network that excels in extracting compact and meaningful representations of data. By promoting sparsity, it encourages the activation of only a small subset of neurons in the hidden layer, which

not only reduces the data dimensionality but also enhances interpretability.

Given that, MgCSL includes an SAE-based module to gain causal abstraction from micro-variables to macro-ones. However, conventional SAE uses a shared encoder for all inputs, entangling the contributions of micro-variables to macro-ones. To ensure the independence of micro inputs during the coarse-graining process, we construct d encoders to separately encode each input, and define the encoder and decoder as:

$$\begin{aligned} \mathbf{Z} &= \sum_{i=1}^d \sigma(\sigma(\mathbf{X}_{[:,i]}(\mathbf{W}_{\text{encoder}}^{(1)})_{[:,i]})\mathbf{W}_{\text{encoder}}^{(2)})_{[:,i]} \\ \mathbf{Y} &= \sigma(\sigma(\mathbf{Z}\mathbf{W}_{\text{decoder}}^{(1)})\mathbf{W}_{\text{decoder}}^{(2)}) \end{aligned} \quad (2)$$

where $\mathbf{Z}=(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_q)$ represents the encoded data of the latent macro-variables, σ is an activation function, and biases are omitted for clarity. $\mathbf{W}_{\text{encoder}}^{(1)} \in \mathbb{R}^{d \times 1 \times m_1}$, $\mathbf{W}_{\text{encoder}}^{(2)} \in \mathbb{R}^{d \times m_1 \times q}$, $\mathbf{W}_{\text{decoder}}^{(1)} \in \mathbb{R}^{q \times m_1}$, $\mathbf{W}_{\text{decoder}}^{(2)} \in \mathbb{R}^{m_1 \times d}$ are parameters of encoder and decoder, m_1 is the number of neurons in the first hidden layer. Although we consider all micro-variables as inputs, in our coarse-grained strategies, not every micro-variable contributes to the generation of macro-variable representations. To determine which micro-variables make contribution, we define the contribution as the path product on the encoder:

$$\mathbf{A} = |\mathbf{W}_{\text{encoder}}^{(1)}||\mathbf{W}_{\text{encoder}}^{(2)}| \quad (3)$$

We say that a path is inactive if at least one weight along the path is zero. Therefore, when the path product from input $i=1, 2, \dots, d$ to output $j=1, 2, \dots, q$ is non-zero, i.e. $\mathbf{A}_{i,j} \neq 0$, v_i is one of the constituents of macro-variable u_j . To clarify the composition of macro-variables from micro-ones, we introduce an $l_{1,1}$ -norm regularization on the parameter matrices of the encoder, which encourages sparsity of \mathbf{A} . This allows some $\mathbf{z}_i \in \mathbf{Z}$ to approach the zero vector, achieving dynamic variability in the number of macro-variables. In addition, we incorporate a data reconstruction loss, which helps retain causal interpretation. In other words, the compressed macro-variables preserve a significant portion of essential causal information, allowing for the faithful reconstruction of the original data. Then the loss function of variable abstraction module is defined as follows:

$$\mathcal{L}_1 = \mathcal{J}(\mathbf{X}, \mathbf{Y}) + \alpha_1 \|\mathbf{W}_{\text{encoder}}\|_{1,1} \quad (4)$$

where $\mathcal{J}(\cdot)$ denotes the cost function, here we use mean squared error, that is $\mathcal{J}(\mathbf{X}, \mathbf{Y}) = \frac{1}{2n} \|\mathbf{X} - \mathbf{Y}\|_F^2$, n is the number of samples. The usage of the regularization term may lead to a reduction of variables contributing to abstractions. In such cases, to ensure approximate data reconstruction, SAE prefers preserving the inputs of parent variables, as they encode more information than child ones. As above, MgCSL extracts the macro-representation vectors \mathbf{Z} from \mathbf{X} , which will be used for subsequent multi-granularity causality orientation.

Causality Orientation

How to orient causal edges between variable nodes is a critical task in causal discovery. Early approaches determine the

orientation based on specific structures or asymmetry. The differentiable acyclicity constraint of NOTEARS (Zheng et al. 2018) enables the usage of standard numerical techniques to search DAGs. However, NOTEARS and its variants are typically operated at the micro-level. In this work, we need to identify nonlinear dependencies among variables and enforce acyclicity across multiple granularities. For this purpose, we propose a simplified acyclicity constraint to improve search efficiency. We first construct an MLP for each micro-variable, taking other variables and latent macro-variables as inputs, to model the underlying causal mechanisms:

$$\hat{\mathbf{x}}_i = \text{MLP}_i(\mathbf{X}_{-i} \oplus \mathbf{Z}; \mathbf{W}_{\text{MLP}_i}) \quad (5)$$

