



Using Bayesian Networks to Investigate the Role of Arctic Variability in Midlatitude Circulation

A thesis submitted for the degree of *Doctor of Philosophy*

Nathanael Harwood

Department of Life Sciences

Institute of Environment, Health and Societies

Brunel University London

October 2019

Abstract

Recent enhanced warming and sea ice depletion in the Arctic have been put forward as potential drivers of severe weather in the midlatitudes. Evidence of a link between Arctic warming and midlatitude atmospheric circulation is growing, but the role of Arctic processes relative to other drivers remains unknown. Arctic-midlatitude connections in the North Atlantic region are particularly complex but important due to the frequent occurrence of severe winters in recent decades. Here, Dynamic Bayesian Networks with hidden variables are introduced to the field to assess their suitability for teleconnection analyses. Climate networks are constructed to analyse North Atlantic circulation variability at 5-day to monthly timescales during the winter months of the years 1981-2018. The inclusion of a number of Arctic, midlatitude and tropical variables allows for an investigation into the relative role of Arctic Amplification as a driver compared to internal atmospheric variability and other remote drivers.

A robust covariability between regions of amplified Arctic warming and two definitions of midlatitude circulation is found to occur entirely within winter at submonthly timescales. Hidden variables incorporated in networks capture periodic shifts between average and anomalously slow stratospheric polar vortex flow. An increase in predictive skill is achieved with the inclusion of hidden variables, but a number of caveats to their usage are demonstrated. The influence of the Barents-Kara Seas region on the North Atlantic Oscillation is found to be the strongest link at 5- and 10-day averages, whilst the stratospheric polar vortex strongly influences jet variability on monthly timescales.

Contents

Abstract	1
Contents.....	2
a) Key Acronyms.....	5
b) Data Dictionary	6
Chapter 1: An Introduction to Arctic-Midlatitude Weather Linkages.....	7
1.1 Research Aims	7
1.2 The Process of Arctic Amplification	8
1.3 Arctic-midlatitude Linkage Mechanisms	9
1.4 Regional Arctic-midlatitude Linkages.....	11
1.4.1 Eurasia.....	11
1.4.2 North America.....	12
1.4.3 The Euro-Atlantic Region	12
1.4.4 Tropical Influences	14
1.5 Midlatitude Circulation Proxies	15
1.6 Motivation	16
Chapter 2: Bayesian Networks	18
2.1 Bayesian Networks	18
2.1.1 Inference	19
2.1.2 Structure Learning	20
2.1.3 Parameter Learning	21
2.2 Dynamic Bayesian Networks	22
2.2.1 Hidden Markov Models.....	22
2.2.2 Dynamic Bayesian Networks.....	23
2.2.3 Hidden Variables	24
2.3 Summary.....	25
2.3.1 Graphical Models in Climate Science	26
2.3.2 Terminology Differences between Climate and Computer Sciences	27
2.3.3 Chapter Outline.....	28

Chapter 3: Learning Climate Network Structures with Static Bayesian Networks	29
3.1 Introduction	29
3.2 Data	30
3.2.1 (1) 6-month 'draft' Bayesian Network	30
3.2.2 (2) 35-year Bayesian Networks	32
3.2.3 Preparation of Spatial Data	33
3.3 Methods.....	34
3.3.1 Training and Testing Datasets.....	35
3.3.2 BIC Score	35
3.4 Results and Discussion	35
3.4.1 (1) 6-Month BN and (2a) 35-Year BNs: Daily.....	35
3.4.2 (2b) 35-Year BNs: Monthly.....	39
3.5 Conclusion.....	46
Chapter 4: Exploring Dynamic Bayesian Networks with Hidden Variables to Detect Arctic-Midlatitude Linkages	50
4.1 Introduction	50
4.2 Data	52
4.2.1 Data Development	55
4.3 Methods.....	58
4.3.1 Experiments	60
4.3.2 Model Development	61
4.4 Results	62
4.4.1 Model Comparison.....	64
4.5 Discussion and Conclusion	66
Chapter 5: Using Dynamic Bayesian Networks to Investigate the Relative Impact of Arctic Amplification on Midlatitude Circulation in the Euro-Atlantic Basin.....	68
5.1 Introduction	68
5.2 Data	70
5.3 Methods.....	73
5.3.1 Experiments	75

5.4 Results	76
5.4.1 Network Structures: the DAG Results	77
5.4.2 Network Performance: Predictive Accuracy	80
5.5 Discussion	83
5.5.1 Tropical Influence and Stratospheric Teleconnection Pathways	85
5.5.2 The Implications of Hidden State Switches	86
5.6 Literature-based Networks	88
5.6.1 Experiments	88
5.6.2 Tropical Variables: A Lead Time Analysis	89
5.7 Data Split Analysis	92
5.7.1 Experiments	93
5.7.2 Results: Before and During the AA Period	95
5.7.3 Model Performance in the AA period	97
5.8 Conclusion	97
Chapter 6: Conclusion	99
6.1 Research Contributions	99
6.1.1 Dynamic Bayesian Networks with Hidden Variables	99
6.1.2 The Relative Impact of AA on Midlatitude Circulation.....	100
6.2 Study Limitations	101
6.3 Future Work	102
6.3.1 Recommendations	102
6.3.2 A Way Forward	104
Appendix	106
A1 Chapter 3 Figures	106
A2 Chapter 4 Figures	111
A3 Chapter 5 Figures	112
Acknowledgements	126
References	128

a) Key Acronyms

BN Bayesian Network

DBN Dynamic Bayesian Network

HMM Hidden Markov Model

pDAG Partially Directed Acyclic Graph

DAG Directed Acyclic Graph

HV Hidden Variable

AO Arctic Oscillation

NAO North Atlantic Oscillation

ENSO El Niño Southern Oscillation

MJO Madden Julian Oscillation

AA Arctic Amplification

ECMWF European Centre for Medium-Range Weather Forecasts

KNMI *Koninklijk Nederlands Meteorologisch Instituut* or Royal Netherlands Meteorological
Institute

b) Data Dictionary

Variable	Abbreviation	Source	Unit	No. Samples
Arctic Oscillation	AO	NOAA	Geopotential height (m)	12755
North Atlantic Oscillation	NAO	NOAA	Sea level pressure	13239
El Niño Southern Oscillation	ENSO	NOAA	Sea surface temperature	13239
Snow Cover Extent	SCE	NOAA	Fraction of snow-covered area	12755
Sea Ice Extent	SIC	ECMWF ERA-Interim	Fraction of sea ice area	12755
Zonal component of wind	uWind	ECMWF ERA-Interim	Metres per second (m/s)	182
Jet latitude	Jetlat	ECMWF ERA-Interim	Degrees (°)	13239
Jet Speed	Jetspeed	ECMWF ERA-Interim	Metres per second (m/s)	12755
850hPa temperature	850hPa	ECMWF ERA-Interim	Temperature (°C)	13239
Stratospheric polar vortex	PoV	ECMWF ERA-Interim	Geopotential height (m)	13239
Meandering Index	MI	ECMWF ERA-Interim	-	13239
Madden-Julian Oscillation	MJO	NOAA	MJO amplitude	13239

Data Dictionary of all datasets used in this doctoral project, listing the variable name, its abbreviation, source and unit of measurement, as well as the upper limit for the number of samples – i.e. the original daily data acquired from sources before any processing listed in the data sections of following chapters.

Chapter 1: An Introduction to Arctic-

Midlatitude Weather Linkages

The literature and theories surrounding the Arctic-midlatitude research area are dynamic as the field has only been extensively published on since around 2012. Nonetheless, this chapter is designed to provide information on our current understanding of the physical processes underlying Arctic-midlatitude linkage mechanisms, and how they may differ geographically across the Northern Hemisphere. Important climate phenomena that occur in the Arctic, midlatitudes and the tropics are presented to underpin the analyses conducted in later chapters. Finally, a 'motivation' for this research project is given to close the chapter.

1.1 Research Aims

This project applies Dynamic Bayesian Networks (DBNs) with hidden variables (HV) to the North Atlantic and European midlatitude circulation research area. Structure-learning algorithms are employed to identify regions of AA that might influence winter jet stream variability. The inclusion of a number of Arctic, midlatitude and tropical variables allows for an investigation into the relative role of AA as a driver compared to internal atmospheric variability and other remote drivers. Other graphical model approaches have fallen short in this regard as they focus either entirely on potential Arctic drivers of midlatitude circulation responses (Kretschmer et al, 2016; Barnes and Simpson, 2017; Samarasinghe et al, 2019) or on possible tropical teleconnections like the MJO-NAO link (Barnes et al, 2019). The main aim of this thesis is to establish how effective DBNs with structure learning algorithms are for investigating this research area, and to measure the impact of hidden variables on model accuracy which is a priority due to the low signal-to-noise ratio of AA linkages and their intermittent nature (Overland et al, 2016).

Finally, the findings presented here are motivated in part by a desire to contribute to the knowledge base of midlatitude weather prediction. Arctic variables like sea ice concentration provide predictive skill for the winter NAO in both atmospheric (Scaife et al, 2014) and statistical models (Hall et al, 2017; 2019; Wang et al, 2017), but a great deal of uncertainty remains surrounding the relative importance of AA as a driver. If AA has played a significant role or increased the frequency of Arctic and midlatitude interactions, then achieving a synthesis on this important area of science and maximising subseasonal-to-seasonal midlatitude weather prediction skill has the potential to impact millions of people living in the midlatitudes.

1.2 The Process of Arctic Amplification

Arctic Amplification (AA) describes the rate with which Arctic temperatures are increasing regionally compared to the rest of the globe (Fig. 1.1). The Arctic has outpaced the global rise in temperatures at more than twice the rate since the mid 20th century, and by over six times the pace between the years 1998-2012 (Huang et al, 2017). The latest IPCC report highlighted a decline in sea ice extent of 3.5-4.1% per decade between 1979 and 2012, and a mean annual snow cover extent decrease of up to -53% in June for the years 1967-2012 (Vaughan et al, 2013). This equates to the loss of approximately half a million square kilometres per decade since the beginning of the satellite era. The greatly increased fraction of open water that results from this lowers the albedo of the Arctic Ocean's surface, i.e. the

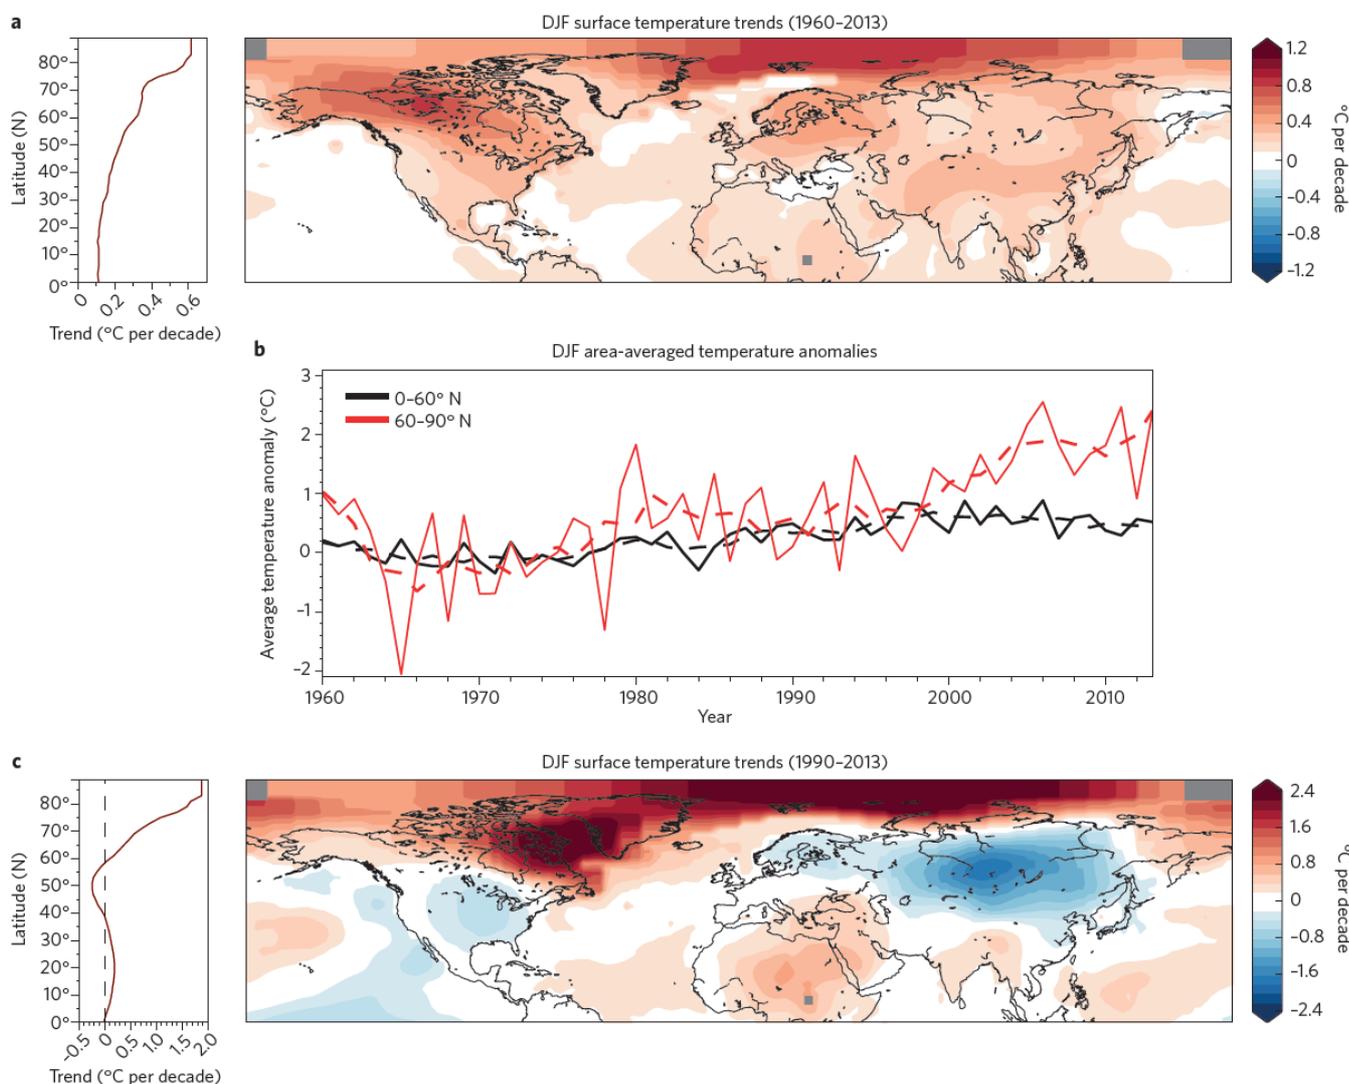


Figure 1.1 Arctic Amplification – trends in winter temperatures. A) Linear trend in SAT from 1960-61 to 2013-14. B) Area-averaged surface temperature anomalies. C) As in a) but 1990-91 to 2013-14. Note differences in temperature by latitude in graphs on left hand side. Source: Cohen et al (2014).

difference between the reflectance of dark water as opposed to snow-covered sea ice. The sea ice albedo feedback (Perovich et al, 2008) describes the excess heat absorbed by the darker ocean surface during the summer melt season which is released during autumn through ocean-atmosphere heat fluxes. The ocean-atmosphere temperature gradient increases in autumn, resulting in a warming of the lower troposphere. This in turn delays the formation of sea ice due to warmer and moister masses of air present late on in the year (Vaughan et al, 2013).

Whilst sea ice plays an important role in AA, other physical mechanisms have contributed to the AA process in significant ways. AA is the result of a complex combination of local sensible heat fluxes, evaporation and the remote transport of heat and moisture from lower latitudes (Cohen et al, 2018a). Remote transport describes the poleward advection of heat and moisture into the Arctic from lower latitudes which has recently been identified as an important driver of sea ice loss (Woods and Caballero, 2016; Kapsch et al, 2016; Gong et al, 2017). Midlatitude atmospheric circulation plays an important role in this process, as intrusions of moist air from lower latitudes have been found to increase the exchange of heat between the atmosphere and ocean; the effects on sea ice are obvious within several days of the event (Kapsch et al, 2016). Moisture transport seems to be especially pronounced through the North Atlantic pathway, the study region for this project, where Atlantic blocking can drive enhanced poleward transport (Kim et al, 2017; Yang and Magnusdottir, 2017). Recent work with atmosphere-only models has suggested that midlatitude circulation has a strong influence on the low-middle altitude warming trends of the Arctic during winter (Ye and Jung, 2019), further emphasising the importance of the midlatitude-Arctic linkage as a driver of AA. Clearly, when considering amplified Arctic warming, sea ice variability does not adequately capture the whole picture as an isolated driver or proxy.

1.3 Arctic-midlatitude Linkage Mechanisms

The possible pathways through which AA processes can have an impact on midlatitude circulation and persistent weather are discussed here. Given that all analyses in this project investigate winter impacts of AA on midlatitude circulation, and that AA is most pronounced during the winter months (Serreze et al, 2009), this chapter focuses on winter studies and mechanisms.

One prominent, widely debated mechanism is a proposed link between AA and meridional (north-south oriented) jet stream patterns. AA results in a decreased meridional temperature gradient as the Arctic warms faster than lower latitudes, which has led some to hypothesise that a decreased poleward temperature gradient would lead to increases in the wave amplitude of the Polar jet stream (e.g. Francis et al, 2009). Upper level westerly winds could therefore be expected to become weaker due to the thermal wind relation which dictates

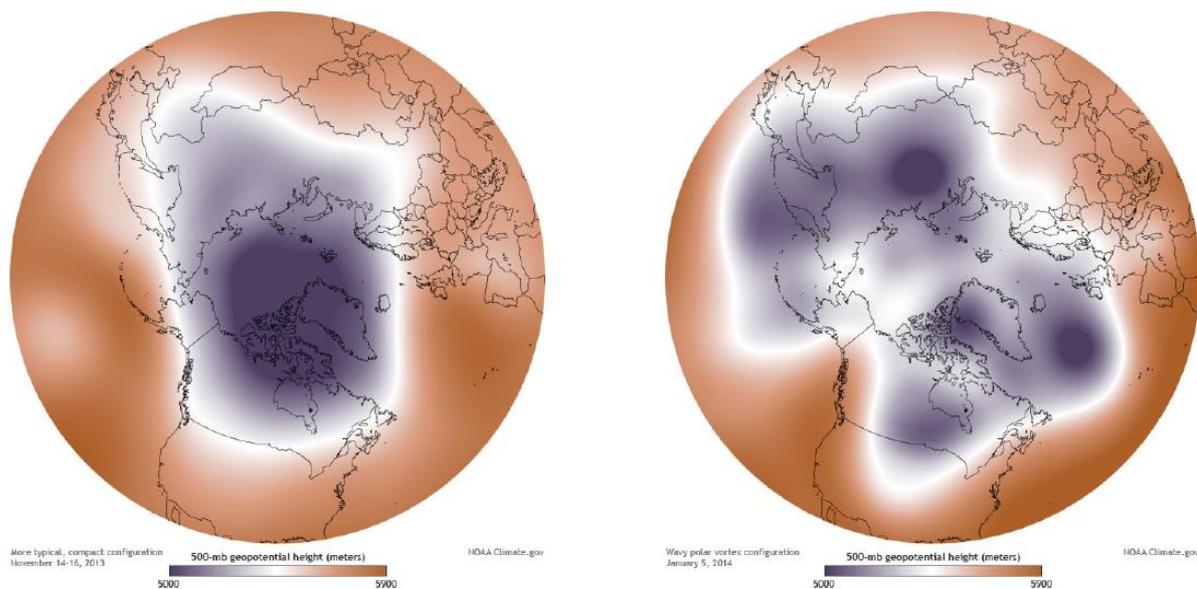


Figure 1.2 Sample 500hPa geopotential height fields showing a zonally oriented jet stream (left) and a wavier flow with multiple outbreaks of low pressure (right). The tropospheric polar vortex is contained within the Arctic on the left. White contours delineate the jet stream. Source: Cohen et al (2018a).

the change of the geostrophic wind with pressure (and height) (e.g. Overland and Wang, 2010; Francis and Vavrus, 2012). Finally, Rossby waves would propagate eastwards at a slower rate as the north-south amplitude of the flow increases in magnitude, causing weather systems to progress eastwards more slowly and persistent weather patterns to dominate (Francis and Vavrus, 2012). This mechanism has been called into question by similar studies that found wave amplitude and blocking frequency results depended strongly on the method used to calculate these metrics, and that the trends are still within the expected range of internal variability (Barnes, 2013; Screen and Simmonds, 2013).

Anomalous circulation patterns and increases in jet stream amplitude can drive the southerly displacement of cold Arctic air masses as well as heat waves and flooding. Months with weather extremes often have Rossby wave patterns with amplified magnitudes when compared to average conditions when zonal circulation patterns dominate (Screen and Simmonds, 2014; Coumou et al, 2014).

Another mechanism thought to impact midlatitude flow is a jet stream split which is thought to be favoured by a weaker meridional temperature gradient and AA, whereby double jet regimes weaken the propagation rates of high-amplitude Rossby waves through a process known as 'quasi-stationary wave amplification' (Petoukhov et al, 2013; Coumou et al, 2014). Given that the process is thought to occur entirely within summer, work published on double jet configurations and trapped atmospheric waves does not inform the aims and analyses of this project.

The stratosphere is also thought to play a role in possible linkages between the Arctic and midlatitudes. Anomalously low sea ice conditions, especially in the Barents and Kara Seas, has been put forward as a potential source region for vertical wave propagation which initiates stratospheric polar vortex weakening and stratospheric warming events (e.g. Kim et al, 2014; Kretschmer et al, 2016). The propagation of Rossby waves vertically into the stratosphere can weaken the stratospheric polar vortex through a process known as wave-breaking, which in turn promotes tropospheric circulation anomalies that resemble the negative phase of the Arctic Oscillation (AO) and a weaker jet stream (Cohen et al, 2014). The stratospheric pathway has thus been hypothesised to cause more severe winter conditions and cold events in the midlatitudes, associated with the negative phase of the AO. Recent papers have questioned whether the link between Barents-Kara sea ice and Eurasian cold extremes seen in observational studies (Kim et al, 2014; Kretschmer et al, 2016; Zuo et al, 2016) represents a causal physical mechanism as the finding is not replicated in model simulations (Kolstad and Screen, 2019; Fyfe, 2019; Warner et al, 2020). Instead, this observed statistical relationship may be caused simply by internal climate variability (Kolstad and Screen, 2019) or by remote tropical drivers acting as an uncaptured driver of both variables (Warner et al, 2020). The experimental design of chapters 4 and 5 includes tropical indices in an attempt to capture non-Arctic remote drivers of midlatitude circulation variability, and section 5.6 runs these network results at different tropical variable lead times.

1.4 Regional Arctic-midlatitude Linkages

1.4.1 Eurasia

Cold anomalies over Eurasia have occurred in recent years leading to severe winters despite enhanced warming over the Arctic. Observational analyses show a consistent link between cold Eurasian extremes and Barents-Kara sea ice variability (Kim et al, 2014; Kretschmer et al, 2016) as already noted, however numerical models have failed to capture a robust remote response to sea ice depletion across studies, leading some to propose internal atmospheric variability (McCusker et al, 2016; Sun et al, 2016) or Ural blocking anomalies (Peings, 2019) as the main drivers of the observed Eurasian cooling.

The stratospheric pathway detailed above may act as an important mechanism for possible linkages (Cohen et al, 2018a). Barents-Kara sea ice loss in early winter initiates the propagation of vertical Rossby waves which weaken the stratospheric polar vortex, resulting in the downward propagation of wave energy into the troposphere in mid-late winter (Kim et al, 2014). Kretschmer et al (2016) use a graphical model approach to support a stratospheric pathway, finding that the Siberian High strengthens and intensifies inducing cold air advection from the Arctic and anomalous cold Eurasian winters. Constructive interference with climatological wave 1 and 2 patterns is thought only to occur in the Barents-Kara region (Kim

et al, 2014; Screen, 2017). Climate models, meanwhile, lack a robust signal which might be a product of the current generation of models and their ability to represent stratospheric processes, a known source of uncertainty (Zhang et al, 2018). Similarly, the magnitude of the atmospheric response to sea ice loss may be underestimated in models; accounting for this led to a robust sea ice signal in atmospheric general circulation models (GCMs) in a recent study (Mori et al, 2019), although Screen and Blackport (2019) caution that the comparison of models and observations made in this paper are misleading and do not constitute evidence that ~44% of Eurasian cooling is attributable to sea ice loss as claimed by Mori et al (2019). As is the case with much of the Arctic-midlatitude research area, the disparity between the results of observational and model approaches is likely only to be resolved by coordinated, multi-model experiments (Cohen et al, 2018a).

1.4.2 North America

A multitude of studies exist on connections between the North American Arctic and midlatitudes, but only the impacts on the east coast of North America and Greenland are relevant to the aims of the project. Linkage mechanisms identified as potential connections are governed by their interactions with large-scale waves in the jet stream, and whether those interactions are constructive or destructive (Cohen et al, 2018a), making them highly intermittent in nature (Overland et al, 2016).

AA processes may contribute to the formation and duration of blocking patterns (Chen and Lou, 2017; Ballinger et al, 2018). Higher geopotential heights in the Greenland and Baffin Bay regions increase the likelihood of Greenland blocking events, leading to an increase in the meridional amplitude of the jet and in turn the persistence of cold events on the eastern US coast (Chen and Lou, 2017). Severe cold outbreaks on the eastern US coast may not be triggered by sea ice variability or surface warming in the Arctic, but such processes can increase regional geopotential heights which reinforces the location of high amplitude jet patterns (Overland and Wang, 2018). Given that more research is required on the potential for regions of AA to support Greenland blocking patterns and the disruption of midlatitude circulation, this project focuses on the North Atlantic midlatitudes as a potential 'impact' region.

1.4.3 The Euro-Atlantic Region

The Euro-Atlantic region, which denotes the North Atlantic and western Europe, is dominated by the phase of the North Atlantic Oscillation (NAO) and NAO variability, which has a number of potential drivers including cryospheric variables (sea ice and snow cover), tropical sea-surface temperatures (SSTs) in the Pacific and Atlantic oceans, stratospheric circulation variability and the Quasi-Biennial Oscillation (QBO) (Hall et al, 2015). The NAO describes the north-south orientation in the dipolar structure of the North Atlantic pressure field, characterised by a centre of low pressure over Iceland and high pressure over the Azores

(Hurrell, 1995). The NAO effectively reflects the variability of the storm tracks and the Polar jet stream (Overland et al, 2015), describing it through phase shifts whereby the positive phase is associated with milder European winters and a further poleward jet displacement, and the negative a severe winter with a southerly displacement of the jet and the predominance of blocking regimes (Hall et al, 2015).

A large amount of uncertainty exists regarding the response of the NAO to AA processes due to the complex set of drivers associated with NAO variability and nonlinear linkages between the two. A number of studies find that reductions in Arctic sea ice promote the negative phase of the NAO (NAO-) in both observations and models (Petoukhov and Semenov, 2010; Kim et al, 2014; Mori et al, 2014), dependent on the pattern and amplitude of sea ice anomalies. The Barents and Kara seas are frequently found to be an important region of sea ice loss for the NAO response (Kim et al, 2014; Mori et al, 2014), hence this project has a strong data focus on the Barents-Kara region alongside the North Atlantic. Screen (2017) points out that the increased frequency of NAO- events may be misleading in terms of a surface severe cold response, as the temperature of such events remains stable on the seasonal scale due to the conflicting effects of dynamical cooling (from NAO- events) and thermodynamical impacts (i.e. the warming trend). Added to this, the simulated response to sea ice loss is generally a decrease in severe cold events in winter (Ayarzaguena and Screen, 2016; Blackport et al, 2019).

Whilst the increased frequency of cold events due to sea ice loss is considered an unlikely response, AA processes may still contribute to atmospheric circulation variability. Atmospheric models suggest robust responses to sea ice reduction including the increased intensity of NAO- events and a weakening in the zonal wind on the poleward side of the jet stream during the winter months (Screen et al, 2018). Most recently, coupled model simulations revealed that a robust weakening and equatorward migration of the jet in response to Arctic sea ice loss occurs predominantly in the Atlantic basin (Blackport and Screen, 2019). Here, the winter atmospheric circulation response was driven by sea ice variability within the winter months, rather than as a lagged response to autumn sea ice as proposed in studies using observation-based approaches (Kim et al, 2014; Kretschmer et al, 2016; Hall et al, 2017). A weak but robust influence of regional AA on the jet stream's position and speed was found in a combined climate projection and reanalysis dataset using Granger causality, matching the equatorward shift of the Atlantic jet in response to near-surface Arctic warming in the North Atlantic region of the Arctic (Barnes and Simpson, 2017). These two factors – a potential atmospheric response to winter AA, and the use of near-surface temperature to represent AA rather than sea ice – inform the study design of later chapters.

1.4.4 Tropical Influences

Tropical modes of variability including the El Niño Southern Oscillation (ENSO) and the Madden-Julian Oscillation (MJO) have an influence on midlatitude circulation in the Euro-Atlantic region. Pacific SSTs can trigger large-scale Rossby waves through intense convection and precipitation which can propagate into the midlatitudes and impact midlatitude circulation (Trenberth et al, 1998; Scaife et al, 2017). Tropical variability is a source of predictive skill for both the wintertime NAO (Scaife et al, 2016) and the AO (Sun and Ahn, 2015), and tropical rainfall has been found to explain a high degree of NAO variability in correlative studies (Scaife et al, 2016; Hall et al, 2017). In a similar mechanism to the ENSO impact, Rossby wave trains which originate in the tropics can reach the midlatitudes, initiated by tropical convection associated with the MJO (Sardeshmukh and Hoskins, 1988; Frederiksen and Lin, 2013). The tropospheric pathway is considered to be the more impactful for MJO teleconnections as MJO-driven Rossby wave trains have been associated with Polar jet stream strength and position (Moore et al, 2010), the positive phase of the NAO in winter (Lin et al, 2015) and blocking occurrence over the Euro-Atlantic basin (Henderson et al, 2016). Whilst studies have demonstrated a response in North Atlantic circulation via the stratosphere for both indices (e.g. Baldwin and Dunkerton, 2001; Jiang et al, 2017), the dominant pathway for strong El Niño events to influence the NAO is also thought to be the troposphere, with La Niña events shown to influence stratospheric polar vortex strength and then subsequently NAO variability (Hardiman et al, 2019).

Stratospheric polar vortex (SPV; the proxy dataset for the SPV is also referred to as 'PoV' in later chapters) conditions are nonetheless an important element of the climate system. A weak SPV can cause anomalous easterlies to propagate downwards into the troposphere which induces a weakening in the zonal component of the Polar jet stream (Baldwin and Dunkerton, 2001; Kidston et al, 2015). Because the wintertime stratospheric polar vortex projects onto the AO pattern, shifting the sign to a negative AO, it is thought to be central to the development of the long-lead predictive capability of Northern Hemisphere midlatitude circulation (Robertson et al, 2015; Scaife et al, 2016). In addition, recent work has highlighted that MJO influence on NAO variability through the stratosphere seems to be contingent on the background state of the SPV, whereby a robust NAO response is found only when the vortex and NAO states are aligned (Barnes et al, 2019). The stratosphere clearly exerts a strong influence on tropospheric midlatitude circulation variability, meaning that potential SPV drivers including sea ice loss (Kim et al, 2014), ENSO (Baldwin and Dunkerton, 2001) and the MJO (Jiang et al, 2017; Barnes et al, 2019), as well as the SPV itself, represent important variables through which to investigate potential stratospheric linkage pathways.

1.5 Midlatitude Circulation Proxies

The metric used to calculate midlatitude circulation, whether it summarises large-scale circulation in an index like the NAO (Hurrell, 1995), approximates the speed and position of the Polar jet stream (Woollings et al, 2010), or measures the degree of waviness (Di Capua and Coumou, 2016), is central to study results and no scientific consensus currently exists on which particular metric is most appropriate. An example that highlights this point is the disparity in the results of Francis and Vavrus (2012), who reported an increase in wave amplitude over North America and the North Atlantic, and Barnes (2013) who found no significant increase with a modified wave definition.

Many studies focusing on impacts in the Euro-Atlantic region use the NAO index as a proxy for midlatitude circulation (e.g. Lin et al, 2015; Wang et al, 2017; Barnes et al, 2019). Whilst the NAO is useful as an indicator of jet variability and is generally not as noisy as other midlatitude circulation datasets, Screen (2017) shows that the increase in NAO- events is not matched by an increase in severe cold event frequency over Northern Europe in simulations. The NAO can therefore be misleading if the study focus is cold waves and the surface temperature response to AA in the midlatitudes.

The variability of the jet stream itself derived as a proxy from reanalysis and model output can be used to describe midlatitude circulation. Jet latitude and speed, calculated using the technique pioneered by Woollings et al (2010), use low-level winds to exclude the higher subtropical jet and can be calculated over regions like the North Pacific to account for differences in seasonal dynamics across the Northern Hemisphere. Many studies focusing on midlatitude circulation variability in response to remote drivers have made use of this proxy (e.g. Hall et al, 2017; Barnes and Simpson, 2017; Samarasinghe et al, 2019).

Secondly, metrics have been developed to describe the meridional component of atmospheric circulation based on a number of characteristics. These methods capture the waviness of circulation by differentiating the geopotential height contour at each timestep from zonal conditions in terms of wave amplitude (e.g. Francis and Vavrus, 2012) and sinuosity (Cattiaux et al, 2016; Di Capua and Coumou, 2016). An increase in waviness was found only for the North American sector in climate projections using these approaches, with no robust trends in the meridional amplitude of circulation for other Northern Hemisphere regions (Cattiaux et al, 2016; Di Capua and Coumou, 2016). Whilst they certainly represent useful circulation proxies, some criticism has been directed at these approaches for being derived geometrically rather than being based on physics (Cohen et al, 2018a).

1.6 Motivation

A number of obvious gaps in knowledge and scientific community consensus are identified which inform the data, methodological approach and conclusions of this project. Firstly, Arctic Amplification is a recent phenomenon, and the AA signal is only distinguishable from internal variability from approximately 1995 onwards at the surface and 2000 onwards in the lower troposphere (Francis and Vavrus, 2015). A low signal-to-noise ratio effectively means that linkages are hard to detect, especially with a limited time range of reliable satellite observations (~40 years). Secondly, midlatitude responses to Arctic forcing are likely to be highly intermittent (Overland et al, 2016). A pre-existing meridional jet stream configuration may be a pre-requisite of Arctic-midlatitude links over North America; as an example, amplified regional Arctic warming in December 2016 had no impact on the midlatitudes due to a highly zonal jet stream, in stark contrast to the December 2017 US Cold wave (Overland and Wang, 2018).

As has been shown above, a number of remote drivers may impact midlatitude circulation in the North Atlantic region including cryospheric variables, tropical sea-surface temperatures and convection, and stratospheric circulation variability (Hall et al, 2015; 2017; Smith et al, 2016). The extraction of a robust Arctic signal therefore becomes a complex problem due to the diverse range of drivers, which themselves may only link intermittently and act to cancel out or reinforce the impact of others. Sea ice concentration has been used as an indicator of AA in both observational (Kim et al, 2014; Kretschmer et al, 2016) and model-based approaches (Mori et al, 2014; 2019; McCusker et al, 2016; Screen, 2017), where sea ice concentrations can be manipulated to isolate a circulation response to sea ice only. This approach only captures sea ice variability however, missing a range of other important AA drivers including the remote transport of heat and moisture into the Arctic through amplified circulation patterns (Ye and Jung, 2019). There is a clear need for another metric which adequately captures the full AA signal; this is a focus of the analysis of chapters 4 and 5.

Alongside these data-based problems, all methodological approaches have to some degree a number of drawbacks which should be assessed at the point of use and addressed at the point of publication. Findings based entirely on correlation analysis are subject to autocorrelation bias and may be the result of indirect links or a common driver entirely unaccounted for in the analysis (Runge et al, 2014). Linear relationships are also directionless, so offer less information than graphical models. Atmospheric model approaches, meanwhile, constitute the bulk of studies able to find robust teleconnection mechanisms due to the large number of years simulations can be run for. Model studies, whilst well regarded as tools for identifying causal linkages, are not immune from potential shortcomings: they may not accurately represent ocean-atmosphere coupling in the Arctic (Cohen et al, 2018a), may

respond too weakly to sea ice forcing (Screen et al, 2018; Mori et al, 2019), may underperform in terms of stratosphere-troposphere coupling (Zhang et al, 2018), and focus on the impact of sea ice removal which may not capture the complex intermittencies thought to define Arctic-midlatitude linkages (Overland and Wang, 2018).

Chapter 2: Bayesian Networks

The fundamentals of the Bayesian Network method are presented here to facilitate understanding of the experimental design of all three data chapters. The structure-learning and parameter-learning processes and the algorithms that drive them are covered, with model evaluation and a number of science applications discussed alongside to demonstrate their usage. Finally, a definition of hidden variables inferred from input data is given as they are central to the model accuracy findings detailed throughout this work.

2.1 Bayesian Networks

Graphical models provide an excellent tool for examining relationships between variables of climatological importance, due to their properties as a visualisation tool for result communication. Bayesian Networks are a form of probabilistic graphical models that provide a useful mechanism for statistically modelling relationships that occur in the real world allowing for their visualisation in a 'graph', or network. BNs combine elements of graph theory with statistics to explore relationships between a set of independent variables (Friedman et al, 2000). BNs broadly function in two ways: a graph-based structure, either defined by the user or constructed from the data using an algorithm, represents relationships between variables which are then defined using conditional probability tables (CPTs) for discrete data and conditional probability distributions (CPDs) for continuous data.

A Bayesian Network (BN) encodes a joint probability distribution, whereby probabilities are assigned for all possible outcomes over a set of random variables taken as input (Friedman et al, 2000). BNs accomplish this by constructing directed acyclic graphs (DAGs) which exploit conditional independence relationships, i.e. that two variables in the graph are independent if knowing the state of one event does not change the probability of the second. A DAG comprises a graph with the variables as 'nodes', links between nodes as 'edges' and 'parent' and 'child' nodes defining their position within the graph.

BNs 'exploit' conditional independence relationships in the network parameterisation phase, as a variable which is conditionally independent from another in the graph structure does not need to be parameterised, resulting in a significant boost to computational efficiency. In Fig. 2.1a for example, x and z are conditionally independent given the third node y therefore no arc between x and z is needed or parameterised within the BN. The CPDs of the BN are then worked out with respect to the DAGs conditional independence relationships. Each variable has a CPD which encodes the probability of observing values given the values of its 'parents', i.e. the variables which it is conditionally dependent on. The BN factorises all of these CPDs in a joint distribution, with each variable represented as x_i :

$$p(x_1 \dots x_n) = \prod_{i=1}^n p(x_i | Pa(x_i)) \quad (1)$$

where n is the total number of variables and $Pa(x_i)$ denotes the parent set of x_i (the variable).

The DAG graphically represents the BN with nodes (the variables) and conditionally dependent relationships indicated with edges (links between the variables). When a node has two or more parents (a ‘collider’ structure, Fig. 2.1b), an increase in probability for parent A which results in a decrease in another B can be said to ‘explain away’ the likelihood of B being a driver. To give a hypothetical example from the climate science field, strongly anomalous values for the stratospheric polar vortex in a given timestep might ‘explain away’ the need for Barents-Kara sea ice concentration as a driver of jet stream variability if both were parent nodes of the jet. This process is referred to as being caused by a ‘v-structure’ in some work (e.g. Verma and Pearl, 1991), and is central to the method of chapters 4 and 5, where links without direction are removed to allow networks to function correctly.

2.1.1 Inference

BNs can also be used for inference, once all the CPDs of the model are defined. As previously mentioned, a node state can change depending on the states of other nodes in the graph and the propagation of probability through the DAG. This makes it possible to manually change the state of a node in the network to see how that variable changes the states of other nodes, or to replace the dataset with new data to examine the posterior probabilities (Koller and Friedman, 2009). In this project, the *junction tree* algorithm was used for inference, which

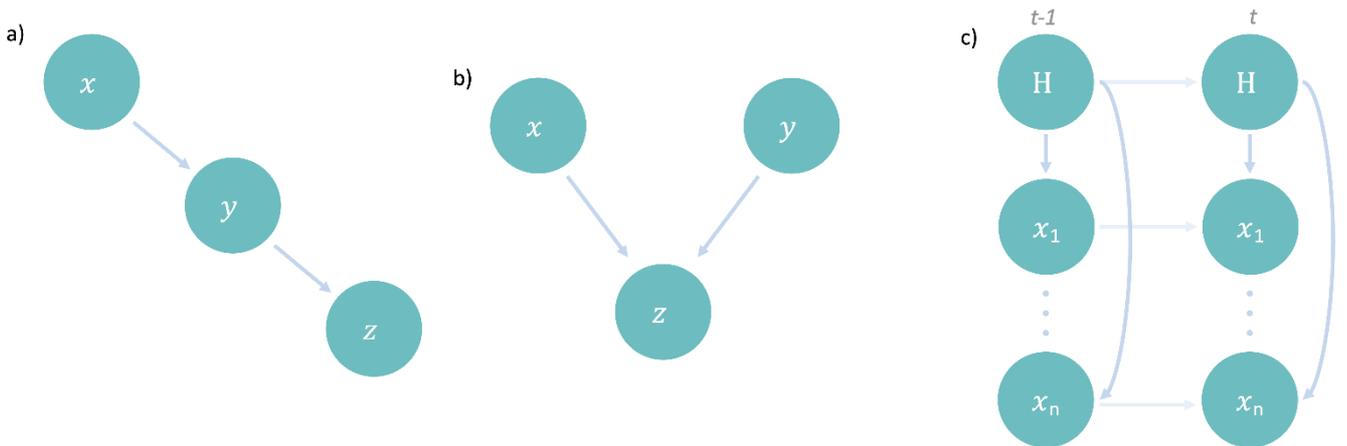


Figure 2.1 Example DAG structures: a) x and z are conditionally independent given y therefore no arc between x and z is needed or parameterised; b) an example of a collider structure, where node z has two parent nodes, x and y ; c) a DBN with a single hidden variable rolled out over two time slices with autoregressive links between all variables.

makes use of a greedy search procedure to find the optimal order for variable elimination (Murphy, 1998). The BN framework employed in chapters 4 and 5 makes use of the junction tree algorithm to predict the node values in testing datasets, i.e. a dataset independent of the input dataset, to examine the fit of the model. The results of this are referred to as the 'predictive accuracy' of each model, discussed in terms of their relative performance for each dataset.

The purpose of the node 'predictions' referred to in this project are therefore simply to test the accuracy of the BN, as section 2.3.2 explains below. Data at time t is inferred from the data at time $t-1$, based on the model structure and parameters that have been fitted. The n -step ahead forecasting method iterates between entering the observations for all variables at time $t-1$ and applying inference to calculate their posterior distributions at time t , then repeats this step n times. From the output of the node prediction process, the proportion of values that were predicted correctly can be ascertained for the whole test dataset, and compared against the actual test dataset values which were excluded from the structure- and parameter-learning phases of model construction. This is then visually and statistically assessed throughout this project; the predictions are plotted against observed values for the test dataset, and the sum of squared error (SSE) is calculated to compare model accuracy with ease.

2.1.2 Structure Learning

The graph structure of a BN can be user-defined through the use of past studies and expert knowledge, or learned from the data using an algorithm. Given that this project aims to test the Arctic-midlatitude linkage hypothesis, that relationships between Arctic and midlatitude processes can be discerned using large modes of variability, the latter approach is employed here. Structure-learning algorithms can be broadly split into three categories; constraint-based, score-based, and hybrid (not used in this work). All return DAGs or pDAGs (partially Directed Acyclic Graphs) which graphically display the conditional dependence relationships between input variables.

The structure-learning mechanism primarily used in this project is the PC algorithm, which is a simple but effective constraint-based structure-learning algorithm developed by Spirtes and Glymour (1991). It works by connecting all nodes in a network initially with undirected edges, and iteratively deleting edges by taking a pair of nodes (X , Y) and trying to find a set of nodes S (exclusive of nodes X and Y) so that X and Y are conditionally independent given S (Ebert-Uphoff and Deng, 2012a). If no such set exists, the edge is preserved. Next, 'collider' structures (also called v structures) are identified – where a pair of edges meet a single node, such that the node has two 'parent' nodes that need not be related – and as many directed edges as possible are added that satisfy the constraints which dictate that loops (cycles) or further addition of collider structures are not allowed. As the PC algorithm

results in a pDAG, bidirectional edges that appear as a result of the structure-learning stage need to be removed as Dynamic Bayesian Networks (DBNs) require fully acyclic graphs. In later chapters, arcs are removed to allow for the creation of a DBN and change a pDAG into a DAG. Bidirectional arcs are given a single direction such that only the directed arcs pointing towards major collider nodes are preserved. Doing otherwise significantly reduces the accuracy of the DAG structure and undermines the DBN as a probabilistic model, as the removal of collider structures impacts the individual probability distributions that make up a BN. This preserves the ‘explain away’ effect produced by collider structures described above.

Given that bidirectional edges have to be removed in order for the static BN to be used in a DBN, and that the arc needs to be preserved in the direction of collider nodes such as jet latitude and the NAO in Chapter 5 (see 5.3.1), a product of this is that networks in this study are built to examine the impact of variables on midlatitude circulation and not the other way around. For example, whilst the networks of chapters 4 and 5 repeatedly find a covariability between Arctic and midlatitude nodes in their static BN phase, the strength of the midlatitude-Arctic relationship cannot be quantified due to the directed nature of the graphs required by DBNs. Conditional dependence between variables in DBNs is explicitly non-causal (Milns et al, 2010) as causal sufficiency is usually not satisfied in observational climate science. Any discussion of Arctic-midlatitude influences therefore refers to external influence on the jet and not a simple cause-and-effect interpretation. This is an important distinction considering the central role anomalous midlatitude circulation plays in sea ice loss and AA (Kapsch et al, 2016).

A second structure-learning algorithm is used throughout this project to ensure robust DAG findings. Score-based algorithms generate a set of possible network structures, and each is scored based on the model fit. The Hill Climbing (HC) algorithm implements this using a greedy search that iteratively adds, removes and reverses arcs between all variables, finding the DAG with the best fit (Russell and Norvig, 2014). A number of networks are created with a single change to a link as the algorithm maximises the score at each stage. Network scoring is achieved with the Bayesian Information Criterion (BIC), which penalises models for overfitting and approximates the DAG’s posterior probability (Schwarz 1978). The BIC is used in this project as it takes the number of observations into account, in contrast to other scoring criteria.

2.1.3 Parameter Learning

Once the structure is defined, the parameters of a distribution can be estimated for each node in the BN and every configuration of its parents. The Expectation-Maximisation (EM) algorithm (Bilmes et al, 1998) was used to parameterise the DBNs presented in chapters 4 and 5 using the junction tree inference engine. This includes the estimation of parameters

for both hidden and observed variables. The EM alternates between finding the expected sufficient statistics using the log-likelihood function, and maximising the estimated likelihood function until a local maximum is converged upon and the parameter estimates are returned (Dempster et al, 1977). Once all the CPDs of the BN are defined, the model can then be used to predict the node values in a test dataset to determine the model fit; this process is used to compare predictive accuracy amongst different model types in later chapters.

BNs can suffer from over-parameterisation which can lead to inaccurate variable prediction, as shown in many of the networks in chapters 3 and 4. The term ‘overfitting’ is used when a statistical model describes too much of the noise rather than the underlying relationship, i.e. the structure. In machine learning, noise in the ‘train’ dataset is learned to the extent that that model performance is negatively impacted when it is applied to a new dataset, i.e. the ‘test’ data (Brownlee, 2016). This could involve the model having too many parameters relative to the number of observations, as is frequently the case in Chapter 3. The predictive performance therefore decreases because the model overreacts to small noise fluctuations in the training data. Throughout this project, efforts are taken to reduce the level of overfitting such as the development of more complex network types, the reduction of variables and the forced deletion of arcs to investigate the impact on predictive accuracy, as in section 5.7.3.

2.2 Dynamic Bayesian Networks

A Dynamic Bayesian Network (DBN) is an extension of a BN over time, whereby nodes are variables in a given time slice. Directed graphical models are used to model time series. Although edges between nodes can be directed and undirected, DBNs require fully directed edges both within and between time slices, meaning that the CPD of each variable can be estimated for each node independently (Murphy, 2002). The structure selection and parameter estimation of DBNs is an extension of the method used for BNs with no time element, called ‘static BNs’, described in sections 2.1.2 and 2.1.3. The terminology regarding graph structure, i.e. nodes, edges and probabilities, is also the same for DBNs, with the addition of ‘intra-slice’ and ‘inter-slice’ topologies which are used to refer to the graph structure within a time slice and between two time slices respectively.

2.2.1 Hidden Markov Models

Hidden Markov Models (HMMs) can be thought of as the simplest form of DBNs. They consist of a set of observed variables, the time series or input dataset, and a set of hidden nodes. A sequence of hidden variables (explained below) are predicted from the set of observed variables. A HMM is a BN unrolled through time, with the sequence of observations used to predict the best sequence of hidden states. The probability that a state will remain in the same state or switch to another is dictated by the state transition probabilities, such that the next state in a HMM is dictated only by the current state and not by any of the other past

or future states of the model, which satisfies the *Markov* property. This refers to the assumption that the future is independent of the past given the current state. The internal state changes are not viewable in the case of HMMs, hence they are ‘hidden’, and only the sequence generated by the hidden states is observable.

A HMM factorises the joint probability for a sequence of observed and hidden states like so:

$$p(\{S_t, Y_t\}) = p(S_1)p(Y_1|S_1) \prod_{t=2}^T p(S_t|S_{t-1})p(Y_t|S_t) \quad (2)$$

As described in Gharamani and Jordan (1997), the hidden state is a single random variable that can take on a discrete value, $S_t \in \{1, \dots, n\}$, and the state transition probabilities, referred to as $p(S_t|S_{t-1})$, are given by a $n \times n$ transition matrix. The transition matrix relates to the ‘inter-slice’ links in the graph structure referred to above.

A distinction is made between the HMMs and DBNs produced in this project; HMMs have a fixed structure, whilst a DBN has more flexibility in terms of graph structure (Murphy, 2002). When the probability distributions of the hidden states are not known, they can be estimated with the Expectation Maximisation (EM) algorithm (Dempster et al, 1977). The EM algorithm iterates between the ‘E’ step, which fixes current parameters and computes posterior probabilities over all the hidden states, and the ‘M’ step, which maximises the expected log likelihood of observations as a function of the parameters (Gharamani and Jordan, 1997). Parameter estimation with hidden variables is achieved with the EM algorithm in this project.

2.2.2 Dynamic Bayesian Networks

The term ‘DBN’ is used to refer to all dynamic models that use a structure-learning approach in this project, with ‘HMM’ used to refer to all fixed structure dynamic networks. A DBN makes use of 2 time slices ($t-1, t, \dots$) to ‘unroll’ a BN into T time slices (Fig. 2.1c), such that model structure and parameters do not change over time and the model stays time invariant (Murphy, 2001a). As with HMMs, DBNs represent a *Markov process* because the state of a system t depends only on the preceding timestep and state at $t-1$ (Mihajlovic and Petkovic, 2001). In a DBN, the transition and observation probabilities need to be defined, as well as the initial state distribution. The conditional distributions of a DBN are defined using a two-slice temporal Bayesian Network, or a ‘2TBN’, with the transition and observation models then defined in the 2TBN as a product of the CPDs:

$$p(Z_t|Z_{t-1}) = \prod_{i=1}^N p(Z_t^{(i)}|\mathbf{Pa}(Z_t^{(i)})) \quad (3)$$

where $Z_t^{(i)}$ represents the i 'th node in timestep t and $\mathbf{Pa}(Z_t^{(i)})$ represent the parents of $Z_t^{(i)}$ which can be in either the t or $t-1$ time slices (Murphy, 2002). The parameters of the first slice of a 2TBN are empty, whereas the second slice nodes have an associated CPD.

The initial state distribution can be defined using a static Bayesian Network. This static BN, alongside the 2TBN, make up the DBN. The joint distribution for a DBN with a sequence of length N can then be created by unrolling the network to N slices, with the CPDs then multiplied together:

$$p(Z_{1:N}) = \prod_{t=1}^N \prod_{i=1}^T p(Z_t^{(i)} | \mathbf{Pa}(Z_t^{(i)})) \quad (4)$$

This process is shown graphically in Fig. 2.2, whereby a 2TBN for a HMM (a) is unrolled over timesteps into a sequence of length $T=4$ (b).

The intra- and inter-slice connections can be seen in Fig. 2.2, where the hidden variable nodes are linked between time slices (i.e. H_1, \dots, H_4 in timesteps $t-1, \dots, t+2$), allowing for conditional dependencies between variables at different time slices. DBNs thus allow states at time t to be conditionally dependent on states at time $t-1$ as well as the states of other nodes within time t .

2.2.3 Hidden Variables

An advantage of BNs is that they can be used to model observed and unobserved data, as hidden nodes can be inferred from the values of the observed nodes. A hidden variable (HV) can be used to capture the underlying state of a time series or represent a variable of interest to the network that cannot be directly observed (Murphy, 2012). HVs may represent something of importance theoretically to the modelled system, or a process or driver that shares interdependencies with the variables but was not explicitly constrained within the model structure for one reason or another (Trifonova et al, 2017). This could occur where no

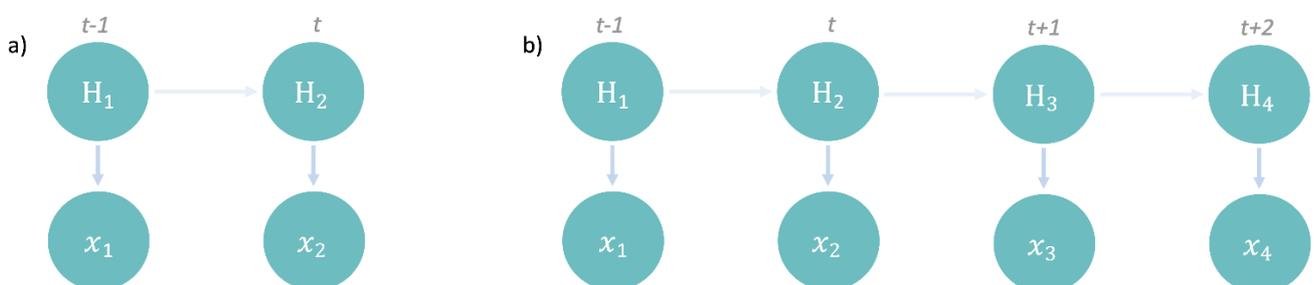


Figure 2.2 An example 2TBN DAG (a) and the same model unrolled over $T=4$ slices (b).

data exists on a phenomenon, or where model and methodological approach constraints dictate the exclusion of certain system components; for example, a model where a set of symptoms are observed, but the disease is unknown (Murphy, 2012). An example of a hidden common cause is cloud cover, which would be the parent node of 'ultraviolet (UV) levels' and 'chance of rain' in a BN as it impacts the probability of both variables. As UV levels cannot impact the chance of rain in a physical sense, it could be wrongly concluded that the UV and rain nodes were causally connected without the inclusion of cloud cover in the network, which could be modelled as a hidden variable.

