

# Benchmarking of Data-Driven Causality Discovery Approaches in the Interactions of Arctic Sea Ice and Atmosphere

CyberTraining: Big Data + High-Performance Computing + Atmospheric Sciences

Yiyi Huang<sup>1</sup>, Matthäus Kleindessner<sup>2</sup>, Alexey Munishkin<sup>3</sup>, Debvrat Varshney<sup>4</sup>,  
Pei Guo<sup>4</sup>, Jianwu Wang<sup>4</sup>

<sup>1</sup>Department of Hydrology and Atmospheric Sciences, University of Arizona

<sup>2</sup>School of Computer Science & Engineering, University of Washington

<sup>3</sup>Department of Computer Science & Engineering, University of California, Santa Cruz

<sup>4</sup>Department of Information Systems, University of Maryland, Baltimore County

Technical Report HPCF-2020-16, [hpcf.umbc.edu](http://hpcf.umbc.edu) - > Publications

## Abstract

The Arctic sea ice has retreated rapidly in the past few decades, which is believed to be driven by various dynamic and thermodynamic processes in the atmosphere. The newly open water resulted from sea ice decline in turn exerts large influence on the atmosphere. Therefore, this study aims to investigate the causality between multiple atmospheric processes and sea ice variations using three distinct data-driven causality approaches: TCDF, NOTEARS and DAG-GNN. We find that the static graphs produced by NOTEARS and DAG-GNN are relatively reasonable. The results from NOTEARS indicate that relative humidity and precipitation dominate sea ice changes among all variables, while the results from DAG-GNN suggest that the horizontal wind fields are more important for driving sea ice variations. However, both of them produce some unrealistic edges. In comparison, the temporal graphs generated by the three methods are not physically meaningful enough. It also turns out that the results are rather sensitive to the choice of hyperparameters of the three methods. As a pioneer study, this work paves the way for us to disentangle the complex causal relationships in the Earth system, by taking the advantage of cutting-edge Artificial Intelligence technologies.

**Key words.** Causality Discovery, time series, Arctic Sea Ice, TCDF, NOTEARS, DAG-GNN.

## 1 Introduction

Warming in the Arctic has been much faster than in the rest of the world in both observations and model simulations, a phenomenon known as the Arctic amplification [25, 57]. Decline in sea ice is believed to be the major driver of Arctic amplification. Over the last few decades, Arctic summer sea ice extent has declined by nearly 50% with accelerated retreat in the early 21st century [58]. These dramatic changes in the Arctic sea ice affect a growing community of diverse stakeholders. Accompanying this growing interest is an urgent demand to increase the pace and scope of the advancements in physical understanding and predictive capabilities. As one of the most important components in the Earth System, the atmosphere actively interacts with the sea ice underneath. On the one hand, the sea ice variations are caused by different dynamic and thermodynamic forcings. On the other hand, sea ice decline in turn exerts large influence on the atmosphere. This will further alter the climate patterns in both Arctic and mid-latitudes, which results in more frequent extreme weather events [14, 64]. These two-way feedbacks are potentially very important in terms of understanding the Arctic warming in the past and future. Therefore, it is vital to analyze both the sea ice retreat's influence on the atmosphere and vice versa.

The traditional way to discover causal relations is to manipulate the value of a variable by using interventions or real-life experiments. All other influencing factors of the target variable can be

held fixed, to test whether a manipulation of a potential cause changes the target variable [44]. Specifically, the typical approach for assessing causal links in climate study is targeted modeling experiments. Such experiments are often computational expensive, time-consuming, or even impossible to carry out. More importantly, the large biases and substantial model spread remain in the state-of-the-art climate models [65], which further introduce some unrealistic causal relations. With the current advances in digital sensing and data assimilation, we have entered a period where earth science tends to be “data rich” in observations [49], allowing us to do data-driven causality discovery [22, 44, 73]. The data-driven causality approach aids scientists in identifying and extracting signals by analyzing statistical properties of purely observational data, which augments targeted model studies and has direct ties to forecasting and prediction. Several different frameworks for observational analysis have been applied to climate science to provide graphical representations of causal relations [10, 13, 20, 30, 42, 59, 60, 72]. Among them, the most relevant topic is the connections between Arctic (60°N northward) and mid-latitude (roughly 30-60°N) climate patterns [38, 41, 54, 62]. However, neither study investigates the relationship between sea ice retreat and the atmospheric dynamic and thermodynamic processes in the Arctic only, which is the focus of this study. It is unclear whether different causality approaches would produce similar results, or whether a particular technique is best suited for this topic as each study employs a different approach.

Thus, the overarching goal of this study is to investigate the causality between multiple atmospheric processes and sea ice variations using data-driven causality approaches. Instead of performing multiple climate model simulations, here we focus solely on an observational-type analysis. Specifically, three distinct data-driven causality approaches, Temporal Causality Discovery Framework (TCDF) [45], Non-combinatorial Optimization via Trace Exponential and Augmented lagRangian for Structure learning (NOTEARS) [74] and Directed Acyclic Graph-Graph Neural Networks (DAG-GNN) [71], will be used and compared to determine whether they are suitable for the particular climate study.

This report is structured in the following sections. Section 2 summarizes the main conclusions from previous studies in terms of causal relations between different atmospheric processes and Arctic sea ice variations; Section 3 lists data sets and data pre-processing methods and steps; Section 4 introduces three data-driven causality discovery frameworks; Section 5 summarizes the results generated by each method and compares those results with a causality graph based on literature review; finally, Section 6 reports the main conclusions and limitations of this study. The codes and data are available on <https://github.com/big-data-lab-umbc/cybertraining/tree/master/year-3-projects/team-6>.

## 2 Causality Between Atmospheric Processes and Arctic Sea Ice Variations

Due to the two-way interactions between the atmosphere and sea ice, studying causality between them is a challenging but important task, which makes it an area of high interest within polar climate community. The sea ice variations can be caused by different dynamical and thermodynamical processes. Important dynamical processes include anomalous surface wind [61, 70], regional atmospheric circulation patterns [48, 53] and abnormal summer storm activities [50]. Cloud [31], radiation [12, 34] and precipitation [3, 67] are the important thermodynamic factors Arctic sea ice trends and variability. On the other hand, sea ice decline in turn exerts large influence on the atmosphere, including cloud [34, 43], surface energy budget [4, 5, 56], precipitation [2, 37] and large-scale circulation [9, 35]. Figure 2.1 depicts the causal relations between key atmospheric variables and




Figure 2.1: The causality graph between key atmospheric variables and sea ice over the Arctic based on literature review. Note that the processes  $a - d$  are well-known atmospheric processes, which can be summarized from multiple textbooks. The processes  $e - i$  are summarized from recent peer-reviewed publications and they are ongoing research. The *sea\_ice* here represents sea ice coverage and/or sea ice thickness; *GH* is the geopotential height; *RH* is relative humidity; *SLP* means sea level pressure; *u10m* and *v10m* represents meridional and zonal wind at 10 meters, respectively; *HFLX* is the sensible and latent heat flux; *Precip* is the total precipitation; *CW* is the total cloud water path; *CC* is the total cloud cover; *SW* and *LW* represent net shortwave and longwave flux at the surface, respectively.

sea ice over the Arctic. The sea ice here represents sea ice coverage and/or sea ice thickness. Note that the processes  $a - d$  are well-known atmospheric processes, including cloud microphysics, thermodynamics, radiation, climate dynamics, which have been studied over the past few decades. The processes  $e - i$  are summarized from more recent publications, which are still under investigation by climate scientists. We will explain processes  $e - i$  in details in the next paragraph.

The arrow  $e$  represents the two-way effect between sea ice and net longwave flux at the surface. Based on global reanalysis [31], surface [15] and satellite observations [29], as well as model simulations [27, 32], the downwelling longwave radiation at the surface is found to enhance sea ice melt in spring, as a result of an increase in cloudiness and humidity in the Arctic Basin. Positive anomalies of longwave flux in spring and early summer initiate an earlier melt onset, hereby triggering several feedback mechanisms that amplify melt during the succeeding months [28, 32]. The sea ice melt increases the air temperature and thus increases the longwave flux at the surface. The downwelling shortwave flux, however, appears only important after melt has started [29, 31]. Once the surface albedo is significantly reduced due to sea ice melt, the tremendous solar radiation could be absorbed by ocean, which further accelerates ice melt in late spring and summer [12, 32, 34]. The feedback