where \oplus is the concatenation operator, $\mathbf{W}_{\text{MLP}_i}$ is parameters of the i -th MLP, \mathbf{X}_{-i} represents the observational data of all variables except \mathbf{x}_i . Actually, the variables involved in the fitting process of \mathbf{x}_i are potential parent variables of v_i . Therefore, we extract the weighted adjacency matrix (WAM) \mathbf{C} on the first layer of $\mathbf{W}_{\text{MLP}_i}$, taking into account the macro-variables:

$$\mathbf{C}_{i,j} = \|(\mathbf{W}_{\text{MLP}_j}^{(1)})_{[i,:]} + \mathbf{A}_{[i,:]}(\mathbf{W}_{\text{MLP}_j}^{(1)})_{[d+1,:,:]}\|_2 \quad (6)$$

where $\mathbf{W}_{\text{MLP}_j}^{(1)} \in \mathbb{R}^{(d+q) \times m_2}$ represents the first parameter matrix of j -th MLP, m_2 is the number of neurons in the first hidden layer. Consequently, we can collectively enforce acyclicity on the multi-granularity variables. Similarly, we can obtain multi-granularity WAM $\mathbf{S} \in \mathbb{R}^{(d+q) \times d}$ as:

$$\mathbf{S}_{i,j} = \|(\mathbf{W}_{\text{MLP}_j}^{(1)})_{[i,:]}\|_2 \quad (7)$$

When a variable participates in joint interactions on other variables, we reduce its individual weight on other variables to exclude redundant causal relations by a redundancy penalty as:

$$\begin{aligned} \mathcal{L}_{\text{red}} = \sum_{j=1}^d \left\| \sum_{k=1}^{m_2} (|\mathbf{A}| \|\mathbf{W}_{\text{MLP}_j}^{(1)}|_{[d+1,:,:]})_{[:,k]} \circ \right. \\ \left. \sum_{k=1}^{m_2} (|\mathbf{W}_{\text{MLP}_j}^{(1)}|_{[:,k]})_{[:,k]} \right\|_1 \end{aligned} \quad (8)$$

where \circ denotes the Hadamard product.

Then we can define the loss function of causality orientation module as follows:

$$\begin{aligned} \min \quad & \mathcal{L}_2 = \mathcal{J}(\mathbf{X}, \hat{\mathbf{X}}) + \mathcal{L}_{\text{red}} + \\ & \alpha_2 \left(\sum_{j=1}^d \|\mathbf{W}_{\text{MLP}_j}^{(1)}\|_{1,1} + \frac{1}{2} \sum_{j=1}^d \|\mathbf{W}_{\text{MLP}_j}^{(1)}\|_F^2 \right) \\ \text{s.t.} \quad & \mathcal{H}(\mathbf{C}) = 0 \end{aligned} \quad (9)$$

$\mathcal{H}(\mathbf{C}) = 0$ is our meticulously simplified directed acyclicity constraint, which will be discussed in more detail later. As sparse parameters lead to more specific direction, we apply an $l_{1,1}$ -norm regularization on each $\mathbf{W}_{\text{MLP}_j}^{(1)}$ to help orientation. We further employs squared Frobenius norm on each $\mathbf{W}_{\text{MLP}_j}^{(1)}$ to facilitate model convergence and enhance generalization. The objective function of MgCSL consists of the two afore-defined loss functions:

$$\begin{aligned} \min \quad & \mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 \\ \text{s.t.} \quad & \mathcal{H}(\mathbf{C}) = 0 \end{aligned} \quad (10)$$

By iteratively optimizing the objective function, we can search for the optimal multi-granularity DAG. Due to the continuous nature of \mathbf{C} and \mathbf{S} obtained after training, it cannot be directly interpreted as a causal graph. Therefore, post-processing procedure on \mathbf{C} and \mathbf{S} is required. We first cut all weights below a certain threshold ϵ , then we iterative eliminate the smallest weight until no directed cycles exist in \mathbf{C} and \mathbf{S} . Finally, we set all remaining non-zero elements to 1.

Optimization

To ensure the acyclicity of \mathbf{C} , an explicit way is to enforce the following condition: $\sum_{k=1}^d \text{tr}(\mathbf{D}^k) = 0$ where $\mathbf{D} = \mathbf{C} \circ \mathbf{C}$ and $\text{tr}(\cdot)$ denotes the trace of \mathbf{D} . NOTEARS (Zheng et al. 2018) suggests that $\sum_{k=1}^d \text{tr}(\mathbf{D}^k) = 0 \Rightarrow \text{tr}(\sum_{k=1}^d \frac{\mathbf{D}^k}{k!}) = 0 \Rightarrow \text{tr}(\sum_{k=0}^d \frac{\mathbf{D}^k}{k!}) - d = 0$, and proposes Theorem 1.

Theorem 1. A WAM $\mathbf{C} \in \mathbb{R}^{d \times d}$ is a DAG if and only if

$$h(\mathbf{C}) = \text{tr}(e^{\mathbf{D}}) - d = 0 \quad (11)$$

This provides a continuous optimization approach to search for the graph structure. Indeed, the computational complexity of matrix exponential makes it challenging to obtain results within an acceptable time on high-dimensional data (Liang et al. 2023a). Here, we introduce an meticulously designed acyclicity constraint to orient edges in an efficacy way.