HV nodes encoded within networks can point to any number of observed variables. The HV then reflects the changes in system interactions between the observed nodes it is linked to, as its value is inferred to maximise the fit (log-likelihood) of the model to the data. As with the observed parameters of HMMs and DBNs in this project, the HV nodes are parameterised with the EM algorithm (Dempster et al, 1977). Discrete HVs with three possible states are used in this work. Inference and variable prediction to check model accuracy was achieved with the junction tree algorithm (Murphy, 1998). Because the HVs are parameterised using the observed variables, the nodes that HVs point to are important for model accuracy (Friedman et al, 1997), and the mean values of the observed variables associated with each hidden state can be calculated for more information on their meaning. Chapter 4 presents an analysis of the impact of different HV configurations, and a hidden state mean value analysis is shown and discussed in Chapter 5.

HV approaches have great potential for climate data applications because they maximise model accuracy, as shown in DBN models constructed to investigate jet stream variability in Chapter 5. Graphical models that incorporate HVs may result in structures that are significantly more similar to the climate system we are trying to model; simpler models are learned which are less prone to overfitting whilst being more efficient for inference (Tucker and Liu, 2004). Given the challenges associated with reproducing Arctic-midlatitude weather linkages, due to noisy internal dynamics and the time-constrained nature of AA processes, potential improvements in model accuracy make HVs worth consideration. Graphical models with HVs inferred from observed data have been largely untouched in climate science studies, but their capabilities have been explored and proven in ecological system analyses (Trifonova et al, 2015; 2017; 2019; Uusitalo et al, 2018).

2.3 Summary

The experimental steps of the full DBN approach taken in chapters 4 and 5 can therefore be summarised as follows. Data were split into training and testing datasets, with the proportions varying from a 50:50 split in Chapter 3 to 80:20 in chapters 4 and 5. For the DBNs, the train dataset structure was learned using either the PC or HC algorithms as a first

step. The PC algorithm uses the *fisher* z test for conditional independence, and a standardised alpha value of 0.01 is used across models unless specified. For the HC algorithm, the BIC is used to score candidate networks and optimise model structure; two algorithms are used to ensure the network structure results were as robust as possible. HVs were coded into the structure after the structure-learning stage for HMM and DBN models, and any bidirectional arcs were removed such that collider structures were preserved in the case of the networks learned with the PC algorithm only. Next, the parameters were estimated from the data using the EM algorithm for both the observed and hidden sets of variables. Finally, predicted test dataset values were generated using inference, and test data predicted values were then compared against observations to validate the networks and select the best performing model. This process was often run multiple times to investigate differences in midlatitude circulation proxies (i.e. 'jet' versus 'MI' DBNs in Chapter 5), HV structures (Chapter 4), DBNs with and without HVs (control runs, Chapter 5), lead times for tropical variables and nonstationary Arctic-midlatitude linkages (both Chapter 5) to see what impact they had on model performance.

Other time series models exist that could be used as a methodological approach for climate teleconnection analyses. State space models incorporate a range of probabilistic graphical models that describe dependence between the hidden variable and the observed measurement, and can be used to reconstruct and predict the state of a dynamical system. HMMs and state space models both give you a posterior distribution over the hidden state at t given the data up to and including time t , but HMMs by convention have discrete hidden states whereas state space models denote models with continuous hidden states (Petris and Petrone, 2011). For this thesis' contribution to the scientific field, HMMs were considered a good starting point because any physical processes targeted by the hidden variable analysis (i.e. the identification of an AA period, or the impact of stratospheric polar vortex states on the propagation of tropical signals to the midlatitudes) were likely to be described more effectively in discrete terms. However, future analyses could certainly make use of the wider range of models available in this area of computer science research, with a particular focus on observation noise in timesteps and the strong degree of random variability one has to account for in the field of atmospheric science.

2.3.1 Graphical Models in Climate Science

Graphical models are increasingly being presented as a robust methodology to investigate relationships between climate variables in observed datasets (Ebert-Uphoff and Deng, 2012a; 2012b; Kretschmer et al, 2016; 2017, Di Capua et al, 2019), as an alternative to correlation analyses. Ebert-Uphoff and Deng (2012a) introduced Bayesian Networks to the climate teleconnection community; relationships between four prominent daily-averaged

indices were demonstrated using the PC algorithm. Using a similar approach to this project, Ebert-Uphoff and Deng (2012a) propose independence graphs using structure learning in a graphical model framework as a way of eliminating indirect connections and returning functional networks, and suggest that BNs that focus on the flow of information throughout the network may perform better than correlation-based networks focusing only on node similarity. This approach was used recently to identify links between the NAO and MJO indices and an apparent state dependence in stratospheric linkages between the two (Barnes et al, 2019). Another seminal study produced evidence for a dynamically-driven impact of sea ice on Eurasian circulation using Causal Effect Networks (CENs) (Kretschmer et al, 2016), which makes use of a range of time lags to determine ‘causal’ drivers of variables. More recently, the combination of a causal discovery algorithm with a response-guided community detection scheme, whereby correlation maps run at different lags provide regions of spatial data likely to be driving variability in the chosen index of interest, has led to the prediction of weak stratospheric polar vortex states up to 30 days before (Kretschmer et al, 2017) and skilful prediction of the Indian Summer Monsoon on seasonal timescales (Di Capua et al, 2019).

2.3.2 Terminology Differences between Climate and Computer Sciences

It is worth pointing out that statistical models, for example probabilistic graphical models like Bayesian Networks, are frequently referred to as ‘models’, and as such some confusion may exist regarding their distinction from numerical atmospheric models. Atmospheric model studies are fundamental to the interpretation of results presented in this work, so are frequently addressed in the discussion sections. This thesis attempts to make a clear distinction between the two, frequently prepending ‘atmospheric’ in front of ‘model’. ‘Model accuracy’ or ‘predictive accuracy’ are however used to describe BN models throughout.

Terms like ‘forecasting’ and ‘prediction’ have established meanings for climate scientists which imply the forecasting of climate conditions into the future. They can also refer to the n -step ahead forecasting method driven by BN inference (see section 2.2.1) that is used to assess network performance and accuracy in this project. The word ‘forecasting’ is therefore used only to mean subseasonal-to-seasonal forecasting (outside Chapter 2), and ‘prediction’ or ‘predictive accuracy’ are frequently used instead to refer to BN performance, assessed using inference. ‘Predictive accuracy’ is an assessment of the model fit: test dataset values are predicted using the structure and parameters learned from the train dataset and then compared against the observed test values. The accuracy of fit is quantified using sum of squared error (SSE) for each variable over all time resolutions. Because inference is used to assess network performance, all variables are predicted, unlike other climate science applications of statistical models where a single variable of interest (for example, jet stream latitude) would be predicted based on other input variables to establish them as predictors of

the response variable. As is explained above, BNs exploit conditional independence graphs and structure-learning, so there is no response variable in this manner.

2.3.3 Chapter Outline

Chapters 3, 4 and 5 are the core data chapters and contain findings from increasingly complex network architectures. Chapter 3 presents a static BN analysis to test the structure-learning phase of BN construction on a number of climate variables. In Chapter 4, DBN methods including HVs and different network architectures are applied to a simple 4-variable dataset to investigate Arctic-midlatitude linkages with a number of BN tools. A full investigation of the impact of Arctic variables on midlatitude atmospheric variability, relative to other remote drivers and internal variability, is carried out in Chapter 5. These findings are backed up by an analysis of tropical lead times and nonstationary Arctic-midlatitude linkages towards the end of Chapter 5. Each data chapter (3-5) has a summary of the results and their implications either for further study presented in later chapters, or for the AA-midlatitude research field as a whole. Chapter 6 therefore concludes the thesis by highlighting the contributions of the research presented, followed by an assessment of the limitations of the method and study design used, and finally a set of recommendations for future work to be carried out based on the findings of chapters 3-5.

Chapter 3: Learning Climate Network

Structures with Static Bayesian Networks

3.1 Introduction

This chapter details the use of static Bayesian Networks (BNs) to detect relationships between geographically disparate climate datasets. A network of 6 months' data and two others of 35 years in length were built as an initial investigation into the appropriateness of the use of BNs for link discovery in a complex nonlinear system. As such, the chapter represents the foundational work of the project undertaken for training in R programming, spatio-temporal (GIS) analysis and Bayesian Network analysis using the simplest form of BNs; static models.

The methodology employed here can be summarised in two main sections. The experiments are numbered as 1, 2a and 2b:

1. A 'draft' network using 6 months of data and two spatially-averaged regions of atmospheric variability
2. Two 35-year networks of (a) daily and (b) monthly summarised data using two jet stream proxy measurements over one region of variability.

Linkages between Arctic Amplification (AA) and midlatitude weather extremes are the subject of continuing debate; clear cause-and-effect relationships are elusive because complex nonlinearities exist in linkages (Overland et al, 2016), presenting a number of problems. Firstly, a large degree of variability in the climate system coupled with the small signal-to-noise ratio of AA due to its recent appearance as a phenomenon (since the late 1990s) masks any simple linear relationships. Intermittent forcing (i.e. AA processes active in certain seasons) and possible dependencies of drivers on background or concomitant processes (i.e. dependency on the Arctic Oscillation phase) are further considerations (Shepherd, 2014). And finally, numerous influences from within the Arctic, midlatitudes and tropics (in particular sea surface temperatures) could be acting upon midlatitude atmospheric flow (Overland and Wang, 2015).

System-level approaches like probabilistic graphical models may address many of these complexities. The purpose of this chapter is to investigate the suitability of static Bayesian Networks for the research field, using networks with different data (variables) to

investigate the fit and climatological plausibility of the results. A description of the variables used follows, with methods detailing static BNs and the steps employed, results showing graphical models and their fit, and a discussion of network performance and subsequent steps for continued research.

3.2 Data

A total of 11 variables are used to determine the dependency relationships between drivers of North Atlantic circulation across several analyses. Variable and spatial scale selection (where appropriate) was based on prior studies and expert knowledge; the purpose of this is to investigate Arctic-midlatitude with a novel graphical model approach, based on the datasets used by a range of studies that have claimed to find linkages between the two (Francis and Vavrus, 2012; Kretschmer et al, 2016; Hall et al, 2017). A decrease in the upper level zonal winds (500 hPa) and a northwards displacement of the jet stream has been found in a number of studies (Francis and Vavrus, 2012; Tang et al, 2013), and whilst these results seem to depend on a number of specific methodological choices (Barnes, 2013; Screen and Simmonds, 2013), the claim that simple indicators of AA like sea ice extent can be linked with amplified upper-level Rossby waves and subsequently persistent midlatitude weather is worth investigating. For summer jet stream variability, sea ice loss during autumn and February snow cover anomalies have been found to impact jet latitude (Hall et al, 2017). During winter, Barents-Kara Seas SIC has been suggested as an important driver of atmospheric circulation through tropospheric-stratospheric coupling (Kretschmer et al, 2016). With this in mind, further investigation of these variables in a static BN framework is undertaken.

3.2.1 (1) 6-month 'draft' Bayesian Network

An initial dataset was constructed using a subset of the full date range (Fig. 3.1), with four cryospheric variables selected to represent the AA processes of sea ice extent decline and loss of snow cover. This was done for the express purpose of developing the method; six months of data is not anywhere near long enough to develop reliable results in a climate science context. Two regions of Arctic sea ice variability are selected: Barents-Kara SIC (BKSeasSIC) and North-East Greenland SIC (NEGreenlandSIC) based on the analyses of Kretschmer et al (2016) and Hall et al (2017). Sea ice concentration data were from the Nimbus-7 SMMR and DMSP SSM/I-SSMIS Passive Microwave dataset, available at <https://nsidc.org/data/NSIDC-0051/versions/1>. Snow cover extent for the North American (NAmericaSCE) and Eurasian (EurasiaSCE) regions was used, provided by NOAA's Climate Data Record (CDR) of Northern Hemisphere Snow Cover Extent and retrieved from <https://data.nodc.noaa.gov/cgi-bin/iso?id=gov.noaa.ncdc:C00756>.

Three major climate indices of importance to Northern Hemisphere atmospheric circulation are included in the dataset. The Arctic Oscillation (AO), which denotes the north-

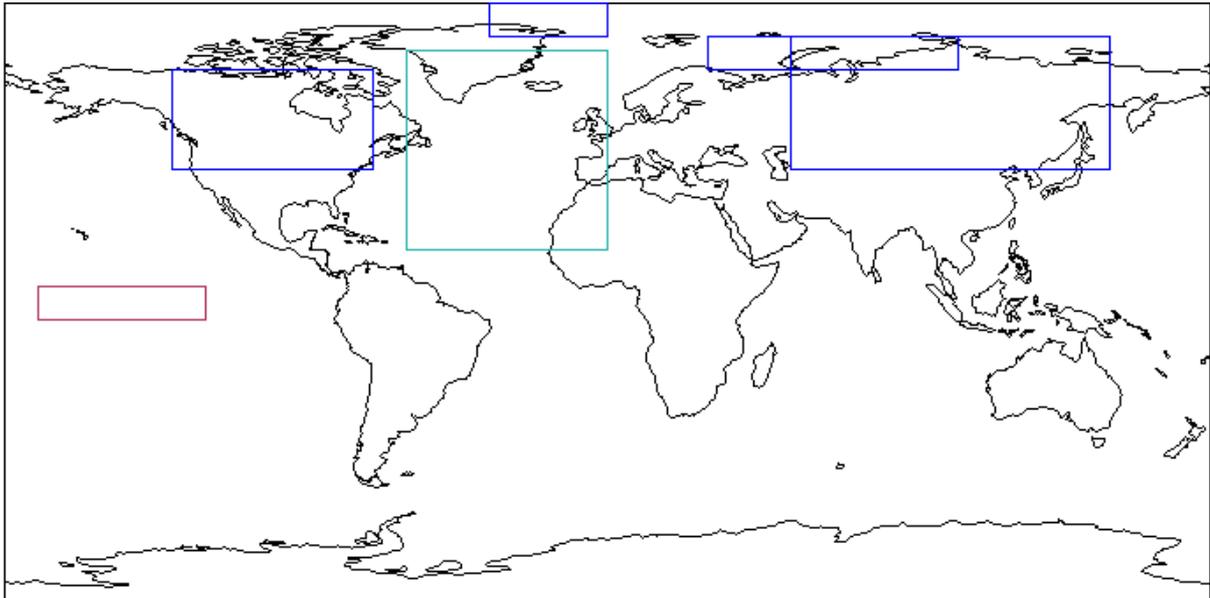


Figure 3.1 Map showing spatial scales used for (1) 6-month dataset: North American and Eurasian SCE in blue, North-East Greenland and Barents-Kara Seas SIC in blue over ocean, North Atlantic and Europe uWind in light blue and Nino 3.4 in red.

south orientation in the dipolar structure of the North Atlantic pressure field (Hurrell, 1995), and the North Atlantic Oscillation (NAO) are included as proxies of midlatitude circulation variability. The NAO is calculated by differencing the sea-level pressure (SLP) of the Subtropical High and the Subpolar Low. The El Niño Southern Oscillation (ENSO) 3.4 index is calculated using SST data from HadISST (Rayner et al, 2003), and provides the potential for tropical Pacific teleconnections which are known to influence winter North Atlantic circulation through the creation of large-scale Rossby wave trains (Trenberth et al, 1998; Scaife et al, 2017). All three indices are made available at <https://climexp.knmi.nl/>.

Finally, for an initial approximation of atmospheric variability at the 500hPa level, the zonal component of the wind, or x-coordinate, was averaged over two regions – the North Atlantic (NAtlanticuWind) and Europe (EuropeuWind). The U component of wind is the eastward speed of air in the horizontal plane measured in metres per second, provided by ECMWF’s ERA-interim reanalysis dataset (<https://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=pl/>). A time range of 2016.01.01 - 2016.06.01 (1st January 2016-30th June 2016) was used for all variables, with data either downloaded or summarised as daily averaged values. For the SCE dataset, this meant using linear interpolation to create daily values for a time series supplied by NOAA at a weekly resolution. In a geometric sense, linear interpolation is the generation of values through the use of a straight line between two adjacent points on a graph. It is appropriate here because the large-scale SCE patterns we are interested in are largely a function of seasonality; a weekly resolution in the unprepared NOAA dataset sufficiently captures SCE variability, but timesteps must match between prepared datasets.

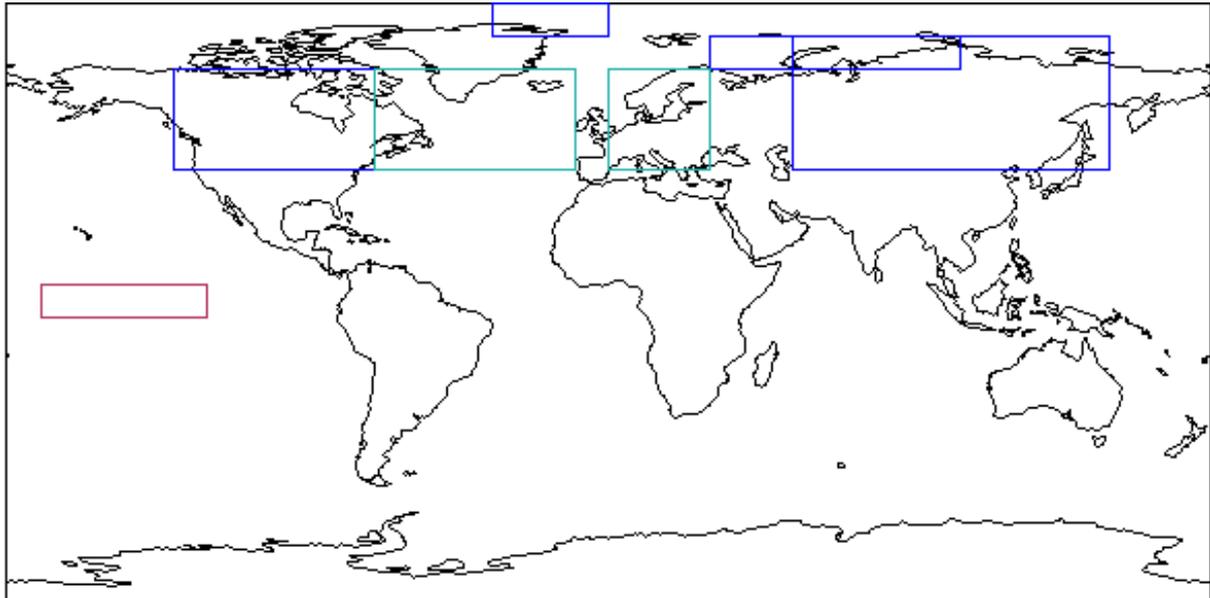


Figure 3.2 Map showing spatial scales used for (2) datasets: North American and Eurasian SCE in blue, North-East Greenland and Barents-Kara Seas SIC in blue over ocean, Jet Latitude and Jet Speed in light blue (single box) and Niño 3.4 in red.

3.2.2 (2) 35-year Bayesian Networks

Two more BNs using daily data summarised over (a) daily and (b) monthly time periods were built (Fig. 3.2) to test appropriate jet stream proxy measurements – jet latitude and jet speed, as described in Woollings et al (2010). Briefly, zonal wind speeds from 900 to 700 hPa were used to identify the Polar Front Jet (PFJ) and exclude the subtropical jet at higher altitudes, and filtered using a 10-day Lanczos low-pass filter with a 61-day window in order to remove synoptic scale variability (Hall et al, 2017). This approach works from zonal winds from 900 to 700 hPa from ERA-Interim data as raw input data (<https://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=pl/>).

All climate indices (AO, NAO and ENSO) and cryospheric regions (SCE and SIC) used in experiment 1 are included in the following datasets, with the exclusion of the zonal component of wind as a midlatitude circulation proxy; jet latitude and speed replaced the two regions of uWind as more advanced indicators of circulation variability used to better effect in several recent studies (Hall et al, 2017; Samarasinghe et al, 2017; Barnes and Simpson, 2017). Both datasets (a) and (b) used a time range of 1981.09.01 – 2016.08.01 (1st September 1981-1st August 2016), approximately 35 years.

Variable	Source	Unit	Spatial Res	Lat	Lon
Arctic Oscillation	NOAA	Geopotential height (m)	-	20-90°N	All
North Atlantic Oscillation	NOAA	Sea Level Pressure	-	-	-
El Niño Southern Oscillation	NOAA	(3.4) Sea Surface Temperature	-	5°N - 5°S	170-120°W
Eurasia Snow Cover Extent	NOAA	Fraction of snow-covered area	0.88 x 0.88	40-80°N	55-150°E
North America Snow Cover Extent	NOAA	Fraction of snow-covered area	0.88 x 0.88	40-70°N	130-70°W
Barents-Kara Sea Ice Cover	ECMWF ERA-Interim	Fraction of sea ice area	0.75 x 0.75	70-80°N	30-105°E
North East Greenland Sea Ice Cover	ECMWF ERA-Interim	Fraction of sea ice area	0.75 x 0.75	80-90°N	35°W-0°
Europe uWind (1)	ECMWF ERA-Interim	Zonal component of wind	0.75 x 0.75	40-70°N	0°-30°E
North Atlantic uWind (1)	ECMWF ERA-Interim	Zonal component of wind	0.75 x 0.75	40-70°N	70-10°W
North Atlantic Jet Latitude (2)	ECMWF ERA-Interim	Degrees (°)	0.75 x 0.75	16-76°N	60°W-0°
North Atlantic Jet Speed (2)	ECMWF ERA-Interim	m/s	0.75 x 0.75	16-76°N	60°W-0°

Table 3.1 Variables used in this analysis, data sources, the unit of measurement, spatial resolution (spatial data only) and regions used (in latitude and longitude) for all data. Note: (1) denotes variables used only in (1) 6-month 'draft' BN, a (2) those used exclusively for (2) 35-year BNs.

3.2.3 Preparation of Spatial Data

All non-index, spatial data was prepared in R Studio in the following way. Data was downloaded from the sources listed in Table 3.1 (or indirectly via KNMI's climate explorer for the three indices), loaded in as a NetCDF file ('ncdf4' package) and converted to raster format ('raster' package) by first extracting the file's dimensions. Once in raster stack format, raster layers were named as their corresponding dates, cropped to appropriate geographical extents ('Lat' and 'Long' columns, Table 3.1) and spatially averaged. This reduces the dimensionality and the variables effectively become an index, allowing for easy storage as a time series in a data frame. The table of variables is split into 'training' and 'testing' datasets (detailed below); training datasets are used for parameter learning, and testing datasets examine the fit of the training set to determine the degree of overfitting that has taken place. In contrast to the

dynamic networks built in subsequent chapters, no further formatting (i.e. detrending, anomalies) was carried out to create this series of static networks.

3.3 Methods

Networks were built in R using the 'bnlearn' package (Scutari, 2010). Model structure learning was achieved with the Hill Climbing (HC) algorithm, which implements a greedy search that iteratively adds, removes and reverses arcs between all variables, finding the DAG with the best fit (Michalewicz and Fogel, 2013). This effectively means that a number of networks are created, each with a single change to a link to maximise the score of the network at each stage. Network scoring is achieved with the Bayesian Information Criterion (BIC), which penalises models for overfitting and approximates the DAG's posterior probability (Schwarz 1978). Other network scoring approaches exist including the Akaike Information Criterion (AIC) (Akaike, 1974), simulated annealing (Bouckaert, 1995) and genetic algorithms (Larranaga et al, 1997). The BIC is used here because unlike the AIC, it uses the number of observations as a basis for the penalisation term. The structure-learning process results in an 'empty' network structure of observed variables; DAG results are shown below for all dataset types. All variables are continuous data, and all edges shown were learned from the data using the HC algorithm as no whitelist or blacklisted arcs (edges that are either forcibly 'set' by the user and hard-coded into the network, or prevented from being set by the algorithm) were set in the model building process. User-determined edges were not necessary in this case as the data was fully observed (no hidden variables are used), and the dependency structures are completely learned from the data to assess the HC algorithm's ability to pick out teleconnections between climate variables.

The dependency relationships returned are spatial relationships that are predictive in an informative way, not in a causal sense (Milns et al, 2010). Whilst the data selection relied on previous knowledge, all structures were intentionally learned from the underlying data rather than through the use of expert knowledge because of the differences in results between varying methodological approaches (e.g. Barnes, 2013). The object of using the HC algorithm was to apply a score-based local optimisation technique; HC is one of the most commonly used of these and the purpose of this chapter was not to investigate differences between score-based algorithms. After learning the DAG structure using the training dataset, the parameters are estimated from an observed sample – in this case, the testing dataset. Parameter estimates are produced using the maximum likelihood estimator, and the quality of estimates is strongly dependent on sample size (Scutari and Denis, 2014) which in this case is quite large due to the use of 35 years of daily data. The results are then visualised as plots of predicted against observed values for analysis of model performance.

3.3.1 Training and Testing Datasets

Different splits in training and testing datasets were used to explore differences in model structure and fit. For (1), a split of 5/1 months was used, whereas (2a) split the daily datasets into 1981-1994 for the training, and 1995-2016 for the testing. (2b) includes two BNs built from monthly data, split by 1981-1994/1995-2016 and 1981-1999/2000-2016 respectively. These two training and testing periods differ from each other (and differ from the 80/20 split used in the following chapters) because the dates were purposely chosen to investigate the impact splitting the datasets outside and inside of the 'AA Period' would have on results. The observational record suggests the amplified warming signal in the Arctic relative to the Northern Hemisphere only developed in the late 1990s, with Cohen et al (2014; also see Fig 1.1) suggesting this transition date occurred around 1997.

3.3.2 BIC Score

The Bayesian Information Criterion (BIC) is a criterion for scoring used to select candidate networks (Schwarz et al, 1978). The likelihood function will increase as more parameters are added during network fitting, meaning that BNs and other statistical models can suffer from overfitting. The BIC resolves this problem through the use of a penalty term for the addition of parameters in a model, effectively stopping structure-learning algorithms from creating over-connected BNs by biasing towards less complex models. The operation that returns the lowest BIC score at each step is the operation (or 'edge') that will be present in the final DAG returned by the algorithm. Here, the scoring function is increased in the form of the 'K =' argument to 5 and then 10 to investigate the differences in edges returned in the resulting DAGs, whereby a large penalising term of 10 would leave only the strongest connections in the network.

3.4 Results and Discussion

3.4.1 (1) 6-Month BN and (2a) 35-Year BNs: Daily

The 6-month BN run (Fig 3.3) resulted in a DAG structure that was not consistent with the background literature due to short time series used to create it (2016.01.01 - 2016.06.30). This initial BN, run over relatively few time points, clearly did not achieve a satisfactory level of predictive accuracy (Fig 3.4). Observed over the full daily dataset, Fig 3.5 demonstrates another common problem; an overconnected DAG suggesting the penalising term (the BIC) was not suitable for data of the size used here. Nodes with 5-7 arcs to and from other nodes are common in the network demonstrating this problem, and Fig 3.6 shows that no test set node was accurately predicted. The full dataset shows the midlatitude variables – the jet speed, latitude and NAO – as child nodes, with the NAO as the parent of both jet nodes, consistent with the NAO being an important indicator of North Atlantic midlatitude variability.

Clearly, neither a 6-month subset nor the full time series of daily data over a 35 year period were appropriate for a static BN approach; the influence of the seasonal cycle in the full 12 months of daily data, combined with the number of time points and the over-connectedness of nodes, has led to the structure-learning phase being dominated by overfitting. Whilst increasing the scoring function resulted in fewer edges and less overfitting (Fig. A3.1), no discernible change in predictive accuracy could be achieved with score alone (Fig. A3.2), so the dataset was changed to mitigate some of the problems associated with the use of a dataset with so many time points.

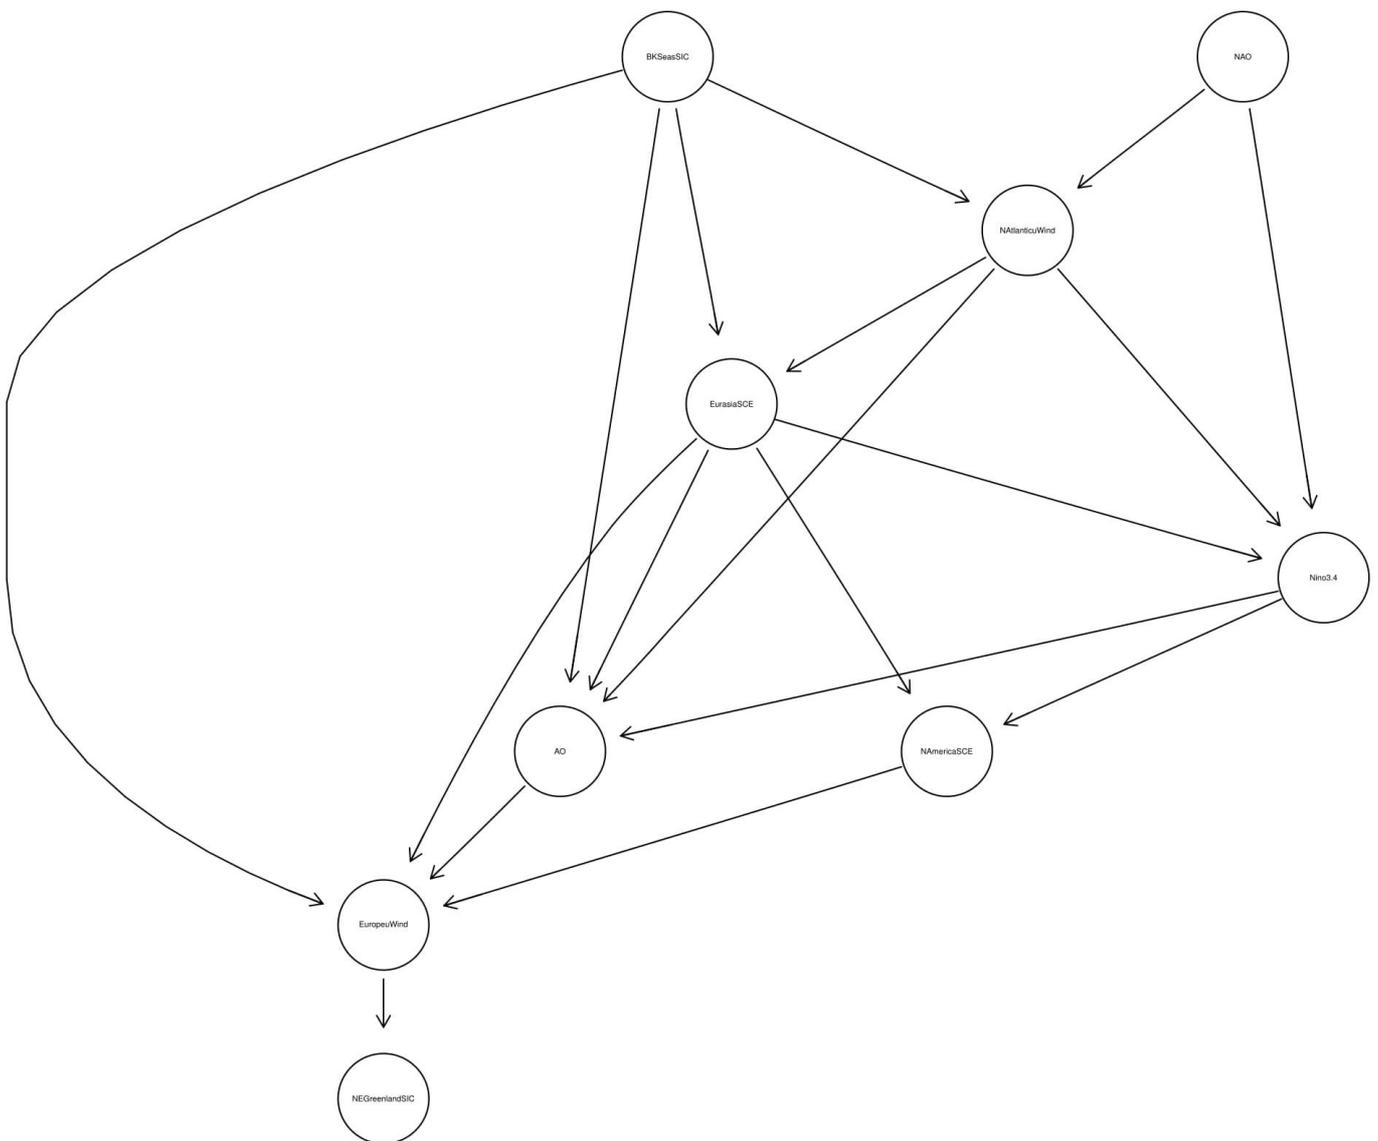


Figure 3.3 (1) 6-Month DAG.

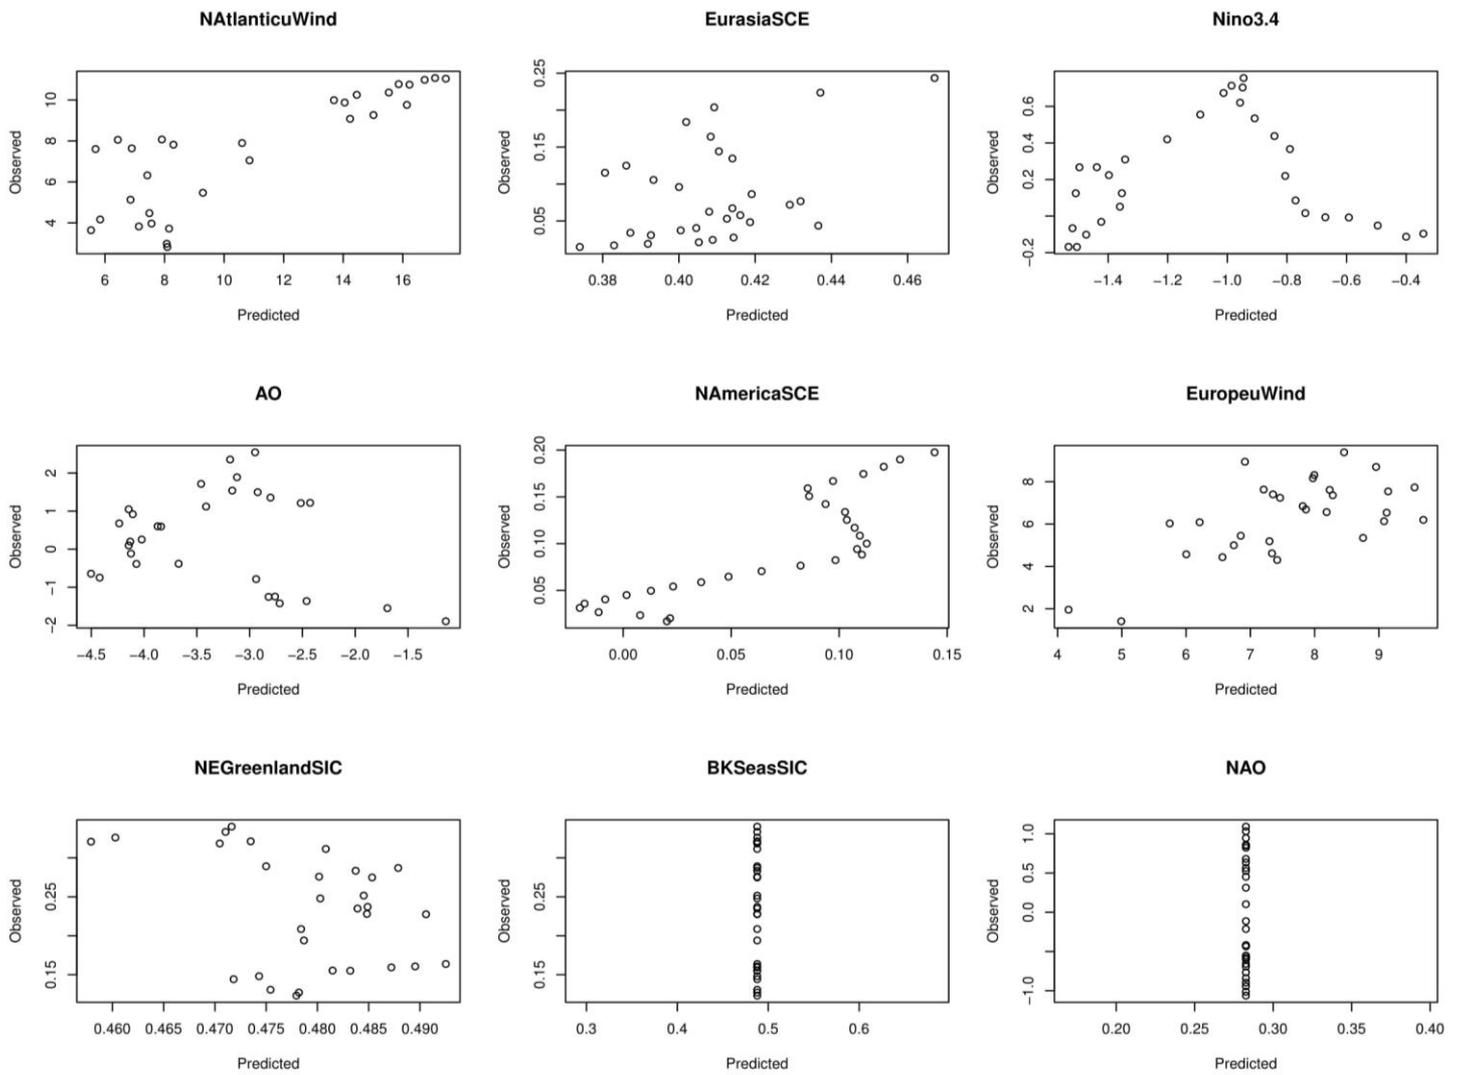


Figure 3.4 (1) Predicted fit for all variables in 6-Month BN.

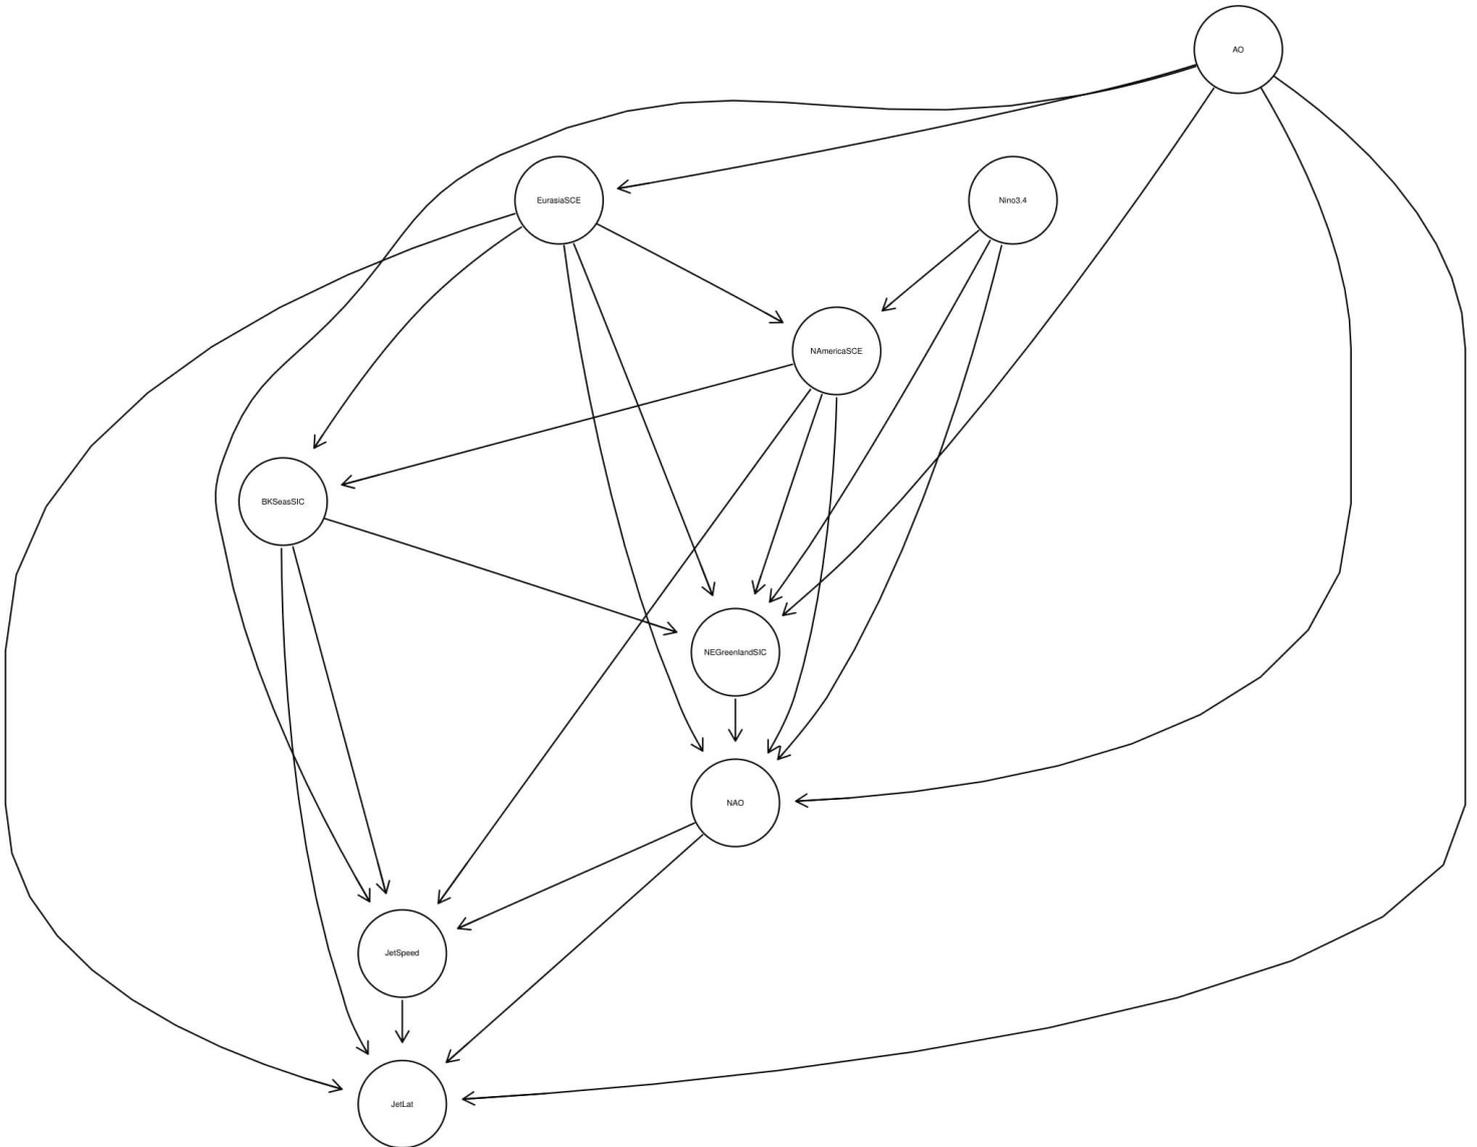


Figure 3.5 (2a) Daily 35-Year BN DAG. Note over-connectedness of nodes.

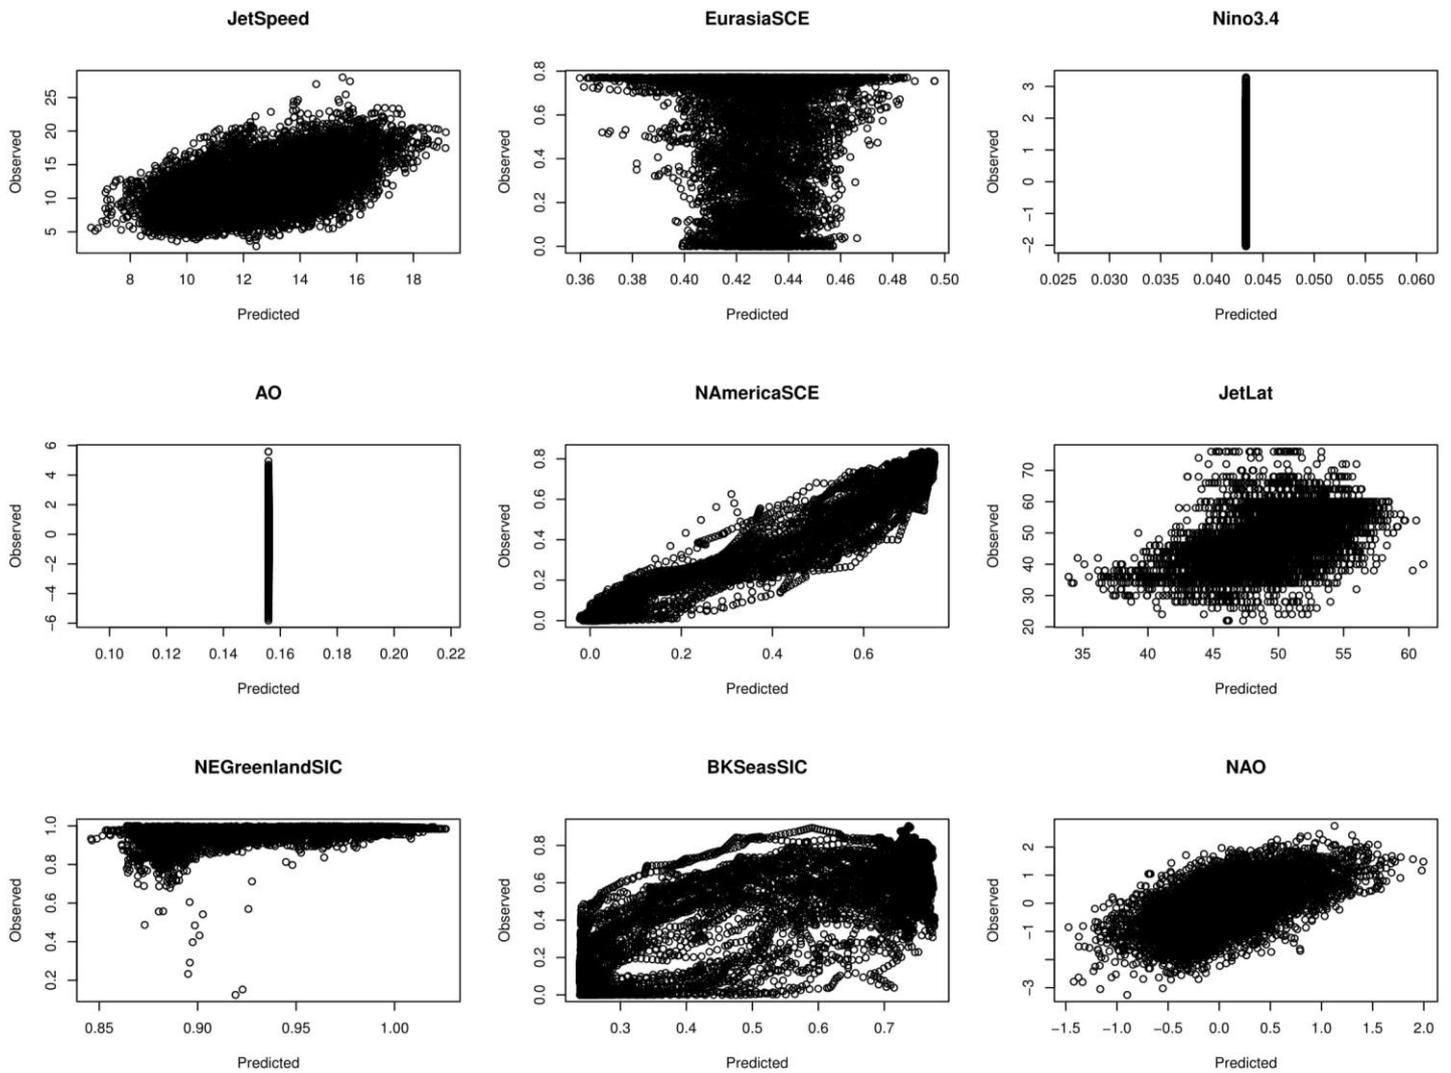


Figure 3.6 (2a) Predicted fit for all variables in Daily 35-Year BN.

3.4.2 (2b) 35-Year BNs: Monthly

The jet stream proxies – jet speed and latitude – are always fundamental to the network structures of Fig. 3.7 onwards appearing at the bottom of the DAGs. It is clear that links between JetLat and both BKSeasSIC and the NAO (as well as NAmericaSCE in cases such as Figs 3.11, 3.12 and 3.13, and indirectly through another node other times) are important, reflecting the conclusions of a number of important linkage studies for both Northern European (Petoukhov and Semenov, 2010; Liptak and Strong, 2014) and Eurasian (Kim et al, 2014; Kretschmer et al, 2016) midlatitude circulation. Whilst the connection to Eurasian extreme cold events is thought to occur primarily through the weakening of the stratospheric polar vortex, Barents and Kara sea ice variability was also found to be an important driver of the AO during winter within the troposphere using a Causal Effect Network (CEN) approach (Kretschmer et al, 2016). This is not captured in any of the static BNs presented here for the North Atlantic region. The link between BKSeasSIC and jet latitude is present in both Figs.

3.7 and 3.9, and arcs between BKSeasSIC and jet speed are present across all score types and data splits. The connection between cold Asian winters and the intensification of the Siberian High due to sea ice depletion in the Barents-Kara Seas is perhaps the most consistently found result in the AA-midlatitude literature (see review in Cohen et al, 2018), although several papers have questioned whether this link represents a causal physical mechanism sooner than internal variability (Kolstad and Screen, 2019; Warner et al, 2020). Cold temperatures are associated with reduced sea ice in the Barents-Kara Seas in numerous studies (Kim et al, 2014; Handorf et al, 2015; Luo et al, 2016), with one study finding that regional sea ice loss more than doubles the chance of severe Eurasian winters (Mori et al, 2014). BKSeasSIC links with North Atlantic jet stream variability have been found for summer circulation with lagged correlation to the preceding autumn (Hall et al, 2017), but no consensus exists on the impact of AA in the North Atlantic region.

The lack of a direct link between EurasiaSCE and jet variability in the midlatitudes is supported by Hall et al (2017), albeit for the summer months. Few studies exist linking snow cover to midlatitude circulation; over North America, spring anomalies in the 1000-500 hPa thickness range occur prominently over high-latitude land suggesting a response to earlier melt in the season (Francis and Vavrus, 2015). For Eurasian extreme weather, the role of snow cover variability is complex and unresolved, but an increase in snow has been linked to increased Siberian SLP in winter (Kretschmer et al, 2016). In contrast to the hypothesis that SLP in central Asia is influenced by late autumn snowfall (Cohen et al, 2014), other graphical model approaches have found no evidence for this link (Kretschmer et al, 2016). NAmSCE, however, is always central to the DAG structures shown in this chapter, and is suggested as an important driver of both jet proxies (Figs. 3.11, 3.12 and 3.13). Whilst the dataset used here features no lead times for potential drivers of midlatitude circulation variability, high (low) February snow anomaly years are associated with high (low) Greenland Blocking Index (GBI) values, promoting meandering jet patterns during summer (Hanna et al, 2013). The DAGs shown here tentatively suggest that snow cover variability in the North American sector are important for North Atlantic midlatitude circulation, and that the influence from Eurasian snow cover is too small to be picked up in this analysis.

Whilst a number of arcs appear across all network types, like BKSeasSIC-JetSpeed, the value chosen to score the network selection process seems to have a greater impact on the resulting structure than the time period chosen to split the data by (i.e. the test dataset). Neither the k score nor the dataset split changed the predictive accuracy significantly for the monthly BNs; the only difference is the AO node, which is predicted poorly in Fig 3.8 and fails to be predicted entirely in Figs 3.10 and A3.3-3.6. Given that the AO is only connected to three variables which effectively describe similar physical processes like the orientation of the North Atlantic pressure field and the zonal characteristics of midlatitude flow (Ambaum et al,

2001), it does not warrant inclusion in the BN investigations of following chapters. This finding contrasts Kretschmer et al (2016), who found that Barents and Kara sea ice in autumn is linked within the troposphere to a weakening of the AO in winter. High sea level pressure in the Arctic should make a negative AO more likely; in Fig. 3.7 however, both BKSeasSIC and the AO are parent nodes of important midlatitude processes like jet speed and latitude, but are not directly linked to each other.

A few notable commonalities exist between the DAG structures produced here. BKSeasSIC is shown to have limited predictive accuracy in the variable scatters, but it is almost always important to the structure of the network. EurasiaSCE is, without exception, always the parent node of NAmericaSCE, with no discernible physical explanation. Importantly, NEGreenlandSIC is never essential to the structure of the DAGs. Its child node status, unlike BKSeasSIC, suggests it is not an important driver of North Atlantic atmospheric variability, in contrast to Hall et al's (2017) study of summer circulation. Hall et al (2017) included this region of sea ice as a result of correlation mapping carried out at the beginning of their study; it is likely the region used to average SIC data over for NEGreenlandSIC does not capture enough of the SIC variability for the North Atlantic sector of the Arctic (Fig. 3.2). Subsequent studies, for example, have used larger areas of near-surface temperature to represent similar physical mechanisms linked to lower latitude circulation for the North Atlantic region (Barnes et al, 2017; Samarasinghe et al, 2017).

As these BNs use datasets which are monthly- or daily-averaged time steps from full-year time series, it is inevitable that variables which are important climatological drivers during a restricted time window would have a weaker yearly signal. Nino3.4 was consistently the first node to be disconnected from the DAG when the penalising term was increased (Figs 3.9, 3.11-3.13), most likely for this reason, despite being an important factor in winter circulation in the literature. ENSO has a complex mixture of effects on Europe, the most consistent of which resembles the negative phase of the NAO in late winter (Bronnimann et al, 2007), with a weaker meridional jet dominating circulation. ENSO is therefore expected to be more important in BNs constrained to the winter months, and could provide valuable insight into tropical influence on North Atlantic jet stream variability.

Recent studies have investigated similar metrics, Jet latitude and speed, using Granger causality methods (Barnes and Simpson, 2017, Samarasinghe et al, 2017). As with the findings presented here, jet speed was found to be related to AA: causal discovery techniques revealed that a stronger jet in the North Pacific is associated with warmer Arctic temperatures (Samarasinghe et al, 2017). By focusing on two-way linkages and feedback mechanisms, Samarasinghe et al (2017) demonstrate a relationship between jet latitude and Arctic temperatures (and a stronger influence vice versa) lagged by <25 days. Whilst this study focuses on the North Pacific jet response, it advances the field in a number of ways: it

uses temperature averaged over 70-90°N rather than sea ice concentrations which are only partially driving AA trends (Francis, 2017), it accounts for two-way linkages found to be important for sea ice decline (Luo et al, 2017), and investigates links over a range of submonthly lags.