between surface net shortwave flux and sea ice, represented by arrow  $f$  in Figure 2.1, has been confirmed by both model simulations [32] and satellite observations [12, 34, 47]. The arrows  $g$  indicate the interactions between the sea ice variations and atmospheric dynamical processes. A series of studies demonstrated that recent loss of Arctic sea ice is triggered by the atmospheric circulation changes such as a tendency toward a dipole pattern in the mean sea level pressure trend with an increase over the Arctic Ocean and a decrease over Siberia. The Arctic dipole anomaly in summer [68], winter [69] and spring [33] produces a strong meridional wind ( $v$ -component) anomaly that drives more sea ice out of the Arctic Ocean. In addition, this dipole anomaly promotes transport of heat and moisture and thus enhances downwelling longwave radiation and control the melt onset [28, 33]. A recent study also pointed out that a stronger anticyclonic circulation over Greenland and the Arctic Ocean in the troposphere may have contributed as much as 60% to the September sea-ice extent decline since 1979, by warming and moistening the lower atmosphere [18]. On the other hand, the reduction in Arctic sea ice extent and increase in open water area in late summer are found to directly contribute to a modification of large-scale circulation patterns in the following autumn through the additional heat stored in the Arctic Ocean and released to the atmosphere during the autumn [48]. The increased 1000-500 hPa thickness in autumn produce anomalous easterly zonal wind component ( $u$ -component), especially over the north of Alaska and Canada. Moreover, a more meridional flow pattern associated with sea ice reduction have an impact on the mid-latitude weather [48]. These conclusions are mainly drawn from model simulations [53, 69], reanalysis and observations [33, 48, 68]. In addition to radiation, the sensible and latent heat flux also plays an important role in the Arctic energy budget. Increases in the downwelling moisture flux triggers the melting of the sea ice in spring. Earlier melt onset and loss of sea ice in the spring enhance warming of the ice-free ocean surface, which in turn leads to an increase of evaporation from the surface into the atmosphere in the autumn. This positive feedback between heat flux and sea ice, indicated by arrow  $h$ , has been confirmed by satellite observations [4, 5] and model simulations [27] during most months of year. The arrow  $i$  represents the influence of precipitation on Arctic sea ice variations. Specifically, the magnitude of precipitation accumulating over the sea ice pack largely determines the depth of the snow layer, which modulates the rate of sea ice growth because of its highly insulating properties [63]. The phase of the precipitation falling on the sea ice pack is also important. As rain, it can instead melt, compact, and densify the snowpack, thus reducing the surface albedo and promoting sea ice melt [51]. These conclusions are mainly drawn from in-situ measurements during field campaign [51, 63] and global reanalysis products [3, 67]. The higher precipitation and snowfall could result in a thicker snowpack that allows less heat loss to the atmosphere. More importantly, modeling studies suggest that increases in Arctic precipitation over the twenty-first century, particularly in late autumn and winter, are due mainly to strongly intensified local surface evaporation (latent heat flux) [2]. Therefore, we believe that Arctic precipitation exerts direct influence on sea ice variations (arrow  $i$ ), while sea ice modulates precipitation mainly through sensible and latent heat flux (arrows  $h, b$ ).

Among these studies, very few of them have demonstrated the delayed impact of one variable on another. Specifically, the net shortwave flux at the surface in early summer (May-July) is found to enhance sea ice melt with a lag of 1 to 4 months [12]. Moreover, the sea ice condition exhibits the delayed impacts on itself, which is called sea ice predictability [16, 21, 26]. The sea ice predictability depends on the predictand (area, extent, volume), region, and the initial and target dates, which can be varied from a few days to a few years [21]. In general, there is predictability for the sea ice area in winter but low predictability throughout the rest of the year in peripheral seas. Based on multiple model simulations, the Labrador Sea stands out among the considered regions, with sea ice predictability extending up to 1.5 years [16].

Note that most of studies mentioned above determine the changes in one variables happening

before another one by applying time series analysis and/or composite analysis based on observations, reanalysis or model output [3, 4, 5, 12, 15, 28, 29, 31, 34, 47, 51, 63, 67, 68]. Among them, some studies use more advanced statistical analysis such as empirical orthogonal function [48, 69] and self-organizing map [33, 53]. Other studies assess the causal links through targeted modeling experiments [2, 16, 18, 27, 32], in order to test whether a manipulation of one variable has an impact on others. And most of the studies focus on relationships between only one or two atmospheric processes with changes in Arctic sea ice. Therefore, in this study, we target to provide a more comprehensive analysis about causality between multiple atmospheric processes and sea ice by applying different data-driven causality approaches.

### 3 Data Sets and Data Pre-processing

In this study, the sea ice concentration was obtained from the Nimbus-7 SSMR and DMSP SSM/I-SSMIS passive microwave data version 1 [8] provided by the National Snow and Ice Data Center (<http://nsidc.org/data/NSIDC-0051>). This dataset was generated from brightness temperature data, and provided daily in the polar stereographic projection with a grid box of 25 km x 25 km since October 1978. The uncertainty of sea ice concentration over the Arctic is within  $\pm 5\%$  during the winter, when the sea ice is relatively thick and the sea ice concentration is high. During the summer, the uncertainty increases to  $\pm 15\%$  when the melt ponds are present [7]. The sea ice extent is further filtered as the total area in the Arctic with sea ice concentration greater than 15% at daily time scale.

The atmospheric variables were obtained from ERA-5 global reanalysis product (<https://cds.climate.copernicus.eu/cdsapp#!/home>). ERA-5 was produced using 4D-Var data assimilation in CY41R2 of European Centre for Medium-Range Weather Forecasts (ECMWF)'s Integrated Forecast System (IFS), with 137 hybrid sigma/pressure (model) levels in the vertical, with the top level at 0.01 hPa [6]. The ECMWF global reanalysis product has been evaluated over the Arctic in the previous studies and it stands out among several global reanalysis products as being more consistent with independent observations [17, 39]. In this study, the variables used in the causality discovery algorithms are listed in Table 3.1. The air temperature has been excluded in this study because it exhibits very high correlation with sea ice concentration. The interactions between air temperature and sea ice could be dominant over all other atmospheric processes based on our tests.

All monthly gridded data during 1980-2018 has been averaged over the Arctic north of 60°N using area-weighted method. Therefore, we created the time series for both sea ice extent and atmospheric variables. In order to eliminate overall impacts of global warming and seasonality during this 39-year time period, we applied detrending and deseasonalizing for each time series. Here we assume the time series is additive and there exist both trend and seasonal components, that is

$$X_t = m_t + s_t + Y_t, \quad (3.1)$$

The time series has been detrended by subtracting the line of best fit from the time series, where the line of best fit was obtained from a linear regression model with the time steps as the predictor. To deseasonalize the time series, we divided the averaged seasonal index from the time series. The seasonal index were calculated from moving averages with 12-month seasonal window in this study [23]. Here, we only kept residual component  $Y_t$ , which fluctuates zero, i.e.,

$$E(Y_t) = 0 \quad (3.2)$$

Then we normalized the  $Y_t$  using max-min method, therefore

$$Y_t \in [0, 1] \tag{3.3}$$

Table 3.1: The atmospheric and sea ice variables considered in this study.

Abbrev.in Fig.2.1	Variable
GH	Geopotential heights averaged from 200 hPa, 500 hPa and 850 hPa
RH	Relative humidity averaged from 1000-300 hPa
SLP	Sea level pressure
u10m	Zonal (u-component) wind at 10 meters
v10m	Meridional (v-component) wind at 10 meters
HFLX	Sensible and latent heat flux
Precip	Total precipitation
CC	Total cloud cover
CW	Total cloud water path
SW	Net shortwave flux at the surface
LW	Net longwave flux at the surface
Sea_ice	Sea ice extent in the Northern Hemisphere

## 4 Data-Driven Causality Discovery Algorithms

In this work, we apply data-driven causal discovery algorithms aiming to find the major causes of some underlying process such as the decrease of Arctic Sea Ice. These algorithms typically assume one process or state, a cause, contributes to the production of another process or state, an effect. The cause is assumed to be partly responsible for the effect, and the effect is partly dependent on the cause. As such the cause affects the effect, but not necessarily the effect affects the cause. Causal discovery aims to discover direct cause-effect relationships for both instantaneous and delayed causes. Here we will investigate three recently proposed causal discovery algorithms: TCDF [45], NOTEARS [74], and DAG-GNN [71].

### 4.1 TCDF

The TCDF algorithm [45] uses a Temporal Causal Discovery Framework (TCDF) that is based on attention-based CNN. The input to the algorithm are time series data and the output produces a causality graph structure with time delay (lag). Figure 4.1 illustrates the architecture of the TCDF method where on the left is the multi-dimensional time-series data and on the right is the produced causality graph. There are four steps to learn a Temporal Causal Graph from the time-series data: Time Series Prediction block, Attention Interpretation block, Causal Validation and Delay Discovery blocks. For multi-dimensional time-series there are  $n$  independent attention-based CNNs, all with the same architecture for each time-series data.

