

WINTER SEASONAL CLIMATOLOGY, VARIABILITY, AND
PREDICTION OF NORTHWEST ATLANTIC EXTRATROPICAL
CYCLONES

by

Rebekah Cavanagh

Submitted in partial fulfillment of the requirements
for the degree of Master of Science

at

Dalhousie University
Halifax, Nova Scotia
February 2022

© Copyright by Rebekah Cavanagh, 2022

To Mom and Dad, my tireless supporters

TABLE OF CONTENTS

List of Tables	vi
List of Figures	vii
Abstract	xi
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Background	2
1.2.1 Role of ETCs in the climate system	3
1.2.2 Energy partitioning in the atmosphere	5
1.2.3 Primary drivers of ETCs formation and development	7
1.2.4 Climatology and variability	8
1.2.5 Developing storm track datasets	10
1.2.6 Predicting ETCs	11
1.3 Objectives	13
Chapter 2 Detection and Tracking	15
2.1 Data	15
2.2 Storm Tracking Algorithm	16
2.2.1 Detection	17
2.2.2 Tracking	19
2.2.3 Pre- and Post-processing	19
2.3 Case Studies	22
2.3.1 New Year Storm 2018	23
2.3.2 White Juan 2004	23
2.4 Summary	25

Chapter 3	Climatology and Variability	26
3.1	Data	26
3.1.1	Storms	26
3.1.2	Atmospheric Variables	27
3.2	Methods	27
3.2.1	Gridded Track Density	27
3.2.2	Climatology	27
3.2.3	Trends	28
3.2.4	Empirical Orthogonal Functions Analysis	28
3.2.5	Self-Organizing Maps	29
3.2.6	Conditional Means of Atmospheric Variables	31
3.3	Extratropical Cyclone Climatology	31
3.3.1	Winter Season Mean	31
3.3.2	Monthly Means	33
3.3.3	Bomb Storms	33
3.3.4	Trends	33
3.4	Variability	36
3.4.1	Empirical Orthogonal Functions Analysis	37
3.4.2	Self-Organizing Maps Analysis	43
3.4.3	Links to atmospheric drivers and effects	50
3.5	Discussion	60
3.6	Conclusions	62
Chapter 4	Prediction	64
4.1	Data	65
4.1.1	Local Weather Observations	65
4.1.2	Predictands: Storm Time Series	68
4.1.3	Predictors: Atmospheric Variables	72
4.2	Prediction Model	74

4.2.1	Multiple Linear Regression	74
4.2.2	Predictor Selection	76
4.2.3	Model Fitting and Validation	81
4.2.4	Seasonal Forecast of Halifax Winter Storms	85
4.3	Summary	88
Chapter 5	Discussion	91
5.1	Summary of findings	91
5.2	Limitations	94
5.2.1	Storm detection and tracking algorithm	94
5.2.2	EOFs and SOMs	95
5.2.3	Prediction model	96
5.3	Conclusion	98
Appendix A	Full set of EOF results	100
A.1	Bomb Storms	101
Appendix B	Additional Conditional Means of Atmospheric Drivers and Ef-	
	fects	111
Appendix C	Further predictand time series analysis	115
Bibliography		121

LIST OF TABLES

3.1	Comparison of EOF results	40
3.2	RMSE analysis of SOM results	46
3.3	Seasonal groupings of SOM pattern 62.	48
4.1	Statistics of Halifax storm timeseries.	72
4.2	Selected predictor summary	81
4.3	Validation results for all eight predictand models.	85
4.4	Model parameters for each type of subseries forecast model.	86
4.5	Storm track activity categories	86
4.6	Comparison of quantitative and qualitative probabilities	87
4.7	Storm activity forecast for 2019 winter season	89
4.8	Storm activity forecast for 2020 winter season	89
4.9	Validation for 2019 and 2020 winter season forecasts	90
C.1	Correlations of the full time series and six subseries.	119
C.2	Highest lagged cross-correlations for time series pairs with max $\rho_x(h)$ at $h \neq 0$	120

LIST OF FIGURES

1.1	Meridional distribution of radiation	3
1.2	Schematic of momentum fluxes in the meridional plane	4
1.3	Schematic of the 4-box model of the energy cycle in the atmosphere	6
1.4	ETC climatology from (Varino et al., 2019)	9
2.1	Study area for storm detection and tracking	16
2.2	Variation of storm centre candidate locations resulting from different distance thresholds	19
2.3	Effect of input field smoothing on time series of algorithm-detected winter storms	20
2.4	Effect of smoothing of the input field on algorithm output	21
2.5	Complete 1979-2019 winter storm track dataset after applying all smoothing and filters	23
2.6	Comparison of algorithm representation of well-known winter storm tracks with observations	24
3.1	Winter season extratropical cyclone climatology	32
3.2	Monthly climatology of total storm density from 1979-2019	34
3.3	Winter season bomb cyclone climatology	35
3.4	Storm track density trends	36
3.5	EOF Spatial patterns	38
3.6	EOF principal component timeseries	39
3.7	Variance explained by EOF modes	39
3.8	Bomb storm EOF spatial patterns	41
3.9	Bomb storm EOF principal component timeseries	42
3.10	Variance explained by bomb storm EOF modes	42
3.11	Variation of quantization error with SOM parameters	45
3.12	Frequency of unique SOM pattern outputs	47
3.13	Anomalous storm counts for SOM-grouped seasons	49

3.14	Important features for the SOM-identified groups of seasons . . .	51
3.15	Surface air temperature mean and anomaly for SOM-grouped seasons	52
3.16	500 mb thickness mean and anomaly for SOM-grouped seasons . .	53
3.17	Mean 250 mb wind for SOM-grouped seasons	54
3.18	Anomalous 250 mb wind for SOM-grouped seasons	55
3.19	500 mb geopotential heights mean and anomaly for SOM-grouped seasons	56
3.20	Precipitable water mean and anomaly for SOM-grouped seasons .	58
3.21	Precipitation mean and anomaly for SOM-grouped seasons	59
3.22	Snowfall mean and anomaly for SOM-grouped seasons	60
3.23	Visual summary of atmospheric variable-storm track density rela- tionships in the SOM framework	61
4.1	Halifax weather station and nearest ERA5 grid point	65
4.2	Comparison of ECCC station data and ERA5 reanalysis outputs . .	66
4.3	Tracks of storms passing Halifax by storm type	71
4.4	Time series of storms passing Halifax by storm type	73
4.5	Visualization of prediction methods using linear regression models	74
4.6	Predictor selection process overview	77
4.7	Initial predictor filtering schematic	78
4.8	Multiple linear regression model training and validation process diagram	79
4.9	Final predictor processing flow chart	80
4.10	Locations of model predictors	82
4.11	Model output time series validation - A	83
4.12	Model output time series validation - B	84
4.13	Probabilistic forecast schematic	87
A.1	Spatial patterns of the first 4 EOF modes for PCA on main storm track	101
A.2	Principal component timeseries of the first 4 EOF modes for PCA on main storm track	102

A.3	Fraction of total variance explained by each principal component for PCA on main storm track	102
A.4	Unique storm centres km ⁻² for storms with origin south of 60°N on a 5 degree grid	103
A.5	Unique storm centres km ⁻² for storms with origin south of 60°N on a 5 degree grid, contoured	104
A.6	Spatial patterns of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts	105
A.7	Spatial patterns of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts	105
A.8	Principial component timeseries of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts	106
A.9	Principial component timeseries of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts	106
A.10	Spatial patterns of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts	107
A.11	Spatial patterns of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts with contours	107
A.12	Principial component timeseries of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts	108
A.13	Principial component timeseries of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts	108
A.14	Spatial patterns of the first 4 EOF modes for PCA on degree gridded bomb storm counts	109
A.15	Principial component timeseries of the first 4 EOF modes for PCA on degree gridded bomb storm counts	110
A.16	Principial component timeseries of the first 4 EOF modes for PCA on gridded bomb storm counts	110
B.1	Mean seasonal relative vorticity at 500 mb for SOM-grouped seasons (unit: 1/second)	111
B.2	Mean seasonal relative vorticity (contour) and anomaly from the climatological mean (colours) at 500 mb for SOM-grouped seasons (unit: 1/second)	112
B.3	Mean seasonal wind direction (arrows) and intensity (colours) at 10 m for SOM-grouped seasons (unit: meters/second)	113

B.4	Anomaly from the climatological mean of seasonal wind direction (arrows) and intensity (colours) at 250 mb for SOM-grouped seasons (unit: meters/second)	114
B.5	Average seasonal mean sea level pressure (contours) and anomaly from the climatological mean (colours) for SOM-grouped seasons (unit: mb)	114
C.1	Autocorrelation function of each subseries with significance shown by dashed line.	116
C.2	Cross-correlation functions for twelve time series combinations.	117
C.3	Cross-correlation functions for the remaining nine time series combinations.	118

ABSTRACT

Winter Extratropical Cyclones (ETCs) are mid-latitude storms that regularly impact the east coast of North America. These storms are characterized by high winds and heavy precipitation (rain, snow, and ice). ETCs are well predicted by NWP models at short- to mid-range forecast lead times (up to 5 days) however, on seasonal time scales these storms are not well predicted. The goal of this thesis is to first understand and then predict the seasonal variability of winter ETC activity along the east coast of North America. It begins with creating a dataset of ETCs from ERA5 mean sea level pressure. After calculating the climatology, the spatio-temporal variability of the dataset is analyzed to draw out the dominant patterns and trends. The covariability between storms and regional atmospheric variables is then examined to look for large scale forcing patterns. Those relationships form the basis of a set of storm type-specific multiple linear regression prediction models. This set of models is applied in a probabilistic forecast framework that predicts characteristics of the winter season experienced in Halifax, Nova Scotia, Canada.

CHAPTER 1

INTRODUCTION

1.1 Motivation

Extratropical cyclones (ETCs), also known as mid-latitude cyclones, wave cyclones, or baroclinic disturbances, threaten life and property in eastern North America every winter. Hurricane-force winds can down power lines, mixed precipitation can cause treacherous road conditions, storm surges can flood coastlines, and countless homes can be snowed in as these powerful synoptic-scale low pressure systems tear through a region. Fortunately, not every storm passing through has these extreme consequences. In fact, all ETCs are climatically critical for horizontal heat, momentum, and moisture transport ([Edmon et al., 1980](#)). System characteristics, intensity, and impacts vary from storm to storm and from winter season to winter season. In the last decade, insured damages from winter storms in the US adjusted for inflation have ranged between 684 million in 2012 to 4.4 billion in 2015 ([Insurance Information Institute, 2021](#)). This variability is a proxy for the combined effects of variability in storm frequency and intensity from season to season.

Furthermore, variations in the usual precipitation type brought by storms in different winter seasons cause very different impacts. For example, in Halifax, NS, the 2000 winter season had more than 150% of the average amount of snowfall and only 80% of the average rainfall. Compare this with the 1994 season that only had about 40% of average snowfall and 110% of average winter rainfall to see the differing impacts of winter season variability. A season with many snow storms will require larger budgets for road maintenance, while a season with more freezing precipitation storms may have a higher number of traffic accidents due to the poor road conditions. A season characterized by rain storms may cause flooding issues, while a winter of many high wind storms is likely to cause power outages.

From 2003-2012 approximately 80% of all power outages were caused by extreme weather events making storms the leading cause of outages (Kenward and Raja, 2014).

Extreme weather events often have the power to cause injury and loss of life, and mid-latitude cyclones are no exception. Deaths may occur through direct impacts such as falling trees which can crush victims, but more often fatalities result from secondary effects. The most common cause of winter storm related deaths is traffic accidents which occur due to poor road conditions, followed by heart attacks experienced when shoveling snow and hypothermia (NOAA National Severe Storms Laboratory, 2021). Power outages are likely to increase the severity of all risks present during storms by increasing the difficulty of contacting emergency services and adding higher risk of hypothermia while awaiting medical attention.

ETC tracks and impacts are typically well predicted by numerical weather prediction (NWP) models and operational forecasters on time scales of a few days. Prediction on this time scale is crucial for emergency preparedness and human safety, however a seasonal forecast would provide the opportunity for more extensive planning, preparation, and damage mitigation in the months leading up to the storm season. Thus, I set out to better understand and quantify the seasonal variability that exists and attempt to quantify the predictability to develop a winter season forecast for ETCs so our communities can be better prepared for whatever mother nature has in store each winter.

1.2 Background

On the scale of individual systems, ETCs are dynamically forced by synoptic scale atmospheric processes. However, to understand the forcing and effects of baroclinic storms collectively in the climate system, we must take a broader viewpoint. These complementary perspectives allow us to follow the story from the movements of an individual parcel within a storm through to the global circulation of the atmosphere. This section provides context for the role of storms in the global atmosphere, works its way down to individual storm dynamics, and ties in the climatology and variability of winter storms in the northwestern Atlantic. The potential predictability of the ETC field and previous efforts to forecast seasonal mid-latitude winter storm activity are also discussed.

1.2.1 Role of ETCs in the climate system

Extratropical cyclones are an integral piece of the global atmospheric circulation pattern, particularly meridional transport (Edmon *et al.*, 1980). Without meridional transport of both energy and momentum, the atmosphere as we currently observe it would be unbalanced. Shortwave and longwave radiation at the surface of the earth (Fig. 1.1) indicates that in the Tropics the net surface incoming solar energy (shortwave radiation) is much greater than the net emitted radiation (longwave radiation). Conversely, at the poles there is less energy emitted (longwave) than gained from the sun (shortwave). In the absence of any other transfers, this would cause the temperature in the tropics to increase and the temperature at the poles to continually decrease until some new equilibrium was reached. However, in reality the annual mean temperature at both the tropics and the poles does not exhibit such behaviour and is relatively stationary (climate variability notwithstanding). A poleward transfer of energy is required to explain the observed latitudinal temperature distribution. Ocean circulation is responsible for some of this energy transport (estimated about 40 %; Talley (2003)). However, atmospheric circulations are responsible for the majority of this poleward flux (Edmon *et al.*, 1980).




Figure 1.1: Meridional distribution of radiation (Pidwirny, 2006). Zonal average net surface shortwave radiation (blue) and longwave radiation (red). Shortwave radiation surplus shaded in yellow and deficit shaded in green.

On average, the atmosphere must conserve its angular momentum to maintain the consistently observed mid-latitude westerlies and tropical easterlies. At lower latitudes

where there are prevailing easterly winds, the angular momentum of the atmosphere (M_{atm}) is less than the solid earth (M_{earth}). Therefore, where they are in contact, the solid earth transfers momentum to the atmosphere. Conversely, at mid- to high-latitudes where the prevailing winds are westerlies, the atmosphere has greater angular momentum than the earth. Therefore, where they are in contact, the earth receives angular momentum from the atmosphere. In both cases, surface friction acts to decelerate the prevailing easterlies or westerlies. However, we observe that the upper level winds maintain their speed, rather than decreasing over time. In order for the winds to maintain their observed speed, there must be an atmospheric flux of angular momentum poleward. This transfer takes momentum from the tropical atmosphere (accelerating the easterlies) and gives it to the atmosphere at mid- to high-latitudes (accelerating the westerlies) to balance the deceleration caused by surface friction (Fig. 1.2).




Figure 1.2: Conceptual diagram of momentum fluxes in the meridional plane. Blue arrows show the transfer of momentum between the earth and atmosphere (from higher momentum to lower momentum) and the atmospheric momentum transfer that allows upper level winds to maintain their velocity.

The atmosphere transfers momentum and energy poleward through fluxes within mid-latitude cyclones. These transfers can be observed on aggregate over a season, but the amount of energy and momentum transferred poleward by an individual baroclinic disturbance makes a great enough contribution that on a single day in the life cycle of a storm the transfers are large enough to be observed (Edmon *et al.* (1980), their Fig. 3). Thus, the

general circulation of the atmosphere depends on mid-latitude baroclinic storms to weaken the meridional temperature gradient through vertical heat fluxes and to accelerate the tropical easterlies and mid-latitude westerlies through meridional momentum fluxes. These eddies force the residual meridional circulation which transfers energy and momentum poleward, thereby maintaining the balance of the atmosphere.

1.2.2 Energy partitioning in the atmosphere

The four box model of energy in the atmosphere (Fig. 1.3) represents the reservoirs or states of energy and the transfers between them. To understand the development, propagation, and decay of extratropical cyclones, it is important to analyze the forms of energy in the atmosphere and their variation in time (Chang and Fu, 2002). Energy in the atmosphere exists in two forms: kinetic energy, which is energy related to motion, and potential energy, which is static.

The potential energy of the atmosphere is due to a combination of gravitational potential energy and internal energy. Dynamics and thermodynamics in the atmosphere transfer energy between kinetic and potential energy which allows these energy pools to vary in time. However, there are physical limits to the amount of energy that can be in either form at a given time. Because of the configuration of the atmosphere, zero gravitational potential energy cannot be achieved, as it would require that the entirety of the atmosphere be contained at the surface. Similarly, zero internal energy cannot be achieved since an atmospheric temperature of absolute zero cannot be reached on the Earth. Therefore, the atmosphere must have a minimum potential energy level, below which values cannot be reached. Only potential energy in excess of this lowest potential energy state is available for conversion into kinetic energy and is, accordingly, called available potential energy (APE; Lorenz (1955)). The study of ETCs is concerned with kinetic energy and transfers into and out of this form, thus the potential energy of interest for the understanding of storm track dynamics is specifically APE. Considering the time mean atmosphere, the main states of energy are a mean background state and the eddy state which consists of energy that is anomalous from the mean state. That is to say that within the total kinetic energy in the atmosphere (KE), energy can be pooled into the mean state kinetic energy (MKE) and as some perturbation from that mean, the eddy component, being the eddy kinetic energy (EKE). As with kinetic energy, APE in the atmosphere can be partitioned into the climatological mean state (MAPE) and the eddy available potential energy (EAPE).

Thus, the traditional 4-box model is developed to describe the transfers of energy in the atmosphere among these partitions of MAPE, EAPE, MKE and EKE (Fig. 1.3).

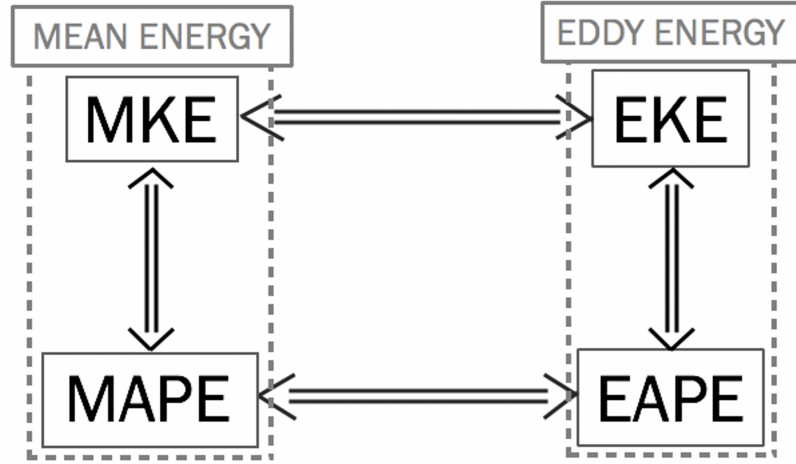


Figure 1.3: Schematic of the 4-box model of the energy cycle in the atmosphere. Energy is partitioned into mean kinetic energy (MKE), eddy kinetic energy (EKE), mean available potential energy (MAPE), and eddy available potential energy (EAPE). The energy of the mean state is the sum of MKE and MAPE and the total eddy energy (TEE) is the sum of EKE and EAPE.

The total eddy energy (TEE) describes the energy related to atmospheric disturbances. Transfers to and from this pool describe how energy is moved between the mean state and storms. The time averaged energy budget for total eddy energy is given in the following equation from [Orlanski and Katzfey \(1991\)](#):

$$\frac{\partial \overline{TEE}}{\partial t} = \frac{\alpha_m}{\theta_m} \overline{\frac{\mathbf{v}'\theta'}{d\theta/dp}} \cdot \nabla\theta_m - \overline{(\mathbf{v}' \cdot \mathbf{v}' \cdot \nabla\mathbf{V}_m)} - \mathbf{V}_m \cdot \nabla\overline{TEE} + \overline{\omega\alpha} - \overline{\mathbf{v}' \cdot \nabla\phi'} \quad (1.1)$$

The first term on the RHS describes the flux due to baroclinic instability occurring around the time of storm genesis, which is associated with transfers between MAPE and EAPE. The second term, describes the (negative) flux of energy due to barotropic conversion as the storm decays. It relates to transfers between MKE and EKE. Additional terms not shown include dissipative loss by mechanical dissipation and diabatic sources and sinks of EAPE. The remaining terms are related to the advection and propagation of total eddy energy describing the redistribution of eddy energy which occurs as the storm evolves between genesis and decay. The first and second term describe how the energy used in storm processes relates to the mean state. Energy is transferred from the

mean state into eddy energy by the first term. This term incorporates the vertical and horizontal temperature gradients ($d\tilde{\theta}/dp$ and $\nabla\theta_m$), and the flux of anomalous temperature by the anomalous wind ($\mathbf{v}'\theta'$). This highlights the importance of temperature gradients and advections in the transfer of energy from mean to eddy available potential energy which occurs at lower mid-latitudes. Through the third collection of terms, energy is redistributed within the pool. Energy is transferred back to the mean state from eddy energy in the second term. This term can be interpreted as Reynolds stresses when averaged ([Orlanski and Katzfey, 1991](#)). This describes how friction and turbulent motion causes storms to spin down and return their energy to the mean state, a transfer that occurs at higher latitudes. Thus, storms transfer energy (both momentum and temperature) poleward to maintain balance in the atmosphere.

1.2.3 Primary drivers of ETCs formation and development

The primary energy source of extratropical cyclones is baroclinicity. A baroclinic zone is a region where density varies with both temperature and pressure, causing isopycnic and isobaric surfaces to be unaligned. Baroclinic zones can be located in areas of temperature advections, i.e. where isotherms intersect with isobars on a horizontal surface. Baroclinic zones are also found in regions of strong horizontal temperature gradients, i.e. where isotherms are tightly packed. Even if the orientation matches that of the isobars, the 3D constant temperature and constant pressure surfaces will intersect to create misalignment. A region of strong temperature gradient where isotherms also cross the isobars is a combination of the two types of baroclinic zones described above. This is where the strongest baroclinic zones are commonly found. Not all baroclinic zones give rise to ETCs. There must also be some mechanism that causes upward motion to initiate the surface low. Most commonly, this is divergence at upper levels. The most common forcings of upper level divergence are positive vorticity advection at upper levels and ageostrophic circulations around upper level jet maxima. When upper level divergence occurs over regions of baroclinicity, advections are initiated which form fronts in a rotation around the surface low pressure centre that forms below the upper level divergence. Thus an extratropical cyclone is formed. The system continues to deepen until the upper level support is lost at which point the thermal wave breaks off from the low pressure centre, the low transitions from having a westward tilt with height to become vertically stacked, and convergence occurs at all levels filling the low and causing it to dissipate.

1.2.4 Climatology and variability

Not surprisingly, most ETCs are found where upper level support from the mechanisms of baroclinic instability is strong. They tend to form under the right entrance and left exit regions of upper level jets, on the lee side of mountain ranges, and on the east side of upper level troughs i.e. regions with strong upper level divergence. ETCs are also found along the east coasts of mid-latitude continental landmasses, regions of consistent temperature gradients with partial meridional orientation i.e. regions of strong baroclinicity.

In North America, high ETC frequency is found on the lee side of the Rocky Mountains, over the Great Lakes, and along the east coast, particularly northeastern USA and Atlantic Canada (Fig. 1.4; [Varino et al. \(2019\)](#); [Hoskins and Hodges \(2002\)](#); [Plante et al. \(2015\)](#)). Storms on the lee side of the Rockies are forced by increased vorticity as air flows over the mountains which leads to upper level divergence. For much of the winter season, the Great Lakes region is a baroclinic zone due to water surface temperatures being warmer than surrounding surface land temperatures. Due to the lee troughing caused by the Rocky Mountains, the Great Lakes region is often located under the east side of an upper level trough providing upper level divergence to trigger ETC genesis in the Great Lakes baroclinic zone. The east side of this upper level trough also frequently sets up over eastern North America supporting ETC genesis along the coastal baroclinic zone.

The winter extratropical cyclone field in the Northwest Atlantic varies by individual storm characteristics, overall storm frequency, and typical storm paths both intraseasonally and interannually ([Chang and Fu, 2002](#); [Hoskins and Hodges, 2019](#); [Varino et al., 2019](#); [Grise et al., 2013](#)). General storm track orientation tends to be northeastward, latitudinal, or slowly moving to the point of being quasi stationary ([Blender et al., 1997](#)). Individual storm characteristics that vary include lifetime, minimum central pressure, precipitation rates, and deepening rates which have been found to differ spatially and temporally ([Gulev et al., 2001](#); [Frankoski and Degaetano, 2011](#)). Spatial and temporal storm frequency has been shown to vary in relation with ENSO ([Plante et al., 2015](#)) and NAO ([Rogers, 1997](#)).

The most extreme ETCs, rapidly deepening systems that undergo explosive cyclogenesis, have been given the nickname bomb cyclones. To be considered in this category, the deepening rate of a mid-latitude cyclone must be at least 12 mb in 12 hours, at 45°N ([Sanders and Gyakum, 1980](#)). These storms are forced similarly to less extreme ETCs. However in bomb storms, the effects of some drivers are considerably larger than they

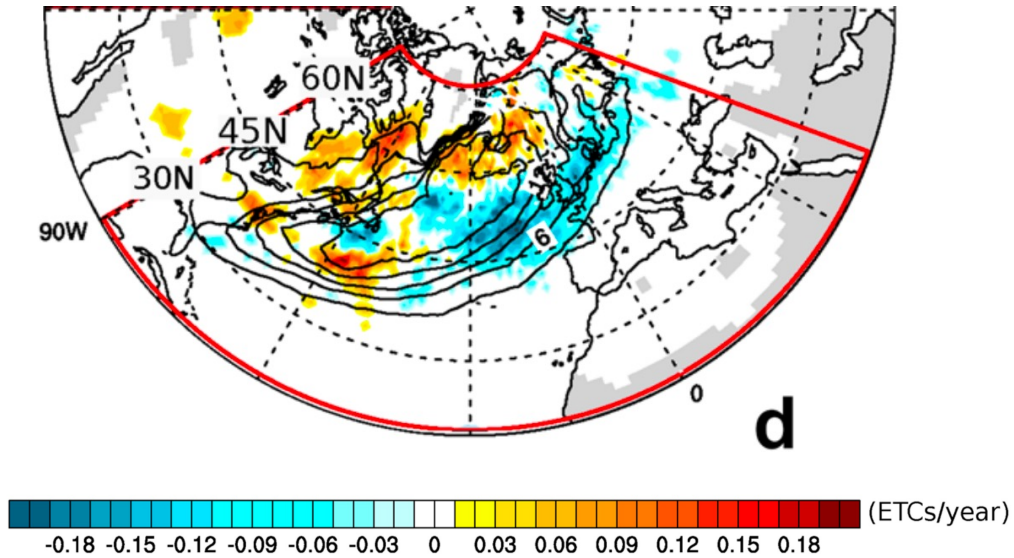


Figure 1.4: Number of Oct-Mar ETCs/year (tracked by 850 hPa vorticity max; black contours) and trends (colour) for 1980-2009 (Fig. 1d [Varino et al. \(2019\)](#)).

would be in a typical ETC leading to the observed rapid deepening of the system. The drivers known to play a large role in the explosive deepening include latent heat release, strong upper-level forcing, stratospheric intrusions of high potential vorticity air, strong upper-level forcing, sensible heating (air-sea fluxes), and differential diabatic heating which enhances baroclinicity ([Wang and Rogers, 2001](#)). Any particular bomb cyclone is forced by a specific combination of these drivers with the dominant mechanism for rapid development varying from system to system ([Uccellini, 1990](#)).

The majority of bomb cyclones in the area of interest track over the Atlantic Ocean ([Gulev et al., 2001](#); [Sanders and Gyakum, 1980](#); [Wang and Rogers, 2001](#)). This common track is collocated with regions of strongest differential diabatic heating, sensible heating, and moisture (necessary for latent heat release) which are found along the coast and over the ocean. Within the main Atlantic bomb storm track, [Wang and Rogers \(2001\)](#) identify two subgroups of storm tracks. One track is located in the west near North America and the other is in the east with storms that dissipate near Iceland. Variable characteristics in the storm duration and deepening process exist between in the two regions which is tied to differences in the forcing fields. Northeastern Atlantic bombs are shorter lived but more violent than their counterparts in the northwestern Atlantic. From winter season to winter season, the number of rapidly intensity Gulf Stream storms can vary by as much as 50%,

however a 3-year running mean gives relatively consistent yearly storm numbers ([Gulev et al., 2001](#)).

1.2.5 Developing storm track datasets

Any study of ETCs begins with developing a dataset of storms to study. This involves some method of detection or identification of storm centres followed by tracking of these centres through time. Prior to the computational era, this had to be done by hand. Storm centres would be identified by analysing a field of atmospheric measurements and the location of a system would be tracked through time on a map.

Modern detection and tracking makes use of computer algorithms to process atmospheric fields and produce a dataset of storm tracks. A plethora of methodologies exist to complete this task which are well explained and evaluated by [Walker et al. \(2020\)](#). There are two main methods for detecting and tracking in ETCs: Lagrangian feature tracking and Eulerian track identification. Eulerian track identification assesses activity in specific fields through a synoptic scale bandpass filter (e.g. 2-6 days). This window only allows activity on synoptic time scales to pass through thus isolating ETCs from the data. A variety of fields may be examined to reveal mid-latitude cyclone activity. These include 500 mb geopotential heights ([Blackmon, 1976](#)), mean sea level pressure, meridional wind at 850hPa, and relative vorticity at 850hPa ([Hoskins and Hodges, 2002](#)). The Eulerian method has the benefit of computational efficiency and smooth output fields.

The Lagrangian method is a two-step process consisting of an identification step and a tracking step. Identifying storm centres involves locating local extrema in an atmospheric field. While atmospheric variables used to locate storm centres vary from algorithm to algorithm, minima in mean sea level pressure (MSLP) and maxima in vorticity are the most commonly used definitions of a storm centre ([Walker et al., 2020](#)). Once a feature has been identified, it must be tracked through time and space. This issue, called the correspondence problem ([Post et al., 2003](#)), presents another opportunity for variation in method. Some algorithms use a nearest neighbour model which connects centres at adjacent time steps based on proximity ([Lakkis et al., 2019](#)), while others employ a cost function ([Massey, 2012](#)). Smoothing and filtering techniques are often used to refine the dataset before analysis begins. This primarily includes adding criteria for minimum duration or spatial extent of the track ([Hoskins and Hodges, 2019](#)). Lagrangian tracking retains details about each system at every time step. This gives the benefit of more detailed study because of

the opportunity to explore multiple aspects of the storm activity e.g. rate of deepening. For a more complete review of methods for developing ETC track datasets, refer to the review done by [Walker et al. \(2020\)](#) mentioned previously or by [Neu et al. \(2013\)](#).

1.2.6 Predicting ETCs

Short term forecasting in the range of 2-3 days generally anticipates extratropical cyclones well, although skill varies depending on the forcing ([Deveson et al., 2002](#)). The short-range forecast is essential. It provides detailed information about the incoming system and allows citizens to make necessary preparations. Unfortunately, some aspects of life related to winter storms require more than a few days of preparations. Seasonal forecasting allows for better resource planning. For example, government budgeting for snow removal and orders for salt or sand to treat the road surfaces must be made in advance. Power grid maintenance can also benefit from a seasonal storm forecast especially for high wind storms. If a major wind storm season is expected, the power company can upgrade or reinforce vulnerable sections of the grid to prevent extended outages during storm events. The insurance industry can also benefit from seasonal outlooks on potential damages caused by extratropical cyclones through a winter season.

Seasonal forecasts can also support short range forecasting and possibly extend lead times or increase confidence. It is common for 3-5 day forecasts to have enough skill to warn of an impending storm, but still have considerable uncertainty in the forecasted track of a system. In atmospheric modelling, the uncertainty of an outcome can be conveyed through probabilistic forecasting. The goal in such a forecast is to model or estimate the probability density function of the output for the atmospheric feature of interest, rather than forecasting a specific single value output. This becomes particularly useful when the nature of the system being forecasted makes it difficult to predict exactly. Ensemble forecasting is a common type of probabilistic forecast in which a model is run multiple times with slightly varying parameters to produce a set of outputs that approximate the probability density function of the forecasted outcome. When forecasting ETCs in the 3-5 day range, an ensemble forecast may indicate a high degree of uncertainty in the storm track. A higher degree of uncertainty may also be inferred from a lack of NWP model agreement or the presence of some confounding atmospheric phenomena. In the case of a forecast with high uncertainty, especially with extreme events, it is common practice among meteorologists to hedge toward the climatological mean. A good seasonal forecast

could provide an alternative to the climatological mean with greater skill. Rather than hedging toward the mean, hedging uncertain systems toward the expected seasonal mean track or activity level may give more accurate short range predictions. Finally, seasonal forecasts are of general interest to the public and can increase overall storm preparedness. A skilled seasonal forecast can have positive impacts in many areas.

In contrast with ETCs, seasonal forecasting of Atlantic hurricanes is a well-established field. The development of seasonal hurricane models has been lead by researchers at Colorado State University for many years ([Klotzbach and Gray, 2009](#)). It began with purely statistical modelling ([Gray, 1984](#)), and more recently, moved into the realm of statistical-dynamical modelling ([Klotzbach et al., 2020](#)). In general, current seasonal hurricane forecasts, such as the outlook produced by NOAA, are skillful and well trusted by the general public.

Seasonal forecasting of ETCs trails behind that of hurricanes. Current seasonal ETC forecasts are mainly derived from global circulation models (GCMs), empirical orthogonal functions (EOFs), or teleconnections such as NAO and ENSO. The most influential teleconnections for east coast winter storms are NAO and ENSO. While they show some effects on seasonal activity, the combined variability explained by these predictors is still small ([DeGaetano et al., 2002](#)). The most predictable area for seasonal ETC forecasting is located in the North Pacific ([Yang et al., 2015](#)). Many studies have identified predictability in this region ([Feng et al., 2019](#); [Yang et al., 2015](#); [Befort et al., 2018](#)). Windstorms near Europe have also been predicted with some skill ([Befort et al., 2018](#)). Most models fail to skillfully predict storms in the western North Atlantic ([Befort et al., 2018](#); [Yang et al., 2015](#); [Feng et al., 2019](#)). In addition, the majority of models are characterized by generality. The spatial extent of the prediction area is large, there is little spatial nuance added in the forecast, and little to no information about precipitation typing or other storm impacts is included. One existing model for extratropical cyclone prediction which specifically focusses along the US coast was developed by [DeGaetano et al. \(2002\)](#). It produces skillful forecasts of overall seasonal storm activity, but does not add detail around the expected storm types. The prediction efforts of my research will therefore focus on building in details about impacts and forecasting storms in areas where GCMs and other prediction models have limited skill: the east coast of North America.

1.3 Objectives

With the overall goal of this thesis being to develop a seasonal ETC forecasting system for storms in the northwest Atlantic, I've identified three primary enabling objectives:

- Develop a method to detect and track storms that occur within the region of interest;
- Determine a storm track climatology and interannual variability including links with regional climate drivers and impacts; and
- Develop a statistical seasonal forecast model of east coast extratropical cyclones using predictors identified above.

Each objective builds on those preceding it to culminate in the development of a regionally specific probabilistic forecast of the characteristics of storm track activity for the northwest Atlantic over an upcoming winter season. Each of the objectives map on to an individual thesis chapter.

Chapter 2 details the storm tracking algorithm, with particular attention to the choice of parameters and pre- and post-processing techniques. These are the steps used to fine-tune the dataset of tracks to be representative of the phenomena this work aims to describe, understand, and ultimately predict. Case studies are included to demonstrate the algorithm outputs.

Chapter 3 validates the dataset by evaluating the climatology of storm track density over the region between 25°N and 70°N and between 110°W and 0°. The full storm dataset is evaluated as well as subsets of storms based on temporal and intensity differences. A study of interannual variability follows by means of empirical orthogonal functions (EOF) analysis and self-organizing maps (SOMs). The clustering of the SOMs is used to connect regional climate drivers and effects of ETCs with variability in the storm field.

Chapter 4 develops the framework of a probabilistic winter season ETC forecast. This model is regionally specific to Halifax, Nova Scotia, Canada. After refining the set of predictands, the development process steps through a skill-based selection of predictors, model fitting, and validation. The final section of this chapter applies the model outputs in a practical probabilistic forecast designed specifically to provide useful, accessible information to the general public.

The efforts of this endeavor are reviewed in a final discussion chapter, which ties together the individual pieces of this work, places it in the context of the broader field of ETC study and applies the proposed probabilistic forecast to the practical field of meteorological forecasting on shorter and longer time scales.

CHAPTER 2

DETECTION AND TRACKING

The study of extratropical cyclones typically begins with a dataset of such storms. While datasets of ETCs are not widely available, extensive high-resolution datasets of mean sea level pressure are publicly accessible. Since ETCs are low pressure systems they can be identified and tracked using these fields. In this Chapter, the steps taken to turn sea level pressure fields into an extensive storm track dataset are explained.

2.1 Data

As in numerous other ETC studies, the storm data used in this study is derived from the outputs of an atmosphere reanalysis. Reanalysis data from the European Centre for Medium-Range Weather Forecasts Reanalysis v5 (ERA5) hourly single levels data from 1979 to present is obtained ([Hersbach et al., 2018b](#)). The variable used is mean sea level pressure, as in [Massey \(2016\)](#), at a temporal resolution of 3 hours, starting at 00:00 UTC. The grid spacing of the model is 0.25 degrees. Data was obtained between 25°N and 70°N and between 110°W and 0° (Fig. [2.1](#)). In this study, the extended winter storm season is defined as November 1 of the first year to March 31 of the following year. Data was obtained from November 1 of 1979 to March 31 of 2019. In total this is 40 winter seasons. The convention used when naming a season is to call it by the year in which it began. For example, the 1979 season spans the period of time from November 1, 1979 to March 31, 1980.




Figure 2.1: Study area for storm detection and tracking (outlined in red).

2.2 Storm Tracking Algorithm

There are two popular methods for detecting and tracking in ETCs: Lagrangian feature tracking and Eulerian track identification. The many implementations of these methods are reviewed well by [Walker et al. \(2020\)](#). Eulerian storm track identification uses bandpass filters to highlight activity in fields such as vorticity, MSLP, meridional wind, and geopotential heights on synoptic timescale (typically 2-6 days; e.g. [Hoskins and Hodges \(2002\)](#)). The variability that occurs in fields that are closely related to ETCs on such time scales can be attributed to the passing of an ETC through the region. This method captures the impacts of large and small storms alike and naturally aggregates the storm activity into a single field. The Lagrangian track identification process involves identifying an individual storm centre, which may be defined as a minimum in MSLP or positive vorticity maximum, at its genesis and then tracking the storm centre through time (e.g. [Catto \(2016\)](#)). The Eulerian method is useful for visualizing a smooth, overall picture of the storm field. The Lagrangian method is useful to observe characteristics of individual storms and investigate different subsets of the data such as storms with genesis in a particular region, storms that have certain length or size, or storms that cause specific effects (e.g. snow) at a given location ([Walker et al., 2020](#)). The Lagrangian tracking method was chosen for this study.

The principles of the storm tracking algorithm used here originate in the eddy tracking algorithm developed by [Chelton et al. \(2011\)](#) as applied by [Oliver et al. \(2015\)](#) which has been adapted here to track atmospheric storms.

2.2.1 Detection

Storms are detected in a mean sea level pressure (MSLP) field that was smoothed with a Gaussian filter with a 1-degree radius to remove noise. The pressure field comes from the ECMWF ERA5 hourly reanalysis dataset which has a $\frac{1}{4}^\circ$ grid spacing. The data is available hourly, however I used 3-hour temporal resolution for detection and tracking to reduce computational costs. At a single point in time the algorithm loops through a series of critical MSLP levels from highest to lowest. The array of critical MSLP levels begins at 1048 mb and goes down to 920 mb at decrements of 4 mb. At each level, the MSLP field is separated into grid cells with values above the critical level and grid cells with values below the critical level.

Storms are then identified as contiguous regions of pixels below the critical level that meet the following criteria:

- i. There is at least 1 local minimum (i.e. a viable storm centre).**
- ii. There are at least 9 pixels and not more than 6000 pixels within a region.**

A minimum of nine pixels is necessary in order to have a central pixel that is not also an edge pixel. Without an upper limit on the number of pixels that can make up a storm, the algorithm would identify the entire input field that is below the first critical level as one storm. On a $\frac{1}{4}^\circ \times \frac{1}{4}^\circ$ grid, a grid cell at 45°N is 560 km^2 so an area of 6000 pixels corresponds to $3.36 \times 10^6 \text{ km}^2$ or roughly the area of a circle with diameter of 2000 km. A synoptic system in the mid-latitudes typically has a length scale on the order of 1000 km ([American Meteorological Society, 2020](#)).