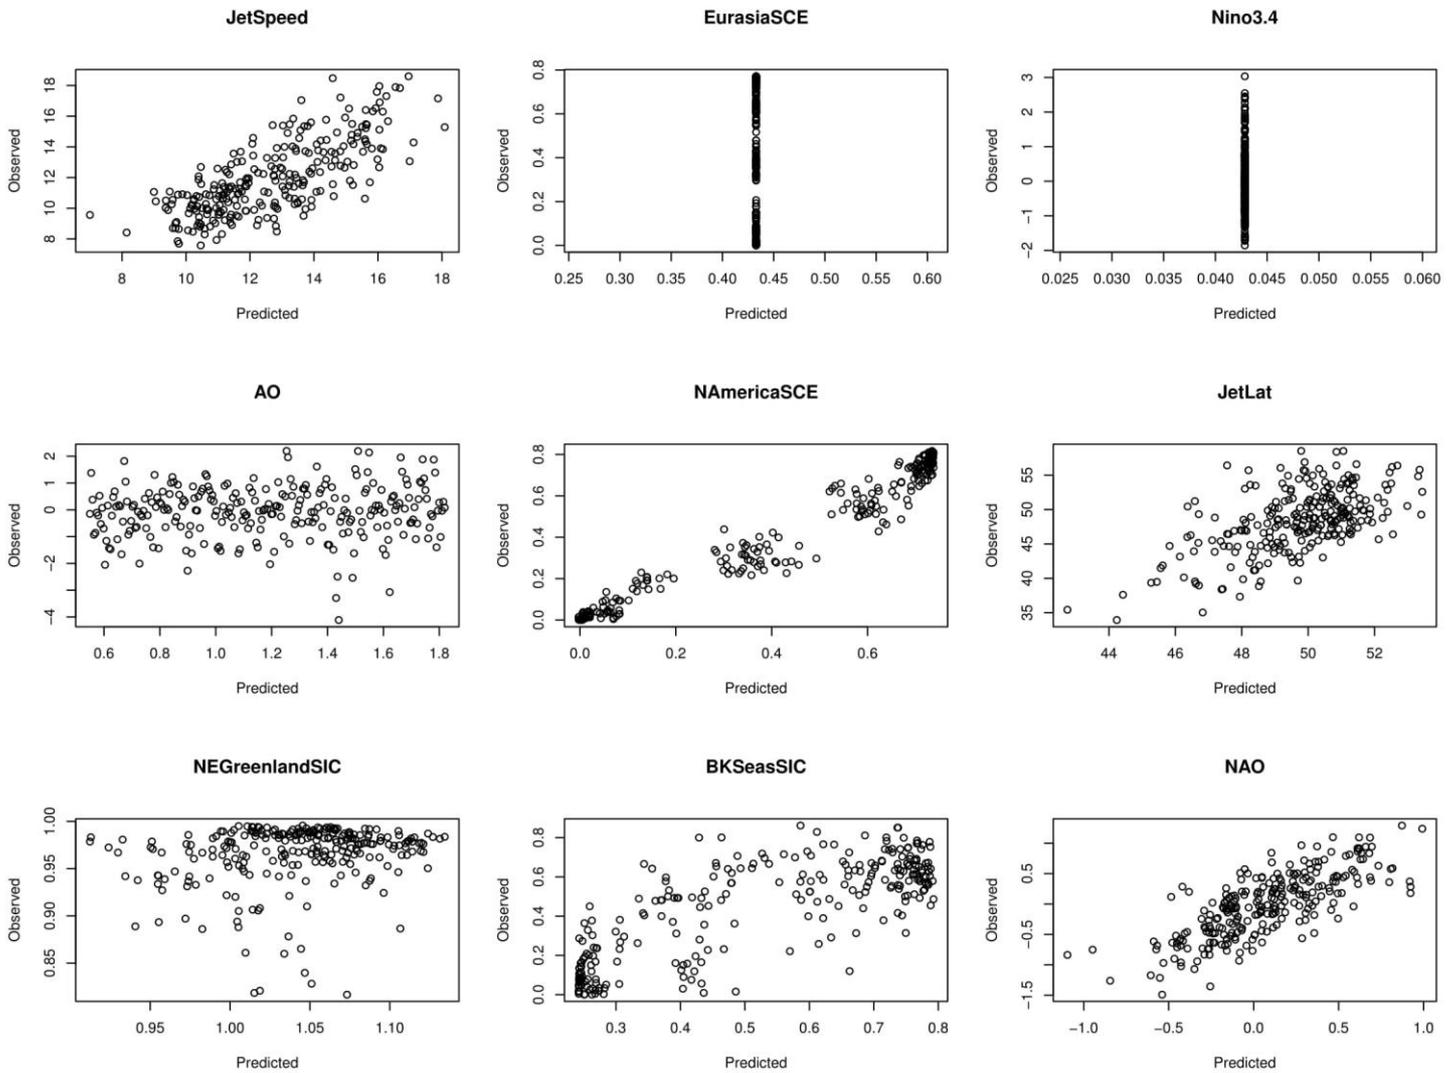


Figure 3.8 (2b) Predicted fit for all variables in Monthly 35-Year BN.

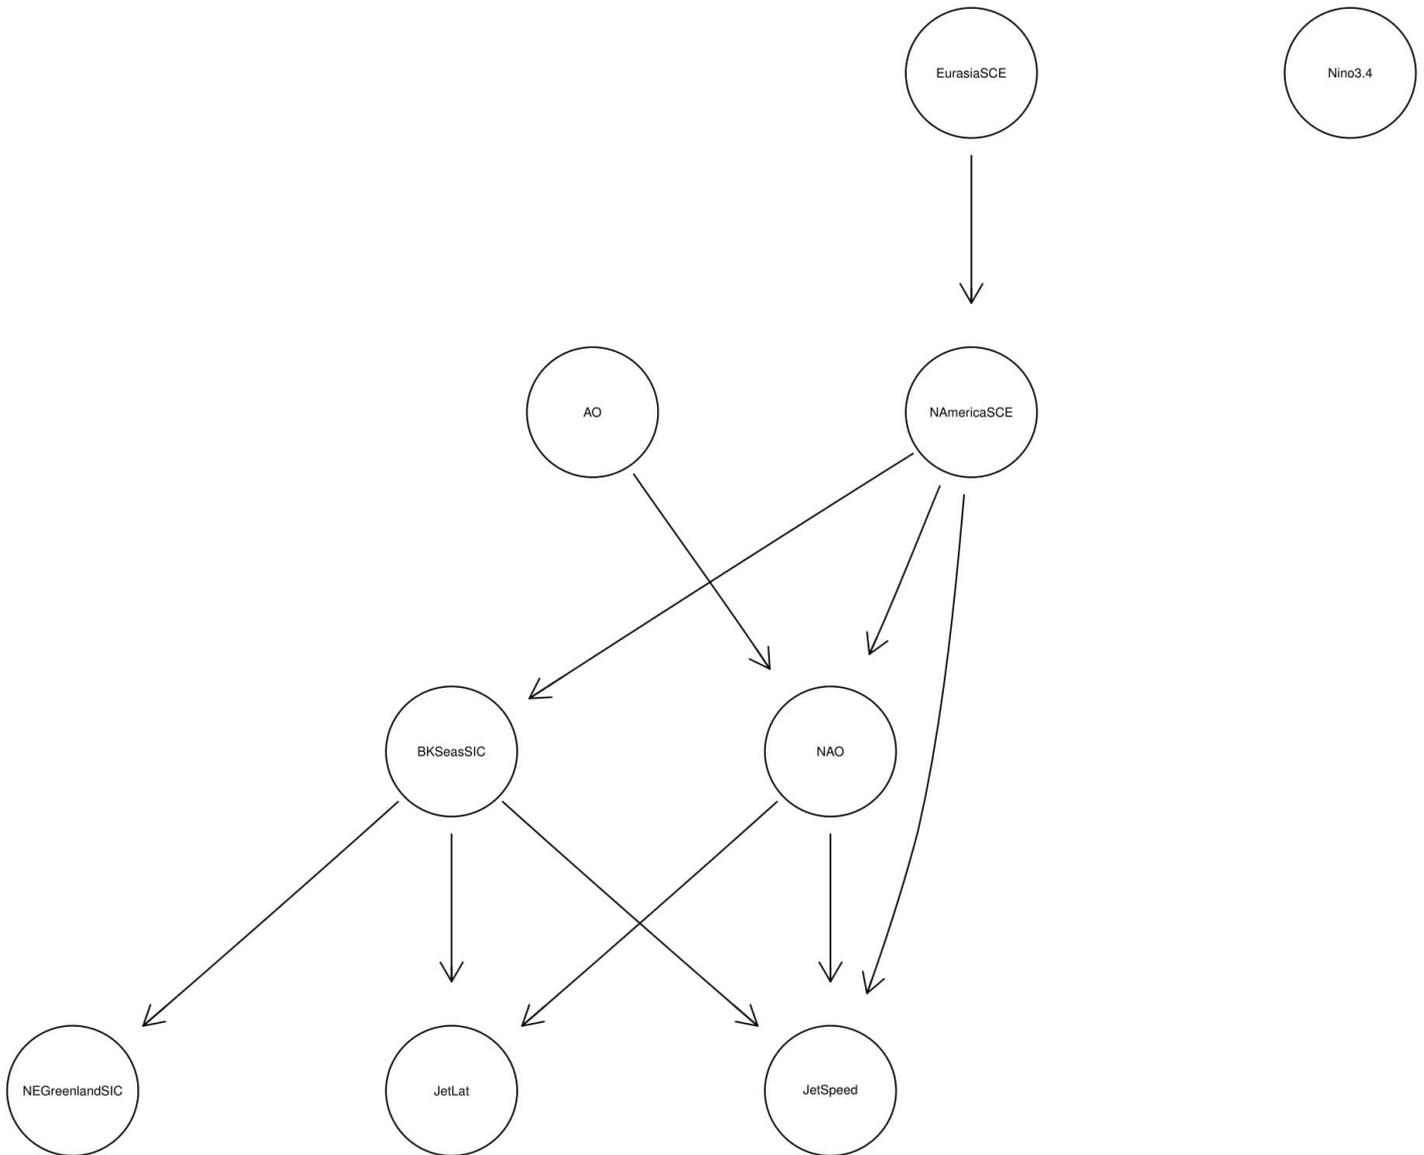


Figure 3.9 (2b) Monthly 35-Year DAG K=5.

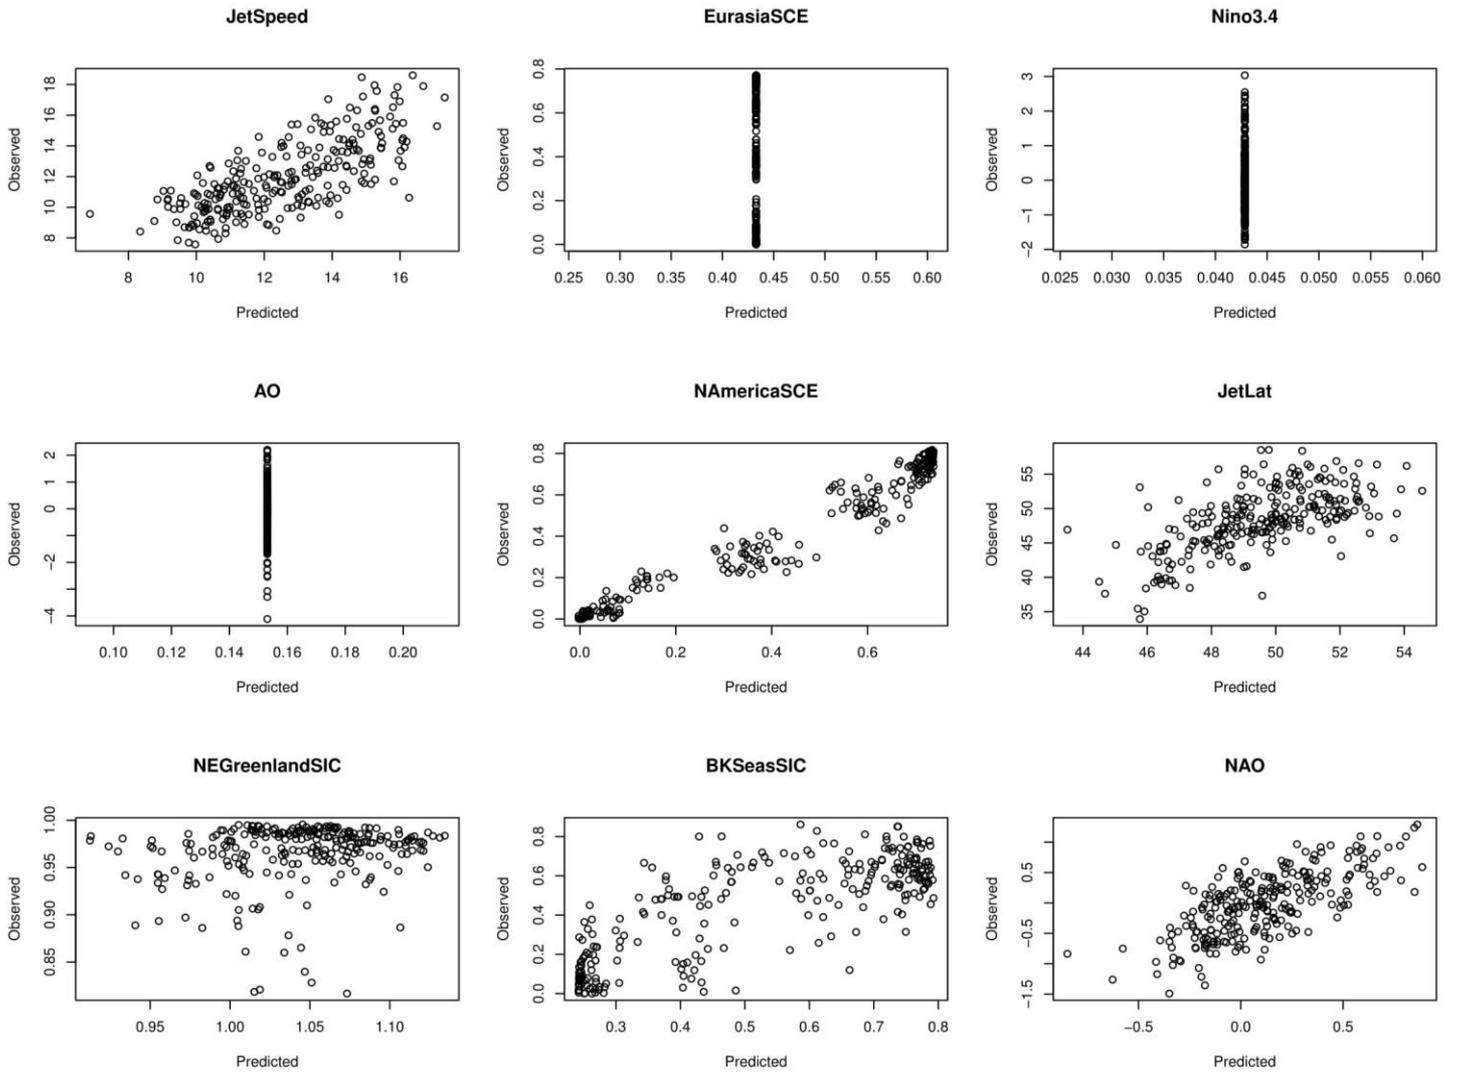


Figure 3.10 (2b) Predicted fit for all variables in Monthly 35-Year BN K=5.

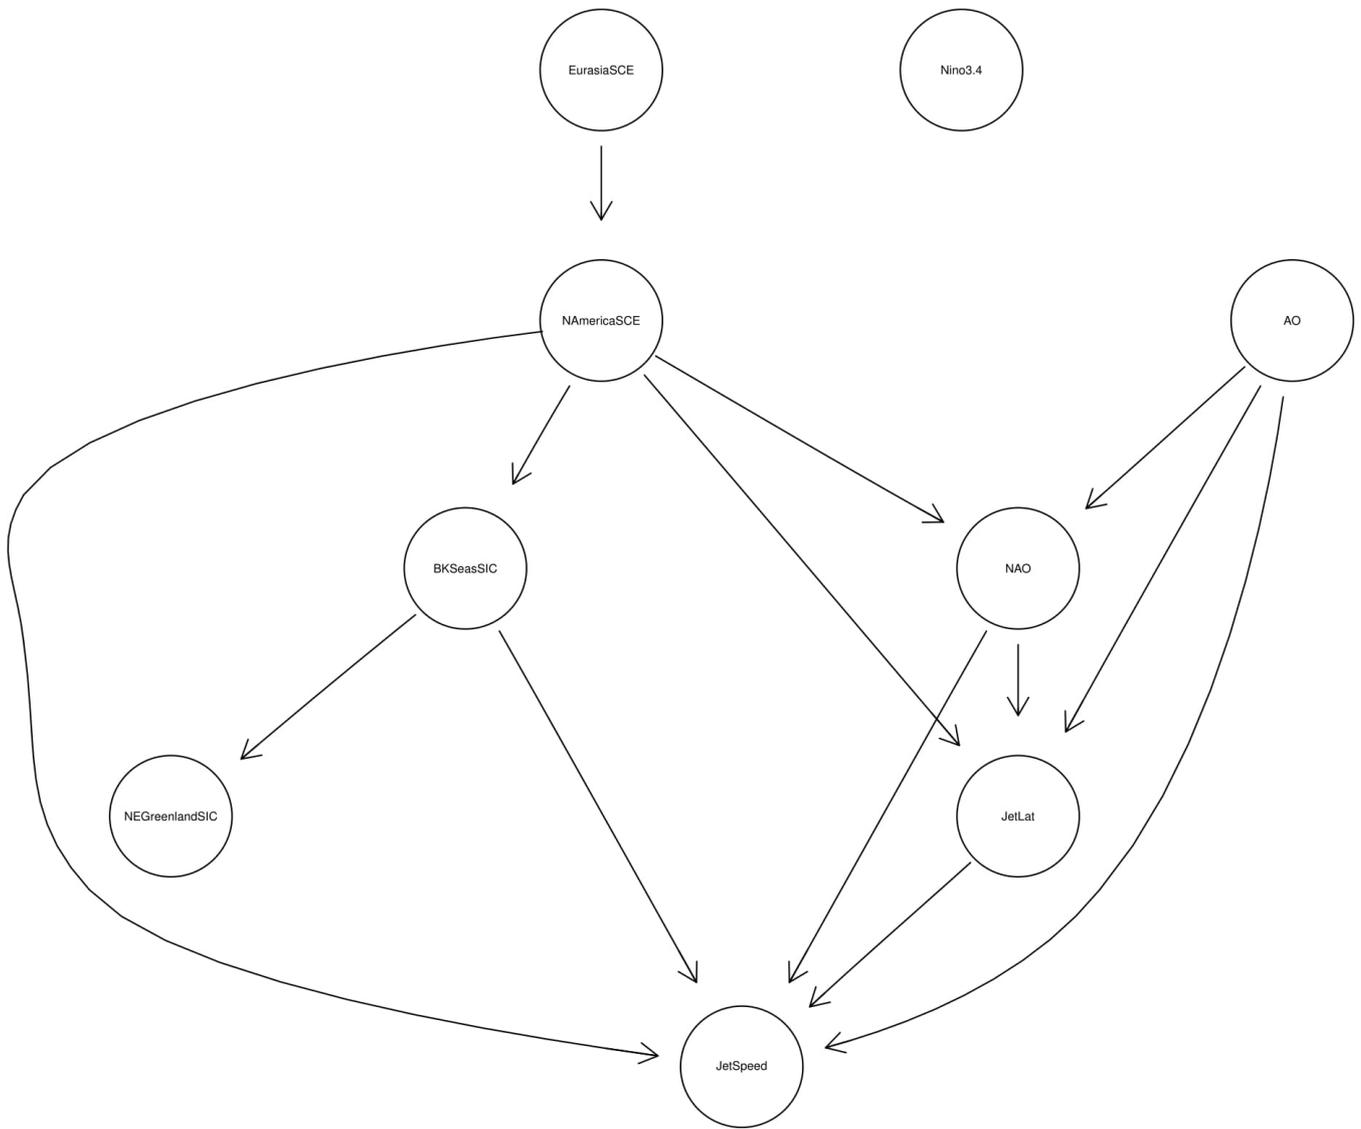


Figure 3.11 (2b) Monthly 35-Year DAG (2000 Split).

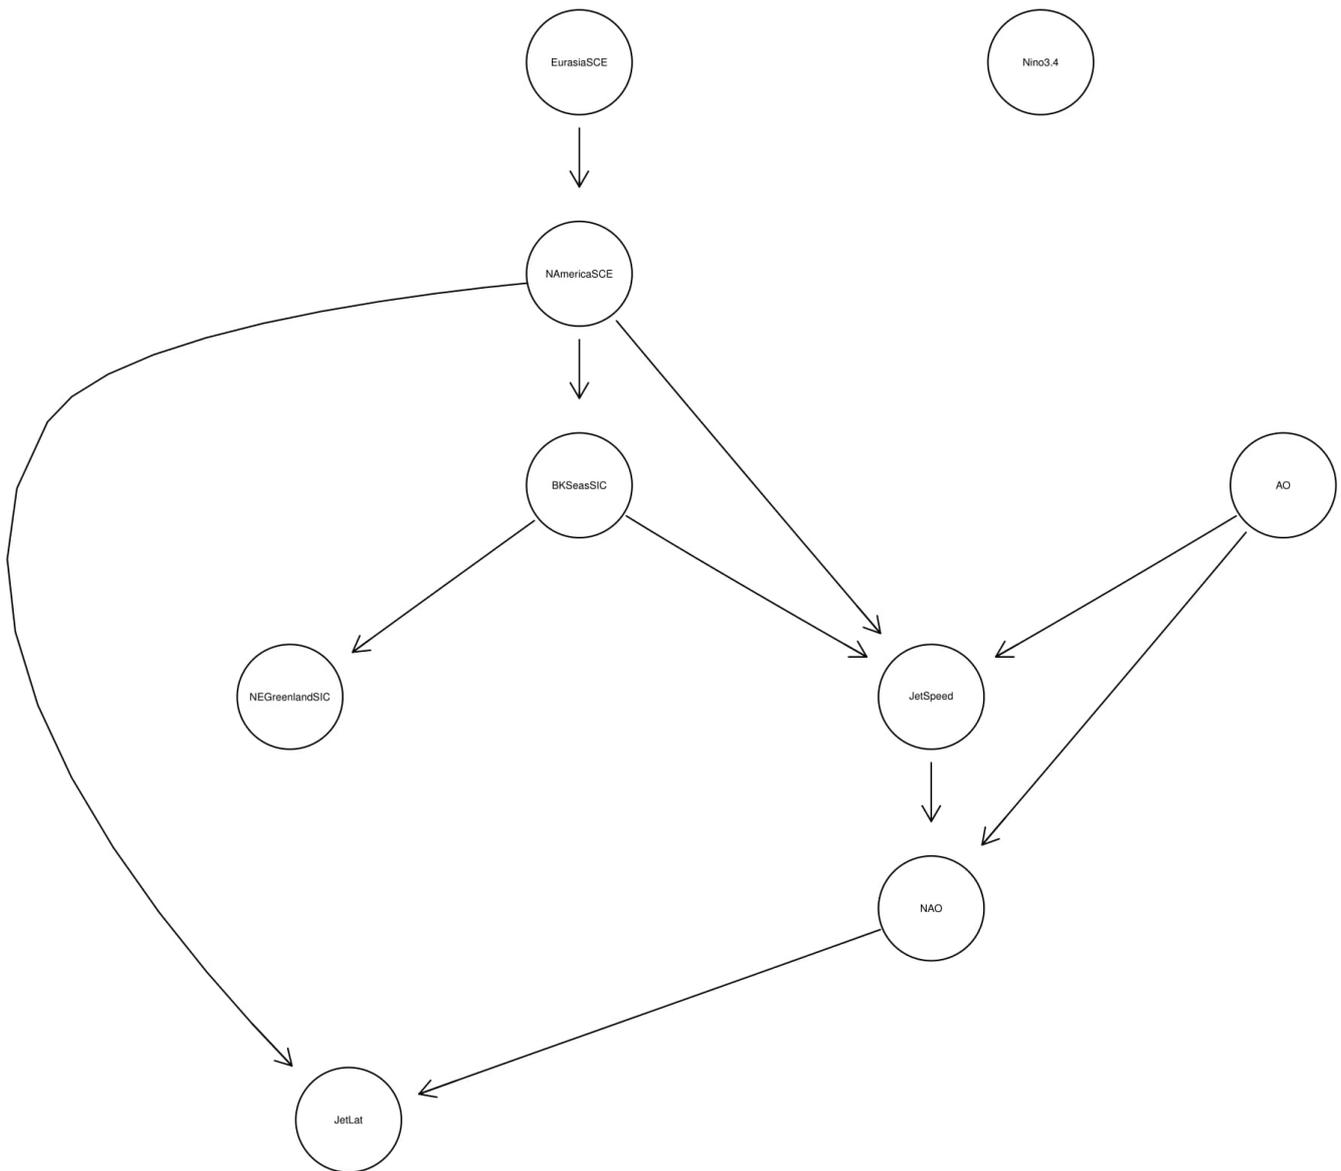


Figure 3.12 (2b) Monthly 35-Year DAG (2000 split) K=5.

3.5 Conclusion

This chapter has detailed an investigation into the suitability of static Bayesian Networks for application to the field of Arctic-Midlatitude linkages. Networks using a range of variables and timescales were constructed to assess how well the results fit our current understanding of cryospheric drivers of midlatitude circulation anomalies and their linkage mechanisms.

Barents-Kara Seas SIC and the NAO were both shown to be central to jet latitude over a robust range of timescales, which is supported by the background literature (Kim et al, 2014; Mori et al, 2014; Handorf et al, 2015; Luo et al, 2016) on Eurasian sea ice depletion and cold extremes in East Asia. A lagged relationship between Barents-Kara SIC and jet latitude has been found for the North Atlantic, whereby summer jet latitude was influenced by the previous autumn's sea ice concentration in this region (Hall et al, 2017). Whilst the season and methodological approach differ for this study, the lagged association points to the need to investigate relationships between SIC and jet stream variables at a range of different lead times. The DBN-based approach taken in following chapters that investigates different time lag associations attempts to resolve this problem. Two studies using similar methodological approaches currently exist for the North Atlantic region: Barnes and Simpson (2017) and

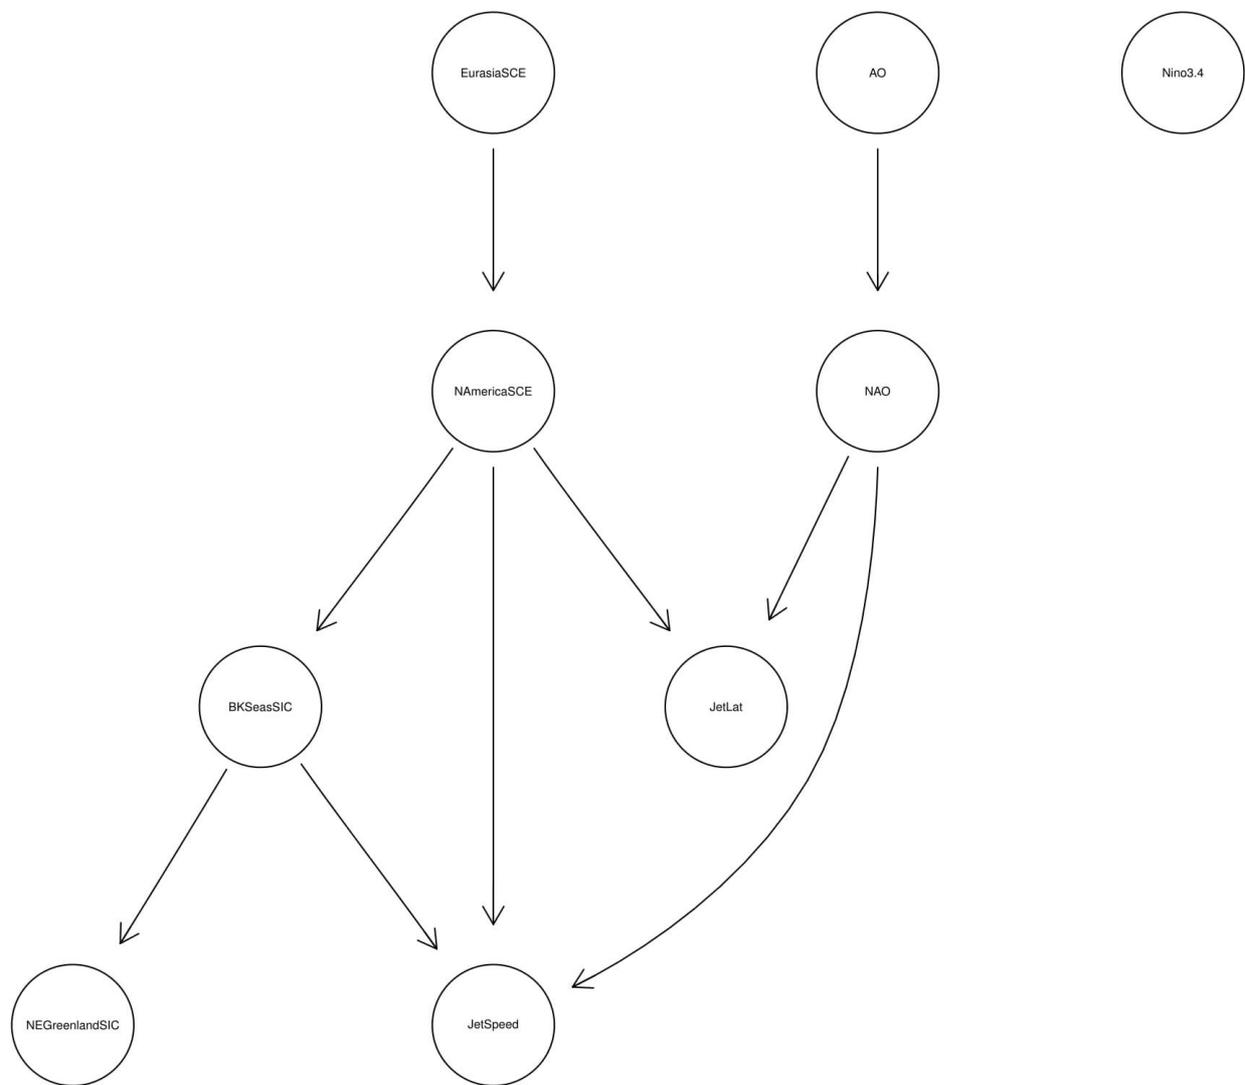


Figure 3.13 (2b) Monthly 35-Year DAG (2000 split) K=10.

Samarasinghe et al (2017) both focus on Arctic-midlatitude connections using a comparatively small number of variables. Subsequently, an important research gap exists in the need for complex datasets with a range of Arctic, midlatitude and tropical variables to determine important drivers of midlatitude circulation variability with graphical model approaches.

A clear theme which emerges from the position of the NE Greenland SIC variable in all BNs is the need for a more advanced approach to selection of cryospheric variables like sea ice regions; NE Greenland SIC was not found to influence North Atlantic atmospheric variability at any of the timescales. Hall et al (2017) used a correlation map approach, whereas this chapter bases variable selection (the data itself and spatial scales used) on the literature (Hall et al, 2017; Kretschmer et al, 2016; Luo et al, 2016). A further investigation of parameters is required in future studies; potential regions of sea ice impact, subdivision of snow cover areas, spatial extent of jet variables and the addition and removal of useful tropical and sub-Arctic data into networks will need to be explored to ensure results are robust.

An important relationship between N American SCE and jet latitude was found only in the DAGs where the training/testing datasets were split up to 2000 (Figs 3.11-3.13); the impact of snow cover on seasonal timescales may be central to this relationship, so the effective use of climatological anomalies to mask the seasonality of the data and reveal the signal is a necessary data preparation step for chapters 4 and 5. In contrast, Eurasian SCE appears not to have a significant impact on N Atlantic jet latitude or speed.

The split of training and testing datasets clearly does not change the network structures in any fundamental way, as both SCE and SIC links are preserved in the 2000 split. This chapter used dataset splits that were approximately 50/50 in terms of training and testing datasets, rather than a more standardised 75/25 one, to investigate the impact of including the beginning 'AA period' (mid-1990s to 2000) in the training dataset. The hypothesis that DAG structures will differ significantly when including the AA period has not been disproved by this analysis; instead, separate BNs created using data outside of (1981-1997) and within (1997-2016) the AA period would address this hypothesis more adequately. Chapter 5 uses dynamic networks and dataset splits to investigate the period of AA and differences in DAG structures more thoroughly.

In summary, the BNs presented here lack several mechanisms that are key to the application of graphical models to climate data analyses. The teleconnections between cryospheric variables (North American SCE and Barents-Kara Seas SIC) and jet latitude and speed found in section 3.4 are central to subseasonal-seasonal forecasting, as the identification of reliable predictors of jet stream variability is a priority for operational weather forecasters (Jung et al, 2014; Overland et al, 2015). However, a number of data preparation steps are missing from this analysis that fundamentally constrain the hypothesis testing and results drawn from it. Given the importance of time lags and seasonally variable linkages in

the study of climate teleconnections, static BNs do not provide the tools necessary to analyse Arctic-midlatitude weather linkages which are regional and intermittent in nature and therefore rely on advanced time series analysis like dynamic graphical models (Cohen et al, 2018). Lagged correlations and causal drivers were central to the methodologies and findings of Hall et al (2017) and Kretschmer et al (2016) respectively, prompting the investigations into appropriate lead times for potential drivers of jet stream variability in chapters 4 and 5. Further chapters aim to address the shortcomings in data and methodology through the introduction of a more advanced data preparation process and Dynamic Bayesian Network tools.

Chapter 4: Exploring Dynamic Bayesian

Networks with Hidden Variables to Detect

Arctic-Midlatitude Linkages

4.1 Introduction

The aim of the following chapter is to explain the application of BN methods in detail, with a particular focus on the application of BNs to time series analysis, to demonstrate how they might be used to investigate climate teleconnections. An attempt is made to show the workings of the approach; data preparation and analysis methods which yielded negative results are included, with findings emphasising DBN development and accuracy in light of these. Both positive and negative results are then used to inform a discussion focusing on the suitability of specific parts of the DBN method for the analysis of observational climate data.

To extend the analysis of Chapter 3 to time series analysis, dynamic networks (HMMs and DBNs) are constructed. Models constructed with the same datasets detailed in Chapter 3 performed poorly in terms of predictive accuracy; Fig. 4.1 shows observed values regressed against predicted for a HMM using the dataset from Chapter 3, with highly dispersed values suggesting that the model fit is inadequate. Subsequently, a number of improvements in data preparation and network construction aimed at investigating model fit are detailed in this chapter.

A simple 4-variable network is used to reveal dependency structures between Arctic warming, North Atlantic jet stream variability and the El Niño Southern Oscillation. Four networks are designed to investigate the impact of increased model complexity on predictive accuracy. Structure-learning and hidden variables (HVs) are incorporated into dynamic networks, with model performance assessed at each stage. DBNs improve on static BNs by allowing for conditional dependencies across time slices when lead times are used for variables, as well as feedback relationships which are of fundamental importance for atmospheric processes.

DBNs with 1 or 2 HVs are constructed, with the value of the HV set at each time point to maximise the model fit (in terms of log-likelihood) to the data. As the HVs encoded within the networks are discrete, the hidden state shifts associated with HV phases point to changes within the system interactions of the particular nodes the HV is connected to (Trifonova et al, 2019). For this reason, adding HVs and connecting them to different sets of variables can have an impact on how the DBN performs. Because HVs maximise the log-likelihood of the

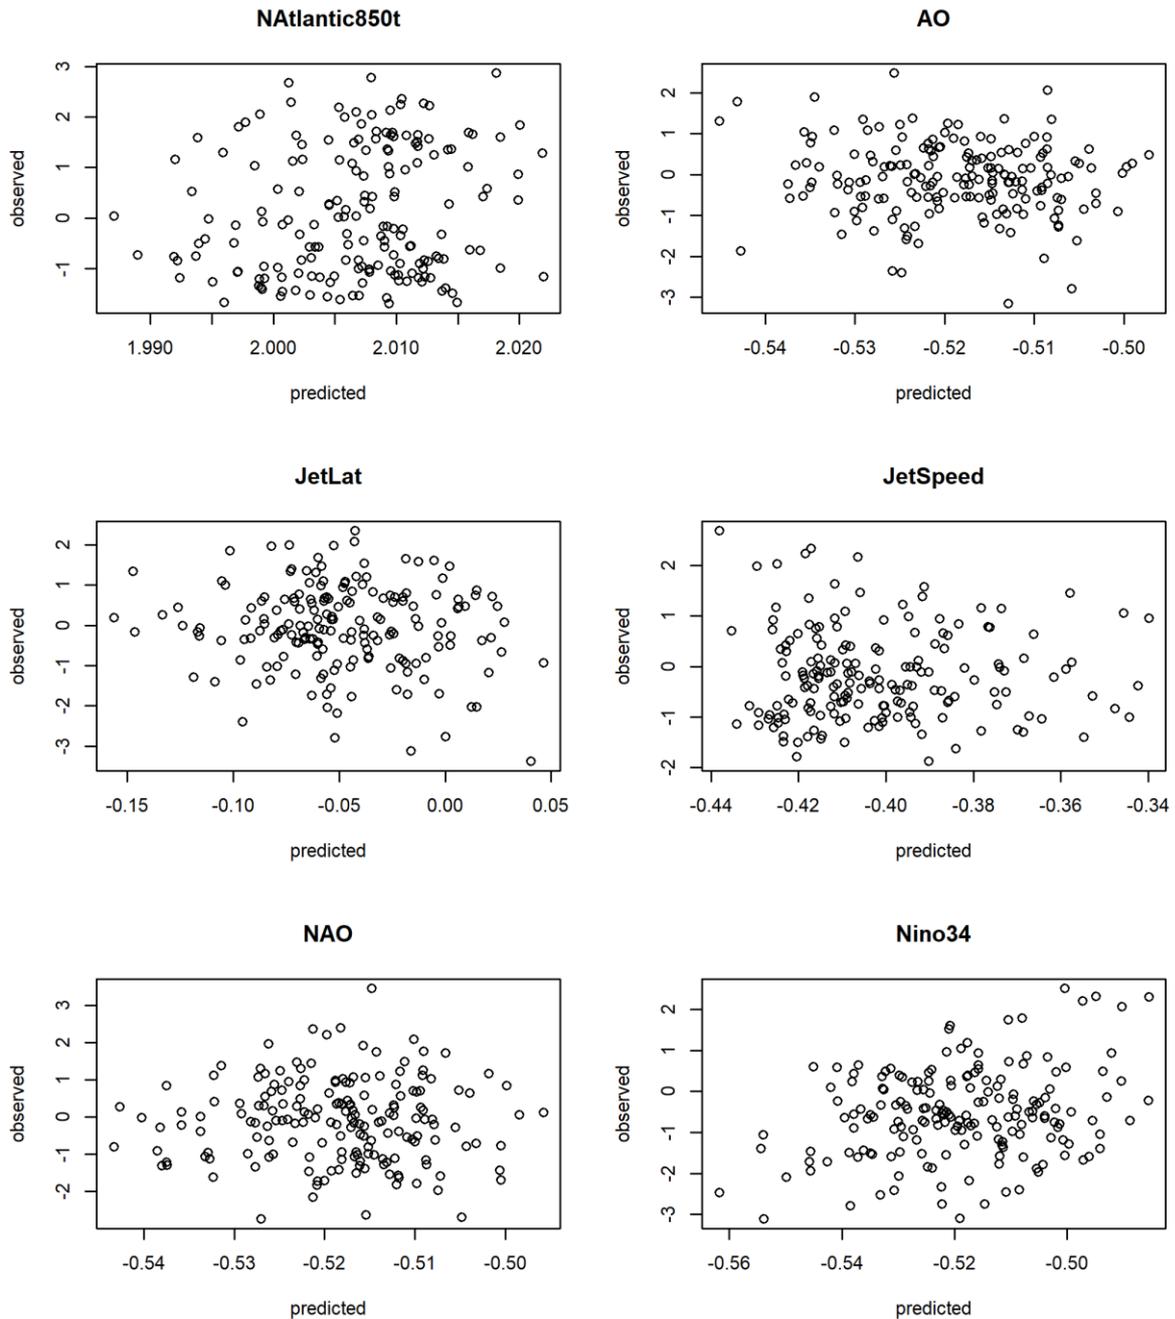


Figure 4.1 HMM predictive accuracy scatters showing predictions plotted against observed monthly data (experiment 2b in section 3.4.2) for all variables used in the analysis of Chapter 3. Model performance was low for all variables in the HMM, leading to a number of changes in data preparation and model construction detailed in this chapter.

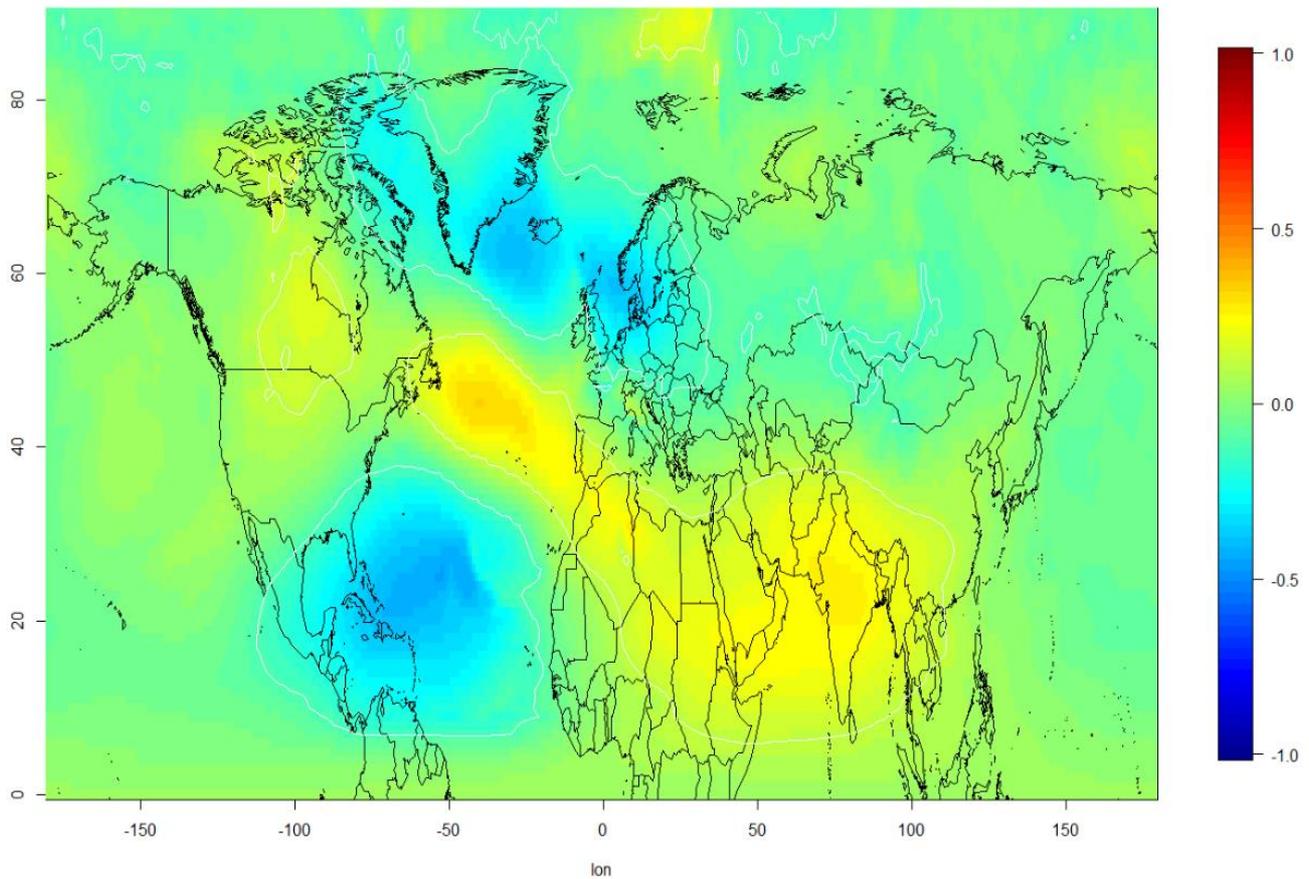


Figure 4.2 Correlation map of 850-hPa temperature to jet latitude over the Northern Hemisphere using 5-day averages for winter (DJF) months of 1981-2018 with no lead time. White contours denote significance at 0.01 (1%) level.

network, graphical model approaches making use of them have great potential for climate data applications. Arctic-midlatitude weather linkages are difficult to reproduce in statistical analyses of the observational record because they are dominated by nonlinear interactions of drivers, noisy dynamics in midlatitude processes, and time constraints as the AA signal only emerges sometime in the late 1990s (Overland et al, 2017). With this in mind, Chapter 4 represents an extensive analysis of time series DBN analysis, structure-learning, and HVs and the effect each has on model accuracy. A number of conclusions and recommendations are drawn from the experiments, many of which are used to structure the study design and analysis of Chapter 5.

4.2 Data

Four variables of importance to Northern European winter weather patterns are selected for this analysis, with a further (the MJO) included to measure in impact of adding unconnected nodes on model structure and accuracy. 5-day averages were used for all data to focus on submonthly variability and to maximise the number of data points as this chapter focuses on changes in accuracy between different models. To this end, a number of

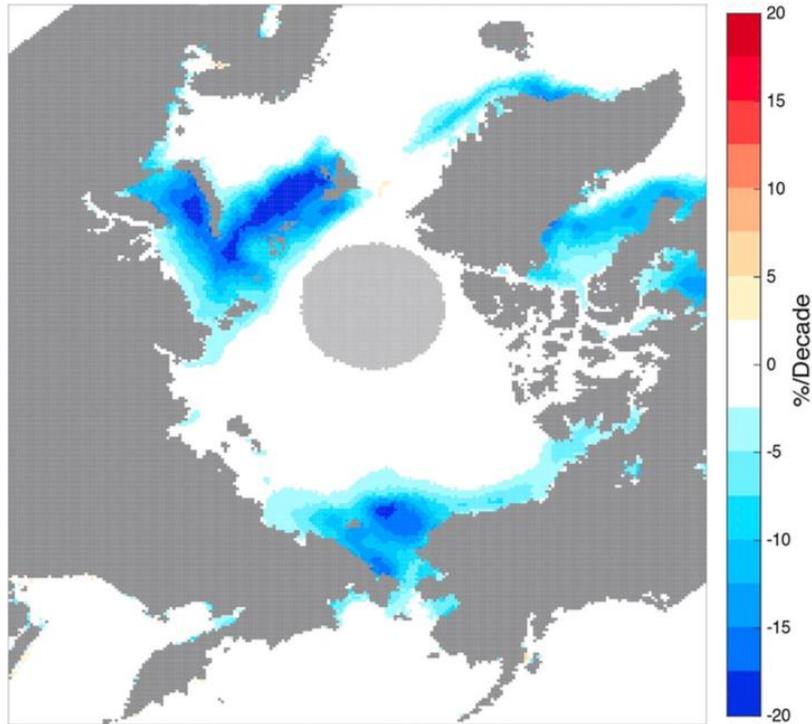


Figure 4.3 NSIDC sea ice concentration data for November of the years 1979-2016. A linear downward trend of 40% is shown for November, when autumn sea ice freeze-up takes place. Source: Overland and Wang (2018).

improvements on data preparation are made from chapter 3. The winter months (DJF) were selected for all variables, with the 29th February removed from each leap year allowing for a total of 90 days in each year and 666 timesteps in total.

Fig. 4.2 suggests the 850-hPa temperature over the south-eastern coast of Greenland is strongly negatively correlated to the latitude of the jet stream in winter months. In the sea ice record of the satellite era (1979-present), Baffin Bay and the east coast of Greenland have both undergone significant reductions in sea ice concentration (Fig. 4.3). The North Atlantic sector (defined as 280-350°E, 70-90°N) was therefore chosen to investigate the impact of regional AA on midlatitude circulation. Near-surface 850-hPa temperature from ECMWF's ERA-Interim global atmospheric reanalysis (Dee et al, 2011) was used to represent regional AA, available at <https://apps.ecmwf.int/datasets/data/interim-full-daily/>. As a first step, the area-weighted spatial average was taken over the region, resulting in a 1-dimensional time series dataset. Next, anomalies were calculated by subtracting each timestep of the multi-year mean (1981-2018) from the value of the matching timestep, and the result was then detrended.

This marks an improvement on chapter 3's dataset in the following ways: the use of 850-hPa temperature ensures that the full AA signal is captured, and the preparation steps are standard practice in climate science to remove unwanted variability on different

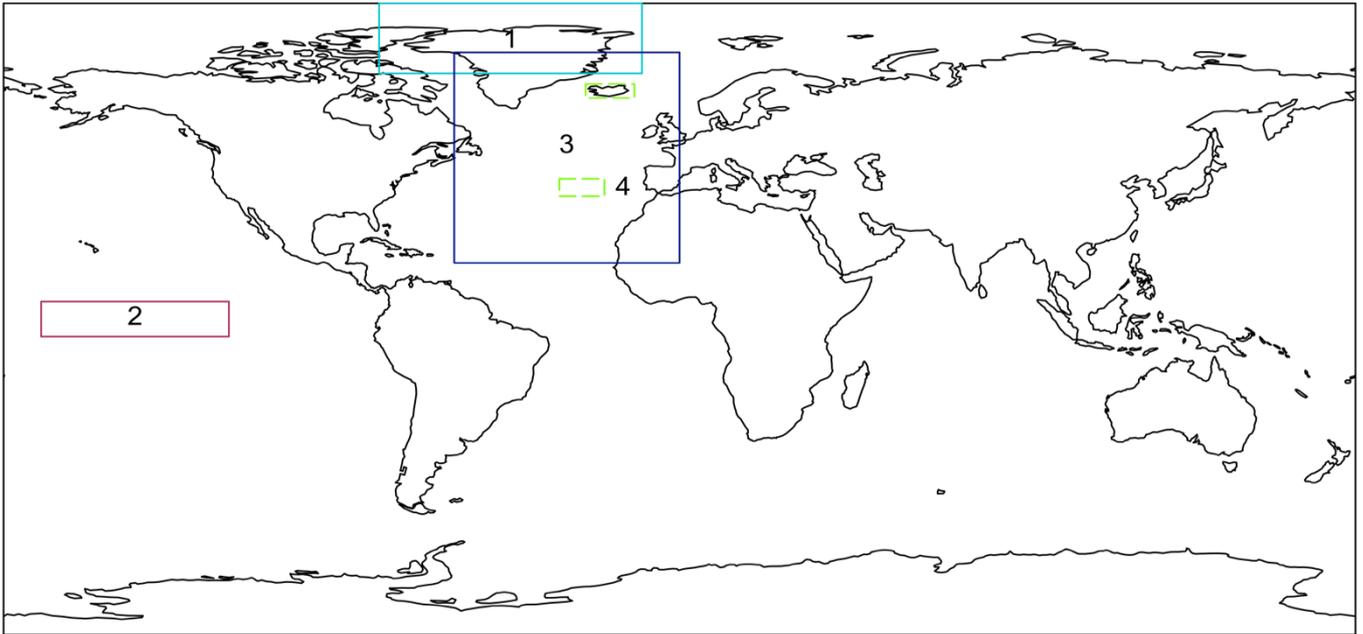


Figure 4.4 Map of variables used in experiments 1-3 below: 1 North Atlantic 850-hPa temperature, NAT; 2 El-Niño Southern Oscillation, ENSO; 3 jet latitude, jet; 4 North Atlantic Oscillation, NAO (dotted green).

timescales, i.e. the seasonal cycle and long-term trends. Firstly, the AA signal is the result of local sensible heat fluxes, evaporation and the transport of heat and moisture into the Arctic from lower latitudes (Cohen et al, 2018a). The downward trend in sea ice concentration and resulting albedo feedback are only partially responsible for AA, whereas near-surface temperature captures the full range of AA drivers and signal. Secondly, using raw data means that the results will be dominated by the influence of the seasonal cycle, so anomalies remove the seasonality allowing data to be expressed in terms of the distance between the value and its mean (Wilks, 2011), which in this case is the multi-year mean of the matching date (timestep). The removal of the trend is similarly central to climate analyses as it distinguishes between ‘climate change’ – i.e. the large upward trend in Arctic temperature over the North Atlantic region of the Arctic – and interannual variability (Wilks, 2011), the natural signal to be considered for climate analyses.

To investigate potential links between AA and midlatitude circulation within a DBN, a proxy for Polar Jet Stream variability was included. Jet latitude was calculated using ERA-Interim data (Dee et al, 2011) and defined over the region 16°–76°N, 0°–60°W. ERA-Interim zonal winds averaged over 900 to 700 hPa were filtered with a 10-day Lanczos low-pass filter using a 61-day window to remove both the synoptic scale variability and the higher-level Subtropical Jet (Woollings et al, 2010). Because the result is consistent across pressure levels, jet latitude is considered a robust way of investigating midlatitude circulation variability (e.g. Hall et al, 2017; Barnes and Simpson, 2017; Samarasinghe et al, 2018). The polar front

jet extends throughout the depth of the troposphere, thus sampling at these lower levels avoids confusion with the subtropical jet which is more vertically restricted.

Two indices known to have an impact on the Northern European climate during the winter were included in the networks. The El Niño Southern Oscillation (ENSO) 3.4 index, calculated with HadISST1 sea surface temperature data (Rayner et al, 2003), was included to represent tropical Pacific variability, retrieved from <https://climexp.knmi.nl/>. Large-scale planetary waves are triggered through intense convection in the tropical Pacific that can influence midlatitude circulation; the ENSO signal is indeed used for seasonal prediction of Atlantic and European weather when ENSO is active (Trenberth et al, 1999; Scaife et al, 2017). The North Atlantic Oscillation (NAO), a metric derived from the surface SLP difference between the Subtropical High and the Subpolar Low, is the primary determinant of Northern European wintertime variability. For the purposes of this analysis, the NAO can be thought of as another proxy of midlatitude circulation, given that it delineates the position of the storm track and the Polar Jet Stream and describes their variability through phase shifts (Hall et al, 2015). The principal component (PC)-based NAO index was obtained from <https://climexp.knmi.nl/>.

Finally, the outgoing longwave radiation (OLR)-based MJO index (OMI) was used as a proxy for tropical convection associated with the MJO's variability, retrieved from <https://www.esrl.noaa.gov/psd/mjo/mjoindex/>. The OMI is created by projecting 20-96 day band-pass filtered OLR onto the two leading EOFs of the 30-96 day eastward filtered OLR (Kiladis et al, 2014). The MJO signal has been linked to NAO forecasting (Lin and Brunet, 2011) and variability (Jiang et al, 2017), and was included in the networks to reveal potential linkages to North Atlantic atmospheric circulation and to investigate the robustness of learned network structures. The OMI is a convection-based index, making it a suitable candidate for identifying teleconnections that originate in tropical convective heating associated with MJO phases (Kiladis et al, 2014; Tseng et al, 2018).

The variables were standardised as a final data preparation step, so that each had a mean of 0 and standard deviation of 1. This ensures the data has equal means and similar ranges for the parameter learning stage of the DBN method, and allows for ease of visualisation of the results.

4.2.1 Data Development

The Arctic Oscillation (AO) was found to frequently be a parent node of midlatitude metrics in Chapter 3, and static BNs showed that the AO-NAO relationship in particular was robust to different structure-learning algorithms and scoring strengths. This is expected, given that the two indices essentially describe a north-south orientation in the dipolar structure of the North Atlantic pressure field, without the Pacific centre of action in the case of the NAO

(Hurrell, 1995; Ambaum et al, 2001). Both were initially included as a proxy definition of North Atlantic storm track variability, to capture the phase shifts of the jet which allow cold Arctic air to spill out into the midlatitudes. The AO was removed from this analysis for two reasons; firstly, the two indices are too similar in the sense that the inclusion of the AO adds nothing of value to the network because of the variability encapsulated by the NAO index, and could in fact capture less variability due to it having a zonal structure with the symmetric variations of an 'annular mode' (Thompson and Wallace, 2000). Secondly, the NAO can be ubiquitously identified in the patterns of local teleconnection and regional EOF analyses, suggesting that the NAO offers a more robust description of Northern Hemisphere variability than the AO (Ambaum et al, 2001), particularly for the North Atlantic region of interest in the present study.

Two Snow Cover Extent (SCE) datasets used in Chapter 3 were initially included in the analysis to investigate the cryospheric influence from land as well as the Arctic Ocean. Eurasian SCE (55-150°E, 40-80°N) and North American SCE (130-70°W, 40-70°N) were processed in the same way as other spatial data detailed above: the area-weighted average of each region was taken; the multi-year mean (1981-2018) was calculated for each timestep and subtracted from the matching date of each timeseries value; and the result was detrended. This method is regarded as an appropriate way of preparing spatial data for timeseries analysis in graphical model studies (e.g. Kretschmer et al, 2016; 2017). Despite this, the preparation steps created timeseries values that were almost always the same as the multi-year mean subtracted from them, meaning that the anomalies were predominantly 0 (or close to 0, with variability in the hundredths or thousandths in decimal places: Fig. 4.5). As with the 850-hPa temperature variable, the data generated was not sensitive to a different ordering of the preparation steps (i.e. spatial averaging and anomalising). The problem is therefore likely caused by the area used for spatial averaging being too large (see Figs. 3.1 and 3.2) for a single SCE value, despite the fact that the same or similar regions were considered appropriate in other studies using different time resolutions (Kretschmer et al, 2016; Hall et al, 2017).