The basic structure of TCDF is for time-series prediction as seen in the first step of the framework in Figure 4.1. After predicting time-series, the output gives attention scores for the attention

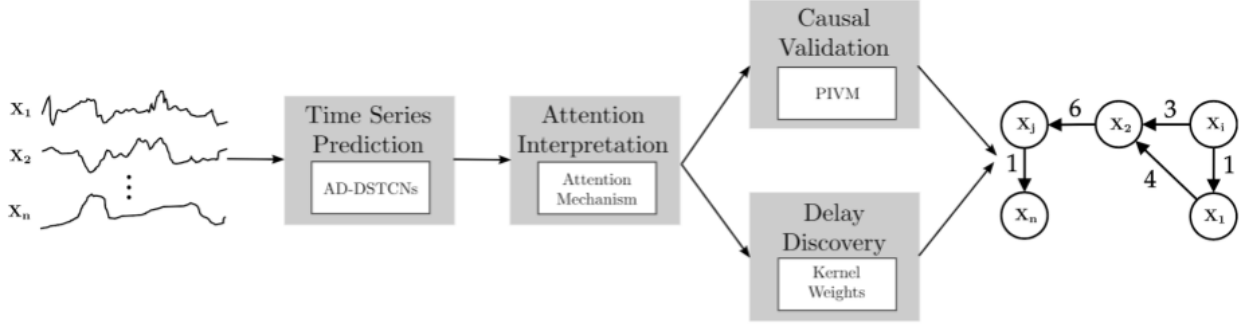


Figure 4.1: Architecture of TCDF [45]

interpretation mechanism: to learn to which time-series a CNN attends to when predicting a time series. Attention is when we let the model learn what to attend to based on the input data and what it has learnt so far. The causality validation reads the final result of the attention scores and applies a permutation importance validation method. The permutation importance is a measurement of how much an error will affect the values of a certain attention score when all scores are randomly permuted. The idea is that permuting a time-series attention score removes potential cause and effect relationships and hence the method can detect real verse fake causal relationships. In parallel the attention scores are feed to the delay discovery to learn potential delay cause and effect relationships.

Another major advantage of the TCDF is to using a CNN verse a traditional RNN such as a LSTM for time-series data. The advantage is that RNN typically have a vanishing gradient problem: long-term information has to sequentially travel through all cells before getting to the present processing cell, which is greatly amplified when the number of layers becomes deep or very deep, typically 10 layers or more [55]. Though a CNN structure might have this problem as well, it is more common in RNN because RNN typically needs much more memory and cells than a CNN structure. With more cells to process, there is a greater chance of obtaining the vanishing gradient problem.

## 4.2 NOTEARS

The NOTEARS algorithm [74] assumes a linear data generating model of the form

$$X_i = \sum_{j:W_{ji} \neq 0} W_{ji}X_j + N_i, \quad (4.1)$$

where  $W$  is the weighted adjacency matrix of the underlying causality graph  $G(W)$ , that is  $j \rightarrow i$  in  $G(W)$  if and only if  $W_{ji} \neq 0$ , and the random variables  $N_i$  are independent noise variables. Given  $n$  i.i.d. observations of the variables  $X_1, \dots, X_d$ , written as matrix  $\mathbf{X} \in \mathbb{R}^{n \times d}$ , a standard estimator for  $W$  is the (regularized) least-squares estimator

$$\widehat{W} = \arg \min_{W \in \mathbb{R}^{d \times d}} \frac{1}{2n} \|\mathbf{X} - \mathbf{X}W\|_F^2 + \lambda \|W\|_1 \quad \text{subject to } G(W) \text{ is a DAG}, \quad (4.2)$$

where  $\lambda \geq 0$  is the regularization parameter. This estimator is theoretically well-studied and satisfies desirable properties such as consistency [1, 40, 66]. However, due to the non-convex, *combinatorial-like* constraint, optimization problems of the form (4.2) are NP-hard to solve [11],

and hence unless the number of variables  $d$  is very small, heuristics such as local search have to be applied [e.g., 24, 52]. The NOTEARS algorithm builds on the insight that

$$G(W) \text{ is a DAG} \Leftrightarrow \text{trace}(\exp(W \circ W)) - d = 0, \quad (4.3)$$

where  $\exp$  denotes the matrix exponential and  $\circ$  the element-wise product. The characterization (4.3) allows to treat the optimization problem (4.2) as an ordinary *continuous* constrained optimization problem and to use any algorithm from the rich literature on continuous optimization to find a locally optimal solution to (4.2). Concretely, the NOTEARS algorithm applies the augmented Lagrangian method [e.g., 46] to search for a locally optimal solution to

$$\arg \min_{W \in \mathbb{R}^{d \times d}} \frac{1}{2n} \|\mathbf{X} - \mathbf{X}W\|_F^2 + \lambda \|W\|_1 \quad \text{subject to} \quad \text{trace}(\exp(W \circ W)) - d = 0. \quad (4.4)$$

After applying the augmented Lagrangian method to (4.4) and obtaining an output  $\widetilde{W}$ , the final step of the NOTEARS algorithm is to “round”  $\widetilde{W}$  and to set all entries of  $\widetilde{W}$  with absolute value smaller than some threshold  $t$  to zero. This yields the final output  $\widehat{W}$  of the NOTEARS algorithm.

### 4.3 DAG-GNN

DAG-GNN [71] can be thought of as an extension to the NOTEARS algorithm [74] in that the proposed method assumes a nonlinear model of the form

$$X = f_2((I - A^T)^{-1} f_1(Z)) \quad (4.5)$$

where  $Z$  is the encoded latent variable of  $X$ . This can be contrasted to the linear model assumed in NOTEARS

$$X = (I - A^T)^{-1} Z \quad (4.6)$$

where Equation (4.6) is a restructured form of Equation (4.1). Further, DAG-GNN builds an inference model to encode  $Z$ , given by

$$Z = f_4((I - A^T) f_3(X)) \quad (4.7)$$

where  $f_3$  and  $f_4$  play a conceptually inverse role for  $f_2$  and  $f_1$  respectively. In particular, this paper assumes  $f_1, f_4$  to be identity functions and  $f_2, f_3$  as Multilayer Perceptrons (MLP). Since MLP is nonlinear, it should in theory capture any nonlinearities in the data better than NOTEARS which is a linear model. This is further explained in Figure 4.2 where  $\hat{X}$  is the regenerated form of  $X$  and MLP has one hidden layer of 64 neurons.




Figure 4.2: Architecture of DAG-GNN [71]



Further, DAG-GNN minimizes the following loss function

$$\min_{A, \theta} f(A, \theta) = -L_{ELBO} \quad (4.8)$$

$$\text{s.t. } h(A) = \text{tr}[(I + \alpha A \circ A)^m] - m = 0 \quad (4.9)$$

where the unknowns include the weight matrix  $A$ , and parameters  $\theta$  for Variational Autoencoder (VAE). Further, ELBO is the Evidence Lower Bound of the VAE adopted from [36] and  $h(A)$  is used to solve the augmented Lagrangian as done in [46].

#### 4.4 Static versus Temporal Model

While TCDF requires time series data as input and explicitly models time delay of causal relations, NOTEARS and DAG-GNN assume to be provided i.i.d. observations of the variables. Similarly to other causal discovery studies in climate research [20], we apply the latter two methods in two ways: in the *static model*, we treat the observations of the variables summarized in Table 3.1 at different points in time as i.i.d. observations and directly feed the data into the two methods. Alternatively, in the *temporal model*, we first augment the data set by adding lagged versions of each variable, that is for each variable  $X$  we additionally consider variables  $X^1, X^2, \dots, X^{12}$ , where  $X^k$  is a version of the variable  $X$  that is measured with a lag of  $k$  time units (in our case: months) compared to  $X$ . We then treat the observations of the various variables at different points in time as i.i.d. observations and run NOTEARS and DAG-GNN, respectively. The graphs produced by these methods, using the augmented data, are assumed to encode the time delay of causal relations. However, in order to obtain a causal graph on the variables of Table 3.1 we generate a “reduced” temporal graph from these “full” temporal graphs by connecting two variables  $X$  and  $Y$  in the reduced temporal graph whenever any of the variables  $X, X^1, \dots, X^{12}$  is connected to any of  $Y, Y^1, \dots, Y^{12}$  in the full temporal graph. The reduced temporal graph is the output of the temporal model.

## 5 Results

In this section we present some results of the three causal discovery algorithms introduced in Section 4. We study how the causality graphs produced by the three methods depend on the choice of hyperparameters and see that the graphs can be quite different for varying hyperparameters. In Section 5.1, we work with the normalized Hamming distance (ignoring the edge weights in the graphs produced by NOTEARS or DAG-GNN) and also compare all graphs to the domain knowledge graph of Figure 2.1. In Section 5.2, we work with the  $l_1$ -distance, which takes edge weights into account. Since for unweighted graphs, the normalized Hamming distance and the  $l_1$ -distance are equivalent (up to the normalization), we do not consider TCDF, which produces unweighted graphs, in that section and also do not compare to the domain knowledge graph of Figure 2.1.

