Numerical and Statistical Modelling of Australian Severe Weather

Joshua Hartigan
Bachelor of Mathematics (Honours)

Thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy
Under the supervision of: James Brown, Lance M. Leslie, and Shev MacNamara

School of Mathematical and Physical Sciences,
Faculty of Science,
University of Technology Sydney

Submitted October 2022

CERTIFICATE OF ORIGINAL AUTHORSHIP

I, Joshua Hartigan declare that this thesis, is submitted in fulfilment of the requirements for the award of Doctor of Philosophy, in the School of Mathematical and Physical Sciences at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise referenced or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

This document has not been submitted for qualifications at any other academic institution.

This research is supported by the Australian Government Research Training Program.

Production Note:
Signature: Signature removed prior to publication.

Date: 30/05/2022

Acknowledgements

This research represents a challenging but rewarding four years where I have had the privilege to apply my background in mathematics to better understand various forms of severe weather, including thunderstorms, for which I have been passionate about for the majority of my life. There are numerous people who have helped make this dream a reality and provided valuable support along the way.

Firstly, I would like to thank my first summer research supervisor, Dr Judy-anne Osborn, who encouraged me to pursue my passion in severe weather and helped set up my honours research project in modelling how severe convection interacts with the sea-breeze air mass. I would also like to thank my honours supervisors, Joshua Soderholm and Robert Warren for taking me on despite the long distance, and supporting me throughout that project. I thank my piano teacher, Luba Totoeva, for her connection with Em/Prof. Alex Novikov, which led me to my PhD supervisors Prof. Lance Leslie and Dr Shev MacNamara. I also thank Luba for the countless hours teaching me piano, providing me a much needed break and emotional outlet throughout my candidature.

To my supervisors: thank you for all of the support you have provided throughout my candidature, and for providing the room necessary for me to grow as a researcher, and explore many of my interests in severe weather. Lance, you provided valuable intellectual stimulation and so much deep knowledge from your career in atmospheric and climate science. Shev, you were a constant source of positivity and interesting insights from a mathematician's point of view. I also thank Prof. James Brown for coming aboard my supervisory panel and providing much needed administrative support.

I would also like to thank my colleagues both at UTS and outside of UTS. The School of Mathematical and Physical Sciences has provided valuable financial and administrative support throughout my candidature. My appreciation goes out especially to the school manager, Elizabeth Gurung Tamang, for always making the time for me to talk and sort out different forms, and providing me with opportunities to work with other areas of the school, the Faculty of Science and the Finance team at UTS. Alex Bishop and Mohammad Mahdi Ahmadian, who completed their PhDs alongside me, provided fun and refreshing conversations every Wednesday that I was on campus, prior to the COVID-19 outbreak. Some of my PhD work required significant amounts of computing power (particularly for the numerical simulations and development of the SCT climatology), to which I am grateful to Mike Lake, the UTS scheme manager of the NCI supercomputing facility, for providing sufficient computation time on the supercomputer and always being available to help with any problems I had running code on it. Further thanks go to Peter Chan and Bruce Buckley at IAG for arranging insurance data for my research project, which has helped with choosing case studies and provided a useful form of verification of my numerical simulations. I would like to thank Ron Holle, John Cramer and Ryan Said of Vaisala Inc. for assisting in my request for GLD360 lightning data, and Joshua Soderholm and Alain Protat for providing weather radar data — both data sets were vital in the development of the severe convective thunderstorm hazards climatology. The weather radar data provided an additional form of verification of my numerical simulations. I was also fortunate to begin collaborating with Dr Milton Speer during my candidature, who has provided a number of additional research avenues to understand Australian severe weather, and whose expertise and meteorological background I truly admire.

I want to thank my friends for supporting me and understanding when I would miss social events so I could work on my project. I would like to acknowledge and thank my parents, Mer-

ren and Peter, for providing love and support, and always encouraging me to follow my dreams. Without both of you, getting to where I am now would have been unachievable. Thanks also goes to my sister, Brea, for always being there for me. Finally I want to thank my partner, Matthew, for being a constant source of love and support, and for keeping me grounded. I hope to provide the same level of support and encouragement as you complete your PhD.

Abstract

Various forms of severe weather frequently impact communities around Australia, including drought, floods, bushfires, severe convective thunderstorms and tropical cyclones. Such events have varied temporal scales and can have a significant effect on society, especially when multiple events occur shortly after each other or at the same time. As the global warming trend continues, the threat of severe weather impacting regions likely will change. Along with using global circulation models and regional climate models to understand how future warming will affect the risk of severe weather, it is important to study severe weather historically, and trends in severe weather and related variables, as this highlights the performance of modelling efforts, and what is currently occurring. Additionally, statistical modelling can be used to assess the attributes of severe weather, at the time scales important to forecasting and planning for these events. However, non-linear statistical methods including machine learning have seldom been used in the atmospheric and climate sciences to help understand severe weather, especially in Australia. Meanwhile, numerical modelling can be applied to past events, such as severe convection, to help understand how the event developed and improve future forecasting of similar events. Yet, there are few numerical modelling case studies of severe convective storms in Australia, particularly for tornadic storms, with much of the knowledge Australian forecasters rely on coming from studies and understandings developed overseas. This thesis applies trend analysis and statistical modelling to improve understanding of drought across southeast Australia, and severe convective storm hazard environments across Australia. Numerical modelling also is applied to a case study of a tornadic storm in Australia's most populous city, Sydney.

First, trends in drought are understood through analysis of change to percentiles, bootstrap resampling of periods within the time series, and permutation testing to understand statistical significance. This has been applied to key areas in southeast Australia, including the Hunter Valley Region, Sydney Catchment Area (SCA), Canberra city, and the Northern Murray-Darling Basin. It was found that, overall, there are concerning changes to precipitation mean and variance particularly over the cool season, suggesting a reduction in streamflow across the region. Coupled with increasing mean temperatures for most seasons across southeast Australia causing an increase in potential evapotranspiration, there appears to be a trend towards more frequent and severe drought. Wavelet analysis was applied to assess potential climate drivers within the time series, and how the relationships with these climate drivers have changed over time. Finally, statistical modelling was applied to further understand the key climate drivers associated with precipitation in the region, and assess how well different modelling techniques perform on prediction over the climate-scale. These models generally performed better than the climatology, and highlighted the influence of global warming on precipitation with Tasman Sea Surface Temperature Anomalies a frequently selected attribute. This is particularly apparent in the SCA, where models performed particularly well on predicting annual precipitation, despite trends in the time series.

On 16 December 2015, a severe thunderstorm and associated tornado affected Sydney, causing widespread damage and insured losses of $206 million. This storm severely impacted the suburb of Kurnell, with significant damage sustained to Sydney's desalination plant, which supplies up to 15% of Sydney water during drought. High resolution numerical simulations of this storm were conducted using the Weather Research and Forecasting (WRF) model on a double nested domain using ECMWF's ERA5 reanalysis data for initial conditions. The results from simulations with the Morrison microphysics scheme are compared to those from the National Severe Storms Laboratory double-moment 4-ice microphysics scheme. Both simulations produced severe convective storms that followed similar paths to the observed storm, but the Morrison scheme did not display the same morphology. On the other hand, the storm simulated with

the NSSL scheme displayed cyclical low- and mid-level mesocyclone development, which was observed in the Kurnell storm. This highlights how supportive the atmosphere was for development of severe rotating thunderstorms with the potential to produce tornadoes. The work presented in this thesis displays the applicability of the WRF model to studying the causes of high-impact Australian thunderstorms.

In the final part of this thesis, statistical modelling is applied to better understand the atmospheric variables most closely associated with severe convective storm hazards (i.e., hail over 2 cm, wind gusts above 90 km/hr or tornadoes). The best performing models were then applied to the Bureau of Meteorology's BARRA reanalysis data set. This highlighted differing regions of Australia that are most likely affected by the individual hazards. There was a high frequency of hail environments over eastern Australia, particularly from the central coast of NSW into southeast Queensland, and the far southeast of Australia. There is a similar peak in the frequency of tornadic environments, with an additional high frequency of environments across much of southern Australia, likely due to cool-season tornadoes. Meanwhile, for environments supportive of severe convective wind, there was a high frequency of environments over part of eastern Australia. There was a very high frequency of environments for all three hazards over much of northern Australia. This is suggested to be an overestimation, except possibly for wind environments, as the majority of data that the statistical models were trained on, was obtained from the southeast of Australia. For wind environments, there were more reports in northern Australia, so the model is likely better trained to assess environment frequency across both regions. These results highlight the importance of developing a spatially balanced data set, along with categorically balanced data, or otherwise breaking a large area into smaller regions, prior to training statistical models. Spatial changes in the frequency of severe thunderstorm hazard environments were assessed between two periods in the BARRA data set. It was found that hail and tornadic environments have increased in frequency over eastern Australia during spring, tornadic environments have also increased in frequency over southeast and northwest Australia during summer, while they have decreased over southwest Australia during autumn and winter. There also is an increase in the frequency of environments supportive of severe wind thunderstorms over Western Australia and eastern Australia during spring, and parts of southwest and southeast Australia during summer.

Supporting Publications

Hartigan, J., MacNamara, S., and Leslie, L.M. Comparing precipitation and temperature trends between inland and coastal locations. ANZIAM J, 60:C109–C126, 2018. DOI: 10.21914/anziamj.v60i0.13967

Hartigan, J., MacNamara, S., and Leslie, L.M. Trends in precipitation and temperature in Canberra. In Elsawah, S. (ed.) MODSIM2019, 23rd International Conference on Modelling and Simulation. Modelling and Simulation Society of Australia and New Zealand, December 2019, pp. 1181–1187, 2019. DOI: 10.36334/modsim.2019.K7.hartigan

Hartigan, J., MacNamara, S., and Leslie, L.M. Application of Machine Learning to Attribution and Prediction of Seasonal Precipitation and Temperature Trends in Canberra, Australia. Climate, 8(6):76, 2020. DOI: 10.3390/cli8060076

Hartigan, J., MacNamara, S., Leslie, L.M., and Speer, M. Attribution and Prediction of Precipitation and Temperature Trends within the Sydney Catchment Using Machine Learning. Climate, 8(10):120, 2020. DOI: 10.3390/cli8100120

Hartigan, J., MacNamara, S., Leslie, L.M., and Speer, M. High resolution simulations of a tornadic storm affecting Sydney. ANZIAM J, 62:C1–C15, 2020. DOI: 10.21914/anziamj.v62.16113

Speer, M., Leslie, L.M., Hartigan, J., and MacNamara, S. Changes in Frequency and Location of East Coast Low Pressure Systems Affecting Southeast Australia. Climate, 9(3):44, 2021. DOI: 10.3390/cli9030044

Speer, M.S., Leslie, L.M., MacNamara, S., and Hartigan, J. From the 1990s climate change has decreased cool season catchment precipitation reducing river heights in Australia's southern Murray-Darling Basin. Sci. Rep., 11(1):1–16, 2021. DOI: 10.1038/s41598-021-95531-4

Speer, M.S., Leslie, L.M., and Hartigan, J. Jet Stream Changes over Southeast Australia during the Early Cool Season in Response to Accelerated Global Warming. Climate, 10(6):84, 2022. DOI: 10.3390/cli10060084

Speer, M.S., Hartigan, J., and Leslie, L.M. Machine Learning Assessment of the Impact of Global Warming on the Climate Drivers of Water Supply to Australia's Northern Murray-Darling Basin. Water, 14(19):3073, 2022. DOI: 10.3390/w14193073

Contents

List of Figures

List of Tables

Abbrevations

- AMO: Atlantic Multidecadal Oscillation
- ARW: Advanced Research WRF
- BARRA: Bureau of Meteorology Atmospheric Reanalysis
- BoM: Bureau of Meteorology
- CAPE: Convective Available Potential Energy
- CIN: Convective Inhibition
- DMC: Deep Moist Convection
- DMI: Dipole Mode Index
- ECMWF: European Centre for Medium-Range Weather Forecasts
- EHI: Energy Helicity Index
- EL: Equilibrium Level
- ENSO: El-Niño Southern Oscillation
- GLD360: Global Lightning Dataset 360
- GlobalSSTA: Global Sea Surface Temperature Anomalies
- GlobalT: Global Temperature Anomalies
- GRIB: General Regularly Distributed Information in Binary form
- IOD: Indian Ocean Dipole
- IQR: Interquartile Range
- IPO: Interdecadal Pacific Oscillation
- JJAS: June, July, August, September (the southern Australian cool season)
- LapR: Lapse Rate
- LCL: Lifted Condensation Level
- LFC: Level of Free Convection
- LI: Lifted Index
- LR: Linear Regression
- LogR: Logistic Regression
- LSM: Land-surface model
- MATLAB: MATrix LABoratory
- MCS: Mesoscale Convective System