- iii. The relative amplitude of the storm is 100 Pa.**

The relative amplitude is defined as the difference between the outermost pressure contour of the storm and its central pressure. Requiring that this value is least 100 Pa eliminates any very weak low pressure regions that might have been detected. The value is based on the relative amplitude threshold used in [Chelton et al. \(2011\)](#) of 1 cm of sea surface height which corresponds to 1 hPa of atmospheric pressure. I

chose not to include an absolute threshold for the central pressure because it has the potential of removing storms in their genesis and decay stages from detection.

iv. The distance between any pair of points is no more than a maximum value that is calculated based on the area of the storm.

The area of a storm is defined as the total area within the outermost closed SLP contour that encloses the identified storm centre. Extratropical cyclones generally have a shape best approximated by an oval. If we assume the storm takes on the shape of an oval with a specified eccentricity, then given the measured area we can calculate the furthest distance apart that any two points in that system should be. This distance is the threshold value. Varying this threshold changes what features are detected as storms. For example, using a threshold that is too high (e.g. 0.98) allows storm centres that are minima in transient features rather than centres of an organized system to enter the dataset. Using a threshold that is too low causes storms that are not circular enough in shape to be discarded. Knowing that the tail of a typical, comma-shaped ETC can extend quite far from the storm centre, we need to be careful regarding the extent to which the eccentricity is restricted. The eccentricity threshold affects storm centre detection (Fig. 2.2). The detection step was run five times with different eccentricity thresholds, and a different set of centres was detected on each run. One can see how using the highest threshold, 0.98, causes the detection of areas that are not true storm centres. The lowest threshold, 0.80, fails to detect any centres that are not nearly circular. If the measured furthest distance between two pixels within the storm exceeds the threshold distance calculated from the eccentricity threshold, this tells us that the area identified is likely an elongated region of contiguous low pressure, rather than an organized storm system. If the algorithm fails to detect a true storm feature because it is contained within an unconventionally shaped area, it will be detected at a lower critical sea level as the algorithm continues.

After a storm is identified, the location of its centre is recorded along with the date and time. The pixels comprising the storm are then removed from consideration at all lower critical MSLP levels to avoid double detection.




Figure 2.2: Variation of storm centre candidate locations resulting from different distance thresholds. The MSLP field (algorithm input) is contoured at 4 hPa intervals and detected storm centres are marked with points. The detection step was run five times with different eccentricity thresholds ranging from 0.8 to 0.98 (legend). Marker colour indicates the eccentricity of the model run during which that storm centre was detected.

2.2.2 Tracking

Once possible storm centres have been identified at each timestep, they are stitched together across time to form tracks. For each storm centre at time t , the algorithm looks for a storm centre at time $t + 1$ within a radius of 240 km. This works out to a maximum propagation speed of 80 km/h. If there is one storm centre found at time $t + 1$ within a 240 km radius of the time t location, the time $t + 1$ and time t storm centres are stitched together as part of a track. If multiple centres are detected, the closest one to the time t location is chosen. If no centres are found, the storm is considered to be terminated. The details of the storm are then saved and no further timesteps can be added to its track. Once the storm track dataset had been created for the whole study period, track locations were interpolated to hourly resolution.

2.2.3 Pre- and Post-processing

As previously stated, the mean sea level pressure (MSLP) field is smoothed by a Gaussian filter with a 1-degree radius prior to detection and tracking. Smoothing levels were tested ranging from 0 (no smoothing) to 5 degrees. Across all smoothing levels, the shape of




Figure 2.3: Effect of input field smoothing on time series of algorithm-detected winter storms. Number of storms in each Nov-Mar storm season from MSLP input fields with no smoothing and smoothing by Gaussian filter with radius of 1-5 degrees.

the timeseries of storm counts per season remains relatively unchanged (Fig. 2.3). The mean number of storms detected per season decreases overall with increased smoothing. However, in a handful of seasons the total number of storms at a higher smoothing level exceeds the number of storms for a lower smoothing level (e.g. 1982 season - no smoothing: 340, 1 degree: 348). The correlation coefficients among the six timeseries range from 0.6898 to 0.9203. The main storm track also retains its shape throughout the smoothing processes (Fig. 2.4). These findings show that this smoothing does not remove the dominant patterns of the variability in the dataset.

After the tracking is complete, the dataset is refined by removing storms that obscure the signal we're interested in. Specifically, storms are removed if they have:

i. Genesis north of 60°N

This is implemented to cut out the signal from the Icelandic low, which is a persistent, quasistationary climatological feature in the NH winter. The algorithm picks up this feature as it is a strong minimum in MSLP. Since it exists in the same region at almost every timestep through the winter, the track densities in this region greatly exceed those at any other point in the study area thereby drowning out the other important signals in the field. One possible solution to this issue would be simply cutting the study area off south of the typical location of the Icelandic low, however

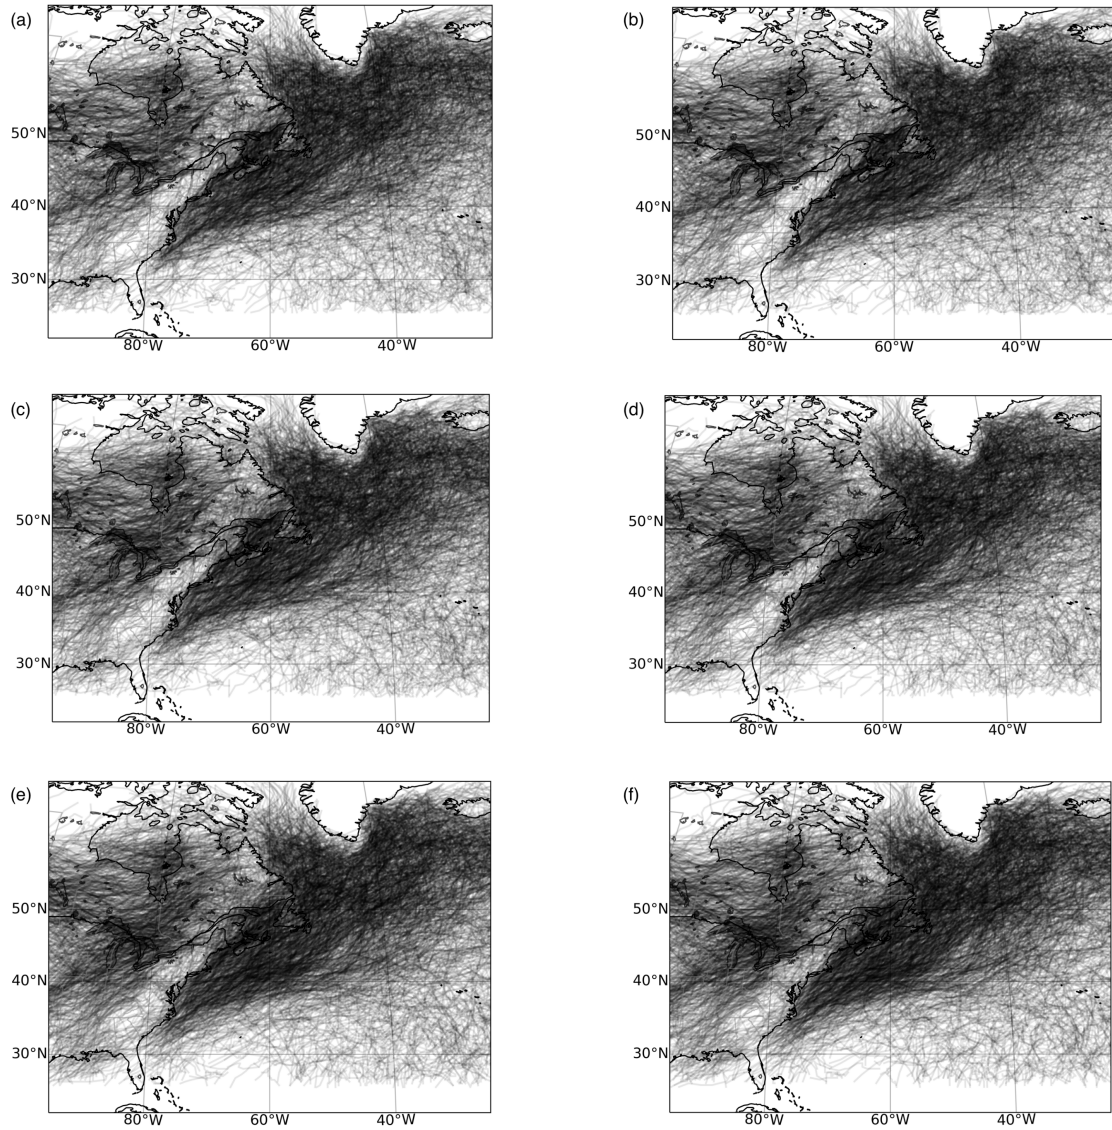


Figure 2.4: Effect of smoothing of the input field on algorithm output. (a) Detected storm tracks from unsmoothed spatial sea level pressure fields and from sea level pressure fields smoothed using a Gaussian filter of (b) one, (c) two, (d) three, (e) four, and (f) five degree radius.

this would result in cutting off the end of many systems within the climatological storm track causing loss of decay stage and termination-related information. Another possible way to remove this feature would be to use a criterion that required a storm to travel at least some minimum distance from genesis to termination, however I chose not to use such a criterion to avoid possible issues that this restriction could create for tracking stalled or very slow-moving systems.

Choosing to remove storms with genesis north of 60°N separates the Icelandic low tracks from the extratropical cyclone tracks. This methodology prevents the Icelandic low from increasing the storm track density in that region, but retains the contribution of extratropical cyclones that pass through that area.

ii. Duration less than 24 hours

This is a common practice used to remove very short term, noisy features that have been picked up by the tracking algorithm ([Hoskins and Hodges, 2002](#); [Massey, 2012, 2016](#); [Hoskins and Hodges, 2005](#); [Hodges et al., 2003](#); [Neu et al., 2013](#); [Pinto et al., 2016](#); [Raible et al., 2008](#)).

iii. Location above 1000 m Above Sea Level

Orographic effects on air circulations in mountainous regions create mesoscale low pressure centres that differ from the larger ETCs which are the focus of this study. For this region, the filtering removes all centres that are detected at a location where the surface elevation is greater than 1000 m above sea level before the tracking begins. Some of these leeward rotational features may develop into ETCs of interest. However if they do, the algorithm picks them up once they are below 1000 m elevation so they are not lost. Using this terrain filter removes the strong signal of lows detected in the mountains that tends to drown out the activity of interest along the climatological east coast storm track and secondary continental tracks.

2.3 Case Studies

The characteristics of individual storms as they are represented in the detection and tracking dataset described above were evaluated in comparison with manual tracking of the same systems based on observations. In general, the algorithm-derived storm tracks matched




Figure 2.5: Complete 1979-2019 winter storm track dataset after applying all smoothing and filters.

well with observed, manually created tracks. The results of this comparison for two well-known storms are included below.

2.3.1 New Year Storm 2018

This storm caused major snowfall throughout the Northeastern States, New Brunswick and Eastern Quebec. High winds also extended throughout Nova Scotia. Overall, this tracking algorithm identifies the New Years Storm of 2018 well. Although there was a break early in the storm track that caused this storm to be represented as two storms, its position and duration were quite accurate. Comparing with observations from the Ocean Prediction Centre, the central pressure of the storm recorded in the storm track dataset was within 10 mb of observations at all timesteps, and usually only a few millibars off.

2.3.2 White Juan 2004

In February 2004, the extratropical cyclone nicknamed White Juan hit the east coast of Canada causing long lasting blizzard conditions, widespread areas of 2-3 feet of snow, and extensive damage due to the high winds and heavy precipitation (Fogarty, 2004). The observed track is much straighter than the track from the dataset, but the overall trajectory

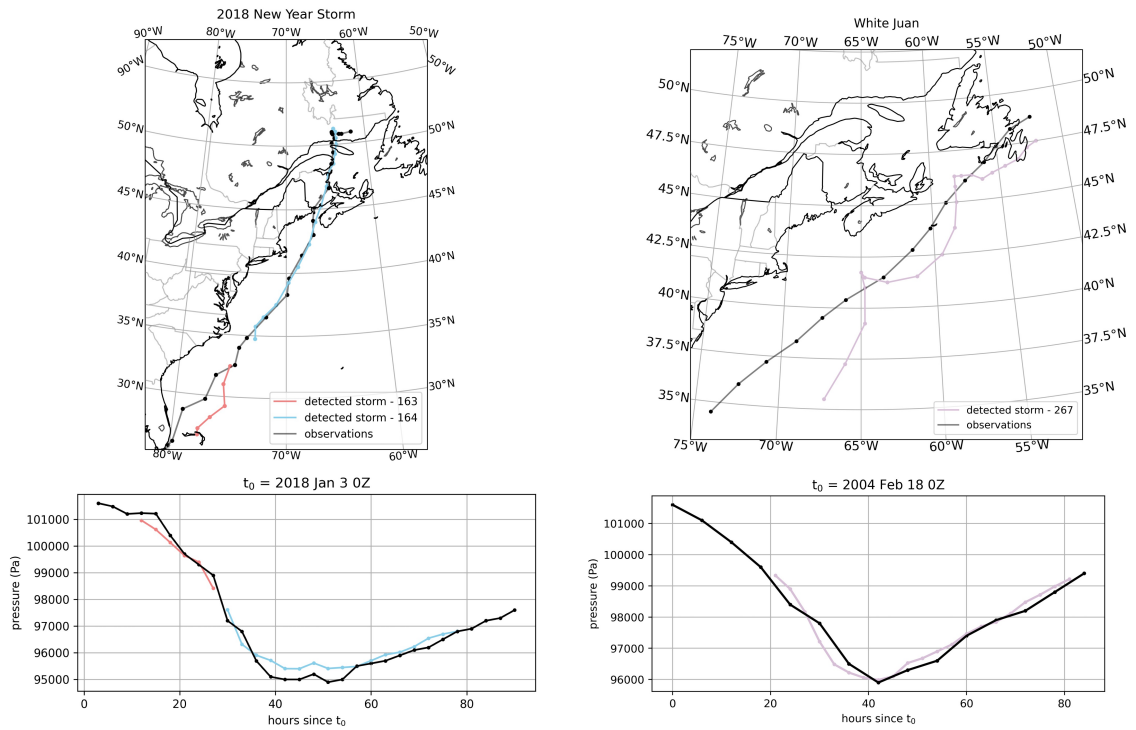


Figure 2.6: Comparison of algorithm representation of well-known winter storm tracks with observations. Observed and algorithm-detected tracks (lower panels) and central pressures (upper panels) of the New Years Storm 2018 (left) and White Juan 2004 (right).

is consistent. The central pressure has good agreement, especially at the lowest point where the observation and dataset have a difference of less than 1 mb. This storm track, as it was represented in this dataset, exhibited a hooking behaviour, travelling towards and then away from the coast as it tracked to the northeast. Upon further investigation it was confirmed that this was not an issue in the detection and tracking algorithm, it exists in the ERA5 MSLP field. This does not appear in the observed storm track (Fogarty, 2004). It was determined that the level of smoothing that would be required to remove the hooking would've smoothed out important features elsewhere in the dataset. Season-by-season analysis by eye revealed only a handful of other tracks in the dataset exhibited such behaviour, thus it was determined to be an insignificant issue which was not dealt with further.

2.4 Summary

The development of the dataset of extratropical cyclones for analysing climatology and variability, and predicting future characteristics of storm activity was described in this section. A Lagrangian feature tracking algorithm was adapted from an existing eddy tracking algorithm (*Oliver et al.*, 2015) to detect and track storm centres as minima in mean sea level pressure fields. Adaptation of the algorithm required the optimization of parameters for the specific purpose of detecting the characteristics of mid-latitude ETCs in Eastern North America and over the North Atlantic Ocean. This included making adjustments to both the detection and tracking portions of the algorithm. The resulting storm track dataset was assessed and the need for further pre- and post-processing was identified. Appropriate processing was added to refine the dataset of storms by removing features that are not pertinent to this study. To verify that the storms in the dataset were generally representative of real ETCs, specific storms were compared with observations and manual tracking techniques. Two of these comparisons were shown in section 2.3. It was determined that individual storms had accurate characteristics in terms of amplitude and geographical positioning. This establishes the dataset that is the foundation for the rest of this work. The climatology and variability of ETCs as they are represented in this dataset will be studied in Chapter 3 which will form the basis of the ETC prediction study in Chapter 4.

CHAPTER 3

CLIMATOLOGY AND VARIABILITY

A dataset of extratropical cyclones has been developed and now the spatiotemporal features present in this dataset are investigated. This begins with analyzing the number of storms and the spatial distribution. These details are compared with other literature to further validate the dataset. The climatology of monthly subsets of storms is then briefly explored in addition to a subset of more extreme ETCs called bomb cyclones.

After establishing the climatology of the system, a variability analysis follows. Both a traditional statistical (empirical orthogonal functions) and machine learning (self-organizing maps) approach are taken to tease apart patterns in both time and space. The identified patterns in storm track variability are then compared with the variability of known drivers and effects of ETCs to build a physical basis for the observed winter season to winter season differences.

3.1 Data

3.1.1 Storms

The data used to analyze the climatology and variability of ETCs is the storm track dataset detailed in Chapter 2. This dataset includes all storms that occurred during the winter seasons of 1979 through 2018 in the eastern North America and North Atlantic region. In total there are 7792 storms detected. The season with the most storms was the 1997 season which totaled 221 storms (Fig. 3.1a). The least active season overall was 1992 with only 159 storms (Fig. 3.1a). On average, there are 194.8 storms per season. Each storm in this dataset records its coordinates, central pressure, size, and the UTC time for each hourly timestep of the storm.

3.1.2 Atmospheric Variables

I compare the patterns identified within the storm field with a set of diagnostic atmospheric variables. The specific variables are mean sea level pressure, 2 m air temperature, 1000-500 mb thickness (calculated from geopotential heights), total precipitable water vapour, wind velocities at 250 mb and 10 m, 500 mb vorticity, hourly precipitation rates, and hourly snowfall rates. The data used is ERA5 hourly data on single levels ([Hersbach et al., 2018b](#)) and on pressure levels from 1979-2019 ([Hersbach et al., 2018a](#)). The spatial extent of this data is restricted to the study area: 25°N to 70°N and between 110°W and 0°W.

3.2 Methods

3.2.1 Gridded Track Density

The storm tracks were gridded on a 1° x 1° grid. The value of a grid cell in the climatology grid was increased by one each time the location of a storm centre from a unique storm was located within the cell. In the dataset developed in the previous section, each track of connected storm centres is considered one unique storm. The storms that passed through the specific grid cell of interest were totaled over each season within the study period. In order to ensure that a storm would not pass over a grid cell without recording a storm centre location within it, the tracks were interpolated down to one hour which, given typical storm speeds and the grid cell size, ensures a storm never skips over adjacent cells, where adjacent is defined as sharing a side or corner. After the climatology grid was populated, the total counts in each grid cell was divided by the area of the grid cell to account for the meridional decrease in grid cell size. Therefore, each grid cell has units of storm counts per square kilometre, which is referred to as storm density.

3.2.2 Climatology

A winter storm season climatology is calculated with the season defined as November 1 to March 31 inclusive where the season is named by the year in which it begins (i.e. the 2015 season is Nov. 1 2015 – Mar. 31 2016). In addition, a monthly climatology is calculated for each month within the winter season. For an individual month, the climatology is determined by summing the storm count totals in that individual month over all seasons. To ensure a storm that occurred in two calendar months was not counted twice in the monthly

climatology analysis, it was stipulated that in order for a storm to contribute toward the monthly storm total, it then must exist in that month for the majority of its lifetime.

Bomb cyclones or bomb storms are classified as systems that experience a central pressure drop of 12 mb in 12 hours (Sanders and Gyakum, 1980). The rate of pressure decrease was calculated for every 12-hour period over the storm's lifetime. If there was at least one 12-hour period over which the central pressure of the storm decreased 12 mb or more, the storm was added to the bomb cyclone dataset. The climatology of bomb cyclones was also calculated as a gridded track density according to the method defined above.

3.2.3 Trends

To analyze long-term, temporal trends in storms, each grid cell was considered individually. The gridding and averaging processes created an annual time series of seasonal storm counts for each grid cell. Using this time series, a simple linear regression was executed to determine the best fit trend line through the timeseries. The slope of this trend line in each grid cell is shown in section 3.3.4. The significance of the slope is assessed using a 90% confidence interval. For each slope calculation, the upper and lower bound of the 90% confidence interval are identified. If a slope of zero is within the range given by the confidence interval, the null hypothesis cannot be rejected and the calculated slope is considered not to be significant at the 10% significance level.

3.2.4 Empirical Orthogonal Functions Analysis

An empirical orthogonal functions (EOF) analysis decomposes a spatiotemporal dataset into a set of statistically independent modes of variability. It has the potential to shed light on predictability, as it gives spatial patterns with associated timeseries for each identified mode (Monahan et al., 2009).

To use EOF analysis, as in Monahan et al. (2009), all grid cells in the field that do not have any storms over the whole season are removed from consideration and then the storm field is converted into anomalies. The anomalous storm count for a grid cell is calculated as the difference in counts of storms in that grid cell from the mean counts of storms per season over the whole field, $Z'_{x,y,t} = Z_{x,y,t} - \bar{Z}_{x,y}$. The anomalous storm field which is a function of latitude, longitude, and time, $Z'(x, y, t)$, is then reduced to a space by time field or N-dimensional vector timeseries, $\mathbf{X}(t)$ where N is equal to the number of active grid cells in the spatial field.

The empirical orthogonal functions of the anomalous storm field are defined as the eigenvectors of the covariance matrix, \mathbf{C} , of $\mathbf{X}(t)$. Given that \mathbf{C} is a real symmetric $N \times N$ matrix, its N eigenvectors form a complete orthonormal basis. One can thus at any time, t , write the anomaly state vector, \mathbf{X} , as a linear combination of the eigenvectors according to

$$\mathbf{X}(t) = \sum_{n=1}^N \alpha_n(t) \mathbf{e}_n, \quad (3.1)$$

where \mathbf{e}_n is the n^{th} eigenvector and the coefficient α_n is the projection of the anomalous state vector at that time onto the eigenvector. Remapping the space dimension of the state vector \mathbf{X} back to latitude and longitude and by the same process mapping each eigenvector onto latitude and longitude gives

$$Z(x, y, t) = \sum_{n=1}^N PC_n(t) EOF_n(x, y). \quad (3.2)$$

where the n^{th} principal component timeseries ($PC_n(t)$) is the time series $\alpha_n(t)$ and the n^{th} EOF spatial pattern ($EOF_n(x, y)$) is the eigenvector \mathbf{e}_n reshaped into dimensions of latitude by longitude. The n^{th} PC time series and EOF spatial pattern pair are together called an EOF mode or simply a mode. Finally, the eigenvalue that corresponds to each mode, λ_n , when divided by the sum of all eigenvalues, gives the fraction of the variance explained by that mode, ρ_n :

$$\rho_n = \frac{\lambda_n}{\sum_{n=1}^N \lambda_n}. \quad (3.3)$$

3.2.5 Self-Organizing Maps

Self-organizing maps (SOM) analysis is a machine learning technique with the primary function of pattern sorting ([Kohonen, 1990](#); [Hewitson and Crane, 2002](#); [Oliver et al., 2018](#); [Liu and Robert, 2011](#)). When given a set of images or patterns, a SOM algorithm identifies common themes within the set and sorts the inputs into groups that are similar. Each of these groups is called a node. Similarities between groups are captured by the node locations relative to each other. The node placement is not random. Each node bears similarities to the nodes to which it is adjacent. The further apart two nodes are in the arrangement, the less similar those groupings are. This set of related nodes is the “map”. Thus, the storm counts fields self-organize into groups that have similar characteristics. I

chose to run a 2 by 2 SOM which sorts the data into four groupings that are arranged in a square formation so that a given node is more similar to the nodes directly adjacent to it than to the node that is diagonal from it. I used the minisom package for python created by Giuseppe Vettigli available on GitHub ([Vettigli, 2018](#)).

The main SOM input is the set of gridded seasonal totals of storm counts for each of the 40 seasons that occur between November 1979 and March 2019. These are the patterns that the algorithm is tasked with sorting into 4 groups. The main SOM output is 4 lists of seasons, one for each node, where the seasons within each list have similar characteristics.

To begin, each node is initialized as a vector or matrix that has the shape of the observation field. I initialize the weights of the nodes with the first four EOF modes determined in the EOF analysis section ([3.4.1](#)). This gives the algorithm a starting point for sorting the inputs. Since the EOF modes are informed by the variability of the field, I expect it would take the algorithm fewer iterations to get to a given level of refinement when the weights are initialized using the EOF modes over random initialization.

With initialization weights set, the algorithm is trained by iterating over the full set of observations. On each iteration, the weights of the nodes are adjusted with the goal of minimizing the euclidean distance between the observation and the weights of the nodes. Over the course of training, the weights are refined to become more and more representative of the patterns of observations in the input dataset. The algorithm requires that the user give parameters pertaining to this refinement process. The initial learning rate, which determines how quickly the weights are adjusted to fit the inputs, and number of iterations to train over were set at 0.4 and 6000, respectively. Together, these parameters determine how and to what extent the nodes are refined before the algorithm gives the final output. To specify how closely each node is related to the nodes around it, a Gaussian neighbourhood function is used with an initial standard deviation of 0.3. The details regarding the selection of these parameters can be found in subsection [3.4.2.1](#).

At the end of training, each node shows a pattern that is characteristic of the set of observations (seasons) associated with that node. The SOM output also includes a list of the “winning node” for each input observation. This is the node to which the input is most similar. This set of outputs gives M groups each containing a list of seasons with similar observations, where M is the number of nodes. There is no requirement for the inputs to be evenly distributed between nodes. The final step is to visualize the SOM groupings.

For each node, the mean storm counts per square kilometre in each grid cell is calculated by averaging over the seasons in that node only. This illustrates the typical pattern of a season that belongs in that grouping. To further highlight the differences among the four groupings, the temporal mean field of storm counts km^{-2} is removed to obtain a grid of anomalous storm counts km^{-2} in each node.

3.2.6 Conditional Means of Atmospheric Variables

To link the patterns of variability identified through the SOM analysis with known drivers and effects of ETCs I calculate conditional means of atmospheric variables based on the SOM nodes. I also set out to explore how related atmospheric variables behave within those same groupings of seasons. For each atmospheric variable considered (section 3.1.2) the seasonal mean was calculated for each winter season of the study period from the ERA5 hourly reanalysis data. For each SOM-identified group, the seasonal mean from the years within the group were averaged to calculate the node mean for that atmospheric variable. This gives four independent conditional means. To further highlight the differences between these four mean states, the full time period mean of the atmospheric variable is subtracted from each conditional mean to get anomalies.

3.3 Extratropical Cyclone Climatology

3.3.1 Winter Season Mean

The annual mean winter storm track density informs us about the basic climatological background state of ETCs and its interannual variability (Fig. 3.1a). The main track begins off the coast of the Carolinas and propagates northeastward following the coast, over Nova Scotia and Newfoundland, then branches west and east around Greenland (Fig. 3.1b). A secondary storm track begins around northwest Texas and propagates northeastward over the Great Lakes and into northern Quebec. Other regions of high track density exist over southern Saskatchewan and Manitoba and southeast Hudson Bay including James Bay. This generally matches the literature across a variety of tracking techniques and storm centre definitions (Hoskins and Hodges, 2002, 2019; Varino et al., 2019; Plante et al., 2015). The time series of total storms per season over the whole study area has a mean value of 194.8 storms with a standard deviation of 13.5 (Fig. 3.1a). Note the season with the most storms is Nov 1997 - Mar 1998 and the season with the fewest storms is Nov 1992

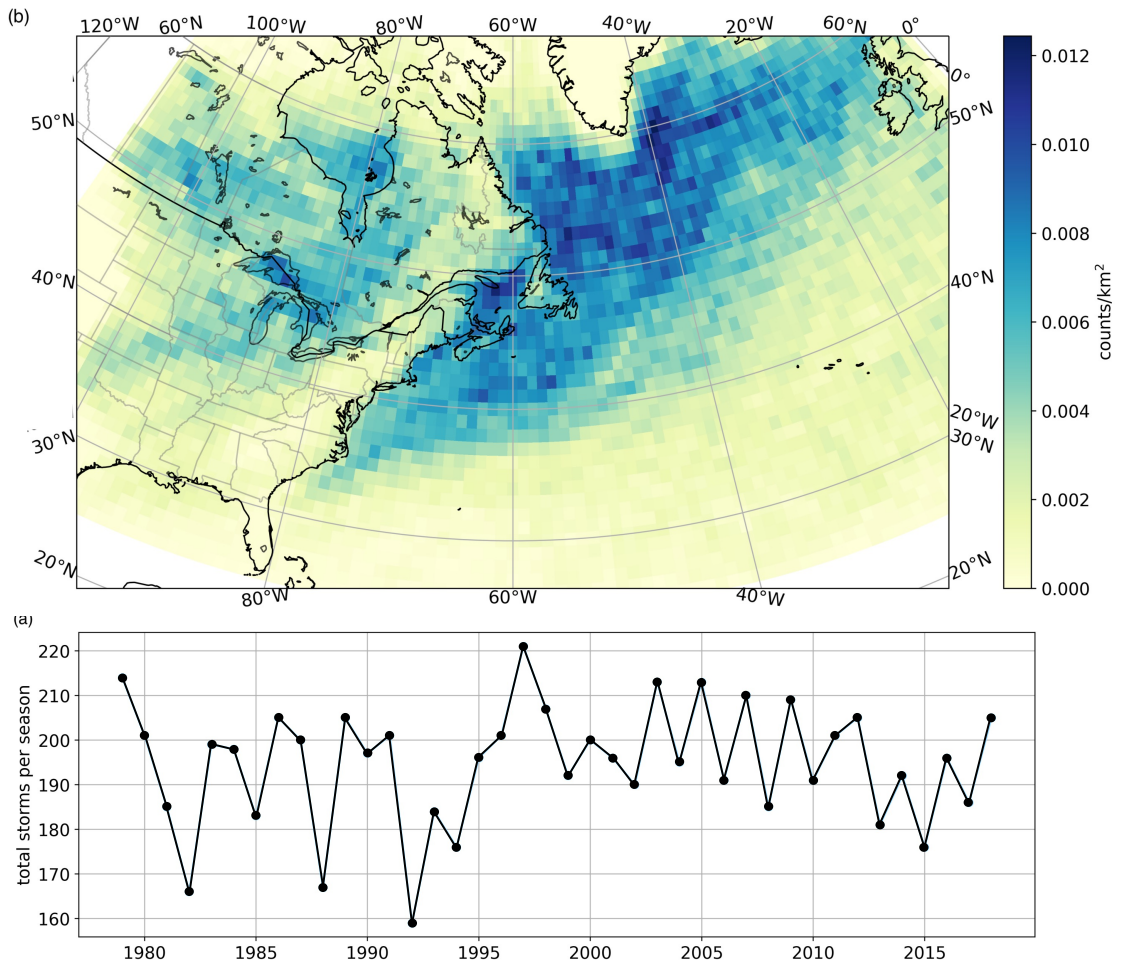


Figure 3.1: Winter season extratropical cyclone climatology. (a) Time series of winter storms per Nov-Mar season. (b) Spatial distribution of total storm density from 1979-2019.

- Mar 1993. The difference between those two seasons is 62 storms. In some cases, there are differences between consecutive seasons of up to 50 storms per season. There appears to also be some long-term trends underlying this high frequency variability, however this was not explored within the scope of this project.

3.3.2 Monthly Means

The monthly climatology reveals intraseasonal variability indicating that storm season characteristics evolve from November to March (Fig. 3.2). The November climatology is characterized by proportionately more storms over the continent and Hudson Bay. As the season progresses, track density increases along the east coast of North America and the climatological track seems to increase in width into the Atlantic Ocean through the later part of the season. The primary storm track extends further south with colder temperatures in January and February.

3.3.3 Bomb Storms

These rapidly deepening storms are concentrated along the east coast of North America from the Mid-Atlantic Bight to Labrador, and, also extending from Newfoundland to east of Greenland (Fig. 3.3b). These systems are not common over the continent. The time series of total bomb storms per season has a mean of 35.2 storms per season and a standard deviation of 6.20 (Fig. 3.3a). A linear fit line has a slope of 0.068 storms per square kilometre per season indicating no trend at the 95% confidence level. Looking at the correlation coefficient, the seasonal variability of these extreme storms is not obviously related to the total seasonal storm counts. The correlation coefficient between the two time series is -0.13 which is not a statistically significant correlation ($p=0.42$). There are multiple seasons with high numbers of bomb cyclones, but average total storm counts (e.g. 1996, 2006) however, there are other seasons in which counts of bomb storms and total counts are both low (e.g. 1982, 1995) or both high (e.g. 1989, 2014).

3.3.4 Trends

The linear trend in storm track density over the study period for each grid cell shows a very heterogeneous pattern (Fig. 3.4). The majority of the trends identified are not statistically significant with a 90% confidence interval. However, I have some confidence in a trend that is found along the east coast of the USA. Noting that the climatological storm track

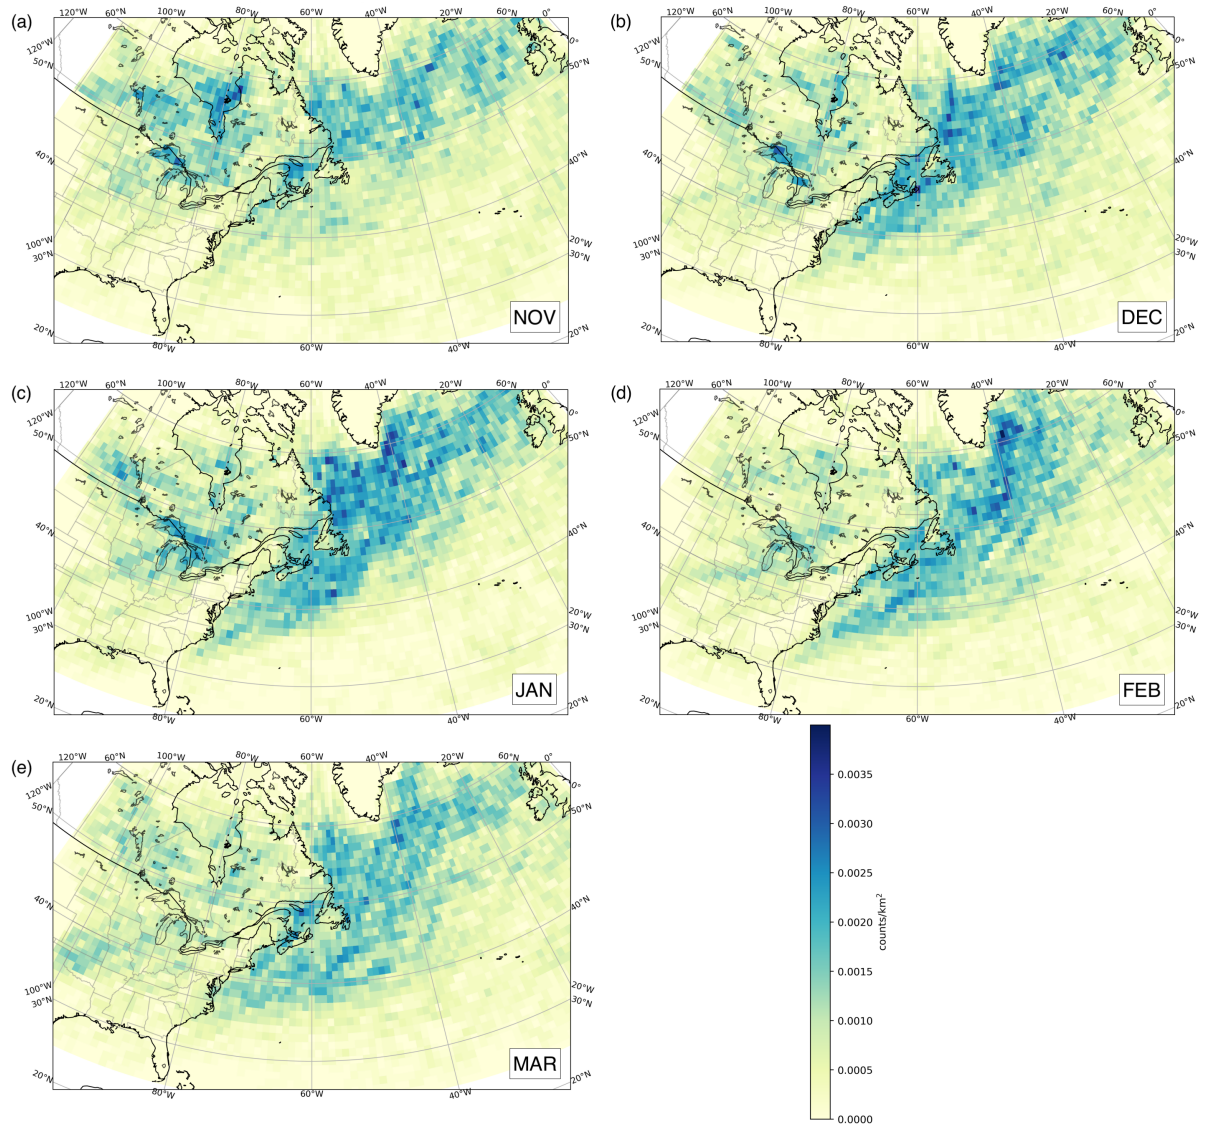


Figure 3.2: Monthly climatology of total storm density from 1979-2019. (a) November, (b) December, (c) January, (d) February, (e) March.

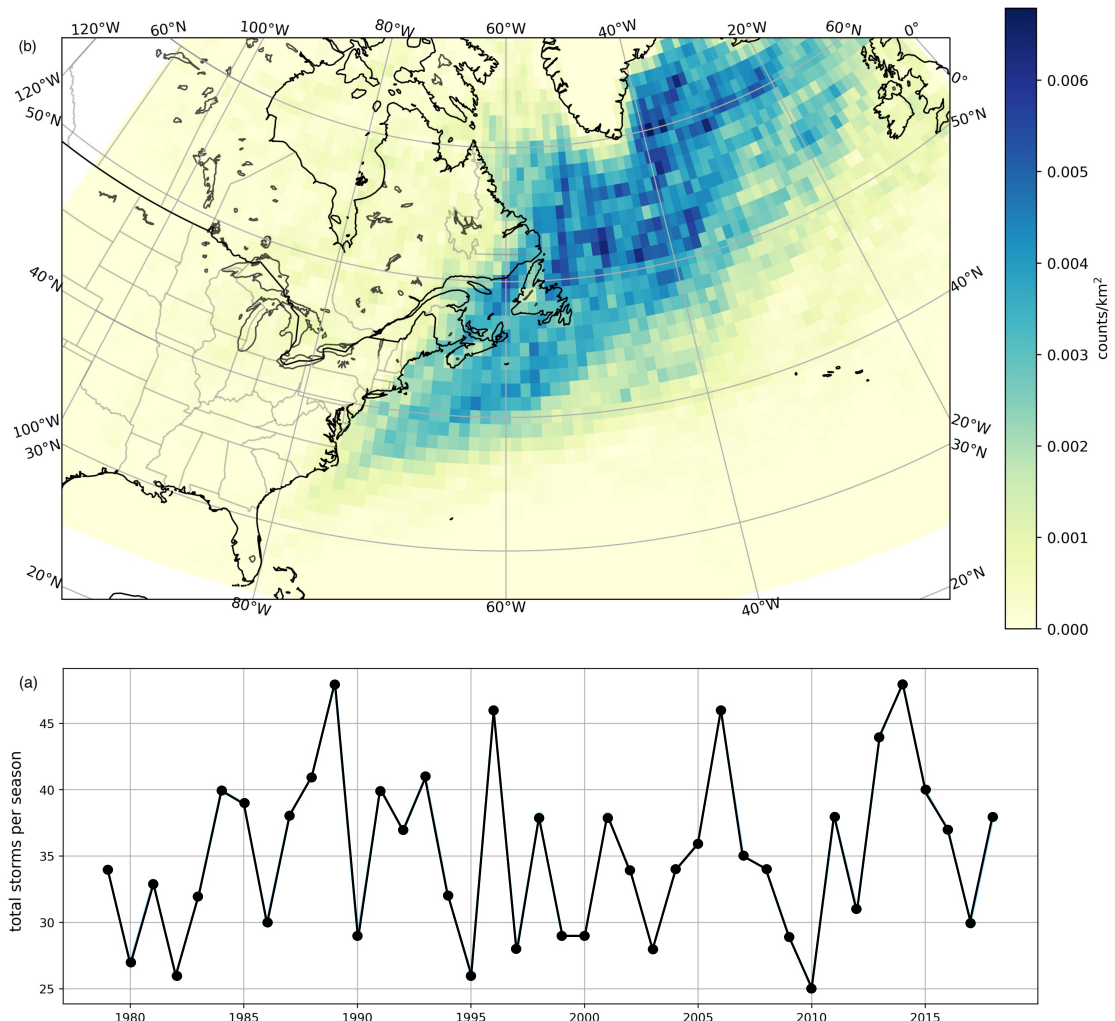


Figure 3.3: Winter season bomb cyclone climatology. (a) Time series of winter bombs per Nov-Mar season. (b) Spatial distribution of total bomb storm density from 1979-2019.