### 5.1 Comparison of the Results for Different Hyperparameters and to the Domain Knowledge Graph with Respect to the Normalized Hamming Distance

In this section, we treat all graphs as unweighted graphs. The normalized Hamming distance is a widely used metric to compare two unweighted graphs on the same set of vertices [19]. Let  $A, B \in \{0, 1\}^{m \times m}$  be the adjacency matrices of two unweighted graphs  $G_A, G_B$  on  $m$  vertices. The normalized Hamming distance between  $G_A$  and  $G_B$  is given by  $\text{dist}_{HD}(G_A, G_B) = \frac{1}{m^2} \sum_{i,j=1}^m \mathbb{1}\{A_{ij} \neq B_{ij}\}$ .

Table 5.1: Distance matrix with respect to the normalized Hamming distance for TCDF. ♣ denotes  $layer = 0, kernel = 4$  are the algorithm’s default hyperparameters. The bottom row compares to the domain knowledge graph of Figure 2.1 (best values in bold).

		<i>Temporal</i>					
		<i>layer = 0</i> <i>kernel = 2</i>	<i>layer = 0</i> <i>kernel = 4♣</i>	<i>layer = 0</i> <i>kernel = 6</i>	<i>layer = 1</i> <i>kernel = 2</i>	<i>layer = 1</i> <i>kernel = 4</i>	<i>layer = 1</i> <i>kernel = 6</i>
<i>Temporal</i>	<i>layer = 0, kernel = 2</i>	0	0.05	0.01	0.02	0.01	0.01
	<i>layer = 0, kernel = 4♣</i>	0.05	0	0.06	0.07	0.06	0.06
	<i>layer = 0, kernel = 6</i>	0.01	0.06	0	0.01	0.01	0.01
	<i>layer = 1, kernel = 2</i>	0.02	0.07	0.01	0	0.02	0.02
	<i>layer = 1, kernel = 4</i>	0.01	0.06	0.01	0.02	0	0
	<i>layer = 1, kernel = 6</i>	0.01	0.06	0.01	0.02	0	0
	Domain knowl.	0.35	<b>0.33</b>	0.34	0.34	<b>0.33</b>	<b>0.33</b>

In the following, for each of the three causal discovery algorithms introduced in Section 4 we compute the normalized Hamming distance between the graphs produced by an algorithm for different values of its hyperparameters. We also compare the graphs to the domain knowledge graph shown in Figure 2.1 which is generated based on current literature.

### 5.1.1 TCDF

Table 5.1 shows the values for the normalized Hamming distance for the TCDF method. Two hyperparameters that were chosen are the kernel size and the number of hidden layers. The number of hidden layers corresponds to the number of hidden CNN layers in the TCDF algorithm. It seems that the addition of a hidden layer leads to far worse results and even produces no causality graphs as is the case for when  $kernel = 4$  and  $kernel = 6$  for  $layer = 1$ . The kernel size is related to how much the TCDF method lags the variables for the causality study. The default setting for the hyperparameters as shown with ♣ in Table 5.1 produce the best result, but that is nowhere close to the domain knowledge graph shown in Figure 2.1. It seems that the TCDF method does not produce good results for our Arctic Sea Ice data.

**Comparison to domain knowledge graph.** Since the TCDF focuses on the time series, only the temporal graph that are closest to the domain knowledge graph is shown (Figure 5.1). There is no cause and effect between sea ice and any atmospheric variables. As for the causality within the atmosphere, only a few edges are generated by TCDF algorithm. Among them, the feedback between u10m and v10m, as well as the impact of SW on CW are not consistent with domain knowledge.

Table 5.2: Distance matrix with respect to the normalized Hamming distance for NOTEARS. ♣ denotes that  $\lambda = 0.1, t = 0.3$  are the algorithm’s default hyperparameters. The bottom row compares to the domain knowledge graph of Figure 2.1 (best values in bold).

		<i>Static</i>				<i>Temporal</i>			
		$\lambda = 0$ $t = 0.2$	$\lambda = 0$ $t = 0.3$	$\lambda = 0.1$ $t = 0.2$	$\lambda = 0.1$ $t = 0.3$ ♣	$\lambda = 0$ $t = 0.2$	$\lambda = 0$ $t = 0.3$	$\lambda = 0.1$ $t = 0.2$	$\lambda = 0.1$ $t = 0.3$ ♣
<i>Static</i>	$\lambda = 0, t = 0.2$	0.0	0.02	0.15	0.15	0.54	0.36	0.16	0.15
	$\lambda = 0, t = 0.3$	0.02	0.0	0.15	0.12	0.53	0.35	0.14	0.12
	$\lambda = 0.1, t = 0.2$	0.15	0.15	0.0	0.02	0.51	0.36	0.09	0.1
	$\lambda = 0.1, t = 0.3$ ♣	0.15	0.12	0.02	0.0	0.52	0.35	0.07	0.08
<i>Temporal</i>	$\lambda = 0, t = 0.2$	0.54	0.53	0.51	0.52	0.0	0.18	0.48	0.51
	$\lambda = 0, t = 0.3$	0.36	0.35	0.36	0.35	0.18	0.0	0.33	0.34
	$\lambda = 0.1, t = 0.2$	0.16	0.14	0.09	0.07	0.48	0.33	0.0	0.03
	$\lambda = 0.1, t = 0.3$ ♣	0.15	0.12	0.1	0.08	0.51	0.34	0.03	0.0
Domain knowl.		0.35	<b>0.33</b>	0.36	0.35	0.54	0.46	0.37	<b>0.35</b>

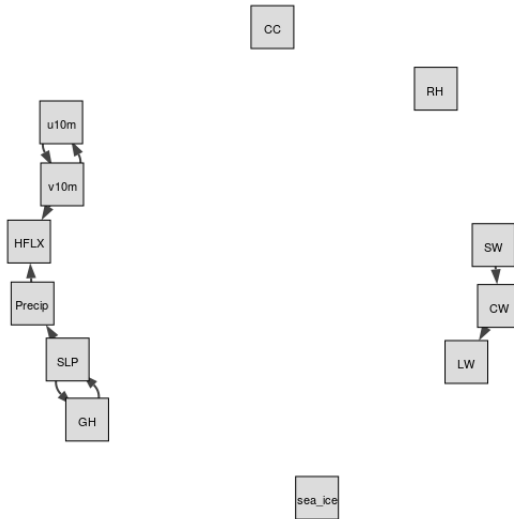


Figure 5.1: The TCDF graph that is closest to the domain knowledge graph of Figure 2.1. The temporal graph for  $layer = 0, kernel = 4$ .

### 5.1.2 NOTEARS

The NOTEARS algorithm has two hyperparameters  $\lambda \geq 0$  and  $t \geq 0$  as explained in Section 4.2: the parameter  $\lambda$  is the regularization parameter (cf. Equation (4.2)) and  $t$  is the threshold for setting edge weights of the preliminary output to zero (cf. end of Section 4.2). There is no default value for  $\lambda$ , but in the main experiment that comes with the NOTEARS code [74], the authors use  $\lambda = 0.1$  and hence we consider that value to be the default value. Furthermore, we observed that choosing a value larger than 0.1 for  $\lambda$  often results in an empty graph as the output of NOTEARS. The default value for  $t$  is  $t = 0.3$ . Indeed, we observed that  $t = 0.3$  yields better results than other values of  $t$ .

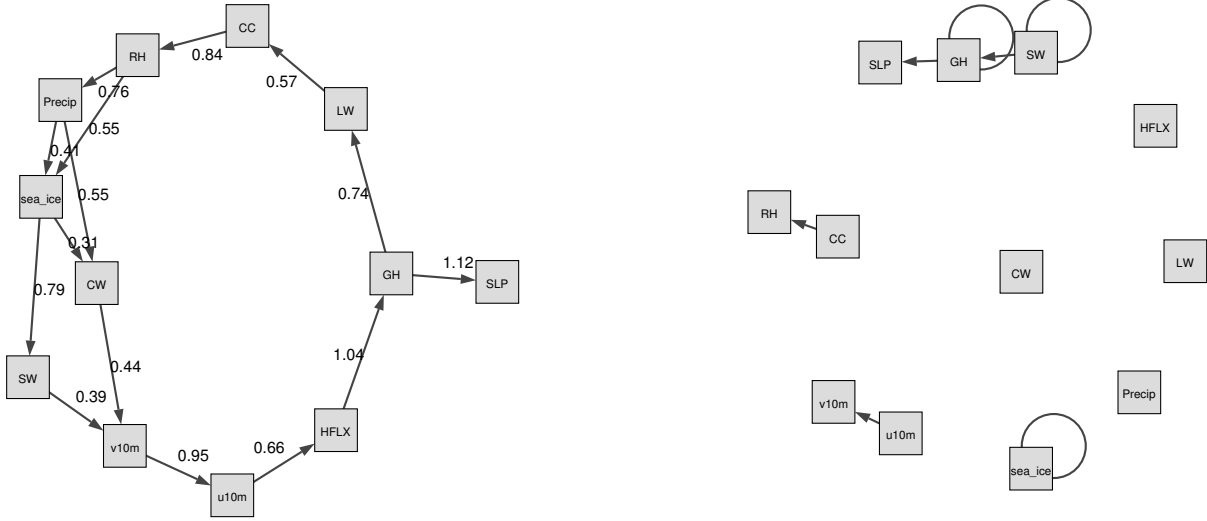


Figure 5.2: The NOTEARS graphs that are closest to the domain knowledge graph of Figure 2.1, with respect to the normalized Hamming distance as shown in Table 5.2. **Left:** The static graph for  $\lambda = 0, t = 0.3$ . Note that in Table 5.2 this graph is treated as an unweighted graph. **Right:** The temporal graph for  $\lambda = 0.1, t = 0.3$ .