Proposition 1. A WAM $\mathbf{C} \in \mathbb{R}^{d \times d}$ is a DAG if and only if

$$\forall i \in [1, d], \lambda_i = 0 \quad (12)$$

where λ_i represents the eigenvalues of \mathbf{D} .

The proof of Proposition 1 is deferred into the Supplementary file (Liang et al. 2023c). Previous effort (Lee et al. 2019) attempts to enforce acyclicity based on the spectral radius, which is unstable and tends to a zero matrix, due to its neglect of global information of eigenvalues and an easily attainable local optimum of the zero matrix. Optimizing the eigenvalues jointly expands the search space but introduces computational burdens. Since a square matrix is necessarily unitarily similar to an upper triangular matrix, we perform *Schur decomposition* on \mathbf{D} :

$$\mathbf{D} = \mathbf{U} \mathbf{R} \mathbf{U}^\top \quad (13)$$

where \mathbf{U} is a unitary matrix, \mathbf{R} is a upper triangular matrix that has the eigenvalues on the diagonal, which are also the eigenvalues of \mathbf{D} , since \mathbf{D} is similar to \mathbf{R} .

Proposition 2. A WAM $\mathbf{C} \in \mathbb{R}^{d \times d}$ is a DAG if and only if

$$\mathcal{H}(\mathbf{C}) = \|\text{diag}(\mathbf{R})\|_2^2 = 0 \quad (14)$$

where $\text{diag}(\cdot)$ denotes the diagonal vector of a matrix.

Optimizing Eq. (14) avoids the intricate computation of matrix exponential and its gradient in Eq. (11), which facilitate an efficient search of causal graphs in the latent graph space. It can also minimize the weights of non-existing edges with the help of reconstruction loss on \mathbf{X} . Therefore, the error caused by transforming \mathbf{D} to \mathbf{R} can be effectively removed through our post-processing procedure.

To this end, Eq. (10) is reformulated into a continuous optimization problem, albeit non-convex due to the non-convex feasible set. Nevertheless, we can employ an augmented Lagrangian approach to replace the original constrained problem Eq. (10) with a sequence of unconstrained subproblems, and then seek stationary points as the solution:

$$\begin{aligned} \min \quad & \mathcal{L} = \mathcal{J}(\mathbf{X}, \mathbf{Y}) + \alpha_1 \|\mathbf{W}_{\text{encoder}}\|_{1,1} + \\ & \mathcal{J}(\mathbf{X}, \hat{\mathbf{X}}) + \mathcal{L}_{\text{red}} + \alpha_2 \left(\sum_{j=1}^d \|\mathbf{W}_{\text{MLP}_j}^{(1)}\|_{1,1} + \right. \\ & \left. \frac{1}{2} \sum_{j=1}^d \|\mathbf{W}_{\text{MLP}_j}\|_F^2 \right) + \frac{\mu}{2} |\mathcal{H}(\mathbf{C})|^2 + \gamma \mathcal{H}(\mathbf{C}) \end{aligned} \quad (15)$$

where μ is the penalty coefficient and γ is the Lagrange multiplier. The parameter update process is as follows:

$$\begin{aligned} \theta^{(\kappa+1)} \leftarrow \arg \min \{ & \mathcal{J}(\mathbf{X}, \mathbf{Y}) + \alpha_1 \|\mathbf{W}_{\text{encoder}}\|_{1,1} + \\ & \mathcal{J}(\mathbf{X}, \hat{\mathbf{X}}) + \mathcal{L}_{\text{red}} + \alpha_2 \left(\sum_{j=1}^d \|\mathbf{W}_{\text{MLP}_j}^{(1)}\|_{1,1} + \right. \\ & \left. \frac{1}{2} \sum_{j=1}^d \|\mathbf{W}_{\text{MLP}_j}^{(\kappa)}\|_F^2 \right) + \frac{\mu}{2} |\mathcal{H}(\mathbf{C})|^2 + \gamma \mathcal{H}(\mathbf{C}) \} \end{aligned} \quad (16)$$

$$\mu^{(\kappa+1)} \leftarrow \begin{cases} \eta \mu^{(\kappa)}, & \mathcal{H}(\mathbf{C}^{(\kappa+1)}) > \rho \mathcal{H}(\mathbf{C}^{(\kappa)}) \\ \mu^{(\kappa)}, & \text{otherwise} \end{cases} \quad (17)$$

$$\gamma^{(\kappa+1)} \leftarrow \gamma^{(\kappa)} + \mu \mathcal{H}(\mathbf{C}) \quad (18)$$

where $\theta \triangleq \{\mathbf{W}_{\text{encoder}}, \mathbf{W}_{\text{decoder}}, \mathbf{W}_{\text{MLP}}\}$, and $\theta^{(\kappa)}$ is the iterative optimized θ in the κ -th iteration. A variety of optimization methods can be applied to Eq. (16). In this work, we apply the L-BFGS-B algorithm (Byrd et al. 1995). The comprehensive procedure of MgCSL is delineated in Algorithm 1 in the Supplementary file.