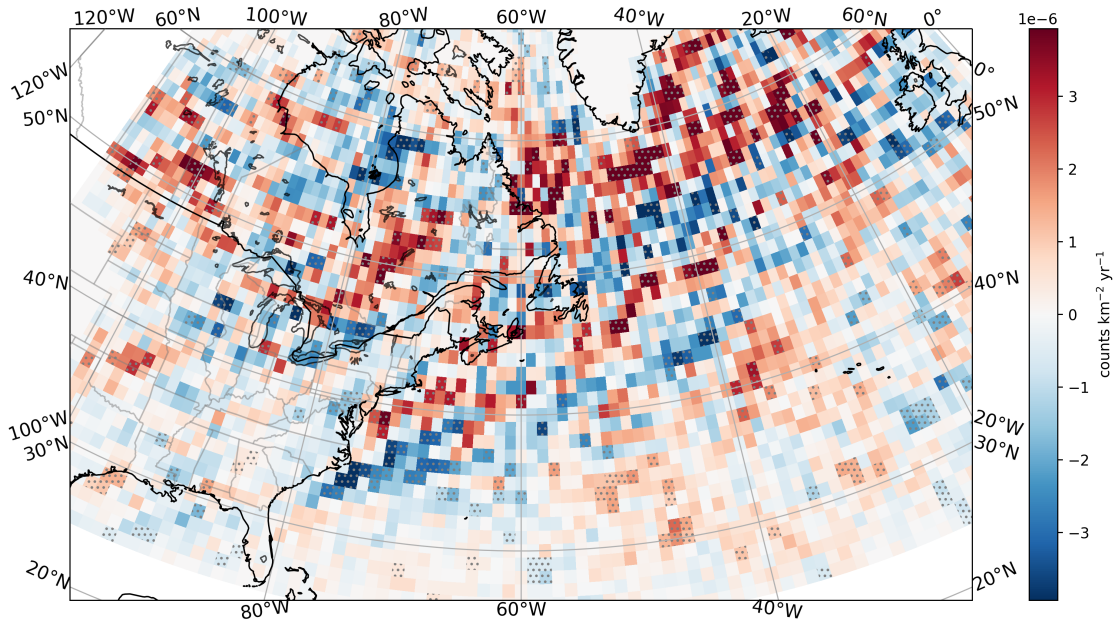


Figure 3.4: Storm track density trends. Slope of linear regression of winter seasonal storm track density time series from 1979-2019 in each grid cell. Shading shows 90% confidence interval.

begins off the coast of Georgia and South Carolina, there is a dipole pattern which shows negative trends to the east of South Carolina and Georgia extending northeastward and positive trends along the northeast US, indicative of a northward shift in the climatological storm track as well as a narrowing of the track along the coast in this region. Note also the statistically significant increases in track density over the subpolar North Atlantic and parts of the continent, namely over Quebec and southeastern parts of the Canadian Prairies.

3.4 Variability

In this section, interannual variability of winter storm activity is investigated with two common techniques. First, it is examined with a statistical technique called Empirical Orthogonal Functions (EOF) analysis to find independent modes of variability in the storm field (refer to section 3.2.4). Then a machine learning technique called Self-Organizing Maps (SOM) is utilized to pick out common winter storm tracks patterns over the years of the study (refer to section 3.2.5). The section is concluded with a study of relevant atmospheric variables from reanalyses that connects the results of the EOF and SOM with observations and further links the identified patterns to potential physical drivers and

climate impacts.

3.4.1 Empirical Orthogonal Functions Analysis

3.4.1.1 All Storms

The EOFs were calculated using the annual fields of gridded storm density of all storms within the study area of 103°W-23°W, 25°N-62°N. The spatial patterns and PCs of the first 4 modes are shown in Figs 3.5 and 3.6, respectively. The variance explained by the leading mode was just 4.28% (Fig. 3.7). The first 15 modes must be considered in order to explain half of the total variance. This means that the pattern associated with the first few modes do not represent a significant amount of the variability in storm track activity. The power of EOFs is that they can reduce the dimensionality of a problem. If there are a few key patterns of variability that explain most of the variability in the full field, an EOF analysis will reveal those patterns. If these patterns together explain a large portion of the variance in the field, these leading patterns can be analysed and discussed to gain insights into the total variability of the field. In this case, a large number of modes is needed to explain most of the variance in the ETC field which counteracts the utility of EOF analysis.

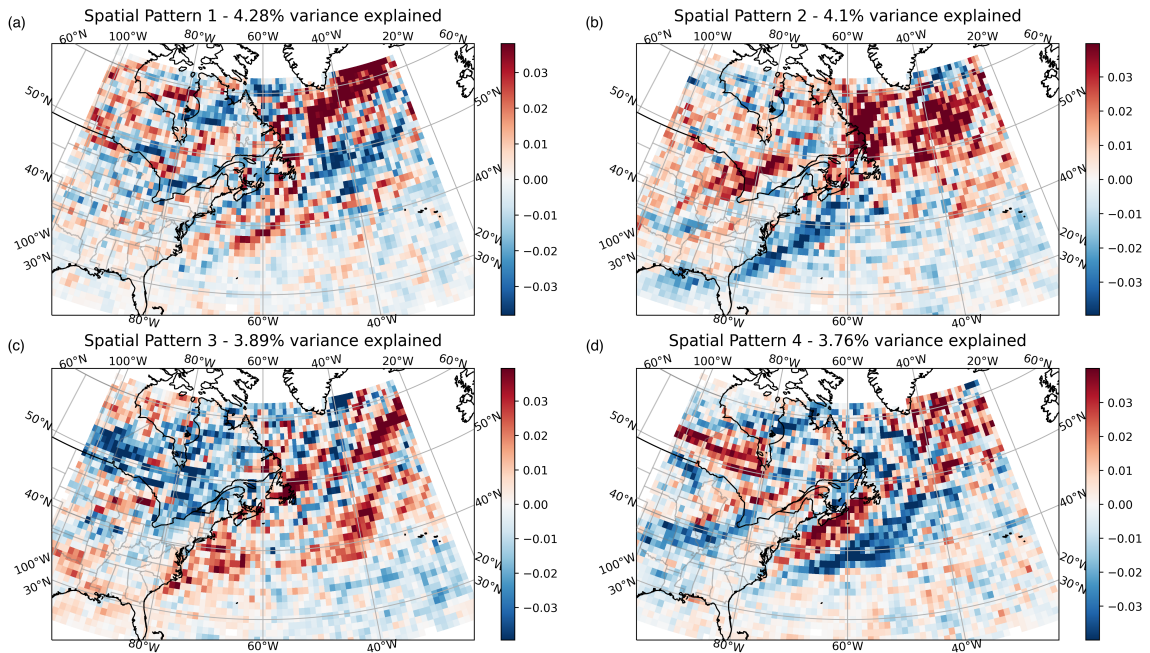


Figure 3.5: Spatial patterns of the first 4 EOF modes of gridded storm densities.

To reduce the noise in the storm track fields and obtain higher explained variance from the leading modes, many adjustments to the input field were tested. These adjustments included significantly reducing the spatial domain to focus only on the main track and increasing grid cell size from the original $1^\circ \times 1^\circ$ cells to $5^\circ \times 5^\circ$ and $10^\circ \times 10^\circ$.

The effects of these adjustments on two important measures of EOF results are given in table [3.1](#). The table shows the variance explained by the leading mode as a percentage of the total variance of the system and the number of modes required to get a cumulative variance explained that surpasses 50% of the total. The full results of these analyses can be found in appendix [A](#).

The spatial domain reduction only slightly increased the explained variance of the leading mode and slightly decreased the modes required to reach 50% variance explained. As one would expect, increasing grid cell size from original to medium and then larger cells improved the variance explained by the leading mode and decreased the number of modes required to explain half the variance of the system.

For the EOF performed on the full field with large (10×10 degree) grid cells, the explained variance of the first 4 modes combined is 52%. This is a significant portion of variance explained by just a few modes, which means analysing and discussing the

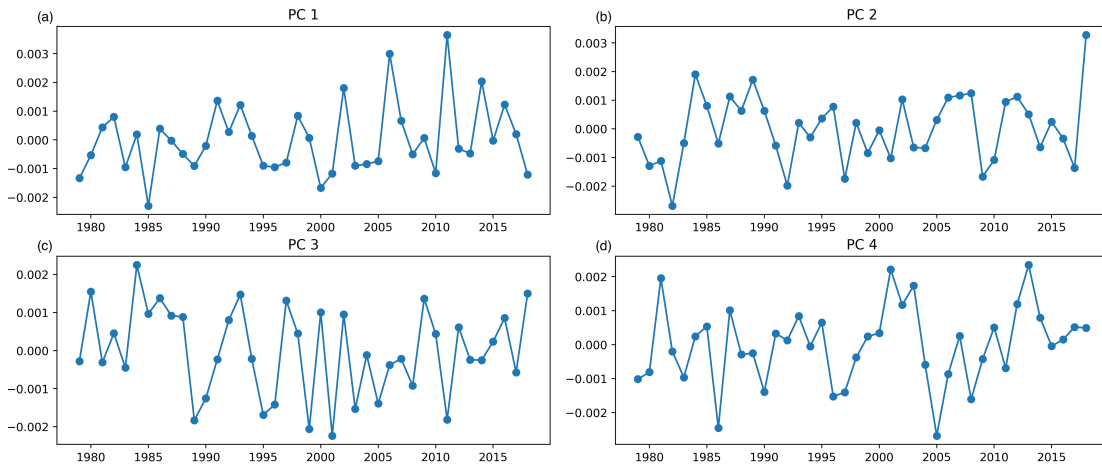


Figure 3.6: Principal component timeseries of the first 4 EOF modes of gridded storm densities.




Figure 3.7: Percent of total variance explained by each EOF mode (blue) and cumulative variance explained (red) of gridded storm densities.

Spatial domain (resolution)	Variance explained by mode 1	Modes required to reach 50% variance explained
original (1x1 deg)	4.28 %	15
main track (1x1 deg)	5.84%	13
medium grid cells (5x5 deg)	11.36%	8
large grid cells (10x10 deg)	18.54%	4

Table 3.1: Comparison of EOF results performed on the gridded storm counts at varying resolutions and over varying spatial extents.

variability in time of these patterns could be useful for understand and possibly predicting the variability of the whole system. However, using grid cells this size causes a loss of much of the spatial detail within the field.

This set of EOF analyses on the field of storm counts over a range of grid cell sizes and spatial extents shows that in general this field is highly variable. Even within the main track, there are many independent modes that contribute to the observed interannual variability. In order to obtain useful EOF results, I must reduce the variability of the field and in this process, important spatial detail is lost. This leads to the conclusion that EOF analysis is not the best method for studying seasonal variability of the western North Atlantic winter storm track. Subsection [3.4.2](#) investigates self-organizing maps as a complementary technique to explore this variability.

3.4.1.2 Bomb Storms

In addition to all storm counts, EOFs were used to analyse the count density fields of bomb storms (Figs. [3.8](#) and [3.9](#)). The first mode, which explains 6.17% of the variability, illustrates a dipole pattern between the early (upstream) and later (downstream) parts of the main storm track. This suggests in a season with many bomb cyclones off the east coast of USA and Atlantic Canada, there are likely to be fewer bomb cyclones in the open ocean south and east of Greenland in that same season. The second and third modes, which explain about the same portion of the total variability (4.9 % and 4.5% respectively) indicate lateral shifting of the track. The track shift occurs primarily between 40 and 60 degrees north. In the third mode, a track passes to the west (east) of Nova Scotia and

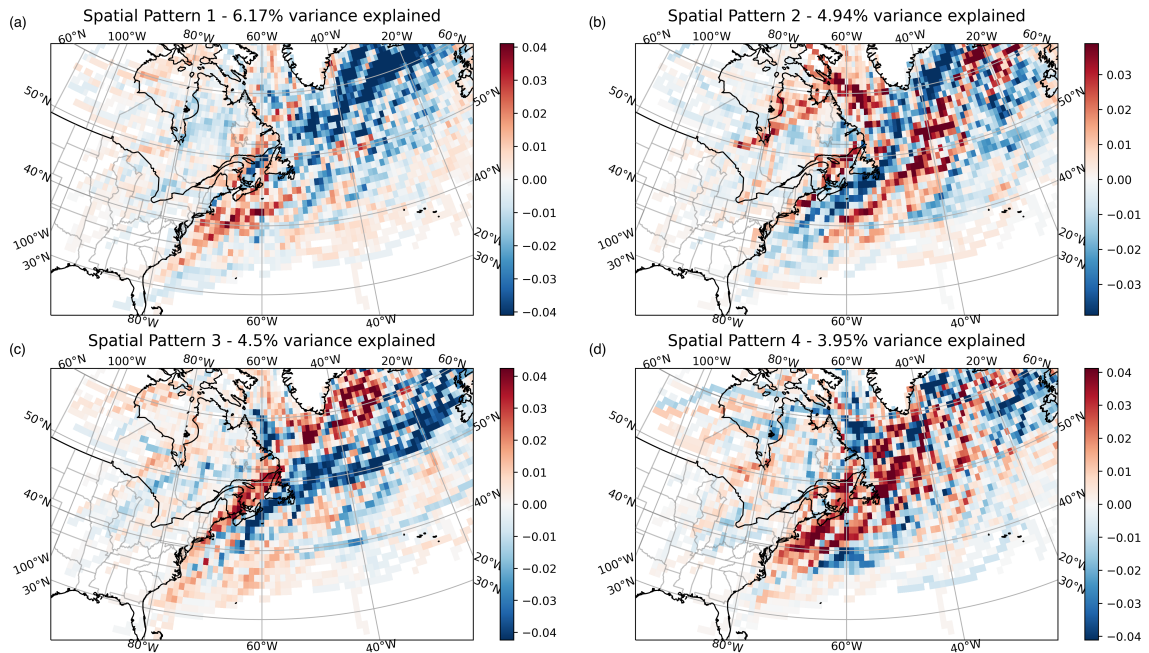


Figure 3.8: Spatial patterns of the first 4 EOF modes of gridded bomb storm densities.

Newfoundland for a positive (negative). One can see a track that hugs the east coast of NS and NFLD for both positive second and third mode. Also, a track that passes further off shore (on shore) occurs in a positive (negative) mode two.

As with the field of all storms, I attempted to reduce variability in the bomb storm field by reducing the spatial extent of the study area. The leading mode of this smaller spatial field EOF analysis explains 5.8% of the variability. See appendix [A](#) for more details. Having identified in the EOF analysis of all storms that increasing the grid cell size causes unwanted loss of spatial resolution, I did not explore EOF analysis of the field of bomb storms on a grid with larger cell sizes.

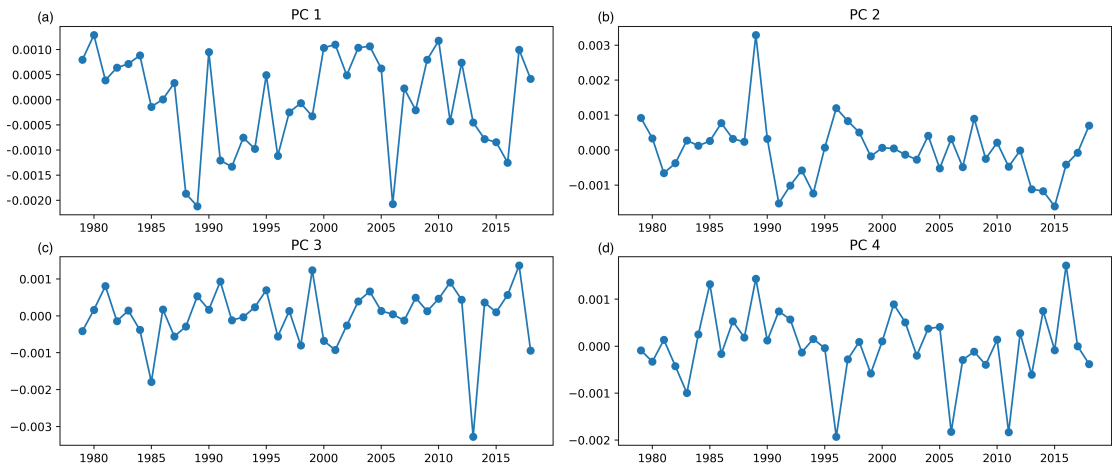


Figure 3.9: Principal component timeseries of the first 4 EOF modes of gridded bomb storm densities.




Figure 3.10: Percent of total variance explained by each EOF mode (blue) and cumulative variance explained (red) of gridded bomb storm densities.

3.4.2 Self-Organizing Maps Analysis

The dimensionality of the highly variable storm track density field is not well reduced using EOF analysis as shown in the previous section. Self-organizing maps (SOMs) is an alternative method of dimensionality reduction that is complementary to the EOF analysis. In EOF analysis, the temporal features can be reconstructed from the decomposed components. Conversely in SOM analysis, the spatial patterns from each time step are preserved but the temporal connections between spatial characteristics do not constrain the analysis. The nonlinear SOM algorithm proves advantageous over the linear EOF analysis for identifying more complex patterns in noisier fields ([Liu and Robert, 2011](#)). SOMs are used to give additional understanding of the variability of the system to further the information obtained with EOFs.

As explained in section [3.2](#), a self-organizing maps algorithm is a machine learning technique that sorts patterns. In this section, SOMs are used to sort the 40 spatial patterns of seasonal storm track density that occur over the study period into groups of seasons with similar activity.

3.4.2.1 Parameter Selection

When using a self-organizing maps algorithm to analyse data, one must choose parameters that set the behaviour of the algorithm. Varying these parameters can give different SOM outputs. These parameters should be chosen carefully to ensure the SOM output is meaningful and reproducible.

This subsection explores the effects of varying three parameters: the initial learning rate (η_{ini}), the initial standard deviation of the Gaussian that defines the neighbourhood function (σ_{ini}), and the number of iterations on which to train the SOM (t_{max}).

At each iteration (t), the set of weight vectors (\mathbf{W}) is adjusted according to the following update equation ([Vettigli, 2018](#)):

$$\mathbf{W}_{new} = \eta(t) \cdot f(\sigma(t)) \cdot (\mathbf{x}_i - \mathbf{W}) \quad (3.4)$$

where \mathbf{x}_i is the input vector on which the algorithm is being trained at that iteration, η is the learning rate, and f is the neighbourhood function which has a Gaussian shape with standard deviation given by the iteration-dependent parameter σ that is centred on the location of the current winning node (\mathbf{w}_{c^i}) of \mathbf{x}_i . The neighbourhood function specifies the

extent to which the nodes surrounding the winning node are adjusted by the current input vector. Effectively this controls the relationships between the final output nodes. Because a Gaussian neighbourhood function was chosen in the 2x2 set up, a given node is more similar to the nodes directly adjacent to it than the node that is diagonally adjacent.

As shown in the update equation (Eq. 3.4) the learning rate is iteration-dependent. At each iteration, the learning rate changes with iterations as follows

$$\eta(t) = \frac{\eta_{ini}}{1 + \frac{2t}{t_{max}}}, \quad (3.5)$$

where t_{max} is the total number of iterations completed during training (Vettigli, 2018).

Similarly, the neighbourhood function decays with iterations according to the decay of its standard deviation which is defined as follows (Vettigli (2018))

$$\sigma(t) = \frac{\sigma_{ini}}{1 + \frac{2t}{t_{max}}}. \quad (3.6)$$

Note the influence of t_{max} on the decay of both η and σ . Stopping a SOM run before it reaches t_{max} does not allow these two parameters to complete their decay process, at which point their value is equal to one third of the initial. For this reason, we only evaluate and compare outputs from runs that have progressed all the way through to $t = t_{max}$. An iteration of the algorithm is defined as each time a vector is given as an input and the update equation is used. Therefore, every N^{th} iteration marks a full cycle through the input vectors. The values of t_{max} are restricted to multiples of N to ensure each input vector is equally accounted for in the final output.

The parameter selection process begins by choosing a testing range for the initial learning rate and initial standard deviation of the Gaussian that defines the neighbourhood function. The quantization error (QE) is a measure of how well the inputs are represented by their individual winning node in the SOM. A small QE means the input vectors fit well within the groups they have been separated into. QE is defined as follows:

$$QE = \frac{\sum_{i=1}^N \|\mathbf{x}_i - \mathbf{w}_{c^i}\|}{N} \quad (3.7)$$

where N is the number of input vectors, \mathbf{x}_i is an individual input vector and \mathbf{w}_{c^i} is the weight vector of the winning node (Wandeto and Dresch-Langley, 2019). QE can be used

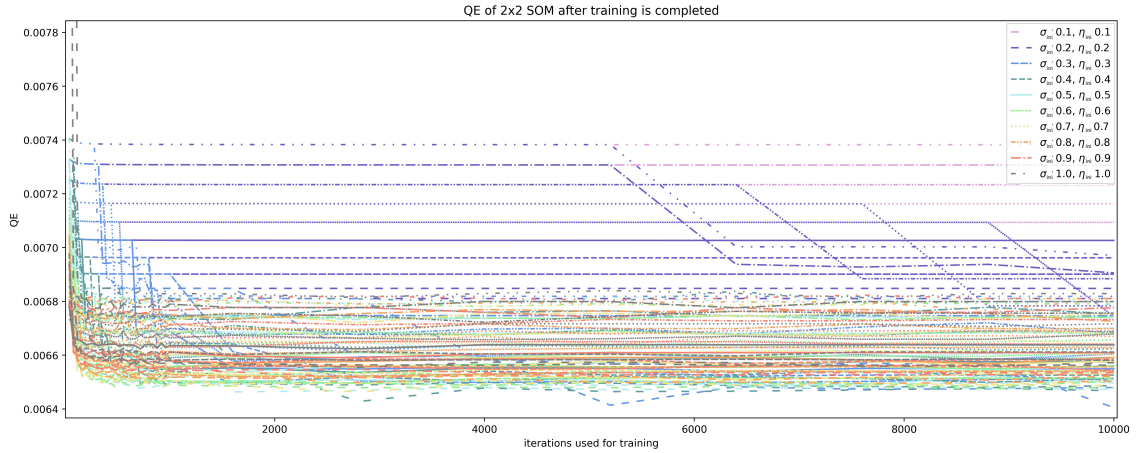


Figure 3.11: Variation of quantization error with SOM parameters. QE of SOM outputs trained over a range of maximum iterations (t_{max}), while varying the standard deviation of the Gaussian neighbourhood function (σ_{ini} , shown by line style) and the learning parameter (η_{ini} , shown by colour).

to narrow down the range of σ_{ini} and η_{ini} that give useful SOM results. The initial range chosen for both parameters is from 0.1 to 1 at intervals of 0.1. The SOM algorithm is run many times varying η_{ini} and σ_{ini} within this range and letting t_{max} vary between 40 to 10 000 specifically at multiples of 40. For each output, the QE is calculated to evaluate the SOM run (Fig. 3.11).

For the full range of parameter values, the QE settles to a relatively constant value for all subsequent iterations once the model is trained with a sufficiently high t_{max} . The sufficient number of max iterations depends on the size of η_{ini} . For smaller initial learning rate, η_{ini} , the SOM requires a higher number of max iterations to reach a steady QE. This is an effect of the decay function (3.6, 3.5). In general, QE is lowest for $\eta_{ini} = 0.3$ (Fig. 3.11, blue). The standard deviation of the Gaussian function in the neighbourhood parameter (σ_{ini}) seems to have less of an effect on QE, but values around 0.4 give lower equilibrium QE than values closer to 1. This preliminary analysis narrows the range of testing on σ_{ini} to [0.4, 0.5, 0.6] and narrow the range of testing on η_{ini} to [0.3, 0.4, 0.5].

Guided by the results of the QE calculations, further investigation narrows down the parameter selection using pattern convergence. If the outputs of many different SOMs converge to a single solution, then the result is reproducible and meaningful. Periods of temporary convergence or near-convergence can be identified for a given (σ_{ini}, η_{ini}) pair by measuring root mean squared error. Root mean squared error (RMSE) can be used

	$\sigma_{ini} = 0.4$	$\sigma_{ini} = 0.5$	$\sigma_{ini} = 0.6$
$\eta_{ini} = 0.3$	41	13	35
$\eta_{ini} = 0.4$	43	4	8

Table 3.2: Number of times RMSE is equal to zero for each $(\sigma_{ini}, \eta_{ini})$ pair over 63 trials

compare the difference between two outputs of the SOM with different initial parameters. When RMSE is zero, it signifies no pattern change. A consistently small RMSE indicates a robust SOM result, because it indicates at least partial elimination of the dependence of the result on the choice of initialization parameters.

RMSE is used to evaluate increasing t_{max} . For a constant $(\sigma_{ini}, \eta_{ini})$ pair, RMSE is calculated between the SOM output from $t_{max} = kN$ and $t_{max} = (k + 1)N$ where k is an integer between 1 and 250 that specifies the number of times the algorithm has looped over the full set of input vectors; and N is 40, the number of input vectors. Testing all SOMs within this range would require the algorithm to be run with 1500 unique parameter combinations (6 $(\sigma_{ini}, \eta_{ini})$ pairs multiplied by 250 possible k values). Due to the computational restrictions, I sampled the range of k values 63 times for each of the 6 $(\sigma_{ini}, \eta_{ini})$ pairs, thereby running 378 SOMs instead of 1500 and calculated the RMSE for each k value.

The goal is to identify a $(\sigma_{ini}, \eta_{ini})$ pair for which the output pattern converges at t_{max} greater than a sufficiently high value. This would reveal the ideal $(\sigma_{ini}, \eta_{ini})$ values and minimum t_{max} . Unfortunately, none of the pairs exhibited perfect convergence in this way. There were, however, regions of temporary convergence i.e. where $RMSE = 0$ for some period of time before the pattern changed again. The total number of k values for which $RMSE = 0$ for each $(\sigma_{ini}, \eta_{ini})$ pair is given in table [3.2](#)

The RMSE is most frequently equal to zero for $\sigma = 0.4$, and $\eta = 0.3$ or 0.4 . This allows the investigation to be reduced to two pairs: A: $\sigma = 0.4; \eta = 0.3$ and B: $\sigma = 0.4; \eta = 0.4$. The RMSE is found for all k values (1 to 250) for pair A and B. Over this range, parameter combination A had 142 instances of RMSE equal to zero, while parameter combination B had 93 instances. This leads me to settle on the parameter pair A: $\sigma = 0.4; \eta = 0.3$ to build the final SOM for analysis.

Having chosen learning rate and neighbourhood parameter, it only remains to choose the number of iterations for training (t_{max}). The SOM temporarily converges on a few different patterns. One of these must be chosen for the analysis. This decision will be




Figure 3.12: Frequency of unique SOM pattern outputs. Each unique output in the set of all SOM results is denoted with a pattern ID number. The pattern frequency is shown as a function of pattern ID with the quantization error of that unique pattern indicated by tone of its marker.

informed by QE and the frequency with which each unique pattern is identified i.e. how many k values give that specific pattern.

Examining how QE changes with number of training iterations for pair A, (Fig. 3.11, dashed teal line) we see that QE settles to a relatively constant mean (6.51×10^{-3}) for $t_{max} \geq 720$ which corresponds to $k = 18$. All t_{max} values below this threshold are neglected. To determine the best t_{max} in this range, I evaluate the patterns that result from running 233 SOMs, using pair A and letting t_{max} range from $18N$ to $250N$. Each unique SOM result was given a unique pattern ID number. For each unique pattern, I calculate how many times a SOM produced that output along with the QE of that pattern (Fig. 3.12).

The most common pattern resulting from the SOM initialized with $\sigma_{ini} = 0.4$ and $\eta_{ini} = 0.3$ is pattern 62 which reoccurs 17 times. Running the SOM algorithm with $t_{max} = 6000$, $\sigma_{ini} = 0.4$ and $\eta_{ini} = 0.3$ produces pattern 62. The QE of this pattern is small at 0.006518. Table 3.3 gives the SOM groupings of the seasons in the study period. Recall that the winter seasons are named according to the year in which they start.

Node 1	Node 2
1983, 1985, 1993, 2000, 2002, 2007, 2009, 2012, 2013, 2018	1980 1982, 1990, 1991, 1995, 1999, 2001, 2003, 2010, 2017
Node 3	Node 4
1979 1984 1987, 1988, 1992, 1997, 1998, 2005, 2008, 2014, 2016	1981, 1986, 1989, 1994, 1996, 2004, 2006, 2011, 2015

Table 3.3: Seasonal groupings of SOM pattern 62.

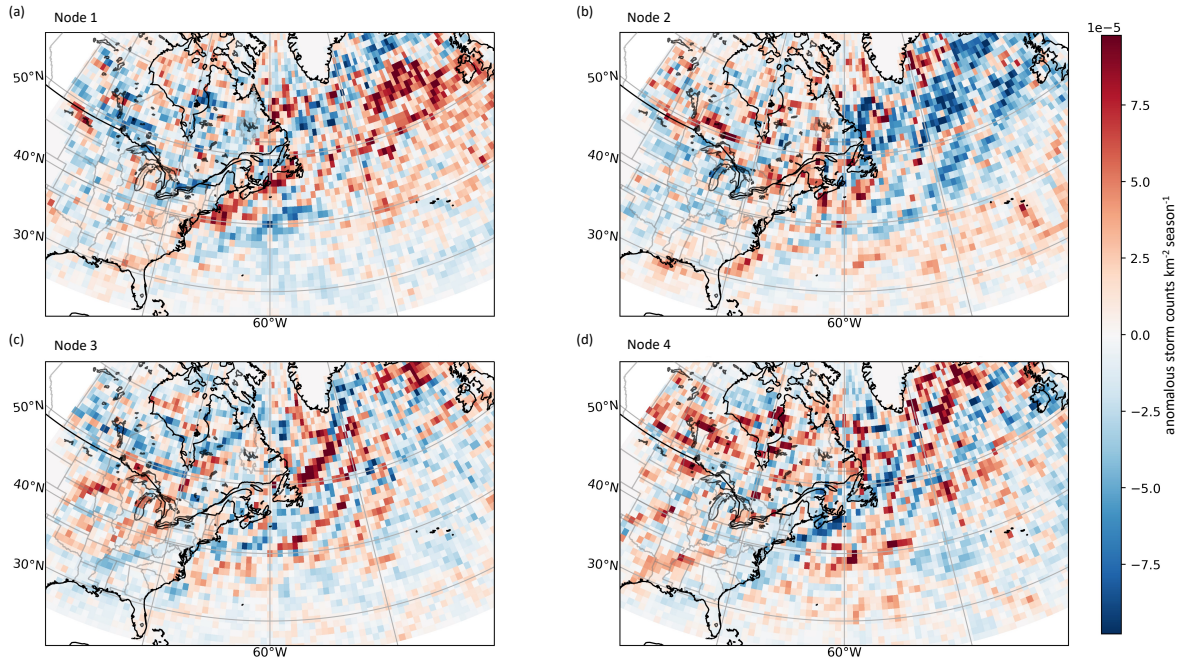


Figure 3.13: Anomalous storm counts for each of the four sets of SOM-grouped seasons. Unique SOM output pattern 62 with parameters ($\sigma = 0.4, \eta = 0.3, t_{max} = 6000$).

3.4.2.2 Results

The SOM was run with the specified parameters (see Section 3.4.2.1) and the 40 storm seasons in our dataset were sorted into 4 groups in a 2x2 arrangement. To visualize the SOM groupings, the mean seasonal storm counts km^{-2} is calculated in each grid cell for the 4 nodes across the seasons grouped into the specific node. This allows us to see the typical pattern of a season that belongs in each grouping. Then, to further highlight the differences among the four groupings, the all time mean field of seasonal storm counts km^{-2} is removed to obtain a grid of anomalous storm counts km^{-2} that spans the study area in each node (Fig. 3.13). The result is further interpreted in Fig. 3.14.

The strongest signal in node 1 (Fig. 3.13a and Fig. 3.14a) is found along the main climatological track identified in Section 3.3. The pattern shows a dipole pattern with an increase in track density extending along the eastern seaboard from Cape Hatteras to southeastern Newfoundland contrasted with a decrease in storms further offshore following a similar orientation as the positive signal. This indicates a shift in the main climatological track towards the shoreline along NE USA and Atlantic Canada during node 1 years. There is also a strong positive signal later in the climatological track located in the North Atlantic to the west of Ireland. The prairie storm track appears to be less active than average in this

node and the Labrador Sea is slightly more active.

The second node (Fig. 3.13b and Fig. 3.14b) shows weak anomalies indicating seasons in this node are generally quite similar to the climatological mean storm track density pattern particularly in the early portion of the track. One can, however, identify a decrease in storms in the northeastern Atlantic as well as in northern Newfoundland and along the coast of Labrador. The St. Lawrence river valley region sees above average activity in this node contrasted with fewer storms around the Great Lakes including around Lake Michigan, Lake Huron and Lake Erie.

The strongest signal in node 3 (Fig. 3.13c and Fig. 3.14c) is found northeast of Newfoundland extending from just offshore toward the southern tip of Greenland. The features offshore from New England and Nova Scotia may indicate a shift of the climatological track away from the coast and into the North Atlantic. In addition, this is the node with the strongest signal of Colorado Lows tracking over the Great Lakes.

The final node (node 4; Fig. 3.13d and Fig. 3.14d) shows below average activity along the main climatological along the northeastern coast of USA, Nova Scotia, and Newfoundland. This is contrasted with above average activity across Saskatchewan, Manitoba, and northwestern Ontario. Another notable feature is a positive anomaly extending from northeastern Texas to Kentucky.

The SOM has organized the winter storm seasons into groups with at least some spatial commonalities as shown by the negative and positive track anomalies identified in the SOM output analysis (Fig. 3.14). The following section (3.4.3) will examine the SOM results further by drawing connections between the signals found within the four nodes and atmospheric variables that are related to ETCs.

3.4.3 Links to atmospheric drivers and effects

If the SOM had separated our inputs into random groups, we would see four nodes with patterns that lack spatial structure. Instead, we see noticeably different patterns of storm track variability across the four nodes. In this section, the SOM outputs are expanded upon by evaluating their physical basis. I look for links between the patterns in storm track activity within the SOM groupings and patterns in two classes of variables: drivers and effects. Drivers are atmospheric variables that are involved in the formation and development of storms while the effects are variables that are typically caused by storms, such as wind and precipitation.

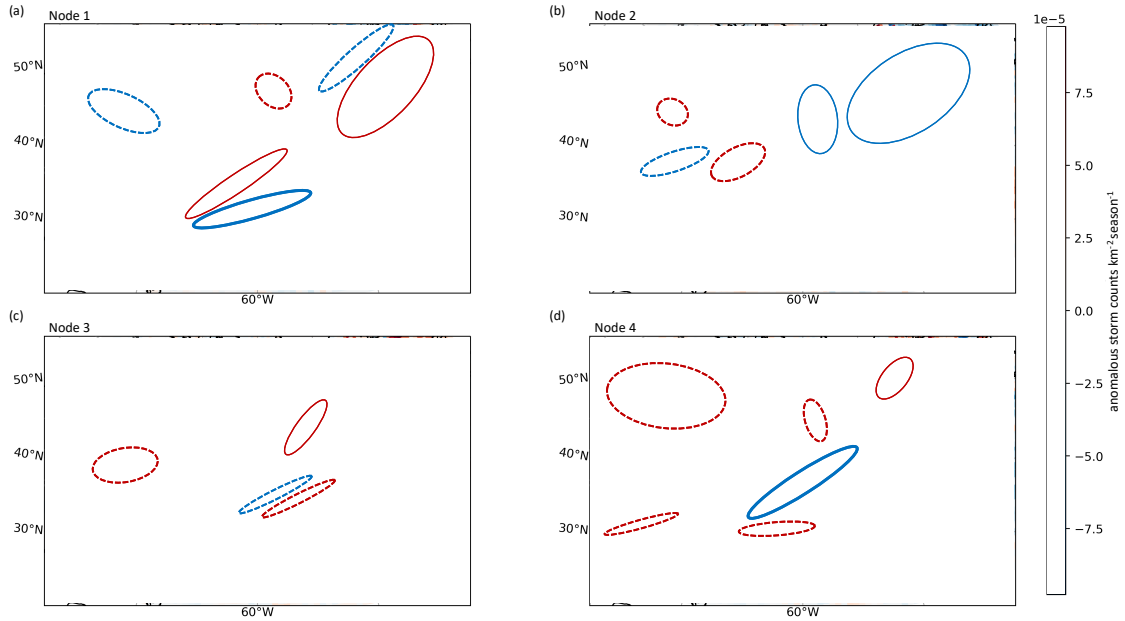


Figure 3.14: Important features within anomalous seasonal storm counts for the SOM-identified groups of seasons. Regions of stronger and weaker positive anomalies indicated by solid and dashed red outlines, respectively. Regions of stronger and weaker negative anomalies indicated by solid and dashed blue outlines, respectively.

3.4.3.1 Drivers

The main driver of ETCs is baroclinic instability. Baroclinic zones occur where there are strong horizontal temperature gradients. The 2m air temp and 500mb thickness are used to identify these zones and how they vary with SOM node (Figs. 3.15 and 3.16, respectively). In the winter months, the air over the ocean typically stays warmer than air over the land. This is due to the greater heat capacity of water. The ocean water itself is warmer than the land, and this keeps the air mass above at a warmer temperature, as well. By contrast, the air over land cools down quickly in the fall and winter. These two air masses meet along the coastline to form a strong temperature gradient. For this reason, most storms are found along this coastline climatologically (Fig. 3.1b). Based on this theory, one can expect an increased coastal temperature gradient to lead to an increase in storm activity.

Across the SOM nodes, a strong relationship with 2m air temperature is observed over land. In node 1, an increase in storms along the main track, east of North America, can be connected with below normal temperatures over land which strengthens the coastal baroclinic zone. In node 2, an increase in 2m air temperatures over Newfoundland and Labrador weakens the coastal temperature gradient in that region. In the 500 mb thickness

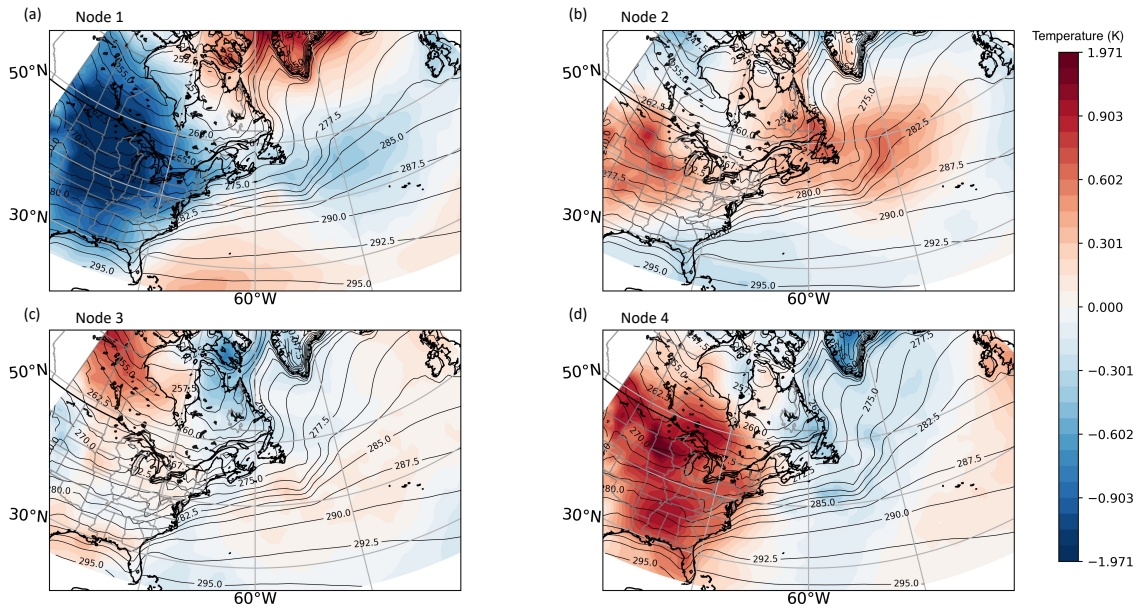


Figure 3.15: Mean seasonal 2 meter air temperature (contours) and anomaly from the climatological mean (colours) for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

field, there is an increase in heights in the later part of the track east of Newfoundland and south of Greenland. In the storm field there is a slight decrease in storms in the region just offshore of Labrador and an overall decrease in the northern part of the North Atlantic. In node 4, the pattern appears to be roughly opposite to node 1. The 2m air temperatures and 500 mb thickness both show warmer continental temperatures which weakens the coastal baroclinicity. As expected, there is a corresponding decrease in the main storm track along the east coast.