- MDB: Murray-Darling Basin
- MESH: Maximum Estimated Size of Hail
- ML: Mixed-Layer
- MSE: Mean Square Error
- MU: Most-Unstable
- NMDB: Northern Murray-Darling Basin
- NSSL: National Severe Storms Laboratory
- NSW: New South Wales
- OT: Overshooting Cloud-Top
- PBL: Planetary Boundary Layer
- PDO: Pacific Decadal Oscillation
- Poly: Polynomial kernel
- PW: Precipitable Water
- QLCS: Quasi-Linear Convective System
- RBF: Radial Basis Function kernel
- RF: Random Forest
- RK3: Runge-Kutta third-order integration scheme
- RMSE: Root Mean Square Error
- RSS: Residual Sum of Squares
- S06: 0–6 km bulk shear
- SA: South Australia
- SB: Surface-Based
- SAM: Southern Annular Mode
- SCA: Sydney Catchment Area
- SCP: Supercell Composite Parameter
- SCT: Severe Convective Thunderstorm
- SD: Standard Deviation
- SHERB: Severe Hazards in Environments with Reduced Buoyancy Index
- SHIP: Significant Hail Parameter
- SOI: Southern Oscillation Index

- SRH: Storm-Relative Helicity
- STA: Severe Thunderstorm Archive
- STP: Significant Tornado Parameter
- SVM: Support Vector Machine
- SVR: Support Vector Regression
- SWEAT: Severe Weather Threat Index
- TMax: Mean maximum temperature
- TMin: Mean minimum temperature
- TPI: Tripole Index for the Interdecadal Pacific Oscillation
- TSSST: Tasman Sea Sea Surface Temperature Anomalies
- TT: Total Totals
- UK: United Kingdom
- USA: United States of America
- UTC: Universal Coordinated Time
- VWS: Vertical Wind Shear
- WA: Western Australia
- WRF: Weather Research and Forecasting

Chapter 1

Introduction

1.1 Motivation

Severe weather occurs globally across a range of spatial and temporal scales, and can have a significant impact on humans and the environment. Many regions of Australia experience a range of severe weather including tropical cyclones, fire, flood, drought and severe convective thunderstorms (SCTs). In order to better prepare communities for the risks of severe weather, detailed understanding of the past and expected future frequency of such severe weather, including environmental and climate attributes, spatiotemporal occurrence and severities, is crucial. Additionally, forecasting skill in the near-future (on the order of days) and on the climate scale (on the order of months or years), is important for preparing both the communities at greatest risk, and the emergency response to severe weater. Improved understanding of severe weather can reduce the impact it has on communities, and of greatest importance, reduce the risk of injury and death.

SCTs can produce a variety of hazards including damaging winds, large hail and tornadoes. They represent a significant forecasting challenge due in part to their relatively small spatiotemporal scale, and the difficulty in understanding where an individual storm may initiate. However, they account for some of the most costly natural disasters in Australia, with almost a quarter of all natural hazard losses associated with SCTs (Middelmann, 2007; Crompton and McAneney, 2008). Many events exceed AU$1 billion in damage, including the Sydney hailstorms of 1999 and 2018, and the Brisbane hailstorm of 2014 (Buckley et al., 2001, 2010; Soderholm et al., 2017b; Insurance Council of Australia, 2019).

Drought occurs over a much longer time scale, on the order of months to years. They have significant societal impacts, particularly for regional and remote communities where alternatives of water supply, such as desalination, typically do not exist. These events can completely destroy such communities and leave many people with long-lasting negative mental health impacts (Stanke et al., 2013; Edwards et al., 2015; Wutich et al., 2020). Additionally, it can act to precondition the environment for devastating fire events, by significantly drying large quantities of potential fuel, as was seen in the Black Summer bushfires of 2019/20 that affected a large area of Australia (Commonwealth of Australia, 2020; Davey and Sarre, 2020). Typically, droughts are broken by significant rain over a season, often resulting in flooding. Although drought is one particular form of severe weather, it is closely associated with additional severe weather, causing the effects of each individual event to often be compounded.

Methodologically, a number of techniques have been used in the past to better understand the atmospheric and climatic conditions associated with severe weather. This includes linear

discriminant analysis, principal component analysis, compositing synoptic maps, and numerical modelling case studies. With the development of atmospheric reanalysis data sets of greater resolution, and more observational data on severe weather events available, there is a significant and ever-growing volume of data available to understand severe weather. This presents a challenge in determining the important features associated with severe weather on a range of time scales. Application of machine learning techniques in the atmospheric sciences is still in its infancy, yet its ability to determine relevant relationships within complex data sets highlights the potential importance of this statistical modelling technique in addressing the challenge of understanding severe weather. Additionally, its low computational cost and adaptability to new data is valuable for improving forecasting skill as more data becomes available. This can prove particularly useful as global warming continues and likely changes the relevance of atmospheric and climatic indicators of severe weather.

While statistical modelling can be valuable for improving our understanding of the variables associated with severe weather, case study analysis also provides valuable understanding of past events that were particularly significant or unusual. Such studies provide useful information for weather forecasting services to better anticipate them in the future, and can assist in understanding how often these events might be expected to occur. Case studies can be performed using observational data or numerical modelling. Numerical modelling of case studies provide a unique insight into the atmospheric conditions that supported a particular severe weather event. These insights often can not be gained from some observational data, with numerical models providing both output of a higher resolution than observations and variables that are not measured, filling data gaps in the observation system. Additionally, numerical modelling can be used to determine the importance of particular features of the environment or geomorphology that supported the severe weather event.

This thesis intends to better understand severe weather in Australia, with a focus on drought and SCT hazards. Trends in drought in southeast Australia are analysed including the trends in modes of variability. Machine learning models are developed to attribute climate drivers to drought in southeast Australia and the predictive skill of these models is assessed for their potential use in forecasting. A numerical modelling case study of a tornado that damaged suburbs in Sydney in December 2015 is performed to assess the ability to simulate the conditions using two different atmospheric reanalysis data sets, and different microphysics packages. Machine learning models are again developed to understand the important environmental attributes of each SCT hazard (wind gusts over 90 km/hr, hail larger than 2 cm or tornadoes), assess the frequency of environments supportive of these hazards across Australia, and determine any trends in SCT environments over a reanalysis data set.

1.2 Drought

Australia is no stranger to drought, with numerous long-term droughts recorded throughout history, including the Federation drought (1895–1902; Bureau of Meteorology, 2009a), World War II drought (1937–1945; Bureau of Meteorology, 2009b) and Millennium drought (1997–2009; Chiew et al., 2014). More recently, a severe drought impacted southeast Australia over 2017–2019, significantly affecting agriculture and water availability. This drought resulted in rapidly decreasing water storages and the lowest inflows ever recorded for the Sydney Catchment Area (SCA), leading to tight water restrictions for Australia's most populous city (Sydney Water, 2020). The three years covering the drought were the driest on record for any three-year period beginning in January over the Murray-Darling Basin and New South Wales (Bureau of Meteorology, 2020). Further, it culminated in the devastating Black Summer bushfires during

the extended 2019–2020 fire season, with more than 24 million hectares burned, 3100 homes destroyed and 33 lives lost directly from the fires (Commonwealth of Australia, 2020; Jalaludin and Morgan, 2021; Davey and Sarre, 2020). Over the course of the summer there were many days where southeast Australia was shrouded in smoke, including the capital cities of Canberra, Melbourne, and Sydney, with at least one site recording $PM_{2.5}$ concentrations exceeding the 95th percentile for 94% of the days in the fire season (Borchers-Arriagada et al., 2020; Simmons et al., 2021). This extreme air pollution resulted in an additional estimated 417–429 deaths and 4456–4753 hospital presentations due to asthma and other cardiovascular and respiratory conditions (Beggs and Zhang, 2020; Johnston et al., 2021; Jalaludin and Morgan, 2021; Linnenluecke and Marrone, 2021; Simmons et al., 2021). The total cost of the disaster is in excess of AUD$10 billion (and potentially over AUD$100 billion when intangible costs are considered; Commonwealth of Australia, 2020; Beggs and Zhang, 2020). This is excluding the human and economic cost of the drought preceding it, with economic costs estimated at $53 billion (Wittwer and Waschik, 2021). The toll on human lives, and regional and remote communities from the drought is not yet fully known.

Though frequently thought of as affecting dry or arid regions, drought can occur anywhere in the world. No truly global definition of drought exists, with even the simplest definition — a deficit of water compared to normal conditions — having complications that will vary by region, including what is considered normal, what constitutes water (e.g., precipitation, land or groundwater reservoirs), and how long or severe a deficit is required for the situation to be considered a drought (Kiem et al., 2016). There are a number of stages to drought that often occur sequentially, beginning with meteorological drought where precipitation deficits accumulate over time (Trenberth et al., 2014; Kiem et al., 2016; Richman and Leslie, 2020). This is followed by agricultural drought, where soil moisture in the root zone impacts crops, then hydrological drought where streamflow, lake and groundwater levels are below average. Ecological drought results from soil moisture levels being low for a sufficient amount of time that the ecosystem begins to fail. Socioeconomic drought is often the final stage, where deficits in water availability are so large that there is a failure to supply crops, water and some services to the public (Richman and Leslie, 2020). Although these stages of drought are useful in understanding the causes of drought and forecasting risk, they exclude the influence of human activity (e.g., water consumption, land surface changes, and water management) on drought that often feedback on each other (AghaKouchak et al., 2021). Human alterations of the water cycle have caused water shortages even during relatively wet years (AghaKouchak et al., 2021), highlighting the significant influence of human activity in producing drought-like conditions.

In addition to the human influence on drought, there are many natural, non-linear effects and positive feedbacks that influence drought onset, severity and longevity. Temperature affects the atmosphere's water-holding capacity, with increasing temperature allowing more water to be held by the atmosphere. This is commonly described using the Clausius-Clapeyron equation:

$$\frac{de_s}{dT} = \frac{L}{T\left(\alpha^v - \alpha^l\right),} \tag{1.2.1}$$

where e_s is the saturation vapour pressure, T is the air temperature, α^v is the specific volume with respect to vapour and α^l the specific volume with respect to liquid (Vallis, 2017). L is the latent heat of evaporation, given by

$$L(T) \equiv h^v - h^l, \tag{1.2.2}$$

where h^v is the specific enthalpy of vapour and h^l is the specific enthalpy of liquid. From the Clausius-Clapeyron equation, the atmosphere can hold approximately 7% more water per 1°C

warmer it is (Trenberth et al., 2003; Pall et al., 2007; Berg et al., 2013; Asadieh and Krakauer, 2015; Berg et al., 2019). However, increasing temperatures causes potential evapotranspiration, and likely true evapotranspiration, to increase, reducing soil moisture levels and water storage. While the atmosphere is able to hold more water, there is a point at which evapotranspiration will begin to decrease as there is little water left to be lost to the atmosphere from soil, vegetation and water reservoirs. This reduces cloud coverage, causing a positive feedback as more solar radiation reaches the ground causing the air temperature to increase futher. Reduced humidity due to large-scale conditions also affect cloud coverage, while higher wind speeds cause the soil to dry out which can eventually reduce evapotranspiration (Trenberth et al., 2014; Richman and Leslie, 2020).

Overall, droughts have a significant and sometimes devastating impact on the human and natural environment. As mentioned earlier, they can increase the likelihood and intensity of bushfires (Trenberth et al., 2014; Richman and Leslie, 2015, 2020). They also can increase the risk of flooding and landslides when rain falls on dry soil (Richman and Leslie, 2015); cause ecosystem damage and loss (Bond et al., 2008); reduce agricultural yields (Cheeseman, 2016); heavily impact tourism (Dube and Nhamo, 2020; Dube et al., 2020); increase energy demand while reducing energy from both hydroelectric and thermoelectric sources (Hardin et al., 2016; van Vliet et al., 2016; Ibrahim et al., 2021); and lead to various effects on human health including decreased sanitation and negative mental health (Dey et al., 2011; Stanke et al., 2013; Edwards et al., 2015; Wutich et al., 2020).