The time complexity of MgCSL comes from two sources, variable abstraction and causality orientation. For the former, the computational complexity of SAE for one iteration step is $\mathcal{O}(nm_1q + ndm_1q + nd^2m_1)$. For the latter, the computational complexity of MLP for one iteration step is $\mathcal{O}(d^2m_2 + nd^2m_2 + d^3)$. Although the *Schur decomposition* of MgCSL and the trace of matrix exponential in NOTEARS share the same time complexity of $\mathcal{O}(d^3)$, eigenvalues are more sensitive to the presence of directed cycles in \mathbf{C} than the trace of matrix exponential, enabling \mathbf{C} to be optimized faster towards a cycle-free direction. This is because under the constraint of the squared Frobenius norm, most entries of \mathbf{C} are smaller than 1, resulting in the trace of matrix exponential approaching zero. Therefore, our approach is approximately one order of magnitude faster, as shown in the experiment section.

Experimental Evaluation

Experimental Setup

We evaluated the proposed MgCSL on random DAG produced by *Erdős-Rényi* (ER) or *Scale-Free* (SF) scheme. We varied the number of variables ($d \in \{20, 50, 100\}$) with edge density (degree=2). For each graph, we generate 10 datasets of $n=1000$ samples following: (i) Additive models with Gaussian processes, and (ii) ANM with Gaussian processes. In addition, on the SF graph with $d=20$, we randomly select $\{2, 4\}$ variables as macro-variables and employ MLP

to decompose them into 8 micro-variables, forming multi-granularity graphs. We consider another well-known dataset Sachs (Sachs et al. 2005), which measures the level of different expressions of proteins and phospholipids in human cells. The ground truth causal graph of this dataset consists of 11 variables and 20 edges. In this work, we test on observational data with 7466 samples.

We compare MgCSL with representative causal structure learning methods, including PC (Spirtes et al. 2000), GES (Chickering 2002), DAG-GNN (Yu et al. 2019), GAE (Ng et al. 2019), GraN-DAG (Lachapelle et al. 2020) and NOTEARS-MLP (NO-MLP) (Zheng et al. 2020). The first two methods focus on combinatorial optimization, and the last four target at general nonlinear dependencies. We refer readers to the Supplementary file (Liang et al. 2023c) for the detailed experimental setup. The code of MgCSL is shared at <http://www.sdu-idea.cn/codes.php?name=MgCSL>.

We evaluate the estimated DAG using five metrics: Precision: proportion of correctly estimated edges to the total estimated edges; Recall: proportion of correctly estimated edges to the total edges in true graph; F1 score: the harmonic average of precision and recall; Structural hamming distance (SHD): the number of missing, falsely estimated or reversed edges; Runtime: the running time required to obtain the results, measured in seconds. When a method fails to return any result within a reasonable response time (100 hours), we mark the entry as '-'. On the case of multi-granularity graphs, we convert the estimated graph and the ground truth graph into the micro-level, i.e., an edge from a macro-variable to a target implies multiple edges from micro-variables encompassed within that macro-one to the same target. For the reported results, $\uparrow(\downarrow)$ means the higher (lower) the value, the better the performance is. The best result is highlighted in **bold** font.

Result Analysis

Results on multi-granularity graph We first compare the performance of our MgCSL with baselines on multi-

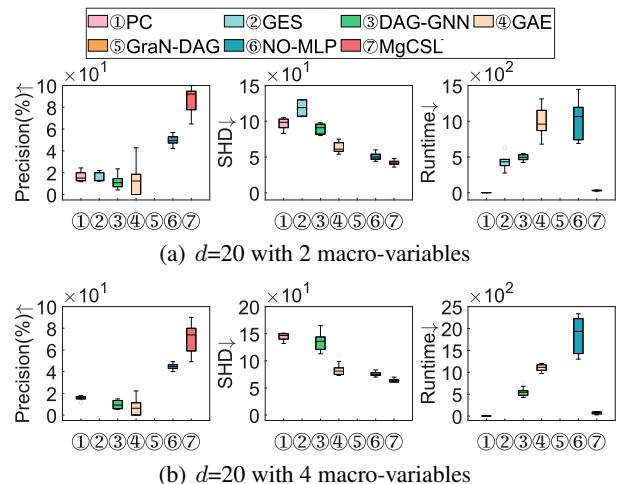


Figure 2: Results on multi-granularity graph.