Within these baroclinic zones, storms are most likely to develop and intensify where there is strong upper level divergence which is often closely tied to vorticity. These areas that are conducive to ETC formation are identified using 250 mb winds (Fig. 3.17 and Fig. 3.18), 500 mb geopotential heights (Fig. 3.19), and 500 mb relative vorticity (Fig. B.1). A common level for diagnosing maximum upper level winds is 250 mb. The winds at this level can be used to identify jet streaks which are areas of especially fast winds within the jet stream. As the air accelerates into (decelerates out of) the region of higher wind speeds the change in speed creates a counter-clockwise (clockwise) ageostrophic circulation in the vertical plane. As a result, there is upper level divergence in the right

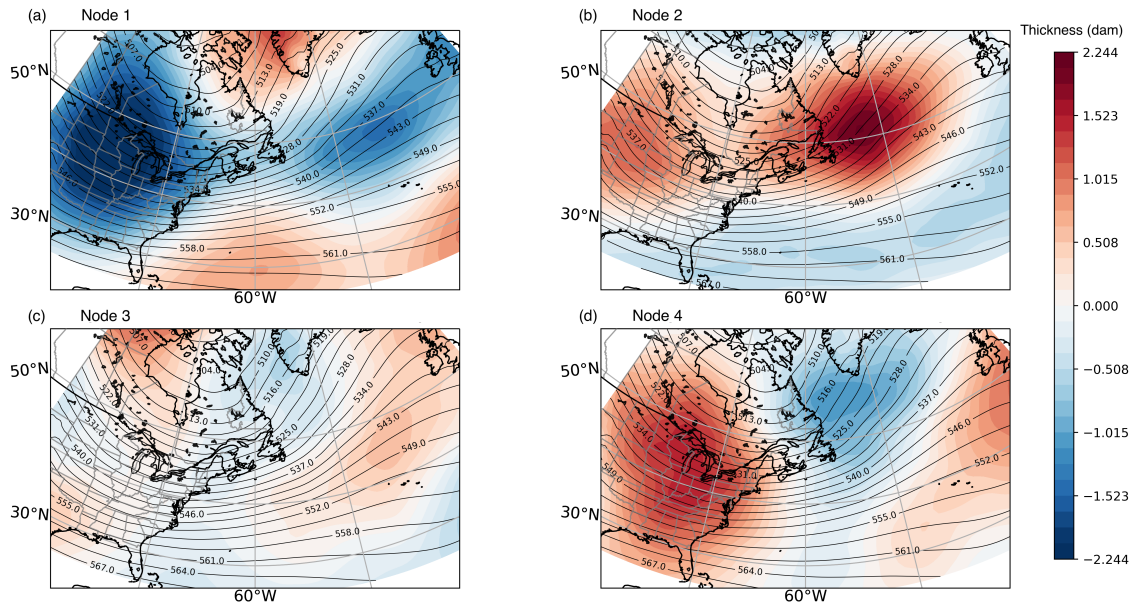


Figure 3.16: Mean seasonal 500 mb thickness (contours) and anomaly from the climatological mean (colours) for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

entrance (left exit) region of the jet streak. Within the 250 mb wind fields, the strongest winds are found in node 1 (Fig. 3.17a). The multi-season mean jet streak located near the Mid-Atlantic Bight has its right entrance region over coastal North Carolina, a common genesis region for ETCs. Its left exit region is located south of Cape Cod. In node 1 of the storm density field, the start of the region of increased storm activity is collocated with the left exit region of the mean jet streak. The other nodes show similar, although weaker, patterns in the 250 mb winds.

Comparing node 1 and 2 of the anomalous 250 mb wind fields in the northern North Atlantic shows opposite patterns (Fig. 3.18a,b). In node 1, the anomalous wind has a counterclockwise circulation and in node 2 there is a clockwise circulation. The anomalous storm field respectively show a corresponding increase and decrease in downtrack storms in node 1 and 2 (Fig. 3.14). Considering the anomalous circulation patterns in the context of the mean 250 mb flow, this tells us there is a strong zonal component to the 250 mb streamlines in seasons with higher storm track density downstream in the eastern North Atlantic. However along the east coast of North America, the anomalous wind fields show that there is an increased meridional component to the 250 mb flow. This is consistent with the physics of ETCs since the presence of a meridional jet increases the likelihood of storm

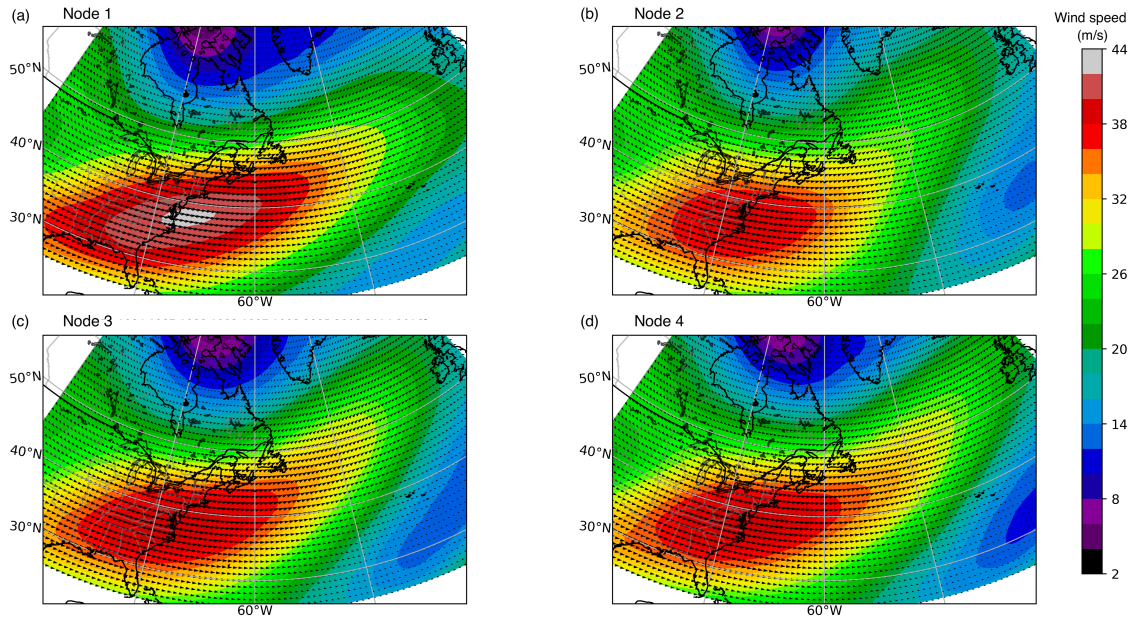


Figure 3.17: Mean seasonal wind direction (arrows) and intensity (colours) at 250 mb for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

development. Node 3 indicates a positive upper level wind anomaly collocated with the positive storm track density anomaly northeast of Newfoundland. Finally, node 4 shows increased wind speeds over eastern Canada with a stronger southeastward component which indicates more cold dry continental air blowing over the Atlantic provinces. There is also a decrease in wind speed along the coast of USA. In the storm density field, node 4 shows a decrease of storms along the eastern seaboard, Nova Scotia, and Newfoundland which could reasonably be the effect of slower 250 mb winds in the typical genesis region leading to weaker divergence aloft and stronger advection of cold dry air over the ocean which weakens the temperature gradient by cooling the air mass over the ocean and lowers the amount of moisture available to condense and release latent heat, an element in storm intensification and persistence.

Next, consider curvature and shear vorticity in the 500 mb geopotential height field. Shear vorticity can be identified in regions of diffluence of geopotential height contours. Curvature vorticity is evident in the curvature of geopotential height contours, such as in the base of a trough. In northern hemispheric winter, the main feature governing the 500 mb field in the mid-latitudes is Rossby wave propagation. These waves propagate across the continent and ocean typically from west to east. The mean state (contours in

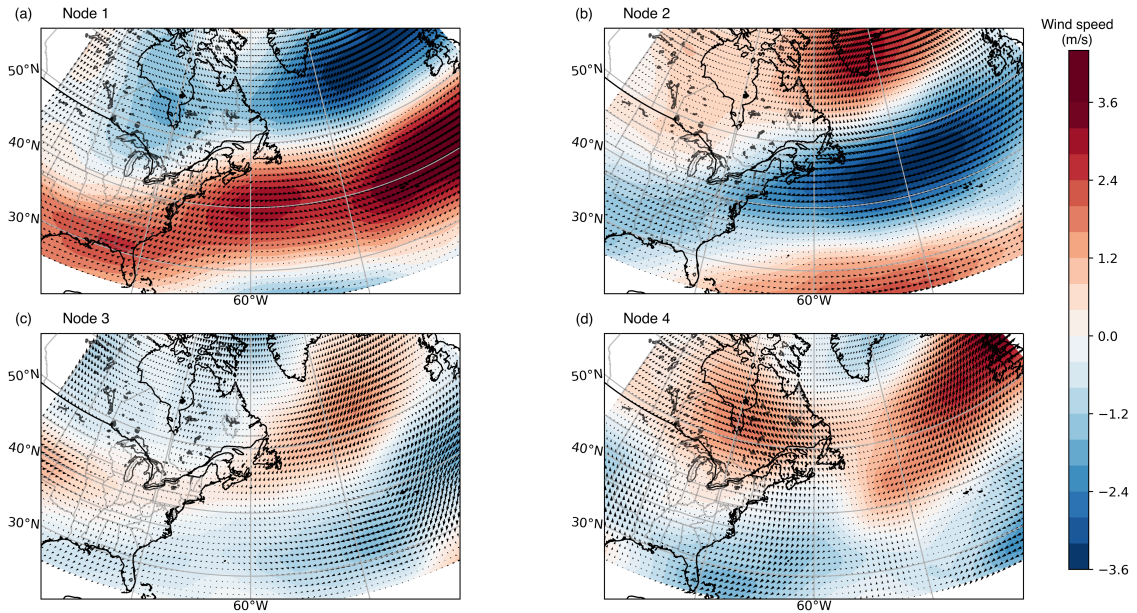


Figure 3.18: Anomaly from the climatological mean of seasonal wind direction (arrows) and intensity (colours) at 250 mb for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

Fig. 3.19) shows heights decreasing with latitude with some slight troughing with an axis over Eastern North America. Day-to-day in the winter, we see a stronger trough than what appears in the mean state, but the axis of the trough varies in location. The strength of the troughing is washed out in the mean state because of the variation in location. Look to the anomalies (colours in Fig. 3.19) to identify more features of the field as it relates to storm activity. Typically, the synoptic situation in which an ETC develops consists of a trough with a north-south axis located over mid to eastern North America. At the base of the trough, there is a maximum of curvature vorticity which leads to positive vorticity advection (PVA) on the right side of the trough. A deeper trough has stronger curvature vorticity. In node 1, the negative anomaly over the continent indicates generally deeper continental troughs in those seasons (Fig. 3.19a). This corresponds with higher east coast storm track density. In node 3, a slight negative anomaly just east of the Rockies could indicate stronger lee-troughing in those seasons (Fig. 3.19c). The track density increases to the right of this negative anomaly. In node 4, higher geopotential heights over eastern North America indicate generally weaker troughs in the region that is important for ETC development. The track density field shows a corresponding decrease in east coast ETCs.

In the eastern portion of our study area, shear vorticity can be deduced from diffluence.

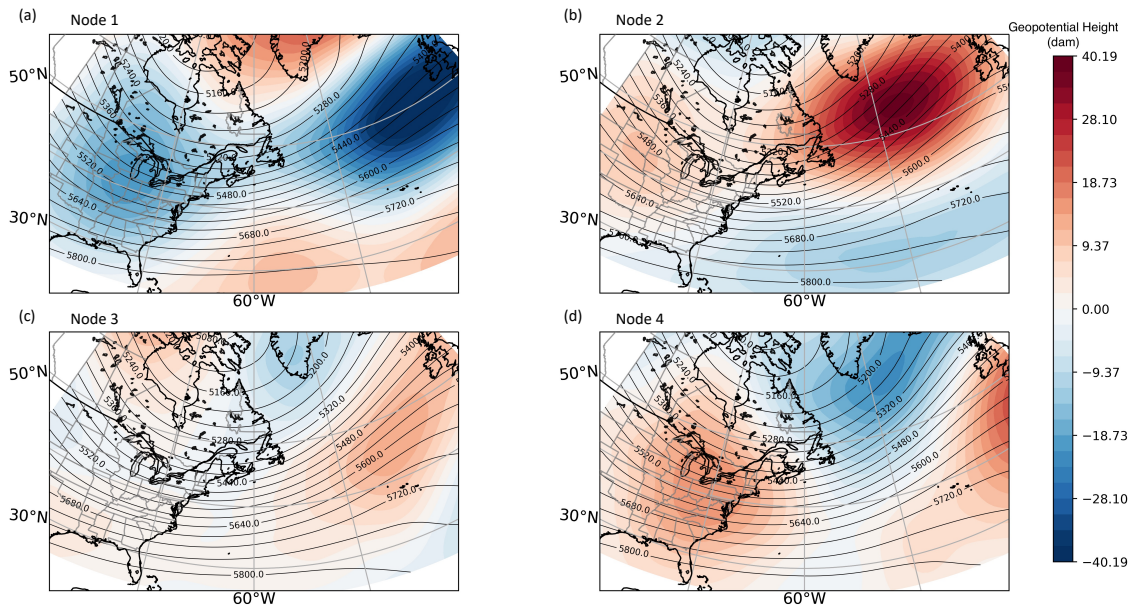


Figure 3.19: Mean seasonal 500 mb geopotential height (contours) and anomaly from the climatological mean (colours) for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

In node 1, the diffluent pattern north of 50° has positive shear vorticity). An increase in storm density also occurs in this region in node 1 (Fig. 3.19a). In node 2, there is generally straight flow north of 50° , while south of this, the diffluence gives rise to negative vorticity or anticyclonic rotation (Fig. 3.19b). The storm track density decreases in the same region for this node. The diffluence in node 3 south of 50° indicates negative vorticity. Storm track density is slightly below or near average in this region for node 3 (Fig. 3.19c). In the far northeast a diffluent pattern that creates positive vorticity corresponds to an increase in storm track density for node 4 (Fig. 3.19d).

The final driver taken into consideration is the total precipitable water vapour (TPWV) field. This field is important for deducing moisture sourcing of storms including locating atmospheric rivers. Having a moisture source is important for the development and propagation of ETCs. A storm with more moisture will have more latent heat release throughout the system which can act to offset the adiabatic cooling in areas of ascent and provide more energy for the persistence of the storm. Greater amounts of TPWV also lead to higher precipitation amounts causing the storm to have greater impacts on human life and the environment. I expect TPWV to be a particularly important field because it acts

as a tracer for many atmospheric processes. TPWV is related to temperature because the amount of water vapour that can be held within a parcel increases as the temperature of the parcel increases. Thus, temperature will be reflected in the TPWV field. Regions of subsidence and upward motion are evident in the TPWV field as well as these processes lead to changes in TPWV. TPWV is also commonly used as a tracer for large scale patterns of movement in the atmosphere to aid in the identification of likely storm tracks.

The TPWV pattern in node 1 indicates a strong moisture source region over the western North Atlantic Ocean associated with more storms along the east coast (Fig. 3.20a). Node 2 shows decreased moisture in the Gulf of Mexico and off the southeast coast of USA which may be connected with fewer storms downstream to the southeast of Greenland (Fig. 3.20b). An overall weak pattern in node 3 can be related to generally weak storm track activity (Fig. 3.20c). The moisture anomaly in the Gulf of Mexico in node 4 can be linked to the storm track density increase in the southern States and continental Canada (Fig. 3.20d). The decrease in moisture over the western Atlantic Ocean and associated decrease in a moisture gradient along the coast may relate to the decrease in storms along the coast of New England, Maine, Nova Scotia and Newfoundland, although this may be more forced by the decreased coastal temperature gradient in this node that is being reflected in the TPWV field.

3.4.3.2 Effects

The SOM method has identified four characteristic types of storm seasons for our study area. In section 3.4.3.1 I found physical connections for the characteristic types. Next, I explore how a storm season of each type affects the weather experienced throughout the winter months.

The precipitation pattern associated with each of the four SOM nodes shows a variation across SOM nodes (Fig. 3.21). The climatological maximum precipitation in the mid-latitudes for our study area extends northeast from the coast of the Carolinas out into the ocean around 40N and 60W. In nodes 1 and 4, the anomalous precipitation pattern shows an off-mean coastal increase/decrease. Conversely, nodes 2 and 3 exhibit anomalously low and high precipitation along the mean, respectively.

As expected, high storm track activity leads to above average precipitation. This is shown in node 1 along the coastline of the eastern States and in the eastern North Atlantic. Positive precipitation anomalies also correspond geographically with the node 2 increase

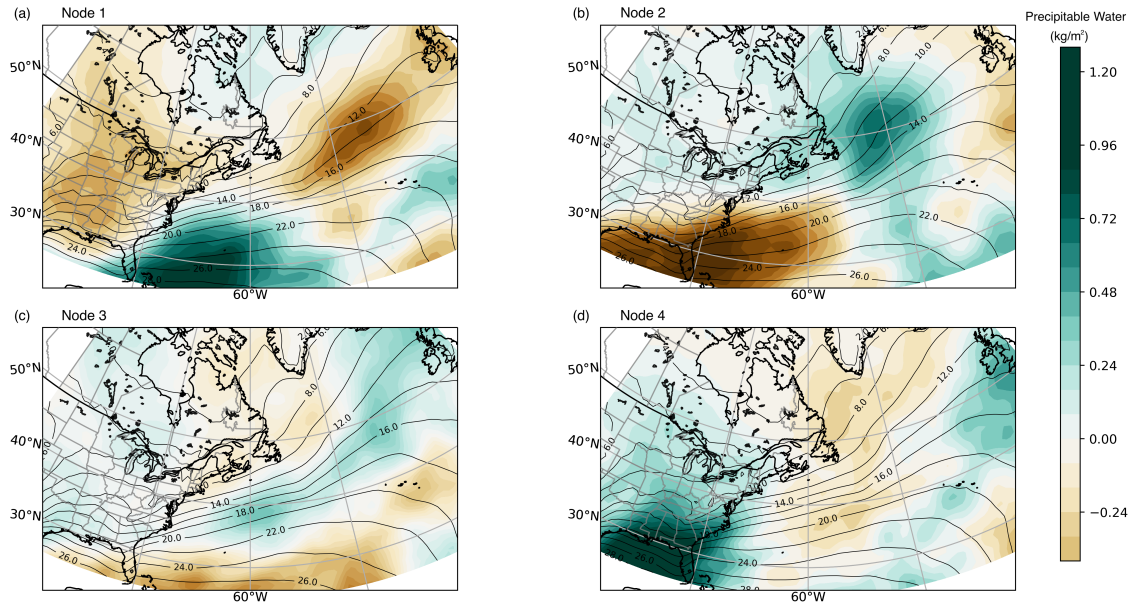


Figure 3.20: Mean seasonal total precipitable water vapour (contours) and anomaly from the climatological mean (colours) for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

in storms along the St. Lawrence River Valley, the node 3 increase in Midwest/Great Lakes storms, and the node 4 increases in southern US and Canadian Prairie storms. Unsurprisingly, the precipitation anomalies and storm track density anomalies are closely related.

Next, consider a subset of the total precipitation field: accumulated snowfall. A season like the node 1 variety, which includes an active east coast storm track, brings a snowy winter season to the northeastern States and Atlantic Canada (Fig. 3.22a). In node 2, the New England region sees below average snowfall as storm densities along the coast are average to below average while densities along the St. Lawrence river valley increase (Fig. 3.22b). This positive storm track anomaly corresponds to increased snowfall levels in Quebec. Throughout Newfoundland and to its East and Northeast there is a decrease in snowfall which corresponds to a decrease in track density. In node 3, the track density increase over the Midwest/Great Lakes corresponds with an increase in snowfall over the Lakes (Fig. 3.22c). Comparison with the total precipitation signature in which the associated region of positive anomaly extends south as far as Northern Texas, suggests storms that track through this region do not produce snow until they pass over the Great Lakes. Knowing that snow is most common in the cold sector of a storm, or the northwest

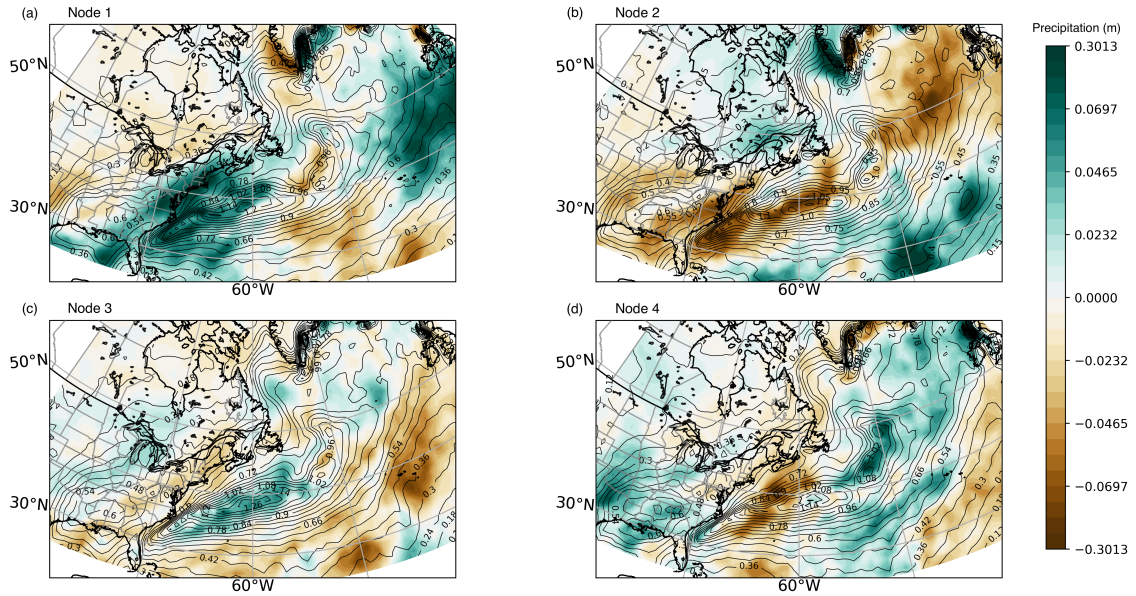


Figure 3.21: Mean seasonal precipitation (contours) and anomaly from the climatological mean (colours) for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

quadrant, the location of the positive snow anomaly northeast of Labrador can be connected with the storm track density increase northeast of Newfoundland. In node 4, which has below average storm track density along the US coast, the accumulated snowfall is below average over northeastern US and maritime Canada (Fig. 3.22d). An increase in storms over the Canadian prairies also matches up with an increase in snowfall in that region. Note that while the extent of this positive storm density anomaly continues south through much of the US, the snowfall field has a negative anomaly. This can be explained by referring again to the 2m temperature field for this node (Fig. 3.15d) which shows that the temperatures are warmer than average. Also noting that the total precipitation field shows a positive anomaly here (Fig. 3.21d), one can deduce that the negative snowfall anomaly is not related to fewer storms, rather it is caused by warmer than average temperatures making the storms more likely to produce rain than snow.

The fields of relative vorticity and 10m winds are also of interest, but due to more complex analysis being required to interpret the results, they were not analysed in this work. Nevertheless, the SOM-clustered seasonal means of relative vorticity and 10m winds are included in Appendix B.

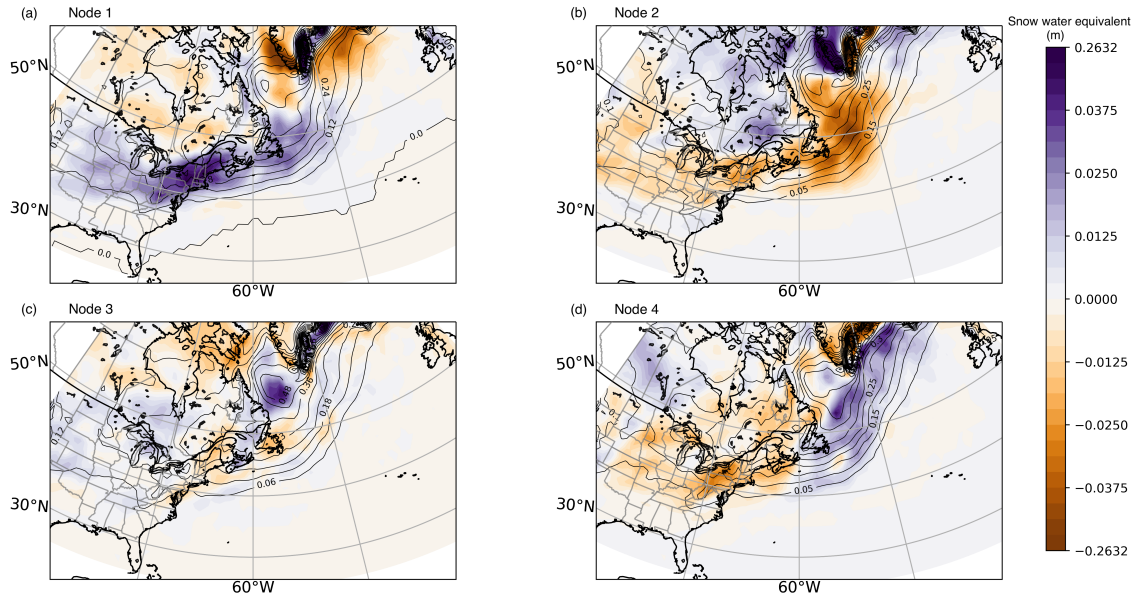


Figure 3.22: Mean seasonal snowfall (contours) and anomaly from the climatological mean (colours) for (a) node 1, (b) node 2, (c) node 3, and (d) node 4 of the SOM-grouped seasons.

3.5 Discussion

The analysis of spatial and temporal connections between atmospheric variables and ETC track density identifies possible predictor-predictand relationships. This also confirms a physical basis for the SOM clustering of storm seasons, which can give confidence in the patterns seen emerging in the mean storm track density SOM results. A summary of the main track density patterns (coloured ovals) and related atmospheric variables (text) is visualized (Fig. 3.23).

The main feature of node 1 is a shift in the climatological storm track closer into the coast. This is deduced from the increase in storms along the east coast of North America from the coast of the Carolinas to Newfoundland coupled with the decrease in storms offshore along a similar axis. The drivers that support this shift and strengthening of the typical storm track include lower 500 mb geopotential heights which increase curvature vorticity over the eastern seaboard, colder continental air temperatures creating a stronger land-sea temperature gradient, and a stronger upper level jet over the mid-Atlantic Bight with an anomalously meridional orientation increasing upper level divergence in key areas along the coast. The supporting effects are an increase in snow to the NW of the positive

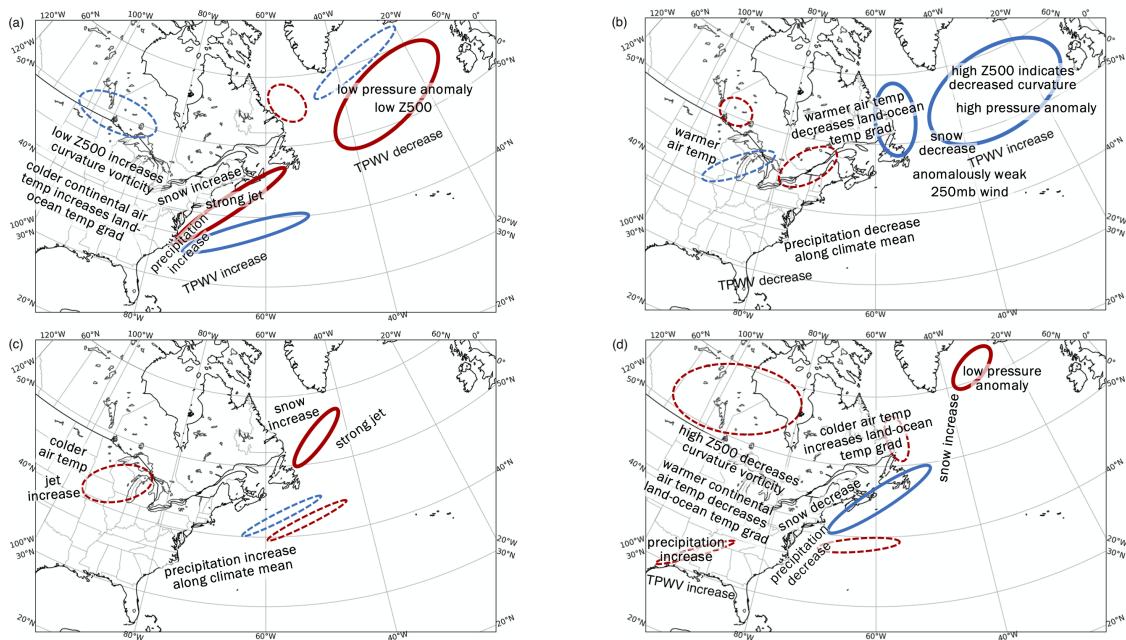


Figure 3.23: Visual summary of atmospheric variable-storm track density relationships in the SOM framework. The locations of main positive (negative) track density signals are shown by red (blue) ovals. Solid (dashed) outline shows stronger (secondary) signals. The atmospheric variables related to the identified signals are given in text in the approximate location that they occur. (a) SOM node 1, (b) node 2, (c), node 3, (d) node 4.

signal in the New England region and increased precipitation along the track.

The main patterns to note in node 2 are no strong anomalies in the track along the east of coast USA and a decrease in storms over the Labrador Sea. This is supported by the lack of strong anomalies in the fields of our drivers, except for anomalously warm temperatures over Labrador which would act to decrease the land-sea temperature gradient, thereby decreasing the available potential energy for storms in the Labrador Sea area. A decrease in snow coincides with fewer storms in that area.

The dominant features in node 3 are a strong positive storm anomaly extending NE from Newfoundland to the tip of Greenland and an anomalously strong Great Lakes storm track. The supporting drivers for both of these features are strong jets increasing upper level divergence which induces upward motion and supports the development of ETCs. The Great Lakes track increase may be supported by colder air temps over the Dakotas increasing the temperature gradient with the warmer air over the lakes and a negative height anomaly indicating troughing over the Plains. Both features are supported by snow increases along their NW side.

Node 4 is characterized by a decrease in storms along the coast of NE USA, Nova Scotia, and Newfoundland - a portion of the main climatological storm track. The drivers that support this signal are the same as node 1, but they act oppositely. The 500 mb heights are higher, the temperature is warmer, and the jet is weaker, all acting to decrease storms along the main coastal track. Supporting effects include a decrease in precipitation along the coast and a decrease in snow over New England and New Brunswick.

3.6 Conclusions

This chapter explored the climatology and variability of our storm track dataset. The gridded climatology showed high track density along the east coast of North America extending west and east of Greenland in a pattern consistent with the known climatology of winter storms in the North Atlantic. The climatology also included higher track density over the US Midwest and Great Lakes along with another continental track across southern Saskatchewan, southern Manitoba, and northern Ontario - well known as the path of Alberta Clippers ([Lawson, 2003](#)).

The monthly climatology highlights some intraseasonal variability within our dataset, but this section focusses more on exploring the annual winter season variability. To quantify

this variability, I used an EOF analysis. The results showed little useful information other than to give the insight that the storm track density field is highly variable and is influenced by many independent modes of variability. Further information was gained through a SOM analysis. The grouping of seasons denoted as Node 1 indicate a common feature of increased storm track density along the northeast coast of the USA and Atlantic Canada. Node 2 showed little anomalous activity in the southern part of the study area, with some decreased density in the Labrador Sea and far northeastern Atlantic. The Great Lakes storm track is strongest in Node 3 along with a region that extends northeastward from Newfoundland to Greenland. The primary feature of Node 4 is a decrease in storm track density along the main climatological east coast storm track. The SOM results were then used to investigate physical connections using conditional means of potential ETC drivers and effects over the SOM-identified groupings of seasons. This analysis revealed many potentially useful predictor-predictand relationships with theoretical bases for the mechanisms at play. Specifically, the association between storm track density and 2m air temperature differences over North America and the Atlantic Ocean was identified as a key relationship. The connections identified in this section will be utilized to inform the prediction model developed in the following Chapter.

CHAPTER 4

PREDICTION

Extratropical cyclone activity in eastern North America and the adjacent North Atlantic region has variable characteristics from season to season as demonstrated in the previous chapter. This seasonal variability is connected with specific atmospheric processes and variables. In this section, those relationships are utilized to develop a forecast model for the seasonal characteristics of ETCs.

The aforementioned variability is noisy and spatially nuanced. This makes prediction difficult and makes it less likely to be able to build a model with any significant skill. The complexity of the problem is reduced by reducing the area of interest to focus on the region around Halifax, Nova Scotia, Canada. The forecast model is developed with the application of operational usage in mind. Therefore, we elect to use predictors from the fall season so they can be observed and processed with adequate time to produce a forecast before the winter season. With a goal of aiding public forecasting, it is desirable to have a model that can give outputs that a typical citizen would value. For this reason, the focus is placed on forecasting frequency of storms and storm types, rather than trying to produce exact track locations or other specifics that are unlikely to have much predictability. The main ETC effects that the average citizen is concerned with are precipitation type, if any, whether or not there will be high winds, and what the overall intensity of the storms will be i.e. whether or not it will be a bomb storm. The storms that pass by Halifax are separated into types according to these effects and develop a multiple linear regression prediction model for each subset of storm type. Model results are evaluated and possible application to seasonal forecasting is explored.

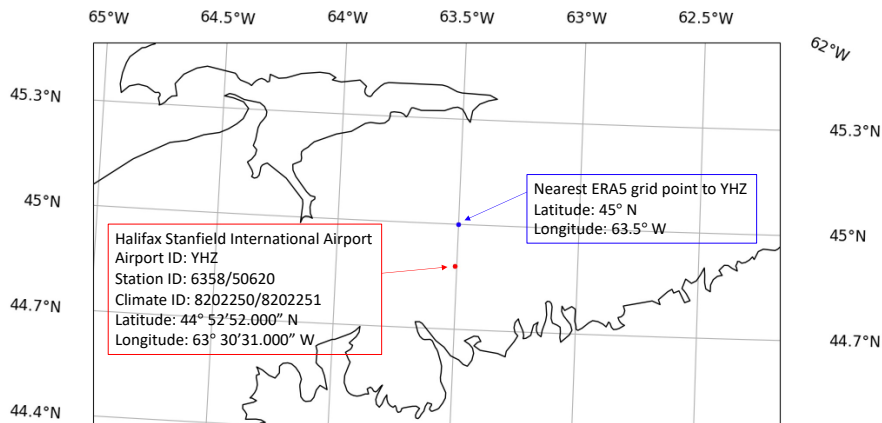


Figure 4.1: The location of the Halifax airport weather station where ECCC observations are gathered and the location of the closest ERA5 Reanalysis grid point.

4.1 Data

4.1.1 Local Weather Observations

To determine the effects of a particular storm at Halifax, observations of the weather are needed. Hourly observations from the Environment and Climate Change Canada (ECCC) weather station at the Halifax Stanfield International Airport are used. The closest grid point to Halifax in the ERA5 reanalysis is located approximately 13.25 km north of the airport weather station (Fig. 4.1).

Reanalyses are known to have biases, specifically in underestimating extremes (Fig. 4.2). Subplots a, b, and c show scatter plots comparing the value of a specific atmospheric variable observed at the Halifax weather station with the ERA5 reanalysis output of that same variable from the nearest grid point. The red line in each subplot shows where the points would lie for a perfect relationship between reanalysis and observational measurements. The correlation coefficient is given in text on each subplot. The fourth panel (d) illustrates the distribution of wind speed measurements from ERA5 and from ECCC. I created 35 bins of equal width for the ECCC wind observations that range from 0 to 93 km/h and 25 bins of equal width for ERA5 measurements ranging from 0 to 48 km/h.

Because the reanalysis assimilates the real station data we expect good agreement between the observations and the reanalysis values. The pressure values indeed show very good agreement with a consistent small overestimation by the reanalysis especially at

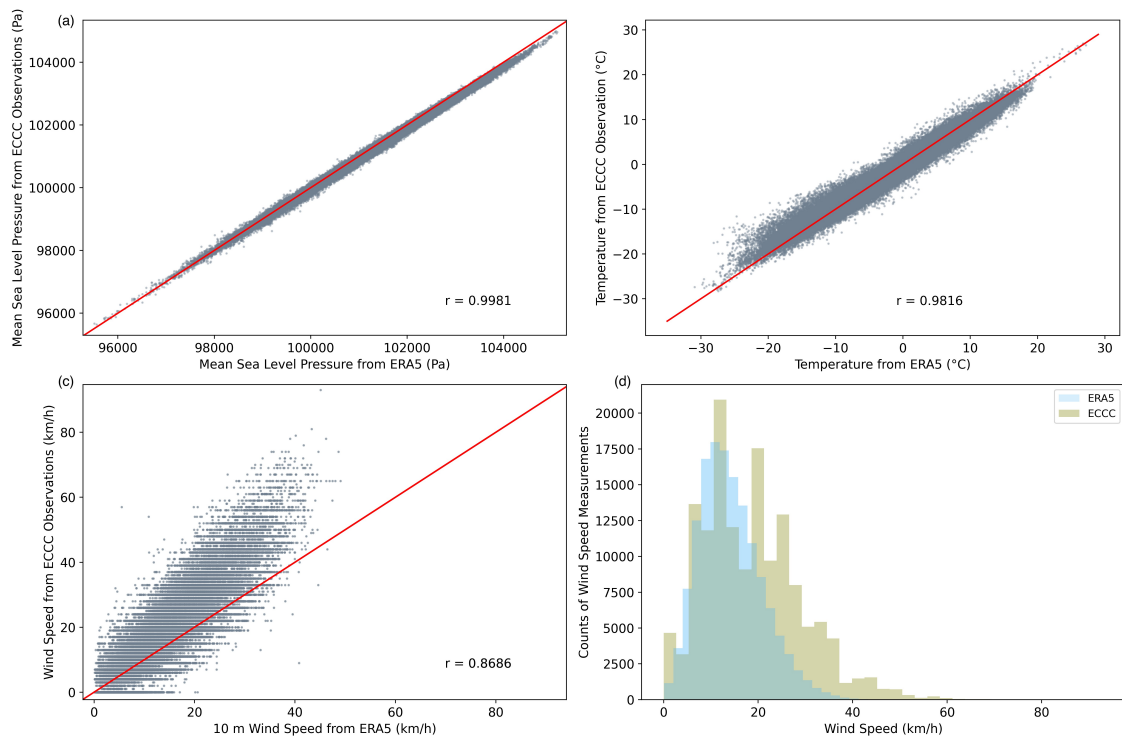


Figure 4.2: Comparison of ECCC station data and ERA5 reanalysis outputs from the nearest grid cell. (a) Pressure observations vs. reanalysis outputs, (b) temperature observations vs. reanalysis outputs, (c) wind speed observations vs. reanalysis outputs, and (d) histogram of wind speed from observations (green) and reanalysis (blue).

higher pressures (Fig. 4.2a). The temperature measurements also show good agreement, however at low temperatures, the reanalysis gives lower temperatures than those measured at the weather station (Fig. 4.2b). The reanalysis has a bias in wind speed, underestimating it except at low speeds (Fig. 4.2c). Looking at the distribution of wind speed measurements from the station (Fig. 4.2d, green) and the reanalysis (Fig. 4.2d, blue) in the bottom right, the reanalysis distribution is like a smoothed version of the observation distribution. There are especially few instances of high winds in the reanalysis dataset.

In the hourly station data obtained from ECCC, precipitation type is recorded in the weather remarks category. This category of the observations includes remarks such as “Rain”, “Freezing Rain”, “Cloudy”, “Snow”, “Blowing Snow”, “Fog”, etc. The ERA5 hourly dataset gives the amount of accumulated precipitation over the hour and the amount of accumulated snowfall over the hour. The combined rainfall and freezing rainfall amounts are inferred by subtracting the snowfall from the total precipitation.