Table 5.2 shows the normalized Hamming distance between the graphs produced by NOTEARS for  $\lambda \in \{0, 0.1\}$  and  $t \in \{0.2, 0.3\}$ , for both the static and the temporal model. The last row of the table shows the normalized Hamming distance between the various graphs and the domain knowledge graph of Figure 2.1. Note that actually none of the graphs considered here is a DAG. We can see that the distances between the various temporal graphs (middle right part of the table) are significantly larger than the distances between the various static graphs (upper left part of the table). We can also see that changing the value of  $\lambda$  from 0 to 0.1 causes a larger difference in the result than changing the value of  $t$  from 0.2 to 0.3 (e.g., the normalized Hamming distance between the static model with  $\lambda = 0, t = 0.2$  and the static model with  $\lambda = 0, t = 0.3$  is only 0.02, while the distance between the static model with  $\lambda = 0, t = 0.2$  and the static model with  $\lambda = 0.1, t = 0.2$  is 0.15).

**Comparison to domain knowledge graph.** In Figure 5.2, we show both the static and the temporal graphs that are closest to the domain knowledge graph of Figure 2.1. Generally, while none of the produced graphs is really close to the domain knowledge graph, the static graph looks more reasonable.

In the static graph, the RH and precipitation seem to dominate the sea ice changes, with weights of 0.55 and 0.41, respectively. In the meantime, the sea ice exerts large influence on SW (weight of 0.79) and CW (weight of 0.31). The causal relations between precipitation, SW and sea ice in the domain knowledge graph of Figure 2.1 are well captured by the NOTEARS algorithm. However, RH and CW are believed to be only indirectly connected with sea ice changes (i.e., in the domain knowledge graph there is no direct connection between RH or CW and sea ice), but in the static graph produced by the NOTEARS algorithm there are direct connections. The causality between each two atmospheric variables is generally reasonable based on the domain knowledge graph, except for the connections between CW and v10m, SW and v10m, and v10m and u10m.

Compared with the static graph, the temporal graph detects only very few edges. It shows

Table 5.3: Distance matrix with respect to the normalized Hamming distance for DAG-GNN. ♣ denotes the algorithm’s default hyperparameters. The bottom row compares to the domain knowledge graph of Figure 2.1 (best values in bold).

		<i>Static</i>				<i>Temporal</i>			
		$\tau = 0$		$\tau = 10^{-7}$		$\tau = 0$		$\tau = 10^{-7}$	
		$t = 0.2$	$t = 0.3$ ♣	$t = 0.2$	$t = 0.3$	$t = 0.2$	$t = 0.3$ ♣	$t = 0.2$	$t = 0.3$
<i>Static</i>	$\tau = 0, t = 0.2$	0.0	0.06	0.04	0.07	0.1	0.12	0.1	0.12
	$\tau = 0, t = 0.3$ ♣	0.06	0.0	0.05	0.01	0.08	0.07	0.08	0.07
	$\tau = 10^{-7}, t = 0.2$	0.04	0.05	0.0	0.06	0.08	0.1	0.08	0.1
	$\tau = 10^{-7}, t = 0.3$	0.07	0.01	0.06	0.0	0.08	0.08	0.08	0.08
<i>Temporal</i>	$\tau = 0, t = 0.2$	0.1	0.08	0.08	0.08	0.0	0.03	0.01	0.03
	$\tau = 0, t = 0.3$ ♣	0.12	0.07	0.1	0.08	0.03	0.0	0.05	0.0
	$\tau = 10^{-7}, t = 0.2$	0.1	0.08	0.08	0.08	0.01	0.05	0.0	0.05
	$\tau = 10^{-7}, t = 0.3$	0.12	0.07	0.1	0.08	0.03	0.0	0.05	0.0
Domain knowl.		0.33	0.33	0.35	<b>0.32</b>	0.35	<b>0.34</b>	0.36	<b>0.34</b>

that the sea ice, SW and GH have delayed impacts on themselves, demonstrating both sea ice and atmosphere have a degree of seasonal to year long climate predictability. As in the static graph, the causality between v10m and u10m is not reasonable.

### 5.1.3 DAG-GNN

Like the NOTEARS algorithm, DAG-GNN has two hyperparameters:  $\tau \geq 0$  and  $t \geq 0$ , where  $\tau$  is similar to  $\lambda$  used in NOTEARS. We noticed that the hyperparameter  $\tau$  is very sensitive, and show the outputs for only two values of  $\tau$  i.e.  $\tau \in \{0, 10^{-7}\}$ . We vary  $t$  similar to NOTEARS, and test for  $t \in \{0.2, 0.3\}$ . Table 5.3 tabulates the normalized Hamming Distance computed between all the graphs obtained by varying these two hyperparameters. Further, we computed the Normalized Hamming Distance between all these graphs and the domain knowledge graph of Figure 2.1. In order to carry out this specific computation, we created unweighted matrices from the weighted outputs of DAG-GNN with the help of absolute thresholding using the hyperparameter  $t$ .

From Table 5.3, we see that the least normalized Hamming Distance with the Domain Knowledge Graph is obtained by  $\tau = 10^{-7}$ ,  $t = 0.3$  for the static model and  $\tau \in \{0, 10^{-7}\}$ ,  $t = 0.3$  for the temporal model. Both the temporal models of  $t = 0.3$  give the same graphs, which is shown in Figure 5.3 on the right. For the static model however, these optimum values of  $\tau$  and  $t$  produce a graph which shows no relation between sea ice and other atmospheric variables. Hence, we reject it. The second most optimum graph showing a dependence of sea ice with atmospheric variables is shown in Figure 5.3 on the left. Its hyperparameters are  $\tau = 0$  and  $t = 0.2$ . Further, we note that for the temporal model, the best graphs are obtained with  $t = 0.3$ , which is one of the default parameters used and suggested by the authors in [71].

**Comparison to domain knowledge graph.** The static and the temporal graphs closest with the domain knowledge graph are shown in Figure 5.3. Compared to NOTEARS, both static and temporal graphs produced by DAG-GNN are more complicated. The dynamical fields (u10m and v10m) dominate the sea ice changes, but with relatively small weights 0.21 and  $-0.25$ , respectively. As for the causal relations between multiple atmospheric processes, they are generally reasonable

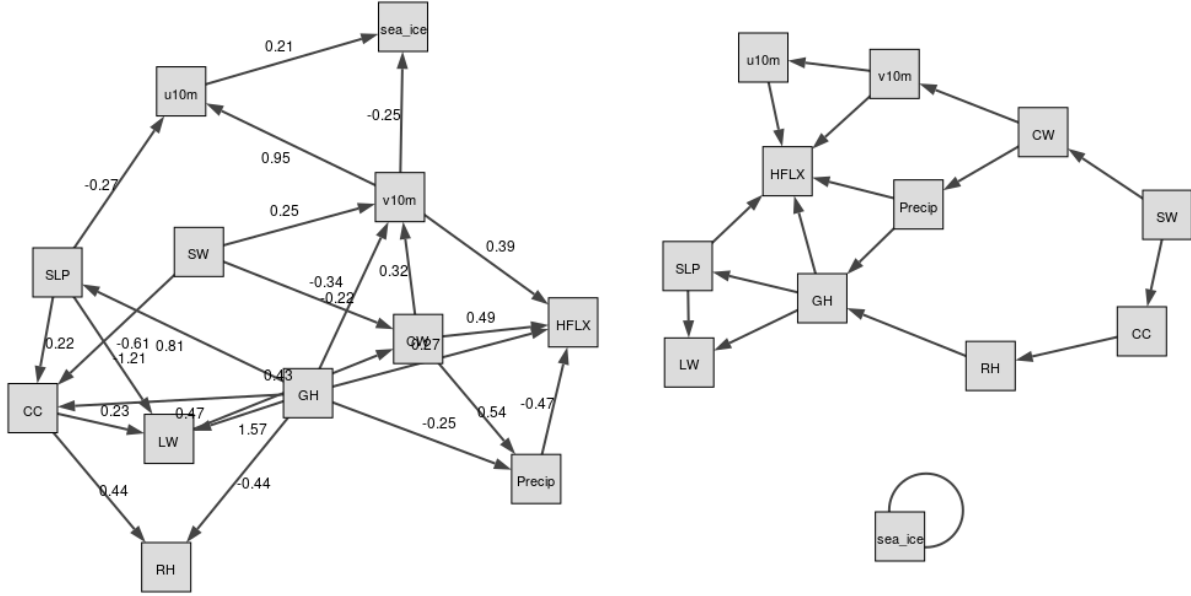


Figure 5.3: The DAG-GNN graphs that are closest to the domain knowledge graph of Figure 2.1. **Left:** The static graph for  $\tau = 0, t = 0.2$ . **Right:** The temporal graph for  $\tau \in \{0, 10^{-7}\}, t = 0.3$ .

compared to Figure 2.1. Similar as NOTEARS, the connections between u10m and v10m (0.95), SW and v10m (0.25), CW and v10m (0.32) do not make any physical sense. Note that the increased CW and CC tend to reflect solar radiation back to the space, leaving less SW reaching at the surface. This negative relationships between CC, CW and SW are captured by the DAG-GNN, however, the direction of arrows are not meaningful. The same issue occurs in the temporal graph. In addition, the sea ice has the delayed impacts on itself, but with no connection with any atmospheric processes in temporal graph. Similar as static graph, the causality between u10m and v10m as well as CW and v10m is not realistic.