Globally, there are numerous regions of similar latitudes to those encompassing southeast Australia that have observed an increase in the frequency and severity of drought including California (Richman and Leslie, 2015) and Cape Town (Richman and Leslie, 2018). This is the result of changing precipitation patterns along with increases in mean temperature causing greater evaporation. It is expected that the global warming trend will continue to cause an increase in the frequency, severity and duration of droughts in these regions (Niang et al., 2014; Diffenbaugh et al., 2015), with dry regions generally getting drier and wet regions getting wetter as a result of changes in the water-holding capacity of the atmosphere (Trenberth et al., 2014). In the region surrounding Australia, there has been an observed poleward expansion of the Hadley cell, causing a shift in the location and strength of the subtropical ridge (Timbal and Drosdowsky, 2013; Post et al., 2014). This has reduced the frequency of rain-bearing systems reaching the Australian mainland, and caused a reduction in precipitation across southern Australia and parts of eastern Australia (Speer et al., 2011; Timbal and Fawcett, 2013; Timbal and Drosdowsky, 2013).

There are many influences on Australian rainfall variability that impact drought (Risbey et al., 2009; Kiem et al., 2016). Atmospheric blocking patterns have a particular influence on southeast Australia, where slow-moving high pressure systems become anchored in the Tasman Sea, stopping low pressure regions from affecting the southeast of the continent (Risbey et al., 2009). However, the influence of these systems on producing dry weather depends on where they are located. Large-scale climate drivers that influence precipitation in the region include the El-Niño Southern Oscillation (ENSO), with El-Niño events generally causing warmer and drier weather over eastern Australia (Risbey et al., 2009; Ummenhofer et al., 2011). The Indian Ocean Dipole (IOD) influences moisture advection, where positive IOD events reduce large-scale moisture advection from the Indian Ocean over the continent (Risbey et al., 2009; Ummenhofer et al., 2009). The Southern Annular Mode (SAM) affects the location of the belt of westerly winds in the southern hemisphere, influencing the position of low pressure systems over southern Australia, with its influence on precipitation varying by season (Hendon et al.,

2007; Risbey et al., 2009). The Pacific Decadal Oscillation (PDO) is a slowly-varying climate driver that has been shown to influence the effect of ENSO thereby influencing precipitation over eastern Australia (Power et al., 1999; Folland et al., 2002; Verdon et al., 2004; Kiem and Franks, 2004), while the Atlantic Multidecadal Oscillation (AMO) affects the global Walker circulation and the resulting climate variability in the tropical Pacific (Kucharski et al., 2011, 2016).

Considering the expected reduction in precipitation over southeast Australia, there is a need to better understand trends in variables that affect drought in key agricultural and population centres in this region. Statistical modelling can prove useful to determine the key climate drivers that impact precipitation variability in these centres, while also determining if climate change has influenced any precipitation trends (Richman and Leslie, 2015; Deo and Şahin, 2015; Bagirov et al., 2017; Richman and Leslie, 2020). Additionally, modelling could be used to predict precipitation in these areas, which would help improve water management and drought preparedness. Such modelling could be tailored to smaller areas to provide more valuable forecasts of precipitation. However, there have been few studies employing statistical modelling for these purposes, leaving a need to assess the applicability of this approach for southeast Australia.

1.3 Convective Storms

1.3.1 Buoyancy

Convective storms are initiated and partially maintained by buoyant processes in the atmosphere that realise instability in the vertical. *Convection* is a means of heat transfer describing the vertical motion of air parcels in the atmosphere (Markowski and Richardson, 2010). This process removes instability that is generated through uneven solar heating of the Earth's surface and ground heating of air in the planetary boundary layer (PBL; Doswell, 2001). The vertical momentum equation may be expressed as

$$\frac{dw}{dt} = \frac{-1}{\rho}\frac{\partial p}{\partial z} + 2\Omega u\cos\phi - g + F_w, \tag{1.3.1}$$

where w is vertical velocity, ρ is density, p is pressure, $2\Omega u\cos\phi$ is the vertical component of the Coriolis acceleration, g is the acceleration due to gravity, and F_w represents viscous effects. On convective scales, excluding mesoscale convective systems that are sufficiently large and longer-lived, the influence of Coriolis acceleration is considered negligible (Markowski and Richardson, 2010). Further, neglecting viscous effects, equation 1.3.1 can be written as

$$\frac{dw}{dt} = \frac{-1}{\rho}\frac{\partial p}{\partial z} - g \tag{1.3.2}$$

$$\iff \rho\frac{dw}{dt} = -\frac{\partial p}{\partial z} - \rho g. \tag{1.3.3}$$

Now consider a horizontally homogeneous base state $\overline{\rho}(z)$, $\overline{p}(z)$ such that $\rho = \overline{\rho}(z) + \rho'$ and $p = \overline{p}(z) + p'$, where the prime symbol denotes a perturbation from the base state. Further, let the field be in hydrostatic balance, so $\frac{\partial \overline{p}}{\partial z} = -\overline{\rho}g$. Then equation 1.3.3 becomes

$$\rho\frac{dw}{dt} = -\frac{\partial p'}{\partial z} - \rho' g \tag{1.3.4}$$

$$\iff \frac{dw}{dt} = \frac{-1}{\rho}\frac{\partial p'}{\partial z} - \frac{\rho'}{\rho}g. \tag{1.3.5}$$

The first term on the right hand side of equation 1.3.5 is the vertical perturbation pressure gradient force, while the second term represents thermal buoyancy. The equation of state is provided by $p = \rho R_d T_v$ where R_d is the gas constant for dry air and T_v is virtual temperature, given by

$$T_v = T\frac{1 + r_v/\varepsilon}{1 + r_v}, \tag{1.3.6}$$

where r_v is the mixing ratio of water vapour and $\varepsilon = R_d/R_v \approx 0.622$. A simplification of the buoyancy term can be obtained by assuming the equation of state holds for the base state, i.e. $\overline{p} = \overline{\rho} R \overline{T}_v$. The result is

$$B = \frac{-\rho'}{\overline{\rho}} g \approx g\left(\frac{T'_v}{\overline{T}_v} - \frac{p'}{\overline{p}}\right). \tag{1.3.7}$$

Convective cloud development is often considered using *parcel theory*. In this, it is assumed that a parcel of air does not mix with its surrounding environment and both pressure perturbations and hydrometeor loading are negligible. These simplifying assumptions allow the thermal buoyancy to be written as

$$B = g\frac{T'_v}{\overline{T}_v}. \tag{1.3.8}$$

Finally, as $T'_v = T_v - \overline{T}_v$, the buoyancy acceleration can be written as

$$B = g\frac{T_v - \overline{T}_v}{\overline{T}_v}. \tag{1.3.9}$$

Consideration of equation 1.3.9 allows a simplified view of convection - when an air parcel, represented by T_v, is warmer than the environment ($\overline{T}_v$), it experiences positive buoyancy and thus an acceleration upwards. Conversely, an air parcel cooler than its environment will experience negative buoyancy and an acceleration downwards.

The assumptions of parcel theory, however, are physically inaccurate for a variety of reasons (Doswell and Markowski, 2004). Air parcels constantly mix with their environment, which can affect buoyancy. For example, evaporative cooling from mixing with dry air can reduce the buoyancy of a parcel. Additionally, hydrometeors provide drag associated with falling at their terminal velocity, and perturbation pressures are ignored twice, with one of these cases being part of the assumption that thermal buoyancy is the only form of buoyancy experienced by a parcel. One final and important point is that this view of buoyancy is purely one dimensional, raising the complication of which environmental base state (or initial conditions) is most appropriate to consider. Two different base states can lead to completely different results under parcel theory. Despite these limitations, parcel theory provides both a useful basis for understanding convection, and the tools to diagnose instability in the atmosphere and forecast the potential for convective storms to occur (Doswell and Markowski, 2004).

1.3.2 Deep Moist Convection

An *adiabatic process* describes the heating or cooling of an air parcel due to expansion or compression, rather than heat being added or removed from the parcel. Lapse rates are used to describe the rate of change of temperature with height and provide a way of understanding the basics of instability within the atmosphere. The *dry adiabatic lapse rate*, Γ_d, is the rate at which a dry air parcel (one containing no water vapour) cools due to expansion when ascending through the atmosphere, and is approximately 9.8 K km^{-1}. The parcel will become saturated after a certain amount of lifting, at the *lifted condensation level* (LCL). At this point, the rate of cooling through adiabatic ascent changes. The rate can be defined by the *reversible moist*

adiabatic lapse rate, in which total water is conserved, or it may be defined by the *pseudoadiabatic lapse rate*, in which the condensate is assumed to be removed immediately. For the purposes of this discussion, we will consider the difference between the two processes to be negligible, and will generically refer to the lapse rate describing moist adiabatic processes as the *moist adiabatic lapse rate*, Γ_m.

Convection within the atmosphere may be either dry or moist. *Deep moist convection* (DMC) concerns moist convection that covers a significant vertical extent of the troposphere and is assocaited with convective storms. An environmental lapse rate, $\gamma = \frac{-\partial \overline{T}}{\partial z}$, is said to be *conditionally unstable* if $\Gamma_d < \gamma < \Gamma_m$. This is one of three conditions considered necessary for DMC to occur in the widely-used ingredients-based approach to DMC (Doswell et al., 1996). The other two ingredients required are the presence of moisture, and some lifting mechanism for an air parcel to be lifted to its *level of free convection* (LFC). The LFC is the point at which a parcel first becomes positively buoyant. Through its vertical ascent, the parcel will eventually reach a point where it is neutrally buoyant, known as the *equilibrium level* (EL). However, conditional instability is concerned with infinitesimal displacements, rather than the finite displacements that are associated with DMC. *Convective available potential energy* (CAPE) is the vertical distribution of buoyancy between the LFC and EL, and is proportional to the kinetic energy a parcel may gain from the environment due to buoyancy, defined as

$$\text{CAPE} = g \int_{LFC}^{EL} \frac{T_v(z) - T'_v(z)}{T'_v(z)} dz, \tag{1.3.10}$$

where $T_v(z)$ is the virtual temperature of the parcel, and $T'_v(z)$ is the virtual temperature of the environment at height z. CAPE is typically used to assess instability in the atmosphere, as it considers both available moisture and lapse rates (Schultz et al., 2000; Sherwood, 2000). Other indices that can be used to assess atmospheric stability include the Lifted Index (LI; Galway, 1956), Showalter Index (Showalter, 1953) and Total Totals (TT; Miller, 1972). These indices are simple to compute as they are based on temperature differences between at least two layers of the atmosphere, rather than the entire convective region (note: in the case of TT, moisture is included). However, due to this simplification, diagnosed instability can be significantly influenced by topography in the area of interest and is dependent on the selected layer. Although this can complicate consideration of the aforementioned indices, CAPE does not come without its own complications. One necessary factor to consider, which is still unresolved, is what parcel should be used for the calculation of CAPE (Doswell and Rasmussen, 1994). Possibilities include using a parcel at the surface, the most unstable parcel below a certain height, or taking the average parcel value between two layers. All of these approaches lead to different values of CAPE that can alter the anticipation of severe weather on a given day. These factors for both CAPE and the other thermodynamic indices highlights the importance of considering regional effectiveness of specific indices in determining the likelihood of convection (Schultz, 1989).

Convective inhibition (CIN) describes the amount of energy required for a parcel to overcome negative buoyancy and reach its LFC, given by

$$\text{CIN} = -g \int_{0}^{LFC} \frac{T_v(z) - T'_v(z)}{T'_v(z)} dz. \tag{1.3.11}$$

The presence of CIN is why a lifting mechanism is one of the three ingredients necessary for convection to occur. There are many different types of lifting mechanisms including orographic lifting (Kottmeier et al., 2008; Barthlott and Kirshbaum, 2013), airmass boundaries such as the sea-breeze front (Kingsmill, 1995; Wilson and Megenhardt, 1997), and even the collision of such boundaries (Wilson and Schreiber, 1986; Kingsmill, 1995). If there is sufficient lift in the

environment for a parcel to attain its LFC, then the presence of CAPE becomes a sufficient condition for DMC (Schultz et al., 2000). Lifting mechanisms are mesoscale processes, though they may be driven by processes occurring on the synoptic scale. Meanwhile the processes that make the environment favourable for convection are synoptic in scale (Doswell, 1987).

When all three ingredients are available, DMC can occur. The life cycle of a *single-cell thunderstorm*, or *convective cell*, was first described by Byers and Braham (1948). The convective cell was found to have three stages. The first stage, the *cumulus stage*, is characterised by the development of a small cumulus cloud with a single updraft that draws in buoyant air. Once the air reaches its LCL, it condenses and releases latent heat through this process, which further increases parcel buoyancy. Additional contributions to buoyancy may be made through latent heat released if particles freeze resulting in positive buoyancy, and precipitation loading reducing buoyancy. After some time, the cumulus cloud will have developed into a cumulonimbus composed of a significant number of hydrometeors (cloud, graupel, hail, ice, rain, snow and/or water vapour) and have reached heights exceeding 25,000 feet (Byers and Braham, 1948). Eventually the weight of the hydrometeors will be too significant for the updraft to keep them suspended. Evaporation of rain and the melting/sublimation of graupel, hail and snow further contributes negative buoyancy that will induce a downdraft. This marks the *mature stage* of the thunderstorm. As the downdraft reaches the surface, it spreads laterally as a density current known as a *cold pool* with its leading edge called a *gust front* or *outflow boundary*. Secondary convection may form along the gust front, however the cold pool area eventually will become too large for the original updraft to continue sourcing positively buoyant air, marking the *dissipating stage* of the convective cell. The remaining precipitation falls out, with entrainment of air causing evaporation of the cloud at lower levels. The *anvil* cloud in the upper levels may remain as the air at that level does not readily mix out. The whole process, from buoyant air parcel to dissipating single-cell thunderstorm generally takes no longer than one hour.