DBNs were run with both SCE variables included in them but returned no learned connections to any other data (not shown). The lack of SCE links may be due to state-dependent or longer-duration relationships that are not showing up as a result of the lead times discussed in 4.3. A link between snow cover anomalies, particularly in October, and a negative AO response in the following winter has been found in observational studies (Takaya and Nakamura, 2008; Honda et al., 2009). Positive snow anomalies in October have been hypothesised to increase upward planetary wave activity, leading to the disruption of the stratospheric polar vortex and a resulting negative AO (NAO) pattern and equatorward shift of

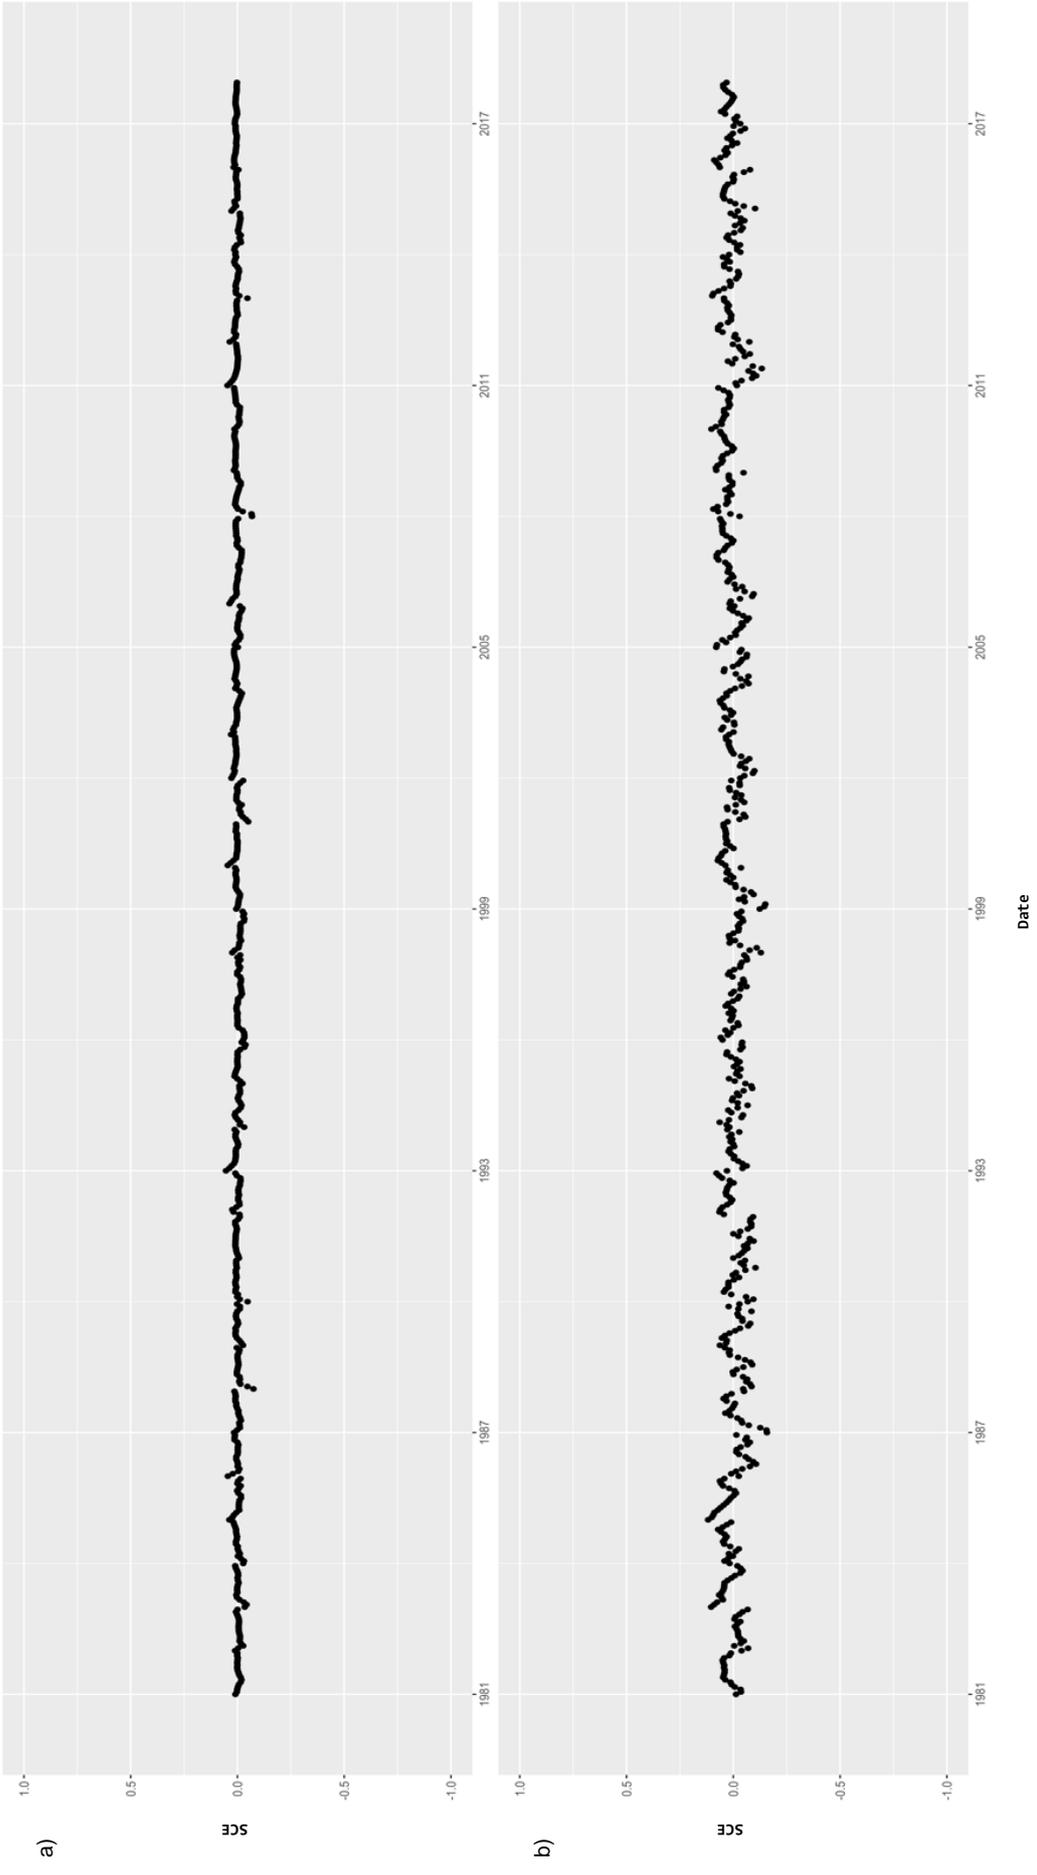


Figure 4.5 Time series plots of 5-day averaged snow cover anomalies spatially averaged over a) Eurasia and b) North America for the years 1981-2018.

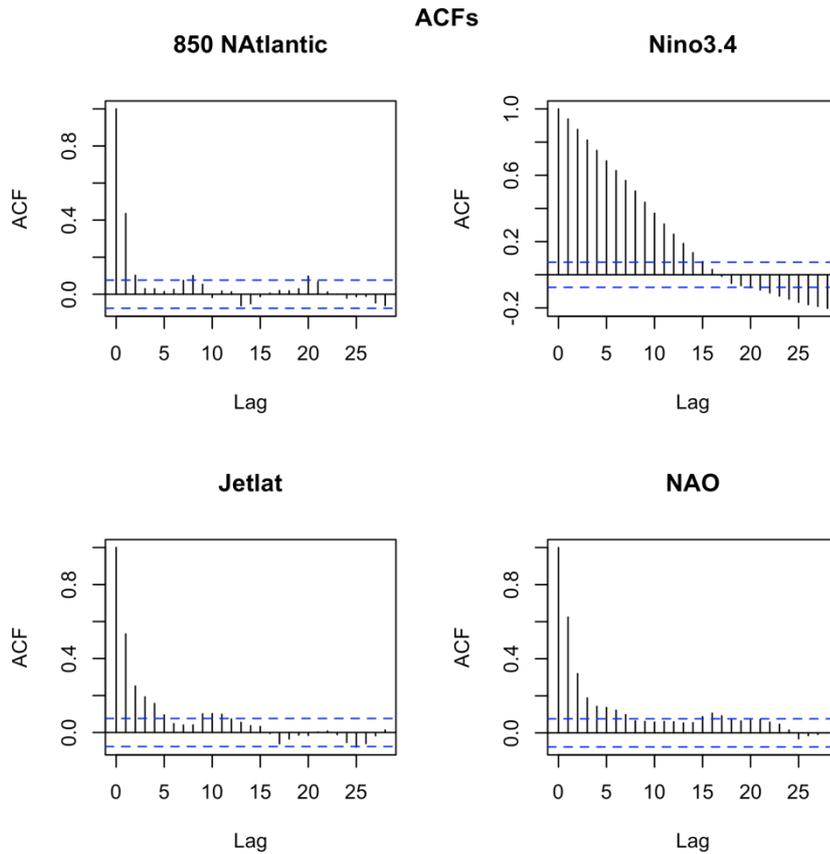


Figure 4.6 Autocorrelation Function (ACF) plots for variables in experiments 1-3, detailed below. Timesteps are 5-day averaged data.

the Atlantic jet stream in winter (Hall et al, 2015; Smith et al, 2016). Furthermore, this link has been shown to be non-stationary in a correlation analysis (Peings et al, 2013). Both of these factors could have effectively prevented links from being identified in networks; the following analyses are run without the SCE variables due to the unsuitability of the prepared SCE data and lack of connectivity across models.

4.3 Methods

A series of networks of increasing complexity were constructed, with BN techniques such as structure learning and hidden variables added at each stage. The impact of adding these approaches to networks is not found to increase the predictive accuracy for models using a small dataset of 4 variables; possible reasons for this are discussed below. The impact of adding unconnected nodes on network structure and accuracy is also investigated. All models were constructed using the Bayes Net Toolbox (BNT) for MATLAB (Murphy, 2001b), with data preparation and plotting carried out in R

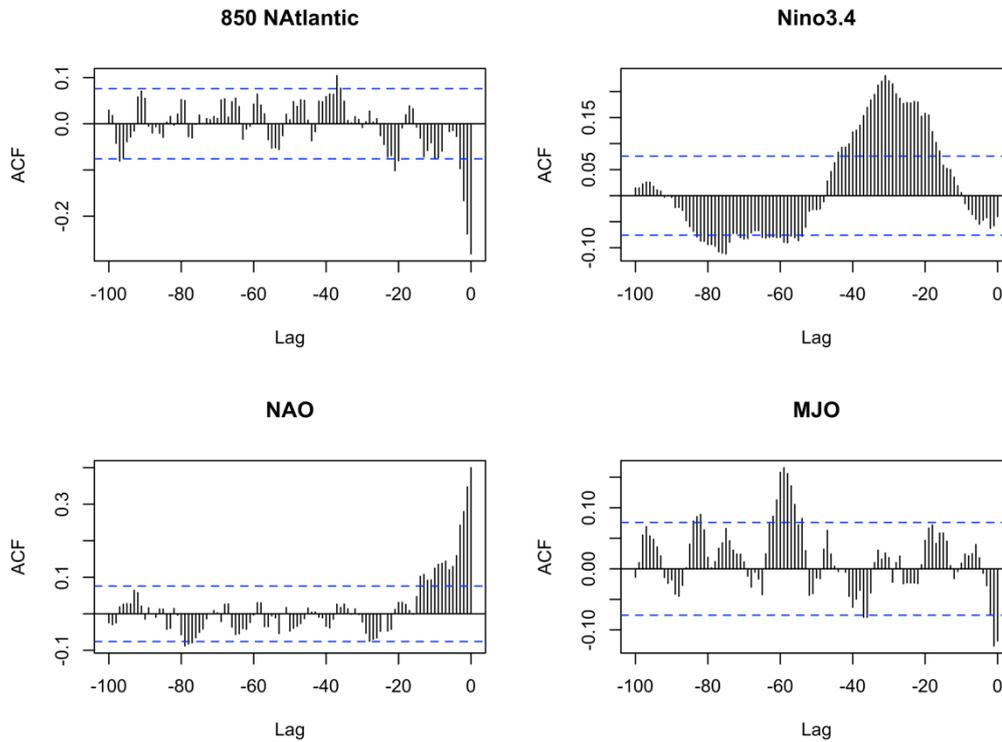


Figure 4.7 Cross-correlation Function (CCF) of all variables plotted against jet latitude. Timesteps are 5-day averaged data.

As a first step, auto-correlation functions (ACFs) were plotted showing that all variables were strongly correlated to values lagged at a single timestep (Fig. 4.6) with 5-day timesteps. For this reason, autoregressive links are coded into the ‘inter’ matrix which dictates how the nodes are connected between time slices. Cross-correlation functions (CCFs) were plotted for each variable against jet latitude (Fig. 4.7) to determine whether lagged correlations existed which may result in conditional dependencies across time slices. These were interpreted carefully and in combination with previous studies to prevent spurious correlations over physically implausible time lags from being used for analysis. In the case of the MJO, which demonstrates a sinusoidal correlation signal, this includes a lead time of -59 timesteps or 295 days.

The ENSO CCF shows a clear, significant positive correlation which peaks at -31 timesteps suggesting late June ENSO variability is most strongly linked in the case of early December jet latitude variability. This lead time is lengthier than most studies investigating ENSO impacts on European weather, where a common lead time is around 2-3 weeks (Baldwin and Dunkerton, 2001; Scaife et al, 2017), but has been found by others investigating jet latitude specifically (Hall et al, 2019). 3 lead times were initially tried for the MJO variable

(-59, -36 and -1 based on the CCF plot, Fig. 4.7), with none connected to the networks returned by the structure-learning stage. The MJO variable was left unlagged to instead investigate whether adding unconnected nodes to the network would change the structures returned. For a full investigation into different lead times for tropical variables, see section 5.6.

4.3.1 Experiments

After being prepared with the steps listed above, the data was loaded into BNT and split 80/20 into 'train' and 'test' datasets in order to build the DBN code based on the training dataset to be tested on the remainder of the data; the reason for the change from a 50/50 split in Chapter 3 is discussed in section 4.3.2 below. Four types of networks are constructed:

- 1) HMM – a simple 4-variable autoregressive HMM with a single HV linked to all nodes and an imposed structure; the HMM was not run through the structure-learning stage.
- 2) DBN1HV – a 4-variable autoregressive DBN with a single HV linked to all nodes to investigate the impact of learning the structure and a HV from the data on predictive accuracy.
- 3) DBN2HV – a 4-variable autoregressive DBN with two HVs: the 1st linked to the jet, NAO and NAt variables, the 2nd linked to all nodes. A second HV is learned from the data to explore its impact on accuracy in small networks.
- 4) DBNMJO – a 5-variable autoregressive DBN with a single HV linked to all nodes to analyse how robust the network structures returned in (2) and (3) are to the addition of a second unconnected node (the MJO and ENSO).

For experiments 2-4, the first step was to learn the dependency relationships from the data using a structure-learning algorithm. In the case of experiment 1, a simple structure was coded into the model as 'control' run which allowed for comparison between the predictive accuracy of BNs with the structure learned from the data (2-4), and those without (1). The DBN structures used the PC algorithm (Spirtes and Glymour, 1991) and an alpha value of 0.01 to ensure that only significant relationships were identified. The structure-learning phase was guided with the '*fisher z*' test for conditional independence. HVs with autoregressive links were coded into DBNs to address the hypothesis that model accuracy would improve with additional parameters in a simple DBN structure. HVs allow for structures that may be significantly more similar to the complex nonlinearities of the climate system that we are trying

to model (Tucker and Liu, 2004), as HVs may act to reduce noise in highly variable child nodes by ‘explaining away’ extreme outliers that lie well outside the SD of the inferred hidden states.

Because DBN structures learned using the PC algorithm return pDAGs, or ‘Partially Directed Acyclic Graphs’, bidirectional arcs had to be removed. DBNs required fully directed graphs, DAGs, to run. The jet latitude variable was consistently found to be a child node with multiple parents, so edges pointing towards it were preferentially retained to preserve the ‘explain away’ effect and create a usable DAG. DBN code cannot run with undirected arcs as DBNs require acyclic graphs (Scutari and Denis, 2015), so the removal of important collider node structures would undermine the individual probability distributions that make up a BN.

Different HV structures were experimented with for the DBN2HV, some of which did not work during the testing phase (not shown), but the learned structure remained unchanged regardless of which nodes each HV was connected to. Experiment 3 has 1 HV connected to the jet, NAO and NAt nodes and a 2nd connected to all to represent a set of HV states with and without the tropical influence (ENSO). A 2nd HV was initially added due to the DBN runs having relatively low predictive accuracy, and the results are shown below. Next, the parameters and HV states were inferred from the data using the EM algorithm, and the resulting DBN was then tested on the data not included in the training stage (the test dataset). Model validation was achieved by comparison of predicted values to the observed (test) dataset, and by calculating the sum of squared error (SSE) for each variable across experiment runs according to:

$$SSE = \sum(\text{predicted} - \text{observed})^2$$

4.3.2 Model Development

A number of steps were taken to develop the methodology and improve the practicality of the models including the shift from HMMs to DBNs, i.e. from an imposed structure to a learned DAG, and the full preparation of data for a more robust set of results. In Chapter 3, static BN performance was assessed with a 50/50 split in training and testing datasets, whereas a more standard 80/20 split is used to build and check the networks shown here (Chen and Pollino, 2012). The motivation for this is simply to reduce overfitting; dividing the dataset into two parts prevents the overfitting of the model to the full dataset. Fitting to the training dataset and then validating it allows for an unbiased estimate of the generalisation error, as the two datasets have as similar distribution as possible and are independent of each other (Shalizi, 2013). This is a well-used approach because no external validation set needs to be found which is identical to the input dataset, and because it ensures that the test data are not part of the network construction.

To have non-autoregressive models, the ‘inter’ relationships between nodes at different timesteps would have to be learned from the data as with the ‘intra’ ones learned

using the PC and HC algorithms. All HMMs and DBNs presented in Chapters 4 and 5 are autoregressive firstly because the ACFs justify this choice of structure for all variables (Fig. 4.6), and secondly due to method limitations. Only one structure-learning algorithm exists for DBNs in BN toolbox for MATLAB (Murphy, 2001b): the REVEAL algorithm developed by Liang et al (1998). REVEAL requires all nodes to be discrete and observed, meaning that all input data would have to be discretised into bins. Discretisation is not a practice applied to observed climate data; discretisation of continuous data can add imprecision to a model by removing variability within an interval for example, and the interval choice itself requires subjective supervision to ensure the bins fit the objectives of the experiment (Chen and Pollino, 2012). Because of these drawbacks, full structure learning in a DBN including the inter-slice connections was considered unsuitable for application to climate data, and autoregressive models with learned intra structures and fixed inter connections were used instead.

4.4 Results

The DAG structure results show either the hard-coded architecture in the case of the HMM, or the graph returned by the PC algorithm for the DBNs. In these plots, grey nodes denote hidden variables and blue the observed variables. In the DBNs, dotted lines represent hard-coded edges (HV links) and solid ones show those returned by the structure-learning stage (non-HV links). Edges represent conditionally dependent relationships between variables rather than explicitly showing causal drivers (Scutari and Denis, 2015). The DBN DAGs show only the intra-slice structure rather than a 2 time-slice DBN because they are autoregressive models with user-defined inter-slice connections where only the intra topology was learned.

The structure-learning phase initially returned a pDAG due to the use of the PC algorithm, with undirected edges between the NAT and jet nodes and between the jet and NAO. As DBNs require fully directed edges however, detailed above, cyclical links were removed in the direction of jet-NAT and jet-NAO to preserve the collider structures that make the jet node central to the network's probability distributions. The undirected edges returned by the PC algorithm may be capturing a positive feedback between the Arctic and midlatitudes, reflected in other machine learning approaches for Eurasian (Kretschmer et al, 2016) and North Pacific (Samarasinghe et al, 2019) circulation anomalies. In a study which used an approach similar to that employed in this chapter, Samarasinghe et al (2019) find a positive feedback loop between 850-hPa temperature averaged over the entire Arctic region and jet speed and latitude over the North Pacific basin using the PC algorithm. They suggest that a southward shift of the jet in response to Arctic warming is responsible for this relationship, as jet displacement results in meridional atmospheric circulation which can in turn warm the Arctic. Recent studies have emphasised the midlatitude-Arctic link as an important driver of

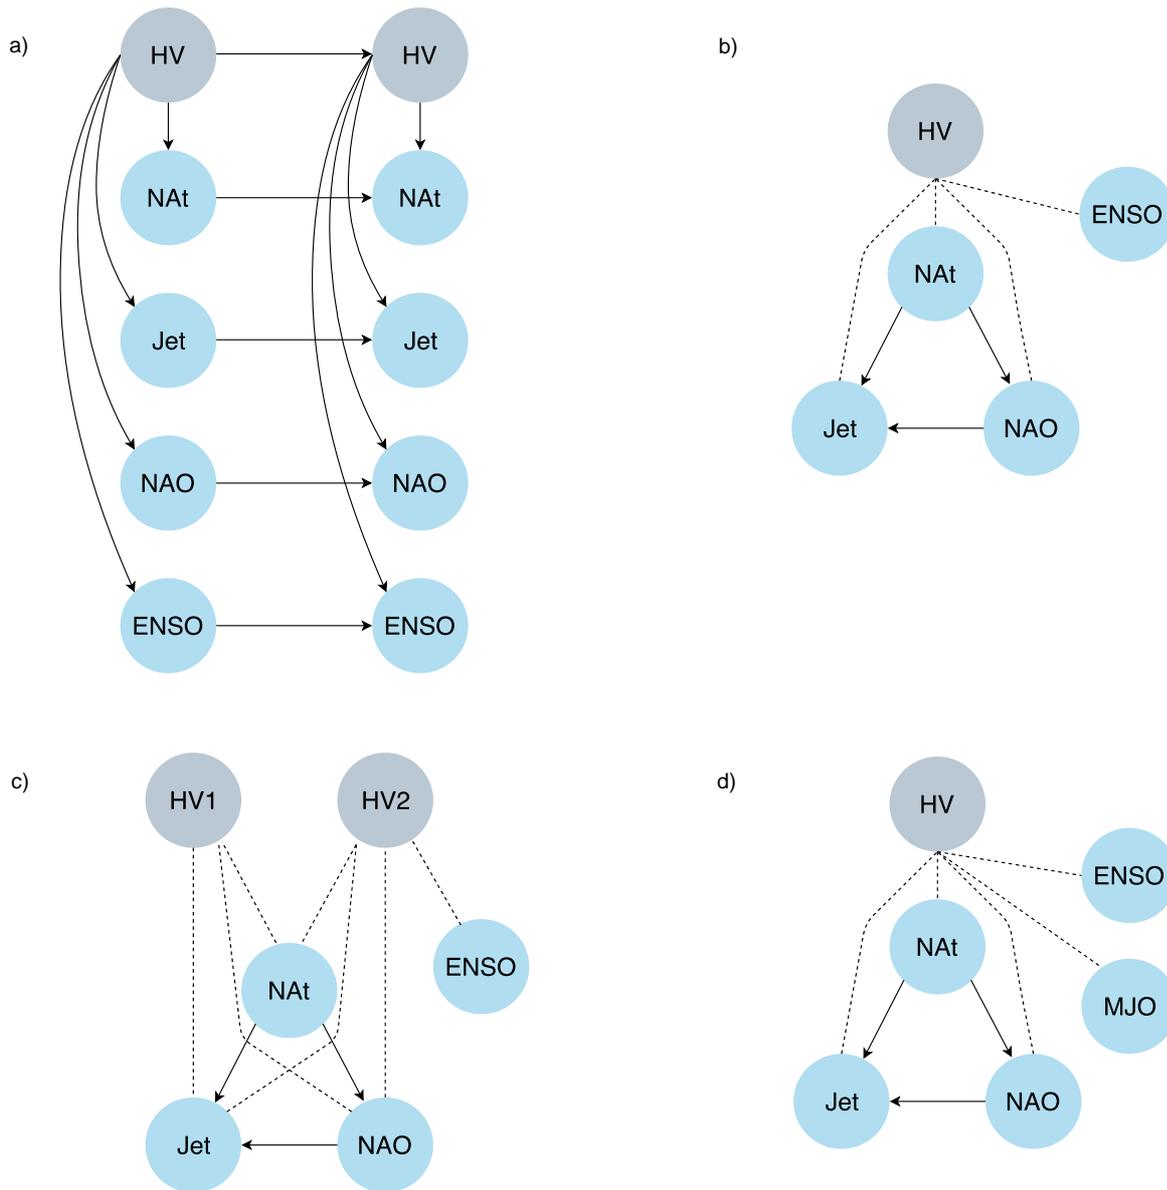


Figure 4.8 The DAG structures for all model types: a) a hidden Markov model (HMM); b) Dynamic Bayesian Network with one HV (DBN1HV); c) DBN with two HVs (DBN2HV); d) Dynamic Bayesian Network with one HV with the addition of the MJO node (DBNMJO).

AA processes; advection of warm air and moisture from the midlatitudes significantly contributes to AA and seems to be especially pronounced through the North Atlantic pathway (Kim et al, 2017; Yang and Magnúsdóttir, 2017). Atmosphere-only models also point to midlatitude circulation trends as a strong influencer of Arctic warming during winter, particularly at non-surface (low-middle) layers of the atmosphere, using relaxation experiments (Ye and Jung, 2019).

No link between either of the tropical variables and NAt, the jet and NAO were found in the networks using the structure-learning process. Despite this, the results of the DBNMJO

experiment are included (Fig. 4.8d) to demonstrate how robust the PC algorithm is to the inclusion of unconnected data. Other graphical model approaches like Causal Effect Networks (CENs) can return different graphs based on the choice of settings during network construction (Kretschmer et al, 2016) or on the inclusion of new variables. In contrast, the DBNs produced DAGs that were consistent despite the introduction of the MJO; the graph effectively had two 'unrelated' variables, and yet the dependency relationships between the NAt, jet and NAO nodes were identical despite the addition of HVs and variables. Fig. 4.8 clearly recommends DBNs as a robust tool to learn relationships from data for climate science applications that are consistent regardless of the number of unconnected nodes. Chapter 5's results suggest that the main difference between the small 4-variable DBNs shown here and those using larger datasets of 5 or more variables is the predictive accuracy of the DBN itself rather than the structure returned by the PC algorithm.

4.4.1 Model Comparison

The HMM seems to predict all variables with the most accuracy and SSE generally increases for variables with model complexity (Table 4.1). The DBN1HV offers similar predictive accuracy and even outperforms the HMM in the case of the NAt and NAO nodes, but is significantly less accurate at predicting the variability of the jet which is central to the analysis. This means that the addition of structure and hidden variables actually decreased the predictive capability of the models, seen most obviously for the jet variable. There are two possible explanations for this. The structure learning process may have returned a graph structure that decreased the accuracy of the model because the dependency relationships contained within it did not recreate those seen in the climate system. The jet variable, for example, loses accuracy when connected to the NAt node in the DBNs shown here, the opposite of what one would expect if jet variability was dependent on AA processes occurring in the North Atlantic region. The Jet and NAt are indeed the worse predicted variables in all 3 DBNs.

The SSE results suggest that the predictive accuracy of the DBNs may have been influenced by overfitting (Table 4.1; also see appendix Table A4.1 for a comparison of train and test dataset SSE). In this case, pursuing a more complex model has resulted in a loss of accuracy when a simpler model (the HMM) better predicted a new sample (the test dataset). Rather than an error in the structure used to estimate the conditional probabilities and parameterise the model, which would have more of an impact with BNs with a greater number of variables, the problem may rest with the high number of edges relative to the small number of variables included. Apart from the DBNMJO, all models have 4 input variables and the number of edges increases significantly with model complexity: 4 for the HMM, 7 for the DBN1HV and 10 for the DBN2HV. 10 edges may simply be too many in a DBN with only 4

observed variables; the ENSO variable is significantly better predicted than the rest of DBN2HV possibly as it has fewer parents.

Chen and Pollino (2012) point out that ‘cases’, or events in which nodes change according to the variability of the parent node, vary in observed datasets even when the same combination of parent states occur multiple times. This is because the climate system is the result of complex, nonlinear processes governed by stochastic forcing (Franzke and O’Kane, 2017), making any meaningful relationships hard to capture when using long timeseries. Accordingly, the accuracy of the BN’s conditional probabilities can increase as a result of a larger number of cases (Chen and Pollino, 2012); Cain (2001) recommends a minimum of 20 cases for each combination of parent node states to prevent overfitting.

The DBN analysis in Chapter 5, conducted on a 7-variable dataset, suggests that the dependency relationships are not the reason for the drop in predictive accuracy shown in Table 4.1 because an almost identical pDAG structure is returned between the NAt, jet and NAO nodes (see Fig. 5.3). As the only difference between the pDAGs of Fig. 4.8b and Chapter 5’s 5-day jet DBN is the NAO-jet edge seen in Fig. 4.8b, the overfitting of edges seems like the only logical explanation for the under-performance of the DBNs shown here. Overfitting was indeed a problem for the DBNs run with 2 HVs which are detailed briefly in Chapter 5, whereby the test code would not even function to allow for model assessment. No loss of predictive accuracy occurs among the equivalent DBNs of chapter 5 with a single HV.

	HMM	DBN 1HV	DBN 2HV	DBN 1HV MJO
850 NAtlantic	1708.86	1474.08	1019.55	1487.97
ENSO	26.77	27.41	0.58	27.57
Jetlat	59.8	185.87	391.29	236.96
NAO	60.94	55.1	32.35	56.27
MJO	NA	NA	NA	104.48

Table 4.1 Sum of squared error (SSE) of all climate variable predictions generated by the HMM, DBN1HV, DBN2HV and DBNMJO.

Clearly, the relatively small number of nodes included in this analysis prevents networks of increasing complexity from improving their accuracy; in other words, there were too few nodes for the number of edges found in the structure-learning stage and overfitting occurred. In another study with a similar methodological approach, Trifonova et al (2019) found that model performance scaled with network complexity. In Trifonova's study, a DBN with 2 HVs outperformed all other network architectures; the application of DBNs to the field of ecosystem modelling however meant that 28 observed variables were required for the analysis.

4.5 Discussion and Conclusion

In this chapter, networks were extended to DBNs to investigate the use of time series BN analysis, structure-learning and hidden variable models on climate data. A simple 4-variable network was used to demonstrate a positive feedback mechanism between Arctic warming and North Atlantic jet stream variability. Undirected edges from both the jet and NAO nodes to the NAt may be indicative of a positive feedback relationship which points to the midlatitude-Arctic pathway as being of importance. Currently, the transport of heat and moisture into the Arctic from midlatitude and even tropical sources is understated in the body of research on Arctic-midlatitude weather linkages (Cohen et al, 2018a). This may be because traditional methods used to research it, such as correlation and trend analyses (e.g. Francis and Vavrus, 2012; Tang et al, 2013) and model studies studying midlatitude responses to sea ice forcing (e.g. Mori et al, 2014), are unable to capture feedback mechanisms. The importance of remote heat transport into the Arctic relative to other factors like local heat fluxes and evaporation has not been definitively quantified but is understood to be a significant contributor to AA (Cohen et al, 2018a). Graphical model techniques seem to identify this feedback process at three major longitudinal regions of interest; the North Atlantic sector of the Arctic, as shown here, the North Pacific (Samarasinghe et al, 2019) and Eurasia (Kretschmer et al, 2016).

Secondly, data preparation steps were improved significantly from Chapter 3, as lead times were investigated with CCFs and added for ENSO, with appropriate steps taken to prepare spatial data for analysis with graphical models. Calculating anomalies to extract the seasonal cycle and isolate the signal rendered the SCE data unusable at a daily resolution for both the North American and Eurasian datasets. The spatial extent of the areas used to average both SCE datasets over is very likely to be the cause of this; a large amount of the spatial variability of snow cover is lost when averaged over such regions, particularly obvious for the larger Eurasian SCE region (Fig. 3.1). Feature selection algorithms may prove to be a useful tool to select regions of SCE which have an impact on the response variable, in this case the jet, without biases. The Response-Guided Causal Precursor Detection (RGCPD)

scheme, which identifies regions of spatial data that precede changes in the response variable, has been used to identify predictors of the stratospheric polar vortex at sub-seasonal lead times (Kretschmer et al, 2017), and features that provide predictability of Indian summer monsoon rainfall (Di Capua et al, 2019). Smaller sub-regions of SCE variability might be detected as drivers using a causal precursor detection algorithm, which would allow for meaningful anomalies to be calculated and the impact of SCE on midlatitude variability to be investigated in DBN or CEN models.

Contrary to expectation, an increase in model complexity resulted in a decrease in accuracy as the HMM outperformed all subsequent DBNs. The overfitting of parameters in DBNs with only 4 variables led to a loss of predictive accuracy, despite the addition of structure-learning and HVs intended to improve it. The dependency relationships shown in Fig. 4.8 are likely to be valid, as Chapter 5 suggests, so an increase in SSE with model development points to over-parameterised models as the cause, most obvious in the results of the DBN2HV. The dependency relationships returned by the PC algorithm were found to be robust across all DBN architectures and did not change when 2 unconnected nodes were included in the network (i.e. the DBNMJO).

Robust results in terms of structure for Arctic and midlatitude variables and an obvious bias in overfitting suggests that the application of DBNs to climate teleconnection analyses is worth further investigation. Chapter 5 expands on this analysis by adding a number of variables thought to have an impact on winter midlatitude circulation, by averaging the data over a number of time resolutions to analyse submonthly drivers, and through the addition of an extensive assessment of DBN robustness, HVs and data selection. In doing so, Chapter 5 also details a vast improvement in model performance when compared to the results of the simple 4-variable DBNs presented here.

Chapter 5: Using Dynamic Bayesian Networks

to Investigate the Relative Impact of Arctic

Amplification on Midlatitude Circulation in the

Euro-Atlantic Basin

5.1 Introduction

The Arctic has warmed at more than twice the speed of the global average since the mid-20th century, and at more than six times the average between 1998 and 2012 (Huang et al, 2017). This process, known as ‘Arctic Amplification’ (AA), is particularly strong in boreal winter. Concurrently, the heavily populated regions of western Europe and the eastern coast of the US have experienced several cold outbreaks during the winters of recent years. A multitude of studies are supportive of a link between AA and midlatitude circulation (e.g. Samarasinghe et al, 2019; Blackport and Screen, 2019), cold air outbreaks (e.g. Kim et al, 2014; Chen and Luo, 2017) and the North Atlantic Oscillation (NAO) (e.g. Pedersen et al, 2016). The strength of the connection between AA and midlatitude circulation remains uncertain; the importance of the linkage relative to other factors such as internal midlatitude variability, tropical forcing and the stratospheric polar vortex (Messori et al, 2018) currently represents a striking gap in our knowledge. In addition, the regional and intermittent nature of linkages means that direct effect attribution studies using Arctic processes like sea ice loss will not provide a way forward for the research area (Overland et al, 2016). Establishing the Arctic’s impact on jet stream variability, relative to these factors, is a complex but essential research endeavour as observational analyses often consider these factors in isolation.

Warming of tropical oceans, in particular above-average sea surface temperatures (SSTs) in the Pacific, are known to impact midlatitude flow through intense convection and latent heat release which generates planetary-scale Rossby waves (Trenberth et al, 1998). The El Niño-Southern Oscillation (ENSO) (Scaife et al, 2017a) and other tropical Rossby wave source regions (Scaife et al, 2017b) provide predictive skill in seasonal midlatitude circulation forecasting, and ENSO has a stronger role in winter. Arctic sea ice concentration has also been put forward as an important driver of meridional jet stream configurations (Francis and

Vavrus, 2015). Low sea ice concentrations expose more open water which absorbs additional heat, leading to a greater exchange of heat and moisture between the ocean and atmosphere in autumn and thus an anomalously warm Arctic. More recently, studies have emphasised the importance of considering warming over the Arctic as opposed to sea ice trends (Barnes and Simpson, 2017; Cohen et al, 2018a), chiefly because AA is the result of a complex combination of local sensible heat fluxes, evaporation and the remote transport of heat and moisture from lower latitudes (Cohen et al, 2018a). Anomalous midlatitude circulation drives intrusions of warm, moist air into the Arctic which play an important role in the feedback between Arctic warming and sea ice retreat (Rigor et al, 2002; Zhang et al, 2008), with transport thought to be particularly pronounced along the North Atlantic pathway due to Atlantic blocking deflecting midlatitude cyclones polewards (Kim et al, 2017; Yang and Magnusdottir, 2017). Intrusion events have measurable impacts on sea ice within several days of an event (Kapsch et al, 2016), in turn strengthening AA and any potential midlatitude feedback.

Sea ice loss in the Barents-Kara Seas, a region of pronounced variability, can expand and intensify the Siberian High through the initiation of vertically propagating Rossby waves and the disruption of the stratospheric polar vortex (Kim et al, 2014; Kretschmer et al, 2016). The southward flow of Arctic air that results from this has been associated with intensified cold events over East Asia (Overland et al, 2015). Over North America meanwhile, AA processes increase the likelihood of Alaskan and Greenland blocking events which can reinforce and prolong cold events (Chen and Luo, 2017; Overland and Wang, 2018). It is becoming increasingly clear that AA has an impact on the strength and position of the North Atlantic eddy-driven jet in winter (Barnes and Simpson, 2017; Blackport and Screen, 2019), although the existence of numerous potential drivers of the jet itself (Hall et al, 2015; Smith et al, 2016) means that the importance of the AA contribution remains unresolved.

As robust linkages are difficult to detect, novel statistical analysis of the observational climate record has been put forward as a potential way to move the Arctic-midlatitude field forwards (Overland et al, 2016; Kretschmer et al, 2016; Cohen et al, 2018a), acting as a supportive tool for large coordinated modelling projects. Studies which rely on correlation analysis are subject to autocorrelation bias, as well as misleading results due to indirect links or common drivers of correlated variables which were unaccounted for in the analysis (Runge et al, 2014). Such linear relationships are also directionless, so offer less information than graphical models. Atmospheric model studies, whilst well regarded as tools for identifying causal linkages, are not immune from potential shortcomings: they may not accurately represent ocean-atmosphere coupling in the Arctic (Cohen et al, 2018a), may respond too weakly to sea ice forcing (Screen et al, 2018; Mori et al, 2019), may underperform in terms of stratosphere-troposphere coupling (Zhang et al, 2018), and focus on the impact of sea ice

removal which may not capture the complex intermittencies thought to define Arctic-midlatitude linkages (Overland and Wang, 2018).

This chapter applies Dynamic Bayesian Networks (DBNs) with hidden variables (HVs) to the North Atlantic and European midlatitude circulation research area. Structure-learning algorithms are employed to identify regions of AA that might influence winter jet stream variability. A number of Arctic, midlatitude and tropical variables are included to investigate the relative role of AA as a driver compared to internal atmospheric variability and other remote forcings. Other graphical model approaches focus either entirely on potential Arctic drivers of midlatitude circulation responses (Kretschmer et al, 2016; Barnes and Simpson, 2017; Samarasinghe et al, 2019) or on tropical teleconnections like the MJO-NAO link (Barnes et al, 2019). The aim of this study is to establish how effective DBNs with structure learning algorithms are for investigating this research area, and to measure the impact of hidden variables on model accuracy which is a priority due to the low signal-to-noise ratio of AA linkages and their intermittent nature (Overland et al, 2016). We demonstrate a feedback relationship between North Atlantic midlatitude circulation and two important regions of Arctic warming occurring in winter at submonthly timescales. Finally, the implications for further study are discussed.

5.2 Data

A number of climatological variables which are understood to have an influence on European midlatitude weather during the winter months are included in this analysis. Four nonoverlapping time average resolutions consisting of 5-day, 10-day, 15-day and monthly averages were used to ensure a robust set of conclusions that include variables that may act at a range of timescales. The 29th February was removed from each leap year, allowing for a total of 90 days in each year using only the winter months (December to February, DJF) from the years 1981-2018.

To investigate the relative impact of the Arctic on midlatitudes, tropical indices formed part of a network of non-local drivers (Fig. 5.1). Both ENSO (Brönnimann, 2007) and the Madden-Julian Oscillation (MJO) (Lin et al, 2015) have been found to explain significant amounts of North Atlantic Oscillation (NAO) variability during wintertime through Rossby wave excitement (Trenberth et al, 1998). A standardised Niño 3.4 Index is included, calculated using data from HadISST1 (Rayner et al, 2003), retrieved from the KNMI climate explorer (<https://climexp.knmi.nl/>). The Real-Time Multivariate Madden-Julian Oscillation (MJO) index devised by Wheeler and Hendon (2004, available at <http://www.bom.gov.au/climate/mjo/>) is generated by identifying the PCs of two combined empirical orthogonal functions (EOFs) of equatorially-averaged 850-hPa zonal wind, 200-hPa zonal wind and outgoing longwave radiation (OLR) anomalies. The two PC time series (RMM1 and RMM2) are used to define the

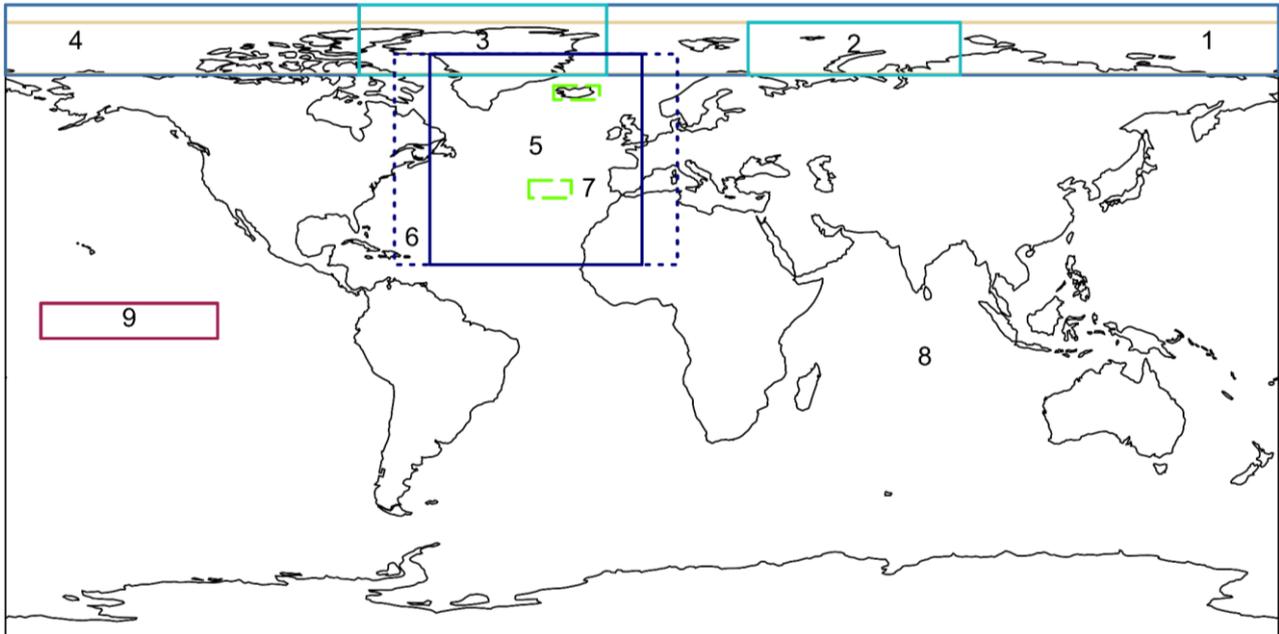


Figure 5.1 Map of all variables used in DBN experiments: 1 Arctic (dark blue); 2 Barents-Kara Seas, BK; 3 North Atlantic, NAT; 4 stratospheric polar vortex, POV (brown); 5 jet latitude, Jet; 6 Meandering Index, MI (dotted blue); 7 North Atlantic Oscillation, NAO (dashed green boxes); 8 Madden Julian Oscillation, MJO; 9 El-Niño Southern Oscillation, ENSO.

MJO amplitude, equal to $\sqrt{(RMM1^2 + RMM2^2)}$. The principal-component (PC)-based NAO Index (Hurrell, 1995), based on the difference in surface sea-level pressure (SLP) between the Subtropical High and the Subpolar Low, provides a secondary indicator of North Atlantic atmospheric variability and was also obtained from KNMI's climate explorer.

To measure the impact of Arctic variability, near-surface 850-hPa temperature from ECMWF's ERA-Interim global atmospheric reanalysis (Dee et al, 2011) was included in the networks, available at <https://apps.ecmwf.int/datasets/data/interim-full-daily/>. 850-hPa temperature is used in place of sea ice concentration to capture the full effect of AA. Arctic near-surface warming is thought to be a better indicator of AA than sea ice variability because the AA signal is made up of a number of factors which include sea ice retreat, and the impact of heat and moisture transport from lower latitudes is fully accounted for (Cohen et al, 2018a). All data were compiled as 1° x 1° gridded spatio-temporal data and then processed. The entire Arctic region was prepared alongside important subsections as regional sea ice loss is known to be an important factor in midlatitude circulation responses (Pedersen et al, 2016; Screen, 2017). Three regions of Arctic 850-hPa temperature were therefore used: the Arctic (180°W-180°E, 70-90°N), the North Atlantic sector (280-350°E, 70-90°N) and the Barents-Kara Seas region (30-90°E, 70-85°N). These capture three key areas of sea ice concentration loss: Baffin Bay; the east coast of Greenland; and the Barents-Kara Seas (see Fig. 2 in Overland and Wang 2018). An area-weighted spatial average was taken over these regions, and anomalies

Variable	Abbreviation	Source	Unit
Arctic 850hPa temperature	Arctic	ECMWF ERA-Interim	Temperature (°C)
Barents-Kara Seas 850hPa temperature	BK	ECMWF ERA-Interim	Temperature (°C)
North Atlantic 850 hPa temperature	NAt	ECMWF ERA-Interim	Temperature (°C)
Stratospheric polar vortex	PoV	ECMWF ERA-Interim	Geopotential height (m)
Jet latitude	Jet	ECMWF ERA-Interim	Degrees (°)
Meandering Index	MI	ECMWF ERA-Interim	-
North Atlantic Oscillation	NAO	NOAA	Sea level pressure
Madden-Julian Oscillation	MJO	NOAA	MJO amplitude
El Niño Southern Oscillation	ENSO	NOAA	Sea surface temperature

Table 5.1 Variables used, their abbreviations, source and unit of measurement.

were calculated from the resulting univariate time series by subtracting each timestep of the multi-year mean (1981-2018) from the value of the matching timestep. This was done for each time resolution, with each then detrended to create a time series for the variable.

As a proxy for stratospheric variability, the polar vortex (PoV) was included to examine the impact that stratospheric circulation might have on tropospheric midlatitude circulation. The PoV index uses ERA-Interim geopotential height anomalies from the Arctic region (180°W-180°E, 65-90°N) averaged over 6 pressure levels from 10 to 100 hPa, with the resulting time series used to create anomalies at each time average and detrended. As such, the PoV opens up the potential for linkages through tropospheric and stratospheric pathways, and allows for comparison with previous studies that identify a link between the PoV and midlatitude circulation (Kim et al, 2014; Kretschmer et al, 2016).

Two metrics were used to represent North Atlantic midlatitude circulation in the networks. The first, jet latitude, was calculated from ERA-Interim data and defined over the region 16°–76°N, 0°–60°W. Jet latitude was determined using the approach taken by Woollings et al (2010): zonal winds were height-averaged over 900 to 700 hPa and filtered with a 10-day Lanczos low-pass filter using a 61-day window to ensure that synoptic scale variability is excluded. The use of lower level winds isolates the eddy-driven jet as the data is not contaminated with the signal of the subtropical jet. Jet latitude output is consistent across pressure levels, and many jet-focused studies have made use of this metric (Hall et al, 2017; Barnes and Simpson, 2017; Samarasinghe et al, 2019).

The second metric of jet variability is the Meandering Index (MI), a measure of tropospheric circulation variability that uses geopotential height contours to capture the maximum waviness at each timestep, taking into account the full spatial position of each contour (Di Capua and Coumou, 2016). The input data used was ERA-Interim 500-hPa geopotential height (gph) for the region 16°–76°N, 10°E–70°W. The MI calculates the length of each isohypse (gph contour) on a 2-D grid, which is then normalised to the Earth's circumference at 60°N. The maximum value of this calculation on the vertical profile (e.g. from 4800m to 5600m) is then taken as the MI, allowing for accurate differentiation between strongly meridional (north-south) deviations and consistently zonal configurations. A full description of the MI can be found in Di Capua and Coumou (2016). The following analysis included both the jet latitude and MI to ensure linkages revealed by the DBN technique were robust to the use of different metrics. The two metrics also describe different aspects of midlatitude circulation; the jet latitude by definition describes the latitudinal position of the jet on a given day, whilst the MI focuses on waviness of the middle troposphere, and thus may reveal the waviness of the jet but cannot be used as an indicator of the jet core location.

All data were standardised as a final stage of data preparation, so that each variable had a mean of 0 and a standard deviation of 1. This was done to maximise model accuracy in the parameter learning stage by giving all variables equal means and similar ranges.

5.3 Methods

To test the hypothesis that a number of important climatological relationships exist between North Atlantic atmospheric variability and remote Arctic and tropical drivers, we built a series of networks of increasing complexity. This ensured that any relationships captured by the structure-learning algorithm were consistent across models, and robust in terms of their predictive accuracy. All networks were constructed using the Bayes Net Toolbox (BNT) for MATLAB (Murphy, 2001b), with all data preparation and plotting carried out in R.

Auto-correlation Functions (ACFs) were used to determine that all variables depended linearly on their values from the previous timestep (not shown). Thus, autoregressive links

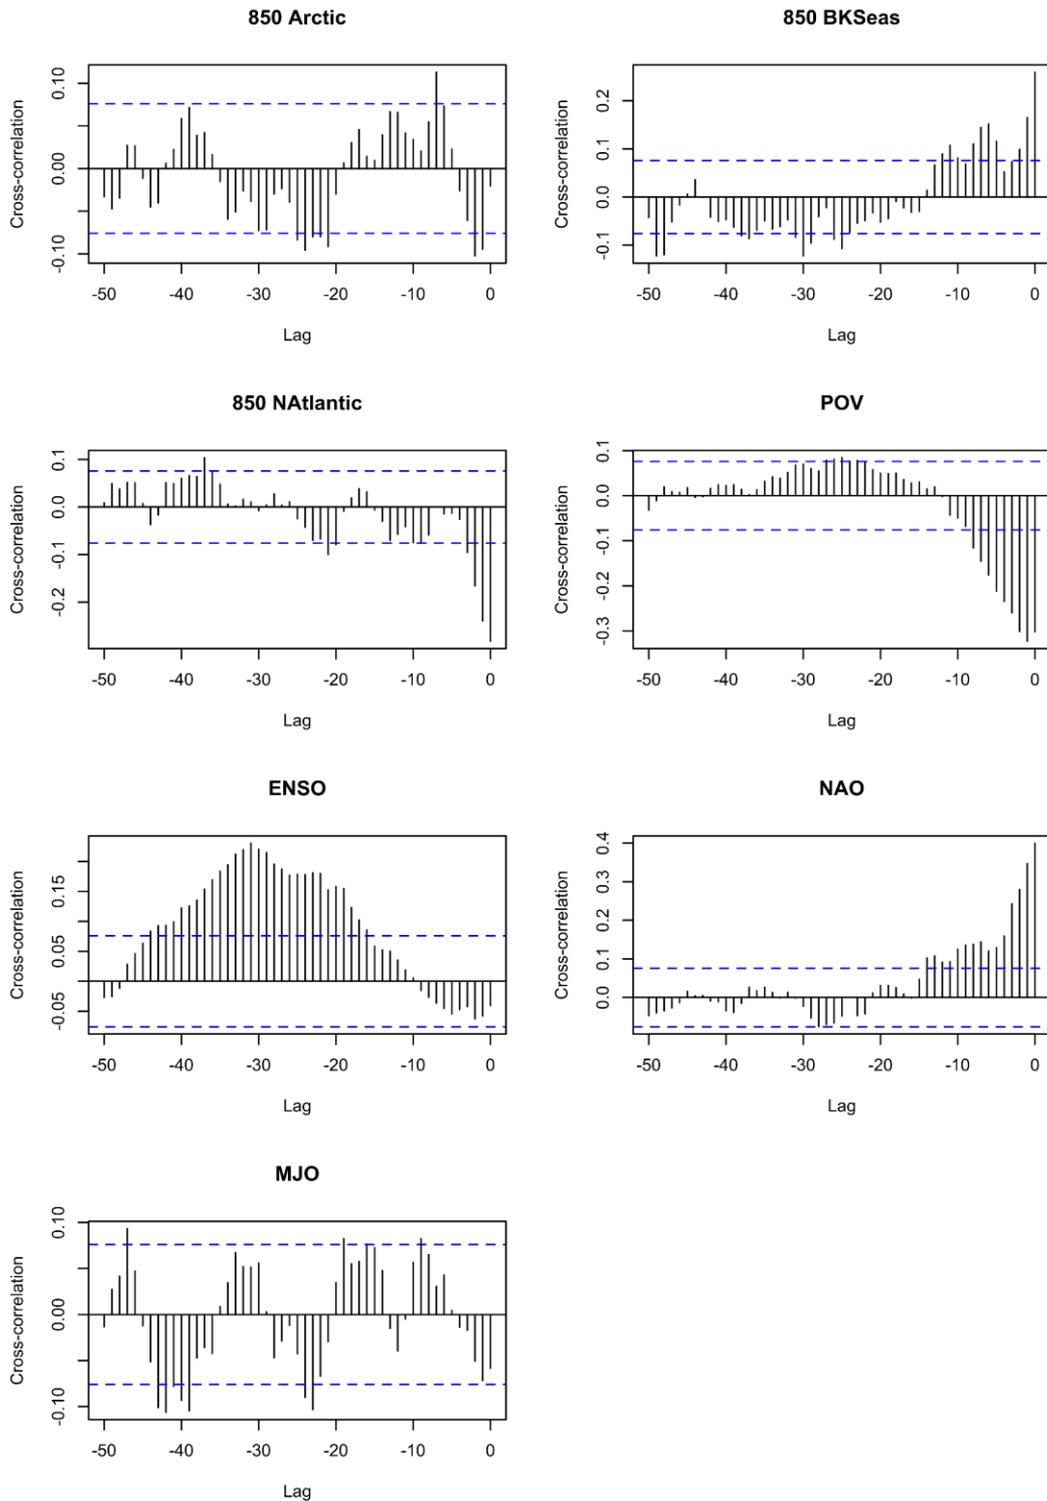


Figure 5.2 Cross-correlation functions (CCFs) between all variables and jet latitude, used to determine the lead time, if any, for each dataset. Blue dotted line denotes significance at the 0.05 level.

were coded into the models for all variables. Cross-correlation Functions (CCFs) were then plotted for each variable against jet latitude to determine the number of timesteps (if any) at which each variable should be lagged (Fig. 5.2). The appropriate lead time for each variable

was then used to build the datasets used for all analyses; where no clear maximum lag value was returned in the CCFs, the nearest significant lag was selected for the dataset. For the MJO, with a sinusoidal pattern of correlations, a lead time of 45 days was chosen. This lead is slightly longer than links between the MJO signal to NAO forecasting (Yu and Lin, 2016) and variability (Jiang et al, 2017). Section 5.6 below is dedicated to investigating different lead times for the MJO and ENSO using expert-guided lags taken from previous literature. Further 5-day networks were thus run with a range of lead times to ensure tropical influences were not missed simply as a result of CCF lead selection.