Comparison of reanalysis precipitation and station precipitation data showed that precipitation is recorded by the reanalysis dataset on many more occasions than the Halifax weather station observations indicate that rain occurred. Precipitation poses unique challenges for reanalysis owing to the discontinuous nature of this field. Overall, the observations and reanalysis matched 59% of the time for rain/freezing rain, 56% of the time for snow, and 69% of the time for hours with multiple precipitation types recorded. Considering the cases in which rain or freezing rain is observed at Halifax, the reanalysis also indicates rain for 99% of those timesteps. If snow is observed at Halifax, the reanalysis also indicates snow for 97% of those timesteps. However, the reanalysis indicates precipitation at more than 3 times as many timesteps as the observations. This could be explained by spatial averaging in the reanalysis. These values are meant to be representative of the surrounding area while observations at the weather station are only concerned with the weather at that specific point. Due to the precipitation discrepancies, the observations made by a human observer at the weather station are used rather than the ERA5 reanalysis data.

Using weather observations also provides more descriptive information about the conditions experienced. At each timestep, the written remarks at the station include comments about the amount and characteristics of the weather conditions, such as “Heavy Drizzle”, “Moderate Snow”, or “Freezing Fog”, which can be used to categorize and describe storm

types more specifically than if the only information considered was total precipitation and snowfall amounts as in the reanalysis.

4.1.2 Predictands: Storm Time Series

The model takes a regionally specific approach by focussing on storms that affect Halifax, Nova Scotia. Such a dataset must be defined and built in order to develop a prediction model.

Storms are selected from the full storm track dataset (section 2) for inclusion in the predictand time series based on a single criterion: that their storm centre location is within 750 km of Halifax for at least one timestep (Fig. 4.3a). Given this set of storm tracks, an annual time series of total Halifax winter storms is calculated (Fig. 4.4a). Each unique storm track can contribute to the total counts in the time series exactly one time, regardless of the number of times it enters or exits the 750 km radius.

The dataset is then separated into seven subsets based on precipitation and winds at the Halifax International Airport. It is assumed that the weather experienced at Halifax is due to the storm when the storm is within 1000 km of the weather station given the typical length scale of ETCs is on the order of 1000 km. When assessing hours of precipitation due to the storm or hours of high winds due to the storm, the precipitation and/or wind are assessed at Halifax airport only for the timesteps when the storm is within 1000 km of the station. The combination of 750 and 1000 km radii allowed us to minimize the number of storms included that had little effect on the conditions at Halifax while retaining sufficient precipitation information from storms that did have an impact. The seven subsets of storms that pass within 750 km of Halifax are defined as follows:

i. Storms with at least 3 hours of precipitation (total precip)

If a storm is within 1000 km of Halifax for at least 3 hours, and precipitation is recorded at the Halifax airport for at least three of the hours that it is within that radius, the storm is determined to be a precipitating storm (Fig. 4.4b). Remarks that qualify as precipitation include the words “Rain”, “Drizzle”, “Freezing”, “Snow”, or “Ice Pellets”.

ii. Snow storms (snow)

All storms with at least 3 hours of precipitation are further classified as either a snow, rain, or mixed storm. To be a snow storm, the observation remarks must

include either “Snow” or “Ice Pellets” for at least 90% of timesteps during which precipitation is recorded, with the restriction that precipitation can only be recorded when the storm is within 1000 km of Halifax (Fig. 4.4c).

iii. *Rain storms (rain)*

To be classified as a rain storm, the observation remarks must include either “Rain” or “Drizzle” for at least 90% of timesteps at which precipitation is recorded when the storm is within 1000 km of Halifax (Fig. 4.4d).

iv. *Mixed precipitation storms (mixed)*

A storm can be grouped into the mixed precipitation category one of two ways. If a precipitating storm does not fit into the snow or rain storm categories, it is marked as a mixed precipitation storm because it is not predominantly one type or the other. Alternatively, a storm will be placed in this category if it records a “Freezing” remark in at least 90 % of the timesteps at which precipitation is recorded within 1000 km of Halifax (Fig. 4.4e).

v. *Storms with less than 3 hours of precipitation (no precip)*

Any system with less than 3 hours of precipitation recorded in the weather remarks at the Halifax airport when it is within 1000 km of the station is categorized as a storm with no precipitation (Fig. 4.4f). This category separates out storms that are geographically close enough to Halifax that they could have an impact on the region, but either aren’t large enough or developed enough to have a significant precipitation impact. These storms may have other impacts such as storm surge or high winds, although in our dataset there are no high wind storms with no precipitation. It also catches storms that are not within 1000 km of Halifax for at least 3 timesteps, since such a storm does not have enough timesteps to meet the precipitation requirements.

vi. *High wind storms (wind)*

If an hourly sustained wind speed greater than 44 km/h is recorded at Halifax while the storm is within 1000 km, it is classified as a high wind storm (Fig. 4.4g). The threshold of 44 km/h is the value of the 98th percentile of all wind measurements at Halifax throughout the study period. This category is not a subset of the precipitation

storms, therefore these storms may or may not also be found in a precipitation category or in bomb storms.

vii. *Bomb storms (bomb)*

A storm is considered to be of the bomb variety if it has a deepening rate of at least $12 \text{ mb} \cdot \frac{\sin\phi}{\sin 45^\circ}$ in 12 hours, where ϕ is the latitude of the storm centre (Sanders and Gyakum (1980); Fig. 4.4h). This category has no requirements pertaining to precipitation or wind speeds, although bomb storms are known to be associated with high impact weather.

Maps of the identified storm tracks show patterns that vary with storm type (Fig. 4.3). Theory of ETCs states that the cold sector is to the NW of the storm centre while warm air is typically to the E and SE. Assuming snow is the more likely precipitation type in the cold sector and rain is the more likely precipitation type in the warm sector, the results are consistent with warm and cold sector theory of ETCs. The vast majority of snow storms track to the south and east of Halifax (Fig. 4.3c) leaving this region in the cold sector of the storm and bringing snow. Conversely, most rain storms stay to the north of Halifax and the region is impacted by warm sector rain (Fig. 4.3d). Mixed storms tend to track almost directly over Halifax, between the preferential tracks of the snow and rain storms (Fig. 4.3e). Not surprisingly, the no precipitation storms are not near Halifax (Fig. 4.3f). It is likely that some of these storms are in the early stages of development and not yet causing precipitation, while others are smaller systems that are too far away for their effects to be felt in Halifax. High wind storms appear to be distributed throughout the range with possibly a slight concentration of tracks that follow the coastline from Long Island/Cape Cod through the Bay of Fundy into the Gulf of St. Lawrence (Fig. 4.3g). The majority of bomb storms track over the open water (Fig. 4.3h), consistent with the climatology given last chapter (Fig. 3.3). The moisture sourcing and lack of surface friction provided by the ocean can be important factors in the bomb deepening process.

Annual time series of storm counts are calculated for each of these categories or types as they will often be referred to (Fig. 4.4). Rain, snow, and mixed storms sum to the precipitation storms. The precipitation storms and the no precipitation storms combine to make up the total storms. Each storm within the high wind storms and bomb storms categories will also be found in either the precipitation or no precipitation storms categories

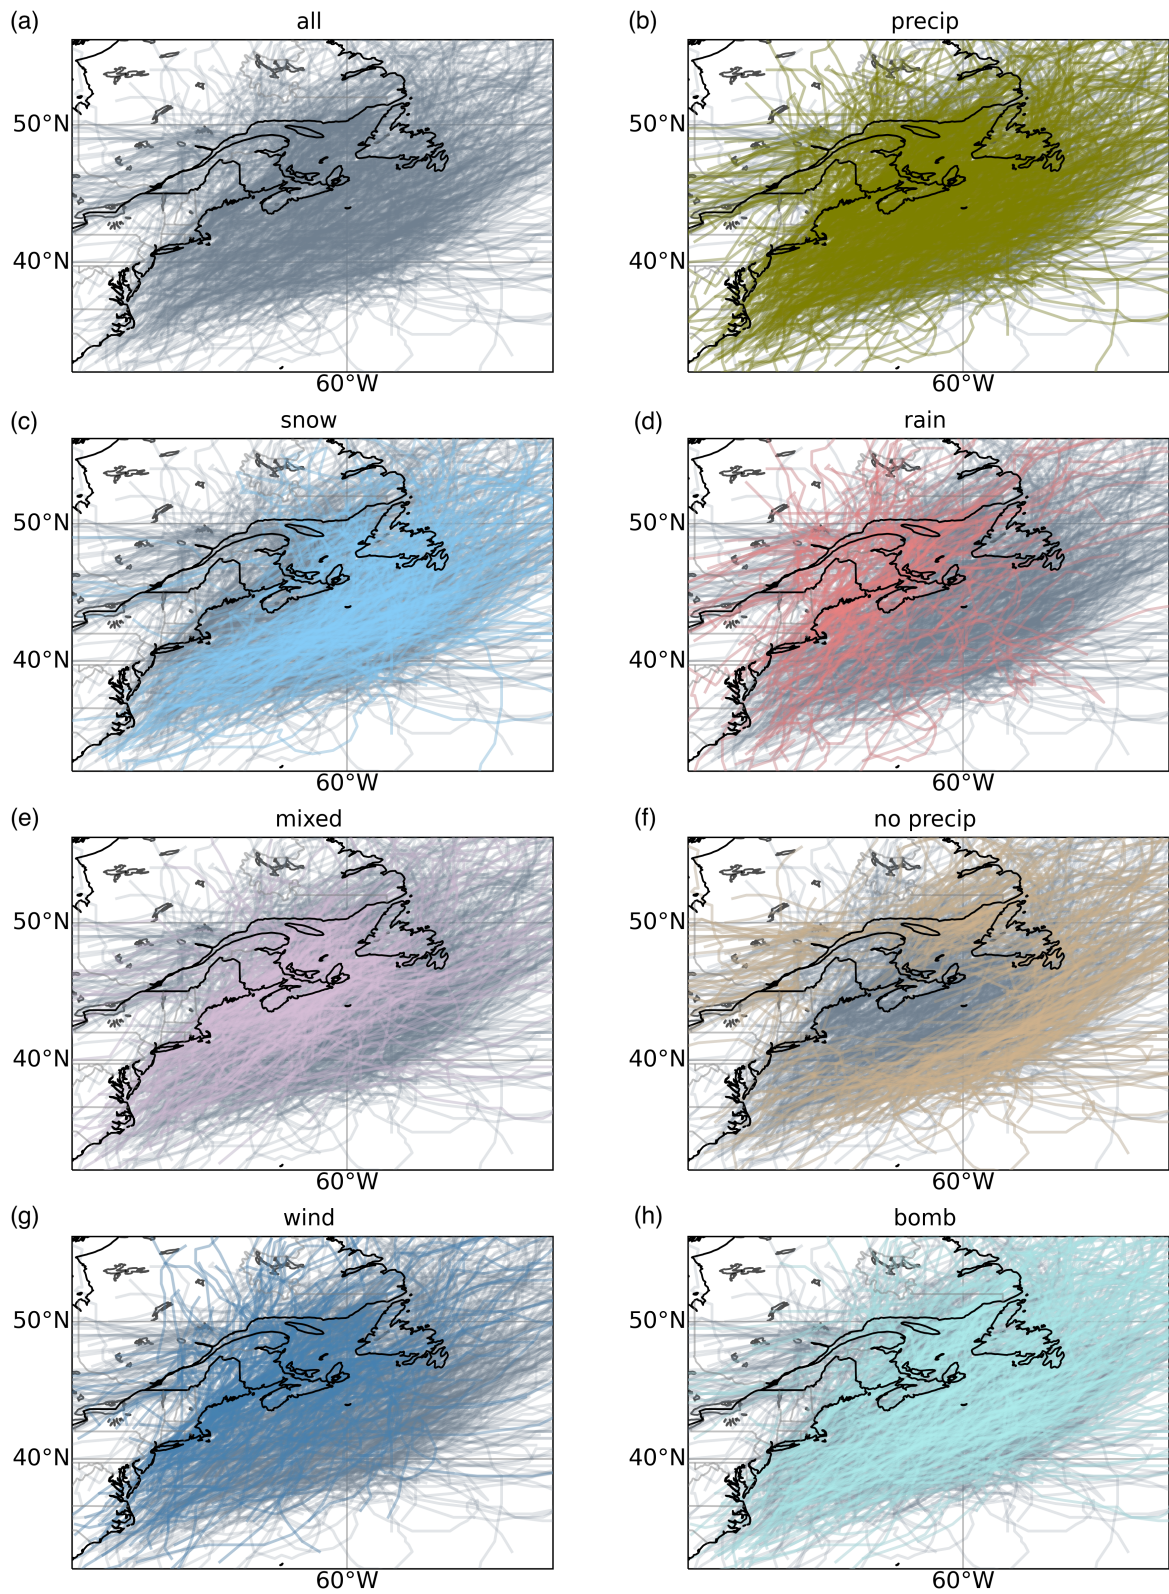


Figure 4.3: Tracks of all storms passing within 750 km of Halifax in (grey) overlaid with storms of each subseries (colour). (a: total, b: total precip, c: snow, d: rain, e: mixed, f: no precip, g: high wind, h: bomb)

Type	Mean	Variance	Slope of Linear Trend	P-Value	Percentage of Total Storms
total	32.2	16.91	0.01	0.8364	100
total precip	24.85	20.23	-0.03	0.6155	77
snow count	10.32	11.62	-0.04	0.4404	32
rain count	5.75	6.69	0.02	0.6365	18
mixed count	8.78	4.07	-0.01	0.6713	27
no precip count	7.35	8.23	0.04	0.2754	23
wind count	7.38	10.43	0.16	0.0002	23
bomb count	11.05	8.15	0.04	0.3573	34

Table 4.1: Statistics of Halifax storm timeseries.

and possibly in one of the subsets of precipitation types. The basic statistics of these time series are summarized in Table 4.1. None of the time series have a statistically significant linear trend except for the wind storm time series ($p = 0.0002$) which increases at a rate of 0.16 storms per season.

4.1.3 Predictors: Atmospheric Variables

The model predictors are chosen from the same set of ERA5 reanalysis atmospheric variables used in the previous chapter. See 3.1.2 for more details.

The use of time-lagged seasonal storm counts as potential predictors was also considered. For example, we might consider using the storm counts from one season to forecast the number of storms in the following season, which would be considered a lag-1 predictor-predictand relationship. The viability of these predictors was assessed through autocorrelation (ACF) and crosscorrelation (CCF) analysis. This revealed very few significant correlations, but some interesting cyclic nature in the ACFs. The use of lagged storm time series as predictors is ruled out, but the ACF and CCF plots are included in Appendix C for the readers interest.

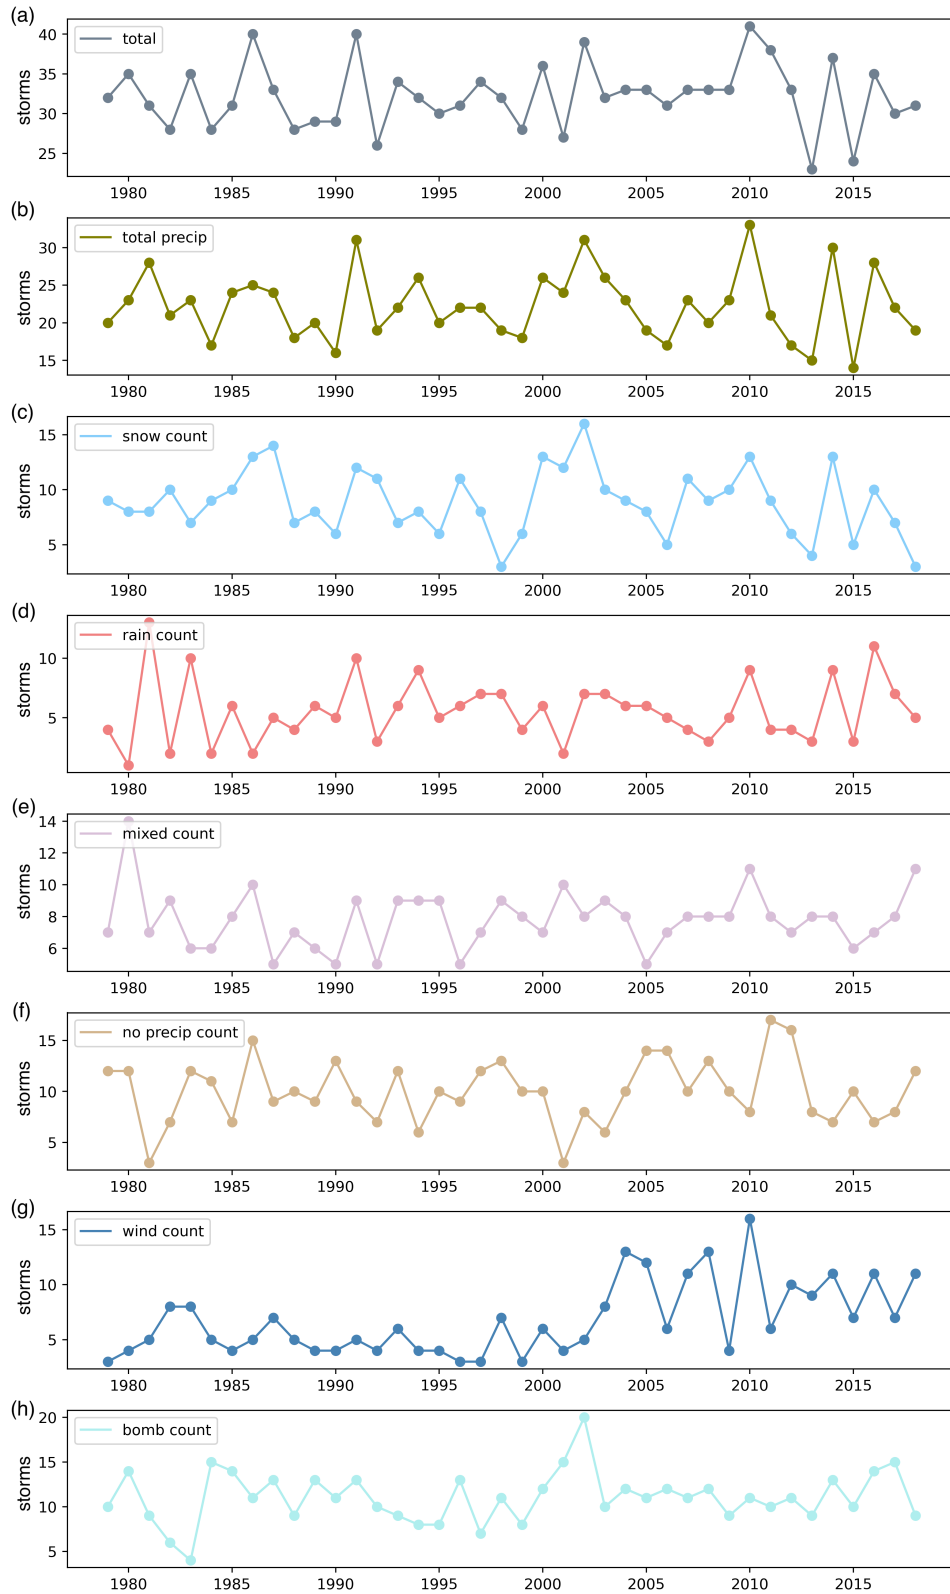


Figure 4.4: Time series of (a) all winter storms affecting Halifax and the seven subseries of storms (b: total precip, c: snow, d: rain, e: mixed, f: no precip, g: high wind, h: bombs).




Figure 4.5: Visualization of three different methods of prediction using linear regression models.

4.2 Prediction Model

4.2.1 Multiple Linear Regression

A multiple linear regression (MLR) model is a simple, but often powerful method of predicting future behaviour within a system that combines the effects of multiple independent variables to predict an outcome. The basic components of a linear regression are one or more predictors and a predictand with a regression coefficient that gives the linear relationship between them.

Given a predictand (y) and a predictor (x) timeseries, a linear regression can take on three basic forms (Fig. 4.5). Model 1 leverages a time lag relationship between the predictand and itself. This model type is particularly useful for forecasting a time series with a strong periodic nature. It takes the following form:

$$y(t + \tau) = \beta_o + \beta y(t) + \varepsilon \quad (4.1)$$

where t indicates current time, τ is some lag so that $(t + \tau)$ is a future time, β is the regression coefficient, β_o is the intercept, and ε is the residual or the portion of $y(t + \tau)$ not explained by a linear relationship with $y(t)$.

The second form (Model 2) uses the relationship between a predictor at the current time (t) and the predictand at some time in the future ($t + \tau$). The model can be described by

the following equation:

$$y(t + \tau) = \beta_o + \beta x(t) + \varepsilon \quad (4.2)$$

This model uses information independent from y (i.e. x), to forecast y at a future time.

The third model type (Model 3) is a two-step model which combines Models 1 and 2. The first step is to use Model 1 to forecast x at a future time based on its present value. The second step is to use a zero-lag relationship between the predictor, x , and predictand, y , to obtain the value of the predictand at the future time. Step one is given by

$$x(t + \tau) = \kappa_o + \kappa x(t) + \epsilon \quad (4.3)$$

where κ is the regression coefficient and κ_o is the intercept. Step two relies on a zero-lag relationship between x and y which can be obtained from the following regression:

$$y(t) = \beta_o + \beta x(t) + \varepsilon \quad (4.4)$$

Assuming β is constant with time, the predicted value of $x(t + \tau)$ from step one (Eq. 4.3) is used to complete step two which can be written as

$$y(t + \tau) = \beta x(t + \tau) + \varepsilon \quad (4.5)$$

Step one could also be completed using a model of a different type, e.g. an existing forecast model from ECMWF. While substituting eq. 4.3 into eq. 4.5 shows that mathematically model 2 and 3 are equivalent, the important difference lies in the lag at which the predictor and predictand are related. In model 3, a lag-0 predictor-predictand relationship is used to obtain the regression coefficient β . However in model 2, the regression coefficient is obtained from a lagged predictor-predictand relationship. If the predictor is closely related to the predictand concurrently but not at a lag, model 3 may be the better choice even though it requires an extra step which adds a second error term.

The choice of model should be made based on performance and practicality of use. To test performance completely, I would need to fully develop all the model types and compare results. This would take a considerable amount of time, so I attempt to narrow down our choice with a preliminary analysis. Assessing the ACF of the eight predictand time series revealed no strong autocorrelations for any predictand (see Appendix C). Therefore,

Model 1, which relies on autocorrelation, is eliminated. Having a multistep model of type three adds more room for error. Any uncertainty in the first step propagates through to the second step and is compounded by the error of the second model. Error in step one could be minimized by relying on an existing seasonal forecast model of the predictors, however I did not have easy access to such model outputs.

This project proceeds, therefore, with model type two (Eq. 4.2), which incorporates a lag in time between predictands and predictors. With practical usage of the prospective model in mind, a short time lag is chosen. It is likely that dynamical predictors from the previous season are too far removed in time to have predictive power over the next season. I decide to use predictors that are averaged over the month of September immediately preceding the winter storm season (Nov-Mar). All the data required to forecast storm track activity in the upcoming season would, in practice, be available before the first day of the season and the time lag is small enough that one can reasonably expect the dynamics of the system to be relevant for the storm season.

4.2.2 Predictor Selection

In section 3.4.3, I identified five theoretical drivers of extratropical cyclones and showed empirical connections between these drivers and east coast winter storm activity. This is evidence that these drivers may be useful predictors of the Halifax storm season. I use them to build a pool of potential predictors for our model. Each potential predictor is a time series of a single variable at a specific location that are expected to have some usefulness in predicting storms. The variables included in the pool are 2m air temp, 500mb thickness, mean sea level pressure, 500 mb geopotential height, wind at 250 mb (u-component, v-component, and magnitude), and total precipitable water vapour. In addition, the spatial gradient of each of these variables was calculated using second order central differencing and added to the pool to make a total of 16 potential variables. Each of these variables is considered over a spatial extent that is bounded by longitudes of 103°W and 22°W and latitudes of 25°N and 63°N. Within this region, the data is available on a 1° by 1° grid which results in 3 198 possible predictor locations for each variable. The time series in this set are not independent. In fact, in some cases they are very highly correlated especially for possible predictors at proximate locations. Removal of this correlation will be addressed in the selection process. With 16 variables at over 3 000 locations, there are over 50 000 possible predictors in the predictor pool, $\mathbf{P}^P = \{\mathbf{p}_i^P | i = 1, 2, \dots, N\}$ where i is the predictor

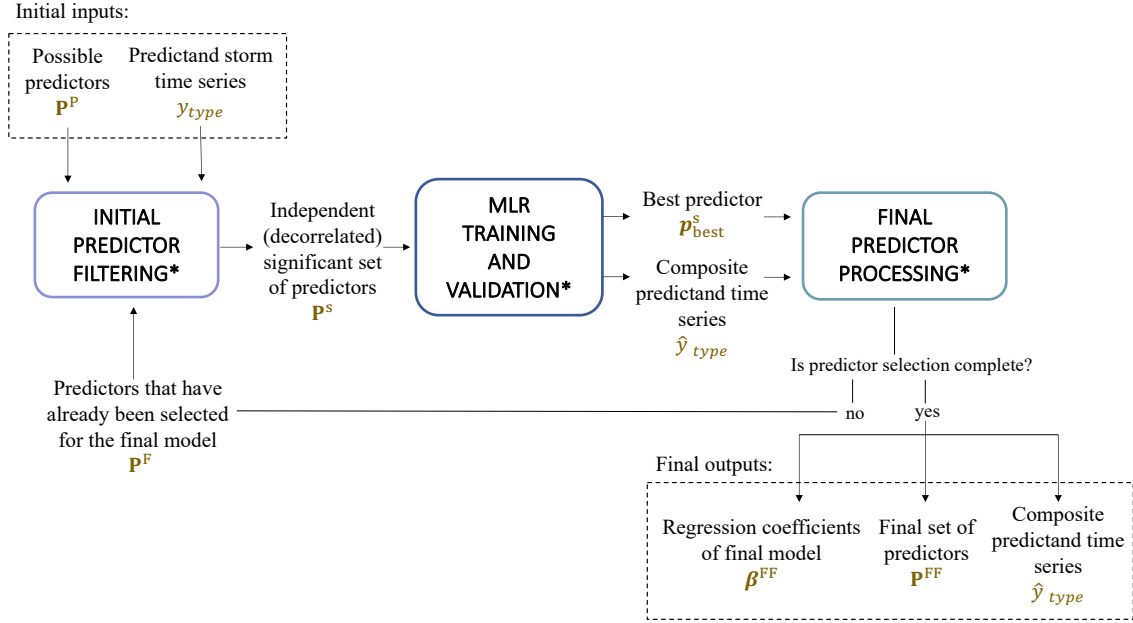


Figure 4.6: Predictor selection process overview. Processes indicated with an asterisk are further explained in subsequently in Figs. 4.7 (initial predictor filtering), 4.8 (MLR training and validation), and 4.9 (final predictor processing).

index and N the number of possible predictors in the pool. Each possible predictor, $\mathbf{p}_i^P = [p_1^i, p_2^i, \dots, p_{40}^i]$, is a 40 year time series (1979-2018) of the mean September value of the variable, p_i^i at its specific location. With the number of possible predictors being greater than the number of observations, the MLR problem is underdetermined. It is not possible to develop an MLR using the entire predictor pool. To solve this problem, cross-validation is used to select an appropriate set of predictors for the model. The process of selecting the best predictors for each storm type model from this broad pool involves three main steps: initial predictor filtering, MLR training and validation, and final predictor processing (Fig. 4.6). These steps are repeated until the addition of any remaining possible predictors fails to improve the model.

The initial pool of predictors used when building the model is common to all the subseries of storms (Fig. 4.7). However, the pool undergoes an initial filtering process as the first step of predictor selection for each subseries specific model. We remove the correlations between the possible predictors and all predictors that have already been selected for use in the model. This prevents the inclusion of redundant information in the model and makes the final coefficients clearer to interpret. To do so, the linear relationship between the set of previously chosen predictors (\mathbf{P}^F) and each possible predictor (\mathbf{p}_i^P) is

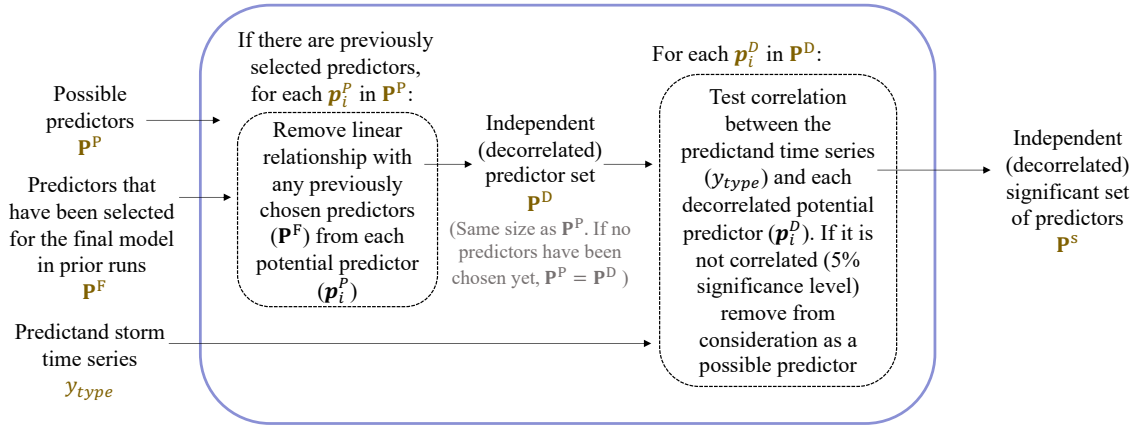


Figure 4.7: Initial predictor filtering schematic.

obtained through ordinary least squares regression. This relationship is then subtracted from the predictor time series (\mathbf{p}_i^P) to leave the independent portion of the time series (\mathbf{p}_i^D) for possible use as a predictor. This step is only necessary when selecting the second or greater predictor for the model. Each predictor time series in the decorrelated pool (\mathbf{P}^D) is then correlated with the predictand time series of that subseries model (y_{type}). Any possible predictor that does not have a statistically significant correlation with the predictand at a 95% confidence level is removed from the pool. The output from the initial predictor filtering step is a reduced set of possible predictors (\mathbf{P}^S) that are independent of any previously selected predictors and significantly correlated with the predictand time series (y_{type}).

The multiple linear regression model is then built by selecting predictors one by one from the filtered pool of possible predictors (Fig. 4.8). The process of selecting the best predictors from our large predictor pool is based on their correlation with predictands. In the first round of predictor selection, the model does not yet have any predictors chosen ($\mathbf{P}^F = 0$). Thus, the process starts by testing each possible predictor (\mathbf{p}_i^S) in a single variable linear regression. In the second round of predictor selection, a two variable linear regression model is used with the previously chosen predictor ($\mathbf{P}^F = \mathbf{p}_1^F$) and the possible predictor (\mathbf{p}_i^S) as the independent variables. In subsequent selection rounds, the number of independent variables in the multiple linear regression continues to grow with \mathbf{P}^F . The performance of the possible predictor is evaluated based on the RMSE of the predictions made from the test MLR. The test model is trained over a 30 year period (Fig. 4.8, training), and then used to forecast storm activity in the remaining 10 years of our 40 year study

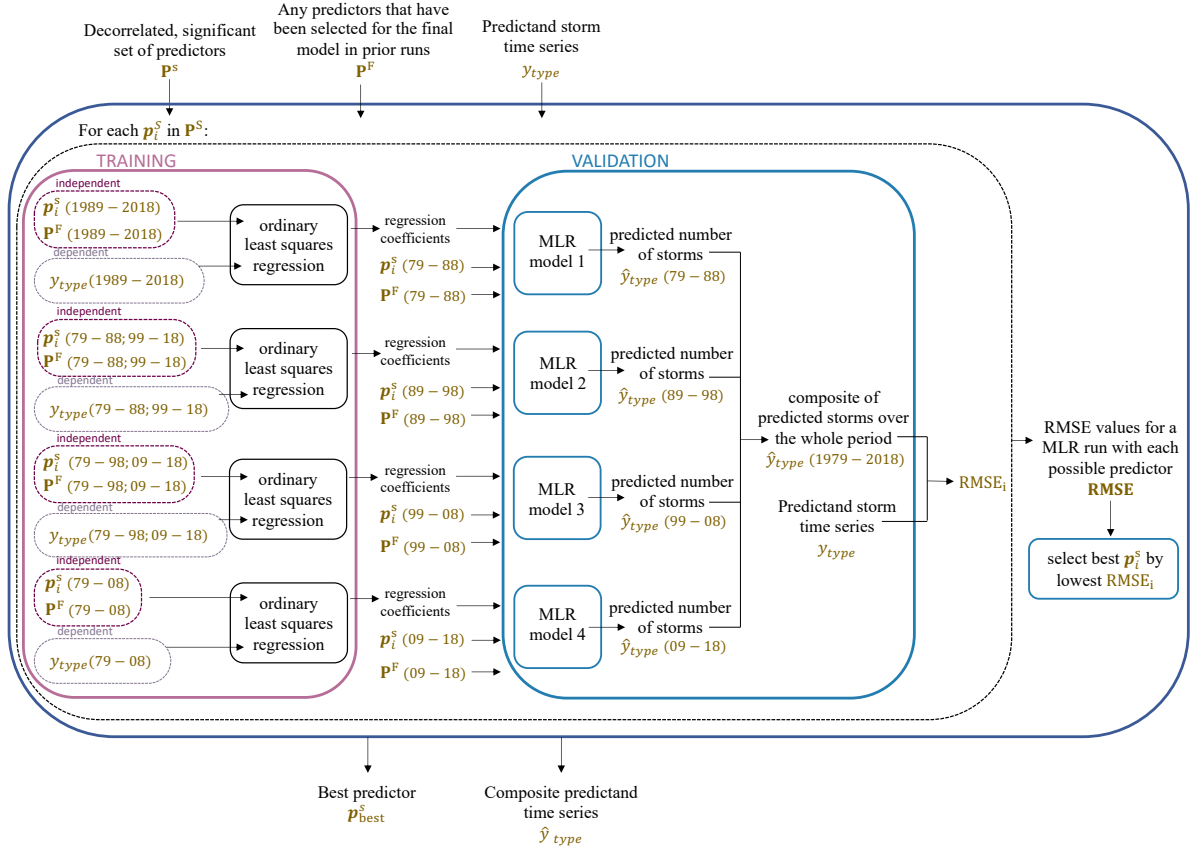


Figure 4.8: Multiple linear regression model training and validation process diagram.

period (Fig. 4.8, validation). This process is repeated four times with unique fitting and prediction periods each time, which allows for the predicted storm counts for each year in the 40 year period (\hat{y}_{type}) to be obtained without the forecast of any individual season being informed by the observations from that season (y_{type}). It also prevents choosing a predictor that by chance does very well over one time period, but not over another and is therefore not a robust predictor. The 40-year composite of forecasted storm counts over the validation time periods (\hat{y}_{type}), and the observed storm counts over the study period (y_{type}) are used to calculate the root mean squared error of the forecasts for each possible predictor ($RMSE_i$). The possible predictor (p_i^S) that combines with the previously chosen predictors (\mathbf{P}^F) to form the MLR with the lowest composite RMSE is chosen as the “best predictor” (p_{best}^S).

The last step is final predictor processing (Fig. 4.9). In this step, the procedures either prepare to select another predictor or finalize the model. If the addition of the best predictor improves the RMSE of the model, p_{best}^S is added to \mathbf{P}^F . Regardless of its usefulness, the

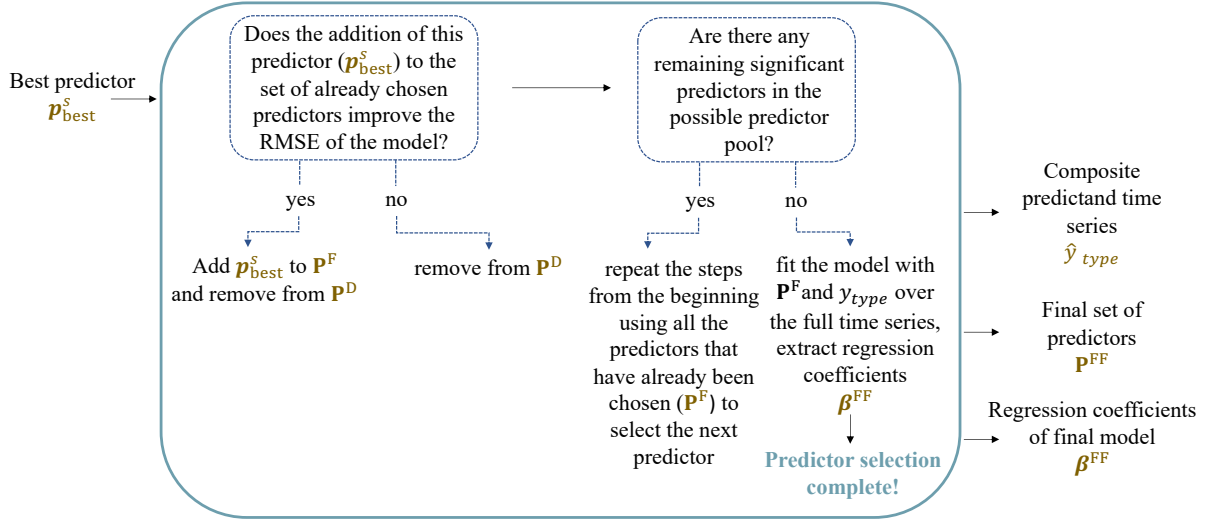


Figure 4.9: Final predictor processing flow chart.

best predictor is removed from future consideration in the possible predictor pool, \mathbf{P}^P . If any possible predictors remain in the possible predictor pool, the next step is to return to the beginning of the selection process to choose another predictor for the model. This process continues, adding new predictors to the model until eventually there are no possible predictors left. At this point, all original possible predictors that have not been added to the model cannot decrease the RMSE of the model if they are added as another predictor. The final model parameters can then be determined. The final set of predictors (\mathbf{P}^{FF}) is the set of predictors that have been selected until this point in the process. The final regression coefficients (β^{FF}) result from using ordinary least squares to fit a model with the full time series of \mathbf{P}^F as the independent variables and y_{type} as the dependent variable.

The set of predictors (\mathbf{P}^{FF}) chosen for each subseries model are given in table 4.2. The models range in size from two to four predictor models. The most commonly selected predictor is gradient of total precipitable water vapour and the second most common predictor field is the gradient of 2m air temperature. The frequency of these predictors is not surprising. The pattern of total precipitable water vapour in the atmosphere is a reflection of upper level flow and integrates many atmospheric fields. This makes it a likely choice in the prediction model since it combines the effects of many predictive features, such as temperature, moisture, vorticity, and mid-level winds into a single predictor. Recalling that the primary energy source of ETCs is baroclinicity explains the prevalence of 2m temperature gradient among the final predictors. Baroclinic zones exist where there

Predictand (y_{type})	Predictor 1 (P_1^{FF}) [lat, lon]	Predictor 2 (P_2^{FF}) [lat, lon]	Predictor 3 (P_3^{FF}) [lat, lon]	Predictor 4 (P_4^{FF}) [lat, lon]
total	gradient TPWV [57, 309]	gradient TPWV [49, 327]		
total precip	gradient TPWV [57, 310]	gradient T2M [29, 258]		
snow count	gradient T2M [48, 332]	gradient V250 [42, 331]	gradient TPWV [53, 311]	gradient TPWV [26, 277]
rain count	gradient TPWV [37, 331]	gradient TPWV [63, 284]	gradient TPWV [48, 321]	
mixed count	gradient T500 [63, 313]	gradient TPWV [29, 293]	gradient T2M [40, 310]	gradient U250 [29, 277]
no precip count	gradient TPWV [42, 279]	gradient T2M [29, 264]	gradient V250 [40, 319]	
wind count	T2M [28, 291]	gradient U250 [27, 281]	T2M [49, 310]	
bomb count	gradient TPWV [30, 292]	gradient V250 [55, 298]	WND250 [32, 309]	T500 [61, 322]

Table 4.2: Location and variable type of predictors selected for each type of subseries seasonal storm count model.

are strong spatial differences in temperature which is represented in the 2m temperature gradient field. The locations of the predictors are quite variable, but many are located downstream from Halifax in the climatological storm track (Fig. 4.10).