## 5.2 Comparison of the Results for Different Hyperparameters with Respect to the $l_1$ -Distance

In case of the static model, the NOTEARS algorithm and the DAG-GNN method output weighted graphs, where an edge weight represents the strength of a causal effect. In order to also see how the edge weights change as we change the hyperparameters of NOTEARS and DAG-GNN, in this section we present analogous results as in Section 5.1, but with the  $l_1$ -distance instead of the normalized Hamming distance. If  $A, B \in \mathbb{R}^{m \times m}$  are the adjacency matrices of two weighted graphs  $G_A, G_B$  on  $m$  vertices, the  $l_1$ -distance between  $G_A$  and  $G_B$  is given by  $dist_{l_1}(G_A, G_B) = \sum_{i,j=1}^m |A_{ij} - B_{ij}|$ .

### 5.2.1 NOTEARS

Table 5.4 shows the  $l_1$ -distance between the graphs produced by NOTEARS for  $\lambda \in \{0, 0.1\}$  and  $t \in \{0.2, 0.3\}$ , where the graphs for the static model are weighted graphs and the graphs for the temporal model are unweighted (cf. Figure 5.2). Since for unweighted graphs, the  $l_1$ -distance and the normalized Hamming distance are equivalent up to the normalization factor used in the latter distance, the bottom right part of Table 5.4 coincides with the middle right part of Table 5.2 up

to a factor of  $1/12^2 = 1/144$ . The most interesting part of Table 5.4 is the upper left part, which compares the various static graphs with each other. As for the normalized Hamming distance, we can see that changing the value of  $\lambda$  from 0 to 0.1 causes a larger difference in the result than changing the value of  $t$  from 0.2 to 0.3 (e.g., the  $l_1$ -distance between the static model with  $\lambda = 0, t = 0.2$  and the static model with  $\lambda = 0, t = 0.3$  is only 0.8, while the  $l_1$ -distance between the static model with  $\lambda = 0, t = 0.2$  and the static model with  $\lambda = 0.1, t = 0.2$  is 12.54).

Table 5.4: Distance matrix with respect to the  $l_1$ -distance for NOTEARS. ♣ denotes that  $\lambda = 0.1, t = 0.3$  are the algorithm’s default hyperparameters.

		<i>Static</i>				<i>Temporal</i>			
		$\lambda = 0$ $t = 0.2$	$\lambda = 0$ $t = 0.3$	$\lambda = 0.1$ $t = 0.2$	$\lambda = 0.1$ $t = 0.3$ ♣	$\lambda = 0$ $t = 0.2$	$\lambda = 0$ $t = 0.3$	$\lambda = 0.1$ $t = 0.2$	$\lambda = 0.1$ $t = 0.3$ ♣
<i>Static</i>	$\lambda = 0, t = 0.2$	0.0	0.8	12.54	12.37	77.58	51.58	16.73	14.2
	$\lambda = 0, t = 0.3$	0.8	0.0	12.27	11.58	77.36	51.36	15.93	13.41
	$\lambda = 0.1, t = 0.2$	12.54	12.27	0.0	0.69	77.34	52.46	11.24	9.8
	$\lambda = 0.1, t = 0.3$ ♣	12.37	11.58	0.69	0.0	77.6	52.29	10.55	9.11
<i>Temporal</i>	$\lambda = 0, t = 0.2$	77.58	77.36	77.34	77.6	0.0	26.0	69.0	73.0
	$\lambda = 0, t = 0.3$	51.58	51.36	52.46	52.29	26.0	0.0	47.0	49.0
	$\lambda = 0.1, t = 0.2$	16.73	15.93	11.24	10.55	69.0	47.0	0.0	4.0
	$\lambda = 0.1, t = 0.3$ ♣	14.2	13.41	9.8	9.11	73.0	49.0	4.0	0.0

## 5.2.2 DAG-GNN

Table 5.5 tabulates the  $l_1$  distance between all the static and temporal outputs of DAG-GNN. As expected, the static graphs are closer to each other (top left section of Table 5.5 has lower values than the top right section) and temporal graphs are closer amongst themselves (bottom right section of Table 5.5 has lower values than the bottom left section). Also, comparing to NOTEARS algorithm, the differences between two graphs generated by different hyperparameters are relatively small.

Table 5.5: Distance matrix with respect to the  $l_1$ -distance for DAG-GNN. ♣ denotes the algorithm’s default hyperparameters.

		<i>Static</i>				<i>Temporal</i>			
		$\tau = 0$ $t = 0.2$	$\tau = 0$ $t = 0.3$ ♣	$\tau = 10^{-7}$ $t = 0.2$	$\tau = 10^{-7}$ $t = 0.3$	$\tau = 0$ $t = 0.2$	$\tau = 0$ $t = 0.3$ ♣	$\tau = 10^{-7}$ $t = 0.2$	$\tau = 10^{-7}$ $t = 0.3$
<i>Static</i>	$\tau = 0, t = 0.2$	0.0	9.0	6.0	10.0	14.0	17.0	14.0	17.0
	$\tau = 0, t = 0.3$ ♣	9.0	0.0	7.0	1.0	11.0	10.0	11.0	10.0
	$\tau = 10^{-7}, t = 0.2$	6.0	7.0	0.0	8.0	12.0	15.0	12.0	15.0
	$\tau = 10^{-7}, t = 0.3$	10.0	1.0	8.0	0.0	12.0	11.0	12.0	11.0
<i>Temporal</i>	$\tau = 0, t = 0.2$	14.0	11.0	12.0	12.0	0.0	5.0	2.0	5.0
	$\tau = 0, t = 0.3$ ♣	17.0	10.0	15.0	11.0	5.0	0.0	7.0	0.0
	$\tau = 10^{-7}, t = 0.2$	14.0	11.0	12.0	12.0	2.0	7.0	0.0	7.0
	$\tau = 10^{-7}, t = 0.3$	17.0	10.0	15.0	11.0	5.0	0.0	7.0	0.0

## 6 Conclusions and Discussion

The Arctic has undergone dramatic changes in the past few decades, and sea ice decline is believed to be a key driver for the Arctic warming. On the one hand, the sea ice is melted by mixed effects of atmospheric dynamical and thermodynamical processes. These processes, on the other hand, can be largely affected by sea ice melt. Therefore, this study investigate the causality between multiple atmospheric processes and sea ice variations using three data-driven causality discovery approaches (TCDF, NOTEARS and DAG-GNN). As shown in previous sections, one advantage of utilizing these approaches is they not only generate causal graphs, but also provide quantified information on causal strength weight or time lag. Another advantage is that these approaches can take all relevant variable into consideration to find potential important causal relationship, which is different from most related studies that only analyze pair-wise causality between two variables. Instead of performing computationally expensive climate model simulations, here we focus solely on an observational-type analysis. Specifically, we examine the sensitivity of causality graphs produced by three methods to different hyperparameters and then compare those graphs with domain knowledge graph.

We found that the outputs of the three algorithms are rather sensitive to the choice of hyperparameters. For example, choosing an only slightly too large regularization parameter can result in NOTEARS or DAG-GNN producing empty graphs, that is not discovering any causal relationships at all. Also the values of the other parameters turned out to be important and outputs for different choices of the hyperparameters can be quite different. Hence, some care must be taken when applying data-driven causality discovery approaches and domain knowledge is indispensable for assessing whether their produced outputs are reasonable.

Compared to domain knowledge graph, the static graphs produced by NOTEARS and DAG-GNN are relatively reasonable. The results from NOTEARS suggest that RH and precipitation dominate sea ice changes among all variables. In the meantime, the sea ice has a large impact on SW and CW. The graph generated by DAG-GNN, however, indicate that the zonal (u10m) and meridional (v10m) wind fields are more important for driving sea ice variations than other variables. And there are no atmospheric variables being affected by the sea ice. Note that the edges between u10m and v10m, SW and v10m, CW and v10m are not realistic, even though they are produced by both NOTEARS and DAG-GNN. As for the temporal graphs, very few edges can be found in TCDF and NOTEARS. In comparison, the DAG-GNN is able to produce more complicated and meaningful results. The sea ice is found to have a delayed impact on itself, but with no causal relationship with any atmospheric processes.