5.3.1 Experiments

After being prepared with the steps listed above, the data was loaded into BNT and split into training and testing datasets (80/20) to allow for an unbiased estimate of the generalisation error and prevent overfitting (Shalizi, 2013). All DBNs were run first with jet latitude, then with the MI swapped in its place. Four main sets of networks are constructed which broadly follow the experimental design of Chapter 4, but are aimed at investigating the impact of networks with or without a learned structure, with or without a hidden variable and Arctic variables averaged over the entire Arctic or smaller subregions:

- 1) Hidden Markov Model – a HMM with autoregressive links and a fixed structure imposed on the network.
- 2) A ‘control’ run – a DBN with autoregressive links and no HV to act as a control run and examine the impact of adding a HV on predictive accuracy.
- 3) The entire Arctic – a DBN with autoregressive links and one HV linked to all variables with the full Arctic area (70-90°N) included and Arctic subregions removed (Barents-Kara Seas and North Atlantic), to examine differences in structure and model accuracy.
- 4) Arctic regions – A DBN using Arctic subregions (full Arctic removed) with autoregressive links and one HV linked to all variables by default, to identify state shifts in Arctic-Midlatitude linkages and model the learned structure as accurately as possible using the full suite of tools that graphical model approaches provide.

As a first step for the DBNs, the structure was learned using the PC algorithm, using the ‘*fisher z*’ test for conditional independence and an alpha value of 0.01. Results were robust to the choice of alpha between 0.01 and 0.05. In the case of (3) and (4), a HV with links to all variables and to itself in each time slice is coded into the structure, as in Uusitalo et al (2018). Next, any bidirectional (cyclical) links are removed from the learned structure by necessity as DBNs require acyclical graphs. DAGs by definition cannot include loops and have bipartite

structures (Scutari and Denis, 2015). Bidirectional arcs occurred for the jet latitude variable and were removed such that only those arcs pointing towards the jet were preserved. The structure revealed the jet node to be a multiple collider node, and as such the links into it needed to be kept to preserve the ‘explain away’ effect for the parent nodes. Doing otherwise would have significantly reduced the accuracy of the DAG structure; removing arcs from collider nodes with multiple inbound conditionally dependent relationships undermines the individual probability distributions that make up a BN.

The parameters are then learned with the EM algorithm and the model is tested on the remainder of the data (the test dataset). BNs perform prediction using inference (Friedman et al, 2000); predicted values are used for model validation by plotting predictions against the (observed) test dataset. Model fit is shown in the predictive accuracy scatter plots below, and assessed using sum of squared error (SSE) calculated for each variable at each time resolution for all network types using the following:

$$SSE = \sum(\text{predicted} - \text{observed})^2$$

It is worth noting that different HV setups were investigated, and a DBN with two HVs and multiple HV structures was initially incorporated into the analysis (not shown). This resulted in the model being overly parameterised, with not enough data to fit the model, as is the case with Chapter 4’s ‘DBN2HV’ results. Clearly for the data and resulting graph structures presented here, a single HV is the limit.

As a final step, an analysis of variance (ANOVA) was conducted on all networks of experiment (4) to estimate the relative importance of each relationship found. Linear regressions, using all parent nodes as independent variables to predict dependent variable (child node) values, were calculated for all edges in networks across all time resolutions. An incremental sum of squares table was produced using ANOVA for each regression model, and the proportion of variance worked out by dividing the sum of squares for each independent variable by the total sum of squares scaled by 1 (Montgomery, 2012). Strength estimates are limited by the fact that independent variables were determined using the directions set at the DAG stage of network construction, and that ANOVA examines only linear relationships. Here, they are used to assign edge weights (arrow widths) to the DAG plots below.

5.4 Results

Graphical models provide an easily interpretable interface, allowing for visualisation of climate teleconnections as DAG structures. Nodes representing the variables are coloured to indicate their relative geographical position for ease of interpretation. Faint edges represent relationships hard-coded into the model, i.e. the HV connections, and solid edges show those

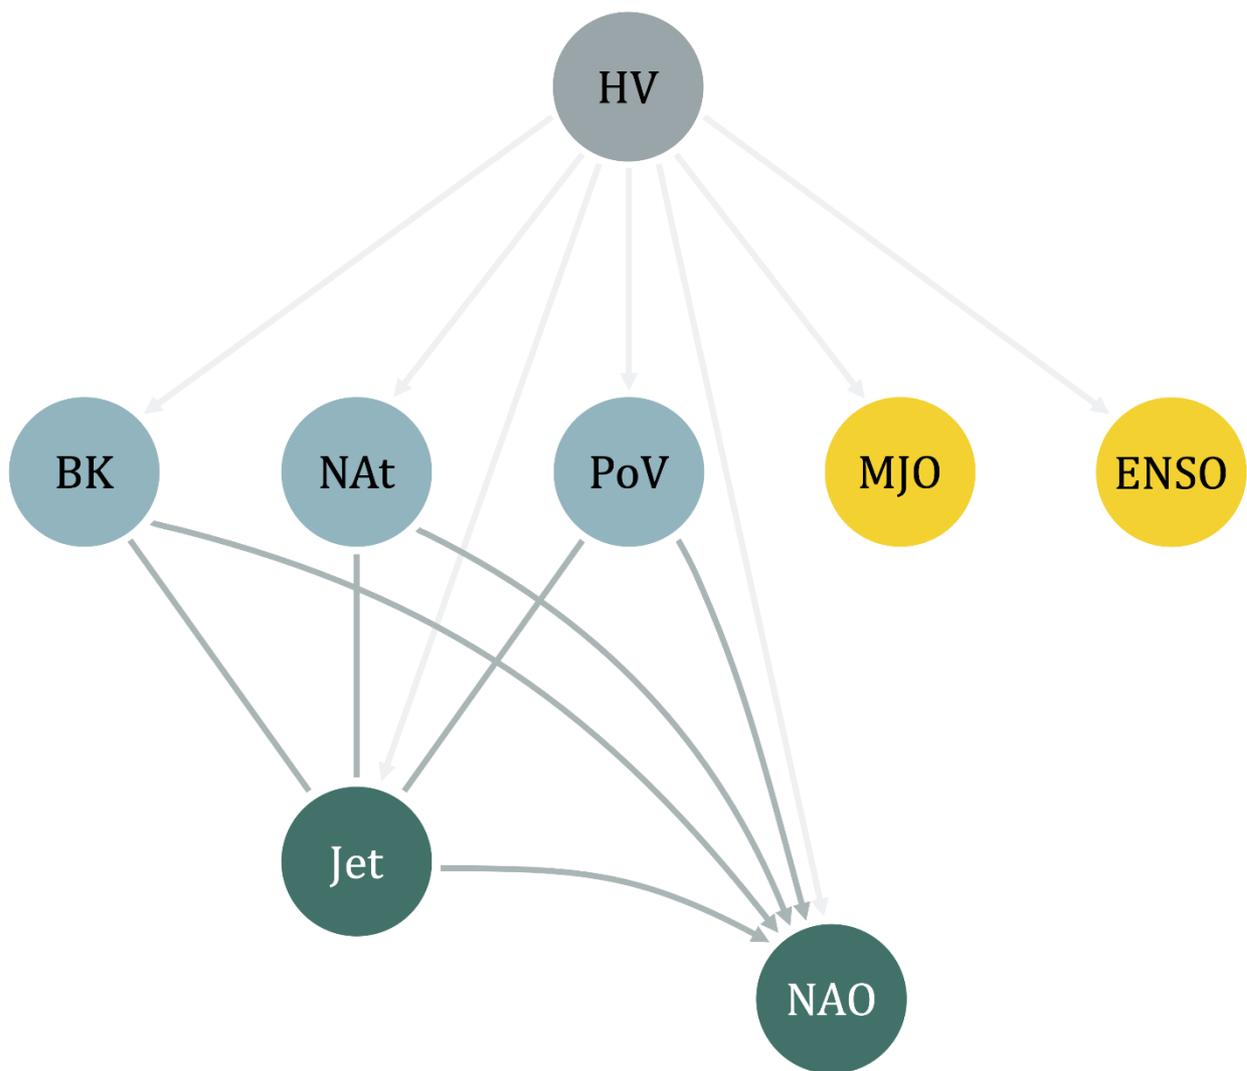


Figure 5.3 pDAG returned by the PC algorithm for the 5-day Jet HV DBN. Faint lines indicate edges coded into the model (i.e. HV edges), solid represent learned edges. Nodes are coloured by their relative geographical location: the tropics (yellow), midlatitudes (green) and the Arctic (blue). Only the jet DBN is shown here; the MI DAG plot is identical at this stage.

learned by the structure-learning algorithm in the case of all non-HV links. Two nodes connected by the latter can be said to be conditionally dependent on each other, but BN arcs do not show cause-and-effect relationships (Scutari and Denis, 2015). Two stages of the network are shown: the partially directed graph (pDAG) returned at the structure-learning stage which includes the HV and tropical nodes (MJO and ENSO; Fig. 5.3), and DAGs showing learned relationships used to run the DBNs in the second stage (no HV or tropical variables; Figs. 5.4 and 5.5).

5.4.1 Network Structures: the DAG Results

The control run (no HV) DBNs generated the same learned structures as seen in Fig. 5.3; model accuracy changes with the addition of HVs is discussed below. Arctic-midlatitude linkages were only captured by the DBNs when regions of AA were used (BK and Nat in Fig.

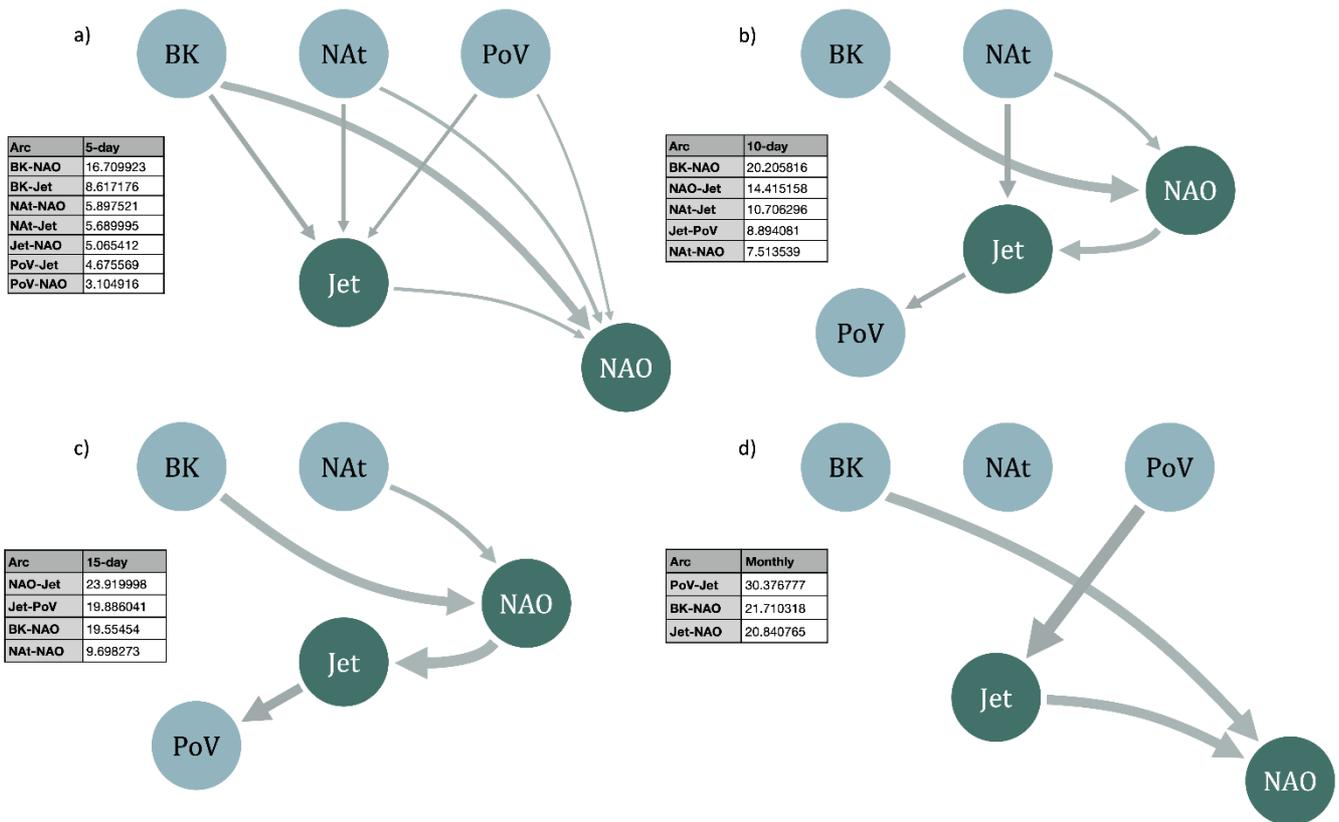


Figure 5.4 DAGs for the Jet HV DBN at a) 5-day, b) 10-day, c) 15-day and d) monthly time averages, with the HV and unconnected tropical variables removed for visualisation purposes. Arc direction is indicated by arrows, and inset tables show relative strength in terms of proportion of variance used to define arrow width.

5.1). For experiment 3, DBNs using the entire Arctic revealed no links between the Arctic and the jet or MI, showing only an Arctic-NAO linkage over 5-day averages (not shown). Midlatitude circulation responses to AA are known to be sensitive to the regions selected (Pedersen et al, 2016; Screen, 2017). The results for the entire Arctic therefore suggest that Arctic-midlatitude linkages are sensitive to the location of sea ice loss and amplified warming.

In the full networks of experiment 4 (Fig. 5.3), both the jet and MI DBNs capture a relationship between BK, NAt and jet latitude for 5-day averages, pointing to a covariability between important regions of AA and the jet stream's latitudinal position. Prior to the removal of cyclical links, the stage of network construction which allows DBNs to function as detailed above, these links were undirected which may indicate that simplistic cause-and-effect relationships do not explain interactions between AA regions and jet variability (Overland and Wang, 2018). This is not the case for the MI DBNs; the meridional component of the jet stream is conditionally dependent on NAt, a region of rapid warming and sea ice loss which includes Baffin Bay, and seems to impact temperature over the Barents-Kara Seas at 5-day resolutions (Fig. 5.5a). BK is identified as a potential driver of the MI at 10- and 15-day averages, when the direct link between NAt and the MI is lost.

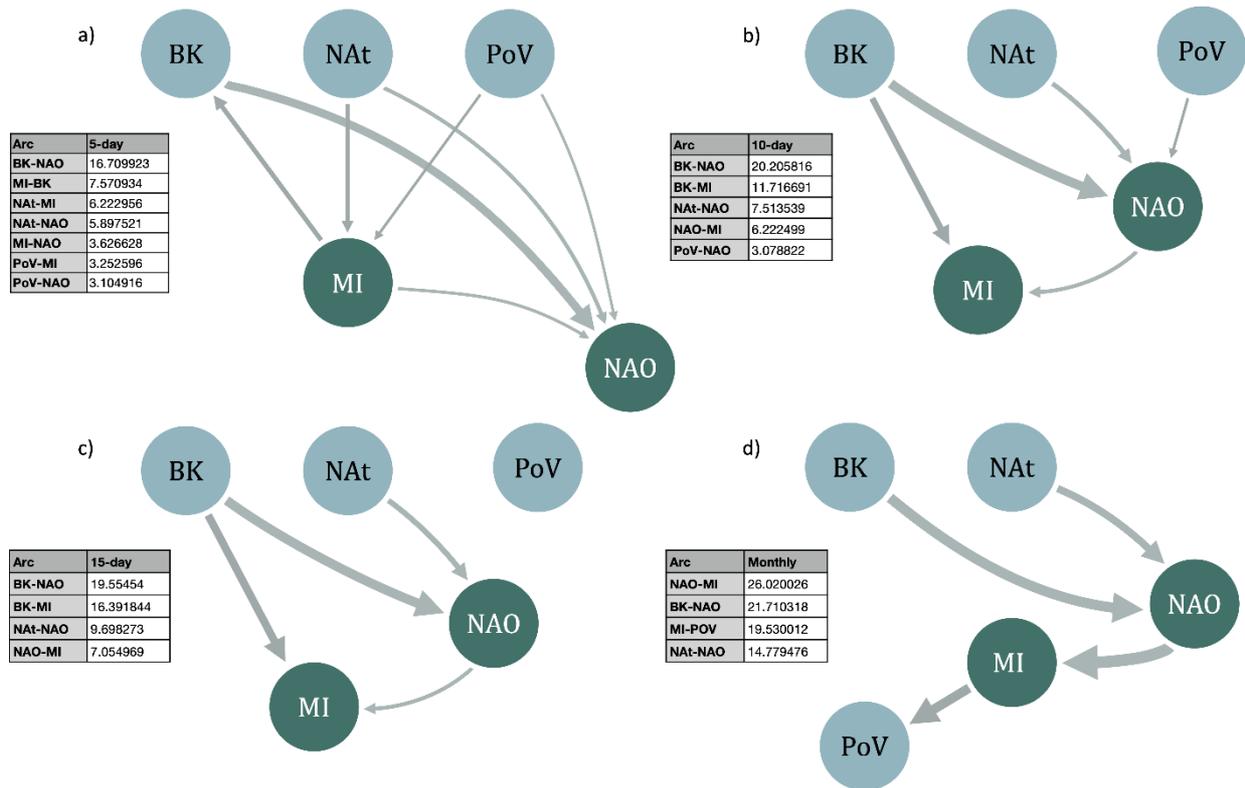


Figure 5.5 As with Fig 5.4, but for MI HV DBNs at a) 5-day, b) 10-day, c) 15-day and d) monthly time averages.

Evidence of tropospheric-stratospheric coupling with short lead time of 5-10 days is also found across all DAGs excluding the 15-day MI DBN (Fig. 5.5c). Indirect links via a stratospheric pathway are missing however as no links pass through the PoV, which is either a parent or child node depending on the time resolution (Figs. 5.4 and 5.5). The PoV becomes the strongest link as a driver of jet variability at monthly intervals, but the vast increase in PoV SSE at 10- and 15-day intervals (Table 5.2a) suggests the direction of the jet-PoV arc is incorrect, as in monthly MI SSE (Table 5.2c). The DAG structures between Arctic (BK, NAT), midlatitude (jet, MI and NAO) and stratospheric (PoV) variables (Fig. 5.3) were found to be robust to the addition of extraneous variables with no learned links; i.e. they did not change when unconnected variables were deliberately added. Whilst these structures were found across networks, there is a noticeable difference between submonthly timescales where AA variables seem to have a stronger link to midlatitude circulation, and monthly DAGs which suggest that the PoV connection is strongest (see proportion of variance tables inset in Figs. 5.4 and 5.5).

Neither of the tropical variables are found to be conditionally dependent on any other variable for both the jet and MI DBNs over all time averages. This result was also robust to the choice of alpha value between 0.01 and 0.05, choice of score or constraint-based algorithms (only results for PC are presented here), and particular MJO index used (the RMM

index is shown here). Figs. 5.4 and 5.5 thus remove the HV and tropical links for ease of visualisation. This finding is consistent across model runs and is discussed in detail below.

5.4.2 Network Performance: Predictive Accuracy

Comparison of predictive accuracy at different time averages shows a notable decrease in accuracy towards lower time resolutions. Figs. A5.1-A5.8 (see Appendix) show the predicted values plotted against the observed values for each variable from the test dataset to assess model fit, as the structure learning and model parameterisation used only the training dataset. It is obvious from Figs. A5.1-A5.4 that the 5-day DBN performed the most accurately, with both the jet and PoV variables losing accuracy in Figs. A5.2 and A5.3 and the slope of the BK, NAt and NAO variables showing an increasingly worse fit. A similar decrease is obvious for the MI DBNs (Figs. A5.5-A5.8). This decrease in accuracy suggests that the structure of the 5-day DBNs most accurately describes the underlying conditional dependencies of the data, as the BK, NAt and PoV variables seem to optimise model performance as parent nodes (Fig. 5.4a). Given that the main difference in the 10- and 15-day DAGs was the loss of the BK (10-day) and BK and NAt (15-day) nodes as parents of the jet variable, alongside the reversed direction of the PoV-jet arc, Figs. A5.2 and A5.3 affirm the importance of Arctic-midlatitude teleconnections for network analyses of midlatitude variability.

Interestingly, the MJO and ENSO remain relatively well predicted throughout the resolution steps for both the jet and MI DBNs, despite only being connected to the network through the HV in all DBN DAGs. Highly autoregressive nodes continue to be more easily predicted throughout the loss of resolution, and this is especially true of ENSO which exhibits long periods of stability in terms of its periodicity compared with the rest of the time series' lengths.

In contrast to Chapter 4, where the autoregressive HMM outperformed DBNs for a simple 4-variable network, the jet and MI HMMs were by far the worst performing models in terms of SSE (Table 5.2). Clearly, the structure provides significant improvements in the predictive capability of the networks as the SSE greatly increases when structure is removed, particularly significant for the 5-day averaged results for the two HMMs.

Obvious disparities in SSE may also provide information about structure. Large increases in SSE can be seen in Table 5.2a for the PoV variable for 10- and 15-day time averages, where the PoV-jet relationship shown in the 5-day averaged data is reversed (Fig. 5.4). This suggests that the jet-PoV direction is the one that provides the most predictive power for this set of DBNs at submonthly time resolutions. The same is true for the MI-BK relationship, which seems to impede accuracy for the 5-day averaged MI DBN runs, but not for 10- and 15-day where the PC algorithm returned the same relationship in the opposite

	5 day	10 day	15 day	monthly
a) Jet HMM				
850 BKSeas	80.77	34.56	36.12	14.6
850 NAtlantic	87.01	60.18	52.59	28.54
POV	44.88	11.56	14.44	14.27
Jet	49.19	26.89	26.21	12.62
NAO	56.56	39.05	41.64	21.38
MJO	71.99	40.81	28.96	21.67
ENSO	44.76	17.02	11.16	9.9
b) MI HMM				
850 BKSeas	102.89	50.13	52.94	19.96
850 NAtlantic	114.23	66.24	58.82	28.26
POV	68.62	10.41	19.13	16.23
MI	145.95	54.48	58.55	32.99
NAO	87.92	38.05	80.94	35.84
MJO	95.46	41.12	28.22	22.28
ENSO	58.03	16.95	16.32	7.66
c) Jet HV DBN				
BK	25.46	26.67	18.41	13.46
NAt	44.52	35.4	43.97	10.95
POV	1.68	20472.46	34495.65	4.84
Jet	28.12	258.23	31.21	9.38
NAO	34.49	16.54	35.78	23.66
MJO	7.67	15.77	12.35	10.76
ENSO	0.37	0.1	0.08	0.54
d) Jet no HV DBN				
BK	34.46	31.79	22.56	9.9
NAt	52.29	53.18	40.37	21.15
POV	1.66	166086.63	608.01	9.4
Jet	31.11	235.46	85.98	9.18
NAO	26.44	31.03	28.03	31.87
MJO	11.05	17.09	17.93	21.43
ENSO	0.56	0.66	0.85	1.68
e) MI HV DBN				
BK	17627.71	25.29	25.33	14.87
NAt	44.13	46.65	44.34	9.15

POV	0.71	3.65	16.51	259469.81
MI	96.25	154.68	505.25	371.19
NAO	383.63	31.26	51.61	26.54
MJO	4.56	13.41	9.52	14.53
ENSO	0.11	0.41	2.81	2.95
f) MI no HV DBN				
BK	1887.33	31.79	22.56	9.9
NAt	52.29	53.18	40.37	21.15
POV	1.66	4	5.22	16.68
MI	70.55	140.54	56.51	34.06
NAO	81.94	32.18	28.03	26.31
MJO	11.05	17.09	17.93	21.43
ENSO	0.56	0.66	0.85	1.68

Table 5.2 SSE for each variable at the 5-day, 10-day, 15-day and monthly time resolutions for a) Jet HMM b) MI HMM c) jet DBN with a HV d) jet DBN no HV e) MI DBN with a HV and f) MI DBN no HV.

direction (i.e. BK-MI; Fig. 5.5a). In essence, Table 5.2 shows us that the skill of networks can be reduced by falsely imposing the direction of relationships.

The DBN with a HV (Fig. 5.4) had the most accurate performance for the jet DBNs in terms of SSE (Table 5.2), indicating that the inference engine was able to accommodate the increase in model complexity. MI DBNs without a HV performed marginally better overall although a few variables were predicted with less accuracy, most notably the MJO and ENSO nodes. The drop in SSE for the MI DBNs seems to be a result of PoV variability dominating the hidden state switches, discussed below in the HV analysis; Fig. 5.5 demonstrates that the PoV node is less important for MI variability in contrast to jet latitude, where the PoV is connected throughout the time resolutions. The DBN without a HV outperformed it because the strongest influence on the hidden state means and switches was the PoV, which is not as central to the learned structure of the MI models as the jet ones. This explanation, paired with the generally worse predictive accuracy of the MI DBNs, also suggests that the MI-BK link in 5-day averaged data is more accurate the other way round (i.e. BK-MI) as in the jet 5-day DAG, and points towards the jet DBN with a HV being the most accurate model learned in this analysis. Clearly, there are caveats to using HVs as model accuracy is not guaranteed to increase across all variables; this finding is reflected in Trifonova et al's (2015) SSE results for biomass prediction and the 'DBN2HV' results in Chapter 4 which suffered from overfitting.

HV ‘states’ are the three possible values the discrete HVs can take, inferred from the observed data to maximise model fit. Timeseries of all variables were split into each HV state, and the three states were then summarised in terms of mean and standard deviation (sd) to investigate what they represent (Table 5.3). For the jet DBNs, state 3 represents ‘average’ conditions with values close to the mean values of the whole training dataset timeseries for the BK, NAt, PoV and jet nodes. State 1 is associated with higher than average PoV geopotential height anomalies (slower than average stratospheric polar vortex), lower temperatures over the NAt region and a higher average jet latitude. Conversely, state 2 indicates a negative PoV geopotential height anomaly pattern (faster than average polar vortex), higher NAt temperatures and average jet conditions. Both states 1 and 2 have values within one standard deviation of the mean, but the HV nonetheless captures three distinct states defined mainly by the PoV, NAt and jet variables. The MI HV states are less clear, but both sets of models are characterised by switches between states 1 and 3 (Fig. 5.6); two obvious differences are that the MI stays in state 1 for longer on average, and that the values for state 1 of the MI HV are closer to each variable’s mean value for the whole dataset than those of the jet HV.

HV state	BK	NAt	POV	Jet/MI	NAO	MJO	ENSO
<i>a) Jet DBN</i>							
1	-0.44	-0.48	440.47	49.06	0.2	1.47	-0.04
2	-0.94	0.98	-490.58	47.83	0.3	1.39	0.09
3	-0.01	0.25	114.27	47.68	0.19	1.36	-0.01
<i>b) MI DBN</i>							
1	-0.3	-0.31	-41.52	1.67	0.24	1.32	0.01
2	0.15	0.81	-178.6	1.64	0.05	1.37	-0.03
3	-0.14	0.22	289.96	1.73	0.19	1.42	-0.02

Table 5.3 Hidden state mean values of each variable for a) Jet DBN and b) MI DBN.

5.5 Discussion

The DBNs provide a relatively robust set of results in terms of network structure and model accuracy. The results described hereafter generally focus on 5-day averages unless otherwise specified. One of the most consistent linkages picked up by the structure-learning

phase points to the relationship between the Arctic regions of the North Atlantic and Barents-Kara seas, and the jet, MI and NAO nodes. Atmospheric responses in Northern Europe to anomalous sea ice loss in the Barents-Kara region were thought to develop through changes in surface turbulent heat fluxes (Petoukhov and Semenov, 2010; Liptak and Strong, 2014), although recent evidence casts doubt on this (Blackport et al, 2019). However, the bidirectional edges of the BK and NAt to the jet (Fig. 5.3) point to a positive feedback of Arctic-midlatitude impacts, matching Samarasinghe et al's (2019) findings over similar timescales (5 to 15 days) for the Pacific, which captures the thermodynamic effect of temperature and moisture advection from lower latitudes and the modulation of Arctic temperatures (Kapsch et al, 2016; Gong et al, 2017; Cohen et al, 2018a). This feedback occurs on timescales of <5 days (Kapsch et al, 2016), suggesting Fig. 5.3 captures a pathway thought to occur at similar time resolutions in model studies; one that may dictate the nature of the two-way relationship between AA and the midlatitudes.

The NAt-jet linkage, found in both the jet and MI 5-day DBNs and for 10-day averages in the jet, may represent an Arctic-jet connection known to develop intermittently and cause cold outbreaks through the assisted formation and support of blocking patterns (Chen and Lou, 2017; Ballinger et al, 2018). Higher geopotential heights in the Greenland and Baffin Bay regions increase the likelihood of Greenland blocking events, contributing to the increased waviness of jet stream patterns (higher MI values) and in turn the persistence of cold events on the eastern coast of the US (Chen and Luo, 2017; Overland and Wang, 2018). The meridional component of the jet stream (the MI) is conditionally dependent on AA over the North Atlantic region, a region of rapid warming and sea ice loss which includes Baffin Bay, and seems to impact temperature over the Barents-Kara Seas at 5-day resolutions (Fig. 5.5a). Interestingly, BK is then identified as a potential driver of meridional patterns at 10- and 15-day averages, when the direct link between NAt and the MI is lost.

The NAO, an index reflecting both jet latitude and speed variability (Woollings and Blackburn, 2012), is central to the network structures in both sets of DBNs. Whilst the NAO is likely influenced by a complex set of nonlinear drivers (Smith et al, 2016), near-surface Arctic temperature seems to have a significant impact on the NAO via the troposphere. The NAO is conditionally dependent on the high sea ice variability region of the Barents-Kara Seas, estimated to be the strongest link in 5- and 10-day jet DAGs and all submonthly MI intervals (Figs. 5.4 and 5.5). The response of atmospheric circulation indices to AA is constrained entirely within winter in this study as in Blackport and Screen (2019); in contrast, others have found that autumn sea ice conditions provide predictive skill of winter NAO variability in statistical models (Wang et al, 2017; Hall et al, 2017). Fig. 5.5 shows that NAt and MI are conditionally independent given the NAO node, but that they are connected via the NAO

throughout the time averages as the BK and NAt nodes impact NAO variability even to monthly resolutions. This is almost identical for jet latitude up to 15-day intervals.

This conditional independence, added to the strength of the AA-NAO links, implies that NAO phase shifts summarise the AA-midlatitude connection well at lower time resolutions. The negative phase of the NAO is indicative of a southerly displacement of the jet, which has been linked to colder, more severe winters in northern Europe and the eastern US and warmer conditions over Greenland and the Barents-Kara region (Cohen et al, 2018b). Whilst the overall effect of AA and a warming world should be that cold outbreaks become less intense (Ayarzaguena and Screen, 2016), AA may favour a shift towards meridional circulation patterns which can promote the increased frequency and persistence of cold air outbreaks (Cohen et al, 2018a). However, a shift to negative NAO patterns is by itself a potentially misleading trend as the warming trend may offset any dynamical cooling influence (Screen, 2017).

5.5.1 Tropical Influence and Stratospheric Teleconnection Pathways

No conditionally dependent relationships are found for the tropical variables; this was robust across all time averages, jet descriptors, and lead times used. This likely reflects on aspects of the data and the model design process, rather than simply pointing to a weak tropical influence on Arctic temperature and jet variability relative to Arctic-midlatitude covariability. Combined effects and long-duration teleconnections through the stratosphere may have been effectively masked by the time series approach. A consistent lack of ENSO influence across networks is significant given that skilful prediction of the wintertime NAO (Scaife et al, 2016) and AO (Sun and Ahn, 2015) can be achieved with the inclusion of tropical variability, and tropical rainfall has been found to explain a significant amount of NAO variability in linear relationships (Scaife et al, 2016; Hall et al, 2017). Likewise, an increase in MJO intensity in the Indian Ocean region has been linked to the positive NAO phase in winter averaged (DJF) data (Lin et al, 2015), and composite analyses of blocking over the North Atlantic and Europe have found that MJO phases play a significant role in blocking occurrence during winter (Henderson et al, 2016).

The stratospheric polar vortex has been shown to project onto the NAO and increase the likelihood of Atlantic blocking during vortex weakening events through tropospheric-stratospheric coupling (Kidston et al, 2015). The stratosphere exhibits much more stability than the troposphere below it, requiring Rossby wave activity from below to cause significant disruptions to the flow of the polar vortex. A large body of studies favour the tropics as the predominant source of Rossby waves necessary to trigger sudden warming events (Liu et al, 2014; Hitchcock and Simpson, 2014; Scaife et al, 2017b; Jiang et al, 2017; Hardiman et al, 2019), NAO variability (Scaife et al, 2017a) and Arctic warming itself (Yoo et al, 2012). It is

highly probable that uncaptured tropical teleconnections contribute to the strength of the PoV to jet, MI and NAO linkages; El Niño events can force the Aleutian low to strengthen and move eastwards, causing enhanced vertical wave transport and negative NAO conditions through the disruption of PoV flow (Scaife et al, 2017b; Hardiman et al, 2019). Similarly, enhanced tropical convection associated with the MJO may weaken the PoV and disturb the NAO (Liu et al, 2014; Jiang et al, 2017) depending on the background state of the PoV (Barnes et al, 2019), although NAO events do not depend solely on tropical convection. Only the final stage of these teleconnections via the stratosphere are captured in DBN structures.

5.5.2 The Implications of Hidden State Switches

The HV state switches did not show a clear delineation between an 'AA' or 'pre-AA' period (Fig. 5.6a and 5.6b); i.e. a shift to an amplified warming state signalled in the BK and NAt variables. This is precisely the type of underlying state switch one would expect to be identified in a HV but may have been filtered out by the detrending step of data preparation. The HVs also hint that the proportion of variance estimates (Figs. 5.4 and 5.5) may not fully represent the importance of stratospheric variability at submonthly timescales as the PoV seems to dictate the hidden state switches by shifting from average (state 3) to slower anomalous PoV flow in state 1 (Fig. 5.6c). Added to this, Fig. 5.4 shows the PoV is connected throughout the time resolutions, and the lack of arcs between the PoV and MI in Fig. 5.5b and 5.5c may have caused the drop in model performance (Table 5.2c) with the addition of a HV to the MI DBN. The HVs maximise model fit by inferring their state from the observed data. The low predictive accuracy of the MI variable in Table 5.2c, particularly towards lower time resolutions, could therefore be a result of the lack of links between the PoV and MI combined with the strong influence of the PoV on state switches.

It is worth noting that recent studies have suggested the covariability between regions of AA and midlatitude flow characteristics could simply be a result of internal variability, with a forced component originating in the tropics which was uncaptured in the networks. The relationship between BK and the NAO was found to be a recent and weak feature in reanalyses and ensemble simulations (Kolstad and Screen, 2019). Warner et al (2020) concluded that internal variability in the Atlantic sector combined with tropical Pacific forcing was more likely to cause BK-NAO covariability than sea ice forcing. Non-simulated contributions from remote forcing could not be ruled out in these studies, but strong anomalous tropical forcing continues to provide skill for forecasting in contrast to extratropical drivers, which remain limited by their low signal-to-noise ratio resulting from natural variability (Trenberth et al, 1998).

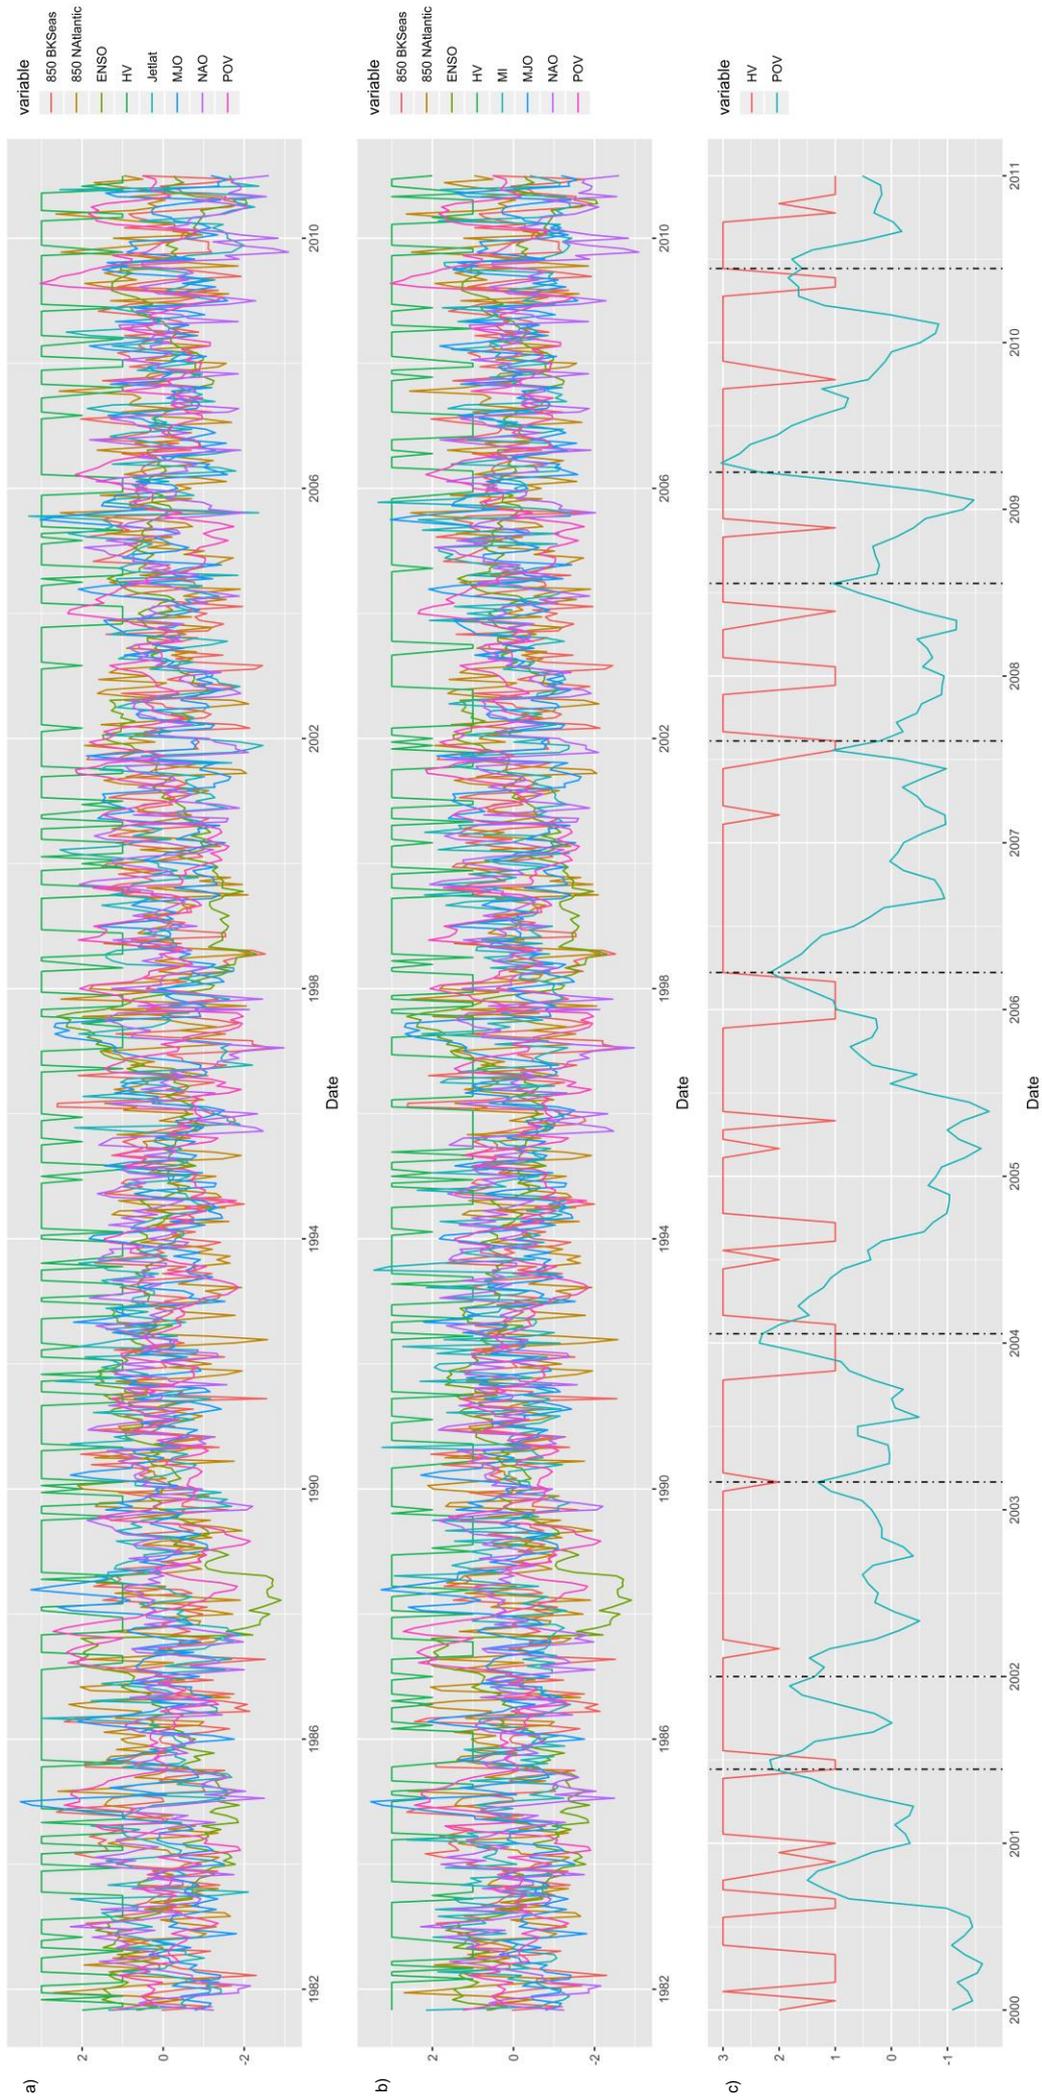


Figure 5.6 Winter (DJF) time series of 5-day train dataset data showing observed and hidden variables (green) for a) Jet HV DBN and b) MI HV DBN, c) is a subset of a) showing the final decade (2000-2010) of the HV and PoV variables from the train dataset with the onset dates of sudden stratospheric warming events (SSWs) plotted as vertical lines (dashed). SSW onset dates were taken from Cao et al (2019). All data in a), b) and c) are left as standardised for visualisation purposes.

5.6 Literature-based Networks

As discussed above, the lack of a tropical influence from the MJO and ENSO variables was an unexpected but consistent finding across all network types. Whilst the CCF analysis was conducted simply to investigate appropriate lead times for variables in relation to the jet and MI datasets, and was not used in the core BN analysis, CCFs can still return spurious relationships due to autocorrelation bias (Runge et al, 2014). An example of this is the CCF for the MJO variable, which demonstrates a sinusoidal pattern of correlations (Fig. 5.2). CCFs provided consistent results for other variables however, including the finding of a possible feedback relationship between the AA variables (BK and NAt) and the jet within the winter season on 5-day timescales. Another way of ensuring the appropriate lead times are found for all variables is to run the same network on multiple datasets with different lead times, which has been used in other graphical model approaches (e.g. Kretschmer et al, 2016; 2017).

To this end, multiple lead times were investigated for the tropical variables by constructing three new datasets with differing lead times based on several impactful studies. For ENSO impacts on North Atlantic circulation during the winter period, a response time of 2 weeks has been found for both stratospheric (Baldwin and Dunkerton, 2001) and tropospheric (Scaife et al, 2017) teleconnection pathways. It is important that both of these pathways are considered, as the stratospheric pathway is thought to impact NAO variability through the stratospheric polar vortex during La Niña events, whilst the tropospheric pathway is most likely the dominant of the two for strong El Niño events known to trigger Rossby wave trains into the North Atlantic (Hardiman et al, 2019).

For the MJO, Jiang et al (2017) found the MJO influences NAO events at timescales of up to 30 days, whilst Tseng et al (2018) suggest that a robust teleconnection between MJO phases 1, 2, 5 and 6 and 500hPa geopotential height over the North Pacific and adjacent regions provides predictive skill for forecasting at 2-3 week lead times. Finally, high amplitude circulation patterns and an increase in blocking frequency are associated with phase 3 of the MJO for the Atlantic region with a lead time of 10-15 days (Henderson et al, 2016).

5.6.1 Experiments

Using the same experimental design and data as detailed in this chapter, further analysis was carried out to investigate the robustness of the DBN structures shown here. CCFs were substituted with expert knowledge taken from the papers detailed above, which was used to determine appropriate lead times for the MJO and ENSO variables. Three datasets with differing lead times were constructed with this intention:

1. 10-day MJO lead time and 14-day ENSO lead time
2. 15-day MJO lead time and 14-day ENSO lead time

3. 30-day MJO lead time and 14-day ENSO lead time

The lead times were applied to daily data before preparation for the winter (DJF) months of the years 1981-2018. All data was sourced and prepared in exactly the same way as above for the 5-day averaged dataset: the spatio-temporal variables were spatially averaged, summarised as 5-day averages, the multi-year mean of 5-day averages was subtracted from the time series, and the result was then detrended (Table 5.1). 7-variable BNs were constructed using the BK, NAt, PoV, jet, NAO, MJO and ENSO variables listed in Table 5.1. For the sake of comparison, the previous lead times used were 45 days (9 timesteps) for the MJO and 155 days (31 timesteps) for ENSO at 5-day averages, based on the nearest significant value in the CCFs to jet latitude for each variable. The lead times identified in the literature that are used here are significantly shorter, especially for ENSO, adding to the need for further analysis based on expert knowledge.

The same training and testing dataset split was used (80/20), and BNs were run first with jet latitude and then with MI as the midlatitude circulation proxy, although the results for MI were identical in terms of the tropical variables so are not shown here. All models were constructed in the 'bnlearn' package in R (Scutari, 2010). As in Chapter 3's experimental analysis, the structure was learned with the PC algorithm using an alpha of 0.05, but networks were not extended to DBNs through parameter-learning and predictive accuracy analysis due to the lack of tropical links. For this reason, bidirectional arcs returned by the PC algorithm were also not corrected for as no further DBN construction was necessary. Results were robust to the use of two algorithms: the HC (hill-climbing) algorithm, a score-based algorithm detailed and used in Chapter 3, was used to check whether tropical links could be identified with hill-climbing optimisation.

5.6.2 Tropical Variables: A Lead Time Analysis

Partially directed acyclic graphs, or pDAGs, are shown for the three datasets used, whereby nodes are labelled as the corresponding variable, and relationships between variables represent conditional dependencies between nodes either in the form of arrows (directed arcs) or lines (undirected edges). As these pDAGs are essentially static BNs, they are comparable to the 'intra' structure findings detailed above, and no 'inter' structure needs to be defined as the networks were not unrolled to DBNs.

The structure between variables where the lead time was unchanged matches the findings for 5-day averaged jet DBNs shown above, before removal of cyclical structures (i.e. the jet to AA links). A positive feedback mechanism is again obvious for the AA variables (BK

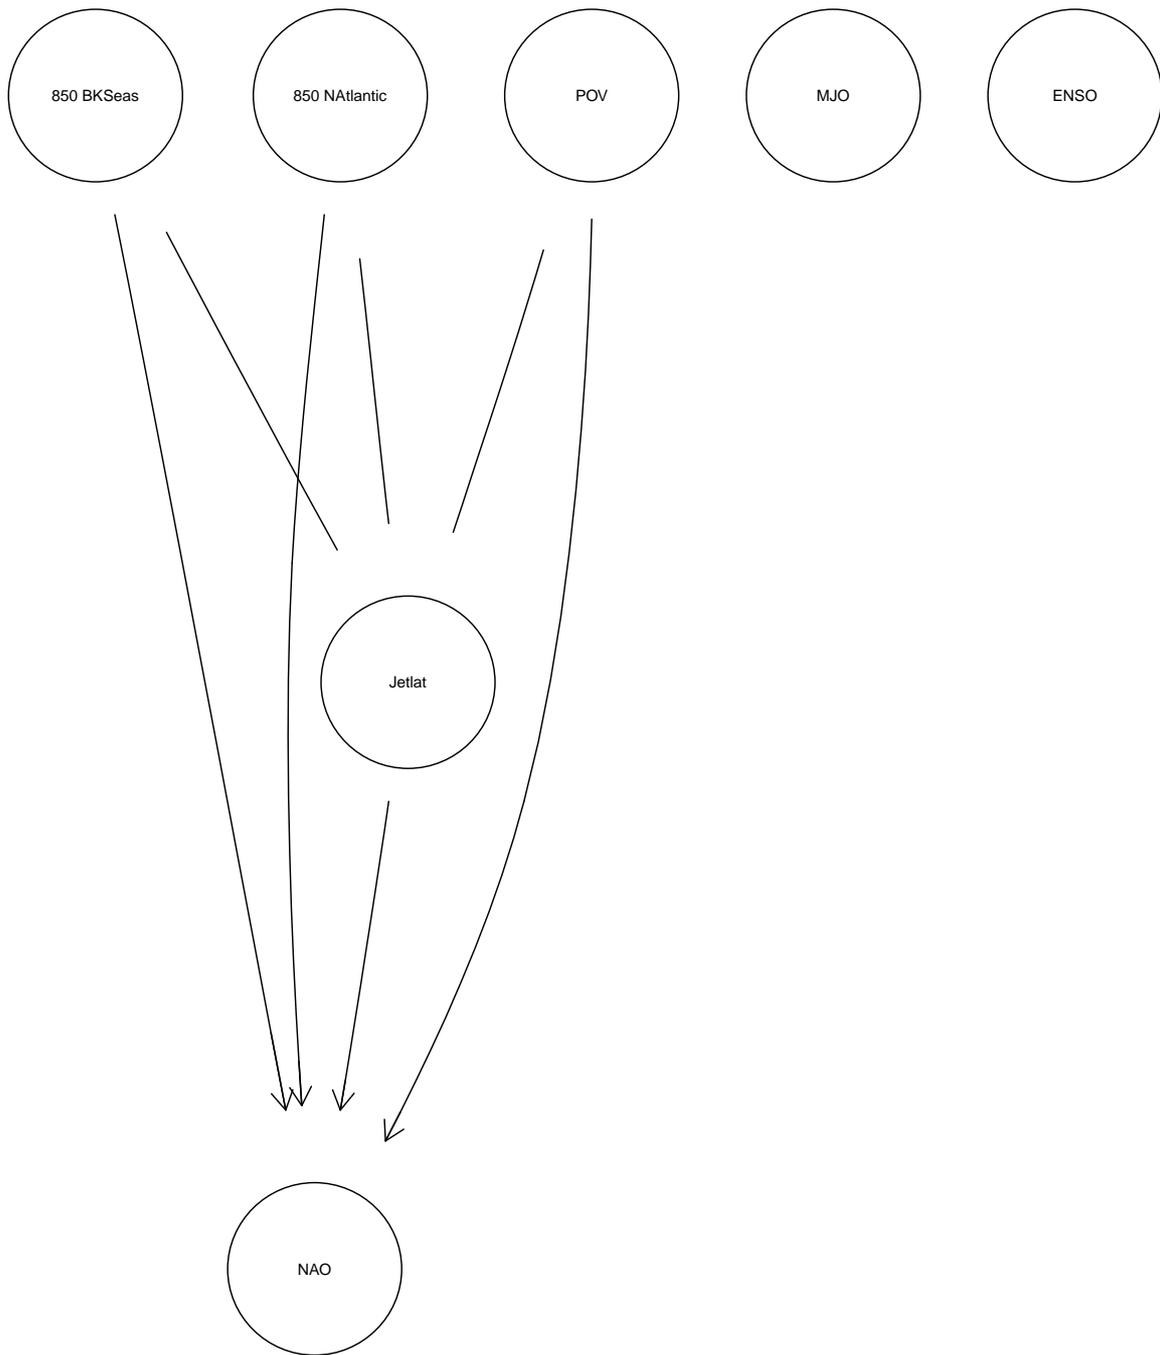


Figure 5.7 pDAG for experiments 1 and 2: 10- and 15-day lead time for MJO, 14-day for ENSO.

and NAt) and the jet latitude node, matching Samarasinghe et al's (2019) findings for North Pacific jet variability. The proxies of midlatitude circulation (jet latitude and the NAO) are again identified as child nodes. Just as above, however, the MJO and ENSO are completely unconnected from the rest of the graph structure for experiments 1, 2 (Fig. 5.7) and 3 (Fig.

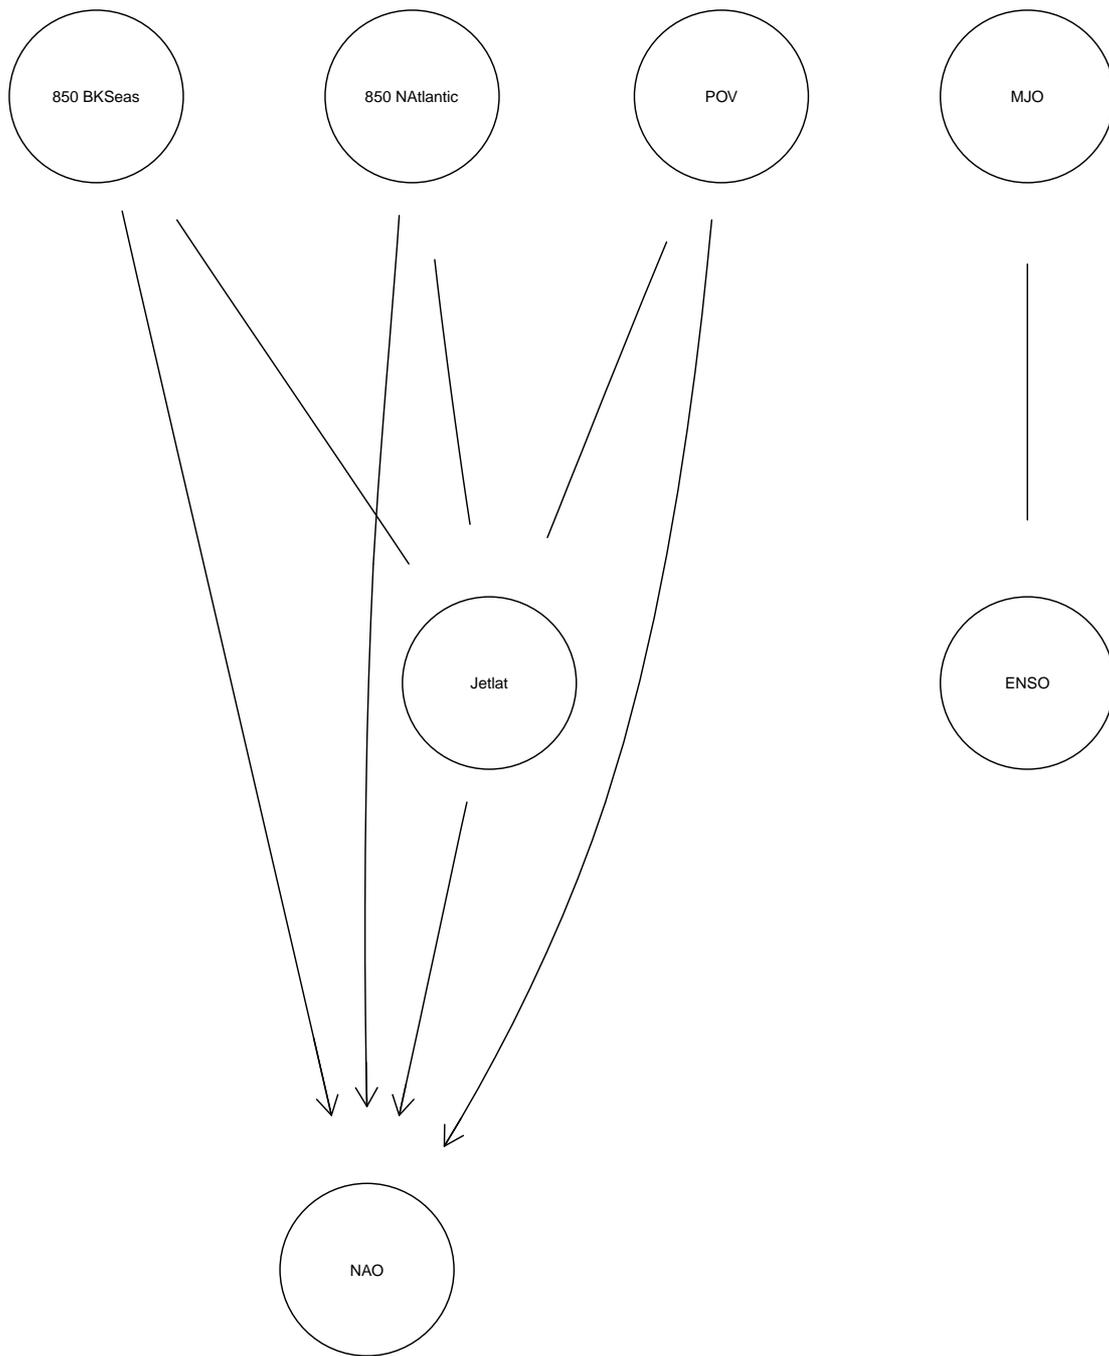


Figure 5.8 pDAG for experiment 3: 30-day lead time for MJO, 14 for ENSO.