4.2.3 Model Fitting and Validation

To assess the performance of the final models (Table 4.2), I evaluate the composite predicted time series (\hat{y}_{type}). For each predicted value two prediction intervals are also calculated (66 and 90 % levels). A prediction interval (PI) gives the range within which a predicted value is expected to fall based on the uncertainty of the model’s ability to predict a specific value rather than simply predict the mean. It incorporates the sample uncertainty typically expressed by a confidence interval (CI) associated with the prediction of the mean in addition to the uncertainty associated with the new prediction. Due to this difference, a PI is always wider than a CI. In this model, PI is calculated for a given significance level (α) and time step (i) according to the following equation:

$$PI_{\hat{y}} = \hat{y} \pm t_{df}^{\alpha} \hat{\sigma}_{\hat{y}_i}, \quad (4.6)$$

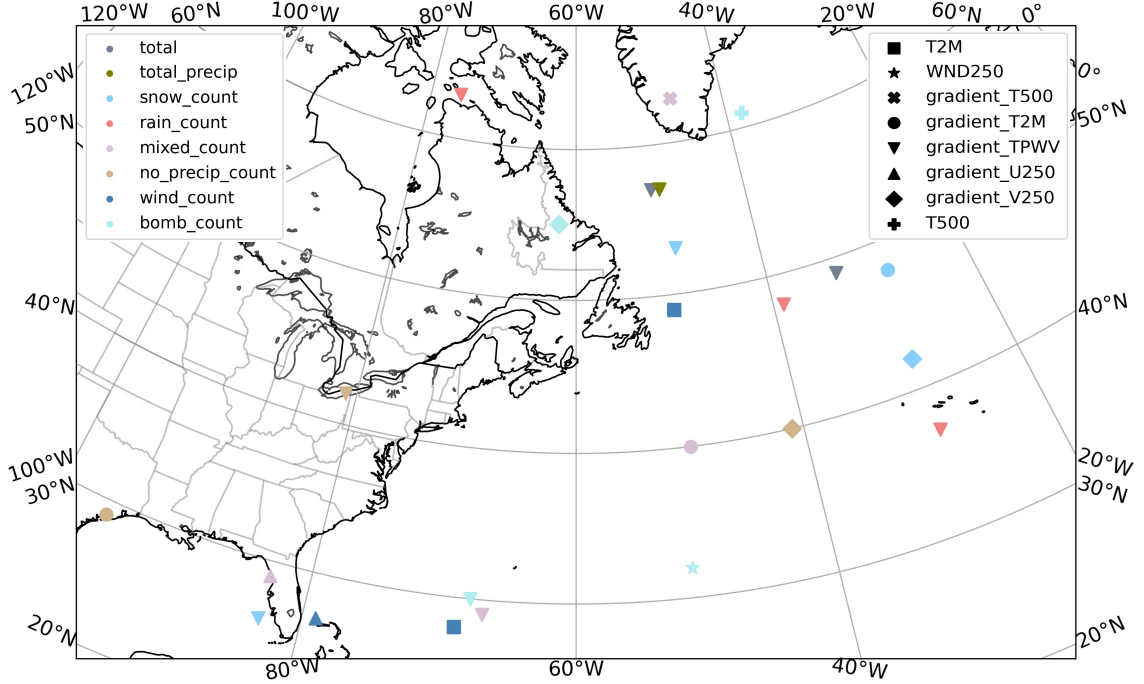


Figure 4.10: Locations of predictors chosen for each seasonal storm count model. Predictand storm type is shown by colour, predictors are indicated with marker shape.

where \hat{y} is the predicted value of storms in the season; t_{df}^{α} is the critical t -value for the specified significance level and the degrees of freedom (df) of the model; and $\hat{\sigma}_{\hat{y}}$ is the standard error of the i^{th} predicted value (Helwig, 2017). The standard error is defined mathematically as follows:

$$\hat{\sigma}_{\hat{y}} = \sqrt{MSE(1 + (P_i^{FF})^T((P^{FF})^T P^{FF})^{-1}P_i^{FF})} \quad (4.7)$$

where MSE is the mean squared error and P_i^{FF} is the set of predictor values at the i^{th} time step.

To validate the model, the 40-year forecasted storm activity time series and prediction intervals are compared with the observed storm activity over those 40 seasons (Figs. 4.11 and 4.12). We must remember that because the MLR problem was underdetermined, the relationships between error measurements that hold for a conventional MLR will not hold in this case. Thus RMSE, normalized root mean squared error (NRMSE), and cross-correlation are used to quantify the fit of each subseries prediction model (Table 4.3). All predicted model timeseries are correlated with their corresponding observation time series at a value of at least 0.78. Our best model as determined by cross-correlation

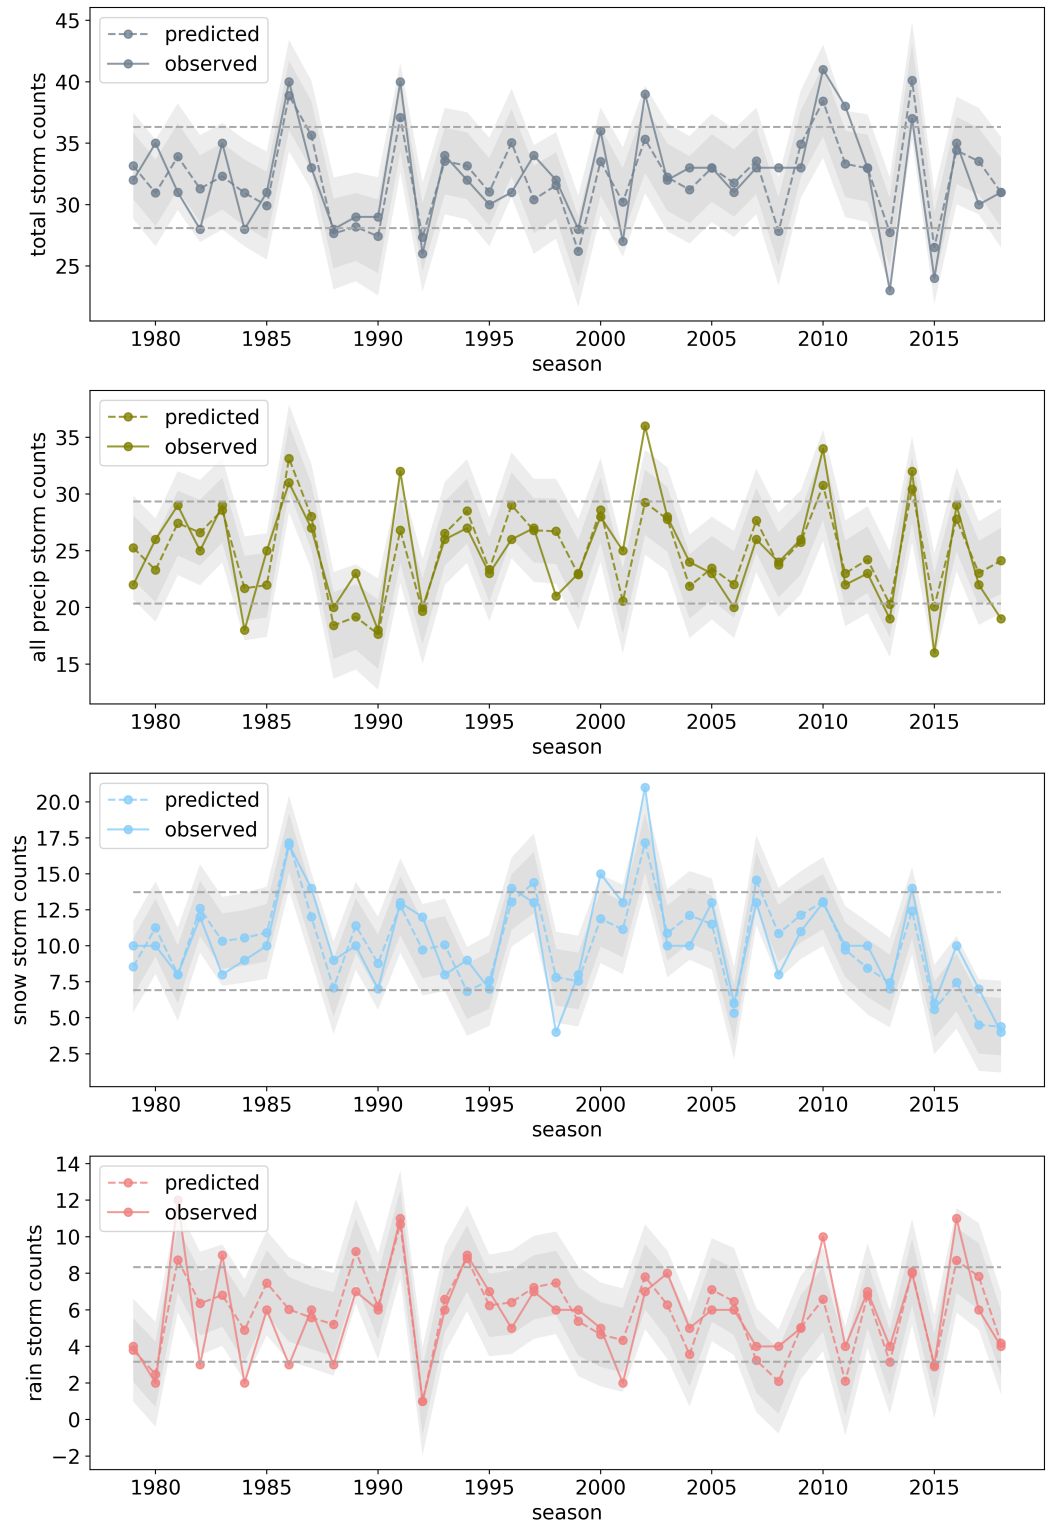


Figure 4.11: Model output time series validation for (a) all storms, (b) precipitation storms, (c) snow storms and (d) rain storms. Forecasted storm count timeseries (dashed line) shown with a 66 % prediction interval indicated by the dark grey shaded area and a 90 % prediction interval indicated by the entire light and dark grey shaded area. Observed storm counts (solid) shown for comparison. For each type, dashed lines indicate one standard deviation above and below the time mean observed storm tracks per season.

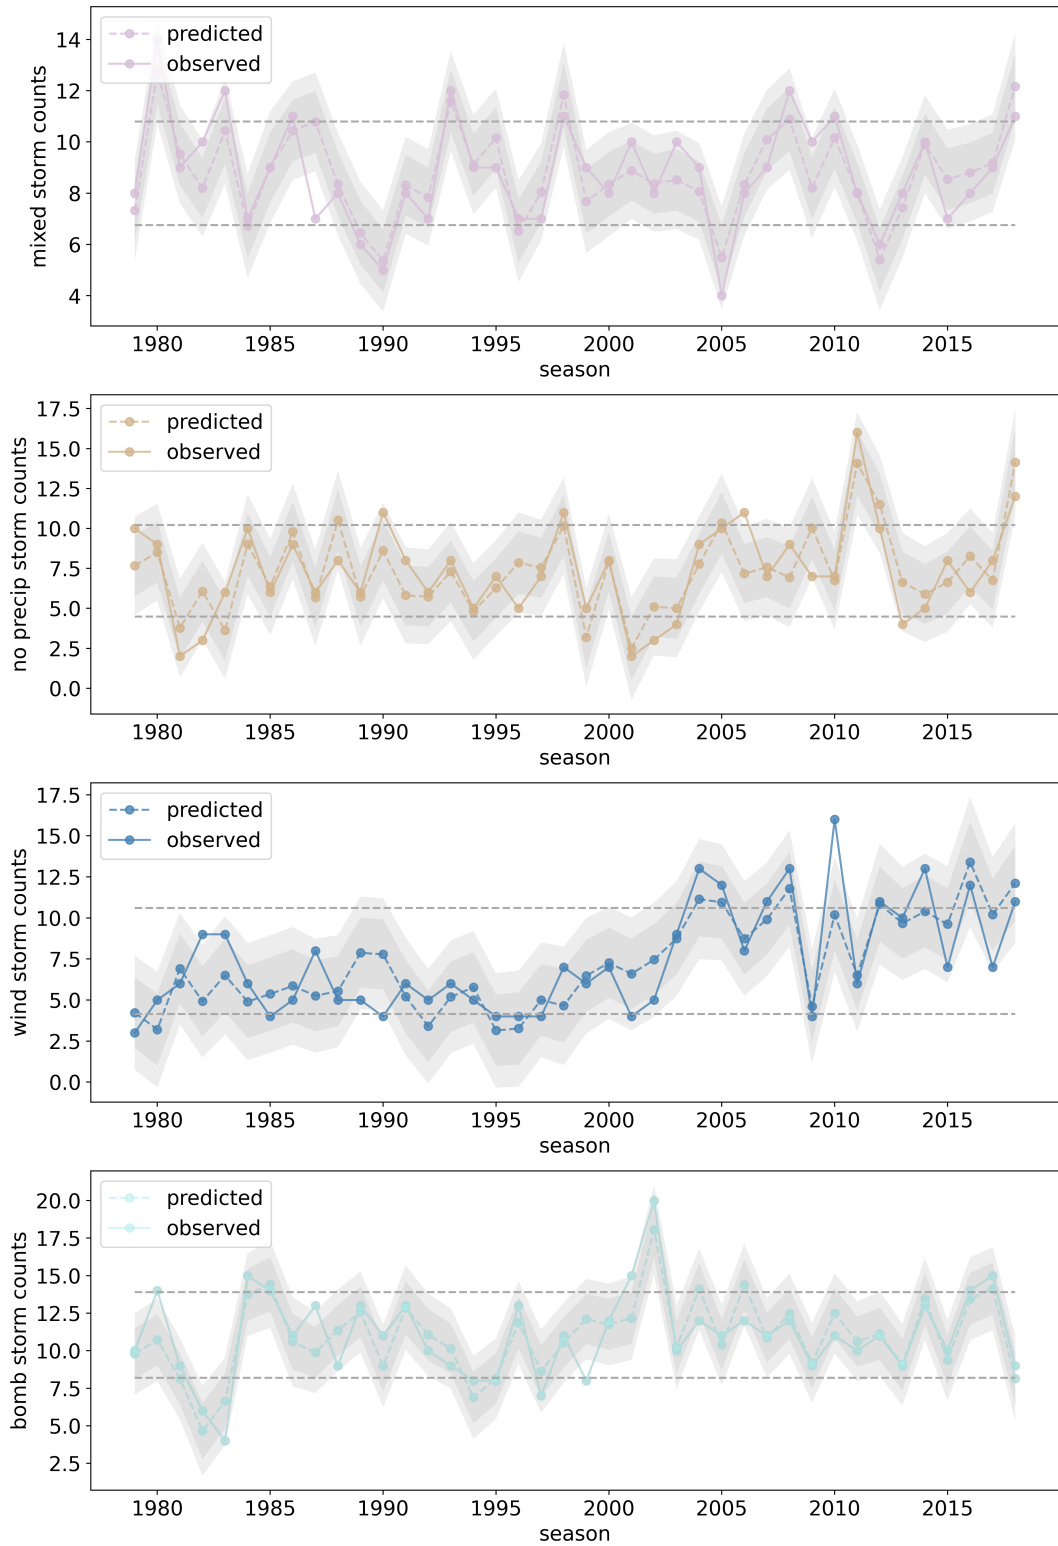


Figure 4.12: As in Fig. 4.11 but for (a) mixed precipitation storms, (b) storms without precipitation, (c) wind storms, and (d) bomb storms.

Predictand	RMSE	NRMSE	corr(obs, pred)
total	2.53	0.76	0.7888
total precip	2.65	0.72	0.8086
snow count	1.76	0.58	0.8580
rain count	1.61	0.74	0.7849
mixed count	1.09	0.62	0.8421
no precip count	1.72	0.68	0.8038
wind count	1.98	0.72	0.7923
bomb count	1.53	0.61	0.8442

Table 4.3: Validation results for all eight predictand models.

is the snow storm model at a correlation of 0.8580. The lowest RMSE is recorded by the mixed storm model. However, to compare RMSE across models, the value should be normalized. The NRMSE is calculated by dividing the RMSE by the standard deviation of each subseries (σ_{type}). When this is taken into account, one can see that while results are comparable across models, the snow storms model is best with the NRMSE equal to 0.58. Overall, the models forecast seasonal storm activity with some skill showing the chosen predictors have some predictive value. However, the large width of the prediction interval on each model may pose issues for model usage.

For future use of these models, I produce eight final equations for forecasting winter seasonal storm activity in the Halifax area. The regression coefficients are obtained from fitting the predictors and predictand time series over the whole 40-year study period (Table 4.4). For six of the eight MLRs, the intercept (β_0^{FF}) is similar to the mean storm activity value for that storm type. The value of the intercepts for the wind and bomb storms however, are much lower to account for the positive trend in these time series.

4.2.4 Seasonal Forecast of Halifax Winter Storms

A forecast made directly from the output of the MLRs has considerable uncertainty as shown by the prediction intervals on the model outputs. While this precludes the use of the models for deterministic forecasting, the model outputs can still be utilized for a probabilistic forecast. The framework for this usage is outlined in this subsection. Rather than forecast an exact number of storms, I create categories of storm track activity that give context to the forecast outputs. These categories are above average, average, and below average. The threshold values that separate these categories are defined based on the mean

Predictand	Intercept (β_0^{FF})	Coeff. 1 (β_1^{FF})	Coeff. 2 (β_2^{FF})	Coeff. 3 (β_3^{FF})	Coeff. 4 (β_4^{FF})
total	+42.0	-10.9	-3.48	-	-
total precip	+ 29.1	-13.8	-9.02	-	-
snow	+ 20.9	-20.4	-4.13	+3.24	+1.42
rain	+ 8.86	-3.77	+4.04	-2.00	-
mixed	+ 11.3	-0.0350	-1.36	+3.37	+0.966
no precip	+ 6.86	+2.60	-4.34	-2.57	-
wind	-1825.	+4.96	+2.43	+1.23	-
bomb	-176.	+2.35	-2.95	-0.547	+0.00357

Table 4.4: Model parameters for each type of subseries forecast model.

Storm track activity category	Range
Above Average	$\hat{y} > T_{type}^{up}$
Average	$T_{type}^{lo} < \hat{y} < T_{type}^{up}$
Below Average	$\hat{y} < T_{type}^{lo}$

Table 4.5: Classifications of storm activity. The upper threshold (T_{type}^{up}) separating above average and below average is one standard deviation above the mean and the lower threshold (T_{type}^{lo})

which separates average and below average is one standard deviation below the mean.

and standard deviation of the observed storm time series for each storm type. The upper threshold is one standard deviation above the mean ($T_{type}^{up} = \bar{y}_{type} + \sigma_{type}$) and the lower threshold is one standard deviation below the mean ($T_{type}^{lo} = \bar{y}_{type} - \sigma_{type}$; illustrated with dashed lines in Figs. 4.11 and 4.12). The above average storm season range lies above the upper threshold and the below average category lies below the lower threshold. The range considered average lies between the two thresholds i.e. within one standard deviation of the mean.

The probabilistic forecast determines the likelihood that the number of storms in a given season falls within each of these three ranges. To do so, the prediction interval (PI) is utilized. The PI follows a Student's t-distribution which is symmetric and dependent on the degrees of freedom of the model. At large degrees of freedom ($df > 30$), the t-distribution approaches a normal distribution (Heckert et al., 2002). Thus, the probability density function of the forecast can be represented using a normal distribution centred on the forecasted value (\hat{y}_{type}) with width given by the standard error of the prediction ($\hat{\sigma}_{\hat{y}_{type}}$).




Figure 4.13: Probabilistic forecast schematic for an example storm season. The probability density function of the forecast is illustrated with the curve and compared with the storm activity categories. The percentage of the area under the curve that lies within each category is used to make the written forecast statements.

To determine the likelihood that the number of storms in a storm season will fall within each of the three activity categories, assess the percent of the forecast distribution that falls above, between, and below the thresholds. This gives a probability as a percentage for each possible activity level (Fig. 4.13). The numerical probabilities are also categorised. The language of our proposed probabilistic forecast follows the IPCC (2014) probability language (Table 4.6).

The resulting language of the forecast would be, for example, it is *very likely* that Halifax has an *above average* number of *rain storms* this season. The italicized text indicates the

Probability terminology	Numerical probability
Very likely	$\geq 90\%$
Likely	66 – 89%
About as likely as not	34 – 63%
Unlikely	11 – 33%
Very unlikely	$\leq 10\%$

Table 4.6: Relationship between probability language and statistical likelihood.

parts of the forecast that would change based on the model output each season and the predictand being examined. Mathematically, the example above is stating that 90 % of the probability density function of the forecast lies greater than one standard deviation above the mean number of rain storms recorded over the 40 year study period.

This prediction model is applied to forecast the likelihood of above average, average, and below average storm activity in the 2019-20 and 2020-21 storm seasons, which were not used in any step of the model development process (Tables 4.7 and 4.8). The forecasts are also compared with the observed storm activity in those seasons (Table 4.9). The probabilistic forecast shows some skill with an accurate activity category being predicted more than half the time. The model seems most accurate when forecasting that the activity for the season will be average. There are some situations where the predicted and observed number of storms for a given type are within a 2 counts of one another, but the forecast is marked as a miss because the activity category threshold is between the predicted and observed values. This occurred in 2019 for both the mixed and no precipitation storm types. The probability distribution of the forecast is narrow for these projections which causes only a small portion of it to extend into the adjacent activity categories. Therefore, in the probabilistic forecast, the proximity of the prediction to the threshold between categories is not as well communicated as one might hope. The observed activity category (average) was the second most probable in the forecast, but was only given a probability of “unlikely”.

4.3 Summary

Forecasting eastern North American ETCs began with narrowing down our spatial focus. The previous sections showed that the field of ETCs in eastern North America is noisy and highly spatially variable. This makes it unlikely that a single model could successfully forecast over the entire broad region explored in the climatology and variability analyses. More accurate results can be gained from regionally specific or single location-focussed models. The Halifax, NS area was chosen as the focus of this study primarily due to its status as the largest Canadian urban centre located along the main climatological storm track. Winter storms are an important feature of local weather and culture in this area.

I used a multiple linear regression model and focussed on obtaining results that could be usefully applied to forecasting and give meaningful results for a citizen of Halifax. I assembled a large pool of possible predictors based on driving forces of ETC activity and

2019 Forecasts	Probability of Above Average	Probability of Average	Probability of Below Average
total	very unlikely 0.084	likely 0.8825	very unlikely 0.0335
total precip	unlikely 0.3066	likely 0.6908	very unlikely 0.0026
snow count	very unlikely 0.0696	very likely 0.9172	very unlikely 0.0132
rain count	very unlikely 0.02	likely 0.8435	unlikely 0.1365
mixed count	very unlikely 0.0001	unlikely 0.2564	likely 0.7435
no precip count	likely 0.7081	unlikely 0.2918	very unlikely 0.0001
wind count	very likely 0.9879	very unlikely 0.0121	very unlikely 0.0
bomb count	about as likely as not 0.6145	about as likely as not 0.3854	very unlikely 0.0001

Table 4.7: Probabilistic forecast of 2019 winter storm activity for 8 subseries. The most probable activity category for each type is italicized. Bold indicates when the observed storm activity falls within the forecasted highest probability category.

2020 Forecasts	Probability of Above Average	Probability of Average	Probability of Below Average
total	very unlikely 0.0302	likely 0.8742	very unlikely 0.0956
total precip	very unlikely 0.0011	about as likely as not 0.5898	about as likely as not 0.4091
snow count	about as likely as not 0.4003	about as likely as not 0.5994	very unlikely 0.0003
rain count	very unlikely 0.0124	likely 0.7762	unlikely 0.2114
mixed count	very unlikely 0.0	very unlikely 0.0867	very likely 0.9133
no precip count	very unlikely 0.0398	very likely 0.8962	very unlikely 0.064
wind count	about as likely as not 0.5767	about as likely as not 0.4225	very unlikely 0.0007
bomb count	very unlikely 0.0027	likely 0.7572	unlikely 0.2402

Table 4.8: As in Fig. 4.7 but for 2020 winter season

2019	y_{hat}	y_{obs}	Validation	2020	y_{hat}	y_{obs}	Validation
total	32.78	32	Hit	total	31.46	41.0	Miss
total precip	27.97	22	Hit	total precip	20.98	27	Hit
snow count	11.01	10	Hit	snow count	13.27	10	Hit
rain count	4.96	5	Hit	rain count	4.53	6	Hit
mixed count	5.9	7	Miss	mixed count	5.11	11	Miss
no precip count	11.23	10	Miss	no precip count	7.15	14	Miss
wind count	16.09	8	Miss	wind count	11.02	15	Hit
bomb count	14.37	7	Miss	bomb count	9.35	10	Hit

Table 4.9: Validation of 2019 and 2020 winter season predictions. Comparison of predicted (y_{hat}) and observed (y_{obs}) storm counts for each type with accuracy of the forecasts given in the validation column. The observed storm counts must fall within the most probable predicted activity category for the forecast to be considered a hit.

selected predictors from this pool based on their performance in the model. The models developed with these predictors showed the predictive power of atmospheric variables in the late summer/early fall on winter ETC activity, especially the spatial gradient of total precipitable water vapour and 2m air temperature. Though the models had skill, there remained some uncertainty. This uncertainty combined with the goal of usability for the general public lead me to suggest a more qualitative, probabilistic seasonal forecast model based on the MLR forecasts and the calculated prediction intervals to be used in operational forecasting.

Though the model could be further developed and refined, I believe I have presented a good basis for a simple model that has reasonable applicability to seasonal forecasting of ETCs in Halifax and the surrounding area.

CHAPTER 5

DISCUSSION

The work presented in this thesis complements the efforts of other scientists in the field to understand and predict the seasonal variability of extratropical cyclones. At the outset, I identified three primary objectives on the path toward the main goal of developing a skilled, practical seasonal ETC forecast. These objectives were addressed in three main science chapters. First, I developed a method to detect and track extratropical cyclones in eastern North America (Ch. 2). Then I explored the climatology and variability of the ETCs that had been detected and tracked (Ch. 3). Finally, informed by the relationships explored in the variability analysis, I developed a probabilistic winter season forecast model to forecast the frequency and types of storms that impact Halifax, Nova Scotia, Canada (Ch. 4).

5.1 Summary of findings

I developed a similar detection, tracking, and smoothing techniques to many other studies to build a dataset of ETCs (Neu et al., 2013). I was able to produce a dataset that well represents storms that impact the east coast of North America. The raw output of the detection and tracking algorithm contained some features that were not desirable for the analysis. This included features that did not impact eastern North America, were persistent low centres rather than transient eddies, or had characteristics that needed refinement, such as zig-zagged rather than direct tracks. Careful tuning of the algorithm as well as pre- and post-processing were required to ensure the dataset captured true ETCs and not these other features. The dataset obtained has similar characteristics to observations and other algorithm-tracked datasets both on aggregate (Figs. 1.4 and 3.1) and when comparing individual storm outputs (Fig. 2.6).

I employed the very common technique of EOF analysis to the winter storm field, as others have done ([Parding et al., 2019](#)). The individual modes of variability identified by this technique explained only small portions of the overall variability of the storm field. I then turned to SOMs which have been applied to atmospheric phenomena for many years ([Liu and Robert, 2011](#); [Hewitson and Crane, 2002](#)). Similarly, cluster analysis has often been used to group ETCs ([Gaffney et al., 2007](#); [Arnott et al., 2004](#)). However, the application of SOMs for clustering ETCs is novel. In the SOMs analysis and subsequent comparison with conditional means of atmospheric variables, I found that anomalies in storm track density within SOM-clustered winters can be clearly related with anomalies in drivers of ETCs within those same years. Node 1 showed increased storm density along the northeast states and Atlantic Canada connected with an increased land-sea temperature difference and a strong jet. Node 2 had no strong anomalies along the early part of the main climatological storm track which was tied to widespread average values of common ETC drivers, however later in the track, between Newfoundland/Labrador and Greenland, related to a weakened land-sea temperature difference. In node 3, there was a strong signal of higher storm density northeast of Newfoundland which was supported at the upper levels by a strong jet. In the final node (4), fewer storms were found along the northeast States and Atlantic Canada connected with warmer land temperatures and weaker troughing at 500 mb. I also identified clear spatial patterns in fields of ETC precipitation effects that coincide with increased and decreased storm activity in the SOM-clustered winter seasons.

I have proposed a probabilistic forecast for all winter extratropical cyclones that affect Halifax separated into seven subsets of storm types, based on extensive investigation of the climatology and variability of the ETC field on the east coast of North America. Storm types included categories based on precipitation type, wind, and storm intensity. The EOF analysis made evident the high degree of spatial and temporal variability exhibited throughout the study area. Therefore, the predictand field was simplified to be regionally focussed on Halifax, NS, allowing the temporal variability in the field to be quantified and predicted without complications of spatial variability. I developed 8 MLR models, one to forecast all storms and seven subset models each to forecast a unique storm type. Using the September mean values of ETC drivers identified in the SOMs analysis, the most skilled predictors were methodically selected for each model. Most often, the best predictors were found to be the gradient of total precipitable water vapour and 2m air

temperature and its gradient field. The best model as determined by R^2 is the snow storm model ($R^2 = 0.7342$). The forecast is built upon this set of 8 moderately skilled, storm type-specific MLR models. Through the whole development process, the intent was to create a forecast that communicates information of value to the general public. This was ultimately achieved with probabilistic descriptions that compare projected activity to usual winter storm season activity in clear, digestible statements. The forecast uses the probability distribution of the model output for a given season compared with typical seasonal storm counts for that storm type to determine the probability of an average, below average, or above average storm season. While this project focussed on Halifax, the framework presented could be applied at any location.

The ETC research field is currently lacking in seasonal prediction of ETCs along much of the east coast of North America, especially regionally nuanced forecasting. Storm characteristics, such as precipitation type or severity, are also not currently well forecasted. Existing seasonal prediction models for mid-latitude cyclones often utilize EOFs as both predictors and predictands as in [Parding et al. \(2019\)](#). The time series of the leading few principal components are used as predictands and because of the spatiotemporal relationship inherent to EOF analysis, predicting these time series provides spatial information as well. The predictors are then chosen from the leading EOFs of common predictor fields such as temperature gradient or 500 mb geopotential heights. Global climate models (GCMs) provide another way to obtain seasonal ETC forecasts. The skill of various GCMs have been assessed for their representation of frequency of ETCs during a winter season ([Yang et al., 2015](#); [Feng et al., 2019](#); [Renggli et al., 2011](#)). Overall, existing seasonal forecasting has shown skill forecasting ETCs in the North Pacific ([Yang et al., 2015](#); [Befort et al., 2018](#); [Feng et al., 2019](#)). It is thought that this predictability is related to ENSO ([Feng et al., 2019](#)), which has been identified as a predictor of ETCs in many studies. While [Plante et al. \(2015\)](#) and others have identified relationships between ENSO and eastern North American storm activity, these connections have not been extended to prediction. Some skill has also been identified in GCM forecasts of windstorms over Europe which is thought to be related to the NAO ([Befort et al., 2018](#)). In general, little predictability has been identified along the east coast of North America - the main area of interest of this study ([Yang et al., 2015](#); [Befort et al., 2018](#)). Even if some predictability exists, forecasts made from climate models, EOF-based predictions, and

teleconnection-based predictions have limited spatial resolution and limited information about storm characteristics making it difficult to get meaningful results when applying them to a specific location or region.

The need for a focussed, digestible seasonal ETC forecast is filled by the development of the regionally specific, multi-type probabilistic forecast model presented in chapter 4 which provides skillful forecasts of ETC activity differentiated according to storm impacts. Existing models can make broad sweeping statements such as, “The storm track will experience a northward shift this winter season.” The models presented here add spatial and impact-related detail to make more specific and applicable forecasts such as, “It is likely to be an above average storm season this winter. It is highly likely to be above average in terms of snow storms, likely to be below average for rain storms, and highly likely to be average in numbers of wind storms and bomb cyclones.” Focussing the spatial extent of the predictand and subtyping storms according to impacts has allowed for increased detail and enhanced usefulness for the general public. As previously explained, accurate seasonal forecasting of ETCs can increase disaster preparedness and mitigate human and economic loss. This regionally specific seasonal probabilistic forecast model framework gives a practical projection of likely winter storm season characteristics one month ahead of the upcoming season providing the lead time to allow for ample preparation to minimize losses.

5.2 Limitations

5.2.1 Storm detection and tracking algorithm

As discussed in the introduction, the primary types of storm detection and tracking algorithms are Lagrangian and Eulerian. A Lagrangian algorithm was chosen to allow for analysis of the characteristics of individual storms. Because of this choice, the storm field was likely noisier than if an Eulerian method had been used. Though this caused some issues in the variability analysis, which is elaborated on subsequently, the analysis was still successfully completed. The choice of Lagrangian tracking enabled the separation of storms into subgroups based on effects for a more detailed and applicable forecast in Chapter 4. Overall, the Lagrangian tracking technique was an appropriate choice for this study.

The precision of the track outputs depends on the model processes. Within the detection

step of the algorithm, the MSLP field was analysed for the presence of storms at 4 hPa intervals. This interval was chosen because 4 hPa intervals are the meteorological standard for contouring MSLP fields. A meteorologist who is manually marking the centre of a low pressure system on a surface map then uses scattered point observations accurate to the tenth of a hectopascal and surface wind fields to place the low centre within the innermost contour. In the algorithm, the low is placed at the pixel with the lowest pressure. For future use, if the intended use of the dataset is highly sensitive to the location of the low centre, the algorithm could be adapted to include surface wind fields in the placement step. It may also be beneficial to adjust the algorithm to analyse MSLP fields at smaller intervals if the exact area of a low pressure system is a key piece of the study. The area is calculated as the region inside the outermost closed MSLP contour that contains the low. Adding more contour levels will increase the precision of this calculation.

The tracked storms do not include storm velocity at each time step. The output tracks are composed of a series of coordinates that specify the location of the storm every hour. Using this information, it would be straightforward to calculate storm velocity vectors at each timestep. Storm speed is important for forecasting wind and precipitation impacts of a system. A slow-moving storm poses a flood risk whereas a fast-moving system is more likely to cause stronger winds particularly on the east side of the low. This information may be incorporated into the forecasting portion of this study. Aggregate or mean fields of storm velocity would be useful for inferring momentum fluxes and illustrating the climatological storm track.

5.2.2 EOFs and SOMs

The initial variability analysis using EOFs had limited success. This method did not identify modes of variability that explained any substantial amount of the storm field indicating the field is highly variable both in space and time. In order to obtain leading EOFs that explained more than 10% of the variance of the field, I had to increase grid cell size. Reducing the spatial extent to focus on the main track did not improve the results. This suggests that the issue was low sample sizes in each grid cell, a consequence of the high resolution of the grid. Many successful EOF studies on ETCs have used an Eulerian tracking method to produce the storm dataset which gives less noisy fields and is not limited by sample size ([Rogers, 1997](#)). Other studies that calculated meaningful EOFs on Lagrangian-tracked datasets used larger grid spacing, or populated cells based on the

number of storms that passed near as well as through a cell ([Parding et al., 2019](#)) to avoid sample size limitations. As an alternative to EOFs, I used SOMs in a novel application to the ETC track density field. By calculating conditional means of ETC drivers based on the SOMs clustering of seasons, the validity of the SOM-identified variability was established by tying anomalous storm density features with forcing by atmospheric effects. Conditional means of ETC effects were also used as evidence to support the presence of the storm track density patterns. A traditional statistical significance analysis could complement this physical validation approach. The SOMs analysis confirmed that annual winter variability exists in the mid-latitude ETC system and that the variability can be connected to possibly predictive atmospheric variables, particularly 2m air temperature.

[Gore et al. \(2020\)](#) are also working on understanding ETCs using SOMs. Contrary to the methodology used in this work, they obtain clusters by applying the SOMs algorithm to large scale meteorological patterns and evaluating the ETC differences that occur within the groups that were distributed according to the atmospheric patterns. Conversely, I choose to apply the SOMs algorithm directly to the ETC field, then evaluate atmospheric patterns within the ETC-directed groupings. While the synoptic typing-like methodology of [Gore et al. \(2020\)](#) may be easier to apply operationally, I surmise that allowing the storms themselves to dictate the groupings results in cleaner separation of unique storm types.

Future work on this topic should focus on quantifying the observed relationships between drivers, storm density, and effects. This would help to differentiate the more influential forcings from the secondary factors. Further, the quantified relationships could be applied in a predictive manner to determine variations in the storm field.

5.2.3 Prediction model

The physical interpretation of the predictor-predictand relationships in the proposed model is ambiguous. Chapter [3](#) demonstrated the concurrent relationships between ETCs and their known drivers within the same time frames. The predictor fields were selected based on those concurrent physical relationships between the storms and atmospheric variables. However, the dependence of predictands on the predictors in the model design is lagged in nature. The predictors exhibit skill which demonstrates that some sort of relationship does exist at this lag, however I have not developed a physical basis for the mechanisms or processes underpinning those relationships. A further investigation

into the predictor-predictand relationships at the lag used in the forecast model would give better insights into the mechanisms at play and build a stronger theoretical basis for the model. Such an analysis could reveal that predictors at a slightly different time lag have a stronger physical connection or that spatial averaging or other filtering of the predictor field gives more robust predictor-predictand relationships. Alternatively, a model of type three (see Fig. 4.5) would be easier to interpret physically based on the variability analysis presented in Chapter 3 because this model type uses no-lag predictor-predictand relationships. Choosing to build a model of this type instead of the type two MLR that I proposed may lead to a forecast with stronger or more straightforward physical roots.

Alternative methods of storm selection may be employed to better refine the predictand time series so they are more representative of impactful storms. The forecast models are predicated on the assumption that storms that are within 750 km of Halifax for a portion of their life cycle are relevant to the weather at Halifax. After testing multiple radii this distance was chosen because it balanced reducing the number of storms in the no precipitation category with keeping storms in other categories that affected weather at Halifax. Based on my analysis, the storms in the “no precip” category do not cause major impacts at Halifax (precipitation or high winds). These storms meet the 750 km criterion, but seem to have little to no impact at the location of interest. It may be informative to explore the effects of only selecting storms when the system is in the later developing and mature stages of its life cycle (as in [Plante et al. \(2015\)](#)). It is apparent that some of the storms in the “no precip” category begin their life cycle within the area, but since precipitation and high winds do not occur until later life cycle stages, the effects of that system are not experienced at Halifax. Within the storm dataset, the detected area of each system at can be found at every time step. This area measurement could be used to determine whether or not a storm is affecting Halifax at any given time step by determining the extent of the storm based on the area and assumed circular or elliptical shape of the system. Further accuracy may be gained by saving the actual shape of the system in the detection process rather than just the area.

The forecast models developed here include no intermodel relationships that would consider different model types jointly. Additionally, they do not consider counts from previous seasons of other storm types as possible predictors because this introduces varying predictor-predictand lags to the models. Since by definition there are mathematical

relationships between the categories of storms (e.g. precip + no precip = total storms), the models could be constrained to preserve these relationships. This could mean letting the total storms forecast simply be the sum of the precipitation and no precipitation forecasts, rather than forecasting total storms by its own individual model. Because the actual use of the model outputs is in a probabilistic forecast, the outputs don't need to be perfectly congruent. However, if an exact quantitative forecast of the storm types was derived from the MLRs, the number of storms in an aggregate category (e.g. precip storms) should be equal to the sum of its constituents (e.g. snow, rain, and mixed storms).

I have proposed a probabilistic forecast based on a set of MLR models. I believe the framework of the probabilistic forecast is well designed. It uses publicly available ERA5 data, has a practical lead time, can be readily applied at any location, and gives forecast statements that incorporate both context and uncertainty in a manner that is digestible for the general public. However in this thesis, the forecasting framework is applied using a set of MLRs that have not been extensively developed. Future work focussed on this endeavor could significantly improve the forecast system.

5.3 Conclusion

From the beginning, the intent of this thesis has been to develop a seasonal ETC forecasting system for storms in the northwest Atlantic. It started with developing detection and tracking method to build a dataset of ETCs in the area of interest. This endeavor was successful in creating a set of storms that both agreed with individual storm tracks and collectively matched with known climatology of ETCs in this region. Next, the interannual variability of the storm track density field was analysed using SOMs to cluster similar winter storm seasons. Using these clusters, I identified spatial and temporal connections between storm track variability and atmospheric drivers, most notably, with the 2m air temperature field. The development of the storm track dataset and the investigation of the climatology and variability of the ETC field culminated in the development of a winter season ETC forecast regionally focussed to Halifax, Nova Scotia, Canada. I established the framework of a probabilistic forecast that delivers estimates of storm type and frequency for the upcoming winter season with one month of lead time. The forecast system leverages the probability distribution of the forecast and the typical characteristics of a winter storm season to give a digestible forecast within a familiar frame of reference that incorporates

mathematical uncertainty in an accessible way. The main success within the forecasting study was not necessarily creating a well-refined deterministic forecast model, but in constructing a framework that takes a model output and its uncertainty and communicates it clearly to the general public with context. Sharing such a forecast ahead of an upcoming winter season can enable mitigation of losses, improved financial planning, and better overall preparedness for the impending trials of Mother Nature.

APPENDIX A

FULL SET OF EOF RESULTS

The EOF analysis in section [3.4.1.1](#) outlined the adjustments made in an attempt to improve the variability explained by the EOF output. The complete analysis for each of these adjustments is included in this appendix.

