



PDF Download
3616855.3635827.pdf
30 December 2025
Total Citations: 2
Total Downloads: 564

Latest updates: <https://dl.acm.org/doi/10.1145/3616855.3635827>

RESEARCH-ARTICLE

Interpretable Imitation Learning with Dynamic Causal Relations

TIANXIANG ZHAO, Pennsylvania State University, University Park, PA, United States

WENCHAO YU, NEC Laboratories America, Inc., Princeton, NJ, United States

SUHANG WANG, Pennsylvania State University, University Park, PA, United States

LU WANG, East China Normal University, Shanghai, China

XIANG ZHANG, Pennsylvania State University, University Park, PA, United States

YUNCONG CHEN, NEC Laboratories America, Inc., Princeton, NJ, United States

[View all](#)

Open Access Support provided by:

NEC Laboratories America, Inc.

Pennsylvania State University

East China Normal University

Published: 04 March 2024

[Citation in BibTeX format](#)

WSDM '24: The 17th ACM International Conference on Web Search and Data Mining

March 4 - 8, 2024
Merida, Mexico

Conference Sponsors:

SIGKDD
SIGMOD
SIGIR
SIGWEB

Interpretable Imitation Learning with Dynamic Causal Relations

Tianxiang Zhao
The Pennsylvania State University
State College, USA
tkz5084@psu.edu

Lu Wang
East China Normal University
Shanghai, China
luwang@stu.ecnu.edu.cn

Yanchi Liu
NEC Laboratories America
Princeton, USA
yanchi@nec-labs.com

Wenchao Yu
NEC Laboratories America
Princeton, USA
wyu@nec-labs.com

Xiang Zhang
The Pennsylvania State University
State College, USA
xzz89@psu.edu

Wei Cheng
NEC Laboratories America
Princeton, USA
weicheng@nec-labs.com

Suhang Wang
The Pennsylvania State University
State College, USA
szw494@psu.edu

Yuncong Chen
NEC Laboratories America
Princeton, USA
yuncong@nec-labs.com

Haifeng Chen
NEC Laboratories America
Princeton, USA
haifeng@nec-labs.com

ABSTRACT

Imitation learning, which learns agent policy by mimicking expert demonstration, has shown promising results in many applications such as medical treatment regimes and self-driving vehicles. However, it remains a difficult task to interpret control policies learned by the agent. Difficulties mainly come from two aspects: 1) agents in imitation learning are usually implemented as deep neural networks, which are black-box models and lack interpretability; 2) the latent causal mechanism behind agents' decisions may vary along the trajectory, rather than staying static throughout time steps. To increase transparency and offer better interpretability of the neural agent, we propose to expose its captured knowledge in the form of a directed acyclic causal graph, with nodes being action and state variables and edges denoting the causal relations behind predictions. Furthermore, we design this causal discovery process to be state-dependent, enabling it to model the dynamics in latent causal graphs. Concretely, we conduct causal discovery from the perspective of Granger causality and propose a self-explainable imitation learning framework, CAIL. The proposed framework is composed of three parts: a dynamic causal discovery module, a causality encoding module, and a prediction module, and is trained in an end-to-end manner. After the model is learned, we can obtain causal relations among states and action variables behind its decisions, exposing policies learned by it. Experimental results on both synthetic and real-world datasets demonstrate the effectiveness of the proposed CAIL in learning the dynamic causal graphs for understanding the decision-making of imitation learning meanwhile maintaining high prediction accuracy.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WSDM '24, March 4–8, 2024, Merida, Mexico

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0371-3/24/03...\$15.00
<https://doi.org/10.1145/3616855.3635827>

CCS CONCEPTS

• **Computing methodologies** → **Reinforcement learning; Learning from demonstrations; Statistical relational learning.**

KEYWORDS

imitation learning, temporal sequences, interpretation

ACM Reference Format:

Tianxiang Zhao, Wenchao Yu, Suhang Wang, Lu Wang, Xiang Zhang, Yuncong Chen, Yanchi Liu, Wei Cheng, and Haifeng Chen. 2024. Interpretable Imitation Learning with Dynamic Causal Relations. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*, March 4–8, 2024, Merida, Mexico. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3616855.3635827>

1 INTRODUCTION

In imitation learning, neural agents are trained to acquire control policies by mimicking expert demonstrations. It circumvents two vital deficiencies of traditional DRL methods: low sampling efficiency and reward sparsity. Following demonstrations that return near-optimal rewards, the imitator can prevent a vast amount of unreasonable attempts during explorations and has been shown to be promising in many real-world applications [11, 19, 25, 34, 46]. However, despite the high performance of imitating neural agents, one problem persists in the interpretability of control policies learned by them. With deep neural networks used as the policy model, the decision mechanism of the trained neural agent is not transparent and remains a black box, making it difficult to trust the model and apply it on high-stake scenarios like the medical domain [35, 41].

Many efforts have been made to increase the interpretability of policy agents. For example, Reference [44] and [30] compute saliency maps to highlight critical features using gradient information or attention mechanism; [45] models interactions among entities via relational reasoning; [28] designs sub-tasks to make decisions with symbolic planning. However, these methods either provide explanations that are noisy and difficult to interpret [30, 44], only in the instance level without a global view of the overall policy or make too strong assumptions on the neural agent and lack generality [28].

To increase the interpretability of learned neural agents, we propose to explain it from the cause-effect perspective, exposing causal relations among observed state variables and outcome decisions. Inspired by advances in discovering directed acyclic graphs (DAGs) [49], we aim to learn a self-explainable imitator by discovering the causal relationship between states and actions. Concretely, taking observable state variables and candidate actions as nodes, the neural agent can generate a DAG to depict the underlying dependency between states and actions, with edges representing causal relationships. For example, in the medical domain, the obtained DAG can contain relations like “Inactive muscle responses often indicates losing of speaking capability” or “Severe liver disease would encourage the agent to recommend using Vancomycin”, as shown in the case study Figure 6. Such exposed relations can improve user understanding of the policies of the neural agent from a global view, and can provide better explanations of decisions made by it.

However, designing such interpretable imitators from a causal perspective is a very challenging task, mainly due to two reasons: 1) It is non-trivial to identify causal relations behind the decision-making of imitating agents. Modern imitators are usually implemented as a deep neural network, in which the utilization of features is entangled and nonlinear, and lack interpret-ability; and 2) Imitators need to make decisions in a sequential manner, and latent causal structures behind it could evolve over time, instead of staying static throughout the produced trajectory. For example, in a medical scenario, a trained imitator needs to make sequential decisions that specify how the treatments should be adjusted through time according to the dynamic states of the patient. As indicated in [4, 8], there are multiple stages in the states of patients w.r.t disease severity, which would influence the efficacy of drug therapies and result in different treatment policies at each stage. However, directly incorporating this temporal dynamic element into causal discovery would give too much flexibility in search space, and can easily lead to over-fitting.