1.3.3 Vertical Wind Shear and Severe Thunderstorms

While the presence of instability, moisture and a lifting mechanism are necessary for a thunderstorm to develop, a number of environmental factors influence storm organisation, longevity and the potential for severe weather. These include the effect of environmental temperature and CAPE (McCaul et al., 2005; Kirkpatrick et al., 2011), humidity above the cloud base (James and Markowski, 2010) and heights of the LCL and LFC (McCaul and Cohen, 2002). However, *vertical wind shear* (VWS) is the most important contributor to storm organisation (Weisman and Klemp, 1982, 1984; Markowski and Richardson, 2010). VWS is the vector difference between horizontal wind speeds at two layers within the atmosphere. The *0–6 km bulk shear* (herein referred to as *bulk shear*) is a commonly used measurement describing VWS, which is the difference between the wind at 6 km and that at the lowest observation level (close to $z = 0$ km).

The magnitude of VWS in the environment provides useful indications to how convective storms will organise and what mode of convection an individual storm might take (Weisman and Klemp, 1982, 1984; Markowski and Richardson, 2010). When it is low, ordinary single-cell convection develops as described in the Byers and Braham (1948) model. When VWS increases, it acts to separate the storm updraft from its downdraft so the two are maintained independently. When there is moderate bulk shear, the development of *multicell* convection occurs, where stronger low-level storm-relative winds enhance convergence along the gust front, causing successive development of ordinary cells there. While the individual cells have a typical lifetime of approximately one hour, the lifetime of the whole system can last significantly

longer. Depending on the shear and thermodynamic profile, multicell convection may consist of a relatively small-scale cluster, or a much larger-scale system known as a *mesoscale convective system* (MCS). MCSs are often organised into long, unbroken lines of convection, with severe weather tending to occur early in their development. However, a severe wind-producing MCS known as a *derecho*, and heavy precipitation events, are more often associated with mature MCSs (Doswell, 2001). These low-level winds also constrain the gust front, stopping it from cutting the updraft from its supply of warm positively buoyant air. When bulk shear is strong, the environment is typically characterised by significant environmental horizontal vorticity. This horizontal vorticity can be stretched and tilted into a thunderstorm updraft, increasing vorticity and causing updraft rotation. *Supercell* convection results, where the storm is characterised by a relatively long life, and a deep, persistently rotating *mesocyclone* (Davies-Jones, 1984). Supercell motion tends to deviate from the mean wind in the cloud-layer. In the Southern Hemisphere, a supercell with a cyclonically rotating updraft will move to the left of the mean wind, while an anti-cyclonically rotating updraft will move to the right (Weisman and Klemp, 1982). The updrafts of supercells are often very strong owing to the dynamic vertical pressure gradients associated with strong-shear environments interacting with the updraft and the vertical perturbation pressure gradient accelerations due to rotation within the updraft. While ordinary convection is almost entirely buoyancy driven, supercell convection is largely dynamically driven.

Convective mode plays an important role in storm longevity and the potential for severe weather to occur. Storms with a longer life time have higher potential to inflict severe weather as a result of being more organised. They also have a better chance of being reported due to travelling relatively longer distances and therefore are more likely to encounter man-made structures that could be damaged. *Severe convective thunderstorms* (SCTs) are defined by the Bureau of Meteorology (2017) as those storms which produce at least one of hail greater than or equal to 2 cm in diameter, wind gusts of at least 90 km h^{-1}, a tornado, or heavy rain conducive to flash flooding. Most weather organisations do not consider heavy rainfall as a characteristic of SCTs as non-convective weather can also be conducive to flash flooding. Further, heavy rainfall events can be the result of *quasi-stationary convective systems* that form through repeated convection initiation over one location and are often not tied to other severe phenomena or more organised systems (Chappell, 1986). While high-precipitation supercells and other forms of severe convection include heavy rain events, often other severe weather accompanies these storms. As such, this study will not include convective events that produce heavy rain or flash flooding.

In Australia, it is estimated that SCTs account for over a quarter of the total yearly cost of natural hazards (Middelmann, 2007; Crompton and McAneney, 2008). Some of the most costly natural disasters in Australia have been the result of severe convective weather, with numerous events exceeding AU$1 billion in damage (Buckley et al., 2001, 2010; Soderholm et al., 2017b; Insurance Council of Australia, 2019). Of the storm types that exist, supercells are the most prolific producers of severe weather, with damaging winds, large hail and tornadoes all a possibility. Their extensive life increases the chance of severe weather being experienced at any given location as they also generally cover a much larger swathe of land. Severe weather, particularly damaging winds and large hail, and occasionally tornadoes may also be experienced in bow echoes, a type of MCS organised into a relatively long, narrow line (Przybylinski, 1995).

Thunderstorms occur over a broad range of environments, with a continuous spectrum of possible storm types. This has prompted the development of various indices that diagnose both the potential of storms occurring and the different storm modes that might occur. We have

seen purely thermodynamic indices such as CAPE, LI and TT, that are often used to forecast the potential for storm development. As VWS plays a primary role in storm organisation, kinematic indices such as bulk shear are often used to forecast the potential storm mode on a given day. Another often-used kinematic parameter is the *storm-relative helicity* (SRH), which measures the streamwise vorticity in the storm-relative flow over some inflow layer, given by

$$\mathrm{SRH} = \int_0^h (\mathbf{v} - \mathbf{c}) \cdot (\boldsymbol{\nabla} \times \mathbf{v})\, dz, \tag{1.3.12}$$

where h is the depth of the inflow layer, $\mathbf{v}$ is the wind velocity vector, and $\mathbf{c} = (c_x, c_y, 0)$ is the storm motion vector. Environments with considerable SRH are associated with supercell thunderstorms (Davies-Jones, 1984). However, it often is the case that these environments are also associated with values of bulk shear that already suggest supercells are a possibility (Markowski and Richardson, 2010).

Although VWS is important for storm severity, sufficient environmental buoyancy is required to compensate for large VWS. If buoyancy is too low relative to VWS, the initial updraft will not develop into a convective storm as the updraft becomes significantly tilted (Weisman and Klemp, 1984). This highlights the importance of developing indices that consider at least buoyancy and VWS. As other aspects of the environment also influence storm mode and the potential for severe weather, many indices that have been developed take various inputs from the environment. Some have also been developed to diagnose the potential of a specific severe weather phenomenon, such as the significant tornado parameter (STP; Edwards and Thompson, 2000; Thompson et al., 2004).

1.4 Severe Thunderstorm Climatologies

1.4.1 What is a Climatology?

Climatologies describe the average atmospheric conditions at some location, often throughout the year. SCT climatologies are developed by observations of these weather phenomena and will often describe the typical number of occurrences of these events for each individual month, as well as the diurnal cycle of such events, atmospheric conditions commonly associated with these events and any inter-annual variability associated with climate drivers such as ENSO. Such climatologies are of high importance not only to meteorologists, but also the emergency management sector and broader society (Concannon et al., 2000).

One key aspect in climatological studies is observation of the atmospheric variable, or phenomena, of interest. In Australia, reports of SCTs are maintained in the publicly accessible severe thunderstorm archive (STA; Bureau of Meteorology, 2012), which contains reports from 1795 to present. It is worth noting, however, that the routine collection of severe weather data only began with the creation of the severe weather section at the BoM in 1987. As such, there are some inconsistencies in the data collected prior to 1987 (Griffiths et al., 1993; Allen and Allen, 2016). As the severe weather associated with convective storms and tornadoes is of a small scale (Doswell et al., 2006), there is a clustering of reports around the more densely populated cities and east coast of Australia and a lower number of SCT reports in rural areas, with most of these reports concentrated near major road networks (Doswell, 1980; Allen et al., 2011). This inconsistent data collection leads to gaps in our knowledge of SCTs and is an issue in the development of accurate SCT climatologies worldwide (Brooks et al., 2003).

Remote sensing of convective weather has the potential to lower some of the bias to population

centres and improve accuracy in the climatological record. This includes the use of radar data (e.g., Lukach et al., 2017; Soderholm et al., 2017a), lightning flash counts (e.g., Schultz et al., 2011), and satellite observations (e.g., Cecil and Blankenship, 2012; Ferraro et al., 2015; Bedka et al., 2018; Shikhov and Chernokulsky, 2018). The main issue surrounding these methods for Australia is the lack of coverage over the entire continent (Walsh et al., 2016), and a temporal record that is significantly shorter than the STA records. The use of remotely sensed SCT data in climatology development for Australia is only beginning to be realised (e.g., Soderholm et al., 2017a; Bedka et al., 2018; Warren et al., 2020; Dowdy et al., 2020), but is expected to become highly useful in the future as the length of the record increases and improved remote-sensing technology is employed (Doswell, 2001; Walsh et al., 2016).

While there are many problems in creating SCT climatologies, their benefits can not be understated. Beyond an understanding of the seasonal and diurnal variation of SCTs, climatologies can be developed to improve understanding of the regional environmental conditions supportive of severe convection (Brooks et al., 2003; Allen et al., 2011). This may lead to the development of environmental proxies that can be applied in forecasting how supportive an environment is to the formation of SCTs. Climatologies also can be developed to improve understanding of the synoptic patterns associated with severe convection (e.g., Carlson and Ludlam, 1968; Miller, 1972; Doswell, 1980). Such climatologies provide a broader understanding of the formative environments of SCTs as the synoptic-scale often is associated with the mesoscale influences that assist in convective initation (Doswell, 2001). The impact of climate drivers can also be studied in long-term climatologies, providing beneficial information that can be used in the seasonal prediction of severe convective weather (Lu et al., 2015; Childs et al., 2018; Dowdy et al., 2020).

1.4.2 Severe Convective Storms in Australia

The spatial distribution of thunderstorms provides a picture of what the SCT climatology might look like over Australia. Thunderstorm frequency was originally estimated based on the numbers of days thunder was observed at an observing station with the highest frequencies found across northern Australia and a secondary maximum extending from southeast Queensland across central and eastern NSW and into northeastern Victoria (Kuleshov et al., 2002). These maxima occur over the warmer spring and summer months. The maximum over northern Australia is associated with the monsoon and is thus less likely to be severe due to generally weaker VWS leading to less organised thunderstorms, while thunderstorms over southeast Australia are more likely to be severe. It was also found that both Perth and western Tasmania have maxima in thunderstorm activity during the winter, which is associated with strong cold fronts passing through these areas (Kuleshov et al., 2002). The dataset from which this climatology was built is spatially inconsistent, with more observing stations clustered around the coast and less stations in remote Australia. Further issues arise with the lack of a long-term record used, leaving inter-annual variability of storms inadequately captured.

Automatic detection of lightning flashes using satellite instruments remove observational inconsistencies, more accurately covering Australia. Using this data, Dowdy and Kuleshov (2014) confirm the distribution of maxima across northern Australia and from southeast Queensland into central and eastern NSW during the warm season, found in Kuleshov et al. (2002). A bi-modality in the monthly distribution also was identified, with a second peak along the east coast associated with east coast lows in the cooler winter months (Dowdy and Kuleshov, 2014). Interestingly, a cool season peak for Perth and western Tasmania was not found in this study.

While lightning flashes are useful in identifying thunderstorm seasonality and regions where thunderstorms most likely occur, there is not a one-to-one correspondence between thunderstorm and SCT activity. Specifically, the thunderstorm climatology would place an upper bound on the number of SCTs that occur in a region. Until recently, most of the SCT climatologies developed have had only one region of Australia in mind. A particular focus has been the east coast of Australia and the cities of Sydney and Brisbane (Matthews and Geerts, 1995; Yeo, 2005; Schuster et al., 2005). Understandably, there is a large population along the east coast of Australia, but SCTs can and do occur elsewhere (e.g., Buckley et al., 2010; Allen, 2012). It is certainly important to include some level of regional analysis when developing SCT climatologies, because the environments influencing severe convection in one region are likely to be different to those in another region (Brooks et al., 2003; Brooks, 2009). However, this should not be done at the expense of other regions, as SCTs can occur anywhere in Australia.