		Erdős-Rényi graph				Scale-Free graph			
d		Precision(%) \uparrow	Recall(%) \uparrow	SHD \downarrow	Runtime \downarrow	Precision(%) \uparrow	Recall(%) \uparrow	SHD \downarrow	Runtime \downarrow
20	PC	60.22 \pm 8.38	54.50 \pm 7.43	25.90 \pm 4.15	0.54	41.20 \pm 9.26	42.70 \pm 9.60	33.70 \pm 5.23	0.69
	GES	58.14 \pm 5.80	59.00 \pm 6.15	26.00 \pm 3.06	12.28	43.72 \pm 5.88	55.68 \pm 6.39	33.60 \pm 3.37	14.75
	DAG-GNN	89.57 \pm 8.44	42.50 \pm 6.01	24.70 \pm 3.09	2512.59	85.02 \pm 11.23	38.92 \pm 5.29	23.90 \pm 2.64	2880.47
	GAE	83.24 \pm 12.07	30.50 \pm 17.59	29.30 \pm 5.42	1469.75	83.93 \pm 22.12	37.03 \pm 19.49	24.60 \pm 7.97	1999.70
	GraN-DAG	97.13 \pm 5.20	51.25 \pm 6.80	20.10 \pm 2.85	311.23	95.12 \pm 3.69	59.19 \pm 12.32	16.10 \pm 4.46	355.68
100	NO-MLP	89.86 \pm 4.59	69.75\pm7.12	14.80 \pm 3.05	45.35	80.62 \pm 5.62	73.24\pm5.32	15.30 \pm 2.87	47.10
	MgCSL	98.61\pm2.30	65.75 \pm 8.42	14.10\pm3.21	9.57	96.19\pm7.36	70.54 \pm 7.48	12.00\pm3.59	7.58
	PC	31.06 \pm 2.79	40.60 \pm 3.16	235.30 \pm 13.67	12.74	27.39 \pm 2.23	38.12 \pm 2.61	263.00 \pm 13.43	27.15
	GES	42.10 \pm 3.77	60.78 \pm 5.01	213.44 \pm 18.53	79857.07	-	-	-	-
	DAG-GNN	88.61 \pm 4.75	34.95 \pm 6.12	135.50 \pm 10.05	1637.52	88.80 \pm 4.10	31.22 \pm 5.38	140.10 \pm 8.17	3654.05
	GAE	50.15 \pm 42.26	13.80 \pm 10.42	211.20 \pm 55.11	5527.23	64.56 \pm 35.86	15.94 \pm 8.78	203.00 \pm 62.33	1991.81
	GraN-DAG	11.70 \pm 3.62	33.00 \pm 9.84	653.50 \pm 208.96	2089.45	6.31 \pm 1.10	23.60 \pm 5.11	828.40 \pm 122.29	1575.67
	NO-MLP	80.77 \pm 4.88	64.50\pm5.13	95.80\pm12.69	1923.61	69.28 \pm 4.38	62.44\pm3.40	122.60 \pm 12.27	1688.14
	MgCSL	94.61\pm3.33	49.65 \pm 7.18	106.10 \pm 11.53	430.81	89.20\pm4.97	55.28 \pm 7.36	98.20\pm8.09	394.56

Table 1: Results on *nonlinear models with additive Gaussian processes*.

granularity synthetic datasets. As shown in Figure 2, the presence of macro-variables has an impact on the precision of the baselines. Even on small graphs, they exhibit a precision as low as 0.5 or even lower. PC and GES achieve the highest SHD, because they struggle to capture complex nonlinear relationships through either CI tests or score functions. GES even fails to produce any results when the number of macro-variables is 4. DAG-GNN, GAE and GraN-DAG possess the capability to learn nonlinear dependencies, but they fail to achieve satisfactory results. In particular, GraN-DAG cannot obtain valid estimated graphs across all datasets. This limitation might be attributed to the introduction of noise through the micro-level representations of macro-variables, which affects the DAG search process. NO-MLP achieves promising results in terms of SHD. However, its low precision and high execution time still make it challenging to directly apply in multi-granularity graphs. MgCSL is capable of extracting valuable information from data mixed with multi-granularity variables for DAG learning, thereby achieving the optimal performance in terms of precision and SHD with a modest expenditure of time. We apply the Wilcoxon signed-rank test (Demšar 2006) to check the statistical difference between MgCSL and other compared methods across these metrics and datasets, all the p -value are smaller than 0.02. This demonstrates its ability to discover multi-granularity causal structures.

Results on typical causal graph We also conduct experiments on typical causal discovery synthetic datasets. As shown in Table 1 and S3 in the Supplementary file, our MgCSL outperforms its opponents across most metrics within a short period of time. Due to page limit, we have moved the results of 50 variables and F1 metric to the Supplementary file. This does not affect the conclusions.