Based on our analysis, it is still very challenging to directly apply these state-of-the-art data-driven causality discovery approaches to this specific climate topic. However, there are several limitations with current study, which potentially has a large influence on our results. (1) There are large uncertainties in the domain knowledge graph and thus cannot be considered as ground truth. The climate scientists strives to investigate the complex feedbacks between atmosphere and sea ice, but our knowledge in this field is still very limited. (2) We average the atmospheric and sea ice variables within the pan-Arctic domain (north of 60°N) and our analysis is only based on the time series. However, the causal relationships between atmosphere and sea ice could be regionally dependent. (3) We use the full monthly atmosphere and sea ice records during 1980-2018. The feedbacks between atmosphere and sea ice are highly variable with season, and may occur at shorter time scales (e.g. daily). (4) The weights among different atmospheric variables are much higher than those between atmosphere and sea ice. Thus, the edges in the latter category could have been filtered out.

Nevertheless, this is a pioneer study in the application of data-drive causality discovery ap-



proaches in the interactions between atmosphere and sea ice. This study will pave the way for us to disentangle the complicated causal relationships in the earth system, by taking the advantage of cutting-edge data science and Artificial Intelligence technologies. It also provides a good opportunity for climate scientists, data scientists and computer scientists to work together to solve the puzzle in the nature, which will eventually advance our understanding of polar climate system and global climate change.

## Acknowledgments

This work is supported by the grant “CyberTraining: DSE: Cross-Training of Researchers in Computing, Applied Mathematics and Atmospheric Sciences using Advanced Cyberinfrastructure Resources” from the National Science Foundation (grant no. OAC–1730250). The hardware used in the computational studies is part of the UMBC High Performance Computing Facility (HPCF). The facility is supported by the U.S. National Science Foundation through the MRI program (grant nos. CNS–0821258, CNS–1228778, and OAC–1726023) and the SCREMS program (grant no. DMS–0821311), with additional substantial support from the University of Maryland, Baltimore County (UMBC). See [hpcf.umbc.edu](http://hpcf.umbc.edu) for more information on HPCF and the projects using its resources.

## References

- [1] B. Aragam, A. Amini, and Q. Zhou. Learning directed acyclic graphs with penalized neighbourhood regression. arXiv:1511.08963, 2017.
- [2] R. Bintanja and F. Selten. Future increases in arctic precipitation linked to local evaporation and sea-ice retreat. *Nature*, 509(7501), 2014.
- [3] L. Boisvert, M. Webster, A. Petty, T. Markus, D. Bromwich, and R. Cullather. Intercomparison of precipitation estimates over the arctic ocean and its peripheral seas from reanalyses. *J. Climate*, 31(20), 2018.
- [4] L. N. Boisvert and J. C. Stroeve. The arctic is becoming warmer and wetter as revealed by the atmospheric infrared sounder. *Geophys. Res. Lett.*, 42(11), 2015.
- [5] L. N. Boisvert, D. Wu, and C. Shie. Increasing evaporation amounts seen in the arctic between 2003 and 2013 from airs data. *J. Geophys. Res. Atmos.*, 120(14), 2015.
- [6] C. C. C. S. (C3S). Era5: Fifth generation of ecmwf atmospheric reanalyses of the global climate. Technical report, Copernicus Climate Change Service Climate Data Store (CDS), 2017.
- [7] D. Cavalieri, J. Crawford, M. Drinkwater, W. Emery, D. Eppler, L. Farmer, M. Goodberlet, R. Jentz, A. Milman, C. Morris, and R. Onstott. Nasa sea ice validation program for the dmsp ssm/i: final report, nasa technical memorandum 104559. Technical report, National Aeronautics and Space Administration, 1992.
- [8] D. Cavalieri, C. Parkinson, P. Gloersen, and H. Zwally. arxiv. Technical report, Istituto Dalle Molle di Studi sull’Intelligenza Artificiale.
- [9] R. Chemke, L. Polvani, and C. Deser. The effect of arctic sea ice loss on the hadley circulation. *Geophys. Res. Lett.*, 46(2), 2019.

- [10] X. Chen, Y. Liu, H. Liu, and J. G. Carbonell. Learning spatial-temporal varying graphs with applications to climate data analysis. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- [11] D. Chickering. Learning bayesian networks is NP-complete. In *Learning from Data*, volume 112 of *Lecture Notes in Statistics*. Springer, 1996.
- [12] Y. Choi, B. Kim, S. Hur, S. Kim, J. Kim, and C. Ho. Connecting early summer cloud-controlled sunlight and late summer sea ice in the arctic. *J. Geophys. Res. Atmos.*, 119(19), 2014.
- [13] T. Chu, D. Danks, and C. Glymour. Data driven methods for nonlinear granger causality: Climate teleconnection mechanisms. Technical Report CMU-PHIL-171, Department of Philosophy, Carnegie Mellon University, 2005.
- [14] J. Cohen, J. Screen, J. Furtado, and Coauthors. Recent arctic amplification and extreme mid-latitude weather. *Nature Geosci.*, 7, 2014.
- [15] C. J. Cox, T. Uttal, C. N. Long, M. D. Shupe, R. S. Stone, and S. Starkweather. The role of springtime arctic clouds in determining autumn sea ice extent. *Journal of Climate*, 29(18):6581–6596, 2016.
- [16] R. Cruz-García, V. Guemas, M. Chevallier, and F. Massonnet. An assessment of regional sea ice predictability in the arctic ocean. *Climate dynamics*, 53(1-2):427–440, 2019.
- [17] D. Dee and Coauthors. The era-interim reanalysis: configuration and performance of the data assimilation system. *Q. J. Roy. Meteor. Soc.*, 137:553–597, 2011.
- [18] Q. Ding, A. Schweiger, M. L’Heureux, D. S. Battisti, S. Po-Chedley, N. C. Johnson, E. Blanchard-Wrigglesworth, K. Harnos, Q. Zhang, R. Eastman, et al. Influence of high-latitude atmospheric circulation changes on summertime arctic sea ice. *Nature Climate Change*, 7(4):289–295, 2017.
- [19] C. Donnat and S. Holmes. Tracking network dynamics: A survey using graph distances. *Annals of Applied Statistics*, 12(2):971–1012, 2018.
- [20] I. Ebert-Uphoff and Y. Deng. Causal discovery for climate research using graphical models. *J. Climate*, 25(17):5648–5665, 2012.
- [21] V. Guemas, E. Blanchard-Wrigglesworth, M. Chevallier, J. J. Day, M. Déqué, F. J. Doblas-Reyes, N. S. Fučkar, A. Germe, E. Hawkins, S. Keeley, et al. A review on arctic sea-ice predictability and prediction on seasonal to decadal time-scales. *Quarterly Journal of the Royal Meteorological Society*, 142(695):546–561, 2016.
- [22] R. Guo, L. Cheng, J. Li, P. R. Hahn, and H. Liu. A survey of learning causality with data: Problems and methods. *ACM Computing Surveys (CSUR)*.
- [23] J. D. Hamilton. *Time series analysis*. Princeton University Press, 1994.
- [24] D. Heckerman, D. Geiger, and D. Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20:197–243, 1995.
- [25] M. M. Holland and C. M. Bitz. Polar amplification of climate change in coupled models. *Clim. Dyn.*, 21(3-4), 2003.