Proximity soundings are one way in which the relationship between SCTs and the large-scale environment might be determined. These are where rawinsonde launches are made near SCT events both spatially and temporally, providing a vertical slice of atmospheric data including environmental temperature, dewpoint temperature, pressure and wind speed and direction at various heights. These studies aim to find a few parameters that may distinguish between the occurrence of a particular type of weather event or lack thereof. Rawinsonde proximity soundings are limited, especially in Australia, due to the sparse distribution of stations that record such data, and the infrequent release of rawinsonde balloons. As the environment varies throughout the day and conditions can be significantly different even 50 km away, the data obtained in these proximity soundings may not be representative of the environment that the storm developed in (Ryan, 1992; Soderholm et al., 2017b). To mitigate this problem, some studies use *pseudo-proximity soundings* derived from atmospheric reanalysis data (Lee, 2002). These are similar to proximity soundings, although the vertical resolution often is reduced. On the other hand, pseudo-proximity soundings are also available over a finer spatial and temporal resolution, so may more accurately represent the environment a storm develops in than a true proximity sounding. From these pseudo-proximity soundings, meteorological covariates (Brown and Murphy, 1996) are related to severe weather, which are in turn used to estimate the frequency of SCT environments in long-term reanalysis data.

Early work attributing pseudo-proximity soundings to SCT reports in the USA developed a covariate relationship between CAPE and bulk shear (Brooks et al., 2003), which was then applied to global reanalysis data. Brooks et al. (2003) found a maximum frequency in SCT environments across eastern and northern Australia. This approach was adapted by Allen et al. (2011), using Australian SCT reports to develop a covariate relationship between CAPE and bulk shear. This relationship was then applied to a longer reanalysis dataset, confirming the high frequency of severe environments across northern and eastern Australia, but also showing that on average, severe environments occur across the entire continent each season (Allen and Karoly, 2014). Two main limitations of the approach by Allen et al. (2011) include the short sample from which the covariate relationship was developed, and the assumption that the discriminant developed holds over all environments across Australia, despite the large area of Australia and the variety of different climate types across the continent.

Remote sensing provides a new opportunity for understanding the SCT climatology globally, especially in regions such as Australia where observations of SCTs are sparse. Bedka et al. (2018) applied a filtering technique on satellite observations to identify overshooting cloud-tops (OTs; where the cloud extends beyond the tropopause for a period of time) in the Australian region. Such cloud-tops are indicative of strong thunderstorm updrafts where severe weather

hazards might be located (Bedka et al., 2018). Further filtering was applied to identify OTs occurring in environments favourable for hail growth. This identified a region from north of Brisbane down to Sydney, and the Goldfields-Esperance region as the most hail-prone areas of Australia (Bedka et al., 2018). Similar filtering techniques may be applied to determine regions prone to tornadoes and damaging convective winds.

Remotely sensed observations also allow the development of a climatology that considers convective initiation, which is currently missed in reanalysis-based climatologies. The first of these efforts in Australia was developed by Soderholm et al. (2017a) utilising radar and surface-based observations. This climatology found the sea-breeze air mass to highly influence the development of hail storms in southeast Queensland. Synoptic-scale conditions also were considered, finding the most favourable environment for hail storms is when a southeasterly change occurred on a sea-breeze day. A more recent study added to the radar climatology of Brisbane while presenting the first radar-based climatology for the Sydney region (Warren et al., 2020). By utilising insurance data, it found the Brisbane region to experience an average of 26 damaging hail days per annum, while Sydney experienced 32 damaging hail days. Further, it was highlighted that mesoscale circulations interacting with the Great Dividing Range likely play a role in initiation of many of these events (Warren et al., 2020). An additional radar-based climatology for hail and lightning was developed for 10 locations in Australia, finding hail events typically had larger measures of wind shear, while lightning events had a higher freezing level (Dowdy et al., 2020). Unfortunately, radar coverage over Australia is limited. However, the climatologies of Soderholm et al. (2017a), Warren et al. (2020) and Dowdy et al. (2020) demonstrate the benefits of using this data to investigate the SCT climatology and, where possible, to consider the problem of convection initiation sources in particular regions.

Prior studies have displayed the performance of machine learning in determining skillful covariate relationships (e.g., Mercer et al., 2009; Shafer et al., 2010). However, its application to understanding SCTs in Australia has been somewhat limited. One study has applied logistic regression to understand severe convective wind hazards in Australia (Brown and Dowdy, 2021). Using a combination of STA reports and convective wind gust observations (determined using lightning observations), the Bureau of Meteorology's Atmospheric Reanalysis (BARRA; Su et al., 2019) and ERA5 (Hersbach et al., 2020) were used to develop a range of diagnostics related to severe convection, from which logistic regression models were trained. They found the most skillful diagnostics to be the TT, the CAPE-shear relationship from Allen and Karoly (2014), and the derecho composite parameter (Evans and Doswell III, 2001) for indicating the observed environments. There were a number of variables associated with convective wind gusts including measures of instability, low-level moisture and environmental wind speed or shear, with selection varying depending on the type of observation or reanalysis data set used to calculate the variable (Brown and Dowdy, 2021). Model fitting has not yet been extended to different forms of machine learning or SCT hazard in Australia. However, based on the results of previous studies, it likely will prove useful in determining environmental variables associated with each SCT hazard unique to Australia, and the resulting frequency of environments supportive of these hazards.

1.5 Thesis Aims

This thesis aims to better understand particularly important aspects of severe weather that affect much of Australia — drought and SCTs. The focus for Australian drought will be in southeast Australia, where drought is expected to increase in the future, associated with a decrease in precipitation and higher temperatures in the region. It is important for trends in rel-

evant climate variables associated with severe weather including drought to be assessed. This is particularly the case given the non-stationary state of the climate as global warming continues. This investigation will assess precipitation and temperature in the southeast Australian region, along with changes in periodicity within the time series that might indicate a changing influence of climate drivers associated with these variables. Previous studies have analysed the influence of climate drivers on drought in southwest Australia (Richman and Leslie, 2020), where there is a significant decreasing trend in precipitation, using machine learning. The present work will extend this to southeast Australia to assess the ability of these non-linear statistical modelling techniques in attributing climate drivers associated with drought, and their predictive skill on short precipitation time series that exhibit trends relative to the training data. Chapters 3–6 analyse the change in drought and important attributes to drought across southeast Australia by considering case study regions of the Hunter Valley, Sydney Catchment Area, Canberra city, and Northern Murray-Darling Basin. Although these Chapters all relate to drought over southeast Australia, Chapters 3, 4 and 5 are published papers, and Chapter 6 is in review. Therefore, these are standalone Chapters that will be discussed again in a broader context in Chapter 9.

Given their relatively small spatiotemporal extent, SCTs can inflict significant damage to the communities they impact, and are amongst the most expensive natural disasters experienced in Australia. However, research efforts into understanding SCTs in Australia have been limited, with much of the knowledge based on research performed overseas, particularly in the USA, where some of the most frequent and damaging severe convective weather in the world occurs. Australian meteorologists have adapted much of the methods of forecasting SCTs, including the environmental parameters used, to be more suitable to local conditions (Deslandes et al., 2008). This assumes that all environments that give rise to SCTs in Australia have been sampled in the USA (Brooks, 2009), which likely is not the case. It is therefore vital to utilise data from Australia to improve the Australian understanding of SCTs. This likely will help improve forecasting skill, and community and emergency preparedness both now and into the future.

To this end, there have been a limited number of numerical modelling studies of Australian SCTs in the past and, to my knowledge, there have been no numerical modelling studies of tornadic SCTs. This is despite the value in numerical modelling studies, particularly for events of a short spatiotemporal scale as tornadoes often are. Although the frequency of damaging tornadoes impacting major cities is low in Australia, there is a non-zero probability of them occurring, and it is only a matter of time before one devastates an entire community. Therefore, it is important to leverage previous Australian tornadic events to better understand their formation in Australia, rather than rely purely on studies from overseas that are focused on the climate and geographical location of the event in which it occurred. Chapter 7 aims to provide a baseline in understanding tornadic SCTs from numerical simulations of the tornadic event that impacted the Sydney suburb of Kurnell in December 2015. This event was particularly unusual for an SCT, let alone a tornadic SCT, and damaged the desalination plant which could have had serious implications if a drought occurred, like the 2017–2019 drought, during the years it took to repair. High resolution simulations are performed using the Weather Research and Forecasting (WRF) model, with initial conditions from the ERA5 reanalysis data set, and results from various microphysics packages compared.

To advance understanding of Australian SCTs, Chapter 8 aims to assess the ability of machine learning techniques to develop classification models of individual SCT hazards (hail, damaging wind, or tornadoes). The machine learning methods include support vector machines with both radial basis function and polynomial kernels, and random forests. The simpler logistic

regression statistical method is compared, along with the covariate relationship developed in Allen et al. (2011). Models are developed to attribute drivers of each severe hazard from a set of 62 potential covariates. These models are applied to the BARRA reanalysis data set to understand the frequency of environments that support each hazard annually and throughout the year, and are used to assess any change in frequency over the reanalysis data. This work establishes a baseline for further research on classification of SCT hazard environments across Australia using a broad range of statistical techniques.

The data and methods used in trend analysis, statistical modelling and numerical modelling are outlined in Chapter 2, while individual results chapters (Chapters 3–8) discuss in greater detail the data and methodology used in that chapter. Chapter 9 summarises the changes in drought across southeast Australia, the utility of statistical models in understanding Australian drought and SCTs, and the capabilities of high resolution numerical modelling for understanding SCT hazards. Future research directions following this work are also provided in Chapter 9.

Bibliography

AghaKouchak, A., A. Mirchi, K. Madani, G. Di Baldassarre, A. Nazemi, A. Alborzi, H. Anjileli, M. Azarderakhsh, F. Chiang, E. Hassanzadeh, L. Huning, I. Mallakpour, A. Martinez, O. Mazdiyasni, H. Moftakhari, H. Norouzi, M. Sadegh, D. Sadeqi, A. Van Loon, and N. Wanders. Anthropogenic Drought: Definition, Challenges, and Opportunities. Rev. Geophys., 59:1–23, 2021. e2019RG000683.

Allen, J. Supercell Storms: Melbourne's white Christmas 2011. Bulletin of the Australian Meteorological and Oceanographic Society, 25:47–51, 2012.

Allen, J. and E. Allen. A review of severe thunderstorms in Australia. Atmos. Res., 178: 347–366, 2016.

Allen, J. and D. Karoly. A climatology of Australian severe thunderstorm environments of 1979–2011: inter-annual variability and ENSO influence. Int. J. Climatol., 34:81–97, 2014.

Allen, J., D. Karoly, and G. Mills. A severe thunderstorm climatology for Australia and associated thunderstorm environments. Australian Meteorological and Oceanographic Journal, 61:143–158, 2011.

Asadieh, B. and N. Krakauer. Global trends in extreme precipitation: climate models versus observations. Hydrol. Earth Syst. Sci., 19:877–891, 2015.

Bagirov, A., A. Mahmood, and A. Barton. Prediction of monthly rainfall in Victoria, Australia: Clusterwise linear regression approach. Atmos. Res., 188:20–29, 2017.

Barthlott, C. and D. Kirshbaum. Sensitivity of deep convection to terrain forcing over Mediterranean islands. Q. J. R. Meteorol. Soc., 139:1762–1779, 2013.

Bedka, K., J. Allen, H. Punge, M. Kunz, and D. Simanovic. A Long-Term Overshooting Cloud-Top Detection Database over Australia Derived from MTSAT Japanese Advanced Meteorological Imager Observations. J. Appl. Meteorol. Climatol., 57:937–951, 2018.

Beggs, P. and Y. Zhang. The 2020 special reports of the MJA-Lancet Countdown on health and climate change: lessons learnt from Australia's '"Black Summer"' (summary. Med. J. Aust., 213:1–4, 2020.

Berg, P., C. Moseley, and J. Haerter. Strong increase in convective precipitation in response to higher temperatures. Nat. Geosci., 6:181–185, 2013.

Berg, P., O. Christensen, K. Klehmet, G. Lenderink, J. Olsson, C. Teichmann, and W. Yang. Summertime precipitation extremes in a EURO-CORDEX 0.11° ensemble at an hourly resolution. Nat. Hazards Earth Syst. Sci., 19:957–971, 2019.

Bond, N., P. Lake, and A. Arthington. The impacts of drought on freshwater ecosystems: an Australian perspective. Hydrobiologia, 600:3–16, 2008.