MgCSL vs. combinatorial optimization methods: PC can swiftly produce comparable results through CI tests with a few variables. However, as the number of variables increases, the limited sample size renders CI tests unreliable, leading to a sharp rise of SHD. The performance of GES is relatively superior to that of PC because it is less affected by data issues (e.g., noisy or small samples). However, its

adoption of the greedy search strategy makes it challenging to deliver results within a reasonable runtime, particularly when confronted with the exponential growth of the search space as the number of variables increases. In comparison, our MgCSL manifests more remarkable performance across these metrics on datasets of different sizes.

MgCSL vs. gradient-based methods: DAG-GNN achieves consistent results in most cases by leveraging the power of variational inference. However, it encounters a substantial time cost on low-dimensional data, which hinders its applicability to small graphs. GAE presents poor performance, with its precision fluctuating significantly. While GraN-DAG exhibits excellent precision on small graphs, its performance falters as the number of variables increases. Its estimated graphs contain numerous erroneous edges, indicating its limitation in learning causal structures on large graphs. NO-MLP shows commendable performance, but it demands a lot of time cost due to the high complexity involved in matrix exponential calculations. MgCSL surpasses the above rivals across most metrics, maintaining a high precision even on high-dimensional graphs. Moreover, the simplified acyclicity penalty empowers MgCSL to provide effective results within shorter time, ensuring its feasibility for real-world applications.

The impact of different graphs: We further compared the performance of the algorithms on ER and SF graphs. On ER graphs, the degrees of individual variables are relatively balanced, whereas SF graphs contain a few hub variables with degrees significantly higher than the average, making it more likely for multiple variables to jointly influence a single target. As the experimental results show, most algorithms such as PC, GES, and NO-MLP perform worse on SF graphs than that on ER graphs in terms of SHD, especially as the number of variables increases. This could be attributed to the causal discovery algorithm’s demand for sparsity in the estimated graph, resulting in poor performance when dealing with graphs with highly uneven degrees. MgCSL has the capability of abstracting multiple micro-variables into a macro-one and discovering causal edges among multi-granularity variables. This empowers it to achieve competi-

	Precision(%) \uparrow	Recall(%) \uparrow	F1(%) \uparrow	SHD \downarrow
PC	37.04	50.00	42.55	24
GES	23.08	45.00	30.51	29
DAG-GNN	26.92	35.00	30.43	29
GAE	28.57	10.00	14.81	20
GrAN-DAG	50.00	10.00	16.67	18
NO-MLP	16.67	10.00	12.50	22
MgCSL	100.00	30.00	46.15	14

Table 2: Results on Sachs’s protein signaling dataset.

tive results on SF graphs.

Results on real-world data To further verify the effectiveness of MgCSL, we conduct experiments on Sachs’s (Sachs et al. 2005) protein signaling dataset. The results in Table 2 show that MgCSL still exhibits to a competitive performance compared with the rivals. PC, GES and DAG-GNN identify quit a lot of correct edges but at the cost of precision, resulting in the highest SHD. GAE and NO-MLP miss too many edges, leading to a low recall. The sparse inference graph of GrAN-DAG contains fewer errors. This can be attributed to its preprocessing procedure that selects candidate neighbors for each variable, reducing the possibility of making mistakes. MgCSL outperforms the baselines in terms of precision, F1 and SHD while identifying 6 correct causal edges. This signifies its tangible potential for real-world applications.

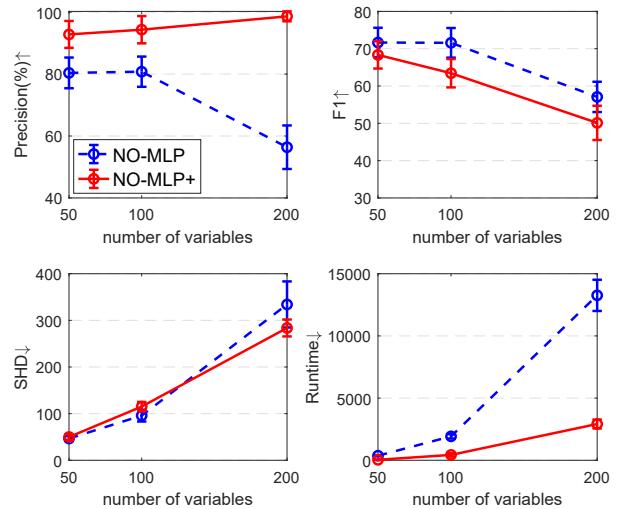
Ablation Study

We conduct ablation study to investigate the effectiveness of the proposed acyclicity constraint. We implement NO-MLP and NO-MLP combined with our proposed acyclicity constraint named NO-MLP+ on the dataset of nonlinear models with additive Gaussian processes. As shown in Fig. 3, NO-MLP+ yields sustained high precision even as the number of variables increases, while NO-MLP discovers excessive erroneous edges. Due to the stronger influence of directed cycles on eigenvalues, NO-MLP+ prunes a significant number of edges, including correct ones, leading to a decrease in F1 score. Nevertheless, in terms of SHD, NO-MLP+ manifests comparable performance to NO-MLP. Most importantly, NO-MLP+ demonstrates a significant decrease in time consumption, about an order of magnitude. This verifies that the proposed acyclicity penalty is beneficial in efficiently and accurately discovering causal DAG.