In our first attempt to reduce the variability of the field and obtain higher explained variability from the leading modes, I reduce the spatial field to focus only on the main track. The area chosen is a quadrilateral with corners at (35°N, 82°W), (35°N, 52°W), (60°N, 28°W), (60°N, 58°W). The spatial patterns and PCs of the first 4 modes are shown in figs [A.1](#) and [A.2](#). This increases the explained variance by the leading mode to 5.84% (fig [A.3](#)). While this is an improvement, it is only slight. I instead test increasing the grid cell size of our input climatology field.

The climatology of winter storms within the study area on a 5x5 degree grid is shown in fig [A.4](#).

The spatial patterns and PC timeseries from the first four modes are shown in figures [A.6](#) and [A.8](#). Using the 5x5 degree grid instead of the 1x1 degree grid for the EOF analysis doubles the explained variance from the first mode: 11.36 % (fig [A.9](#)). The leading 8 modes combine to explain half of the variance. It takes 16 modes to explain 75% of the variance. While this is a better improvement than the main track spatial restriction, the explained variance is still less than I want.

The final EOF analysis for the field of all storms is performed on the climatology at 10x10 degree grid size. Spatial patterns and PC timeseries of the first four modes are shown in figures [A.10](#) and [A.12](#). This gives a leading mode that explains 18.54% of the variance of the field ([A.13](#)). The explained variance of the first 4 modes combined is 52%.

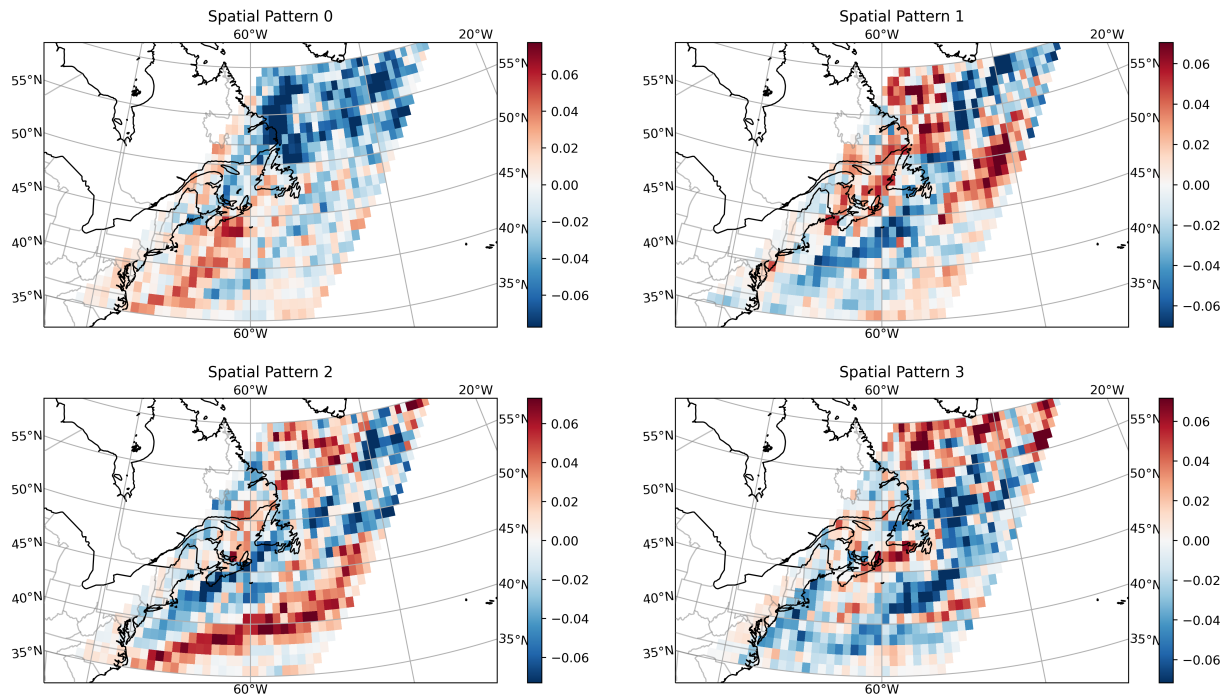


Figure A.1: Spatial patterns of the first 4 EOF modes for PCA on main storm track

This is a significant portion of variance explained by a few modes, which means analysing and discussing the variability in time of these patterns would be useful for understand and possibly predicting the variability of the whole system. However, using grid cells this size causes a loss of much of the detail within the field.

A.1 Bomb Storms

As with the field of all storms, I attempted to reduce variability in the field by reducing the spatial extent of the study area. The leading mode of the reduced spatial field EOF analysis explains 5.8% of the variability (fig [A.16](#)). This is actually less than when I used the full study area. The spatial pattern of the leading mode (fig [A.14a](#)) indicates an increase or decrease in storms over the entire main track, a common first EOF pattern. Modes 2, 3, and 4 (fig [A.14](#)) show lateral track shifts similar to the second and third mode of the full spatial field EOF analysis.

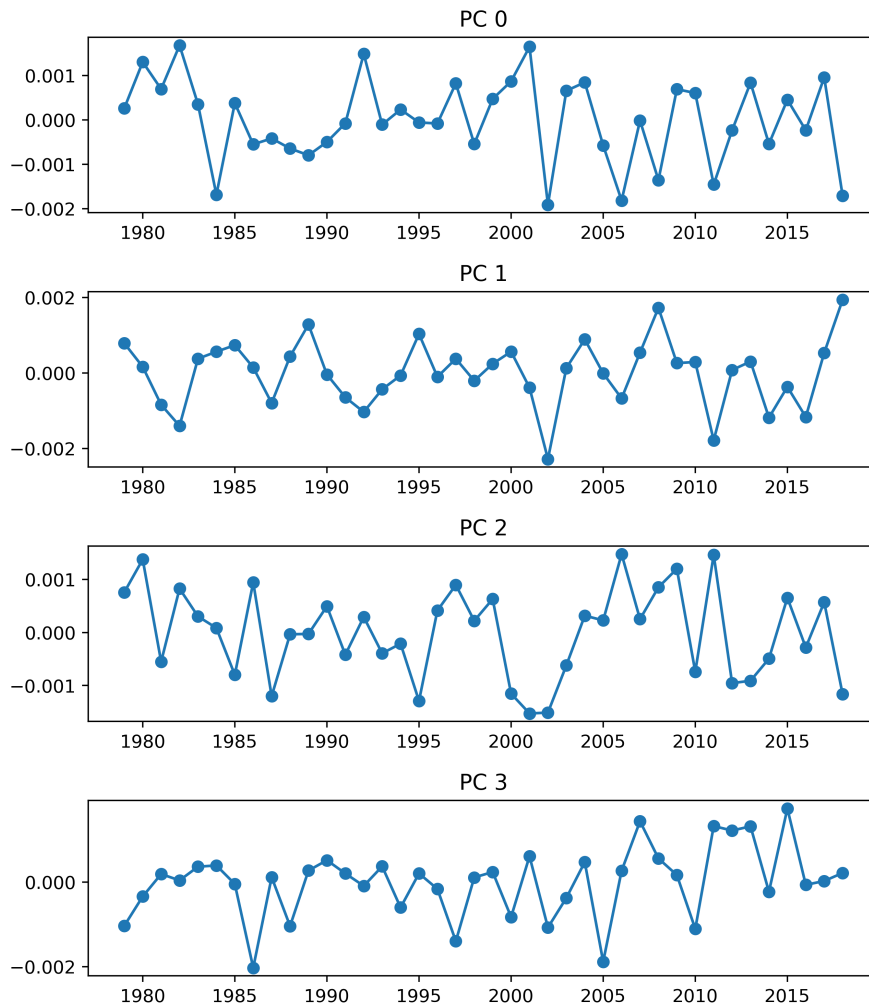


Figure A.2: Principal component timeseries of the first 4 EOF modes for PCA on main storm track




Figure A.3: Fraction of total variance explained by each principal component for PCA on main storm track

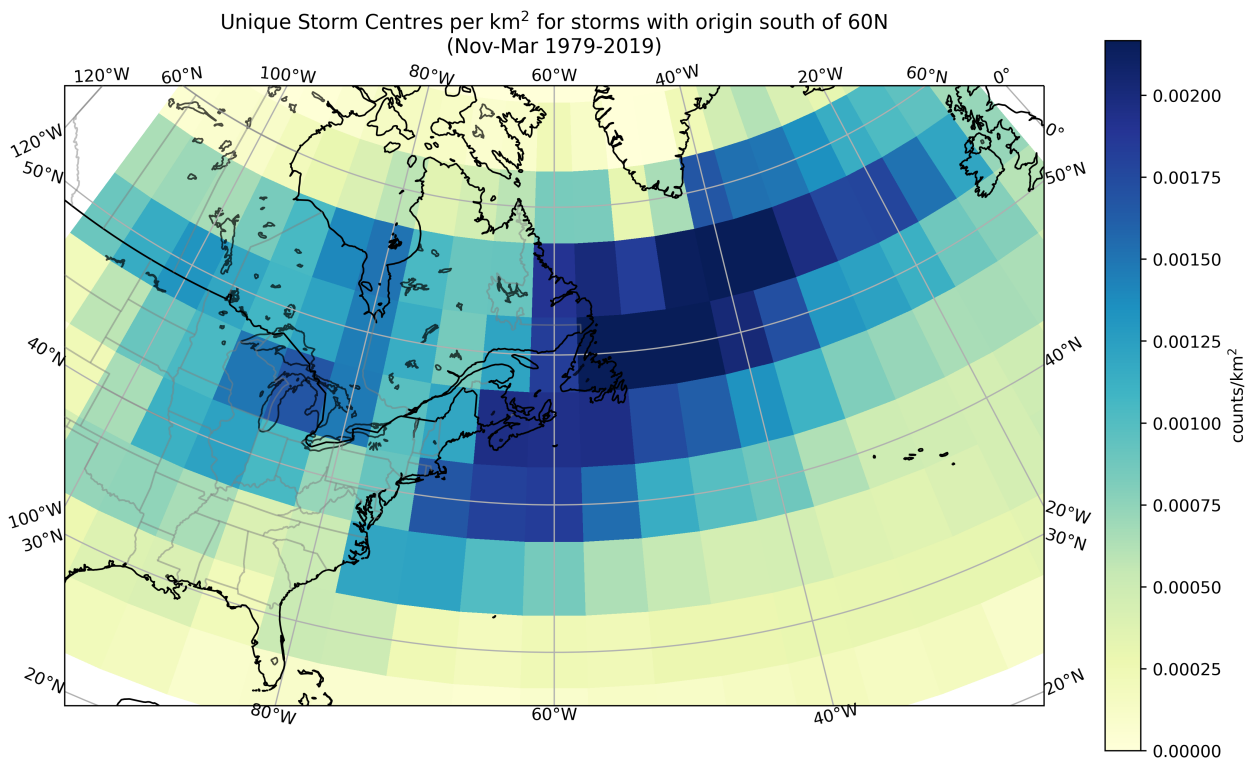


Figure A.4: Unique storm centres km⁻² for storms with origin south of 60°N on a 5 degree grid

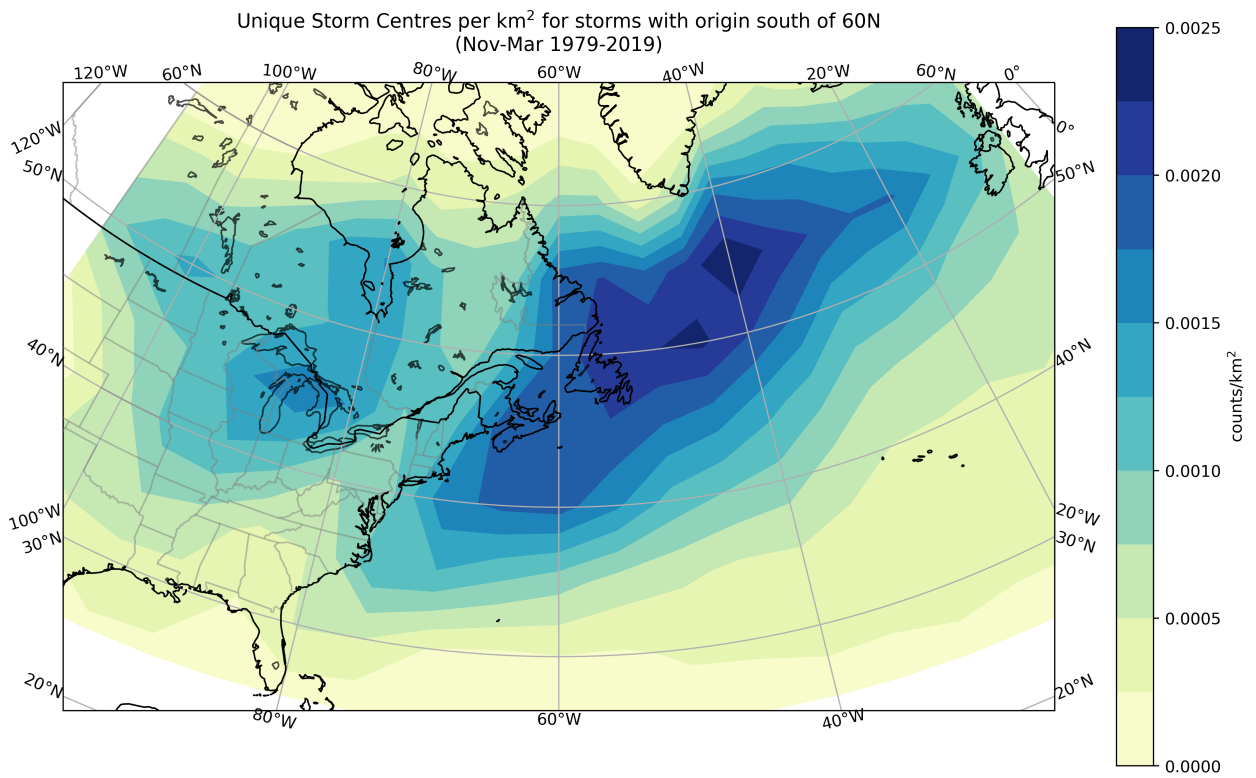


Figure A.5: Unique storm centres km⁻² for storms with origin south of 60°N on a 5 degree grid, contoured

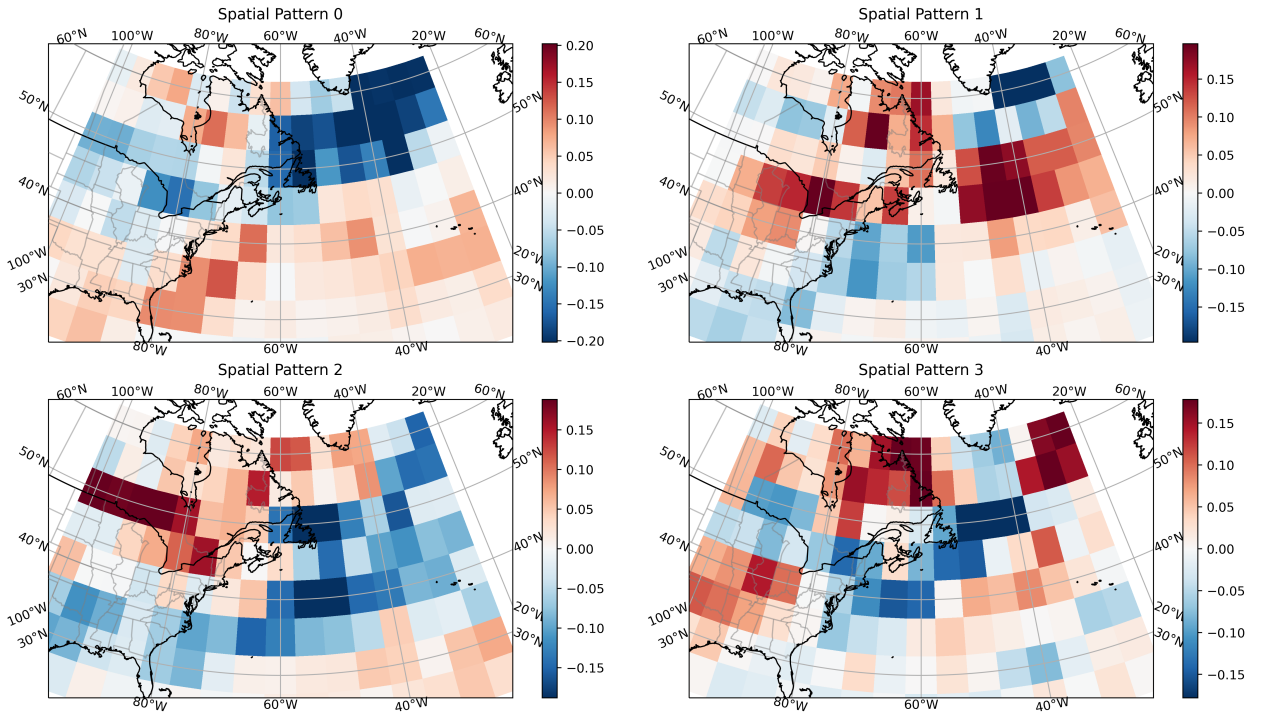


Figure A.6: Spatial patterns of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts

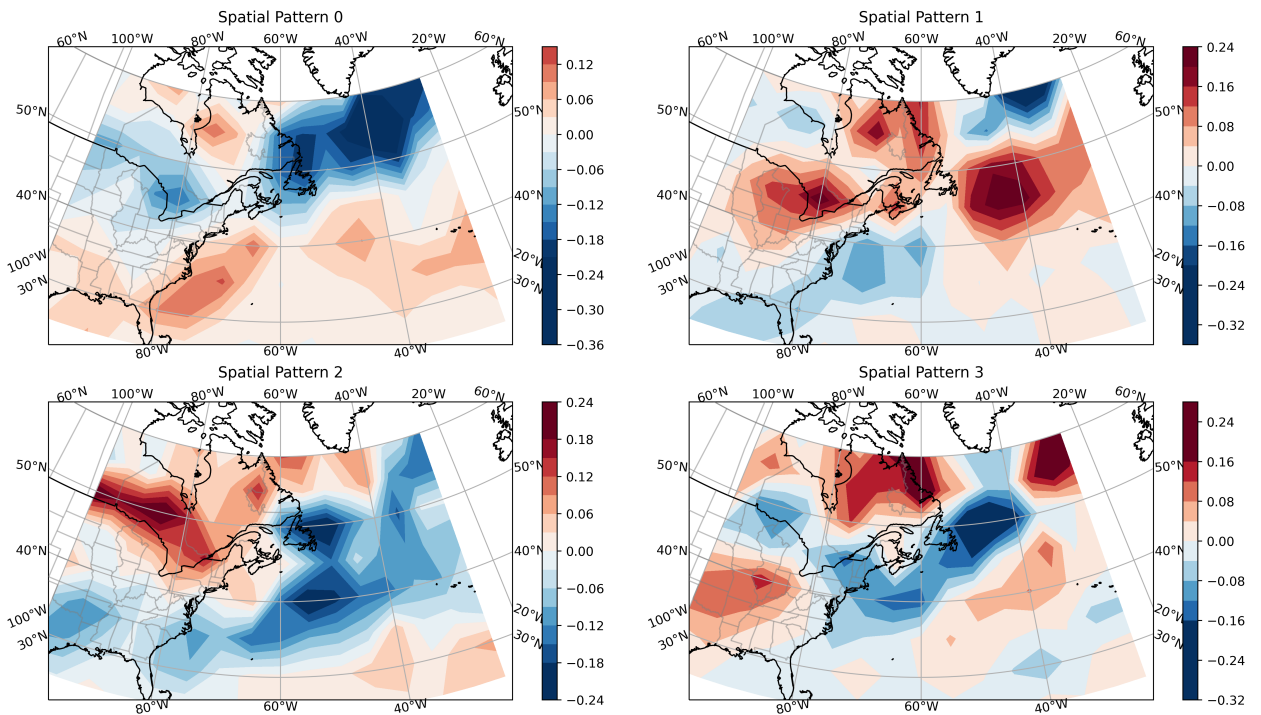


Figure A.7: Spatial patterns of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts

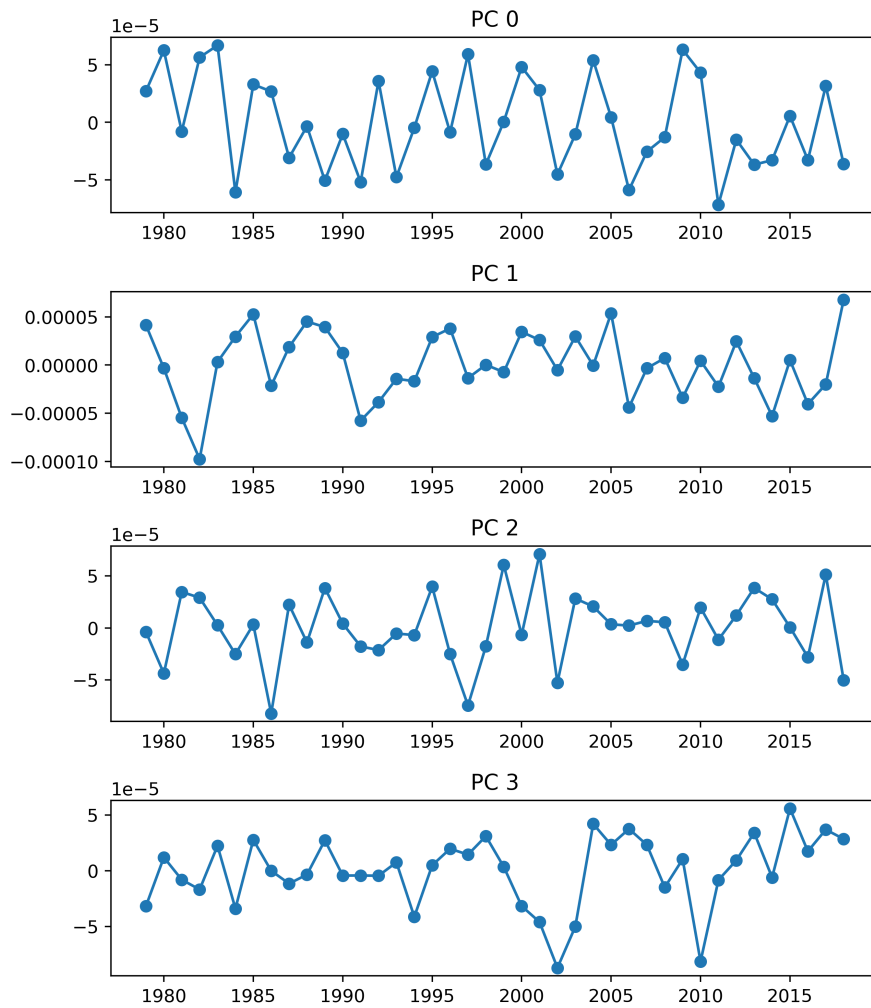


Figure A.8: Principal component timeseries of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts




Figure A.9: Principal component timeseries of the first 4 EOF modes for PCA on 5x5 degree gridded storm counts

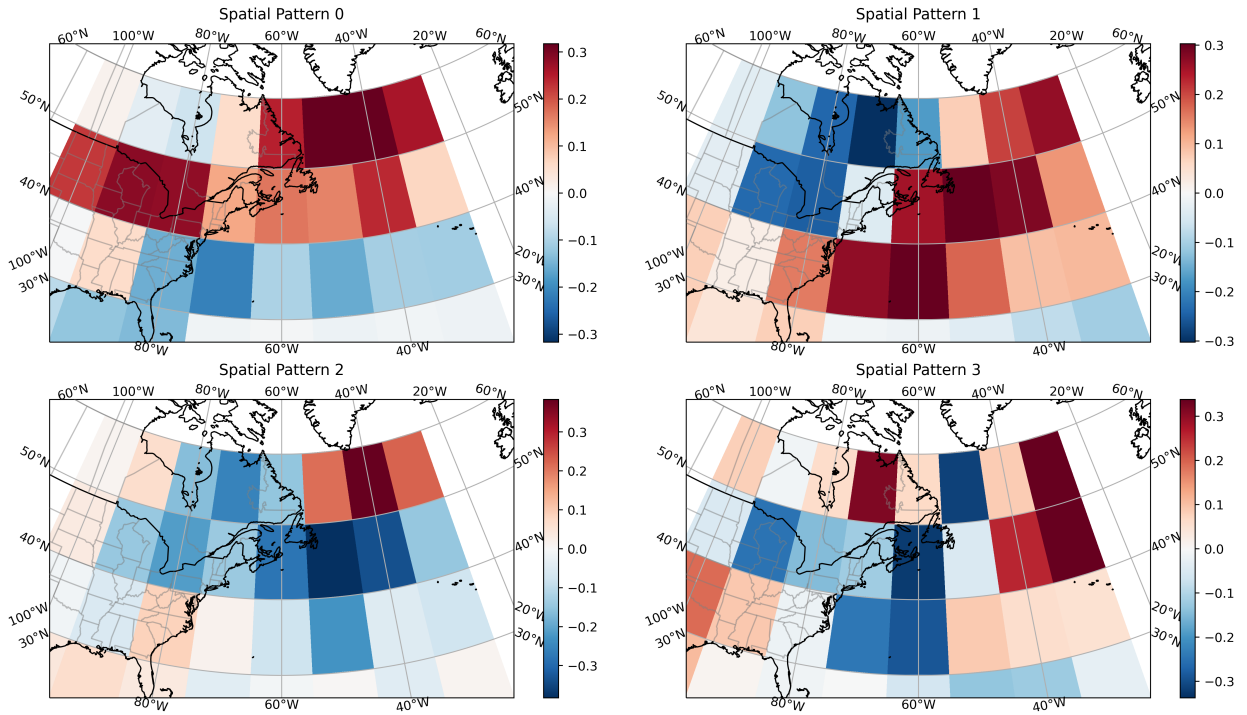


Figure A.10: Spatial patterns of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts

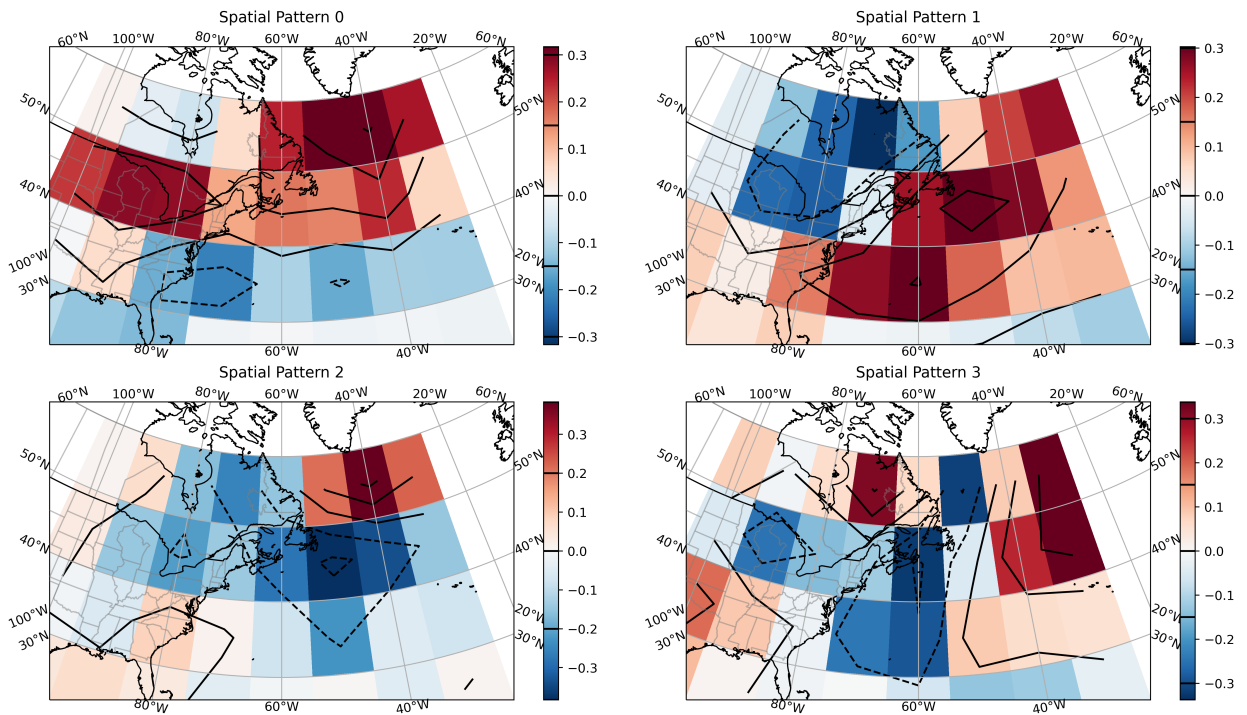


Figure A.11: Spatial patterns of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts with contours

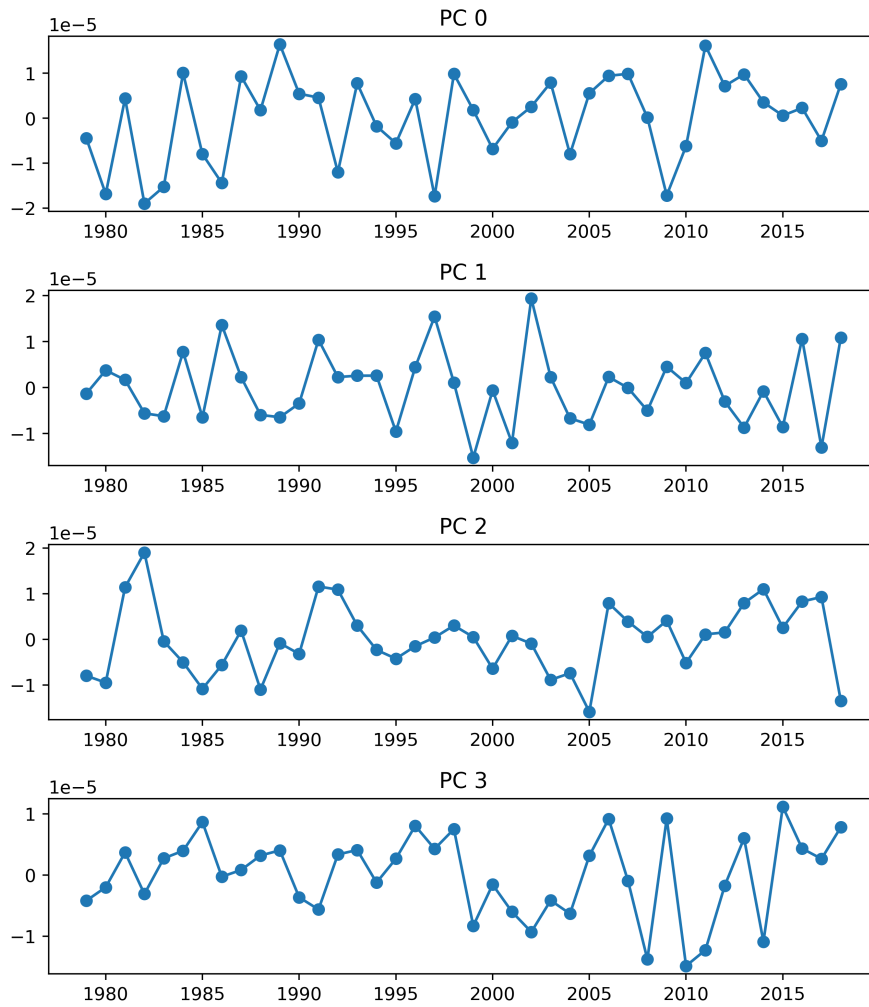


Figure A.12: Principial component timeseries of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts




Figure A.13: Principial component timeseries of the first 4 EOF modes for PCA on 10x10 degree gridded storm counts

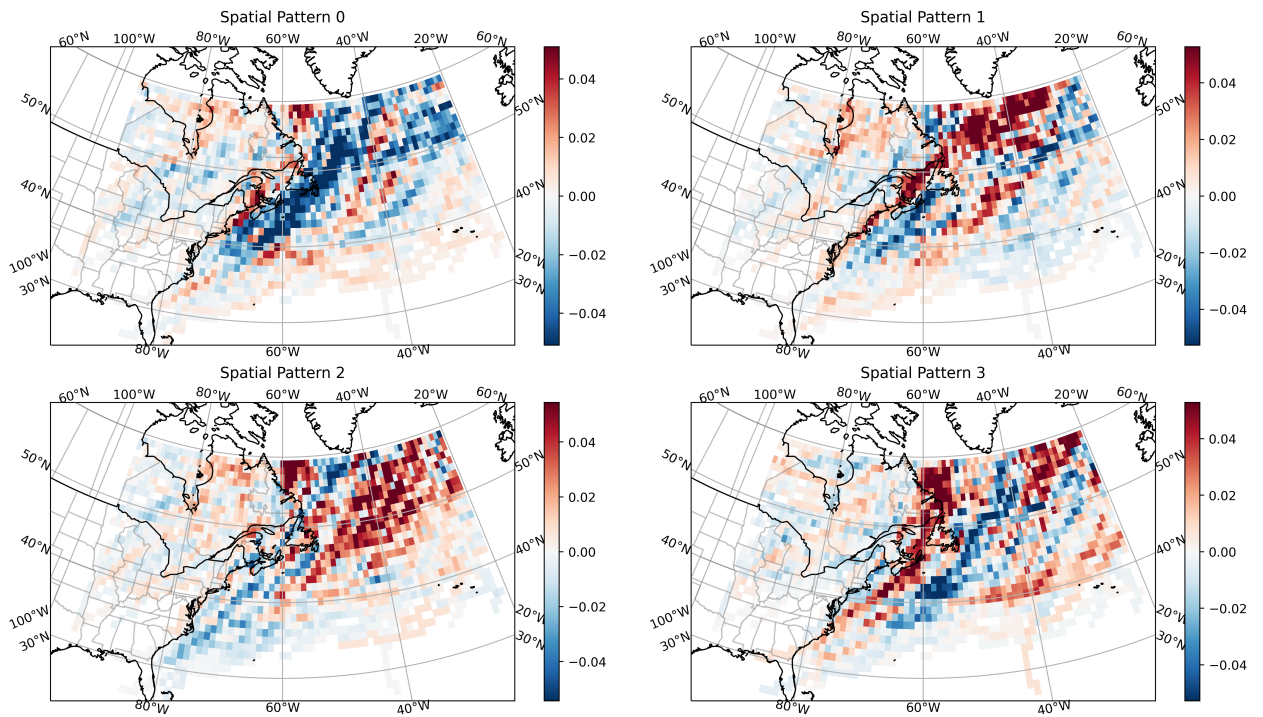


Figure A.14: Spatial patterns of the first 4 EOF modes for PCA on degree gridded bomb storm counts

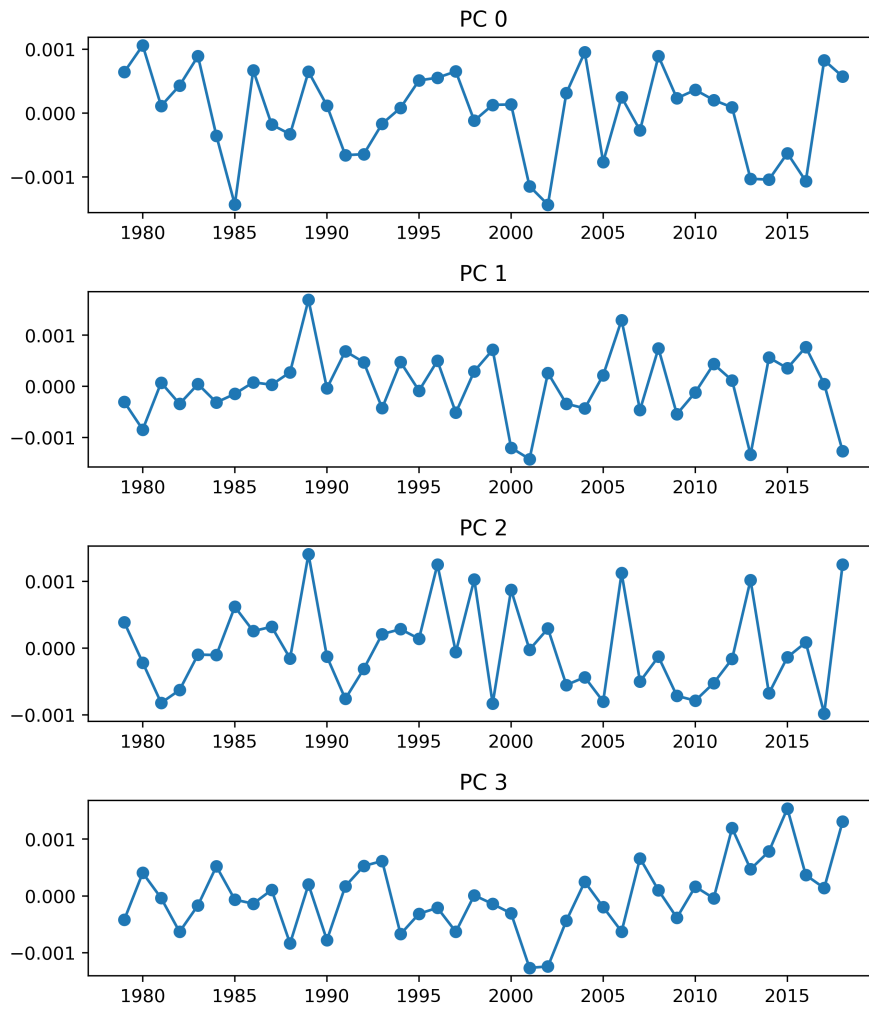


Figure A.15: Principal component timeseries of the first 4 EOF modes for PCA on degree gridded bomb storm counts




Figure A.16: Principal component timeseries of the first 4 EOF modes for PCA on gridded bomb storm counts

APPENDIX B

ADDITIONAL CONDITIONAL MEANS OF ATMOSPHERIC DRIVERS AND EFFECTS

Following the methods outlined in Section [3.4.3](#), this appendix shows supplementary analysis of conditional means using atmospheric variables that were not analyzed in this work but may be of interest to the reader.