Targeting at aforementioned challenges, we build our causal discovery objective upon the notion of Granger causality [3, 10], which declares a causal relationship $s_i \rightarrow a_j$ between variables s_i and a_j if a_j can be better predicted with s_i available than not. A causal discovery module is designed to uncover causal relations among variables, and extracted causes are encoded into embeddings of outcome variables before action prediction. Noted that we intervene on the variables used by the agent during prediction and are interested in how its behavior is affected following the notion of Granger causality, and are not discovering causal relations of real-world data.

Concretely, in this work, we propose to design an imitator which is able to produce DAGs providing interpretations on the control policy alongside predicting actions and name it as Causal-Augmented Imitation Learning (CAIL). Identified causal relations are encoded into variable representations as evidence for making decisions. With this manipulation of inputs, we circumvent the onerous analysis of internal structures of neural agents and manage to model causal discovery as an optimization task. Following the observation that the evolution of causal structures usually follows a stage-wise process [4], we assume a set of latent templates during designing the causal discovery module which can both model the temporal dynamics across stages and allows for knowledge sharing within

the same stage. Consistency between extracted DAGs and captured policies is guaranteed in design, and this framework can be updated in an end-to-end manner. Intuitive constraints are also enforced to regularize the structure of discovered causal graphs, like encouraging sparsity and preventing loops. In summary, our main contributions are:

- We study a novel problem of learning dynamic causal graphs to uncover the knowledge captured as well as latent causes behind agent’s decisions.
- We propose a novel framework CAIL, which is able to learn dynamic DAGs to capture the casual relation between state variables and actions and adopt the DAGs for decision-making in imitation learning;
- We conduct experiments on synthetic and real-world datasets to demonstrate the effectiveness of CAIL in learning the dynamic DAGs for understanding the decision making of imitation learning meanwhile maintain high prediction accuracy.

2 RELATED WORK

2.1 Imitation Learning

Imitation learning is a special case in the reinforcement learning domain, where an agent (policy model) is trained to perform a task and learn a mapping between observations and actions from expert demonstrations [18, 48]. Such expert knowledge avoids acquiring skills from scratch, and reduces the difficulties of learning in complex and uncertain environments. Imitation learning has been found effective in a wide range of applications, such as human-computer interaction [32, 36, 46, 50], self-driving vehicles [1, 6] and robotic arms [16, 23].

Existing imitation learning algorithms can mainly be categorized into two groups, behavior cloning and reinforcement learning with self-defined reward functions [41]. Behavioral cloning establishes a direct mapping between states and actions on expert trajectories, hence rewards can be obtained similar to the supervised learning setting [39, 42]. However, it could suffer from reward sparsity problems due to insufficient demonstrations. The other group designs their own reward scores, like inverse reinforcement learning [51] which learns a reward function that would be maximized by expert trajectories, and generative adversarial imitation learning (GAIL) [15] which trains a discriminator to tell those generated by agents apart from expert trajectories.

Besides these progresses, lack of interpretability is a critical weak point shared by most imitation learning methods. Neural agent is usually implemented as a DNN and treated as a black box. Currently, interpreting imitation learning agents is still an under-explored task. Similar to our idea, Reference [7] also introduces causality into imitation learning. However, it identifies causal relations through targeted intervention and feature masking, which is computation extensive and relies upon domain knowledge.

2.2 Structure Learning from Time-Series

The task of modeling discrete-time temporal dynamics in DAGs, which is also known as learning dynamic Bayesian networks (DBNs), has generated significant interest in recent years [14, 21, 26, 31, 47?]. It has been used successfully in a variety of domains like disease prognosis [40] and speech recognition [29, 33]. Most existing

approaches are designed as *structural vector autoregressive* (SVAR) models, learning cross-time dependency among variables. For example, Reference [31] fits a linear structural equation model (SEM) on temporal sequences. Reference [37] designs an LSTM-based model to latently encode SEM, and is able to model nonlinear relations. However, in these works, conditional dependence among variables is taken as stationary across time points. In many real-world applications, the conditional dependency among variables might change. For example, the impacts of oil price on GDP growth are significantly different when the economy is in a high growth phase versus low growth phase [2]. Initial efforts have been made to allow the modeling of dynamic DAGs [12, 38] via regularization like graph LASSO. More related to ours is the work [17], which learns state-dependent linear SEMs with the assumption of multiple stages.

Our work differentiates from these methods mainly from two perspectives. First, we focus on incorporating causal discovery into the design of imitators to make them more interpretable. Second, we make little assumption on the form of causal models as the decision policy is usually very complex and nonlinear.

3 PROBLEM DEFINITION

Throughout this work, we use \mathcal{S} and \mathcal{A} to denote sets of states and actions, respectively. In a classical discrete-time stochastic control process, the state at each time step is dependent upon the state and action from the previous step: $\mathbf{s}_{t+1} \sim \mathbf{P}(\mathbf{s}|\mathbf{s}_t, \mathbf{a}_t)$. $\mathbf{s}_t \in \mathcal{S}$ is the state vector in time step t , consisting of descriptions over observable state variables. $\mathbf{a}_t \in \mathbb{R}^K$ indicates actions taken in time t , and K is the size of candidate action set $|\mathcal{A}|$. Traditionally, deep reinforcement learning dedicates to learning a policy model π_θ to select actions given states: $\pi_\theta(\mathbf{s}) = \mathbf{P}_{\pi_\theta}(\mathbf{a}|\mathbf{s})$, which can maximize long-term rewards. In imitation learning setting, ground-truth rewards on actions at each time step are not available. Instead, a set of demonstration trajectories $\tau = \{\tau_1, \tau_2, \dots, \tau_m\}$ sampled from expert policy π_E is given, where $\tau_i = (\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1, \mathbf{a}_1, \dots)$ is the i -th trajectory with \mathbf{s}_t and \mathbf{a}_t being the state and action at time step t . Accordingly, the target is changed to learn a policy π_θ that mimics the behavior of expert π_E . A summary of notations is provided in Appendix. A.

In this work, besides obtaining the policy model π_θ , we further seek to provide interpretations for its decisions. Using notations from the causality scope, we focus on discovering the cause-effect dependency among observed states and predicted actions encoded in π_θ . Without loss of generality, we can formalize it as a causal discovery task. Concretely, we model causal relations with an augmented linear Structural Equation Model (SEM) [49]:

$$\mathbf{s}_{t+1}, \mathbf{a}_t = f_2(\mathcal{G}_t \cdot f_1(\mathbf{s}_t, \mathbf{a}_{t-1})) \quad (1)$$

In this equation, f_1, f_2 are nonlinear transformation functions. Directed acyclic graph (DAG) $\mathcal{G}_t \in \mathbb{R}^{(S+\mathcal{A}) \times (S+\mathcal{A})}$ can be represented as an adjacency matrix as it is unattributed. \mathcal{G}_t measures the causal relation of state variables \mathbf{s} and action variable \mathbf{a} in time step t , and sheds lights on interpreting decision mechanism of π_θ . It exposes the latent interaction mechanism between state and action variables lying behind π_θ . The task can be formally defined as:

PROBLEM 1. *Given m expert trajectories represented as τ , learn a policy model π_θ that predicts the action \mathbf{a}_t based on states \mathbf{s}_t , along with a DAG \mathcal{G}_t exposing the causal structure captured by it in the current time step. This self-explaining strategy helps to improve user understanding of trained imitator.*

4 METHODOLOGY

In this section, we introduce the details of the proposed framework CAIL. The basic idea of CAIL is to discover the causal relationships between state and action and utilize the causal relations to help the agent make decisions. The discovered causal graphs can also provide a high-level interpretation of the neural agent, exposing the reasons behind its decisions. An overview of the proposed CAIL is provided in Figure 1. Concretely, we develop a self-explaining framework that can provide the latent causal graph besides predicted actions, which is composed of: (1) a causal discovery module that constructs a causal graph capturing the causal relations among states and actions for each time step, which can help decision of which action to take next and explain the decision; (2) a causal encoding module which models causal graphs to encode the discovered causal relations for imitation learning; and (3) a prediction module that conducts the imitation learning task based on both the current state and causal relation. All three components are trained end-to-end, and this design guarantees the conformity between discovered causal structures and the behavior of π_θ . Next, we will introduce the detailed design of these modules one by one.

4.1 Dynamic Causal Discovery

Discovering the causal relations between state and action variables can help decision-making of neural agents and increase their interpretability. However, for many real-world applications, the latent generation process of observable states \mathbf{s} and the corresponding action \mathbf{a} may undergo transitions at different periods of the trajectory. For example, there are multiple stages for a patient, such as “just infected”, “become severe” and “begin to recovery”. Different stages of patients would influence the efficacy of drug therapies [4, 8], making it sub-optimal to use one fixed causal graph to model policy π_θ . On the other hand, separately fitting a \mathcal{G}_t at each time step is an onerous task, and could suffer from the lack of training data.

To address this problem, we design a causal discovery module to produce dynamic causal graphs. Concretely, we assume that the evolving of a time series can be split into multiple stages, and the casual relationship within each stage is static. This assumption widely holds in many real-world applications, as observed in [4, 8, 12]. Under this assumption, a discovery model with M DAG templates is designed, and \mathcal{G}_t is extracted as a soft selection of those templates.

4.1.1 Causal Graph Learning. An illustration of this causal discovery module is shown in Figure 2. Specifically, we first construct an explicit dictionary $\{\mathcal{G}^i, i \in [1, 2, \dots, M]\}$ as the DAG templates. $\mathcal{G}^i \in \mathbb{R}^{(S+\mathcal{A}) \times (S+\mathcal{A})}$ and these templates are randomly initialized and will be learned together with the other modules of CAIL. They encode the time-variate part of causal relations.

Following existing work [49], we add the sparsity constraint and the acyclicity regularizer on \mathcal{G}^i to make sure that \mathcal{G}^i is directed acyclic graph. The sparsity regularizer applies $L1$ norm on the

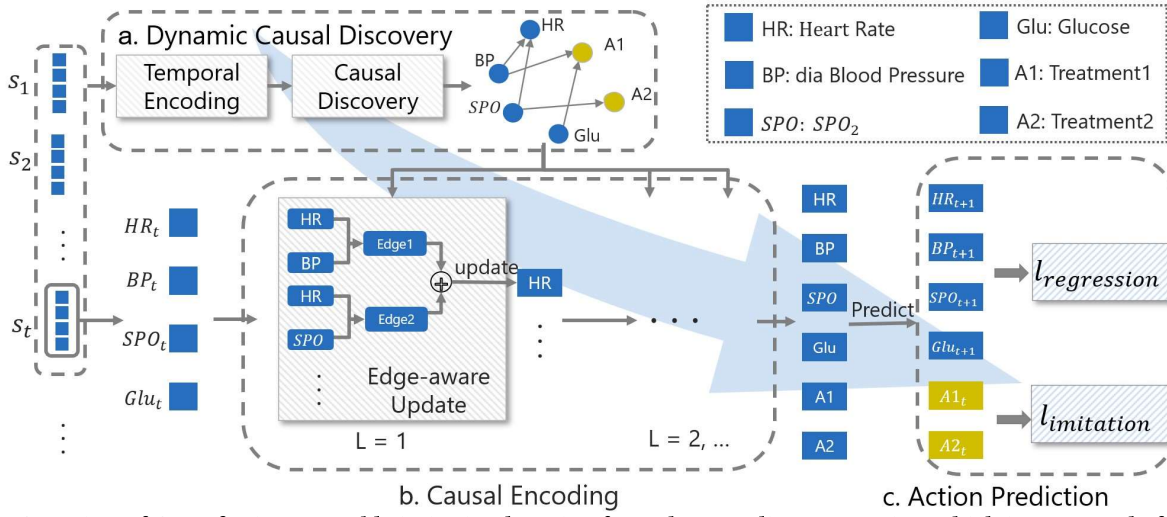


Figure 1: Overview of CAIL for interpretable imitation learning from the causality perspective, which is composed of (a) causal discovery module, (b) causal encoding module, (c) prediction module. The arrow in the back shows the inference order of them.




Figure 2: Overview of the dynamic causal discovery module.

causal graph templates to encourage sparsity of discovered causal relations so that those non-causal edges could be removed. It can be mathematically written as

$$\min_{\{\mathcal{G}^i, i \in [1, 2, \dots, M]\}} \mathcal{R}_{\text{sparcity}} = \sum_{i=1}^M |\mathcal{G}^i|. \quad (2)$$

where $|\mathcal{G}^i|$ denotes number of edges inside it.

In a causal graph, edges are directed and a node cannot be its own descendant. To enforce such constraint on extracted graphs, we adopt the acyclicity regularization in [43]. Concretely, \mathcal{G}^i is acyclic if and only if $\mathcal{H}(\mathcal{G}^i) = \text{tr}[e^{\mathcal{G}^i \circ \mathcal{G}^i}] - (|\mathcal{S}| + |\mathcal{A}|) = 0$, where \circ is element-wise square, $e^{\mathcal{A}}$ is the matrix exponential of \mathcal{A} , and tr denotes matrix trace. $|\mathcal{S}|$ and $|\mathcal{A}|$ are the number of state and action variables, respectively. Then the regularizer to make the graph acyclic can be written as:

$$\min_{\{\mathcal{G}^i, i \in [1, 2, \dots, M]\}} \mathcal{R}_{\text{DAG}} = \sum_{i=1}^M (\mathcal{H}(\mathcal{G}^i) - (|\mathcal{S}| + |\mathcal{A}|)). \quad (3)$$

When \mathcal{R}_{DAG} is minimized to be 0, there would be no loops in the discovered causal graphs and they are guaranteed to be DAGs.