5.8), where a bidirectional conditionally dependent relationship was found between the tropical variables lagged at 30 and 14 days for MJO and ENSO respectively.

It is clear from the pDAG results that the findings detailed above regarding the lack of an influence from the MJO and ENSO are not likely to be a fault of the methodology employed;

specifically, that the CCF results used to determine the lead times for all variables are not causing the DBNs produced to display unconnected tropical nodes. Possible reasons for this have been discussed above, but a likely cause is the time series-based approach used (DBNs) which may mask intermittent tropical teleconnections. Studies which have identified tropical influences often use forms of non-continuous (or 'broken') time series analysis to split the data up into weather events; Henderson et al (2016) for example use only days with an RMM amplitude greater than 1 to assess blocking frequency (as defined by Wheeler and Hendon, 2004), and Barnes et al (2019) discretise MJO data into active and inactive periods to investigate pathways between the MJO and NAO. If the absence of tropical teleconnections is a result of the continuous time series approach masking intermittent but significant linkages, it is important to note that DBNs consistently pick up the AA-midlatitude covariability. Arctic-midlatitude linkages are thought to be similarly intermittent in nature as they rely on the background jet stream pattern to act as a bridge between thermodynamic forcing and persistent midlatitude extremes (Overland et al, 2016; Overland and Wang, 2018; Kolstad and Screen, 2019). In this research, this did not prevent their identification at submonthly timescales.

5.7 Data Split Analysis

The above analyses use 37 years of data and an 80/20 split in training and testing datasets to investigate potential drivers of midlatitude atmospheric circulation. One of the problems associated with this approach is that potential linkages between cryospheric or tropical variables and midlatitude proxies are likely to be nonstationary in nature (Coumou et al, 2018). This has been suggested as the cause for the lack of a robust interdecadal link between ENSO and the Euro-Atlantic sector, where the ENSO signal is much stronger and can be found further poleward in the post-1970s period than before it (Greatbatch et al, 2004). This is also true for the lagged relationship between Eurasian snow cover and the AO, which only emerges as a plausible teleconnection in the observational record from the 1970s onwards (Douville et al, 2017).

The signal of Arctic Amplification itself has only separated from the trend for the rest of the Northern Hemisphere in recent decades (Fig 1.1); enhanced warming of more than six times the global average has occurred since the late 1990s (Huang et al, 2017). Francis and Vavrus (2015) note that the AA signal is distinguishable from the noise of internal variability from approximately 1995 onwards for surface processes and from 2000 onwards for the lower troposphere, leading them to describe an 'AA era' as opposed to a pre-AA era in which natural variability dominates. Given that this is the case, teleconnections between Arctic and midlatitude variables are likely to be highly nonstationary, and the linkages identified may be highly dependent on the start and end dates of the period used as input for DBN analyses.

Kolstad and Screen (2019) point out that nonstationarity may have an impact on the reliability of all machine learning applications; the predictive skill of NAO forecasts using dynamical models varies noticeably over time for example, so the NAO signal itself may be the result of nonstationary linkages.

5.7.1 Experiments

To test the extent to which the linkages identified in the above analyses were nonstationary in nature, the dataset used above was split into a 'pre-AA' and 'AA' period and the experiments were re-run. The same dataset, sourced and prepared as detailed above, was used: 7-variable DBNs were constructed using 5-day averages of the BK, NAt, PoV, jet, NAO, MJO and ENSO variables for the winter months (DJF) of the years 1981-2018 (Table 5.1). The dataset was then split into two half periods:

1. Pre-AA – 1981.12.03 to 2000.01.12 to capture the period before a strong AA signal develops
2. AA – 2000.01.17 to 2018.02.26 to capture the period of AA identified in the observational record

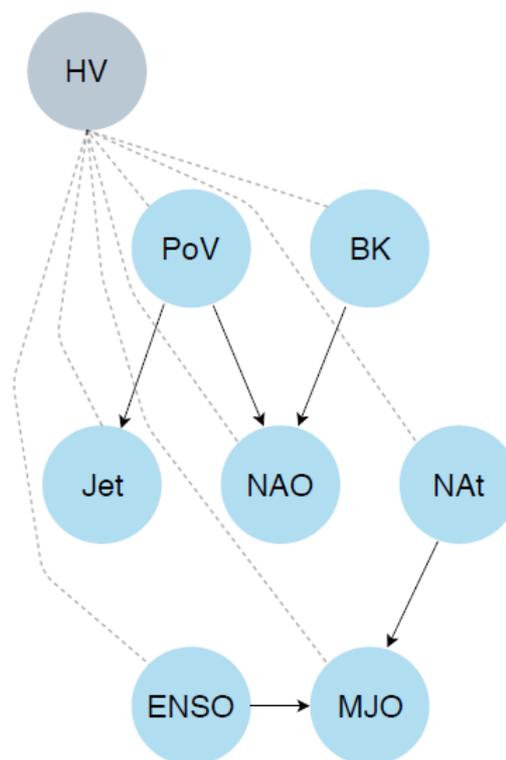


Figure 5.9 DAG of the Pre-AA DBN (1). Dotted lines indicate edges coded into the model (i.e. HV edges), solid represent learned edges.

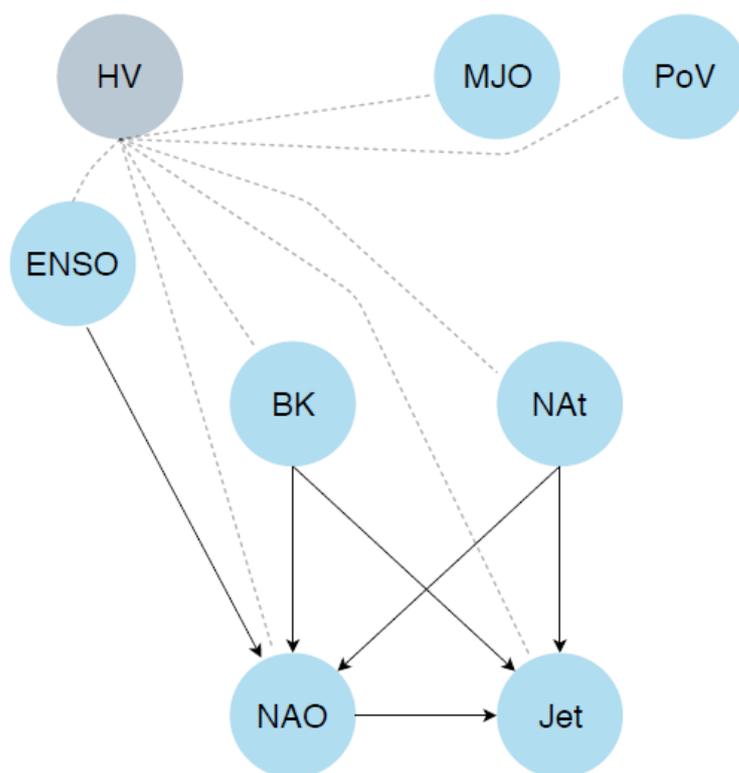


Figure 5.10 DAG of the AA era DBN (2). Dotted lines indicate edges coded into the model (i.e. HV edges), solid represent learned edges.

DBNs were constructed, again using an 80/20 split whereby the network structure and parameters were learned using 80% of the data, and the model fit was tested on the remaining 20%. Splitting the data resulted in fewer timesteps, but the total number in each dataset was equal to the 10-day averages (333 timesteps) shown above, which managed to predict most variables accurately for the jet DBNs with a HV. Network structures were inferred from the data using the PC algorithm using the *'fisher z'* test for conditional independence and an alpha value of 0.01. Results were dependent on the choice of alpha value as significant overfitting of the networks occurred with alpha values of >0.01 . This contrasts the 5-day jet and MI DBNs, which had twice the number of timesteps for the full time series (1981-2018). Bidirectional arcs were removed from the learned structure to preserve the 'explain-away' effect, as above, to create an acyclical structure and run the DBN. The EM algorithm was then used to fit the parameters of the model, and the DBN was tested on the test dataset to examine the degree of overfitting contained within the model. As a final step, model validation was achieved by plotted predictive DBN output against the observed values for the test dataset, and SSE was used to assess model performance in terms of each variable.

5.7.2 Results: Before and During the AA Period

The DAGs produced for 5-day DBNs in the pre-AA (1981-2000) and AA (2000-2018) periods suggest that splitting the data in this way changes the DAG structures in fundamental ways that are similar to averaging the data over different time resolutions. The pre-AA DAG (Fig. 5.9) differs greatly from the 5-day Jet DBN above (Fig. 5.3): all but one (BK-NAO) of the arcs that make up the Arctic-midlatitude feedback relationship are not identified, whilst the links between the PoV and jet latitude and NAO variables are still captured. The structure is therefore most similar to the monthly-averaged jet DBN results (Fig 5.4d).

For the DBN trained on data during the AA period however, the same structure shown by the 5-day jet DBNs (Fig. 5.10) can be seen for the BK, NAt, NAO and jet variables, although unlike Fig. 5.3 the jet node arcs (BK-jet, NAt-jet) were not bidirectional. Fig. 5.10 suggests that the BK and NAt regions can influence North Atlantic circulation at 5-day timescales within the winter months (DJF). No feedback relationship was identified, suggesting that DBNs trained only on the AA period do not capture the thermodynamic effect of temperature and moisture advection into the Arctic driven by high amplitude jet stream configurations (Kim et al, 2017; Cohen et al, 2018a). There is no climatological reason why the Arctic variables should be identified only as parent nodes or drivers with the AA dataset in contrast to the feedback mechanism between the AA and jet variables identified with the full 37-year dataset. Recent work using atmospheric models has indeed pointed to the reverse scenario as more likely for colder midlatitude winters: by attempting to reconcile conclusions drawn from model and observational analyses, Blackport et al (2019) suggest that the covariability between sea ice and midlatitude temperatures occurs due to anomalous midlatitude circulation patterns acting as a driving force of sea ice variability rather than a response to it. Whilst this chapter suggests a covariability between Arctic near-surface temperatures and midlatitude circulation proxies for the North Atlantic, rather than regional temperature anomalies, several recent studies highlight the importance of midlatitude-Arctic linkages (Woods and Caballero, 2016; Kim et al, 2017; Ye and Jung, 2019) and caution against ubiquitously considering Arctic variables as drivers of midlatitude circulation anomalies (Cohen et al, 2018a; Fyfe, 2019).

ENSO is identified as a driver of NAO variability in the AA period DBN, in contrast to all other DAG structures presented in Chapters 4 and 5. As already discussed, ENSO can influence the North Atlantic midlatitude region via the troposphere through planetary wave initiation (Scaife et al, 2017b) and via the stratosphere through the disruption of the stratospheric polar vortex (Baldwin and Dunkerton, 2001; Hardiman et al, 2019). The ENSO-NAO connection identified here is a long-duration teleconnection indicating that late June ENSO variability can impact the early December NAO signal, as Hall et al (2019) found for jet latitude anomalies. The ENSO-NAO link may have appeared only in this DBN due to the

shorter time series used for the AA period (2000-2018); the ENSO influence on the Euro-Atlantic sector is known to be nonstationary in nature (Greatbatch et al, 2004), meaning that time series-based approaches using multi-decadal datasets are unlikely to pick up ENSO teleconnections.

It should be clarified that to fully eliminate the possibility that internal variability created the differences between the two periods, further processing like random sampling 19-year periods across the AA and pre-AA periods would be necessary to determine how large the differences in networks can be outside of the impact of AA. Basing such conclusions solely on a 19-year period would be misleading given events like the Northern Hemisphere continental cooling which occurred during the 1930s-1940s; a pattern that reversed in the following decades (Wegmann et al, 2018). The lack of an Arctic-midlatitude covariability in Fig. 5.10 is evidence that internal variability may be having an impact on network structure, as this covariability is the most consistent finding in the networks of chapters 4 and 5. Caution is therefore advised in the interpretation of this result, as further analysis is needed to quantify the role of internal variability.

	pre-AA	AA 1	AA 2
850 BKSeas	57.21	40.72	45.24
850 NAtlantic	53.42	41.85	45.54
POV	31.9	11.76	9.49
Jetlat	34.88	416406.73	13.34
NAO	46.74	12564.6	22.63
MJO	6134.95	51.68	54.02
ENSO	65.58	22.44	19.14

Table 5.4 SSE for each variable in the pre-AA DBN, AA DBN and AA DBN with arc deletions.

5.7.3 Model Performance in the AA period

The model fit was assessed for all networks using SSE (Table 5.4), as in the above analyses. The MJO node was badly predicted by the pre-AA DBN, suggesting that the NAT-MJO link shown in Fig. 5.9, which does not appear in any other DBN structure, may not be valid as the NAT node provides low predictive skill as a parent to the MJO. For the AA period ('AA 1'), the jet latitude and NAO variables were predicted with significantly less accuracy than the rest of the variables. To investigate this, the ENSO-NAO and NAO-Jet arcs were deleted from the DAG structure (Fig A5.9), and the parameters and SSE were recalculated for the resulting modified DBN named 'AA 2' in Table 5.4. The SSE of both variables increased significantly, suggesting that these arcs drastically reduced the predictive skill of the DBN trained only on the AA period, in particular the midlatitude circulation indicators. The ENSO-NAO teleconnection identified here may therefore be a spurious link, as the midlatitude components were predicted more accurately during the AA period with only the AA variables as parent nodes.

5.8 Conclusion

This study represents the first foray into Arctic-midlatitude weather linkages using a graphical model approach paired with hidden variables. A robust covariability between regions of amplified Arctic warming and midlatitude circulation characteristics is found, suggesting this feedback has a significant influence on variability both in the Arctic and the North Atlantic midlatitudes at submonthly timescales. The two-way nature of the link suggests uni-directional interpretations need revising in favour of one that takes into account the central importance of poleward heat and moisture fluxes into the Arctic. Of the relationships found, BK-NAO has the strongest impact on the DBNs at 5- and 10-day intervals, with the PoV-jet link becoming the strongest at monthly resolutions. Midlatitude circulation responses at submonthly timescales are driven by AA processes within winter rather than a lagged response to sea ice losses during autumn. Links to tropical modes of variability were not identified using a range of lead times and methods, but this in no way implies that their impact on observed jet variability is small. The DBNs found no evidence for a clearly defined 'AA' and 'pre-AA' period through hidden variables included in model architectures.

Clearly, these results only give us part of the picture. Stratospheric polar vortex variability dominated the HV state switches, shifting from average conditions to slower anomalous PoV flow, higher NAT temperatures and slightly higher jet latitude values. The HVs hint that strength estimates may underrepresent the importance of the PoV, with no lagged connections to midlatitude flow identified through the stratosphere. The findings focus on submonthly timescales and as such may miss stratospheric links thought to be important for North Atlantic circulation (Scaife et al, 2017a; Hardiman et al, 2019). As a form of time series

analysis, the DBNs may not have picked up state-dependent linkages or combined effects like those shown in simulation analyses (e.g. Lee et al, 2015), which could have masked potential tropical influences like ENSO-PoV. Furthermore, input data was averaged over large areas chosen to maximise the amount of variability captured, which may not represent the true complexity of midlatitude circulation influences.

A number of improvements are made from the 4-variable DBN analysis in Chapter 4, including the use of more data from Arctic, midlatitude and tropical sources, a proper investigation into appropriate lead times for all variables, an analysis of the hidden state values and their impact on the networks, and finally a study of the AA era and the impact it has on the relationship between Arctic and midlatitude linkages. The lack of a robust tropical influence on the DBNs was found not to be an artefact of the method used to ascertain lead times for variables, as the use of expert-guided lead times for the MJO and ENSO did not change the DAG structure results. Splitting the dataset into a pre-AA and AA period results in a pre-AA DAG that is similar to the monthly-averaged jet DBNs, and an AA DAG that reproduces the 5-day jet DBN Arctic-midlatitude structure without bidirectional links indicating that the North Atlantic jet influences AA processes on 5-day timescales within winter. The reason for this is unclear, but the disparity between the two periods implies that the Arctic-midlatitude linkages shown here for the North Atlantic region are likely to be nonstationary in nature (Coumou et al, 2018; Kolstad and Screen, 2019). This could be a result of the recent appearance of the AA signal in the observational record (Huang et al, 2017), or intermittent teleconnections between the Arctic and the midlatitudes that are dependent on the background atmospheric flow characteristics as a bridge between thermodynamic forcing and amplified midlatitude flow (Overland et al, 2016).

Chapter 6: Conclusion

Summaries of the results and their implications for the research field have already been discussed in the conclusions of the preceding three chapters. This chapter therefore focuses on the contribution of this thesis to the Arctic-midlatitude field and wider climate science research area. This is followed by an assessment of the limitations of the method and study design, and finally a recommendation for future work to tackle a number of gaps in the research field with suggestions for a potential way forward.

6.1 Research Contributions

6.1.1 Dynamic Bayesian Networks with Hidden Variables

The Arctic-midlatitude research area is constrained by a low signal-to-noise ratio of AA to internal variability, and intermittent linkages that may be dependent on the state of the background atmospheric flow in the midlatitudes (Overland et al, 2016). This thesis has presented evidence for a robust wintertime link between important regions of AA and midlatitude circulation on submonthly timescales using Dynamic Bayesian Networks with a single hidden variable. The results are suggestive of a feedback loop between enhanced Arctic warming at the near-surface, and large-scale atmospheric circulation anomalies. In terms of relative strength, the BK-NAO link was the strongest in DBNs at 5- and 10-day intervals, whilst the PoV-jet link seems to be the most important at monthly resolutions. AA-midlatitude covariability occurs entirely within winter (DJF) when the AA signal is strongest (Serreze et al, 2009) in contrast to other studies which find a lagged response to autumn sea ice loss (Kretschmer et al, 2016; Hall et al, 2017), although longer-duration autumn-to-winter DBNs were not specifically constructed. This was because the CCF analysis revealed no long-duration linkages between AA regions and jet variability as the BK and NAt variables were strongly correlated to jet latitude only within a few timesteps (i.e. within winter). On this basis, DBNs with structure-learning algorithms prove to be a useful tool for investigating complex, nonlinear climate teleconnections in a network which includes Arctic, midlatitude and tropical variables, and may prove to be a valuable alternative to correlation-based networks and studies (Ebert-Uphoff and Deng, 2012a).

The application of DBNs and HVs to the Arctic-midlatitude research field also revealed a number of important caveats, which may be pertinent to future graphical model studies. Firstly, when DBNs are used to model climate relationships with a high number of edges and a low number of input variables (≤ 4) they can be over-parameterised, which has a significant impact on model performance and error rate as Chapter 4 shows. Furthermore, the addition of HVs to graphical models does not guarantee that model accuracy will increase. HVs were

added to dynamic networks to optimise model performance and maximise the likelihood that the AA signal would be picked up in the learned network structures. The MI DBNs with a HV in Chapter 5 performed slightly worse than the control run with no HV. The variability in HV state switches was dominated by the stratospheric polar vortex shifting from average to anomalously slow flow. This result, combined with the relative strength analysis which suggests the PoV-jet link is the strongest in monthly jet DBNs, strongly hints that the Chapter 5 results underrepresent the importance of stratospheric variability which may predominantly be driven by tropical influences. In summary, researchers using HVs need to ensure that graph structures returned by a DBN are as accurate as possible by conducting extensive sensitivity analyses. This may include multiple structure-learning algorithms, different proxies of similar physical mechanisms, a range of time resolutions and lead times, and data averaged over different regions.

6.1.2 The Relative Impact of AA on Midlatitude Circulation

Whereas other graphical model studies focus on potential Arctic drivers (Kretschmer et al, 2016; Barnes and Simpson, 2017; Samarasinghe et al, 2019) or possible tropical teleconnections (Barnes et al, 2019) using models with two variables, this project proves that DBNs can accurately model relationships between climate datasets with larger numbers of variables. In further contrast to these studies, a number of Arctic, midlatitude and tropical variables are included which provide insight into the relative role of AA processes. The lack of a robust tropical influence in DBNs, whilst likely not to be a result of a flaw in methodology, may suggest that teleconnections to tropical indices will not be found using a DBN approach based on continuous timeseries due to intermittent or state-dependent linkages. Arctic-midlatitude linkages are also likely to be intermittent in nature and dependent on background flow characteristics (Section 5.7; Kolstad and Screen 2019) yet are consistently picked up in DBNs at submonthly timescales in chapters 4 and 5. This could reflect the importance of poleward heat and moisture fluxes into the Arctic with amplified jet stream patterns and blocking regimes (Kapsch et al, 2016; Yang and Magnusdottir, 2017) rather than an interpretation that favours the unidirectional influence of AA processes on jet stream configurations (as in Francis and Vavrus, 2015). Support for this hypothesis is found in the fact that bidirectional AA-midlatitude links are found in all the networks of chapters 4 and 5 (with the sole exception of section 5.7.2).

For subseasonal-to-seasonal atmospheric circulation prediction, sea ice concentration provides predictive skill for the winter NAO in both atmospheric (Scaife et al, 2014) and statistical models (Hall et al, 2017; 2019; Wang et al, 2017). Considerable uncertainty remains surrounding the relative importance of AA as a driver, however, and DBNs could play a part in identifying regions of AA that provide predictive skill. Regions of sea ice loss, like the

Barents-Kara Seas, can provide predictive skill from the preceding autumn even when no direct causal link can be found in atmospheric modelling studies (Blackport and Screen, 2019). Links between the Barents-Kara region and the NAO are found for up to monthly-averaged timescales in Chapter 5. Furthermore, the North Atlantic and Barents-Kara regions of the Arctic and stratospheric polar vortex variability are found to be robustly linked to midlatitude circulation at submonthly timescales. The proportion of variance analysis conducted in Chapter 5 may point to a difference in submonthly to monthly drivers; at submonthly timescales, the Barents-Kara Seas and NAO (and thus the 'phase' or speed and latitudinal position of the jet stream) share a strong interdependence, whilst at monthly averages the stratospheric polar vortex has a greater influence through flow weakening and downward wave propagation which projects onto a negative NAO pattern. This suggests two paths forward for analysis: firstly, establishing whether the BK-NAO covariability is simply a manifestation of internal variability as suggested in Warner et al (2020), and secondly to quantify the relative importance of stratospheric flow influences which include tropical (e.g. Scaife et al, 2017b) and midlatitude (e.g. Kim et al, 2014) sources.

6.2 Study Limitations

DBNs with a single HV were found to have a relatively high predictive accuracy when used to investigate climate teleconnections, picking out informative relationships between large areas of climatic variability. As with other statistical approaches and atmospheric models however, they are limited by some drawbacks.

- **Data Quality** The biggest problem with graphical model approaches as applied to climate data is the averaging of data over large areas, whereby the weighted spatial average over a 'gridbox' of latitude and longitude is taken. All non-index data (spatio-temporal NetCDF files) was processed in this way by necessity, to produce 1-dimensional time series for BN analysis. Some datasets, most noticeably Eurasian snow cover extent which is used in several important studies (Kretschmer et al, 2016; Hall et al, 2017), are averaged over regions that cannot possibly be informative in terms of the complex patterns of variability they are designed to capture at daily resolutions. As detailed in Chapter 4, the gridboxes used for snow cover meant that the method used for calculating anomalies rendered these datasets useless as indicators of snow variability. Whilst this did not seem to be the case for other datasets, care must be taken when averaging over such large regions.
- **Stratospheric Drivers** No long-lead (>10 days) linkages were identified for the variables included in BNs. The submonthly timescale focus of this research may miss important long-lead relationships, especially those thought to act on midlatitude extremes through the stratosphere (Kim et al, 2014; Kretschmer et al, 2016). None of

the variables were connected to midlatitude circulation through the stratosphere, so the timescale focus may have biased the results in favour of tropospheric linkages, and could be part of the reason that no tropical influences were found.

- **Structural Uncertainties** The ideal reconstruction of the North Atlantic midlatitude climate system is an impossible task, at least using observed data and a network approach. Whilst this project attempted to represent as many potential drivers of Euro-Atlantic midlatitude circulation as possible, results are hampered by noisy data and the complexity of the climate system being modelled, and the selection of variables may have relied too heavily on the literature and expert knowledge.
- **Edge Deletion** Finally, the results of the structure-learning phase with the PC algorithm, whilst interesting from a climatological perspective, meant that the DBN code would not run as the graph was not fully directed. Bidirectional arcs were removed to turn the pDAG produced into a DAG to allow the DBNs to run. This was done manually to ensure that collider structures were maintained. Ebert-Uphoff and Deng (2012a) also used this approach to establish edge direction and create fully directed graphs. 'bnlearn' (Scutari, 2010) has an operator that can infer the direction of an undirected edge based on the lowest p -value, the highest score or highest bootstrap probability, which could be useful if the pDAG does not contain collider structures that make direction choice obvious for the user. Studies that are intent on proving causality with a graphical model approach use multiple lagged copies of the same dataset to use conditional independence tests as indicators of causal relationships (e.g. Barnes et al, 2019). Whilst this method has other problems, direction is clear in the structure results as the express purpose is to look for drivers of response variables.

6.3 Future Work

6.3.1 Recommendations

A number of recommendations for future work seem appropriate given the findings and limitations of the research presented here. To allow for the inclusion of large-scale variables like Eurasian and North American snow cover extent, and to increase the complexity of variable selection, the Response-Guided Causal Precursor Detection (RGCPD) scheme (Kretschmer et al, 2017; Di Capua et al, 2019) could remove the need for large gridboxes of spatial data. RGCPD detects regions of spatial data that drive variability in a response variable, for example the jet stream. A number of tropical datasets could then be identified as causal precursors, potentially including sea surface temperatures (SSTs) and precipitation (Hall et al, 2017) or outgoing longwave radiation (OLR). Data, especially SSTs, may be too noisy for any significant relationships to be picked up. Global gridded precipitation datasets, which could be used as a proxy for tropical variability if detected as a causal precursor by the

RGCPD method, have a number of inconsistencies detailed in Sun et al (2017) which limit their reliability for use in climate teleconnection studies and model validation. RGCPD must also be used in conjunction with expert knowledge, and care must be taken when using it to identify regions of interest as it can return spurious correlations.

Further study looking to build on this analysis would also benefit greatly from better automation of the lead and lag process. Whilst bidirectional edges were found for AA and midlatitude variables throughout this study in the structure-learning phase of networks, a single direction was required for all of them to function as DBNs and perform inference to quantify model accuracy. Running a number of lead and lag times for variables could potentially establish direction; for example, regions of AA could hypothetically impact the jet at short lead times with the jet in turn becoming a driver of AA at short lag times. Indeed, CCF plots may have captured this as the maximum lag value in Fig. 5.2 (i.e. 0 lag), resulting in bidirectional edges. In reality, these kinds of results can be rare in graphical model analyses and users should make sure they are not tampering with parameters until the results meet expectations. Full automation of the lead and lag process however could have removed some of the longer-duration linkages like the ENSO-jet maximum correlation (-155 days, see section 5.3) which likely do not reflect a physical teleconnection in light of the 2-week North Atlantic midlatitude circulation response found in other studies (Baldwin and Dunkerton, 2001; Scaife et al, 2017a; Hardiman et al, 2019; see also section 5.6 above).

A simple extension of this project could be the addition of more variables. Whilst problems with overfitting occurred with relatively small numbers of variables (4 or 5 in Chapter 4), the 7-variable networks of Chapter 5 performed accurately with a single hidden variable. Datasets missing in this study, like the QBO and multi-decadal modes of natural variability, could be included in future analyses (see Smith et al, 2016; Hall et al, 2015; 2017) but the near-term focus must be the identification of a tropical signal using graphical model approaches to fully measure the relative contribution of remote drivers in observational data. More complex networks with a larger number of variables like the 12 predictor datasets used in Hall et al (2017) are possible using a DBN approach, as Trifonova et al (2015; 2017; 2019) show.

Another simple variation on the research presented here would be the use of contiguous (non-continuous) timeseries. Overland and Wang (2018) suggest finding climatological explanations for individual cold midlatitude, anomalously warm Arctic events for example, rather than investigating the seasonal averages of indices. Individual events, or contiguous timeseries of similar events over the past 40 years for example, can be analysed in a DBN setup where individual arcs found in a DAG could be quantified in terms of their strength or score contribution if using 'bnlearn' (Scutari, 2010).

Datasets can also be manipulated in simple ways to reveal potential linkages between midlatitude circulation proxies and remote drivers. Discretising datasets, as Barnes et al (2019) do for the NAO, stratospheric polar vortex and MJO indices, significantly reduces noise which can lead to more relationships between variables in the structure-learning phase. Whilst this may work in the case of the MJO, where amplitude can be summarised to turn it into a binary variable where MJO forcing is either 'on' or 'off', this may be less applicable and even inadvisable for other metrics like jet stream latitude.

6.3.2 A Way Forward

This study can be seen as part of a growing effort to understand the Arctic contribution to jet variability using machine learning approaches (Kretschmer et al, 2016; 2017; Barnes and Simpson, 2017; Samarasinghe et al, 2019; Francis et al, 2018; Barnes et al, 2019). Graphical models paired with structure learning algorithms prove to be a robust tool for investigating complex climate teleconnections in a network which includes Arctic, midlatitude and tropical variables, and may provide a valuable alternative to correlation-based networks and studies (Ebert-Uphoff and Deng, 2012a). DBNs with a single HV were found to have predictive accuracy, subject to a number of important caveats; their addition does not guarantee network accuracy will increase. Network performance was high despite the noisy internal variability of North Atlantic flow and the low signal-to-noise ratio of AA-midlatitude covariability, suggesting any nonlinearity in linkages did not substantially interfere with accurate node prediction in the test dataset. HVs should be considered for inclusion in leading graphical model packages like 'bnlearn' (Scutari, 2010) in R and 'TiGraMITe' (Runge et al, 2015) in the python language; both are modern, fully open source packages and written in open source languages. Studies using them could then increase the reproducibility of their results by uploading the code online. The Bayes Net Toolbox (Murphy, 2001b), solely due to being written in MATLAB, prevents this option for this project. Open source languages and packages are generally considered good practice for scientists to publish in a transparent manner.

Finally, further study should continue to use a definition of AA that captures its full signal and include non-Arctic drivers in analyses with methods suitable for non-linear linkage detection (Cohen et al, 2018a; Overland and Wang, 2018). Sea ice does not adequately capture the full AA signal, so observational and atmospheric model approaches should move beyond simply using sea ice to drive midlatitude circulation variability or extreme surface temperature anomalies. This project has found robust teleconnections using near-surface temperature to represent AA variability, inspired by Barnes and Simpson (2017), suggesting that this might be a more appropriate AA metric.

Atmospheric models will be the primary method through which robust teleconnection mechanisms central to midlatitude circulation and extreme weather will be found, due to the large date ranges simulations can be run for and their ability to reliably detect causal relationships. Coordinated model experiments using atmosphere-ocean coupled models are therefore of great importance to the progression of the Arctic-midlatitude field and the establishment of a consensus regarding the relative importance of Arctic processes, and a number of modelling groups are already working towards this end (Cohen et al, 2018a). DBNs are found to be a consistent and reliable method for analysing climate teleconnections, especially when good practice is observed regarding the use of structure-learning and hidden variables detailed above. Further study seeking to apply DBNs should note they can suffer from over-parameterisation and will lead to inaccurate variable prediction where graph structure deviates significantly from the climate system we are trying to model, as shown in the jet-PoV links in Chapter 5 results. Graphical models would prove themselves more than capable in a supporting role as the results from large coordinated modelling projects are published in the near future.

Appendix

A1 Chapter 3 Figures

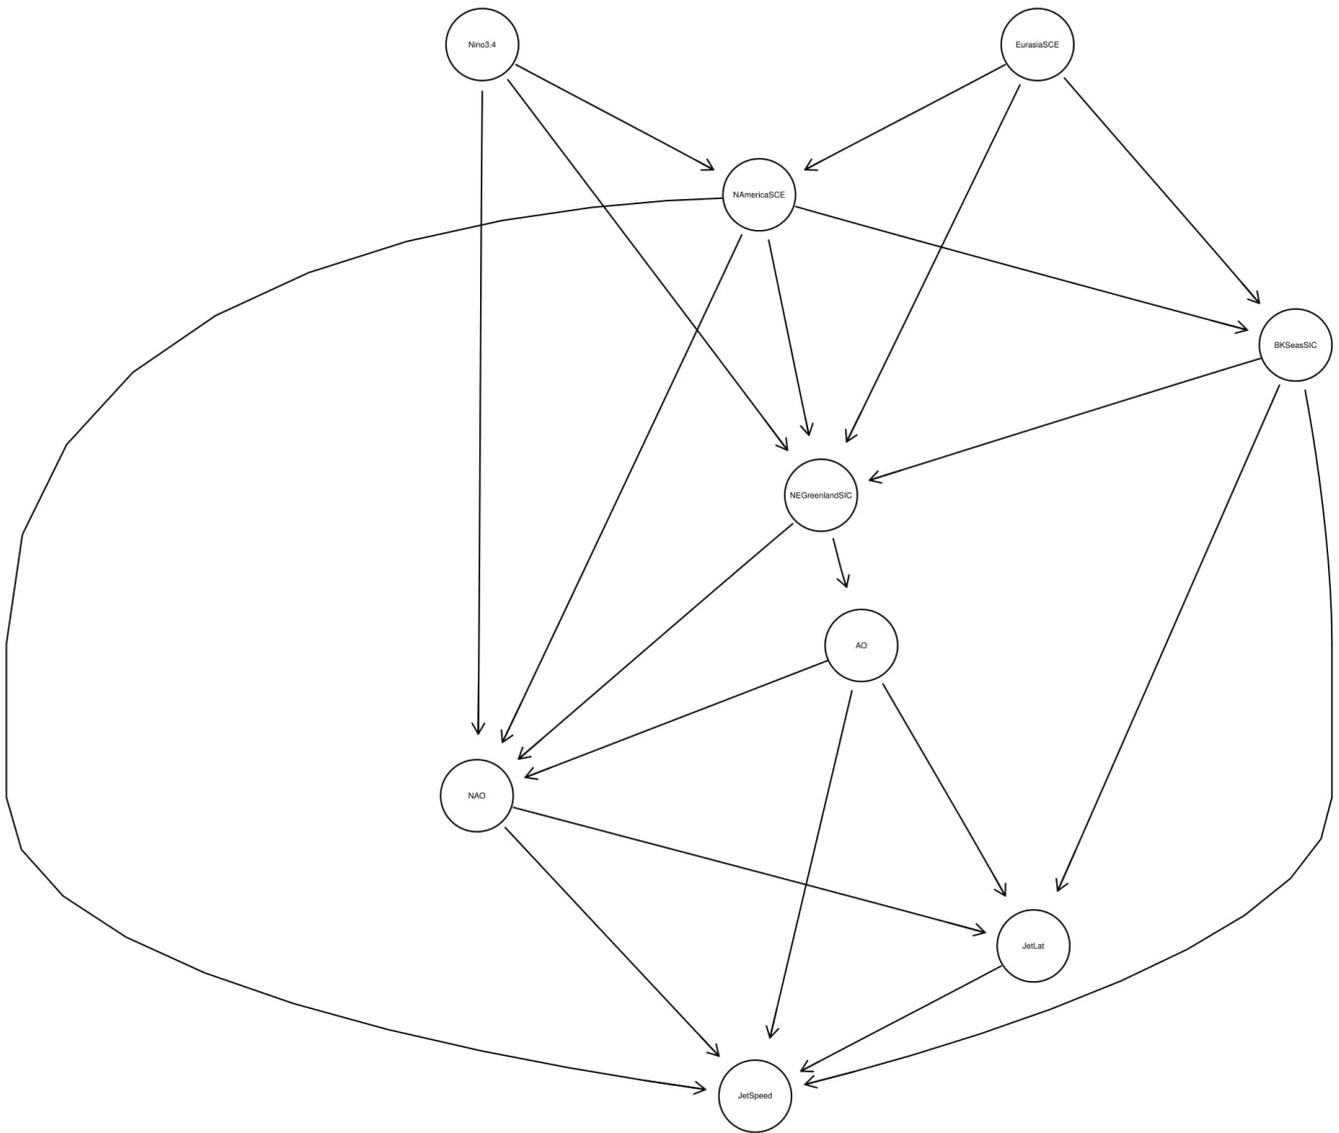


Figure A3.1 Daily 35-Year DAG K=7.

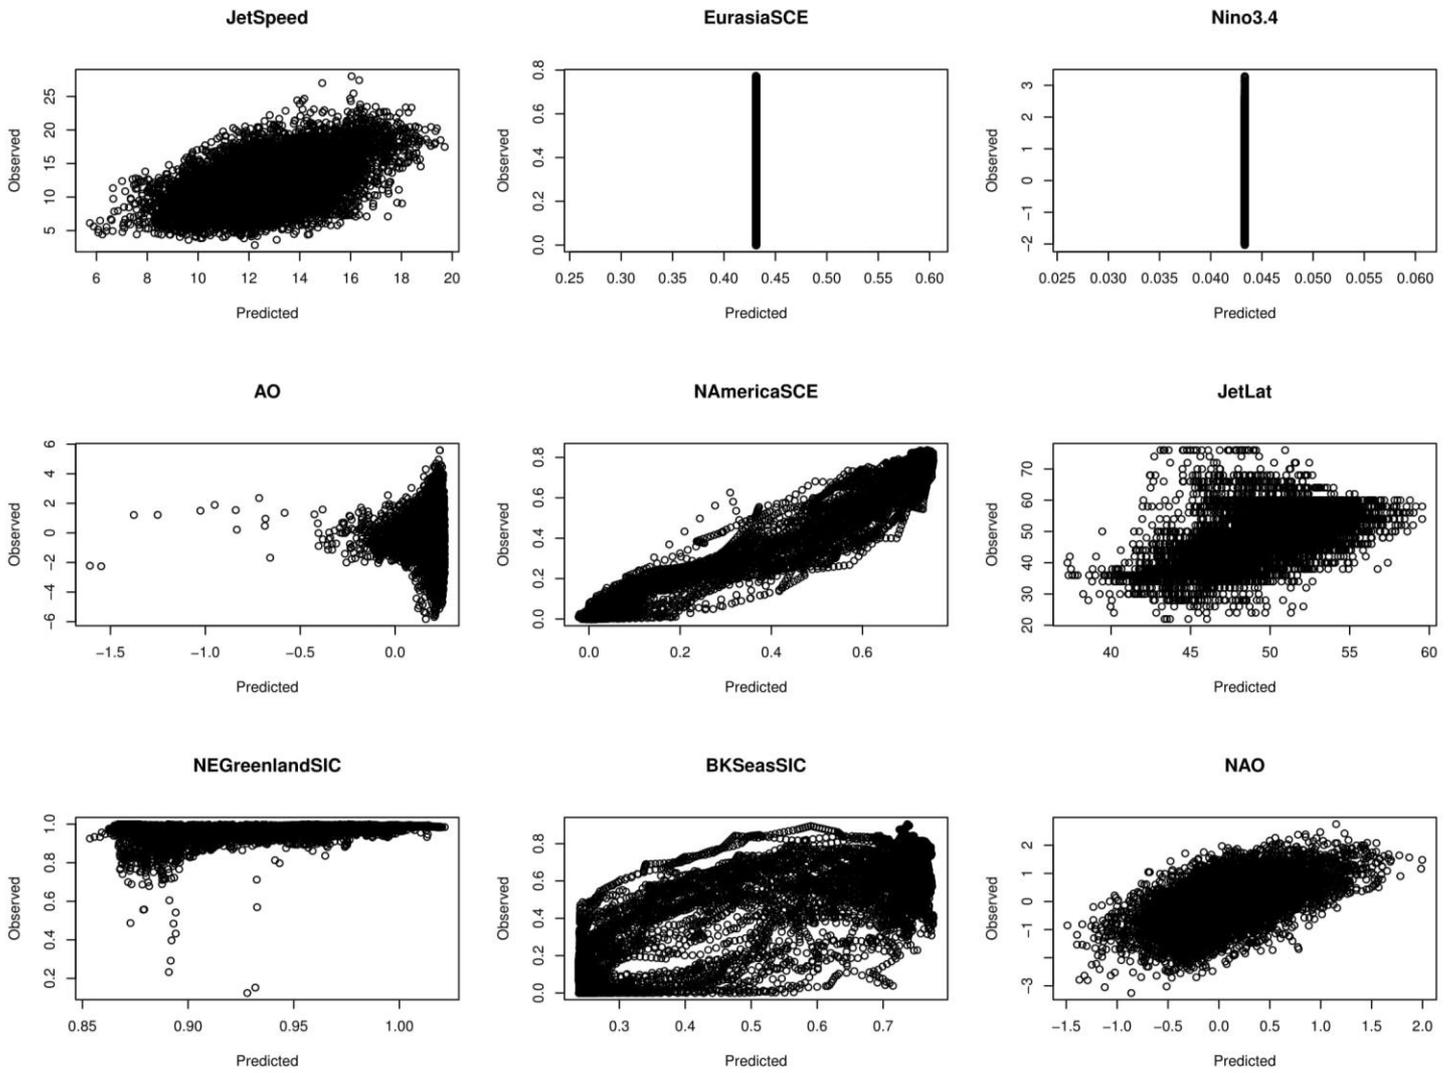


Figure A3.2 Predicted fit for all variables in Daily 35-Year K=7 BN.

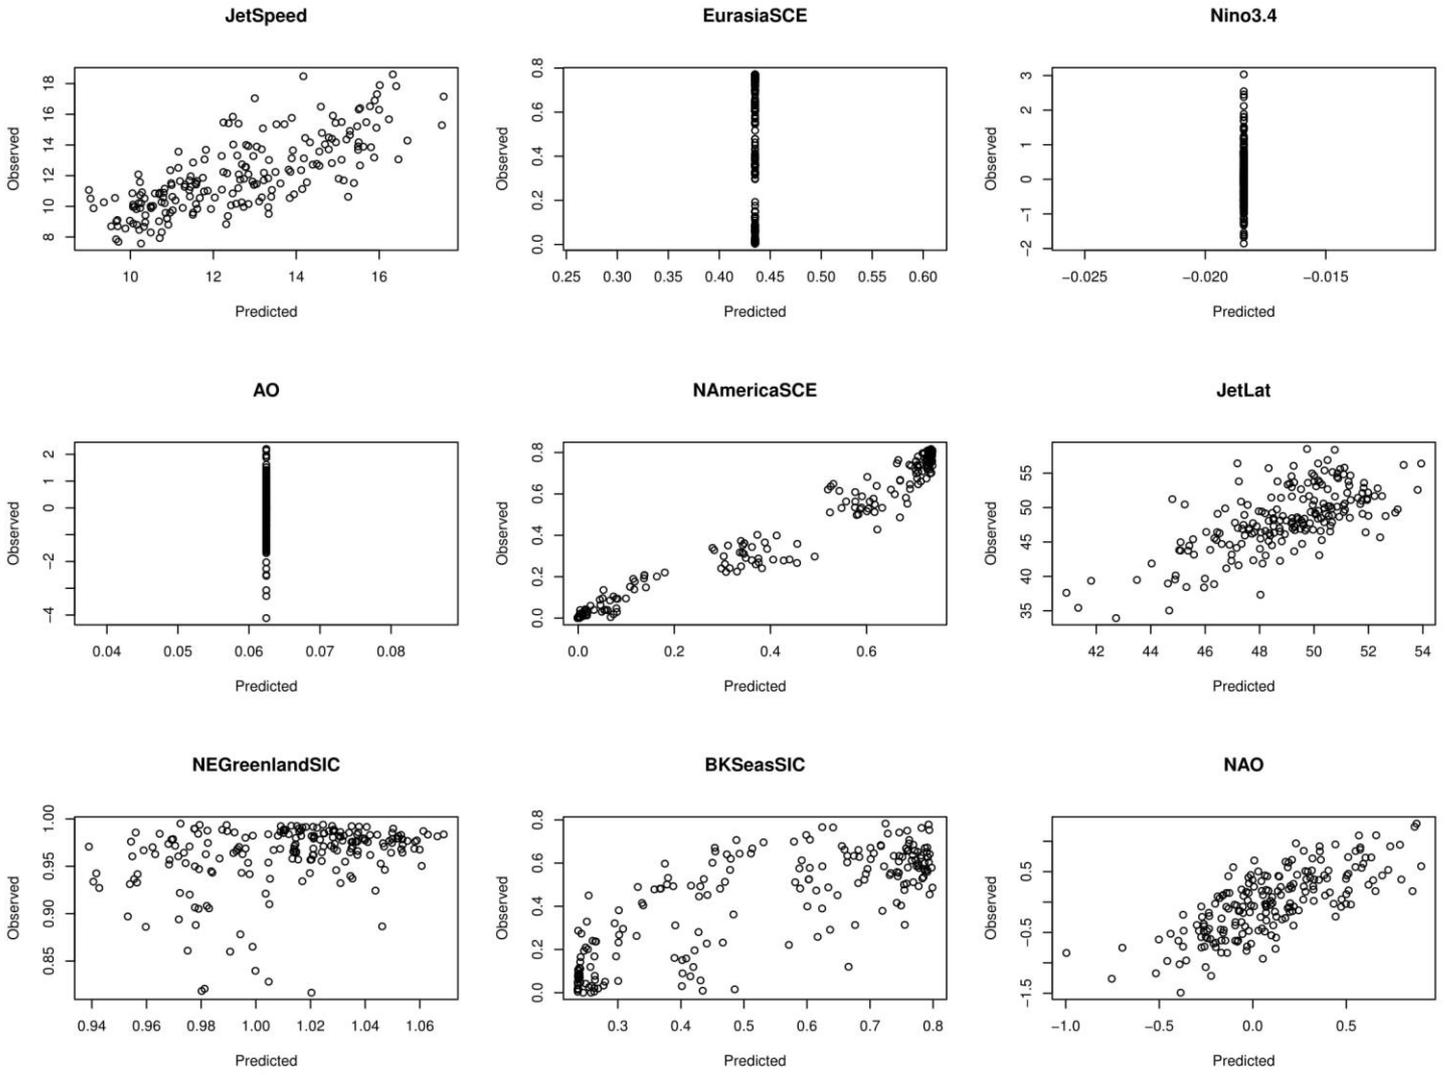


Figure A3.3 Predicted fit for all variables in Monthly 35-Year BN (2000 split).

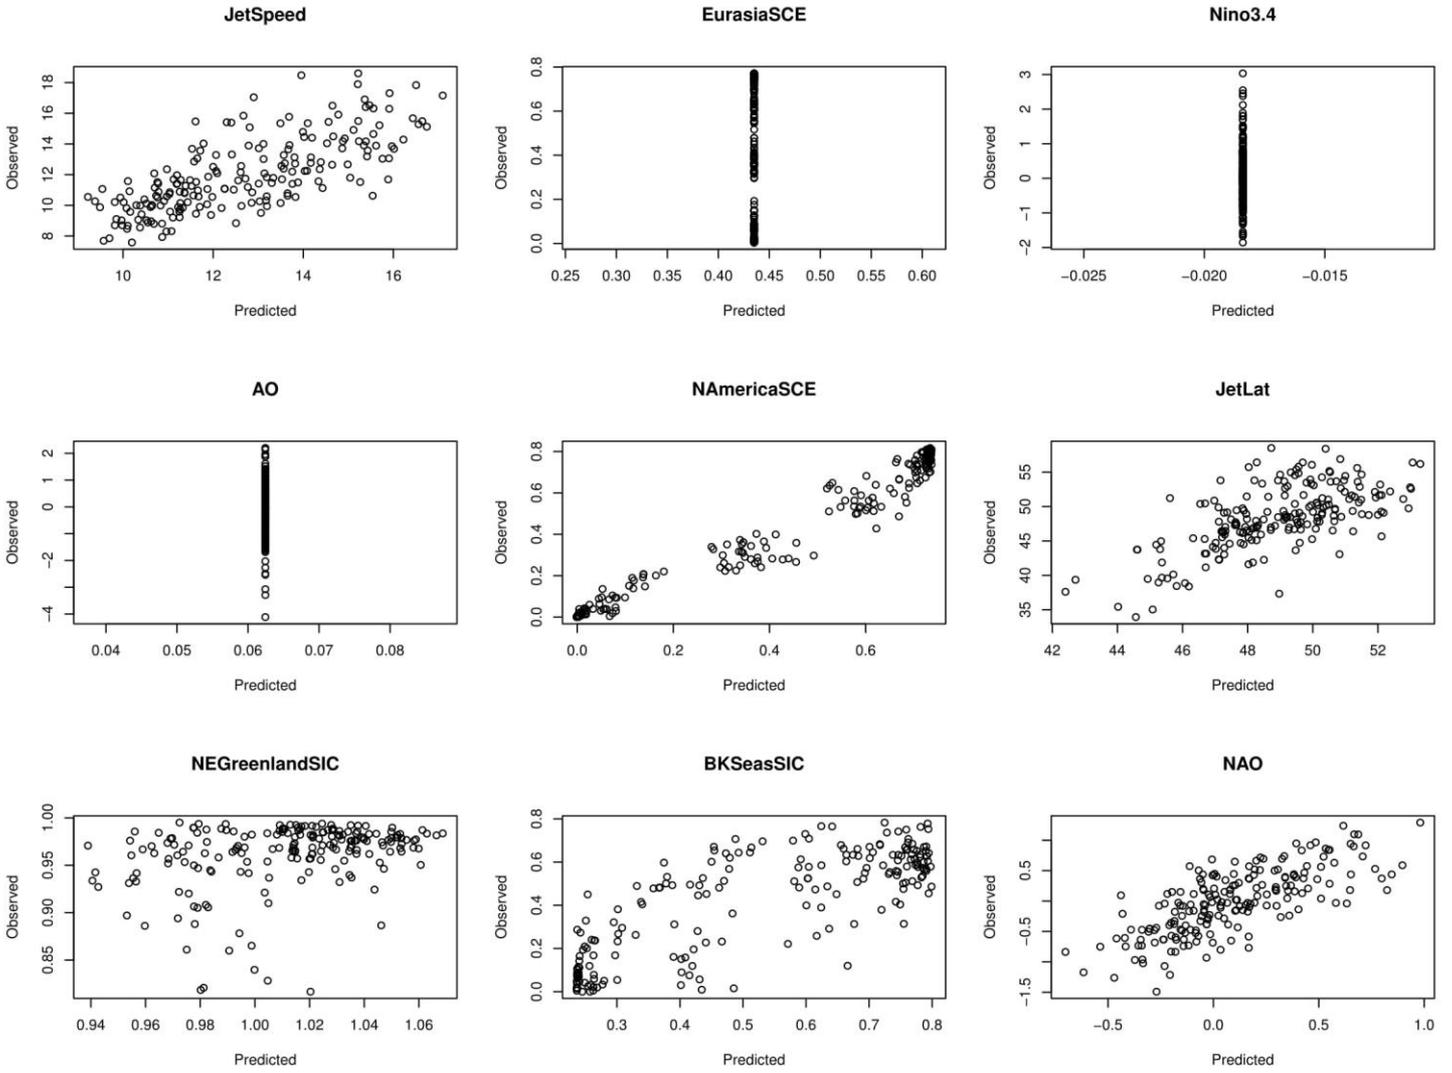


Figure A3.4 Predicted fit for all variables in Monthly 35-Year BN (2000 split) K=5.