Borchers-Arriagada, N., A. Palmer, D. Bowman, G. Morgan, B. Jalaludin, and F. Johnston. Unprecedented smoke-related health burden associated with the 2019–20 bushfires in eastern Australia. Med. J. Aust., 213:282–283, 2020.

Brooks, H. Proximity soundings for severe convection for Europe and the United States from reanalysis data. Atmos. Res., 93:546–553, 2009.

Brooks, H., J. Lee, and J. Craven. The spatial distribution of severe thunderstorm and tornado environments from global reanalysis data. Atmos. Res., 67–68:73–94, 2003.

Brown, A. and A. Dowdy. Severe convection-related winds in Australia and their associated environments. J. South. Hemisphere Earth Syst. Sci., 71:30–52, 2021.

Brown, B. and A. Murphy. Verification of aircraft icing forecasts: The use of standard measures and meteorological covariates. In Preprints, 13th Conf. on Probability and Statistics in the Atmospheric Sciences, San Francisco, CA, Amer. Meteor. Soc, pages 251–252, 1996.

Buckley, B., L. Leslie, and Y. Wang. The Sydney Hailstorm of April 14, 1999: Synoptic description and numerical simulation. Meteorol. Atmospheric Phys., 76:167–182, 2001.

Buckley, B., W. Sullivan, P. Chan, and M. Leplastrier. Two Record Breaking Australian Hailstorms: Storm Environments, Damage Characteristics and Rarity. In 25th Conference on Severe Local Storms, pages 11–14. American Meteorological Society Denver, Colorado, 2010.

Bureau of Meteorology. The "federation drought", 1895–1902. https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought1.html, 2009a.

Bureau of Meteorology. The World War II droughts 1937–1945. https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought3.html, 2009b.

Bureau of Meteorology. Severe Storms Archive, 2012. URL http://www.bom.gov.au/australia/stormarchive/.

Bureau of Meteorology. Severe Thunderstorms, 2017. URL http://www.bom.gov.au/weather-services/severe-weather-knowledge-centre/severethunder.shtml.

Bureau of Meteorology. Previous droughts. http://www.bom.gov.au/climate/drought/knowledge-centre/previous-droughts.shtml, 2020.

Byers, H. and J. Braham, R.R. Thunderstorm Structure and Circulation. J. Meteorol., 5:71–85, 1948.

Carlson, T. and F. Ludlam. Conditions for the occurrence of severe local storms. Tellus, 20: 203–226, 1968.

Cecil, D. and C. Blankenship. Towards a Global Climatology of Hailstorms as Estimated by Satellite Passive Microwave Imagers. J. Clim., 25:687–703, 2012.

Chappell, C. Quasi-Stationary Convective Events. In Mesoscale meteorology and forecasting, pages 289–310. Springer, 1986.

Cheeseman, J. Food Security in the Face of Salinity, Drought, Climate Change, and Population Growth. In Halophytes for food security in dry lands, pages 111–123. Elsevier, 2016.

Chiew, F. H. S., N. J. Potter, J. Vaze, C. Petheram, L. Zhang, J. Teng, and D. A. Post. Observed hydrologic non-stationarity in far south-eastern Australia: implications for modelling and prediction. Stoch. Environ. Res. Risk Assess., 28:3–15, 2014.

Childs, S., R. Schumacher, and J. Allen. Cold-Season Tornadoes: Climatological and Meteorological Insights. Weather Forecast., 33:671–691, 2018.

Commonwealth of Australia. Royal Commission into National Disaster Arrangements Report, 2020. URL https://naturaldisaster.royalcommission.gov.au/system/files/2020-11/Royal%20Commission%20into%20National%20Natural%20Disaster%20Arrangements%20-%20Report%20%20%5Baccessible%5D.pdf.

Concannon, P., H. Brooks, and C. Doswell, III. Climatological Risk of Strong and Violent Tornadoes in the United States. In Preprints, 2nd Symp. on Environmental Applications, Long Beach, CA, Amer. Meteor. Soc, pages 212–219, 2000.

Crompton, R. and K. McAneney. Normalised Australian insured losses from meteorological hazards: 1967–2006. Environ. Sci. Policy, 11:371–378, 2008.

Davey, S. and A. Sarre. Editorial: the 2019/20 Black Summer bushfires. Aust. For., 83:47–51, 2020.

Davies-Jones, R. Streamwise Vorticity: The Origin of Updraft Rotation in Supercell Storms. J. Atmos. Sci., 59:3178–3205, 1984.

Deo, R. and M. Şahin. Application of the Artificial Neural Network model for prediction of monthly Standardized Precipitation and Evapotranspiration Index using hydrometeorological parameters and climate indices in eastern Australia. Atmos. Res., 161:65–81, 2015.

Deslandes, R., H. Richter, and T. Bannister. The end-to-end severe thunderstorm forecasting system in Australia: overview and training issues. Australian Meteorological Magazine, 57: 329–343, 2008.

Dey, N., M. Alam, A. Sajjan, M. Bhuiyan, L. Ghose, Y. Ibaraki, and F. Karim. Assessing Environmental and Health Impact of Drought in the Northwest Bangladesh. J. Environ. Sci. Nat. Resour., 4:89–97, 2011.

Diffenbaugh, N. S., D. L. Swain, and D. Touma. Anthropogenic warming has increased drought risk in California. Proc. Natl. Acad. Sci. U.S.A., 112:3931–3936, 2015.

Doswell, C. Severe Convective Storms—An Overview. In Severe Convective Storms, volume 50, pages 1–26. Springer, 2001.

Doswell, C. and P. Markowski. Is Buoyancy a Relative Quantity? Mon. Weather Rev., 132: 853–863, 2004.

Doswell, C., III. Synoptic-Scale Environments Associated with High Plains Severe Thunderstorms. Bull. Am. Meteorol. Soc., 61:1388–1400, 1980.

Doswell, C., III. The Distinction between Large-Scale and Mesoscale Contribution to Severe Convection: A Case Study Example. Weather Forecast., 2:3–16, 1987.

Doswell, C., III, H. Brooks, and R. Maddox. Flash Flood Forecasting: An Ingredients-Based Methodology. Weather Forecast., 11:560–581, 1996.

Doswell, C., III, R. Edwards, R. Thompson, J. Hart, and K. Crosbie. A Simple and Flexible Method for Ranking Severe Weather Events. Weather Forecast., 21:939–951, 2006.

Doswell, I., C.A. and E. Rasmussen. The Effect of Neglecting the Virtual Temperature Correction on CAPE Calculations. Weather Forecast., 9:625–629, 1994.

Dowdy, A. and Y. Kuleshov. Climatology of lightning activity in Australia: spatial and seasonal variability. Australian Meteorological and Oceanographic Journal, 64:103–108, 2014.

Dowdy, A., J. Soderholm, J. Brook, A. Brown, and H. McGowan. Quantifying Hail and Lightning Risk Factors Using Long-Term Observations Around Australia. J. Geophys. Res. Atmos., 125:1–14, 2020. 2020JD0331101.

Dube, K. and G. Nhamo. Vulnerability of nature-based tourism to climate variability and change: Case of Kariba resort town, Zimbabwe. J. Outdoor Recreat. Tour., 29:1–13, 2020. 100281.

Dube, K., G. Nhamo, and D. Chikodzi. Climate change-induced droughts and tourism: Impacts and responses of Western Cape province, South Africa. J. Outdoor Recreat. Tour., pages 1–10, 2020. 100319.

Edwards, B., M. Gray, and B. Hunter. The Impact of Drought on Mental Health in Rural and Regional Australia. Soc. Indic. Res., 121:177–194, 2015.

Edwards, R. and R. Thompson. RUC-2 Supercell Proximity Soundings, Part II: An Independent Assessment of Supercell Forecast Parameters. In Preprints, 20th Conf. on Severe Local Storms, Orlando, FL, Amer. Meteor. Soc, pages 435–438, 2000.

Evans, J. and C. Doswell III. Examination of Derecho Environments Using Proximity Soundings. Weather Forecast., 16:329–342, 2001.

Ferraro, R., J. Beauchamp, D. Cecil, and G. Heymsfield. A prototype hail detection algorithm and hail climatology developed with the advanced microwave sounding unit (AMSU). Atmos. Res., 163:24–35, 2015.

Folland, C., J. Renwick, M. Salinger, and A. Mullan. Relative influences of the Interecadal Pacific Oscillation and ENSO on the South Pacific Convergence Zone. Geophys. Res. Lett., 29:21–1, 2002.

Galway, J. The Lifted Index as a Predictor of Latent Instability. Bull. Am. Meteorol. Soc., 37: 528–529, 1956.

Griffiths, D., J. Colquhoun, K. Batt, and T. Casinader. Severe Thunderstorms in New South Wales: Climatology and Means of Assessing the Impact of Climate Change. Clim. Change, 25:369–388, 1993.

Hardin, E., A. AghaKouchak, M. Qomi, K. Madani, B. Tarroja, Y. Zhou, T. Yang, and S. Samuelsen. California drought increases CO_2 footprint of energy. Sustain. Cities Soc., 28:450–452, 2016.

Hendon, H. H., D. W. J. Thompson, and M. C. Wheeler. Australian Rainfall and Surface Temperature Variations Associated with the Southern Hemisphere Annular Mode. J. Clim., 20:2452–2467, 2007.

Hersbach, H., B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. The ERA5 global reanalysis. Q. J. R. Meteorol. Soc., 146:1999–2049, 2020.

Ibrahim, N., S. Wan Alwi, Z. Abdul Manan, A. Mustaffa, and K. Kidam. Impact of Drought Phenomenon on Renewable and Non-renewable Energy Systems in the ASEAN Countries. Chem. Eng. Trans., 83:73–78, 2021.

Insurance Council of Australia. Insurance bill for summer catastrophes passes 2.2 billion, 2019.

Jalaludin, B. and G. Morgan. What does climate change have to do with bushfires? Aust. Health Rev., 45:4–6, 2021.

James, R. and P. Markowski. A Numerical Investigation of the Effects of Dry Air Aloft on Deep Convection. Mon. Weather Rev., 138:140–161, 2010.

Johnston, F., N. Borchers-Arriagada, G. Morgan, B. Jalaludin, A. Palmer, G. Williamson, and D. Bowman. Unprecedented health costs of smoke-related $PM_{2.5}$ from the 2019–20 Australian megafires. Nat. Sustain., 4:42–47, 2021.

Kiem, A. and S. Franks. Multi-decadal variability of drought risk, eastern Australia. Hydrol. Process., 18:2039–2050, 2004.

Kiem, A. S., F. Johnson, S. Westra, A. van Dijk, J. P. Evans, A. O'Donnell, A. Rouillard, C. Barr, J. Tyler, M. Thyer, D. Jakob, F. Woldemeskel, B. Sivakumar, and R. Mehrotra. Natural hazards in Australia: droughts. Clim. Change, 139:37–54, 2016.

Kingsmill, D. Convection Initiation Associated with a Sea-Breeze Front, a Gust Front, and Their Collision. Mon. Weather Rev., 123:2913–2933, 1995.

Kirkpatrick, C., J. McCaul, E.W., and C. Cohen. Sensitivities of Simulated Convective Storms to Environmental CAPE. Mon. Weather Rev., 139:3514–3532, 2011.

Kottmeier, C., N. Kalthoff, C. Barthlott, C. Corsmeier, J. Baelen, A. Behrendt, R. Behrendt, A. Blyth, R. Coulter, S. Crewell, P. Girolamo, M. Dorninger, C. Flamant, T. Foken, M. Hagen, C. Hauck, H. Höller, H. Konow, M. Kunz, H. Mahlke, S. Mobbs, E. Richard, R. Steinacker, T. Weckwerth, A. Wieser, and V. Wulfmeyer. Mechanisms initiating deep convection over complex terrain durings COPS. Meteorol. Z., 17:931–948, 2008.

Kucharski, F., I.-S. Kang, R. Farneti, and L. Feudale. Tropical Pacific response to 20th century Atlantic warming. Geophys. Res. Lett., 38:5, 2011.

Kucharski, F., F. Ikram, F. Molteni, R. Farneti, S. Kang, H.-H. No, M. King, G. Giuliani, and K. Mogensen. Atlantic forcing of Pacific decadal variability. Clim. Dyn., 46:2337–2351, 2016.

Kuleshov, Y., G. de Hoedt, W. Wright, and A. Brewster. Thunderstorm distribution and frequency in Australia. Australian Meteorological Magazine, 51:145–154, 2002.

Lee, J. Tornado proximity soundings from the NCEP/NCAR reanalysis data, 2002.

Linnenluecke, M. and M. Marrone. Air pollution, human health and climate change: newspaper coverage of Australian bushfires. Environ. Res. Lett., 16:1–11, 2021. 125003.