4.1.2 Causal Graph Selection. With the DAG templates, at each time stamp t , we select one DAG from the templates that can well describe the causal relation between state variables and actions at the current status. Specifically, we use a temporal encoding network

to learn the representation of the trajectory for input time step t as

$$\mathbf{z}_t = \text{Enc}(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_t). \quad (4)$$

In the experiments, we apply a Temporal CNN as the encoding model. Note that other sequence encoding models like LSTM and Transformer can also be used here. For each template \mathcal{G}^i , we also learn its representation as:

$$\mathbf{u}^i = g(\mathcal{G}^i). \quad (5)$$

As \mathcal{G} is unattributed and its nodes are ordered, we implement $g(\cdot)$ as an MLP with flattened \mathcal{G} as input, i.e., the connectivity of each node. Note that geometric networks like graph neural networks (GNNs) [13, 22] can also be used here.

Since \mathbf{z}_t captures the trajectory up to time t , we can use \mathbf{z}_t to generate \mathcal{G}_t by selecting from templates $\{\mathcal{G}^i\}$ as

$$\alpha_t^i = \frac{\exp(\langle \mathbf{z}_t, \mathbf{u}^i \rangle / T)}{\sum_{i=1}^M \exp(\langle \mathbf{z}_t, \mathbf{u}^i \rangle / T)}, \quad \mathcal{G}_t = \sum_{i=1}^M \alpha_t^i \cdot \mathcal{G}^i \quad (6)$$

where $\langle \cdot, \cdot \rangle$ denotes vector inner-product. Here, we adopt a soft selection by setting temperature T to a small value, 0.1. A small T would make α_t^i more close to 0 or 1.

To encourage consistency in template selection across similar time steps, we design the template selection regularization loss here. Specifically, states and historical actions at each time are concatenated and clustered into M groups beforehand. We use q_t^i to denote whether time steps t belongs to group i , which is obtained from the clustering results. Then, the loss function for guiding the template selection can be written as

$$\min_{\theta} \mathcal{R}_{\text{option}} = - \sum_{i=1}^M \sum_t q_t^i \log \alpha_t^i. \quad (7)$$

where α_t^i is the selection weight of time step t on template i from Eq.(6) and θ is the set of parameters of graph templates, temporal encoding network Enc and $g(\cdot)$.

4.2 Encoding Causal Relations into Embeddings

For the purpose of learning \mathcal{G} to capture causal structures, we need to guarantee its consistency with the behavior of π_θ . In this work, we achieve that on the input level. Specifically, we obtain variable embeddings by modeling the interactions among them based on discovered causal relations, and then train π_θ on top of these updated embeddings. In this way, the structure of \mathcal{G}_t can be updated along with the optimization of π_θ . Next, we will introduce the process of encoding causality into variable embeddings in detail.

4.2.1 Variable Initialization. Let $s_{t,j}$ denote state variable s_j at time t . First, we map each observed variable $s_{t,j}$ to the embedding of the same shape for future computations with:

$$\hat{\mathbf{h}}_{t,j}^0 = s_{t,j} \cdot \mathbf{E}_j, \quad (8)$$

where $\mathbf{E}_j \in \mathbb{R}^{|s_j| \times d}$ is the embedding matrix to be learned for the j -th observed variable. $\hat{\mathbf{h}}_t^0 \in \mathbb{R}^{|\mathcal{S}| \times d}$, d is the dimension of embedding for each variable. We further extend it to $\mathbf{h}_t^0 \in \mathbb{R}^{(|\mathcal{S}|+|\mathcal{A}|) \times d}$, to include representation of actions. Representation of these actions are initialized as zero and are learned during training.

4.2.2 Causal Relation Encoding. Then, we update the representation of all variables using \mathcal{G}_t , which aims to encode the casual relation with the representations. In many real-world cases, variables may contain very different semantics and directly fusing them using homophily-based GNNs like GCN [22] is improper. To better model the heterogeneous property of variables, we adopt an edge-aware architecture:

$$\begin{aligned} \mathbf{m}_{j \rightarrow i} &= [\mathbf{h}_{i,t}^{l-1}, \mathbf{h}_{j,t}^{l-1}] \cdot \mathbf{W}_{edge}^l \\ \mathbf{h}_{i,t}^l &= \sigma \left(\left[\sum_{j \in \mathcal{V}} \tilde{\mathcal{G}}_{j,i} \mathbf{m}_{j \rightarrow i}, \mathbf{h}_{i,t}^{l-1} \right] \mathbf{W}_{agg}^l \right) \end{aligned} \quad (9)$$

where \mathbf{W}_{edge}^l and \mathbf{W}_{agg}^l are the parameter matrices for edge-wise propagation and node-wise aggregation respectively in layer l . $\mathbf{m}_{j \rightarrow i}$ refers to the message from node j to node i . In the experiments, L is set as 2 if not stated otherwise.

4.3 Prediction with Causality-Encoded Embeddings

After obtaining causality-encoded variable embeddings, a prediction module is implemented on top of them to conduct the imitation learning task. Its gradients will be back-propagated through the causal encoding module to the causal discovery module, hence informative edges containing causal relations can be identified. In this section, we will introduce the detailed design of this module, along with its training signals.

4.3.1 Imitation Learning Task. After previous steps, now $\mathbf{h}_{t,j}$ encodes both observations and causal factors for variable j . Then, we make predictions on \mathbf{a}_t , which is a vector of length $|\mathcal{A}|$, with each dimension indicating whether to take the corresponding action or not. Concretely, for action candidate a' , the process is as follows: (1) $\mathbf{h}_{t,a'}$ and a'_{t-1} are concatenated as the input evidence. $\mathbf{h}_{t,a'}$ is the obtained embedding for variable a' at time t , and a'_{t-1} corresponds to the history action from last time. (2) The branch a' of trained policy model π_θ predicts the action a'_t based on $[\mathbf{h}_{t,a'}, a'_{t-1}]$. In our

implementation, π_θ is composed of $|\mathcal{A}|$ branches with each branch corresponding to one certain action variable.

Following existing works [15], the proposed policy model is adversarially trained with a discriminator D to imitate expert decisions. Specifically, the policy π_θ aims to generate realistic trajectories that can mimic π_E so as to fool the discriminator D ; while the discriminator aims to differentiate if a trajectory is from π_θ or π_E . Through such a min-max game, π_θ can imitate the expert trajectories. The learning objective \mathcal{L}_{imi} on policy π_θ is given as:

$$\begin{aligned} \min_{\pi_\theta} \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \rho_{\pi_\theta}} \log(1 - D(\mathbf{s}, \mathbf{a})) - \lambda H(\pi_\theta) \\ - \mathbb{E}_{\tau_i \in \tau} \mathbb{E}_{(\mathbf{s}_t, \mathbf{a}_t) \sim \tau_i} P_{\pi_\theta}(\mathbf{a}_t | \mathbf{s}_t), \end{aligned} \quad (10)$$

where ρ_{π_θ} is the trajectory generated by π_θ and τ is the set of expert demonstrations. $H(\pi) \triangleq \mathbb{E}_{\pi_\theta}[-\log \pi(\mathbf{a}|\mathbf{s})]$ is the entropy that encourages π_θ to explore and make diverse decisions. The discriminator D is trained to differentiate expert paths from those generated by π_θ , whose objective function is:

$$\max_D \mathbb{E}_{\rho_E} \log(D(\mathbf{s}, \mathbf{a})) + \mathbb{E}_{\rho_\theta} \log(1 - D(\mathbf{s}, \mathbf{a})) \quad (11)$$

Our framework is agnostic towards architecture choices of policy model π_θ . In the experiments, π_θ is implemented as a three-layer MLP, with the first two layers shared by all branches. Relu is selected as the activation function.