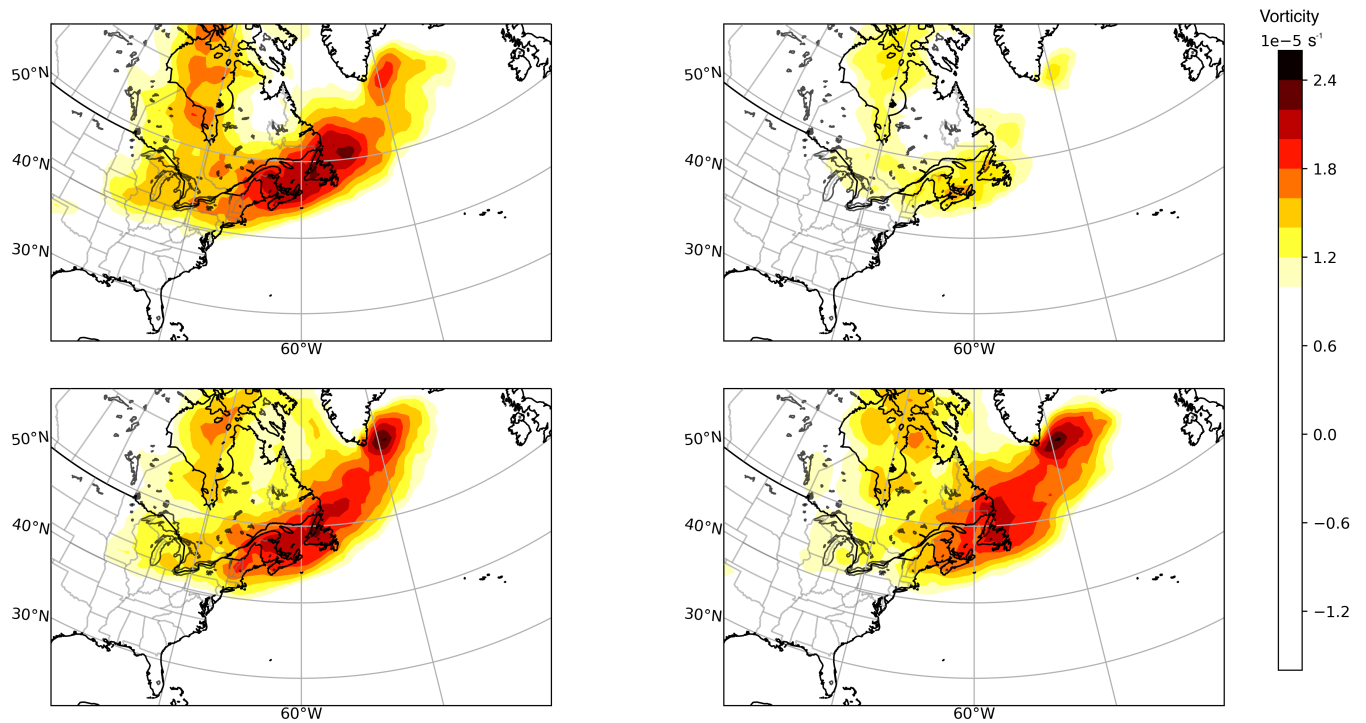


Figure B.1: Mean seasonal relative vorticity at 500 mb for SOM-grouped seasons (unit: 1/second)

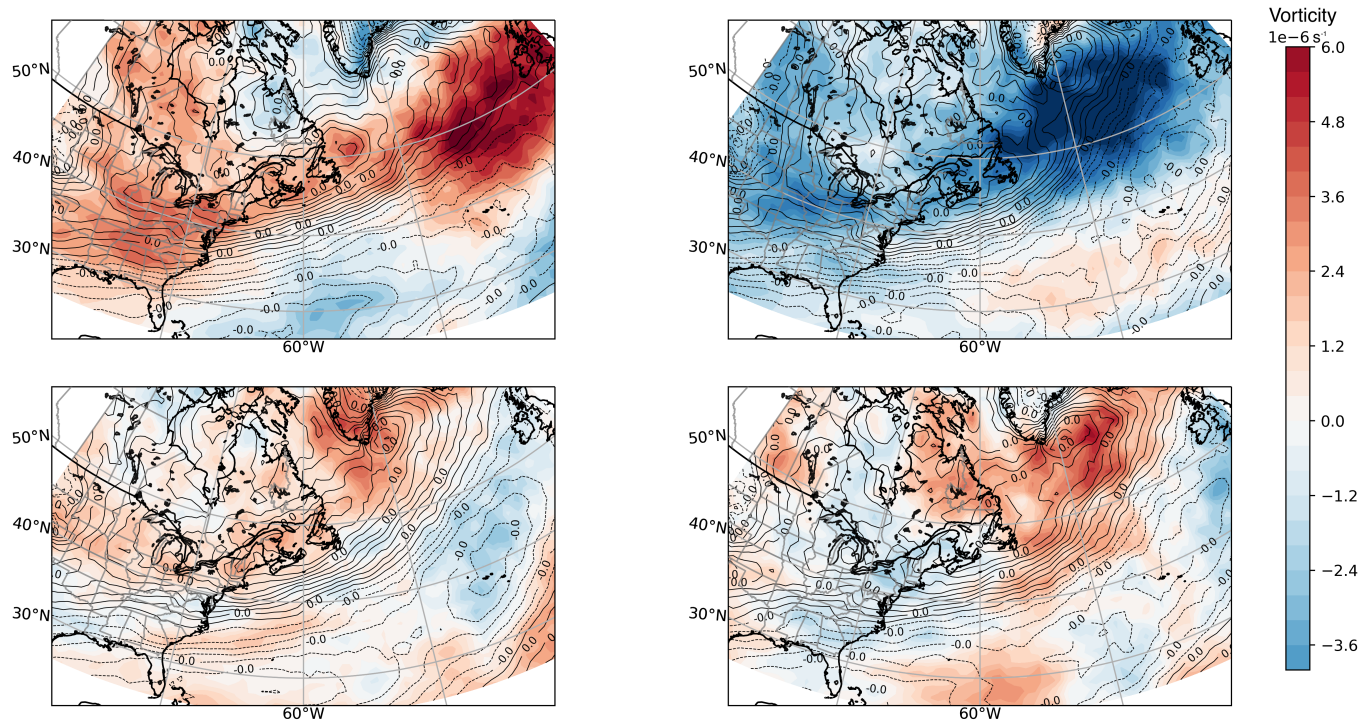


Figure B.2: Mean seasonal relative vorticity (contour) and anomaly from the climatological mean (colours) at 500 mb for SOM-grouped seasons (unit: 1/second)

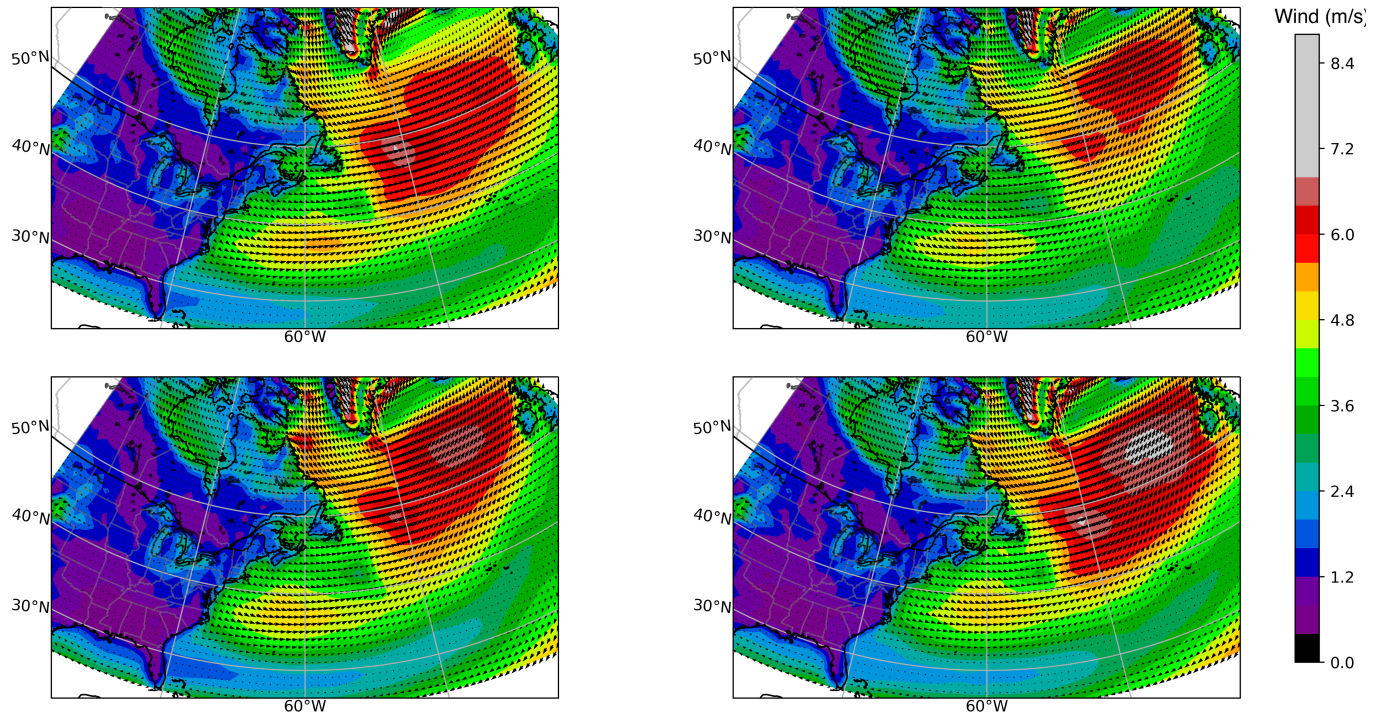


Figure B.3: Mean seasonal wind direction (arrows) and intensity (colours) at 10 m for SOM-grouped seasons (unit: meters/second)

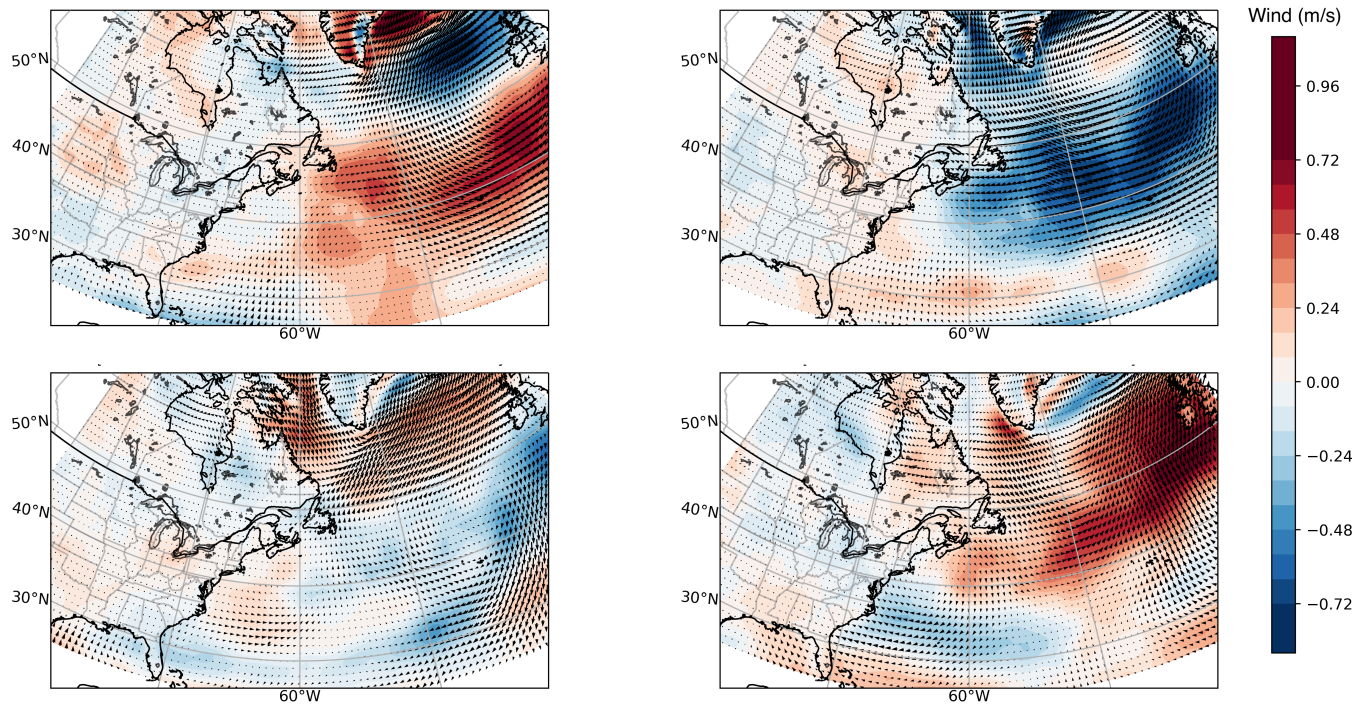


Figure B.4: Anomaly from the climatological mean of seasonal wind direction (arrows) and intensity (colours) at 250 mb for SOM-grouped seasons (unit: meters/second)

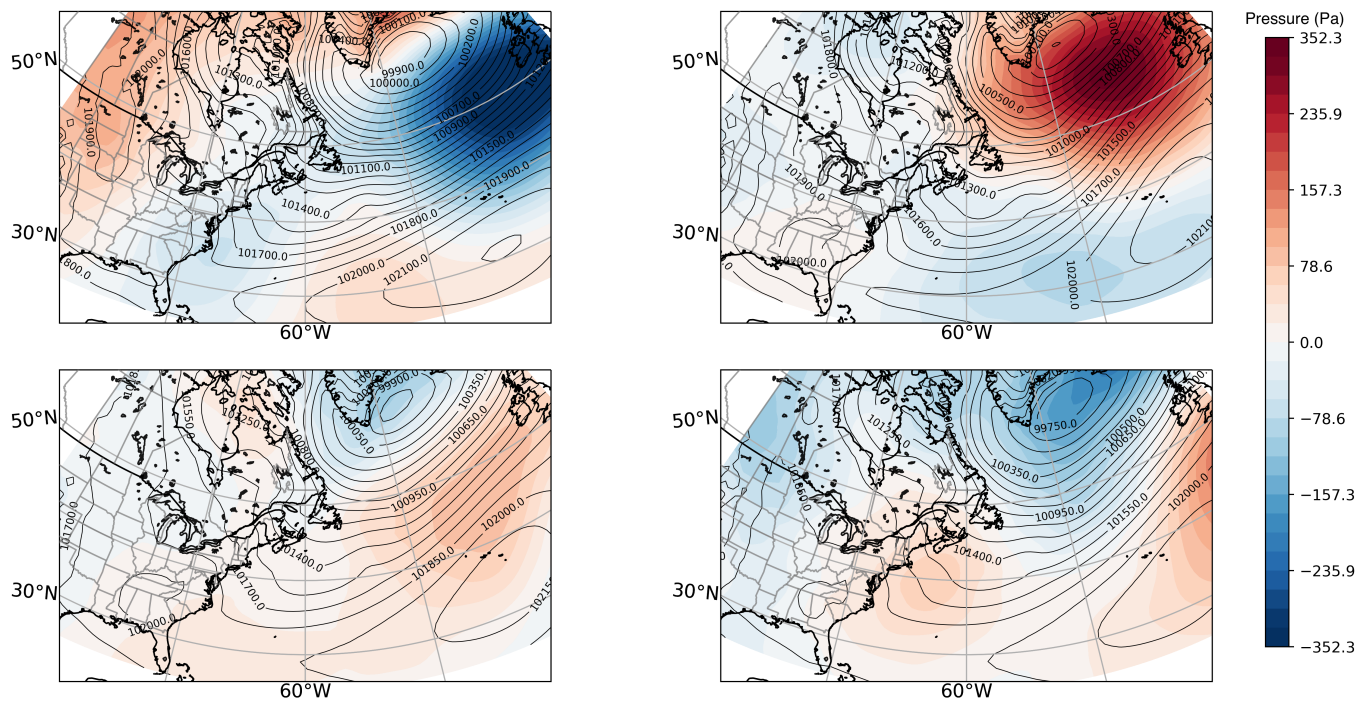


Figure B.5: Average seasonal mean sea level pressure (contours) and anomaly from the climatological mean (colours) for SOM-grouped seasons (unit: mb)

APPENDIX C

FURTHER PREDICTAND TIME SERIES ANALYSIS

The eight time series that are used as predictands can be further analyzed for their properties by means of the autocorrelation function and the cross-correlation function. The results are shown in the following figures and tables.

The ACF of rain counts shows the only significant correlation which is found at lag of one season. Despite having no significant correlations, the other time series show interesting cyclical nature, especially no precipitation and snow count which has a nearly significant signal at lag 4.

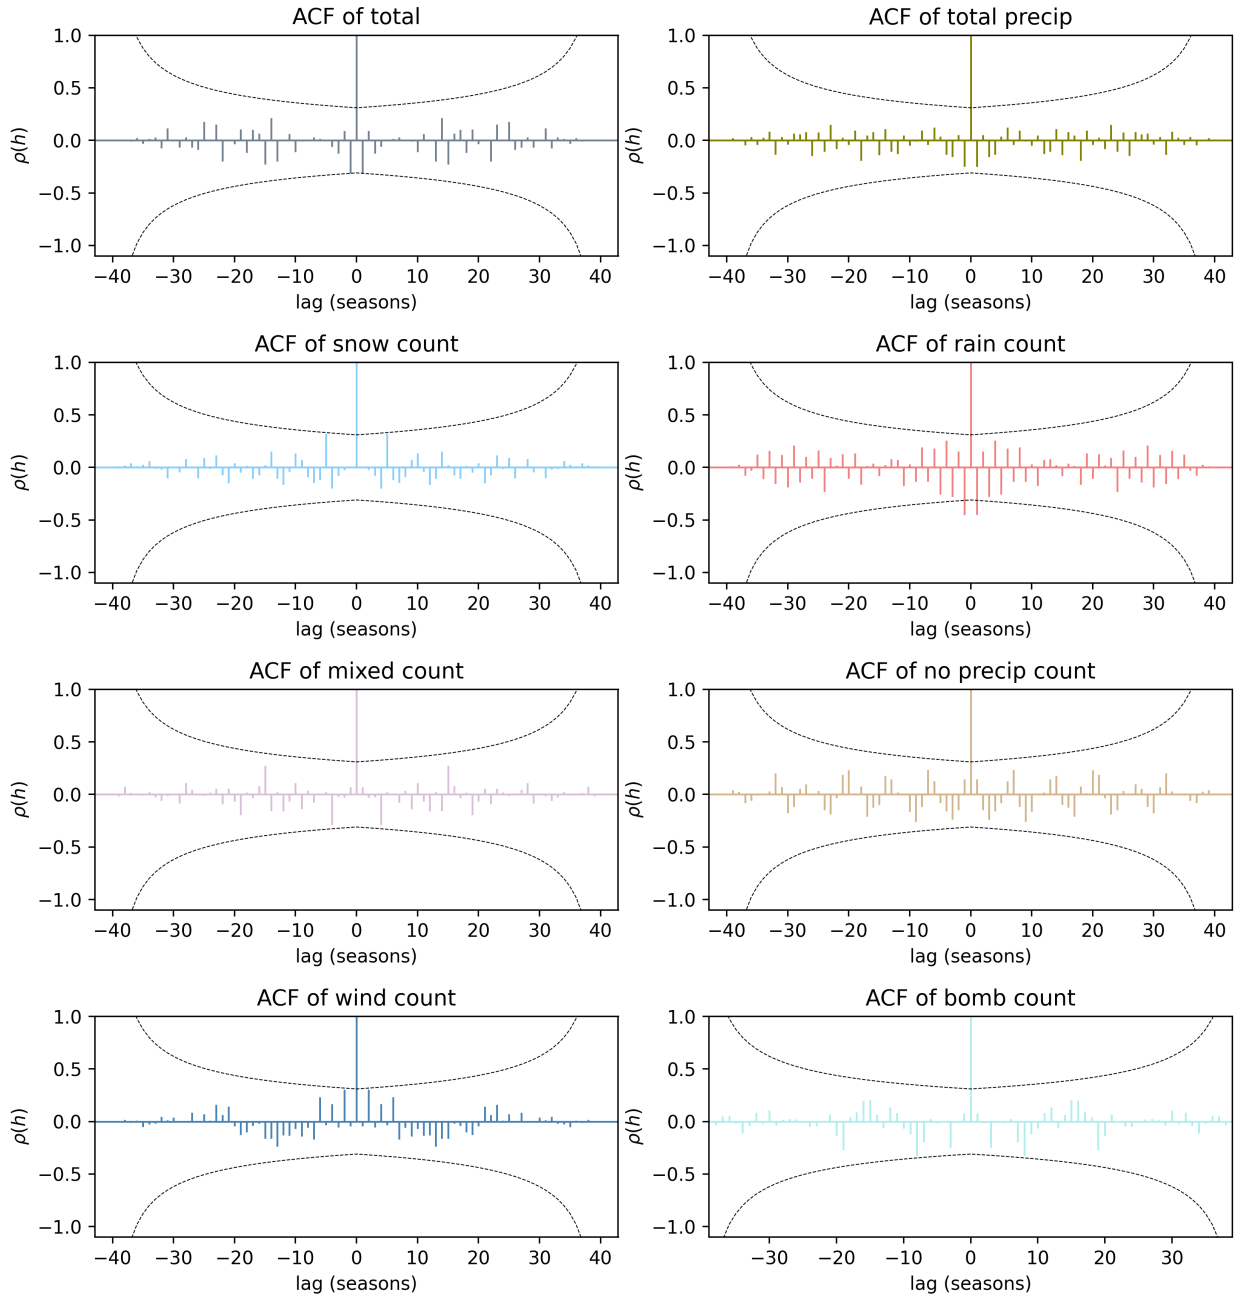


Figure C.1: Autocorrelation function of each subseries with significance shown by dashed line.

The cross-correlation function between the first 7 time series is shown in Figures [C.2](#) and [C.3](#). Note CCF was not calculated for bomb storms.

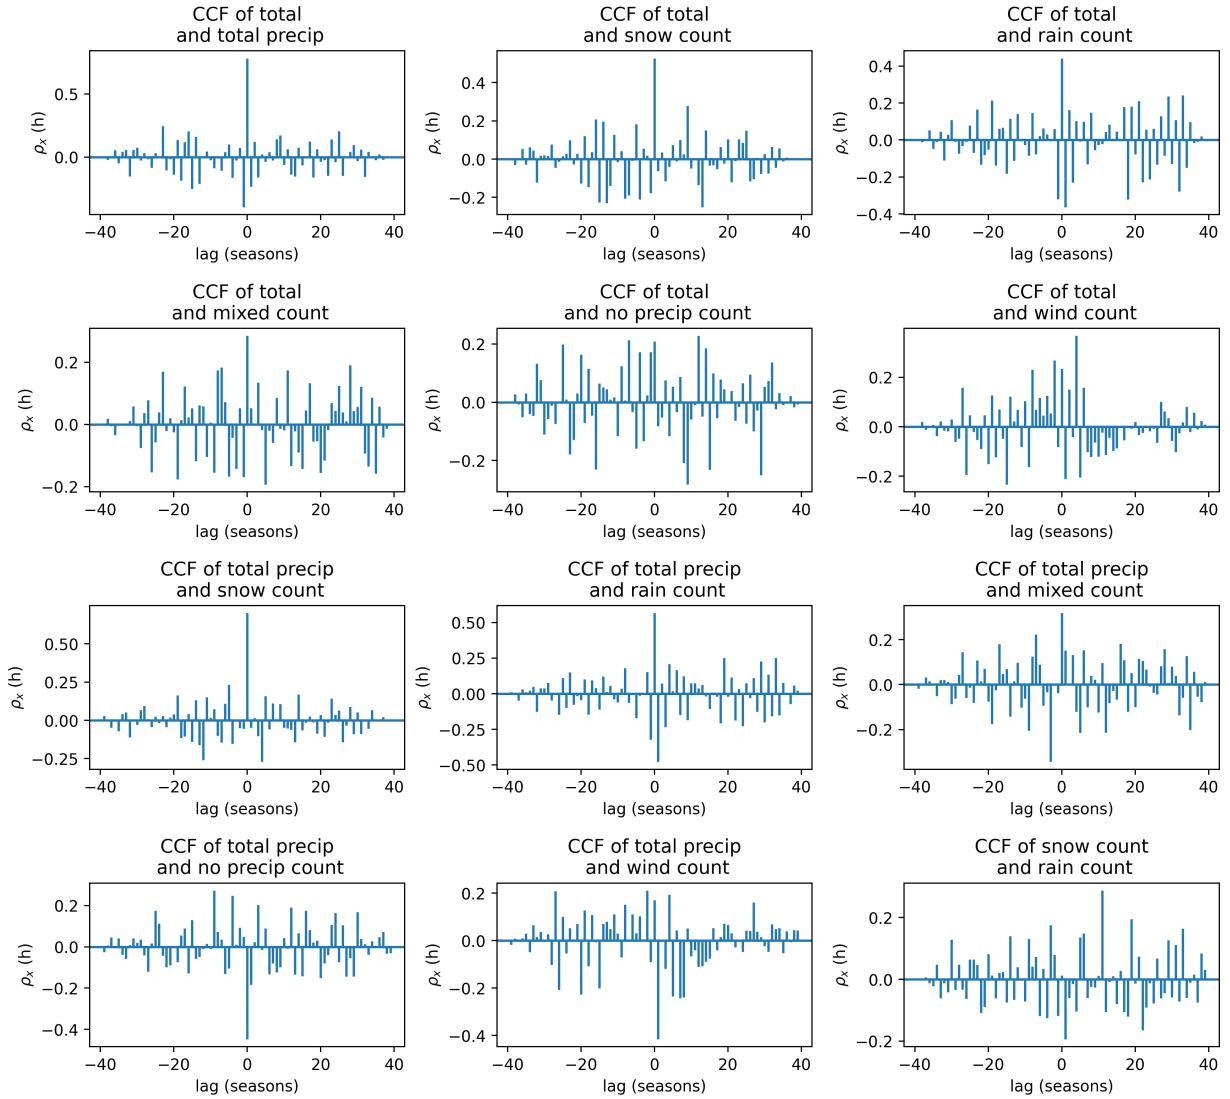


Figure C.2: Cross-correlation functions for twelve time series combinations.

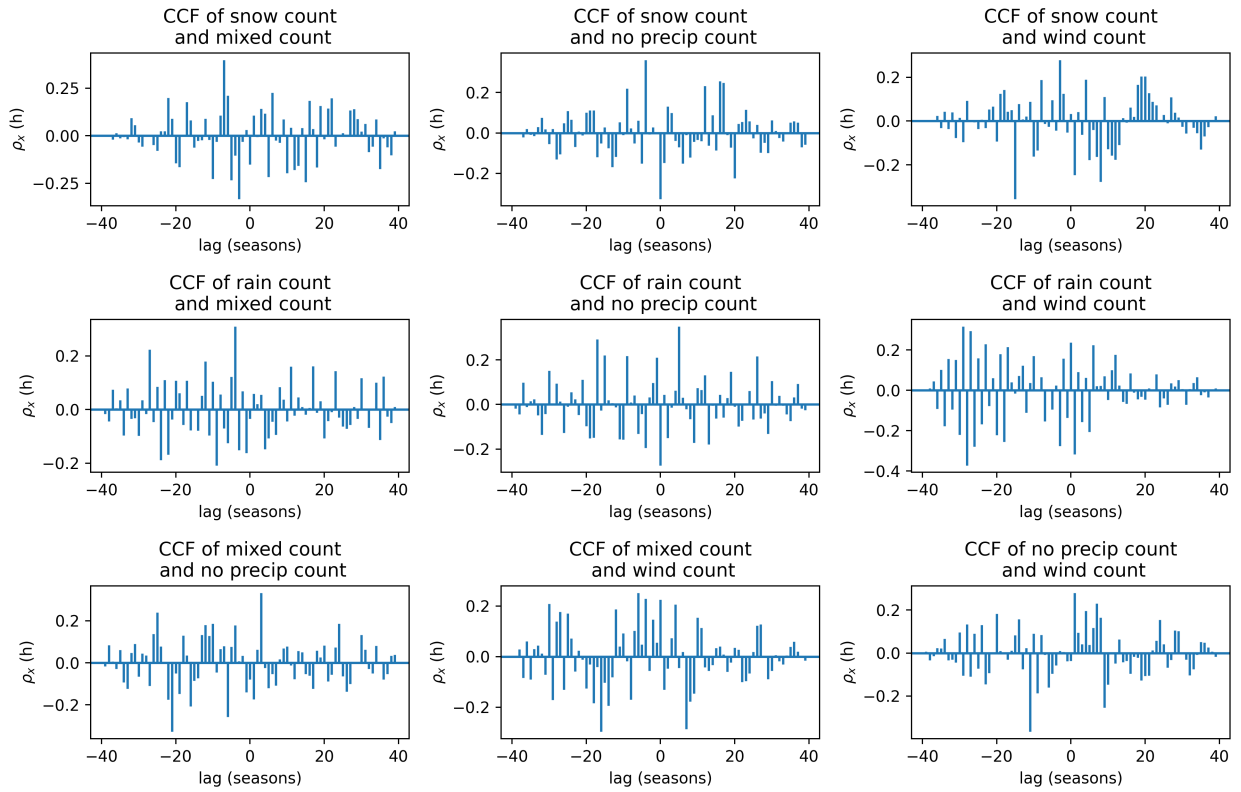


Figure C.3: Cross-correlation functions for the remaining nine time series combinations.

The zero lag correlations are summarized in table [C.1](#).

	total	total precip	snow count	rain count	mixed count	no precip count	wind count
total	1	0.79	0.54	0.44	0.29	0.21	0.26
total precip	0.79	1	0.70	0.58	0.31	-0.44	0.26
snow count	0.54	0.70	1	0.02	-0.15	-0.33	0.03
rain count	0.44	0.58	0.02	1	-0.03	-0.27	0.24
mixed count	0.29	0.31	-0.15	-0.03	1	-0.08	0.22
no precip count	0.21	-0.44	-0.33	-0.27	-0.08	1	-0.04
wind count	0.26	0.26	0.03	0.24	0.22	-0.04	1

Table C.1: Correlations of the full time series and six subseries.

The time series that have higher cross-correlation at a non-zero lag are summarized in the table below

	$\rho_x(h)$ max	lag (h) of max $\rho_x(h)$
total x no precip count	-0.3	9
total x wind count	0.44	4
total precip x mixed count	-0.34	-3
total precip x wind count	-0.43	1
snow count x rain count	0.3	11
snow count x mixed count	0.42	-7
snow count x no precip count	0.36	-4
snow count x wind count	-0.34	-15
rain count x mixed count	0.31	-4
rain count x no precip count	0.35	5
rain count x wind count	-0.37	-28
mixed count x no precip count	0.33	3
mixed count x wind count	-0.3	-16
no precip count x wind count	-0.37	-11

Table C.2: Highest lagged cross-correlations for time series pairs with max $\rho_x(h)$ at $h \neq 0$

BIBLIOGRAPHY

- American Meteorological Society, Cyclonic scale, https://glossary.ametsoc.org/wiki/Cyclonic_scale, 2020, accessed: 2021-03-21.
- Arnott, J. M., J. L. Evans, and F. Chiaromonte, Characterization of extratropical transition using cluster analysis, *Monthly weather review*, 132, 2916–2937, 2004.
- Befort, D. J., S. Wild, J. R. Knight, J. F. Lockwood, H. E. Thornton, L. Hermanson, P. E. Bett, A. Weisheimer, and G. C. Leckebusch, Seasonal forecast skill for extratropical cyclones and windstorms, *Quarterly Journal of the Royal Meteorological Society*, p. qj.3406, 2018.
- Blackmon, M. L., A climatological spectral study of the 500 mb geopotential height of the northern hemisphere, *Journal of the Atmospheric Sciences*, 33, 1607–1623, 1976.
- Blender, R., K. Fraedrich, and F. Lunkeit, Identification of cyclone-track regimes in the North Atlantic, 1997.
- Catto, J. L., Extratropical cyclone classification and its use in climate studies, *Reviews of Geophysics*, 54, 486–520, 2016.
- Chang, E. K. M., and Y. Fu, Interdecadal variations in Northern Hemisphere winter storm track intensity, *Journal of Climate*, 15, 642–658, 2002.
- Chelton, D. B., M. G. Schlax, and R. M. Samelson, Global observations of nonlinear mesoscale eddies, *Progress in Oceanography*, 91, 167–216, 2011.
- DeGaetano, A. T., M. E. Hirsch, and S. J. Colucci, Statistical prediction of seasonal East Coast winter storm frequency, *Journal of Climate*, 15, 1101–1117, 2002.
- Deveson, A. C., K. A. Browning, and T. D. Hewson, A classification of FASTEX cyclones using a height-attributable quasi-geostrophic vertical-motion diagnostic, *Quarterly Journal of the Royal Meteorological Society*, 128, 93–117, 2002.
- Edmon, H. J., B. Hoskins, and M. McIntyre, Eliassen-Palm Cross Sections for the Troposphere, *Journal of the Atmospheric Sciences*, 37, 2600–2616, 1980.
- Feng, X., B. Huang, and D. M. Straus, Seasonal prediction skill and predictability of the Northern Hemisphere storm track variability in Project Minerva, *Climate Dynamics*, 52, 6427–6440, 2019.
- Fogarty, C. T., The great maritimes blizzard of february 18-19, 2004, *Canadian Meteor. and Oceanogr. Soc. Bul.*, 32, 100–106, 2004.
- Frankoski, N. J., and A. T. Degaetano, An East Coast winter storm precipitation climatology, *International Journal of Climatology*, 31, 802–814, 2011.

- Gaffney, S. J., A. W. Robertson, P. Smyth, S. J. Camargo, and M. Ghil, Probabilistic clustering of extratropical cyclones using regression mixture models, *Climate dynamics*, 29, 423–440, 2007.
- Gore, M., C. M. Zarzycki, and M. Gervais, Tying large-scale meteorological patterns to northeastern U.S. extratropical cyclones with climate data and self-organizing maps, in *American Geophysical Union, Fall Meeting 2020, abstract A147-0021*, vol. 2020, pp. A147–0021, 2020.
- Gray, W. M., Atlantic seasonal hurricane frequency. part ii: Forecasting its variability, *Monthly Weather Review*, 112, 1669–1683, 1984.
- Grise, K. M., S. W. Son, and J. R. Gyakum, Intraseasonal and interannual variability in north american storm tracks and its relationship to equatorial pacific variability, *Monthly Weather Review*, 141, 3610–3625, 2013.
- Gulev, S. K., O. Zolina, and S. Grigoriev, Extratropical cyclone variability in the Northern Hemisphere winter from the NCEP/NCAR reanalysis data, *Climate Dynamics*, 17, 795–809, 2001.
- Heckert, N. A., J. J. Filliben, C. M. Croarkin, B. Hembree, W. F. Guthrie, P. Tobias, J. Prinz, et al., Handbook 151: Nist/sematech e-handbook of statistical methods, 2002.
- Helwig, N. E., Multiple Linear Regression, <http://users.stat.umn.edu/~helwig/notes/mlr-Notes.pdf>, 2017.
- Hersbach, H., B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, and J.-N. Thépaut, Era5 hourly data on pressure levels from 1979 to present., Copernicus Climate Change Service (C3S) Climate Data Store (CDS), 2018a, accessed: 2020-04-01.
- Hersbach, H., B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, and J.-N. Thépaut, Era5 hourly data on single levels from 1979 to present., Copernicus Climate Change Service (C3S) Climate Data Store (CDS), 2018b, accessed: 2020-04-01.
- Hewitson, B. C., and R. G. Crane, Self-organizing maps: Applications to synoptic climatology, *Climate Research*, 22, 13–26, 2002.
- Hodges, K. I., B. J. Hoskins, J. Boyle, and C. Thorncroft, A comparison of recent reanalysis datasets using objective feature tracking: Storm tracks and tropical easterly waves, *Monthly Weather Review*, 131, 2012–2037, 2003.
- Hoskins, B., and K. Hodges, The annual cycle of northern hemisphere storm tracks. part i: Seasons, *Journal of Climate*, 32, 1743–1760, 2019.
- Hoskins, B. J., and K. I. Hodges, New perspectives on the Northern Hemisphere winter storm tracks, *Journal of the Atmospheric Sciences*, 59, 1041–1061, 2002.

- Hoskins, B. J., and K. I. Hodges, A new perspective on southern hemisphere storm tracks, *Journal of Climate*, 18, 4108–4129, 2005.
- Insurance Information Institute, ., Facts + Statistics: Winter storms, <https://www.iii.org/fact-statistic/facts-statistics-winter-storms>, 2021.
- IPCC, *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)], IPCC, Geneva, Switzerland, 151 pp, 2014.
- Kenward, A., and U. Raja, BLACKOUT: EXTREME WEATHER, CLIMATE CHANGE AND POWER OUTAGES, *Tech. rep.*, Climate Central, 2014.
- Klotzbach, P., L.-P. Caron, and M. Bell, A statistical/dynamical model for north atlantic seasonal hurricane prediction, *Geophysical Research Letters*, 47, e2020GL089357, 2020.
- Klotzbach, P. J., and W. M. Gray, Twenty-five years of atlantic basin seasonal hurricane forecasts (1984–2008), *Geophysical research letters*, 36, 2009.
- Kohonen, T., The Self-Organizing Map, *Proceedings of the IEEE*, 78, 1464–1480, 1990.
- Lakkis, S. G., P. Canziani, A. Yuchechen, L. Rocamora, A. Caferra, K. Hodges, and A. O’Neill, A 4d feature-tracking algorithm: A multidimensional view of cyclone systems, *Quarterly Journal of the Royal Meteorological Society*, 145, 395–417, 2019.
- Lawson, B. D., Trends in blizzards at selected locations on the canadian prairies, *Natural Hazards*, 29, 123–138, 2003.
- Liu, Y., and H. Robert, A Review of Self-Organizing Map Applications in Meteorology and Oceanography, in *Self Organizing Maps - Applications and Novel Algorithm Design*, InTech, 2011.
- Lorenz, E. N., Available Potential Energy and the Maintenance of the General Circulation, *Tellus*, 7, 157–167, 1955.
- Massey, N., Feature tracking on the hierarchical equal area triangular mesh, *Computers & Geosciences*, 44, 42–51, 2012.
- Massey, N. R., Feature tracking in high-resolution regional climate data, *Computers & Geosciences*, 93, 36–44, 2016.
- Monahan, A. H., J. C. Fyfe, M. H. Ambaum, D. B. Stephenson, and G. R. North, Empirical orthogonal functions: The medium is the message, *Journal of Climate*, 22, 6501–6514, 2009.

- Neu, U., M. G. Akperov, N. Bellenbaum, R. Benestad, R. Blender, R. Caballero, A. Cozza, H. F. Dacre, Y. Feng, K. Fraedrich, et al., Imilast: A community effort to intercompare extratropical cyclone detection and tracking algorithms, *Bulletin of the American Meteorological Society*, *94*, 529–547, 2013.
- NOAA National Severe Storms Laboratory, Severe Weather 101: Winter Weather Basics, <https://www.nssl.noaa.gov/education/svrwx101/winter/>, 2021.
- Oliver, E. C., V. Lago, A. J. Hobday, N. J. Holbrook, S. D. Ling, and C. N. Mundy, Marine heatwaves off eastern Tasmania: Trends, interannual variability, and predictability, *Progress in Oceanography*, *161*, 116–130, 2018.
- Oliver, E. C. J., T. J. O’kane, and N. J. Holbrook, Projected changes to Tasman Sea eddies in a future climate, 2015.
- Orlanski, I., and J. Katzfey, The Life Cycle of a Cyclone Wave in the Southern Hemisphere. Part I: Eddy Energy Budget, *Journal of the Atmospheric Sciences*, *48*, 1972–1998, 1991.
- Parding, K. M., R. Benestad, A. Mezghani, and H. B. Erlandsen, Statistical projection of the North Atlantic storm tracks, *Journal of Applied Meteorology and Climatology*, *58*, 1509–1522, 2019.
- Pidwirny, M., Global heat balance: Introduction to heat fluxes, *Fundamentals of Physical Geography*, 2nd Edition, 2006.
- Pinto, J. G., S. Ulbrich, T. Economou, D. B. Stephenson, M. K. Karremann, and L. C. Shaffrey, Robustness of serial clustering of extratropical cyclones to the choice of tracking method, *Tellus A: Dynamic Meteorology and Oceanography*, *68*, 32204, 2016.
- Plante, M., S. W. Son, E. Atallah, J. Gyakum, and K. Grise, Extratropical cyclone climatology across eastern Canada, *International Journal of Climatology*, *35*, 2759–2776, 2015.
- Post, F. H., B. Vrolijk, H. Hauser, R. S. Laramée, and H. Doleisch, The State of the Art in Flow Visualisation: Feature Extraction and Tracking, *Computer Graphics Forum*, *22*, 775–792, 2003.
- Raible, C., P. Della-Marta, C. Schierz, H. Wernli, and R. Blender, Northern hemisphere extratropical cyclones: A comparison of detection and tracking methods and different reanalyses, *Monthly Weather Review*, *136*, 880–897, 2008.
- Renggli, D., G. C. Leckebusch, U. Ulbrich, S. N. Gleixner, and E. Faust, The Skill of Seasonal Ensemble Prediction Systems to Forecast Wintertime Windstorm Frequency over the North Atlantic and Europe, *Monthly Weather Review*, *139*, 3052–3068, 2011.
- Rogers, J. C., North Atlantic Storm Track Variability and Its Association to the North Atlantic Oscillation and Climate Variability of Northern Europe, *Journal of Climate*, *10*, 1635–1647, 1997.

- Sanders, F., and J. R. Gyakum, Synoptic-dynamic climatology of the “bomb”, *Monthly Weather Review*, 108, 1589–1606, 1980.
- Talley, L. D., Shallow, intermediate, and deep overturning components of the global heat budget, *Journal of Physical Oceanography*, 33, 530–560, 2003.
- Uccellini, L. W., Processes Contributing to the Rapid Development of Extratropical Cyclones, pp. 81–105, 1990.
- Varino, F., P. Arbogast, B. Joly, G. Riviere, M. L. Fandeur, H. Bovy, and J. B. Granier, Northern Hemisphere extratropical winter cyclones variability over the 20th century derived from ERA-20C reanalysis, *Climate Dynamics*, 52, 1027–1048, 2019.
- Vettigli, G., Minisom: minimalistic and numpy-based implementation of the self organizing map, 2018.
- Walker, E., D. Mitchell, and W. Seviour, The numerous approaches to tracking extratropical cyclones and the challenges they present, *Weather*, 75, 336–341, 2020.
- Wandeto, J. M., and B. Dresch-Langley, The quantization error in a Self-Organizing Map as a contrast and colour specific indicator of single-pixel change in large random patterns, *Neural Networks*, 119, 273–285, 2019.
- Wang, C.-C., and J. C. Rogers, A Composite Study of Explosive Cyclogenesis in Different Sectors of the North Atlantic. Part I: Cyclone Structure and Evolution, *Monthly Weather Review*, 129, 1481–1499, 2001.
- Yang, X., G. A. Vecchi, R. G. Gudgel, T. L. Delworth, S. Zhang, A. Rosati, L. Jia, W. F. Stern, A. T. Wittenberg, S. Kapnick, R. Msadek, S. D. Underwood, F. Zeng, W. Anderson, and V. Balaji, Seasonal Predictability of Extratropical Storm Tracks in GFDL’s High-Resolution Climate Prediction Model, *Journal of Climate*, 28, 3592–3611, 2015.