- [26] M. M. Holland, L. Landrum, D. Bailey, and S. Vavrus. Changing seasonal predictability of arctic summer sea ice area in a warming climate. *Journal of Climate*, 32(16):4963–4979, 2019.
- [27] Y. Huang, X. Dong, D. A. Bailey, M. M. Holland, B. Xi, A. K. DuVivier, J. E. Kay, L. L. Landrum, and Y. Deng. Thicker clouds and accelerated arctic sea ice decline: The atmosphere-sea ice interactions in spring. *Geophysical Research Letters*, 46(12):6980–6989, 2019.
- [28] Y. Huang, X. Dong, B. Xi, and Y. Deng. A survey of the atmospheric physical processes key to the onset of arctic sea ice melt in spring. *Climate Dynamics*, 52(7-8):4907–4922, 2019.
- [29] Y. Huang, X. Dong, B. Xi, E. K. Dolinar, and R. E. Stanfield. The footprints of 16 year trends of arctic springtime cloud and radiation properties on september sea ice retreat. *Journal of Geophysical Research: Atmospheres*, 122(4):2179–2193, 2017.
- [30] S. Hussung, S. Mahmud, A. Sampath, M. Wu, P. Guo, and J. Wang. Evaluation of Data-driven Causality Discovery Approaches among Dominant Climate Modes. *Technical Report HPCF-2019-12, UMBC High Performance Computing Facility, University of Maryland, Baltimore County*, 2019.
- [31] M. Kapsch, R. Graversen, and M. Tjernström. Springtime atmospheric energy transport and the control of arctic summer sea-ice extent. *Nat. Clim. Change*, 3, 2013.
- [32] M.-L. Kapsch, R. G. Graversen, M. Tjernström, and R. Bintanja. The effect of downwelling longwave and shortwave radiation on arctic summer sea ice. *Journal of Climate*, 29(3):1143–1159, 2016.
- [33] M.-L. Kapsch, N. Skific, R. G. Graversen, M. Tjernström, and J. A. Francis. Summers with low arctic sea ice linked to persistence of spring atmospheric circulation patterns. *Climate dynamics*, 52(3-4):2497–2512, 2019.
- [34] J. E. Kay and A. Gettelman. Cloud influence on and response to seasonal arctic sea ice loss. *J. Geophys. Res. Atmos.*, 114(D18), 2009.
- [35] C. F. Kennel and E. Yulaeva. Influence of arctic sea-ice variability on pacific trade winds. *Proc. Natl. Acad. Sci.*, 117(6):2824–2834, 2020.
- [36] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In Y. Bengio and Y. LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [37] B. Kopec, X. Feng, F. Michel, and E. Posmentier. Influence of sea ice on arctic precipitation. *Proc. Natl. Acad. Sci.*, 113(1), 2016.
- [38] M. Kretschmer, D. Coumou, J. F. Donges, and J. Runge. Using causal effect networks to analyze different arctic drivers of midlatitude winter circulation. *J. Climate*, 29(11):4069–4081, 2016.
- [39] R. Lindsay, M. Wensnahan, A. Schweiger, and J. Zhang. Evaluation of seven different atmospheric reanalysis products in the arctic. *J. Climate*, 27(7):2588–2606, 2014.
- [40] P.-L. Loh and P. Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *Journal of Machine Learning Research*, 15:3065–3105, 2014.

- [41] N. J. Matthewman and G. Magnusdottir. Observed interaction between pacific sea ice and the western pacific pattern on intraseasonal time scales. *J. Climate*, 24(19):5031–5042, 2011.
- [42] M. C. McGraw and E. A. Barnes. Memory matters: A case for granger causality in climate variability studies. *J. Climate*, 31(8):3289–3300, 2018.
- [43] A. Morrison, J. Kay, H. Chepfer, R. Guzman, and V. Yettella. Isolating the liquid cloud response to recent arctic sea ice variability using spaceborne lidar observations. *J. Geophys. Res. Atmos.*, 123(1), 2018.
- [44] M. Nauta, D. Bucur, and C. Seifert. Causal discovery with attention-based convolutional neural networks. *Machine Learning and Knowledge Extraction*, 1(1):312–340, 2019.
- [45] M. Nauta, D. Bucur, and C. Seifert. Causal discovery with attention-based convolutional neural networks. *machine learning & knowledge extraction*, 19(1), 2019. Code available on <https://github.com/M-Nauta/TCDF>.
- [46] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 2006.
- [47] E. A. Nussbaumer and R. T. Pinker. The role of shortwave radiation in the 2007 arctic sea ice anomaly. *Geophysical research letters*, 39(15), 2012.
- [48] J. Overland and M. Wang. Large-scale atmospheric circulation changes are associated with the recent loss of arctic sea ice. *Tellus A.*, 62(1), 2010.
- [49] J. T. Overpeck, G. A. Meehl, S. Bony, and D. R. Easterling. Climate data challenges in the 21st century. *science*, 331(6018):700–702, 2011.
- [50] C. Parkinson and J. Comiso. On the 2012 record low arctic sea ice cover: Combined impact of preconditioning and an august storm. *Geophys. Res. Lett.*, 40(7), 2013.
- [51] D. Perovich, T. Grenfell, B. Light, and P. Hobbs. Seasonal evolution of the albedo of multiyear arctic sea ice. *Journal of Geophysical Research: Oceans*, 107(C10):SHE–20, 2002.
- [52] J. Ramsey, M. Glymour, R. Sanchez-Romero, and C. Glymour. 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. *International Journal of Data Science and Analytics*, 3:121–129, 2017.
- [53] A. Rinke, E. Knudsen, D. Mewes, W. Dorn, D. Handorf, K. Dethloff, and J. Moore. Arctic summer sea ice melt and related atmospheric conditions in coupled regional climate model simulations and observations. *J. Geophys. Res. Atmos.*, 124(12), 2019.
- [54] S. Samarasinghe, M. McGraw, E. Barnes, and I. Ebert-Uphoff. A study of links between the arctic and the midlatitude jet stream using granger and pearl causality. *Environmetrics*, 30(4):e2540, 2019.
- [55] J. Schmidhuber. Deep learning in neural networks: An overview. Technical report, 2015.
- [56] T. Semmler, R. McGrath, and S. Wang. The impact of arctic sea ice on the arctic energy budget and on the climate of the northern mid-latitudes. *Clim. Dyn.*, 39(11), 2012.
- [57] M. C. Serreze and R. G. Barry. Processes and impacts of arctic amplification: A research synthesis. *Global and planetary change*, 77(1-2):85–96, 2011.

- [58] M. C. Serreze and J. C. Stroeve. Arctic sea ice trends, variability and implications for seasonal ice forecasting. *Phil. Trans. R. Soc. A.*, 373(2045), 2015.
- [59] H. Song, J. Tian, J. Huang, P. Guo, Z. Zhang, and J. Wang. Hybrid causality analysis of enso’s global impacts on climate variables based on data-driven analytics and climate model simulation. *Frontiers in Earth Science*, 7:233, 2019.
- [60] H. Song, J. Wang, J. Tian, J. Huang, and Z. Zhang. Spatio-temporal climate data causality analytics-an analysis of enso’s global impacts. In *Proceedings of the 8th International Workshop on Climate Informatics (CI2018)*, pages 45–48, 2018.
- [61] G. Spreen, R. Kwok, and D. Menemenlis. Trends in arctic sea ice drift and role of wind forcing: 1992–2009. *Geophys. Res. Lett.*, 38(19), 2011.
- [62] C. Strong, G. Magnusdottir, and H. Stern. Observed feedback between winter sea ice and the north atlantic oscillation. *J. Climate*, 22(22):6021–6032, 2009.
- [63] M. Sturm, J. Holmgren, and D. K. Perovich. Winter snow cover on the sea ice of the arctic ocean at the surface heat budget of the arctic ocean (sheba): Temporal evolution and spatial variability. *Journal of Geophysical Research: Oceans*, 107(C10):SHE–23, 2002.
- [64] L. Sun, J. Perlwitz, and M. Hoerling. What caused the recent “warm arctic, cold continents” trend pattern in winter temperatures? *Geophys. Res. Lett.*, 43(10), 2016.
- [65] C. W. Team, R. Pachauri, and L. M. (eds.). Climate change 2014: Synthesis report. contribution of working groups i, ii and iii to the fifth assessment report of the intergovernmental panel on climate change. Technical report, IPCC, Geneva, Switzerland, 2014.
- [66] S. Van de Geer and P. Bühlmann.  $l_0$ -penalized maximum likelihood for sparse directed acyclic graphs. *Annals of Statistics*, 41(2):536–567, 2013.
- [67] C. Wang, R. Graham, K. Wang, S. Gerland, and M. Granskog. Comparison of era5 and era-interim near-surface air temperature, snowfall and precipitation over arctic sea ice: effects on sea ice thermodynamics and evolution. *The Cryosphere*, 13(6), 2019.
- [68] J. Wang, J. Zhang, E. Watanabe, M. Ikeda, K. Mizobata, J. E. Walsh, X. Bai, and B. Wu. Is the dipole anomaly a major driver to record lows in arctic summer sea ice extent? *Geophysical Research Letters*, 36(5), 2009.
- [69] E. Watanabe, J. Wang, A. Sumi, and H. Hasumi. Arctic dipole anomaly and its contribution to sea ice export from the arctic ocean in the 20th century. *Geophysical research letters*, 33(23), 2006.
- [70] B. Wu, J. Overland, and R. D’Arrigo. Anomalous arctic surface wind patterns and their impacts on september sea ice minima and trend. *Tellus A: Dynamic Meteorology and Oceanography*, 64(1), 2012.
- [71] Y. Yu, J. Chen, T. Gao, and M. Yu. Dag-gnn: Dag structure learning with graph neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, 2019. Code available on <https://github.com/fishmoon1234/DAG-GNN>.
- [72] T. Zerenner, P. Friederichs, K. Lehnertz, and A. Hense. A gaussian graphical model approach to climate networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 24(2):023103, 2014.

- [73] K. Zhang, B. Schölkopf, P. Spirtes, and C. Glymour. Learning causality and causality-related learning: some recent progress. *National science review*, 5(1):26–29, 2018.
- [74] X. Zheng, B. Aragam, P. Ravikumar, and E. Xing. Dags with NO TEARS: Continuous optimization for structure learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2018. Code available on <https://github.com/xunzheng/notears>.