4.3.2 Auxiliary Regression Task. Besides the common imitation learning task, we further conduct an auto-regression task on state variables. This task can provide auxiliary signals to guide the discovery of causal relations, like the edge from Blood Pressure to Heart Rate in Figure 1. Similar to the imitation learning task, for state variable s' we use $[\mathbf{h}_{t,s'}, s'_t]$ as the evidence, and use model π_ϕ to predict s'_{t+1} as \mathcal{L}_{res} :

$$\min_{\pi_\phi} -\mathbb{E}_{\tau_i \in \tau} \mathbb{E}_{(\mathbf{s}_t, \mathbf{a}_t) \sim \tau_i} \log P_{\pi_\phi}(\mathbf{s}_{t+1} | \mathbf{h}_{t,s}, \mathbf{s}_t), \quad (12)$$

in which P_{π_ϕ} denotes the predicted distribution of \mathbf{s}_{t+1} .

4.4 Final Objective Function of CAIL

Putting everything together, the final objective function of the proposed CAIL is given as:

$$\begin{aligned} \min_{\pi_\phi, \pi_\theta} \max_D \mathcal{L}_{imi} + \gamma_1 \cdot \mathcal{L}_{res} + \lambda_1 \cdot \mathcal{R}_{sparse} + \gamma_2 \cdot \mathcal{R}_{option} \\ \text{s.t. } \mathcal{R}_{DAG} = 0. \end{aligned} \quad (13)$$

where λ_1, γ_1 , and γ_2 are weights of different losses, and the constraint guarantees acyclicity in graph templates.

To solve this constrained problem in Equation 13, we use the augmented Lagrangian algorithm and get its dual form:

$$\begin{aligned} \min_{\pi_\phi, \pi_\theta} \max_D \mathcal{L}_{imi} + \gamma_1 \cdot \mathcal{L}_{res} + \lambda_1 \cdot \mathcal{R}_{sparse} + \gamma_2 \cdot \mathcal{R}_{option} \\ + \lambda_2 \cdot \mathcal{R}_{DAG} + \frac{c}{2} |\mathcal{R}_{DAG}|^2, \end{aligned} \quad (14)$$

where λ_2 is the Lagrangian multiplier and c is the penalty parameter. The optimization steps are summarized in Algorithm 1. Within each epoch, discriminator and the model parameter θ, ϕ are updated iteratively, as shown from line 2 to line 5. Between each epoch, we use augmented Lagrangian algorithm to update the multiplier λ_2 and penalty weight c from line 6 to line 11. These steps progressively increase the weight of \mathcal{R}_{DAG} , so that it will gradually converge to zero and templates will satisfy the DAG constraint.

Algorithm 1 Full Training Algorithm

Input: Demonstrations τ generated from expert policy π_E , initial template set $\{\mathcal{G}^i, i \in [1, 2, \dots, M]\}$, initial model parameter θ, ϕ , hyperparameters $\lambda_1, \lambda_2, \gamma_1, \gamma_2, c$, initialize $\mathcal{H}_{old} = \text{inf}$, parameter in Augmented Lagrangian: $\sigma = \frac{1}{4}, \rho = 10$

- 1: **while** Not Converged **do**
- 2: **for** $\tau_i \sim \tau$ **do**
- 3: Update parameter of discriminator D to increase the loss of Equation 11;
- 4: Update θ, ϕ with gradients to minimize Equation 13;
- 5: **end for**
- 6: Compute \mathcal{H} with Equation 3;
- 7: $\lambda_2 \leftarrow \lambda_2 + \mathcal{H} \cdot c$
- 8: **if** $\mathcal{H} \leq \sigma \cdot \mathcal{H}_{old}$ **then**
- 9: $c \leftarrow c * \rho$
- 10: **end if**
- 11: $\mathcal{H}_{old} \leftarrow \mathcal{H}$
- 12: **end while**
- 13: **return** Learned templates $\{\mathcal{G}^i, i \in [1, 2, \dots, M]\}$, trained policy model π_θ

5 EXPERIMENT

In this section, we evaluate the prediction accuracy and interpretability of the proposed CAIL on both synthetic and real-world datasets. Specifically, we aim to answer the following questions:

- **RQ1:** Can the proposed approach correctly identify causal relations among state and action variables?
- **RQ2:** Can our proposed method achieve better interpretability without sacrificing performance in imitation learning?
- **RQ3:** How would different hyperparameter configurations influence the effectiveness of proposed method?

5.1 Baselines

To the best of our knowledge, there is no existing work on discovering DAGs to help learn and interpret imitation learning models. To evaluate capacity of CAIL in identifying causal relations, we compare it with representative and state-of-the-art causal discovery methods in time-series data: (1) cMLP/cLSTM [37], which discovers nonlinear causal relations by training an MLP or an LSTM for each effect, and its causal factors are identified through analyzing non-zero entries in the weight matrix. (2) SRU/eSRU [21], which extends cLSTM through training a component-wise time-series predictor based on Statistical Recurrent Units(SRU). (3) DYNOTEARS [31], which designs a score-based approach to capture causal structures in time series. Causal graph is explicitly parameterized. and causal relations are assumed to be linear. (4) SrVARM [17], which is similar to DYNOTEARS in design, but assumes that causal graphs are state-specific and fall into K groups. (5) ACD [27], which trains a single amortized model that can infer causal relations across instances with different underlying causal graphs.

The comparison between our approach and these baselines is summarized in Table 1 in terms of three dimensions: whether discovers dynamic causal relations, whether supports nonlinear causal relations, and whether guarantees acyclicity in discovered causal

Table 1: Comparison with Baseline Causal Discovery Models

Models	Dynamic	Nolinear Causality	Acyclicity
cMLP	×	√	×
cLSTM	×	√	×
SRU	×	√	×
eSRU	×	√	×
DYNOTEARS	×	×	√
SrVARM	√	×	√
ACD	√	√	×
Our Approach	√	√	√

Table 2: Causal discovery performance on static Kuramoto dataset with different number of oscillators.

Models	AUROC		
	Kura5	Kura10	Kura50
cMLP	0.52 ± 0.01	0.48 ± 0.01	0.49 ± 0.02
cLSTM	0.47 ± 0.01	0.49 ± 0.01	0.49 ± 0.01
SRU	0.81 ± 0.04	0.64 ± 0.02	0.53 ± 0.03
eSRU	0.61 ± 0.04	0.54 ± 0.03	0.51 ± 0.02
DYNOTEARS	0.63 ± 0.01	0.57 ± 0.01	0.55 ± 0.02
SrVARM	0.67 ± 0.02	0.66 ± 0.04	0.57 ± 0.03
ACD	0.82 ± 0.04	0.71 ± 0.03	0.63 ± 0.02
Ours	0.95 ± 0.02	0.98 ± 0.01	0.91 ± 0.03

graphs. Following their designs, we train them through conducting auto-regression directly on expert trajectories π_E .