Lu, M., M. Tippett, and U. Lall. Changes in the seasonality of tornado and favorable genesis conditions in the central United States. Geophys. Res. Lett., 42:4224–4231, 2015.

Lukach, M., L. Foresti, O. Giot, and L. Delobbe. Estimating the occurrence and severity of hail based on 10 years of observations from weather radar in Belgium. Meteorol. Appl., 24: 250–259, 2017.

Markowski, P. and Y. Richardson. Mesoscale Meteorology in Midlatitudes. Wiley-Blackwell, 2010.

Matthews, C. and B. Geerts. Characteristic thunderstorm distribution in the Sydney area. Australian Meteorological Magazine, 44:127–138, 1995.

McCaul, J., E.W. and C. Cohen. The Impact on Simulated Storm Structure and Intensity of Variations in the Mixed Layer and Moist Layer Depths. Mon. Weather Rev., 130:1722–1748, 2002.

McCaul, J., E.W., C. Cohen, and C. Kirkpatrick. The Sensitivity of Simulated Storm Structure, Intensity, and Precipitation Efficiency to Environmental Temperature. Mon. Weather Rev., 133:3015–3037, 2005.

Mercer, A., C. Shafer, C. Doswelll, III, L. Leslie, and M. Richman. Objective Classification of Tornadic and Nontornadic Severe Weather Outbreaks. Mon. Weather Rev., 137:4355–4368, 2009.

Middelmann, M. Natural Hazards in Australia: Identifying Risk Analysis Requirements. Technical report, Geoscience Australia, 2007.

Miller, R. Notes On Analysis and Severe Storm Forecasting Procedures of the Air Force Global Weather Central. Technical Report 200, Air Weather Service, United States Air Force, 1972.

Niang, I., O. C. Ruppel, M. A. Abdrabo, A. Essel, C. Lennard, J. Padgham, and P. Urquhart. 2014: Africa. In: Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, 2014. 1199–1265 pp.

Pall, P., M. Allen, and D. Stone. Testing the Clausius-Clapeyron constraint on changes in extreme precipitation under CO_2 warming. Clim. Dyn., 28:351–363, 2007.

Post, D., B. Timbal, F. Chiew, H. Hendon, H. Nguyen, and R. Moran. Decrease in southeastern Australian water availability linked to ongoing Hadley cell expansion. Earth's Future, 2:231–238, 2014.

Power, S., T. Casey, C. Folland, A. Colman, and V. Mehta. Inter-decadal modulation of the impact of ENSO on Australia. Clim. Dyn., 15:319–324, 1999.

Przybylinski, R. The Bow Echo: Observations, Numerical Simulations, and Severe Weather Detection Methods. Weather Forecast., 10:203–218, 1995.

Richman, M. and L. Leslie. The 2015–2017 Cape Town Drought: Attribution and Prediction Using Machine Learning. Procedia Comput. Sci., 140:248–257, 2018.

Richman, M. B. and L. M. Leslie. Uniqueness and Causes of the California Drought. Procedia Comput. Sci., 61:428–435, 2015.

Richman, M. B. and L. M. Leslie. Machine Learning for Attribution of Heat and Drought in Southwestern Australia. Procedia Comput. Sci., pages 3–10, 2020.

Risbey, J., M. Pook, P. McIntosh, M. Wheeler, and H. Hendon. On the Remote Drivers of Rainfall Variability in Australia. Mon. Weather Rev., 137:3233–3253, 2009.

Ryan, C. Dynamical Classification of Australian Thunderstorms. Australian Government Pub. Service, 1992.

Schultz, C., W. Petersen, and L. Carey. Lightning and Severe Weather: A Comparison between Total and Cloud-to-Ground Lightning Trends. Weather Forecast., 26:744–755, 2011.

Schultz, D., P. Schumacher, and C. Doswell, III. The Intricacies of Instabilities. Mon. Weather Rev., 128:4143–4148, 2000.

Schultz, P. Relationships of Several Stability Indices to Convective Weather Events in Northeast Colorado. Weather Forecast., 4:73–80, 1989.

Schuster, S., R. Blong, and M. Speer. A Hail Climatology of the Greater Sydney Area and New South Wales, Australia. Int. J. Climatol., 25:1633–1650, 2005.

Shafer, C., A. Mercer, L. Leslie, M. Richman, and C. Doswell, III. Evaluation of WRF Model Simulations of Tornadic and Nontornadic Outbreaks Occurring in the Spring and Fall. Mon. Weather Rev., 138:4098–4119, 2010.

Sherwood, S. On Moist Instability. Mon. Weather Rev., 128:4139–4142, 2000.

Shikhov, A. and A. Chernokulsky. A satellite-derived climatology of unreported tornadoes in forested regions of northeast Europe. Remote Sens. Environ., 204:553–567, 2018.

Showalter, A. A Stability Index for Thunderstorm Forecasting. Bull. Am. Meteorol. Soc., 34: 250–252, 1953.

Simmons, J., C. Paton-Walsh, A. Mouat, J. Kaiser, R. Humphries, M. Keywood, A. Sutresna, D. Griffith, J. Naylor, and J. Ramirez-Gamboa. The gas and aerosol phase composition of smoke plumes from the 2019–2020 Black Summer bushfires and potential implications for human health. pages 1–29, 2021. Preprint.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Weckwerth, and M. Coleman. An 18-year climatology of hailstorm trends and related drivers across southeast Queensland Australia. Q. J. R. Meteorol. Soc., 143:1123–1135, 2017a.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Wedd, and T. Weckwerth. Diurnal Preconditioning of Subtropical Coastal Convective Storm Environments. Mon. Weather Rev., 145:3839–3859, 2017b.

Speer, M., L. Leslie, and A. Fierro. Australian east coast rainfall decline related to large scale climate drivers. Clim. Dyn., 36:1419–1429, 2011.

Stanke, C., M. Kerac, C. Prudhomme, J. Medlock, and V. Murray. Health Effects of Drought: a Systematic Review of the Evidence. PLOS Curr., 5, 2013.

Su, C.-H., N. Eizenberg, P. Steinle, D. Jakob, P. Fox-Hughes, C. White, S. Rennie, C. Franklin, I. Dharssi, and H. Zhu. BARRAv1.0: the Bureau of Meteorology Atmospheric high-resolution Regional Reanalysis for Australia. Geosci. Model Dev., 12:2049–2068, 2019.

Sydney Water. Sydney's drought: Aquabumps shows just how dry it got. https://www.sydneywater.com.au/SW/about-us/our-publications/Media/sydney-s-drought--aquabumps-shows-just-how-dry-it-got/index.htm, 2020.

Thompson, R., R. Edwards, and C. Mead. An Update to the Supercell Composite and Significant Tornado Parameters. In Preprints, 22nd Conf. on Severe Local Storms, Hyannis, MA, Amer. Meteor. Soc. P, volume 8, 2004.

Timbal, B. and W. Drosdowsky. The relationship between the decline of Southeastern Australian rainfall and the strengthening of the subtropical ridge. Int. J. Climatol., 33:1021–1034, 2013.

Timbal, B. and R. Fawcett. A Historical Perspective on Southeastern Australian Rainfall since 1865 Using the Instrumental Record. J. Clim., 26:1112–1129, 2013.

Trenberth, K., A. Dai, R. Rasmussen, and D. Parsons. The changing character of precipitation. Bull. Am. Meteorol. Soc., 84:1205–1217, 2003.

Trenberth, K., A. Dai, G. van der Schrier, P. Jones, J. Barichivich, K. Briffa, and J. Sheffield. Global warming and changes in drought. Nat. Clim. Change, 4:17–22, 2014.

Ummenhofer, C., A. Gupta, P. Briggs, M. England, P. McIntosh, G. Meyers, M. Pook, M. Raupach, and J. Risbey. Indian and Pacific Ocean Influences on Southeast Australian Drought and Soil Moisture. J. Clim., 24:1313–1336, 2011.

Ummenhofer, C. C., M. H. England, P. C. McIntosh, G. A. Meyers, M. J. Pook, J. S. Risbey, A. S. Gupta, and A. S. Taschetto. What causes southeast Australia's worst droughts? Geophys. Res. Lett., 36, 2009.

Vallis, G. Atmospheric and Oceanic Fluid Dynamics. Cambridge, 2017.

van Vliet, M., J. Sheffield, D. Wiberg, and E. Wood. Impacts of recent drought and warm years on water resources and electricity supply worldwide. Environ. Res. Lett., 11:1–10, 2016. 124021.

Verdon, D., A. Wyatt, A. Kiem, and S. Franks. Multidecadal variability of rainfall and streamflow: Eastern Australia. Water Resour. Res., 40, 2004. W10201.

Walsh, K., C. White, K. McInnes, J. Holmes, S. Schuster, H. Richter, J. Evans, A. Di Luca, and R. Warren. Natural hazards in Australia: storms, wind and hail. Clim. Change, 139: 55–67, 2016.

Warren, R., H. Ramsay, S. Siems, M. Manton, J. Peter, A. Protat, and A. Pillalamarri. Radar-based climatology of damaging hailstorms in Brisbane and Sydney, Australia. Q. J. R. Meteorol. Soc., 146:505–530, 2020. doi: 10.1002/qj.3693.

Weisman, M. and J. Klemp. The Dependence of Numerically Simulated Convective Storms on Vertical Wind Shear and Buoyancy. Mon. Weather Rev., 110:504–520, 1982.

Weisman, M. and J. Klemp. The Structure and Classification of Numerically Simulated Convective Storms in Directionally Varying Wind Shears. Mon. Weather Rev., 112:2479–2498, 1984.

Wilson, J. and D. Megenhardt. Thunderstorm Initiation, Organization, and Lifetime Associated with Florida Boundary Layer Convergence Lines. Mon. Weather Rev., 125:1507–1525, 1997.

Wilson, J. and W. Schreiber. Initiation of Convective Storms at Radar-Observed Boundary-Layer Convergence Lines. Mon. Weather Rev., 114:2516–2536, 1986.

Wittwer, G. and R. Waschik. Estimating the economic impacts of the 2017–2019 drought and 2019–2020 bushfires on regional NSW and the rest of Australia. Agric. Resour. Econ., 65: 918–936, 2021.

Wutich, A., A. Brewis, and A. Tsai. Water and mental health. WIREs Water, 7:1–16, 2020. e1461.

Yeo, C. Severe thunderstorms in the Brisbane region and a relationship to the El niño Southern Oscillation. Australian Meteorological Magazine, 54:197–202, 2005.

Chapter 2

Data and Methodologies

2.1 Precipitation, Temperature and Climate Drivers

The Australian Bureau of Meteorology (BoM) records precipitation and temperature at numerous locations around Australia. These data are fed into numerical weather prediction models to inform forecasts, and can also be used to understand the climate of a city or region. In the early 20th century, this data was recorded maually by volunteers. However, from the mid-20th century onwards, the majority of observing sites have been replaced with automatic weather stations allowing a higher temporal resolution data set. This precipitation and temperature data is contained in the Climate Data Online section of the BoM's website (`http://www.bom.gov.au/climate/data/index.shtml`).

Table 2.1 lists the 24 stations that have been selected from the BoM's climate data to broadly cover the southeast Australia region. These sites were used for analysis of precipitation and temperature trends, attribution of climate drivers with precipitation, and prediction of precipitation in southeast Australia. The stations were selected to contain a relatively long, continuous record of data with the preference of recording both temperature and precipitation. However, not all precipitation sites that were selected in the Sydney Catchment Area also recorded temperature. Meanwhile, in the Northern Murray-Darling Basin (NMDB), the majority of sites that recorded precipitation did not record temperature. For those precipitation sites in the NMDB that did not record temperature, the nearest temperature site was used. From the data available, monthly mean maximum temperature, mean minimum temperature, and monthly total precipitation were analysed.

There is missing data within some of the precipitation and temperature time series used due to factors including equipment maintenance or damage, changes to weather station location or discontinuation of the weather station. Much of the missing data is only for a few months in the year, and was filled using a 4-year moving average centred on the year of the missing data, wherever possible. If the data was missing towards the edges of the time series, then a smaller average centred on the year, or an average from the years following or preceding the data point, would be taken. However, some sites had missing data for a number of years. If there was no nearby station, then this data would be filled using a simple moving average of the previous four years. If there was a nearby station in operation that also had data coinciding with the station currently being used, then the difference in mean precipitation or temperature between the two sites is calculated and the data adjusted by this mean to fill the missing data in the station being used. Finally, if a site has been discontinued, then this data is combined with a nearby site that has a relatively long record following the discontinuation of the other site. If the site is located some distance from the previous site, with a considerable change in

Table 2.1: List of locations where weather station data was obtained to assess trends in precipitation and temperature over regions of southeast Australia.