In addition, we carry out parameter sensitivity study w.r.t. α_1 and α_2 . A small α_1 may introduce considerable noise, while an excessively large α_1 could weaken the representational capacity of macro-variables. Besides, a larger α_2 can improve the performance in terms of precision and running time, but MgCSL tends to achieve an overly sparse estimated graph with a few edges when α_2 is too large. Based on the aforementioned analysis, we set $\alpha_1=0.1$ and $\alpha_2=0.01$.

Case Study

To further appraise the effectiveness of MgCSL in real-world applications, we conducted a case study on the functional magnetic resonance imaging (fMRI) Hippocampus

Figure 3: Precision, F1, SHD and runtime of NO-MLP and NO-MLP combined with our proposed acyclicity constraint named NO-MLP+ on $d=\{50, 100, 200\}$.

dataset, which is consisted of the resting-state signals from six separate brain regions (Poldrack et al. 2015). It was collected from the same person for 84 days and the sample size for each day is 518. We apply our proposed MgCSL on 10 successive days. We consider the top 5 edges discovered by MgCSL as the estimated causal edges: {PHC, Sub} \rightarrow ERC, ERC \rightarrow DG, CA1 \rightarrow Sub, and CA1 \rightarrow DG. We compared the results with anatomical directed connections (Bird and Burgess 2008; Zhang et al. 2017), since theoretically, when two regions have an anatomical connection, there is a high likelihood of causal relationship between them. The ground truths provided by anatomical structures contain cycles, while we assume that there are no directed cycles in the causal graph. The results indicate that the first four edges discovered by MgCSL match the fact that most of the hippocampus’s neocortical inputs come from the PHC via the ERC, with most of its neocortical output directed towards the Sub, which also projects back to the ERC. The collaborative interaction between PHC and Sub on ERC unveils potential cross-area coordination within the brain. The last edge is reversed, possibly attributed to the presence of directed cycles, leading MgCSL to misjudge the direction. In summary, the results demonstrate the effectiveness of MgCSL in identifying causal connections between different brain regions on real fMRI datasets.

Conclusion

In this paper, we study how to acquire multi-granularity causal structure on observational data, which is a practical and significant, but currently under-researched problem. We proposed an effective approach MgCSL that leverages SAE to extract latent macro-variables, and utilizes MLPs with a simplified acyclicity penalty to speed up the discovery of multi-granularity causal structure. The efficacy of MgCSL is verified by experiments on synthetic and real-world datasets.

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

Bird, C. M.; and Burgess, N. 2008. The hippocampus and memory: insights from spatial processing. *Nat. Rev. Neurosci.*, 9(3): 182–194.

Byrd, R. H.; Lu, P.; Nocedal, J.; and Zhu, C. 1995. A limited memory algorithm for bound constrained optimization. *SIAM J Sci. Comput.*, 16(5): 1190–1208.

Chickering, D. M. 2002. Optimal structure identification with greedy search. *JMLR*, 3(11): 507–554.

Colombo, D.; Maathuis, M. H.; et al. 2014. Order-independent constraint-based causal structure learning. *JMLR*, 15(116): 3921–3962.

Demšar, J. 2006. Statistical comparisons of classifiers over multiple data sets. *JMLR*, 7: 1–30.

Elder-Vass, D. 2012. Top-down causation and social structures. *Interface focus*, 2(1): 82–90.

Hauser, A.; and Bühlmann, P. 2012. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *JMLR*, 13(79): 2409–2464.

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.

Hoyer, P.; Janzing, D.; Mooij, J.; Peters, J.; and Schölkopf, B. 2008. Nonlinear causal discovery with additive noise models. In *NeuIPS*, 689–696.

Lachapelle, S.; Brouillard, P.; Deleu, T.; and Lacoste-Julien, S. 2020. Gradient-based neural DAG learning. In *ICLR*.

Lee, H.-C.; Danieletto, M.; Miotto, R.; Cherng, S. T.; and Dudley, J. T. 2019. Scaling structural learning with NO-BEARS to infer causal transcriptome networks. In *PSB*, 391–402. World Scientific.

Liang, J.; Wang, J.; Yu, G.; Domeniconi, C.; Zhang, X.; and Guo, M. 2023a. Gradient-based local causal structure learning. *IEEE TCYB*, 99(1): 1–10.

Liang, J.; Wang, J.; Yu, G.; Guo, W.; Domeniconi, C.; and Guo, M. 2023b. Directed acyclic graph learning on attributed heterogeneous network. *TKDE*, 99(1): 1–12.

Liang, J.; Wang, J.; Yu, G.; Xia, S.; and Wang, G. 2023c. Multi-granularity Causal Structure Learning. arXiv:2312.05549.