5.2 Datasets

To evaluate the performance of causal discovery, we conduct experiments on a publicly available synthetic dataset Kuramoto [27] and a real-world dataset MIMIC-IV [41]. To examine the performance of CAIL in imitation learning task ad answer RQ2, we further test it on three classical gym datasets [5]: MiniGrid-FourRoom, LavaGap, and DoorKey. Descriptions of these datasets are provided in Ap. B.

5.3 Configurations

5.3.1 Hyperparameter Settings. For all approaches, the learning rate is initialized as 0.005 and maximum epoch is set as 1,000. For both our approach and baselines, we use grid search to find hyperparameters with the best performance on each dataset. For synthetic dataset *Kuramoto*, M is fixed as 3 for ease of evaluation. For real-world datasets, Sepsis and Comorbidity, M is also set as 3 with prior knowledge on degree of severity [4]. Train:validation:test ratio is split as 2 : 3 : 5.

5.3.2 Evaluation Metrics. To evaluate the quality of discovered causal relations provided by the policy model, we use AUC-ROC score to measure their alignment with the ground-truth causal relationships. A higher AUC-ROC score means more accurate causal discovery performance, which indicates better interpretability. Besides, to evaluate performance in decision making, we also conduct teacher-forcing test and report the accuracy in action prediction.

Table 3: Causal discovery performance on dynamic Kuramoto dataset with a different number of oscillators.

Models	AUROC		
	Kura5_vary	Kura10_vary	Kura50_vary
SRU-D	0.78 ± 0.03	0.62 ± 0.02	0.54 ± 0.01
DYNOTEAR-D	0.62 ± 0.03	0.57 ± 0.02	0.52 ± 0.03
SrVARM	0.63 ± 0.04	0.64 ± 0.05	0.56 ± 0.02
ACD	0.73 ± 0.03	0.68 ± 0.02	0.58 ± 0.04
Ours	0.82 ± 0.05	0.78 ± 0.04	0.72 ± 0.03

5.4 Performance on Discovering Static DAG

First, we compare the performance of DAG learning of CAIL with baselines in the static setting. Each method is trained until converges to make a fair comparison. We conduct each experiment for 5 times. Both the average performance and the standard deviation are reported in Table 2. From the result, we can observe that:

- Our proposed framework achieves the best performance compared to baselines, which validates its capacity to correctly identify nonlinear causal relations;
- Our proposed framework scales well on this Kuramoto dataset. Compared to baselines like SRU, its performance is much more stable w.r.t graph sizes.
- Due to limitations like only support linear causal relations [17, 31] or unable to guarantee acyclicity [27, 37], baselines fail to achieve satisfactory performance on this task.

5.5 Performance on Discovering Dynamic DAG

Based on the results on static DAG, we select SRU, DYNOTEARS, SrVARM, and ACD as baselines for the scenario with varying DAGs. For SRU and DYNOTEARS, they are designed only for static causal relation discovery. Hence, we take ground-truth DAG index as known for these baselines and train them once for each DAG. We mark them as SRU-D and DYNOTEAR-D respectively. Other approaches can cope with dynamic causal relations and do not require such modification. The results are summarized in Table 3. From the results, we can observe that: (i) Our approach again achieves the best performance, outperforming all baselines with a clear margin; and (ii) It is more difficult to conduct causal discovery when the latent DAGs are dynamic. The performance of all approaches degrades in this setting.

5.6 Performance on Decision Making


To answer RQ2, we compare proposed approach with representative methods in imitation learning, including Behavior Cloning (BC), Adversarial Inverse Reinforcement Learning (AIRL) [9] and GAIL [15]. Performance of our backbone model architecture is also reported as Vanilla. For the vanilla model, there is no causal discovery or causal encoding modules, and the policy model predicts actions based on the full observed states. Train, validation and test sets are split as 3 : 2 : 5. Trained policy models are tested on expert trajectories. For Kuramoto dataset we report MSE difference in predicted trajectories (smaller is better), and for Mimic-IV datasets we report the mean Accuracy and AUROC score (higher is better). For three gym datasets, we report the mean accuracy, mean reward, and macro-F score (higher is better). Results are summarized in Table 4, 5, and Table 6. From the three tables, we can observe that

Table 4: Action prediction performance measured in terms of prediction distance on Kuramoto datasets. Smaller is better

Methods	Kura5	Kura10	Kura5_vary	Kura10_vary
BC	0.0061	0.0069	0.0129	0.0117
AIRL	0.0062	0.0068	0.0133	0.0119
GAIL	0.0059	0.0067	0.0131	0.0128
Vanilla	0.0053	0.0065	0.0122	0.0114
Ours	0.0056	0.0063	0.0125	0.0115

Table 5: Action prediction performance measured in terms of AUC and Accuracy on MIMIC-IV datasets. Higher is better

Methods	Comorbidity		Sepsis	
	AUC	ACC	AUC	ACC
BC	0.8695	0.9507	0.9311	0.9735
AIRL	0.8574	0.9427	0.9216	0.9673
GAIL	0.8716	0.9585	0.9315	0.9774
Vanilla	0.8793	0.9618	0.9402	0.9818
Ours	0.8774	0.9628	0.9317	0.9817

**Figure 3: Sensitivity on weight of sparsity regularization, λ_1**

the proposed approach can provide an explanation with comparable performance in decision making.

5.7 Ablation Study


5.7.1 Analyzing Sparsity Regularization. In this subsection, we analyze the sensitivity of the proposed CAIL on hyperparameters λ_1 . λ_1 controls the importance of sparsity regularization term. We vary it as $\{10^{-6}, 10^{-5}, \dots, 0.1, 1\}$. Evaluations are conducted both on causal discovery (Kura10 and Kura10_vary) and on action prediction (LavaGap and FourRoom), with other configurations remaining the same as main experiment. Each experiment is conducted 3 times, and the average results in causal discovery are shown in Fig 3. From the figure, we can observe that the proposed CAIL performs relatively stable with $\lambda_1 \in [10^{-5}, 10^{-3}]$. Setting λ_1 to a too-large value, e.g. larger than 0.1, will result in a sharp drop in the quality of identified causal edges and prediction accuracy. This is because sparsity constraint can encourage the causal discovery module to remove uninformative edges, but setting it too large would remove correctly identified edges as well.

5.7.2 Analyzing Acyclicity Regularization. In this subsection, we analyze the sensitivity of CAIL on the importance of acyclicity regularization. Results are presented in Ap. C.

5.7.3 Influence of Training Instance Amount. In this subsection, we evaluate the influence of training size on causal discovery and action prediction quality, to obtain an idea of the amount of data needed for a successful learning process. We vary the ratio of instances for

Table 6: Action prediction performance measured on three gym datasets. Besides accuracy and MacroF, we also test the trained agent on the environment to obtain the mean rewards.