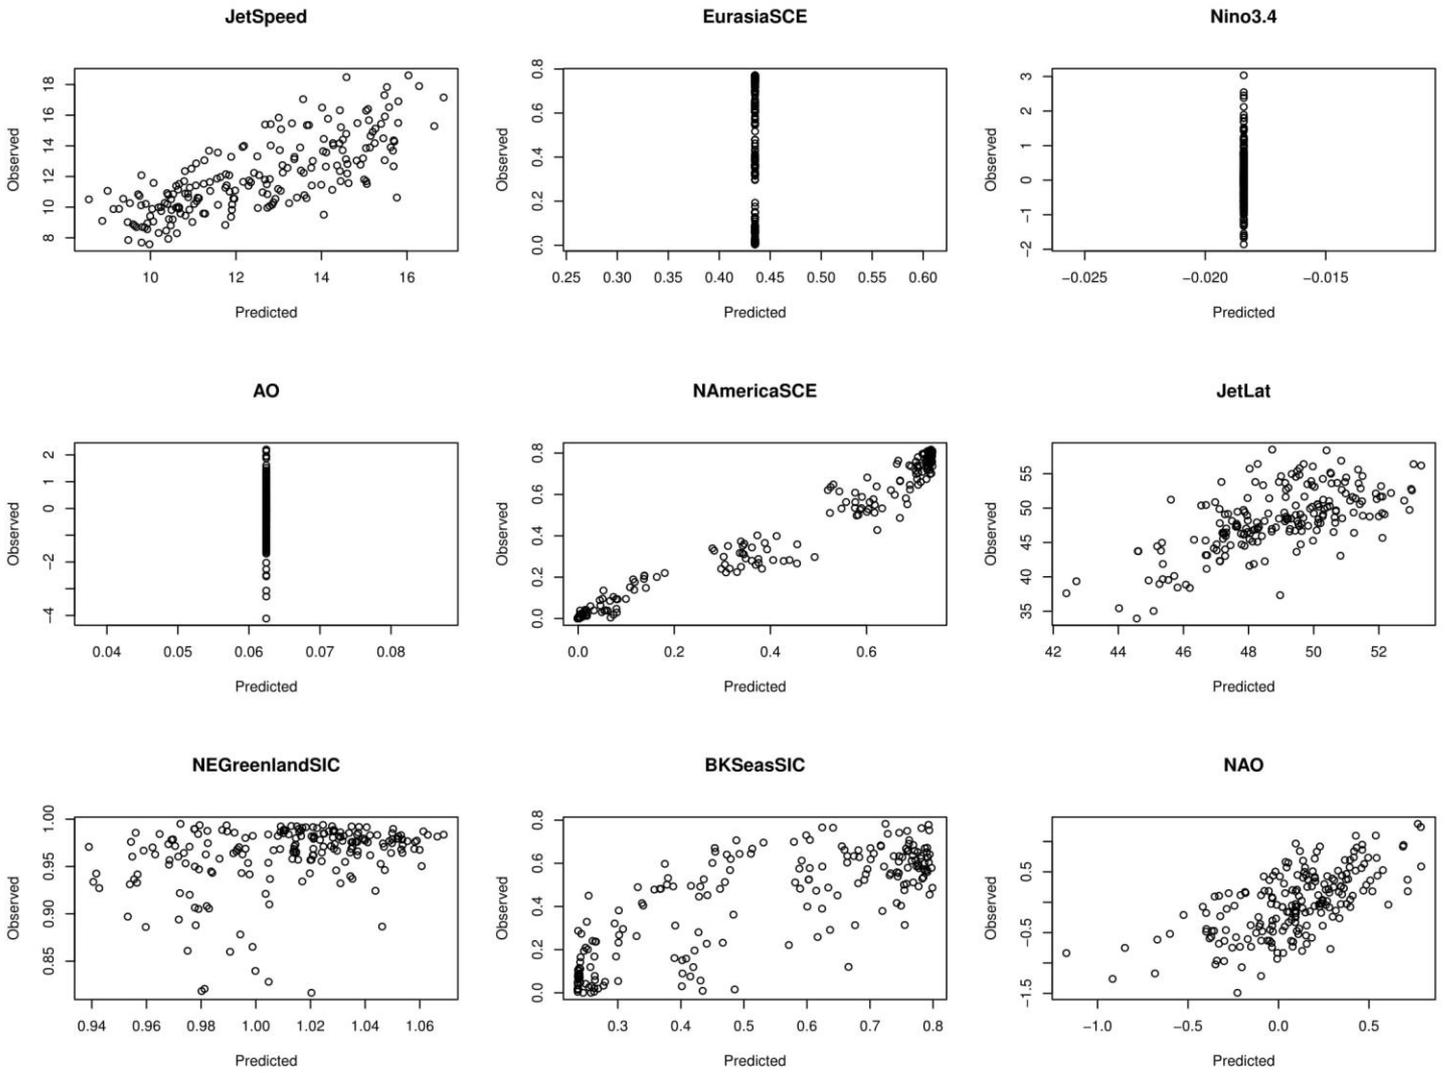


Figure A3.5 Predicted fit for all variables in Monthly 35-Year BN (2000 split) K=10.

A2 Chapter 4 Figures

	HMM	DBN 1HV	DBN 2HV
<i>a) Train data</i>			
850 NAtlantic	0.67	1.23	7.36
ENSO	0	0	0.01
Jetlat	0.15	1.48	4.9
NAO	0.09	0.05	0.09
<i>b) Test data</i>			
850 NAtlantic	3.21	2.77	1.92
Nino3.4	0.05	0.05	0
Jetlat	0.11	0.35	0.74
NAO	0.11	0.1	0.06

Table A4.1 Sum of squared error (SSE) for each variable in the HMM, DBN1HV and DBN2HV networks described in 4.3.1 split into a) Train data and b) Test data. All data is normalised (divided by the sample size) for direct comparison between training and testing datasets.

A3 Chapter 5 Figures

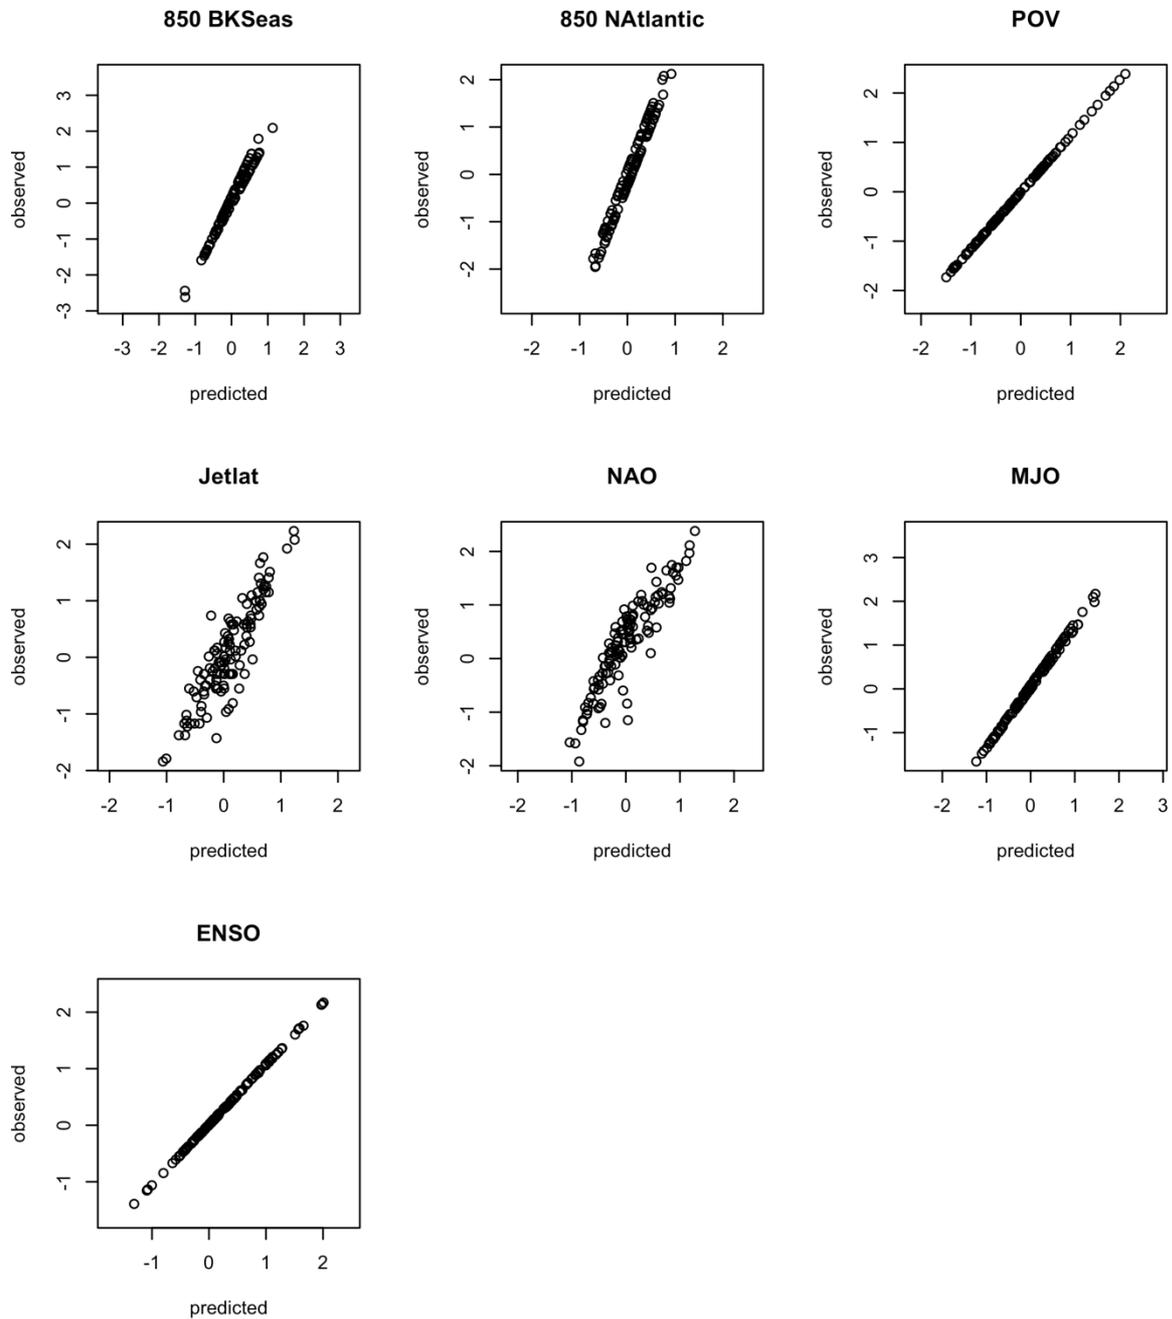


Figure A5.1 Predictive accuracy scatter plots for all variables of the 5-day Jet HV DBN; predicted test dataset values plotted against the observed values to visually assess predictive skill among models.

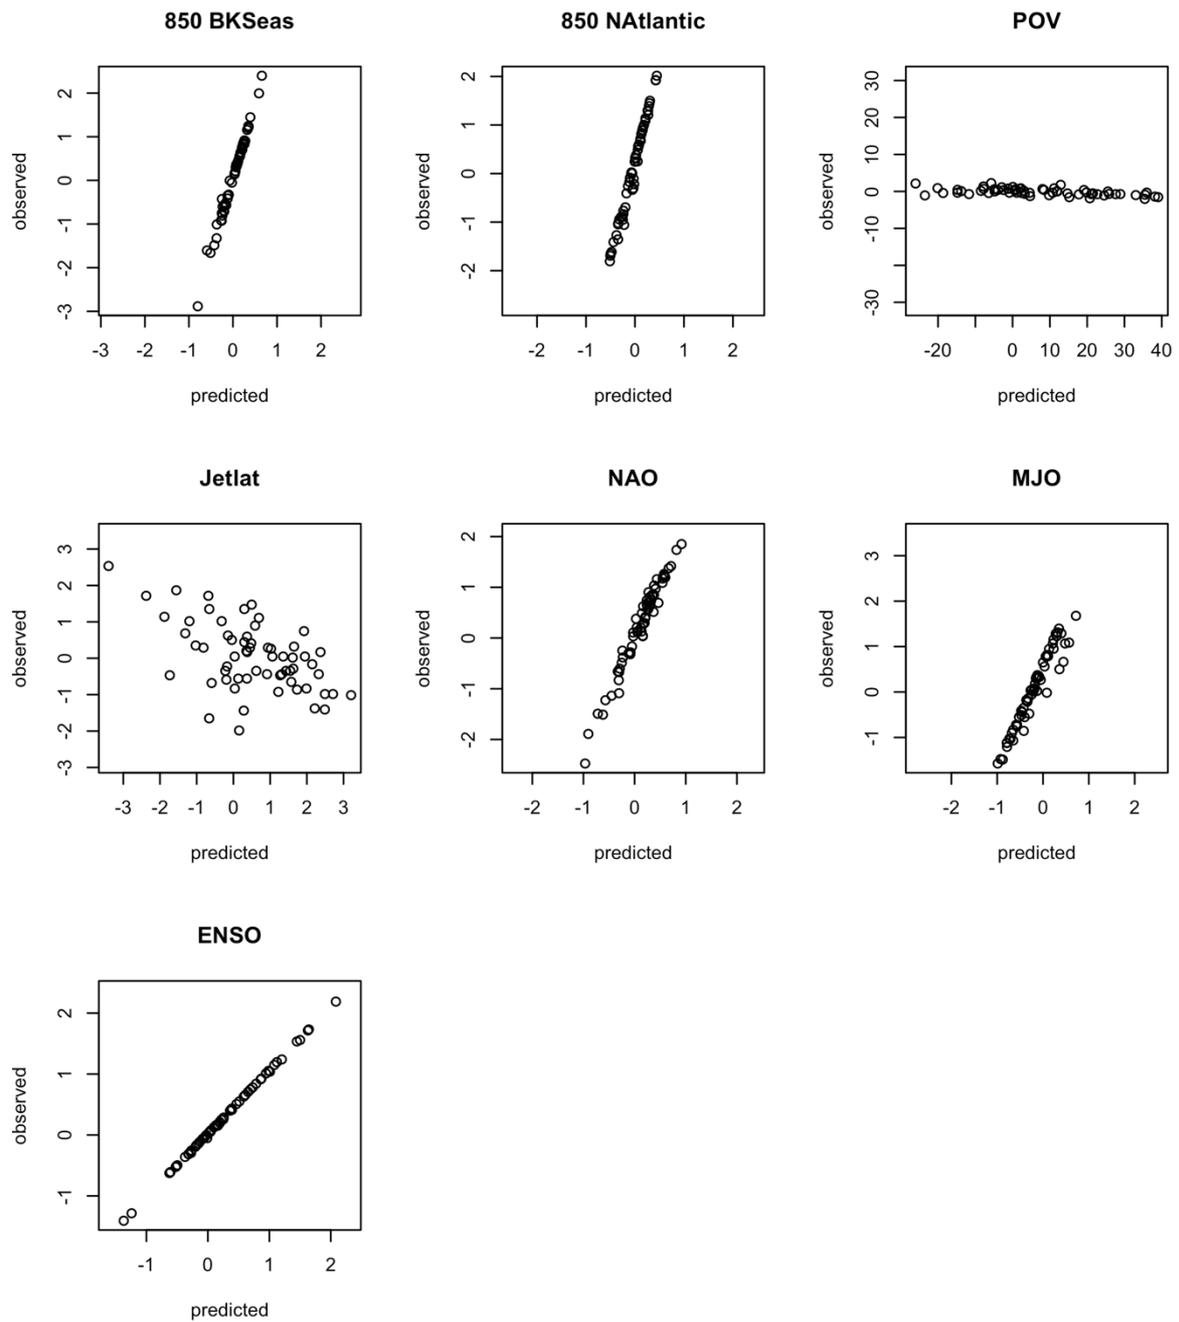


Figure A5.2 As with Fig. A5.1, but for the 10-day Jet HV DBN.

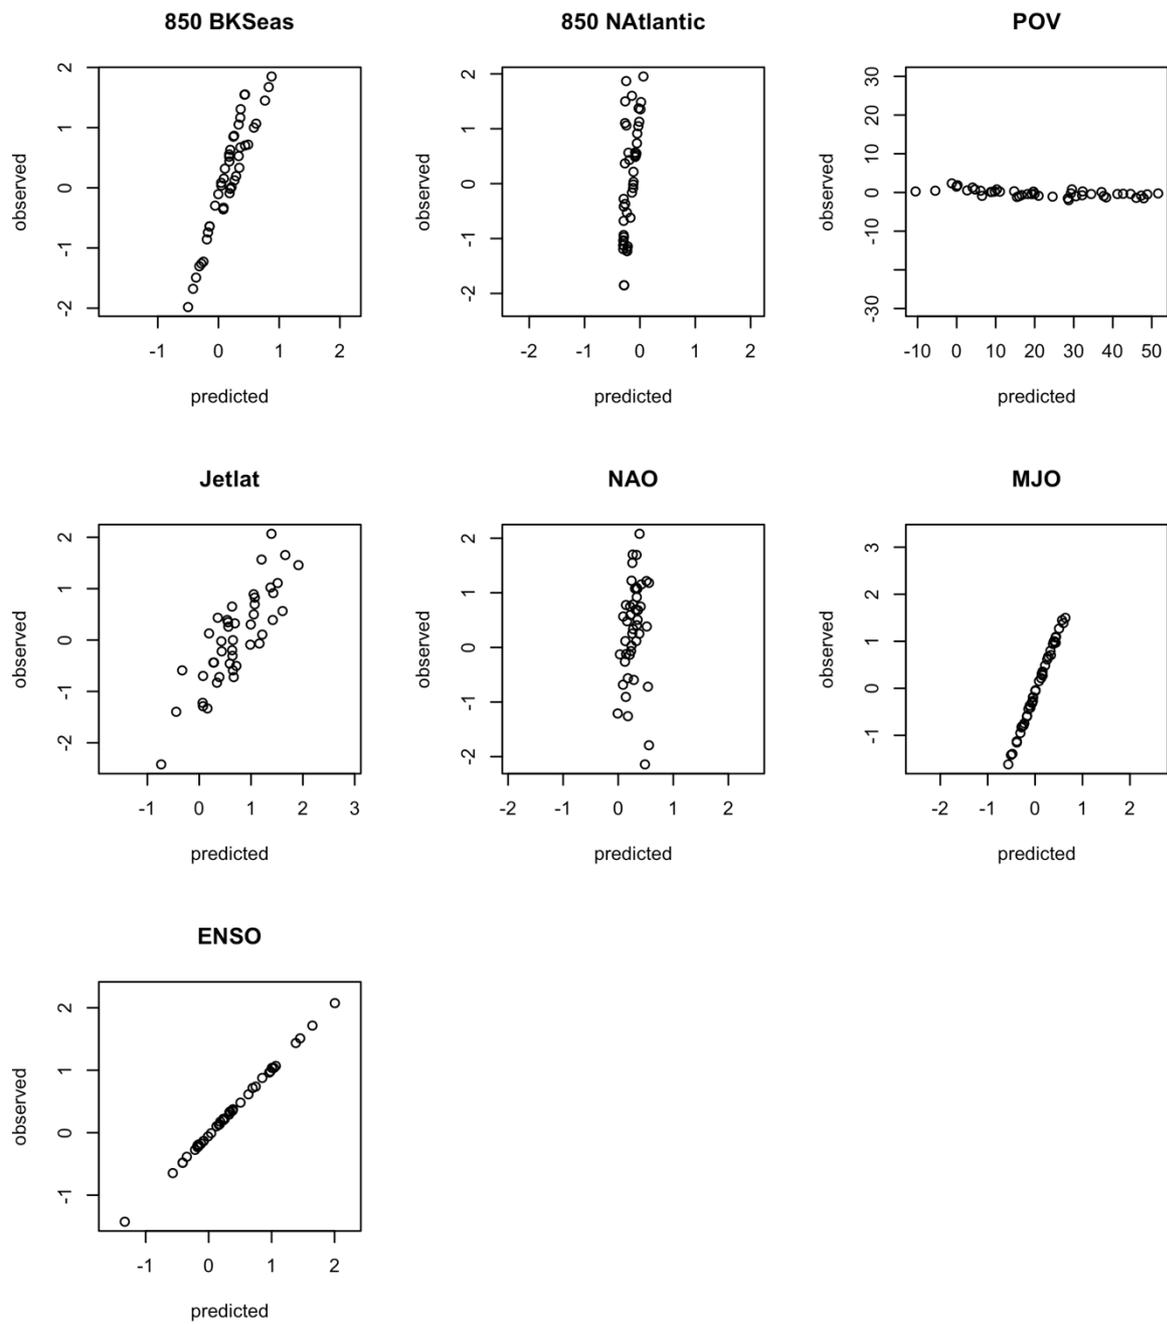


Figure A5.3 As with Fig. A5.1, but for the 15-day Jet HV DBN.

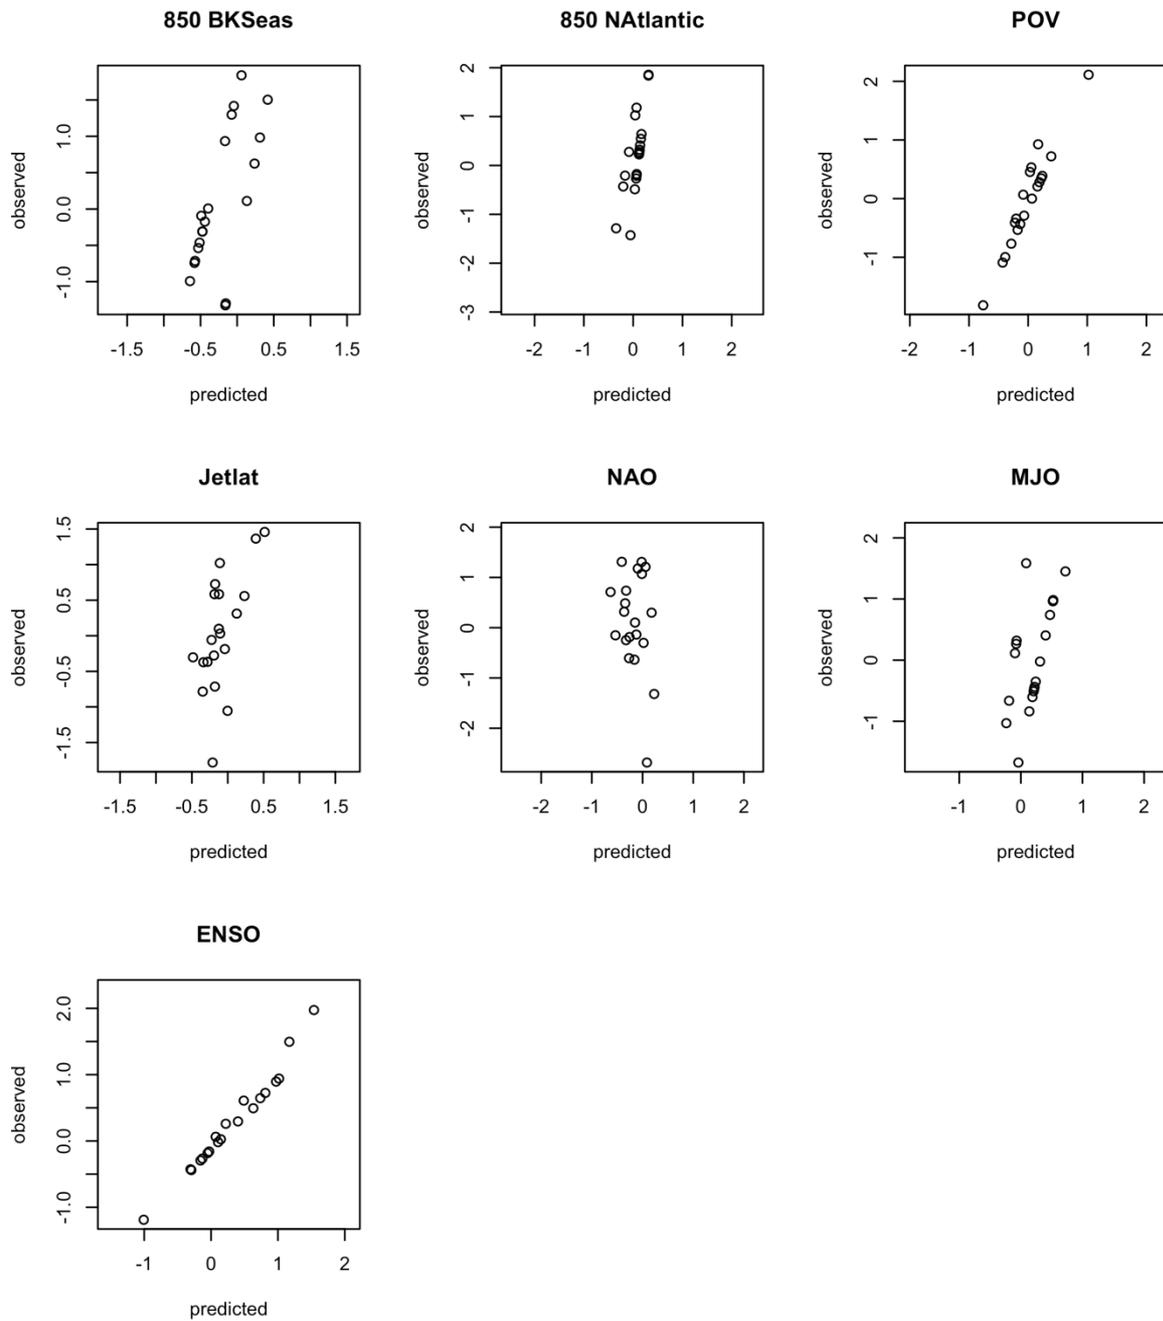


Figure A5.4 As with Fig. A5.1, but for the monthly Jet HV DBN.

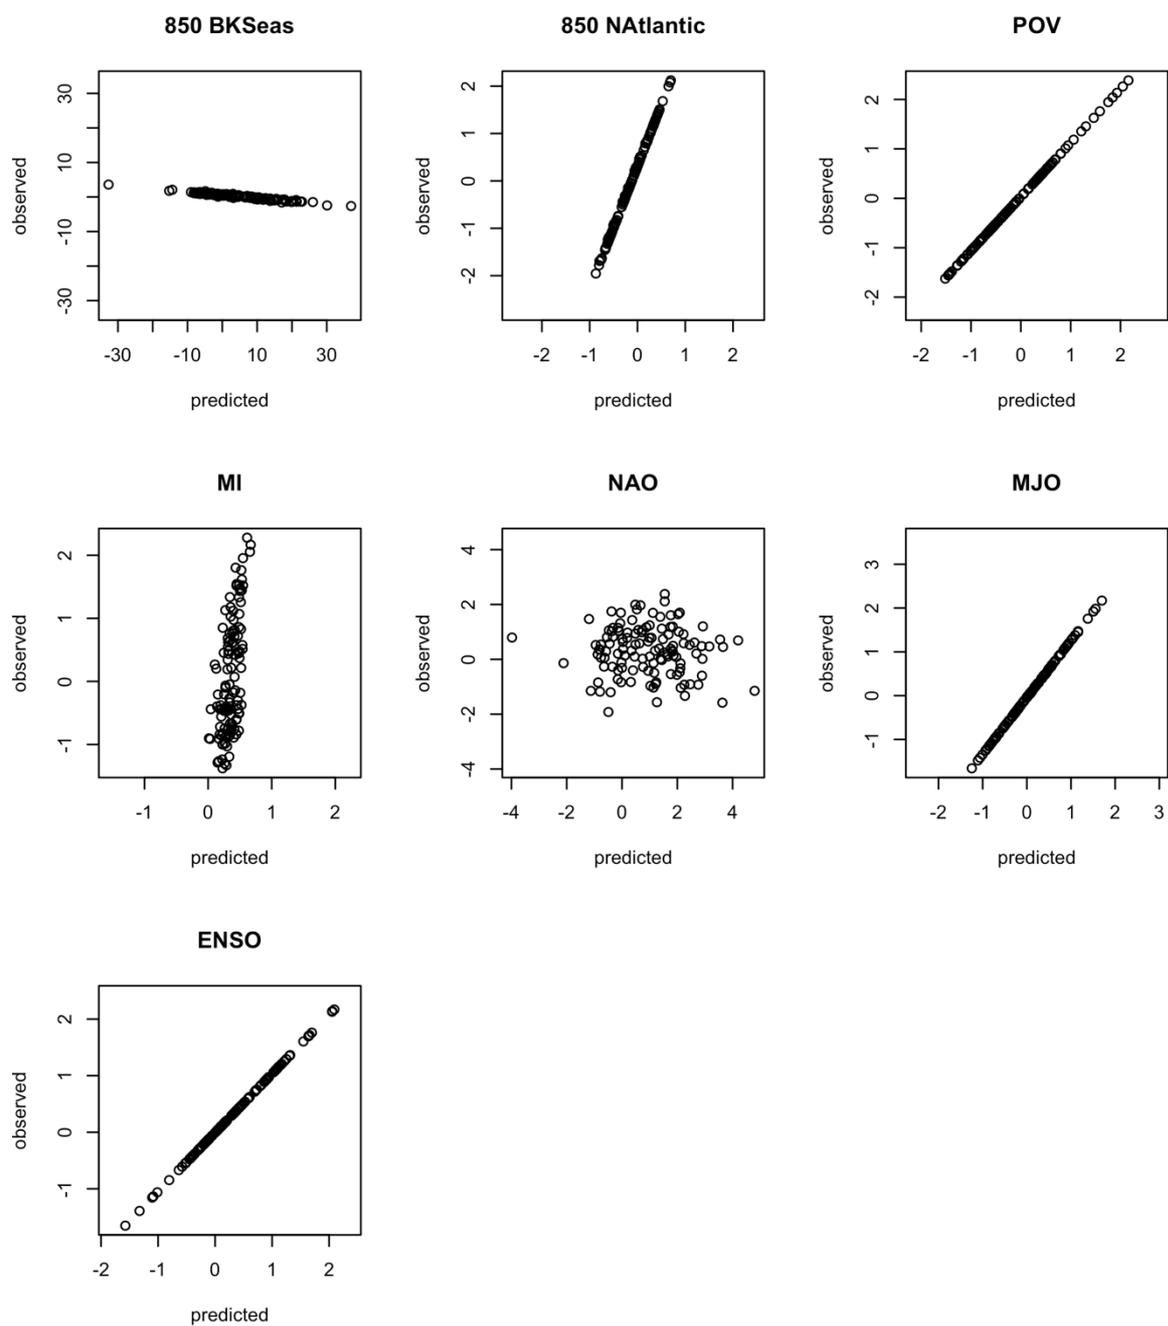


Figure A5.5 Predictive accuracy scatter plots for all variables of the 5-day MI HV DBN; predicted test dataset values plotted against the observed values to visually assess predictive skill among models.

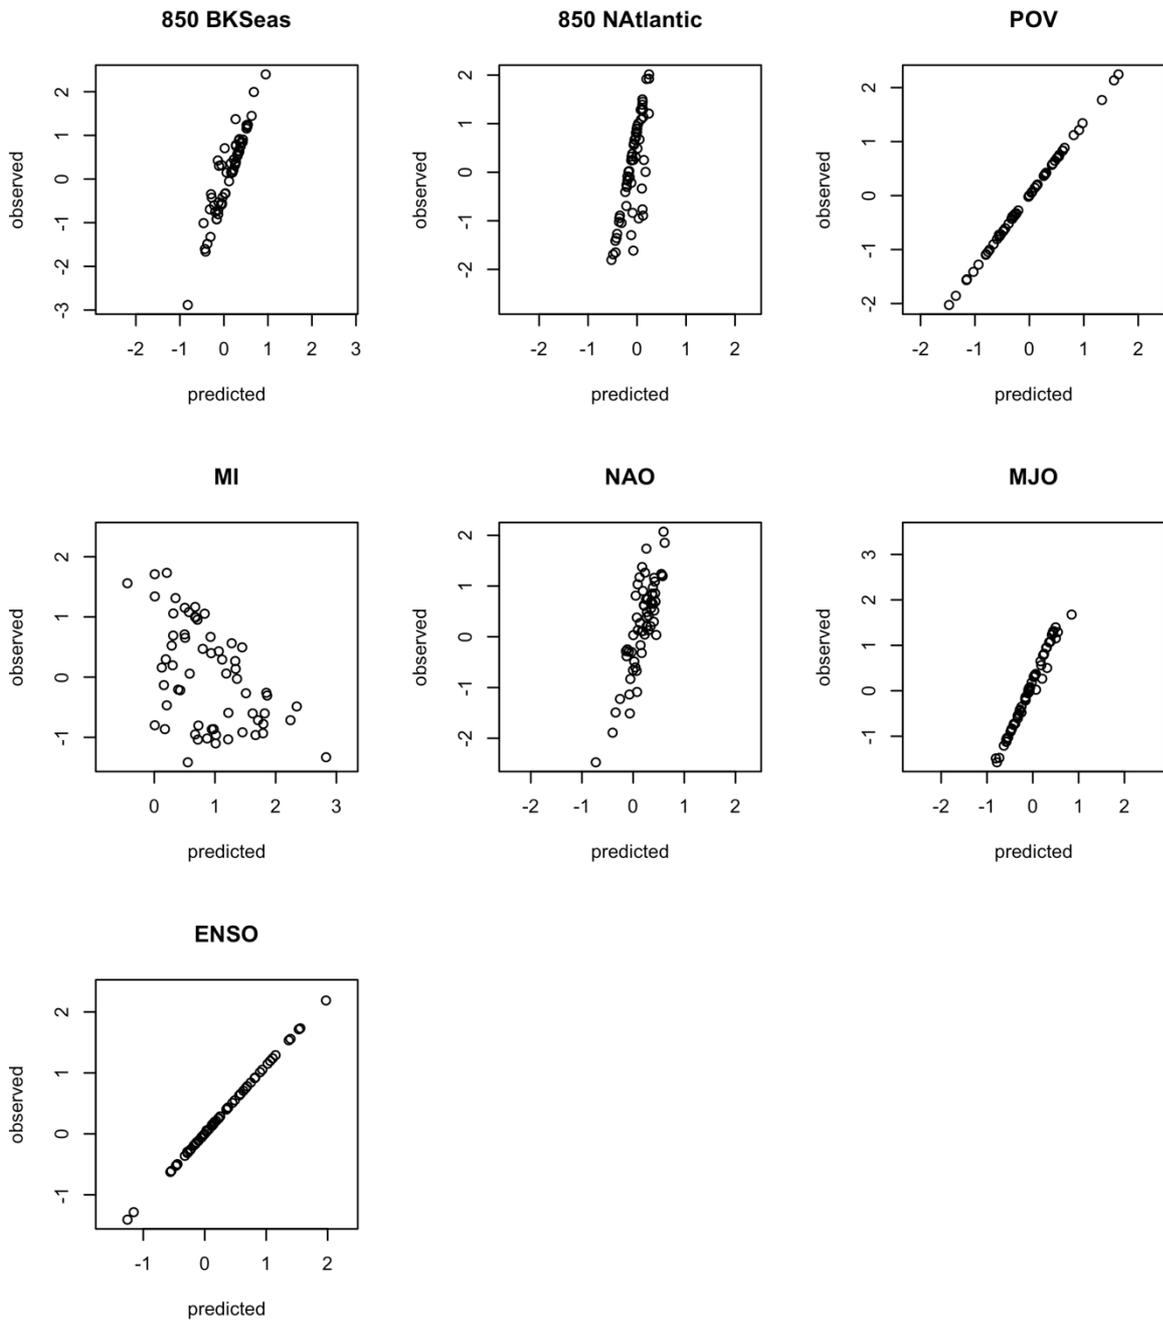


Figure A5.6 As with Fig. A5.5, but for the 10-day MI HV DBN.

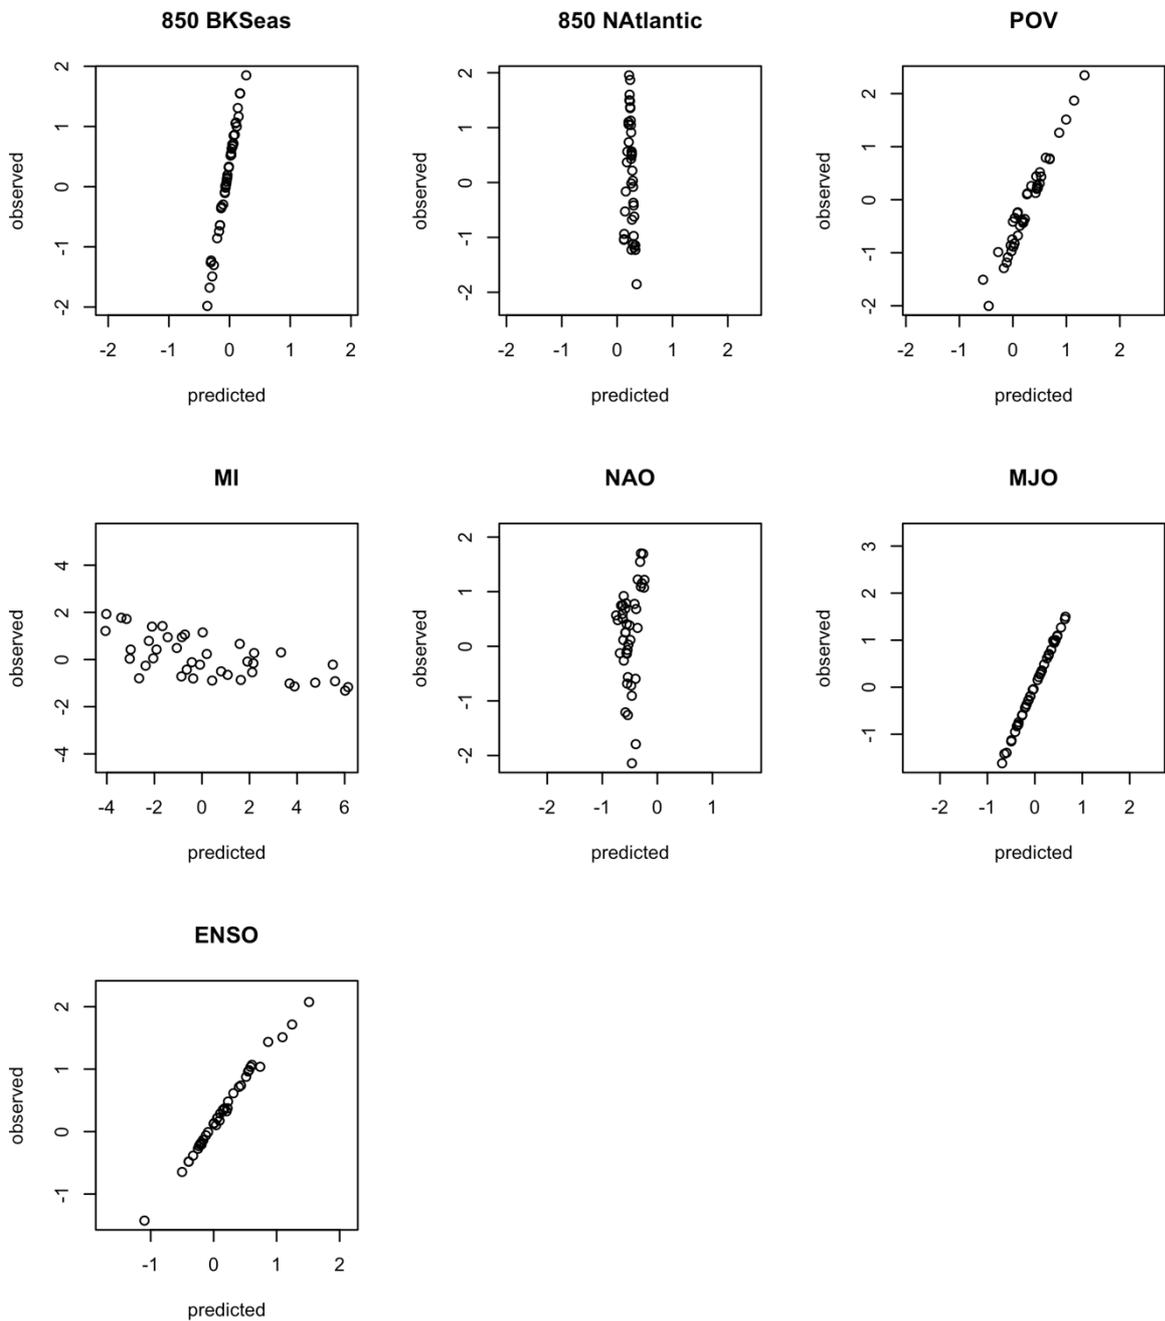


Figure A5.7 As with Fig. A5.5, but for the 15-day MI HV DBN.

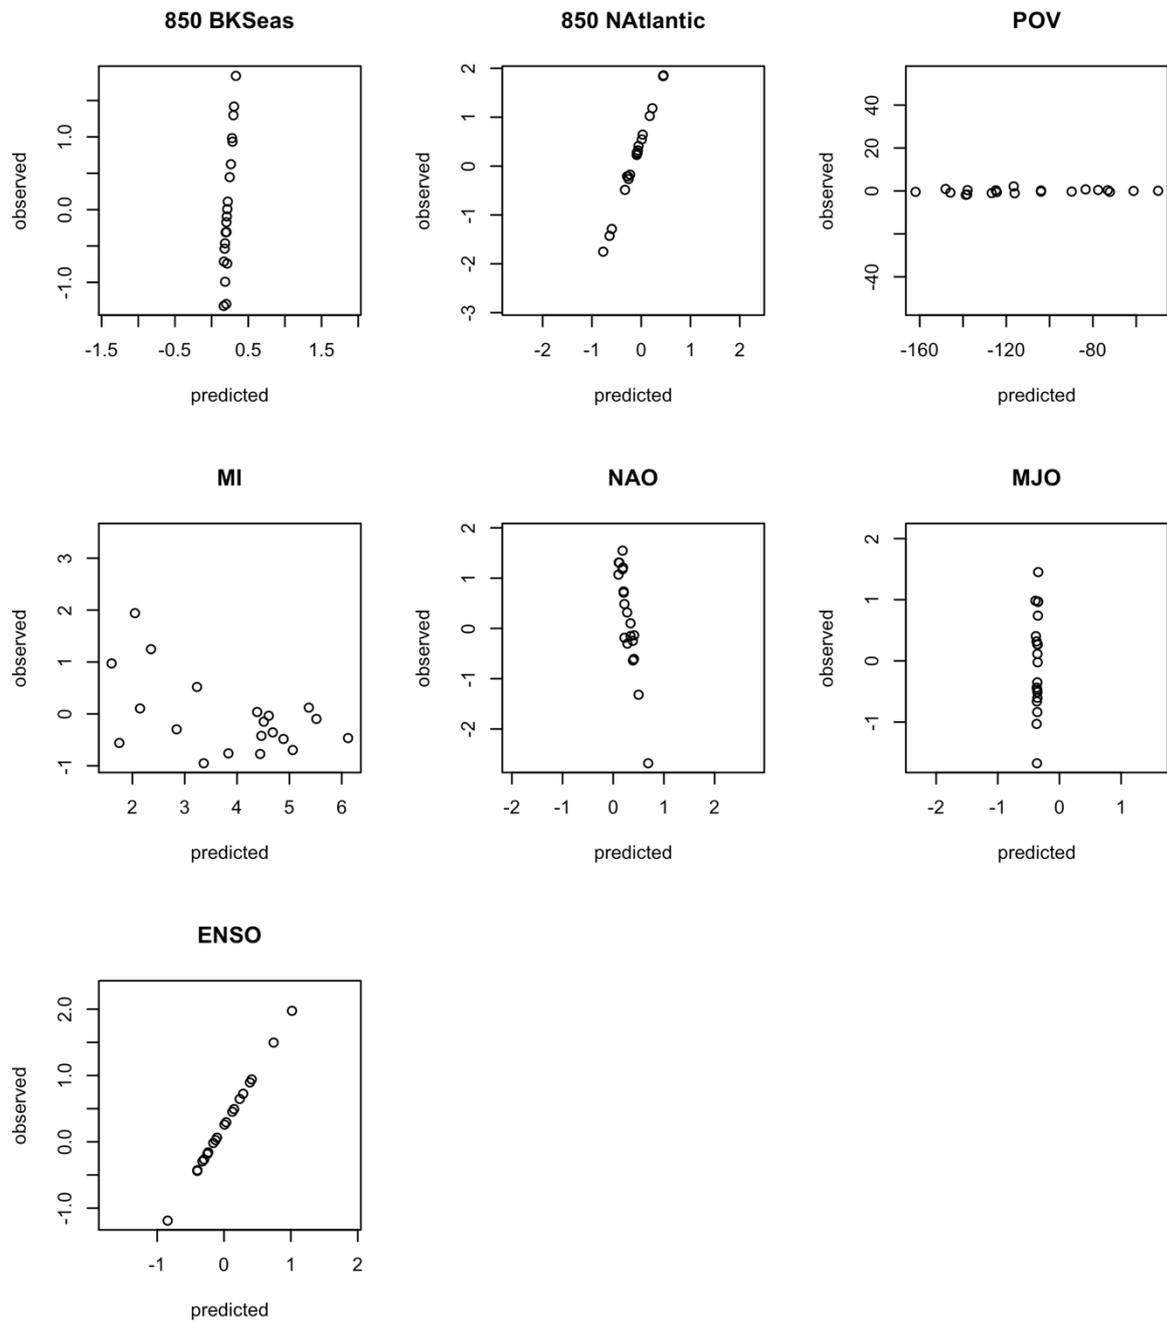


Figure A5.8 As with Fig. A5.5, but for the monthly MI HV DBN.

HV state	BK	NAt	POV	Jet/MI	NAO	MJO	ENSO
<i>a) Jet DBN</i>							
1	3.8	3.81	3295.51	8.16	0.7	0.74	0.8
2	3.91	3.16	2784.57	7.73	0.88	0.69	0.67
3	3.5	3.44	2756.17	7.99	0.77	0.7	0.6
<i>b) MI DBN</i>							
1	3.28	3.62	3109.53	0.44	0.79	0.68	0.64
2	3.47	3.67	2611.93	0.43	0.8	0.75	0.59
3	3.76	3.49	2837.93	0.47	0.74	0.73	0.68

Table A5.1 Hidden state sd values of each variable for a) Jet DBN and b) MI DBN.

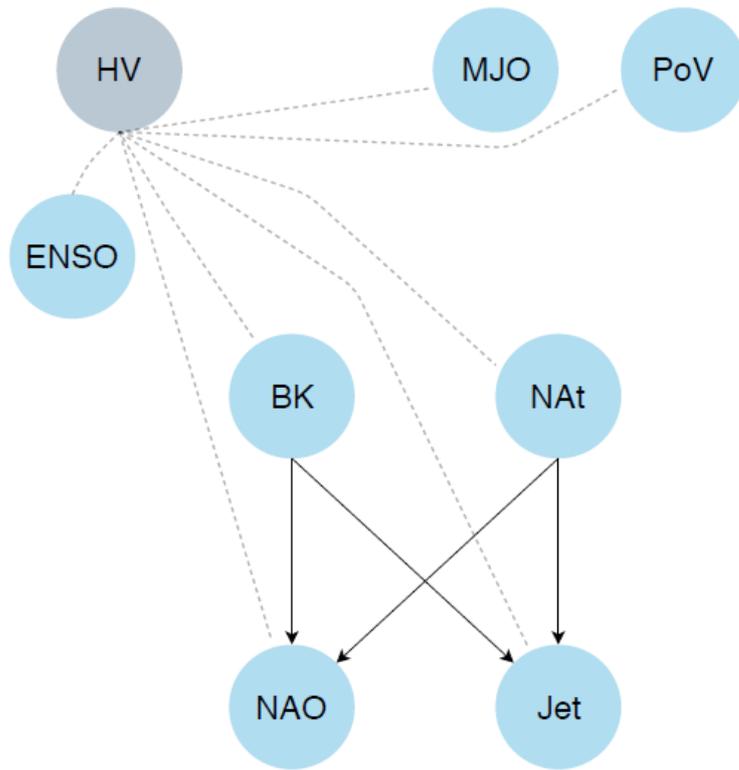


Figure A5.9 DAG of the AA era DBN with edges deleted (5.7.3). Dotted lines indicate edges coded into the model (i.e. HV edges), solid represent learned edges.

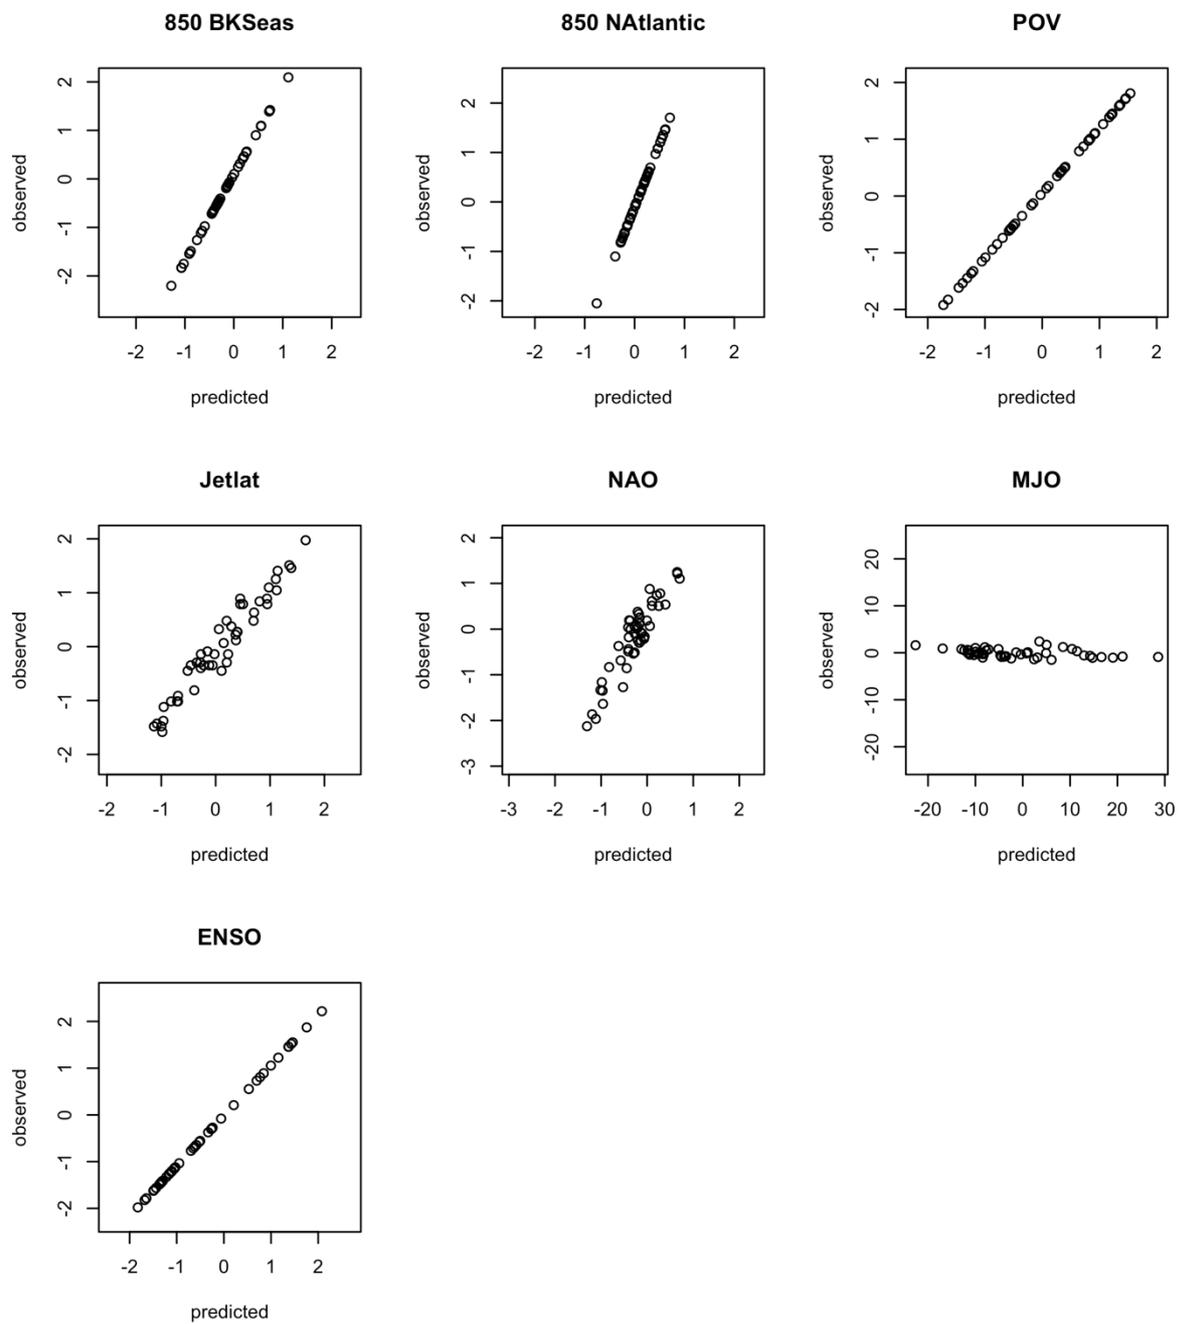


Figure A5.10 Predictive accuracy scatter plots for all variables of the pre-AA DBN (1); predicted test dataset values plotted against the observed values to visually assess predictive skill among models.

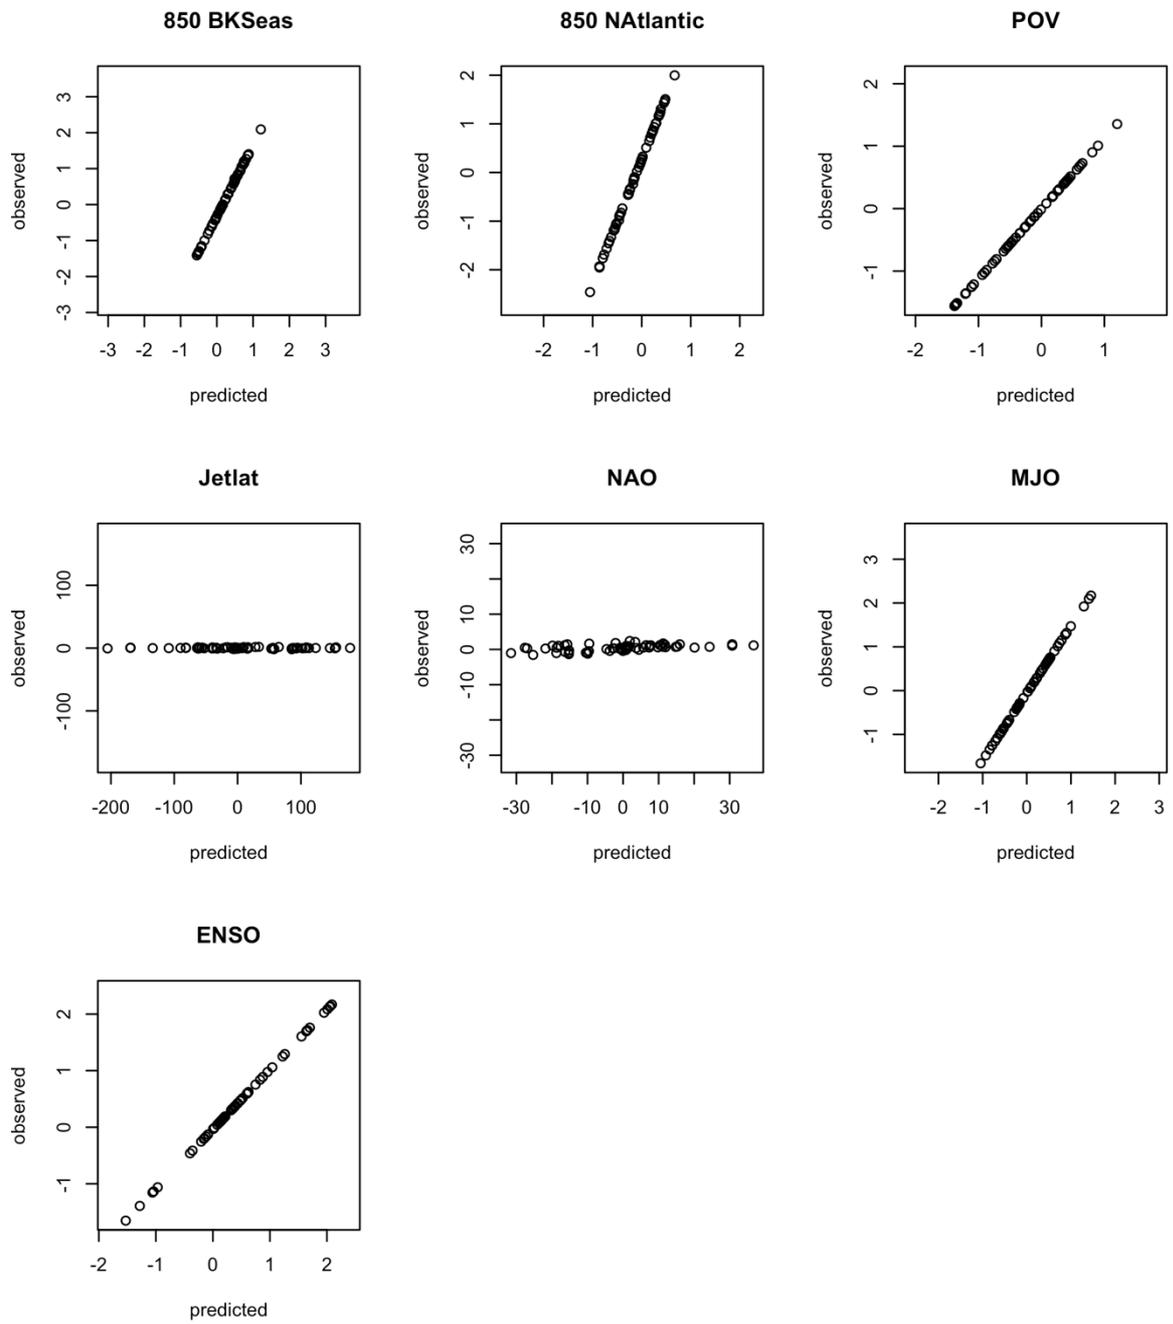


Figure A5.11 As with A5.10, but for the AA DBN (2).