Loewer, B. 2012. The emergence of time’s arrows and special science laws from physics. *Interface focus*, 2(1): 13–19.

Marella, D.; and Vicard, P. 2022. Bayesian network structural learning from complex survey data: a resampling based approach. *SMA*, 31(4): 981–1013.

Ng, A.; et al. 2011. Sparse autoencoder. *CS294A Lecture notes*, 72(2011): 1–19.

Ng, I.; Zhu, S.; Chen, Z.; and Fang, Z. 2019. A graph autoencoder approach to causal structure learning. In *NeurIPS Workshop*.

Noble, D. 2012. A theory of biological relativity: no privileged level of causation. *Interface focus*, 2(1): 55–64.

Peters, J.; Mooij, J. M.; Janzing, D.; and Schölkopf, B. 2014. Causal discovery with continuous additive noise models. *JMLR*, 15: 2009–2053.

Poldrack, R. A.; Laumann, T. O.; Koyejo, O.; Gregory, B.; Hover, A.; Chen, M.-Y.; Gorgolewski, K. J.; Luci, J.; Joo, S. J.; Boyd, R. L.; et al. 2015. Long-term neural and physiological phenotyping of a single human. *Nat. Commun.*, 6(1): 8885.

Qiang, W.; Li, J.; Su, B.; Fu, J.; Xiong, H.; and Wen, J.-R. 2023. Meta attention-generation network for cross-granularity few-shot learning. *IJCV*, 131(5): 1211–1233.

Ramsey, J.; Glymour, M.; Sanchez-Romero, R.; and Glymour, C. 2017. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *JDSA*, 3(2): 121–129.

Richens, J. G.; Lee, C. M.; and Johri, S. 2020. Improving the accuracy of medical diagnosis with causal machine learning. *Nat. Commun.*, 11(1): 3923.

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.

Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634.

Spirites, P.; Glymour, C. N.; Scheines, R.; and Heckerman, D. 2000. *Causation, prediction, and search*. Cambridge, MA: MIT Press, 2nd ed. edition.

Vandenbroucke, J. P.; Broadbent, A.; and Pearce, N. 2016. Causality and causal inference in epidemiology: the need for a pluralistic approach. *Int. J. Epidemiol.*, 45(6): 1776–1786.

Varley, T. F.; Pope, M.; Faskowitz, J.; and Sporns, O. 2023. Multivariate information theory uncovers synergistic subsystems of the human cerebral cortex. *Commun. Biol.*, 6(1): 451.

Von Kügelgen, J.; Karimi, A.-H.; Bhatt, U.; Valera, I.; Weller, A.; and Schölkopf, B. 2022. On the fairness of causal algorithmic recourse. In *AAAI*, 9584–9594.

Wang, G. 2017. DGCC: data-driven granular cognitive computing. *Granul. Comput.*, 2(4): 343–355.

Wang, G.; Fu, S.; Yang, J.; and Guo, Y. 2022. A review of research on multi-granularity cognition based intelligent computing. *Chin. J. Comput.*, 45(6): 1161–1175.

Wang, X.; Du, Y.; Zhu, S.; Ke, L.; Chen, Z.; Hao, J.; and Wang, J. 2021. Ordering-based causal discovery with reinforcement learning. In *IJCAI*, 3566–3573.

Wang, Y.; Liang, D.; Charlin, L.; and Blei, D. M. 2020. Causal inference for recommender systems. In *RecSys*, 426–431.

Wu, M.; Zhang, Q.; Wu, C.; and Wang, G. 2023. End-to-end multi-granulation causality extraction model. *DCN*.

Yang, C.; Zhang, Z.; Ding, J.; Zheng, W.; Jing, Z.; and Li, Y. 2022. A multi-granularity network for emotion-cause pair extraction via matrix capsule. In *CIKM*, 4625–4629.

Yang, D.; Yu, G.; Wang, J.; Wu, Z.; and Guo, M. 2023. Reinforcement causal structure learning on order graph. In *AAAI*, 10737–10744.

Yu, Y.; Chen, J.; Gao, T.; and Yu, M. 2019. DAG-GNN: DAG structure learning with graph neural networks. In *ICML*, 7154–7163.

Zhang, K.; Huang, B.; Zhang, J.; Glymour, C.; and Schölkopf, B. 2017. Causal discovery from nonstationary/heterogeneous data: skeleton estimation and orientation determination. In *IJCAI*, 1347–1353.

Zhang, K.; and Hyvärinen, A. 2009. On the identifiability of the post-nonlinear causal model. In *UAI*, 647–655.

Zheng, X.; Aragam, B.; Ravikumar, P. K.; and Xing, E. P. 2018. Dags with no tears: continuous optimization for structure learning. In *NeuIPS*, 9492–9503.

Zheng, X.; Dan, C.; Aragam, B.; Ravikumar, P.; and Xing, E. 2020. Learning sparse nonparametric dags. In *AISTATS*, 3414–3425.