Methods	FourRoom			LavaGap			DoorKey		
	ACC	Reward	macroF	ACC	Reward	macroF	ACC	Reward	macroF
BC	0.9012	0.2836	0.6737	0.8745	0.9357	0.8634	0.8859	0.8403	0.8618
AIRL	0.8725	0.2617	0.6593	0.8733	0.9342	0.8629	0.8915	0.8672	0.8728
GAIL	0.8468	0.1904	0.6336	0.8659	0.9223	0.8605	0.8733	0.8328	0.8531
Vanilla	0.9037	0.2872	0.6795	0.8793	0.9464	0.8728	0.9035	0.8742	0.8813
Ours	0.9045	0.2863	0.6746	0.8821	0.9487	0.8745	0.9012	0.8756	0.8804

**Figure 4: Influence of training instance amount.****Table 7: Causal discovery performance in terms of AUCROC with different number or type of GNN layers used in enforcing causal graphs to the decision-making process.**

Layer Number	Kura10		Kura10_vary	
	GCN	Edge-aware	GCN	Edge-aware
1-layer	0.95	0.96	0.77	0.76
2-layer	0.71	0.98	0.64	0.78
3-layer	0.58	0.59	0.56	0.59

training as $\{1e-3, 5e-3, 1e-2, 5e-2, 0.1, 0.15, 0.2\}$, and experiment on datasets Kura10, Kura10_vary, LavaGap and FourRoom. The results are summarized in Fig 4. From the figure, we can observe that for the Kuramoto dataset, the benefit of increasing training examples is more clear with an amount less than 0.01 (corresponds to 1,000 transitions) in the static setting. In the dynamic setting, on the other hand, more training examples are needed for successful training. For datasets LavaGap and FourRoom, the performance become more stable with ratio larger then 0.1.

5.7.4 Influence of Model Architectures. Causal encoding module uses discovered causal graphs to guide the message propagation among variables, and back-propagates gradients to the causal discovery module so that the framework can be trained end-to-end jointly. In this subsection, we analyze the model’s performance w.r.t different architectures of it. Particularly, we test two GNN layers: GCN [22] which is one of the most popular GNN layers, and Edge-aware layer which used in the main experiments. Different numbers of GNN layers are also tested here. Experiments are conducted on Kura10 and Kura10_vary, with AUROC score on discovered causal edges reported. Results in causal discovery are summarized in Table 7. From the table, we observe that the performance drops quickly as the GCN goes deep, while the edge-aware layer used in this work performs well for both 1-layer and 2-layer settings. We attribute this observation to the “homophily” assumption of GCN, which is ineffective in modeling complex interactions of causal graphs.

5.8 Case Study

We further conduct case studies on dataset Kuramoto and Mimic-IV. Due to space limitations, we put it in Appendix. D and E.

6 CONCLUSION

In this work, we integrate causal discovery into imitation learning and propose a framework with improved interpretability. Besides learning control policies, the trained imitator is able to provide DAGs depicting captured dependence among state and action variables. With dynamic causal discovery module and the causality encoding module implemented as GNNs, the framework can model complex nonlinear causal relations. Experimental results on both simulation and real-world datasets show the effectiveness of the proposed method in capturing the casual relations for explanation and prediction. There are several interesting directions need further investigation. First, in this paper, we use clustering algorithm to cluster the states into stages and utilize it to supervise template selection. We would like to extend CAIL to learn the stages instead of relying on pre-clustered stages. Second, the identified causal relations could expose distribution shifts across domains. It is promising to utilize them for a more efficient transfer learning algorithm.

ACKNOWLEDGMENTS

This material is based upon work supported by, or in part by the National Science Foundation (NSF) under grant number IIS-1909702, Army Research Office (ARO) under grant number W911NF-21-1-0198, Department of Homeland Security (DHS) CINA under grant number E205949D, and Cisco Faculty Research Award.

REFERENCES

- [1] Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y Ng. 2007. An application of reinforcement learning to aerobatic helicopter flight. *Advances in neural information processing systems* 19 (2007), 1.
- [2] Mehmet Balcilar, Renee Van Eyden, Josine Uwilingiye, and Rangan Gupta. 2017. The impact of oil price on South African GDP growth: A Bayesian Markov switching-VAR analysis. *African Development Review* 29, 2 (2017), 319–336.
- [3] Steven L Bressler and Anil K Seth. 2011. Wiener–Granger causality: a well established methodology. *Neuroimage* 58, 2 (2011), 323–329.
- [4] Jon A Buras, Bernhard Holzmann, and Michail Sitkovsky. 2005. Animal models of sepsis: setting the stage. *Nature reviews Drug discovery* 4, 10 (2005), 854–865.
- [5] Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. 2018. Minimalistic Gridworld Environment for OpenAI Gym. <https://github.com/maximecb/gym-minigrid>.
- [6] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. 2018. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4693–4700.
- [7] Pim de Haan, Dinesh Jayaraman, and Sergey Levine. 2019. Causal confusion in imitation learning. *Advances in Neural Information Processing Systems* 32 (2019), 11698–11709.
- [8] M Sigfrido Rangel Frausto, Didier Pittet, Taekyu Hwang, Robert F Woolson, and Richard P Wenzel. 1998. The dynamics of disease progression in sepsis:

- Markov modeling describing the natural history and the likely impact of effective antiseptic agents. *Clinical infectious diseases* 27, 1 (1998), 185–190.
- [9] Justin Fu, Katie Luo, and Sergey Levine. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.
- [10] Clive WJ Granger. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society* (1969), 424–438.
- [11] Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. 2017. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 3389–3396.
- [12] David Hallac, Youngsuk Park, Stephen Boyd, and Jure Leskovec. 2017. Network inference via the time-varying graphical lasso. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 205–213.
- [13] William L. Hamilton, Zitao Ying, and J. Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *NIPS*.
- [14] Stefan Haufe, Klaus-Robert Müller, Guido Nolte, and Nicole Krämer. 2010. Sparse causal discovery in multivariate time series. In *Causality: Objectives and Assessment*. PMLR, 97–106.
- [15] Jonathan Ho and Stefano Ermon. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems* 29 (2016), 4565–4573.
- [16] Kaijen Hsiao and Tomas Lozano-Perez. 2006. Imitation learning of whole-body grasps. In *2006 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 5657–5662.
- [17] Tsung-Yu Hsieh, Yiwei Sun, Xianfeng Tang, Suhang Wang, and Vasant G Honavar. 2021. SrVARM: State Regularized Vector Autoregressive Model for Joint Learning of Hidden State Transitions and State-Dependent Inter-Variable Dependencies from Multi-variate Time Series. In *Proceedings of the Web Conference 2021*. 2270–2280.
- [18] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. 2017. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)* 50, 2 (2017), 1–35.
- [19] Junqi Jin, Chengru Song, Han Li, Kun Gai, Jun Wang, and Weinan Zhang. 2018. Real-time bidding with multi-agent reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 2193–2201.
- [20] Alistair Johnson, Lucas Bulgarelli, Tom Pollard, Steven Horng, Leo Anthony Celi, and R Mark IV. 2020. Mimic-iv (version 0.4). *PhysioNet* (2020).
- [21] Saurabh Khanna and Vincent YF Tan. 2019. Economy Statistical Recurrent Units For Inferring Nonlinear Granger Causality. In *International Conference on Learning Representations*.
- [22] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [23] Jens Kober and Jan Peters. 2009. Learning motor primitives for robotics. In *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2112–2118.
- [24] Shuang Li, Lu Wang, Ruizhi Zhang, Xiaofu Chang, Xuqin Liu, Yao Xie, Yuan Qi, and Le Song. 2020. Temporal Logic Point Processes. In *International Conference on Machine Learning*. PMLR, 5990–6000.
- [25] Ke Liang, Lingyuan Meng, Meng Liu, Yue Liu, Wenxuan Tu, Siwei Wang, Sihang Zhou, and Xinwang Liu. 2023. Learn from relational correlations and periodic events for temporal knowledge graph reasoning. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1559–1568.
- [26] Ke Liang, Sihang Zhou, Yue Liu, Lingyuan Meng, Meng Liu, and Xinwang Liu. 2023. Structure Guided Multi-modal Pre-trained Transformer for Knowledge Graph Reasoning. *arXiv preprint arXiv:2307.03591* (2023).
- [27] Sindy Löwe, David Madras, Richard Zemel, and Max Welling. 2020. Amortized causal discovery: Learning to infer causal graphs from time-series data. *arXiv preprint arXiv:2006.10833* (2020).
- [28] Daoming Lyu, Fangkai Yang, Bo Liu, and Steven Gustafson. 2019. SDRL: interpretable and data-efficient deep reinforcement learning leveraging symbolic planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 2970–2977.
- [29] Zibo Meng, Shizhong Han, Ping Liu, and Yan Tong. 2018. Improving speech related facial action unit recognition by audiovisual information fusion. *IEEE transactions on cybernetics* 49, 9 (2018), 3293–3306.
- [30] Alexander Mott, Daniel Zoran, Mike Chrzanowski, Daan Wierstra, and Danilo Jimenez Rezende. 2019. Towards Interpretable Reinforcement Learning Using Attention Augmented Agents. *Advances in Neural Information Processing Systems* 32 (2019), 12350–12359.
- [31] Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. 2020. Dynotears: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1595–1605.
- [32] Wei Qin, Zetong Chen, Lei Wang, Yunshi Lan, Weijie Ren, and Richang Hong. 2023. Read, Diagnose and Chat: Towards Explainable and Interactive LLMs-Augmented Depression Detection in Social Media. *arXiv preprint arXiv:2305.05138* (2023).
- [33] Weijie Ren, Lei Wang, Kumpeng Liu, Ruocheng Guo, Lim Ee Peng, and Yanjie Fu. 2022. Mitigating popularity bias in recommendation with unbalanced interactions: A gradient perspective. In *2022 IEEE International Conference on Data Mining (ICDM)*. IEEE, 438–447.
- [34] Weijie Ren, Pengyang Wang, Xiaolin Li, Charles E Hughes, and Yanjie Fu. 2022. Semi-supervised drifted stream learning with short lookback. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1504–1513.
- [35] Weijie Ren, Lei Zhang, Bo Jiang, Zhefeng Wang, Guangming Guo, and Guiquan Liu. 2017. Robust mapping learning for multi-view multi-label classification with missing labels. In *Knowledge Science, Engineering and Management: 10th International Conference, KSEM 2017, Melbourne, VIC, Australia, August 19-20, 2017, Proceedings 10*. Springer, 543–551.
- [36] Xiaoying Ren, Jing Jiang, Ling Min Serena Khoo, and Hai Leong Chieu. 2021. Cross-Topic Rumor Detection using Topic-Mixtures. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. 1534–1538.
- [37] Alex Tank, Ian Covert, Nicholas Foti, Ali Shojaie, and Emily Fox. 2018. Neural Granger Causality. *arXiv preprint arXiv:1802.05842* (2018).
- [38] Federico Tomasi, Veronica Tozzo, Saverio Salzo, and Alessandro Verri. 2018. Latent variable time-varying network inference. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2338–2346.
- [39] Faraz Torabi, Garrett Warnell, and Peter Stone. 2018. Behavioral cloning from observation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 4950–4957.
- [40] Marcel AJ Van Gerven, Babs G Taal, and Peter JF Lucas. 2008. Dynamic Bayesian networks as prognostic models for clinical patient management. *Journal of biomedical informatics* 41, 4 (2008), 515–529.
- [41] Lu Wang, Wenchao Yu, Xiaofeng He, Wei Cheng, Martin Renqiang Ren, Wei Wang, Bo Zong, Haifeng Chen, and Hongyuan Zha. 2020. Adversarial Cooperative Imitation Learning for Dynamic Treatment Regimes. In *Proceedings of The Web Conference 2020*. 1785–1795.
- [42] Bernard Widrow. 1964. Pattern-recognizing control systems. *Computer and Information Sciences* (1964).
- [43] Yue Yu, Jie Chen, Tian Gao, and Mo Yu. 2019. Dag-gnn: Dag structure learning with graph neural networks. In *International Conference on Machine Learning*. PMLR, 7154–7163.
- [44] Tom Zahavy, Nir Ben-Zrihem, and Shie Mannor. 2016. Graying the black box: Understanding dqns. In *International Conference on Machine Learning*. PMLR, 1899–1908.
- [45] Vinicius Zambaldi, David Raposo, Adam Santoro, Victor Bapst, Yujia Li, Igor Babuschkin, Karl Tuyls, David Reichert, Timothy Lillicrap, Edward Lockhart, et al. 2018. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830* (2018).
- [46] Tianxiang Zhao, Lemao Liu, Guoping Huang, Huayang Li, Yingling Liu, Liu GuiQuan, and Shuming Shi. 2020. Balancing quality and human involvement: An effective approach to interactive neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 9660–9667.
- [47] Tianxiang Zhao, Dongsheng Luo, Xiang Zhang, and Suhang Wang. 2023. Towards Faithful and Consistent Explanations for Graph Neural Networks. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 634–642.
- [48] Tianxiang Zhao, Wenchao Yu, Suhang Wang, Lu Wang, Xiang Zhang, Yuncong Chen, Yanchi Liu, Wei Cheng, and Haifeng Chen. 2023. Skill Disentanglement for Imitation Learning from Suboptimal Demonstrations. *arXiv preprint arXiv:2306.07919* (2023).
- [49] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. 2018. DAGs with NO TEARS: Continuous Optimization for Structure Learning. *Advances in Neural Information Processing Systems* 31 (2018).
- [50] Zhi Zheng, Shuvajit Das, Eric M Young, Amy Swanson, Zachary Warren, and Nilanjan Sarkar. 2014. Autonomous robot-mediated imitation learning for children with autism. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2707–2712.
- [51] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, Vol. 8. Chicago, IL, USA, 1433–1438.