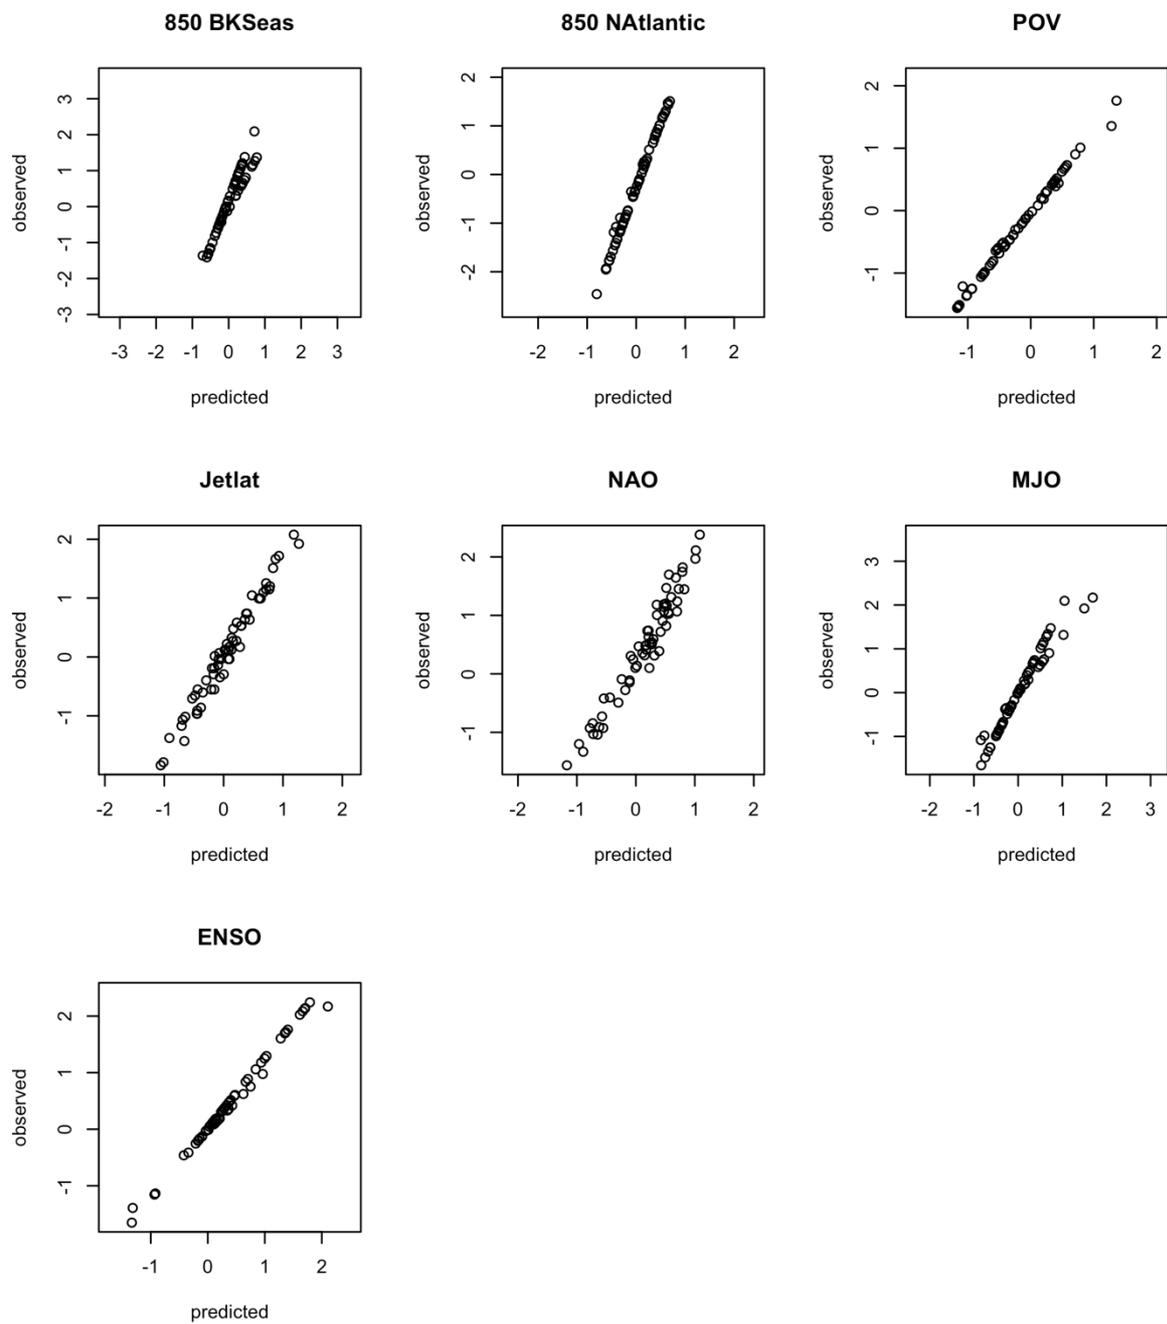


Figure A5.12 As with A5.10, but for the AA DBN with edges deleted.

	5 day	10 day	15 day	monthly
<i>a) Jet DBN train data</i>				
850 BKSeas	0.21	0.4	0.42	0.39
850 NAtlantic	0.32	0.57	0.48	0.65
POV	0.01	54.08	2283.1	0.42
Jetlat	0.24	0.54	12.36	0.65
NAO	0.31	0.62	0.46	0.7
MJO	0.05	0.18	0.24	0.64
ENSO	0	0	0	0.06
<i>b) Jet DBN test data</i>				
850 BKSeas	0.05	0.1	0.11	0.15
850 NAtlantic	0.08	0.13	0.25	0.13
POV	0	77.84	197.12	0.06
Jetlat	0.05	0.98	0.18	0.11
NAO	0.07	0.06	0.2	0.27
MJO	0.01	0.06	0.07	0.12
ENSO	0	0	0	0.01
<i>c) MI DBN train data</i>				
850 BKSeas	168.76	0.29	0.33	0.5
850 NAtlantic	0.33	0.55	0.64	0.68
POV	0.01	0.05	0.17	327.72
MI	0.84	3.21	1.79	2.03
NAO	3.91	0.59	0.42	1
MJO	0.04	0.23	0.25	0.49
ENSO	0	0.01	0.01	0.21
<i>d) MI DBN test data</i>				
850 BKSeas	33.32	0.1	0.14	0.17
850 NAtlantic	0.08	0.18	0.25	0.11
POV	0	0.01	0.09	2982.41
MI	0.18	0.59	2.89	4.27
NAO	0.73	0.12	0.29	0.31
MJO	0.01	0.05	0.05	0.17
ENSO	0	0	0.02	0.03

Table A5.2 SSE for each variable at the 5-day, 10-day, 15-day and monthly time resolutions for a) jet DBN training data b) jet DBN testing data c)MI DBN training data and d)MI DBN testing data. All data is normalised (divided by the sample size) for direct comparison between training and testing datasets.

Acknowledgements

Dr Allan Tucker was central to the workings of this doctoral project, in that he showed great patience and encouragement throughout the gruelling process of model construction. Many hours went into this process, and without his excellent supervision the methods basis of this project would never have left the ground. I am endlessly thankful for the supervision of Dr Richard Hall, who from a distance and without any institutional or other obligation to do so, managed to break a Guinness world record for supervision in the category 'fastest return of a chapter'. Both of these fantastic scientists meticulously checked over code and writing, and were central to the data and experimental design used.

My secondary supervisor, Dr Ilan Kelman of UCL, provided many teaching, reviewing and socialising opportunities for which I am grateful, and never let me get too entrenched within the closed-circuit world of climate and computer science terminology. Thanks to Dr Andrew Russell of the Committee on Climate Change, who allowed me to dream big when we came up with this project and stayed on to supervise and check over work throughout.

I am incredibly grateful for the chats I have had, through many a screen as well as in London, Vienna and Potsdam, with Giorgia Di Capua and Sem Vijverberg who also helped guide the specific focus of this project. Giorgia is an exceptional climate scientist and, perhaps even better, one of the nicest human beings I've had the pleasure of working with. I'm thankful for the departmental guidance that Dr Elisabete Silva has provided, which kept me on track and helped ameliorate anxiety levels.

This project was funded by the Natural Environment Research Council (NERC) via the London NERC Doctoral Training Programme. Being fortunate enough to join the London NERC DTP was the greatest accomplishment of my life to date, and I gained some excellent environmental science colleagues and close friends in the process.

Heartfelt thanks must go to friends and family who had to put up with so much Bayesian Network and climate science chat over the years. Ben Walter, who possibly took the brunt of this, made sure I maintained a life outside science by filling it with good food, good chats and music. Cherry-Beth and Alice also provided enthusiastic assistance towards this end, and I am thankful for every laugh and drink around the house and National Trust property that got me out of London. The office of 136 have been excellent PhD companions, and Phil in particular has always helped me see the funny side of research, as well as the earlier side of 8am for morning fitness antics. For the support of my parents, Tim and Harriet, a lifetime's worth of debt will be accrued. Their support in every step taken towards the PhD, both before

becoming a doctoral researcher and during the project, means the world. Finally, to Sabina, who was there for me during the hardest months of the project in the last two years.

References

- Aguilera, P. A., A. Fernández, R. Fernández, R. Rumí, and A. Salmerón, 2011: Bayesian networks in environmental modelling. *Environ. Model. Softw.*, **26**, 1376–1388, <https://doi.org/10.1016/j.envsoft.2011.06.004>.
- Ambaum, M. H. P., 2012: Frequentist vs Bayesian statistics - a non-statisticians view.
- Ambaum, M. H. P., B. J. Hoskins, and D. B. Stephenson, 2001: Arctic Oscillation or North Atlantic Oscillation? *J. Clim.*, **14**, 3495–3507, [https://doi.org/10.1175/1520-0442\(2001\)014<3495:AOONAO>2.0.CO;2](https://doi.org/10.1175/1520-0442(2001)014<3495:AOONAO>2.0.CO;2).
- Ayarzagüena, B., and J. A. Screen, 2016: Future Arctic sea ice loss reduces severity of cold air outbreaks in midlatitudes. *Geophys. Res. Lett.*, **43**, 2801–2809, <https://doi.org/10.1002/2016GL068092>.
- Baldwin, M. P., and T. J. Dunkerton, 2001: Stratospheric harbingers of anomalous weather regimes. *Science (80-.)*, **294**, 581–584, <https://doi.org/10.1126/science.1063315>.
- Ballinger, T. J., E. Hanna, R. J. Hall, J. Miller, M. H. Ribergaard, and J. L. Høyer, 2018: Greenland coastal air temperatures linked to Baffin Bay and Greenland Sea ice conditions during autumn through regional blocking patterns. *Clim. Dyn.*, **50**, 83–100, <https://doi.org/10.1007/s00382-017-3583-3>.
- Barnes, E. A., 2013: Revisiting the evidence linking Arctic amplification to extreme weather in midlatitudes. *Geophys. Res. Lett.*, **40**, 4734–4739, <https://doi.org/10.1002/grl.50880>.
- Barnes, E. A., E. Dunn-Sigouin, G. Masato, and T. Woollings, 2014: Exploring recent trends in Northern Hemisphere blocking. *Geophys. Res. Lett.*, **41**, 638–644, <https://doi.org/10.1002/2013GL058745>.
- Barnes, E. A., and L. M. Polvani, 2015: CMIP5 projections of arctic amplification, of the North American/North Atlantic circulation, and of their relationship. *J. Clim.*, **28**, 5254–5271, <https://doi.org/10.1175/JCLI-D-14-00589.1>.
- Barnes, E. A., and J. A. Screen, 2015: The impact of Arctic warming on the midlatitude jet-stream: Can it? Has it? Will it? *Wiley Interdiscip. Rev. Clim. Chang.*, **6**, 277–286, <https://doi.org/10.1002/wcc.337>.
- Barnes, E. A., and I. R. Simpson, 2017: Seasonal sensitivity of the Northern Hemisphere jet streams to Arctic temperatures on subseasonal time scales. *J. Clim.*, **30**, 10117–10137, <https://doi.org/10.1175/JCLI-D-17-0299.1>.
- Bilmes, J. a., 1998: A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models. *Int. Comput. Sci. Inst.*, **4**, 126, <https://doi.org/10.1.1.119.4856>.

- Blackport, R., and J. A. Screen, 2019: Influence of Arctic Sea Ice Loss in Autumn Compared to That in Winter on the Atmospheric Circulation. *Geophys. Res. Lett.*, **46**, 2213–2221, <https://doi.org/10.1029/2018GL081469>.
- Blackport, R., J. A. Screen, K. van der Wiel, and R. Bintanja, 2019: Minimal influence of reduced Arctic sea ice on coincident cold winters in mid-latitudes. *Nat. Clim. Chang.*, **9**, 697–704, <https://doi.org/10.1038/s41558-019-0551-4>.
- Blockley, E. W., and Coauthors, 2014: Recent development of the Met Office operational ocean forecasting system: An overview and assessment of the new Global FOAM forecasts. *Geosci. Model Dev.*, **7**, 2613–2638, <https://doi.org/10.5194/gmd-7-2613-2014>.
- Brönnimann, S., 2007: Impact of El Niño-Southern Oscillation on European climate. *Rev. Geophys.*, <https://doi.org/10.1029/2006RG000199>.1.INTRODUCTION.
- Brownlee, J., 2016: *Master Machine Learning Algorithms: discover how they work and implement them from scratch*.
- Cain, J., 2001: Planning improvements in natural resources management: Guidelines for using Bayesian networks to support the planning and management of development programmes in the water sector and beyond. *Ecology*, **44**, 132.
- Cao, C., Y.-H. Chen, J. Rao, S.-M. Liu, S.-Y. Li, M.-H. Ma, and Y.-B. Wang, 2019: Statistical Characteristics of Major Sudden Stratospheric Warming Events in CESM1-WACCM: A Comparison with the JRA55 and NCEP/NCAR Reanalyses. *Atmosphere (Basel)*, **10**, 519, <https://doi.org/10.3390/atmos10090519>.
- Chelton, D. B., and C. M. Risien, 2016: Zonal and Meridional Discontinuities and Other Issues with the HadISST1.1 Dataset. **1**.
- Chen, S. H., and C. A. Pollino, 2012: Good practice in Bayesian network modelling. *Environ. Model. Softw.*, **37**, 134–145, <https://doi.org/10.1016/j.envsoft.2012.03.012>.
- Chen, X., and D. Luo, 2017: Arctic sea ice decline and continental cold anomalies: Upstream and downstream effects of Greenland blocking. *Geophys. Res. Lett.*, **44**, 3411–3419, <https://doi.org/10.1002/2016GL072387>.
- Cohen, J., 2016: An observational analysis: Tropical relative to Arctic influence on midlatitude weather in the era of Arctic amplification. *Geophys. Res. Lett.*, **43**, 5287–5294, <https://doi.org/10.1002/2016GL069102>.
- Cohen, J., and Coauthors, 2014: Recent Arctic amplification and extreme mid-latitude weather. *Nat. Geosci.*, **7**, 627–637, <https://doi.org/10.1038/ngeo2234>.
- Cohen, J., and Coauthors, **2018a**: Arctic change and possible influence on mid-latitude climate and weather. *US CLIVAR Rep. 2018-1*, 41, <https://doi.org/10.5065/D6TH8KGW.CONTRIBUTORS>.

- Cohen, J., K. Pfeiffer, and J. A. Francis, **2018b**: Warm Arctic episodes linked with increased frequency of extreme winter weather in the United States. *Nat. Commun.*, **9**, 1–12, <https://doi.org/10.1038/s41467-018-02992-9>.
- Coumou, D., and S. Rahmstorf, 2012: A decade of weather extremes. *Nat. Clim. Chang.*, **2**, 491–496, <https://doi.org/10.1038/nclimate1452>.
- Coumou, D., J. Lehmann, and J. Beckmann, 2015: The weakening summer circulation in the Northern Hemisphere mid-latitudes. *Science (80-.)*, **348**, 324–327, <https://doi.org/10.1126/science.1261768>.
- Coumou, D., V. Petoukhov, S. Rahmstorf, S. Petri, and H. J. Schellnhuber, 2014: Quasi-resonant circulation regimes and hemispheric synchronization of extreme weather in boreal summer. *Proc. Natl. Acad. Sci. U. S. A.*, **111**, 12331–12336, <https://doi.org/10.1073/pnas.1412797111>.
- Cvijanovic, I., and K. Caldeira, 2015: Atmospheric impacts of sea ice decline in CO2 induced global warming. *Clim. Dyn.*, **44**, 1173–1186, <https://doi.org/10.1007/s00382-015-2489-1>.
- Dee, D. P., and Coauthors, 2011: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Q. J. R. Meteorol. Soc.*, **137**, 553–597, <https://doi.org/10.1002/qj.828>.
- Dempster, A. P., N. M. Laird, and D. B. Rubin, 1977: Maximum Likelihood from Incomplete Data Via the EM Algorithm. *J. R. Stat. Soc. Ser. B*, **39**, 1–22, <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>.
- Di Capua, G., and D. Coumou, 2016: Changes in meandering of the Northern Hemisphere circulation. *Environ. Res. Lett.*, **11**, 094028, <https://doi.org/10.1088/1748-9326/11/9/094028>.
- Di Capua, G., and Coauthors, 2019: Long-lead statistical forecasts of the indian summer monsoon rainfall based on causal precursors. *Weather Forecast.*, **34**, 1377–1394, <https://doi.org/10.1175/WAF-D-19-0002.1>.
- Dielmann, A., and S. Renals, 2007: Dynamic Bayesian Networks. **9**, 25–36.
- Dobricic, S., E. Vignati, and S. Russo, 2016: Large-scale atmospheric warming in winter and the arctic sea ice retreat. *J. Clim.*, **29**, 2869–2888, <https://doi.org/10.1175/JCLI-D-15-0417.1>.
- Douville, H., Y. Peings, and D. Saint-Martin, 2017: Snow-(N)AO relationship revisited over the whole twentieth century. *Geophys. Res. Lett.*, **44**, 569–577, <https://doi.org/10.1002/2016GL071584>.
- Ebert-Uphoff, I., and Y. Deng, **2012a**: Causal discovery for climate research using graphical models. *J. Clim.*, **25**, 5648–5665, <https://doi.org/10.1175/JCLI-D-11-00387.1>.

- Ebert-Uphoff, I., and Y. Deng, **2012b**: A new type of climate network based on probabilistic graphical models: Results of boreal winter versus summer. *Geophys. Res. Lett.*, **39**, 1–7, <https://doi.org/10.1029/2012GL053269>.
- Eguchi, T., 2008: An Introduction to Bayesian Statistics Without Using Equations. *Mar. Turt. Newsl.*, 1–5.
- Feldstein, S. B., and S. Lee, 2014: Intraseasonal and interdecadal jet shifts in the Northern Hemisphere: The role of warm pool tropical convection and sea ice. *J. Clim.*, **27**, 6497–6518, <https://doi.org/10.1175/JCLI-D-14-00057.1>.
- Fettweis, X., E. Hanna, C. Lang, A. Belleflamme, M. Erpicum, and H. Gallée, 2013: Brief communication Important role of the mid-tropospheric atmospheric circulation in the recent surface melt increase over the Greenland ice sheet. *Cryosphere*, **7**, 241–248, <https://doi.org/10.5194/tc-7-241-2013>.
- Fountalis, I., A. Bracco, and C. Dvovolis, 2014: Spatio-temporal network analysis for studying climate patterns. *Clim. Dyn.*, **42**, 879–899, <https://doi.org/10.1007/s00382-013-1729-5>.
- Francis, J., and J. Overland, 2014: Implications of rapid Arctic change for weather patterns in northern mid-latitudes. *Us Clivar*, 1–28.
- Francis, J., 2013: Rapid Arctic Warming and Wacky Weather: Are they linked? *association for women in science*.
- Francis, J. A., W. Chan, D. J. Leathers, J. R. Miller, and D. E. Veron, 2009: Winter Northern Hemisphere weather patterns remember summer Arctic sea-ice extent. *Geophys. Res. Lett.*, **36**, <https://doi.org/10.1029/2009GL037274>.
- Francis, J. A., and S. J. Vavrus, 2012: Francis_2012_Evidence linking Arctic Amplification to Extreme Weather.pdf. *Geophys. Res. Lett.*, **39**, 1–6, <https://doi.org/10.1029/2012GL051000>.
- Francis, J. A., N. Skific, and S. J. Vavrus, 2018: North American Weather Regimes Are Becoming More Persistent: Is Arctic Amplification a Factor? *Geophys. Res. Lett.*, **45**, 11,414–11,422, <https://doi.org/10.1029/2018GL080252>.
- Francis, J. A., and S. J. Vavrus, 2015: Evidence for a wavier jet stream in response to rapid Arctic warming. *Environ. Res. Lett.*, **10**, 014005, <https://doi.org/10.1088/1748-9326/10/1/014005>.
- Franco, C., L. A. Hepburn, D. J. Smith, S. Nimrod, and A. Tucker, 2016: A Bayesian Belief Network to assess rate of changes in coral reef ecosystems. *Environ. Model. Softw.*, **80**, 132–142, <https://doi.org/10.1016/j.envsoft.2016.02.029>.
- Franzke, C. L. E., and T. J. O’Kane, 2017: *Nonlinear and stochastic climate dynamics*. 1–466 pp.

- Frederiksen, J. S., and H. Lin, 2013: Tropical-extratropical interactions of intraseasonal oscillations. *J. Atmos. Sci.*, **70**, 3180–3197, <https://doi.org/10.1175/JAS-D-12-0302.1>.
- Friedman, N., M. Linial, I. Nachman, and D. Pe'er, 2000: Using Bayesian networks to analyze expression data. *Journal of Computational Biology*, Vol. 7 of, 601–620.
- Friedman, N., K. Murphy, and S. Russell, 1998: Learning the structure of dynamic probabilistic networks. *Proc. Fourteenth Conf. Uncertain. Artif. Intell.*, 139–147, <https://doi.org/10.1111/j.1469-7580.2008.00962.x>.
- Fyfe, J. C., 2019: Midlatitudes unaffected by sea ice loss. *Nat. Clim. Chang.*, **9**, 649–650, <https://doi.org/10.1038/s41558-019-0560-3>.
- Ghahramani, Z., 1997: Learning dynamic bayesian networks.
- Gong, T., S. Feldstein, and S. Lee, 2017: The role of downward infrared radiation in the recent arctic winter warming trend. *J. Clim.*, **30**, 4937–4949, <https://doi.org/10.1175/JCLI-D-16-0180.1>.
- Greatbatch, R. J., J. Lu, and K. A. Peterson, 2004: Nonstationary impact of ENSO on Euro-Atlantic winter climate. *Geophys. Res. Lett.*, **31**, <https://doi.org/10.1029/2003GL018542>.
- Hall, R. J., J. M. Jones, E. Hanna, A. A. Scaife, and R. Erdélyi, 2017: Drivers and potential predictability of summer time North Atlantic polar front jet variability. *Clim. Dyn.*, **48**, 3869–3887, <https://doi.org/10.1007/s00382-016-3307-0>.
- Hall, R., R. Erdélyi, E. Hanna, J. M. Jones, and A. A. Scaife, 2015: Drivers of North Atlantic Polar Front jet stream variability. *Int. J. Climatol.*, **35**, 1697–1720, <https://doi.org/10.1002/joc.4121>.
- Hamilton, L. C., and M. Lemcke-Stampone, 2014: Arctic warming and your weather: Public belief in the connection. *Int. J. Climatol.*, **34**, 1723–1728, <https://doi.org/10.1002/joc.3796>.
- Hardiman, S. C., N. J. Dunstone, A. A. Scaife, D. M. Smith, S. Ineson, J. Lim, and D. Fereday, 2019: The Impact of Strong El Niño and La Niña Events on the North Atlantic. *Geophys. Res. Lett.*, **46**, 2874–2883, <https://doi.org/10.1029/2018GL081776>.
- Handorf, D., R. Jaiser, K. Dethloff, A. Rinke, and J. Cohen, 2015: Figure S1. *Geophys. Res. Lett.*, **1218**, 180, <https://doi.org/10.1002/2015GL063203>.Received.
- Hanna, E., T. E. Cropper, R. J. Hall, and J. Cappelen, 2016: Greenland Blocking Index 1851–2015: a regional climate change signal. *Int. J. Climatol.*, **36**, 4847–4861, <https://doi.org/10.1002/joc.4673>.
- Hanna, E., T. E. Cropper, P. D. Jones, A. A. Scaife, and R. Allan, 2014: Recent seasonal asymmetric changes in the NAO (a marked summer decline and increased winter variability) and associated changes in the AO and Greenland Blocking Index. *Int. J. Climatol.*, **2554**, 2540–2554, <https://doi.org/10.1002/joc.4157>.

- Hanna, E., and Coauthors, 2014: Atmospheric and oceanic climate forcing of the exceptional Greenland ice sheet surface melt in summer 2012. *Int. J. Climatol.*, **34**, 1022–1037, <https://doi.org/10.1002/joc.3743>.
- Hanna, E., J. M. Jones, J. Cappelen, S. H. Mernild, L. Wood, K. Steffen, and P. Huybrechts, 2013: The influence of North Atlantic atmospheric and oceanic forcing effects on 1900–2010 Greenland summer climate and ice melt/runoff. *Int. J. Climatol.*, **33**, 862–880, <https://doi.org/10.1002/joc.3475>.
- Hanna, E., S. H. Mernild, J. Cappelen, and K. Steffen, 2012: Recent warming in Greenland in a long-term instrumental (1881–2012) climatic context: I. Evaluation of surface air temperature records. *Environ. Res. Lett.*, **7**, 045404, <https://doi.org/10.1088/1748-9326/7/4/045404>.
- Hassanzadeh, P., and Z. Kuang, 2015: Blocking variability: Arctic Amplification versus Arctic Oscillation. *Geophys. Res. Lett.*, **42**, 8586–8595, <https://doi.org/10.1002/2015GL065923>.
- Hay, J. E., D. Easterling, K. L. Ebi, A. Kitoh, and M. Parry, 2016: Conclusion to the special issue: Observed and projected changes in weather and climate extremes. *Weather Clim. Extrem.*, **11**, 103–105, <https://doi.org/10.1016/j.wace.2015.11.002>.
- Heckerman, D., 1996: A Tutorial on Learning With Bayesian Networks. *Innov. Bayesian Networks*, **1995**, 33–82, <https://doi.org/10.1007/978-3-540-85066-3>.
- Heim, R. R., 2015: An overview of weather and climate extremes - Products and trends. *Weather Clim. Extrem.*, **10**, 1–9, <https://doi.org/10.1016/j.wace.2015.11.001>.
- Henderson, S. A., E. D. Maloney, and E. A. Barnes, 2016: The influence of the Madden-Julian oscillation on Northern Hemisphere winter blocking. *J. Clim.*, **29**, 4597–4616, <https://doi.org/10.1175/JCLI-D-15-0502.1>.
- Hitchcock, P., and I. R. Simpson, 2014: The Downward Influence of Stratospheric Sudden Warmings*. *J. Atmos. Sci.*, **71**, 3856–3876, <https://doi.org/10.1175/jas-d-14-0012.1>.
- Hoeppe, P., 2016: Trends in weather related disasters - Consequences for insurers and society. *Weather Clim. Extrem.*, **11**, 70–79, <https://doi.org/10.1016/j.wace.2015.10.002>.
- Hoshi, K., J. Ukita, M. Honda, T. Nakamura, K. Yamazaki, Y. Miyoshi, and R. Jaiser, 2019: Weak Stratospheric Polar Vortex Events Modulated by the Arctic Sea-Ice Loss. *J. Geophys. Res. Atmos.*, **124**, 858–869, <https://doi.org/10.1029/2018JD029222>.
- Huang, J., and Coauthors, 2017: Recently amplified arctic warming has contributed to a continual global warming trend. *Nat. Clim. Chang.*, **7**, 875–879, <https://doi.org/10.1038/s41558-017-0009-5>.
- Hurrell, J. W., 1995: Decadal trends in the North Atlantic oscillation: Regional temperatures and precipitation. *Science (80-.)*, **269**, 676–679, <https://doi.org/10.1126/science.269.5224.676>.

- Jahn, M., 2015: Economics of extreme weather events: Terminology and regional impact models. *Weather Clim. Extrem.*, **10**, 29–39, <https://doi.org/10.1016/j.wace.2015.08.005>.
- Jiang, Z., S. B. Feldstein, and S. Lee, 2017: The relationship between the Madden–Julian Oscillation and the North Atlantic Oscillation. *Q. J. R. Meteorol. Soc.*, **143**, 240–250, <https://doi.org/10.1002/qj.2917>.
- Jung, T., M. A. Kasper, T. Semmler, and S. Serrar, 2014: Arctic influence on subseasonal midlatitude prediction. *Geophys. Res. Lett.*, **41**, 3676–3680, <https://doi.org/10.1002/2014GL059961>.
- Kapsch, M. L., R. G. Graversen, M. Tjernström, and R. Bintanja, 2016: The effect of downwelling longwave and shortwave radiation on Arctic summer sea ice. *J. Clim.*, **29**, 1143–1159, <https://doi.org/10.1175/JCLI-D-15-0238.1>.
- Kaufmann, R. K., M. L. Mann, S. Gopal, J. A. Liederman, P. D. Howe, F. Pretis, X. Tang, and M. Gilmore, 2016: Spatial heterogeneity of climate change as an experiential basis for skepticism. *Proc. Natl. Acad. Sci.*, 201607032, <https://doi.org/10.1073/pnas.1607032113>.
- Kidston, J., A. A. Scaife, S. C. Hardiman, D. M. Mitchell, N. Butchart, M. P. Baldwin, and L. J. Gray, 2015: Stratospheric influence on tropospheric jet streams, storm tracks and surface weather. *Nat. Geosci.*, **8**, 433, <https://doi.org/10.1038/NGEO2424>.
- Kim, B. M., and Coauthors, 2017: Major cause of unprecedented Arctic warming in January 2016: Critical role of an Atlantic windstorm. *Sci. Rep.*, **7**, <https://doi.org/10.1038/srep40051>.
- Kim, B.-M., S.-W. Son, S.-K. Min, J.-H. Jeong, S.-J. Kim, X. Zhang, T. Shim, and J.-H. Yoon, 2014: Weakening of the stratospheric polar vortex by Arctic sea-ice loss. *Nat. Commun.*, **5**, 4646, <https://doi.org/10.1038/ncomms5646>.
- Kintisch, E., 2014: Into the Maelstrom. *Science (80-.)*, **344**, 250–253, <https://doi.org/10.1126/science.344.6181.250>.
- Koller, D., and N. Friedman, 2009: *Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series)*.
- Konisky, D. M., L. Hughes, and C. H. Kaylor, 2016: Extreme weather events and climate change concern. *Clim. Change*, **134**, 533–547, <https://doi.org/10.1007/s10584-015-1555-3>.
- Kretschmer, M., D. Coumou, J. F. Donges, and J. Runge, 2016: Using causal effect networks to analyze different arctic drivers of midlatitude winter circulation. *J. Clim.*, **29**, 4069–4081, <https://doi.org/10.1175/JCLI-D-15-0654.1>.
- Kretschmer, M., J. Runge, and D. Coumou, 2017: Early prediction of extreme stratospheric polar vortex states based on causal precursors. *Geophys. Res. Lett.*, **44**, 8592–8600, <https://doi.org/10.1002/2017GL074696>.

- Lee, M. Y., C. C. Hong, and H. H. Hsu, 2015: Compounding effects of warm sea surface temperature and reduced sea ice on the extreme circulation over the extratropical North Pacific and North America during the 2013-2014 boreal winter. *Geophys. Res. Lett.*, **42**, 1612–1618, <https://doi.org/10.1002/2014GL062956>.
- Levine, R. A., and D. S. Wilks, 2006: Statistical Methods in the Atmospheric Sciences. *J. Am. Stat. Assoc.*, **95**, 344, <https://doi.org/10.2307/2669579>.
- Liang, S., S. Fuhrman, and R. Somogyi, 1998: Reveal, a general reverse engineering algorithm for inference of genetic network architectures. *Pac. Symp. Biocomput.*,
- Lin, H., G. Brunet, and B. Yu, 2015: Interannual variability of the Madden-Julian Oscillation and its impact on the North Atlantic Oscillation in the boreal winter. *Geophys. Res. Lett.*, **42**, 5571–5576, <https://doi.org/10.1002/2015GL064547>.
- Liptak, J., and C. Strong, 2014: The winter atmospheric response to sea ice anomalies in the barents sea. *J. Clim.*, **27**, 914–924, <https://doi.org/10.1175/JCLI-D-13-00186.1>.
- Luo, D., Y. Xiao, Y. Yao, A. Dai, I. Simmonds, and C. L. E. Franzke, 2016: Impact of Ural Blocking on Winter Warm Arctic–Cold Eurasian Anomalies. Part I: Blocking-Induced Amplification. *J. Clim.*, **29**, 3925–3947, <https://doi.org/10.1175/JCLI-D-15-0611.1>.
- Marcot, B. G., J. D. Steventon, G. D. Sutherland, and R. K. McCann, 2006: Guidelines for developing and updating Bayesian belief networks applied to ecological modeling and conservation. *Can. J. For. Res.*, **36**, 3063–3074, <https://doi.org/10.1139/x06-135>.
- McCusker, K. E., J. C. Fyfe, and M. Sigmond, 2016: Twenty-five winters of unexpected Eurasian cooling unlikely due to Arctic sea-ice loss. *Nat. Geosci.*, **9**, 838–842, <https://doi.org/10.1038/ngeo2820>.
- Messori, G., and Coauthors, 2018: An interdisciplinary approach to the study of extreme weather events: Large-scale atmospheric controls and insights from dynamical systems theory and statistical mechanics. *Bull. Am. Meteorol. Soc.*, **99**, ES81–ES85, <https://doi.org/10.1175/BAMS-D-17-0296.1>.
- Met Office, and CEH, 2014: The Recent Storms and Floods in the UK. *Met Off. UK*, 1–27, <https://doi.org/10.1038/scientificamerican04051913-315>.
- Michalewicz, Z., and D. Fogel, 2013: *How to solve it: modern heuristics*.
- Mihajlovic, V., and M. Petkovic, 2001: *Dynamic bayesian networks: A state of the art*. 1–37 pp.
- Milns, I., C. M. Beale, and V. Anne Smith, 2010: Revealing ecological networks using Bayesian network inference algorithms. *Ecology*, **91**, 1892–1899, <https://doi.org/10.1890/09-0731.1>.
- Montgomery, D. C., 2012: *Design and Analysis of Experiments Eighth Edition*. John Wiley and Sons Ltd,.

- Moore, R. W., O. Martius, and T. Spengler, 2010: The modulation of the subtropical and extratropical atmosphere in the Pacific basin in response to the Madden-Julian oscillation. *Mon. Weather Rev.*, **138**, 2761–2779, <https://doi.org/10.1175/2010MWR3194.1>.
- Mori, M., M. Watanabe, H. Shiogama, J. Inoue, and M. Kimoto, 2014: Robust Arctic sea-ice influence on the frequent Eurasian cold winters in past decades. *Nat. Geosci.*, **7**, 869–874, <https://doi.org/10.1038/ngeo2277>.
- Murphy, K. P., 1998: *Inference and Learning in Hybrid Bayesian Networks*. 18 pp.
- Murphy, K. P., **2001a**: *Introduction to Graphical Models*.
- Murphy, K. P., **2001b**: The Bayes Net Toolbox for Matlab. *Comput. Sci. Stat.*, **33**, 1024–1034, <https://doi.org/10.1.1.25.1216>.
- Murphy, K. P., 2004: Dynamic Bayesian Networks Representation, Inference And Learning - Kevin Patrick Murphy. 1–223.
- Murphy, K. P., 2012: *Machine learning: a probabilistic perspective (adaptive computation and machine learning series)*.
- Nakamura, T., K. Yamazaki, K. Iwamoto, M. Honda, Y. Miyoshi, Y. Ogawa, Y. Tomikawa, and J. Ukita, 2016: The stratospheric pathway for Arctic impacts on midlatitude climate. *Geophys. Res. Lett.*, **43**, 3494–3501, <https://doi.org/10.1002/2016GL068330>.
- Nanjundiah, R. S., 2014: Warming causes cooling! The recent cold event over North America. *Curr. Sci.*, **106**, 339–340.
- Norwegian Meteorological Institute, 2013: *Extreme Weather Events in Europe : preparing for climate change adaptation*.
- Overland, J. E., 2016: Is the melting Arctic changing midlatitude weather? *Phys. Today*, **69**, 38–43, <https://doi.org/10.1063/PT.3.3107>.
- Overland, J. E., and Coauthors, 2016: Nonlinear response of mid-latitude weather to the changing Arctic. *Nat. Clim. Chang.*, **6**, 992–999, <https://doi.org/10.1038/nclimate3121>.
- Overland, J. E., J. A. Francis, E. Hanna, and M. Wang, 2012: The recent shift in early summer Arctic atmospheric circulation. *Geophys. Res. Lett.*, **39**, 1–6, <https://doi.org/10.1029/2012GL053268>.
- Overland, J. E., and M. Wang, 2015: Increased variability in the early winter subarctic North American atmospheric circulation. *J. Clim.*, **28**, 7297–7305, <https://doi.org/10.1175/JCLI-D-15-0395.1>.
- Overland, J. E., and M. Wang, 2018: Resolving Future Arctic/Midlatitude Weather Connections. *Earth's Futur.*, **6**, 1146–1152, <https://doi.org/10.1029/2018EF000901>.

- Overland, J. E., K. R. Wood, and M. Wang, 2011: Warm Arctic-cold continents: Climate impacts of the newly open arctic sea. *Polar Res.*, **30**, 1–14, <https://doi.org/10.3402/polar.v30i0.15787>.
- Overland, J., J. A. Francis, R. Hall, E. Hanna, S. J. Kim, and T. Vihma, 2015: The melting arctic and midlatitude weather patterns: Are they connected? *J. Clim.*, **28**, 7917–7932, <https://doi.org/10.1175/JCLI-D-14-00822.1>.
- Pearce, F., 2013: The jet stream: slowing or no? *New Sci.*, **220**, 38–41, [https://doi.org/10.1016/S0262-4079\(13\)62536-1](https://doi.org/10.1016/S0262-4079(13)62536-1).
- Pedersen, R. A., I. Cvijanovic, P. L. Langen, and B. M. Vinther, 2016: The impact of regional Arctic sea ice loss on atmospheric circulation and the NAO. *J. Clim.*, **29**, 889–902, <https://doi.org/10.1175/JCLI-D-15-0315.1>.
- Perovich, D. K., J. A. Richeter-Menge, K. F. Jones, and B. Light, 2008: Sunlight, water, and ice: Extreme Arctic sea ice melt during the summer of 2007. *Geophys. Res. Lett.*, **35**, <https://doi.org/10.1029/2008GL034007>.
- Petoukhov, V., and V. A. Semenov, 2010: A link between reduced Barents-Kara sea ice and cold winter extremes over northern continents. *J. Geophys. Res. Atmos.*, **115**, 1–11, <https://doi.org/10.1029/2009JD013568>.
- Petoukhov, V., S. Petri, S. Rahmstorf, D. Coumou, K. Kornhuber, and H. JoachimSchellnhuber, 2016: The role of quasi-resonant planetary wave dynamics in recent boreal spring-to-autumn extreme events. *Proc. Natl. Acad. Sci.*, **113**, 6862–6867, <https://doi.org/10.1073/pnas.1606300113>.
- Petoukhov, V., S. Rahmstorf, S. Petri, H. J. Schellnhuber, and H. Joachim, 2013: Quasiresonant amplification of planetary waves and recent Northern Hemisphere weather extremes. *Pnas*, **110**, 5336–5341, <https://doi.org/10.1073/pnas.1222000110/-/DCSupplemental.www.pnas.org/cgi/doi/10.1073/pnas.1222000110>.
- Praetorius, S., M. Rugenstein, G. Persad, and K. Caldeira, 2018: Global and Arctic climate sensitivity enhanced by changes in North Pacific heat flux. *Nat. Commun.*, **9**, <https://doi.org/10.1038/s41467-018-05337-8>.
- Rahmstorf, S., D. Coumou, C. Voss, B. Zhao, S. S. Li, D. Coumou, M. Bianchi, and E. Baulieu, 2012: Correction for Cao et al., Differential regulation of the activity of deleted in liver cancer 1 (DLC1) by tensins controls cell migration and transformation. *Proc. Natl. Acad. Sci.*, **109**, 4708–4708, <https://doi.org/10.1073/pnas.1201491109>.
- Rayner, N. A., D. E. Parker, E. B. Horton, C. K. Folland, L. V. Alexander, D. P. Rowell, E. C. Kent, and A. Kaplan, 2003: Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century. *J. Geophys. Res.*, **108**, 4407, <https://doi.org/10.1029/2002JD002670>.
- Richards, R., M. Sanó, A. Roiko, R. W. Carter, M. Bussey, J. Matthews, and T. F. Smith, 2013: Bayesian belief modeling of climate change impacts for informing regional

- adaptation options. *Environ. Model. Softw.*, **44**, 113–121, <https://doi.org/10.1016/j.envsoft.2012.07.008>.
- Rigor, I. G., J. M. Wallace, and R. L. Colony, 2002: Response of sea ice to the Arctic Oscillation. *J. Clim.*, **15**, 2648–2663, [https://doi.org/10.1175/1520-0442\(2002\)015<2648:ROSITT>2.0.CO;2](https://doi.org/10.1175/1520-0442(2002)015<2648:ROSITT>2.0.CO;2).
- Robertson, A. W., A. Kumar, M. Peña, and F. Vitart, 2015: Improving and promoting subseasonal to seasonal prediction. *Bulletin of the American Meteorological Society*, Vol. 96 of, American Meteorological Society, ES49–ES53.
- Runge, J., and Coauthors, 2015: Identifying causal gateways and mediators in complex spatio-temporal systems. *Nat. Commun.*, **6**, 8502, <https://doi.org/10.1038/ncomms9502>.
- Runge, J., V. Petoukhov, and J. Kurths, 2014: Quantifying the strength and delay of climatic interactions: The ambiguities of cross correlation and a novel measure based on graphical models. *J. Clim.*, **27**, 720–739, <https://doi.org/10.1175/JCLI-D-13-00159.1>.
- Russell, S., and P. Norvig, 2014: *Artificial Intelligence: A Modern Approach, Third edition*. 33–36 pp.
- Samarasinghe, S. M., M. C. McGraw, E. A. Barnes, and I. Ebert-Uphoff, 2019: A study of links between the Arctic and the midlatitude jet stream using Granger and Pearl causality. *Environmetrics*, **30**, 1–14, <https://doi.org/10.1002/env.2540>.
- Sardeshmukh, P. D., and B. J. Hoskins, 1988: The generation of global rotational flow by steady idealized tropical divergence. *J. Atmos. Sci.*, **45**, 1228–1251, [https://doi.org/10.1175/1520-0469\(1988\)045<1228:TGOGRF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1988)045<1228:TGOGRF>2.0.CO;2).
- Scaife, A. A., and Coauthors, 2016: Seasonal winter forecasts and the stratosphere. *Atmos. Sci. Lett.*, **17**, 51–56, <https://doi.org/10.1002/asl.598>.
- Scaife, A. A., and Coauthors, **2017a**: Predictability of European winter 2015/2016. *Atmos. Sci. Lett.*, **18**, 38–44, <https://doi.org/10.1002/asl.721>.
- Scaife, A. A., and Coauthors, **2017b**: Tropical rainfall, Rossby waves and regional winter climate predictions. *Q. J. R. Meteorol. Soc.*, **143**, 1–11, <https://doi.org/10.1002/qj.2910>.
- Screen, J. A., and I. Simmonds, 2010: The central role of diminishing sea ice in recent Arctic temperature amplification. *Nature*, **464**, 1334–1337, <https://doi.org/10.1038/nature09051>.
- Screen, J. A., 2017: Simulated atmospheric response to regional and pan-arctic sea ice loss. *J. Clim.*, **30**, 3945–3962, <https://doi.org/10.1175/JCLI-D-16-0197.1>.
- Screen, J. A., 2014: Arctic amplification decreases temperature variance in northern mid- to high-latitudes. *Nat. Clim. Chang.*, **4**, 577–582, <https://doi.org/10.1038/nclimate2268>.

- Screen, J. A., and Coauthors, 2018: Consistency and discrepancy in the atmospheric response to Arctic sea-ice loss across climate models. *Nat. Geosci.*, **11**, 155–163, <https://doi.org/10.1038/s41561-018-0059-y>.
- Screen, J. A., C. Deser, and L. Sun, 2015: Projected changes in regional climate extremes arising from Arctic sea ice loss. *Environ. Res. Lett.*, **10**, 084006, <https://doi.org/10.1088/1748-9326/10/8/084006>.
- Screen, J. A., and I. Simmonds, 2013: Exploring links between Arctic amplification and mid-latitude weather. *Geophys. Res. Lett.*, **40**, 959–964, <https://doi.org/10.1002/grl.50174>.
- Screen, J. a., and I. Simmonds, 2014: Amplified mid-latitude planetary waves favour particular regional weather extremes. *Nat. Clim. Chang.*, **4**, 704–709, <https://doi.org/10.1038/nclimate2271>.
- Scutari, M., 2010: Learning Bayesian Networks with the bnlearn R Package. *J. Stat. Softw.*, **35**, 1–22, <https://doi.org/10.18637/jss.v035.i03>.
- Scutari, M., and J.-B. Denis, 2015: *Bayesian Networks With Examples in R*.
- Serreze, M. C., A. P. Barrett, J. C. Stroeve, D. N. Kindig, and M. M. Holland, 2009: The emergence of surface-based Arctic amplification. *Cryosphere*, **3**, 11–19, <https://doi.org/10.5194/tc-3-11-2009>.
- Shalizi, C. R., 2013: *Advanced data analysis from an elementary point of view*. 801 pp.
- Shepherd, T. G., 2014: Atmospheric circulation as a source of uncertainty in climate change projections. *Nat. Geosci.*, **7**, 703–708, <https://doi.org/10.1038/NGEO2253>.
- Smith, D. M., A. A. Scaife, R. Eade, and J. R. Knight, 2016: Seasonal to decadal prediction of the winter North Atlantic Oscillation: emerging capability and future prospects. *Q. J. R. Meteorol. Soc.*, **142**, 611–617, <https://doi.org/10.1002/qj.2479>.
- Spirtes, P., and C. Glymour, 1991: An Algorithm for Fast Recovery of Sparse Causal Graphs. *Soc. Sci. Comput. Rev.*, **9**, 62–72, <https://doi.org/10.1177/089443939100900106>.
- Stott, P. A., and Coauthors, 2016: Attribution of extreme weather and climate-related events. *Wiley Interdiscip. Rev. Clim. Chang.*, **7**, 23–41, <https://doi.org/10.1002/wcc.380>.
- Sumata, H., and K. Shimada, 2007: Northward transport of pacific summer water along the Northwind Ridge in the Western Arctic Ocean. *J. Oceanogr.*, **63**, 363–378, <https://doi.org/10.1007/s10872-007-0035-4>.
- Sun, J., and J. B. Ahn, 2015: Dynamical seasonal predictability of the Arctic Oscillation using a CGCM. *Int. J. Climatol.*, **35**, 1342–1353, <https://doi.org/10.1002/joc.4060>.
- Sun, L., J. Perlwitz, and M. Hoerling, 2016: What caused the recent “Warm Arctic, Cold Continents” trend pattern in winter temperatures? *Geophys. Res. Lett.*, **43**, 5345–5352, <https://doi.org/10.1002/2016GL069024>.

- Sun, Q., C. Miao, Q. Duan, H. Ashouri, S. Sorooshian, and K. L. Hsu, 2018: A Review of Global Precipitation Data Sets: Data Sources, Estimation, and Intercomparisons. *Rev. Geophys.*, **56**, 79–107, <https://doi.org/10.1002/2017RG000574>.
- Tang, Q., X. Zhang, and A. Francis, Jennifer, 2014: Extreme summer weather in northern mid-latitudes linked to a vanishing cryosphere. *Nat. Clim. Chang.*, **4**, 45–50, <https://doi.org/10.1038/nclimate2065>.
- Thompson, D. W. J., and J. M. Wallace, 2000: Annular modes in the extratropical circulation. Part I: Month-to-month variability. *J. Clim.*, [https://doi.org/10.1175/1520-0442\(2000\)013<1000:AMITEC>2.0.CO;2](https://doi.org/10.1175/1520-0442(2000)013<1000:AMITEC>2.0.CO;2).
- Trenberth, K. E., G. W. Branstator, D. Karoly, A. Kumar, N.-C. Lau, and C. Ropelewski, 1998: Progress during TOGA in understanding and modeling. *J. Geophys. Res. Ocean.*, <https://doi.org/10.1029/97JC01444>.
- Trifonova, N., A. Kenny, D. Maxwell, D. Duplisea, J. Fernandes, and A. Tucker, 2015: Spatio-temporal Bayesian network models with latent variables for revealing trophic dynamics and functional networks in fisheries ecology. *Ecol. Inform.*, **30**, 142–158, <https://doi.org/10.1016/j.ecoinf.2015.10.003>.
- Trifonova, N., D. Maxwell, J. Pinnegar, A. Kenny, and A. Tucker, 2017: Predicting ecosystem responses to changes in fisheries catch, temperature, and primary productivity with a dynamic Bayesian network model. *ICES J. Mar. Sci. J. du Cons.*, fsw231, <https://doi.org/10.1093/icesjms/fsw231>.
- Trifonova, N., M. Karnauskas, and C. Kelble, 2019: Predicting ecosystem components in the Gulf of Mexico and their responses to climate variability with a dynamic Bayesian network model. *PLoS One*, **14**, 1–23, <https://doi.org/10.1371/journal.pone.0209257>.
- Tseng, K. C., E. A. Barnes, and E. D. Maloney, 2018: Prediction of the Midlatitude Response to Strong Madden-Julian Oscillation Events on S2S Time Scales. *Geophys. Res. Lett.*, **45**, 463–470, <https://doi.org/10.1002/2017GL075734>.
- Tucker, a., and D. Duplisea, 2012: Bioinformatics tools in predictive ecology: applications to fisheries. *Philos. Trans. R. Soc. B Biol. Sci.*, **367**, 279–290, <https://doi.org/10.1098/rstb.2011.0184>.
- Tucker, A., and X. Liu, 2004: A Bayesian network approach to explaining time series with changing structure. *Intell. Data Anal.*, **8**, 469–480, <https://doi.org/10.3233/ida-2004-8504>.
- Uusitalo, L., 2007: Advantages and challenges of Bayesian networks in environmental modelling. *Ecol. Modell.*, **203**, 312–318, <https://doi.org/10.1016/j.ecolmodel.2006.11.033>.
- Uusitalo, L., M. T. Tomczak, B. Müller-Karulis, I. Putnis, N. Trifonova, and A. Tucker, 2018: Hidden variables in a Dynamic Bayesian Network identify ecosystem level change. *Ecol. Inform.*, **45**, 9–15, <https://doi.org/10.1016/j.ecoinf.2018.03.003>.

- Vavrus, S. J., M. Notaro, and D. J. Lorenz, 2015: Interpreting climate model projections of extreme weather events. *Weather Clim. Extrem.*, **10**, 10–28, <https://doi.org/10.1016/j.wace.2015.10.005>.
- Verma, T., and J. Pearl, 1991: Equivalence and synthesis of causal models. *Proc. Sixth Conf. Uncertain. Artif. Intell.*, <https://doi.org/10.1017/CBO9781107415324.004>.
- Vihma, T., 2014: *Effects of Arctic Sea Ice Decline on Weather and Climate: A Review*. 1175–1214 pp.
- Walsh, J. E., 2014: Intensified warming of the Arctic: Causes and impacts on middle latitudes. *Glob. Planet. Change*, **117**, 52–63, <https://doi.org/10.1016/j.gloplacha.2014.03.003>.
- Wang, S. Y. S., Y. H. Lin, M. Y. Lee, J. H. Yoon, J. D. D. Meyer, and P. J. Rasch, 2017: Accelerated increase in the Arctic tropospheric warming events surpassing stratospheric warming events during winter. *Geophys. Res. Lett.*, **44**, 3806–3815, <https://doi.org/10.1002/2017GL073012>.
- Whan, K., F. Zwiers, and J. Sillmann, 2016: The influence of atmospheric blocking on extreme winter minimum temperatures in North America. *J. Clim.*, **29**, 4361–4381, <https://doi.org/10.1175/JCLI-D-15-0493.1>.
- Wheeler, M. C., and H. H. Hendon, 2004: An All-Season Real-Time Multivariate MJO Index: Development of an Index for Monitoring and Prediction. *Mon. Weather Rev.*, **132**, 1917–1932, [https://doi.org/10.1175/1520-0493\(2004\)132<1917:aarmmi>2.0.co;2](https://doi.org/10.1175/1520-0493(2004)132<1917:aarmmi>2.0.co;2).
- Woollings, T., and M. Blackburn, 2012: The north Atlantic jet stream under climate change and its relation to the NAO and EA patterns. *J. Clim.*, **25**, 886–902, <https://doi.org/10.1175/JCLI-D-11-00087.1>.
- Woollings, T., A. Hannachi, and B. Hoskins, 2010: Variability of the North Atlantic eddy-driven jet stream. *Q. J. R. Meteorol. Soc.*, **136**, 856–868, <https://doi.org/10.1002/qj.625>.
- Yang, W., and G. Magnusdottir, 2017: Springtime extrememoisture transport into the Arctic and its impact on sea ice concentration. *J. Geophys. Res.*, **122**, 5316–5329, <https://doi.org/10.1002/2016JD026324>.
- Yates, D., B. Q. Luna, R. Rasmussen, D. Bratcher, L. Garre, F. Chen, M. Tewari, and P. Friis-Hansen, 2014: Stormy weather: Assessing climate change hazards to electric power infrastructure: A sandy case study. *IEEE Power Energy Mag.*, **12**, 66–75, <https://doi.org/10.1109/MPE.2014.2331901>.
- Ye, K., and T. Jung, 2019: How Strong Is Influence of the Tropics and Midlatitudes on the Arctic Atmospheric Circulation and Climate Change? *Geophys. Res. Lett.*, **46**, 4942–4952, <https://doi.org/10.1029/2019GL082391>.
- Yoo, C., S. Lee, and S. B. Feldstein, 2012: Mechanisms of arctic surface air temperature change in response to the Madden-Julian oscillation. *J. Clim.*, **25**, 5777–5790, <https://doi.org/10.1175/JCLI-D-11-00566.1>.

Yu, B., and H. Lin, 2016: Tropical atmospheric forcing of the wintertime North Atlantic Oscillation. *J. Clim.*, **29**, 1755–1772, <https://doi.org/10.1175/JCLI-D-15-0583.1>.

Zhang, J., R. Lindsay, M. Steele, and A. Schweiger, 2008: What drove the dramatic retreat of arctic sea ice during summer 2007? *Geophys. Res. Lett.*, **35**, <https://doi.org/10.1029/2008GL034005>.

Zucchini, W., and I. MacDonald, 2009: Hidden Markov Models for Time Series. *Hand, The*, **110**, 278, <https://doi.org/10.1201/9781420010893>.