Location	Coordinate	Precipitation	Temperature
Augathella	25.8°S, 146.6°E	X	
Bellata	29.9°S, 149.8°E	X	
Bingara	29.9°S, 150.6°E	X	
Blackheath	33.6°S, 150.3°E	X	
Braidwood	35.4°S, 149.8°E	X	
Canberra	35.3°S, 149.1°E	X	X
Charleville	26.4°S, 146.2°E		X
Cunnamulla	28.1°S, 145.7°E	X	
Curlewis	31.1°S, 150.3°E	X	
Darkes Forest	34.2°S, 150.9°E	X	
Goulburn	34.7°S, 149.7°E	X	X
Gunnedah	31.0°S, 150.3°E		X
Inverell	29.8°S, 151.1°E		X
Lithgow	33.5°S, 150.2°E	X	X
Miles	26.7°S, 150.2°E	X	X
Moss Vale	34.6°S, 150.4°E	X	X
Moree	29.5°S, 149.8°E		X
Newcastle	32.9° S, 151.8°E	X	X
Normandy	26.8°S, 148.4°E	X	
Scone	32.1°S, 150.9°E	X	X
St. George	28.0°S, 148.6°E		X
Surat	27.2°S, 149.1°E	X	
Thargomindah	28.0°S, 143.8°E		X
Walgett	30.0°S, 148.1°E		X

geography, then the data is adjusted by the difference in mean precipitation or temperature between the two sites.

Seasonal variations in climate occur as a result of the Earth's tilt from the sun affecting heating at the surface. For example, during the mid-latitude winter, temperatures are much colder than during summer as that hemisphere is farther from the sun. However, the climate of any one particular region also is influenced by the broad-scale circulation of the atmosphere, for which there is variability on time scales of years to decades. Climate drivers refer to specific patterns of change in the broad-scale circulation for different regions on Earth. These are often focused on one ocean basin, although what occurs in one ocean basin often affects another (e.g., Kucharski et al., 2011, 2016; Johnson et al., 2018). The El-Niño Southern Oscillation (ENSO) is an example of a climate driver with particular influence on Australian climate. ENSO is associated with variability in strength of the Walker Circulation in the tropical Pacific Ocean. During an El-Niño phase, the central and eastern tropical Pacific Ocean warms, causing a weakening of the Walker Circulation and increased convection in the central and eastern tropical Pacific. This causes eastern Australia to typically be warmer and drier (McBride and Nicholls, 1983; Nicholls, 1989; Ummenhofer et al., 2011). The La-Niña phase is the opposite phase to El-Niño, causing a strengthened Walker Circulation, with warmer ocean temperatures in the western tropical Pacific, increased cloudiness and precipitation, and cooler temperatures across eastern Australia.

There are numerous climate drivers that have the potential to influence precipitation, temperature and severe weather activity across Australia. These include the Atlantic Multidecadal Oscillation (AMO; Enfield et al., 2001), the Dipole Mode Index (DMI; Webster et al., 1999; Saji et al., 1999), Niño3.4 (Bamston et al., 1997), the Tripole Index for the Interdecadal Pacific Oscillation (TPI; Henley et al., 2015), the Pacific Decadal Oscillation (PDO; Mantua and Hare, 2002), the Southern Annular Mode (SAM; Limpasuvan and Hartmann, 1999), and the Southern Oscillation Index (SOI; Troup, 1965; McBride and Nicholls, 1983). Local sea surface temperatures, such as the Tasman Sea Sea Surface Temperature Anomalies (TSSST, which affects southeast Australia) also can influence precipitation and temperature as the Tasman Sea is a source of moisture for precipitation over the southeast (Holgate et al., 2020) and anomalies in the TSSST likely affect how much moisture is advected. Further, global warming, which is often measured using global temperature anomalies (GlobalT) or global sea surface temperature anomalies (GlobalSSTA), has an influence on the global circulation, thereby influencing smaller scale circulations and weather phenomena (Kushner et al., 2001; Trenberth, 2011; Scher and Messori, 2019). As such, it is important to consider these as potential drivers of Australian precipitation, temperature and severe weather. Time series of these climate driver indices have been obtained from the National Oceanic and Atmospheric Administration's Earth System Research Laboratory (`http://www.esrl.noaa.gov/psd/gcos_wgsp/Timeseries/`), except for local sea surface temperatures that were obtained from the BoM (`http://www.bom.gov.au/climate/change/?ref=ftr#tabs=Tracker&tracker=timeseries`). These time series have been used in this project as inputs to statistical models to determine attributes of precipitation and temperature in regions of southeast Australia.

2.2 Trend Analysis

Although interannual variability of atmospheric phenomena is associated with drivers like ENSO and the DMI, larger changes in the atmospheric circulation due to global warming are also likely to cause changes in the frequency and variability of atmospheric phenomena (Scher and Messori, 2019). Modelling changes to severe weather under a changing climate is valuable for anticipating the changes that could occur. However, trend analysis also is important to measure any changes as they occur, including changes in the variables associated with severe weather, and to compare this against model projected changes. Continued monitoring of trends in severe weather (or lack thereof) is necessary as the response to global warming might vary by region and rates of warming. General trends and changes in variability may initially be understood by plotting time series data against their percentiles. The percentiles can be calculated over specific periods, such as a 30-year period which is representative of a climatology, or the entire time series. The selection of the period over which percentiles are calculated can influence how extreme an observation is perceived to be. In this project, the percentiles are calculated using the entire length of the time series in order to understand changes to precipitation or temperature extremes relative to the complete observation period rather than a particular climate window.

Following this, the data is split into multi-decade intervals to better understand potential trends in the data. Bootstrap resampling is applied with 5000 resamples to create a distribution of means of the climate variable over each period. This enabled box plots to be created that provide an overview of changes present in the data set. Two-sided permutation testing can be applied alongside bootstrap resampling to evaluate statistical significance of any differences between two periods of data, without assuming the underlying distribution. In permutation testing, the sampling distribution is built by permuting the observed data without replacement. The difference between the mean or variance of the two distributions is then tested, with the

null hypothesis being that there is no difference between the two distributions. For this project, a significance level of $\alpha = 0.1$ is considered a statistically significant result, with the direction of associated change determined by the changes in box plots between the two periods.

Wavelet analysis is a method of determining dominant modes of variability in a time series, so it can be used to detect potential climate drivers and analyse how the influence of these climate drivers changes over time. By stretching and translating the local base functions in space and time, the wavelet transform efficiently resolves both high- and low-frequency signals (Lau and Weng, 1995; Torrence and Compo, 1998). A number of local base functions could be used such as the *Mexican hat* wavelet or the *Paul* wavelet. In this project, the *Morlet* wavelet was used as it is one of the most commonly used continuous wavelets in the geophysics field. Additionally, it can detect both the time-dependent amplitude and phase for multiple frequencies in the time series (Lau and Weng, 1995). This is given by

$$\psi_0(\eta) = \pi^{-1/4} e^{i\omega_0 \eta} e^{-\eta^2/2}, \tag{2.2.1}$$

where η is a non-dimensional time parameter and ω_0 is the non-dimensional frequency. From this, the continuous wavelet transform of a time series, x_n, is

$$W_n(s) = \sum_{n'=0}^{N-1} x_{n'} \psi^* \left(\frac{(n'-n)\delta t}{s} \right), \tag{2.2.2}$$

where n is a localised time index, s is the wavelet scale, N is the number of points in a time series and $*$ is the complex conjugate. To improve computational efficiency, the wavelet transform is computed in Fourier space using the discrete Fourier transform of x_n

$$\hat{x}_k = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i k n/N}, \tag{2.2.3}$$

where $k = 0,\, 1,\, \cdots,\, N-1$ is the frequency index. The Fourier transform of $\psi(t/s)$ is given by $\hat{\psi}(s\omega)$. Using the convolution theorem, and taking the Fourier transform, the wavelet transform in Fourier space is given by

$$W_n(s) = \sum_{k=0}^{N-1} \hat{x}_k \hat{\psi}^*(s\omega_k) e^{-\omega_k n \delta t}, \tag{2.2.4}$$

where

$$\omega_k = \begin{cases} \frac{2\pi k}{N\delta t} & : \quad k \leq \frac{N}{2} \\ \\ \frac{-2\pi k}{N\delta t} & : \quad k > \frac{N}{2}. \end{cases} \tag{2.2.5}$$

The local wavelet power spectrum is then defined by

$$P(s) = |W_n(s)|^2, \tag{2.2.6}$$

which displays the time evolution of periodic signals within a time series. The global power spectrum also can be defined as

$$\overline{W}^2(s) = \frac{1}{N} \sum_{n=0}^{N-1} |W_n(s)|^2, \tag{2.2.7}$$

which provides an overview of the dominant amplitudes within the time series.

2.3 Statistical Attribution and Prediction

Wavelet analysis, which was introduced in Section 2.2, could be considered an initial step to attributing different climate drivers to severe weather. The local wavelet power spectrum provides an understanding of the time evolution of periodic signals in the time series, while the global power spectrum displays the dominant amplitudes in the time series. Together with knowledge of the periodicity of different climate drivers, likely attributes of severe weather can be determined. However, attributes are not explicitly determined using this method, as some climate drivers share similar periodicities, and depending on the length and frequency of data used in the wavelet analysis, certain climate drivers might not be considered. In particular are those climate drivers that vary on time scales less than a year, or multi-decadal climate drivers when considering a relatively short time series. As such, it is important to supplement this with other statistical approaches. Additionally, determination of the important environmental attributes of SCT hazards is not possible using wavelet analysis as the response data are categorical and not provided as a uniformly spaced time series. In this project, linear regression, logistic regression, and machine learning methods are also utilised to determine attributes of precipitation, temperature and SCT hazards.

2.3.1 Linear and Logistic Regression

Multiple linear regression (LR) has been used for meteorological and climatological purposes for decades, as it is a simple and computationally inexpensive model to fit p predictors to obtain a forecast model or determine attributes of certain aspects of climate. The multiple LR model is described by the equation

$$Y = \beta_0 + \sum_{i=1}^{p} \beta_i X_i + \epsilon, \tag{2.3.1}$$

where X_i is the ith predictor, β_i are the *regression coefficients*, marking the association between X_i and the response Y, β_0 is the intercept and ϵ represents the model error. A simple LR model is where $p = 1$; however, this is generally not applicable in meteorological settings as many variables (e.g., temperature, moisture, wind speed and direction) all play a role in different features of interest in the atmosphere. Since the regression coefficients are unknown, estimates must be made. The multiple LR model becomes

$$\hat{Y} = \hat{\beta}_0 + \sum_{i=1}^{p} \hat{\beta}_i X_i, \tag{2.3.2}$$

where $\hat{\beta}_i$ are selected to minimise the *residual sum of squares* (RSS):

$$\begin{aligned} RSS &= \sum_{i=1}^{p} \left(Y_i - \hat{Y}_i\right)^2 \\ &= \sum_{i=1}^{p} \left(Y_i - \beta_0 - \sum_{j=1}^{p} \hat{\beta}_j X_j\right)^2 \end{aligned} \tag{2.3.3}$$

Multiple LR is a classical approach of modelling for quantitative response variables. However, there are occasions where the response variable is qualitative. For example, in trying to understand key environmental attributes of SCT hazards, the classifications 'severe hail', 'non-severe storm', and 'no storm' can be used to develop models that attribute environmental variables to severe hail storms (as opposed to non-severe storms, or no storms). Since the regression coefficients are determined by minimising the RSS, the qualitative responses would need to be

converted to numerical values; however, there is no ideal conversion in this setting.

A more appropriate (and popular) model for classification problems is the *logistic regression* (LogR) model. Here, the probability that the response Y belongs to a particular class is predicted. In LogR, a binary prediction is fit to the logit link function as:

$$\ln\left(\frac{p_i}{1-p_i}\right) = \hat{\beta}_0 + \sum_{i=1}^{p} \hat{\beta}_i X_i, \tag{2.3.4}$$

where p_i represents the probability of one outcome, and $1-p_i$ represents the probability of the alternate outcome in the binary classification. This can be re-written as:

$$p_i = \frac{1}{1+\exp\left(-\hat{\beta}_0 - \sum_{i=1}^{p} \hat{\beta}_i X_i\right)}. \tag{2.3.5}$$

In order to determine the regression coefficients $\hat{\beta}_i$, the log-likelihood function is maximised

$$L(\hat{\beta}) = \sum_{i=1}^{n} y_i \left(\hat{\beta}_0 + \sum_{j=1}^{p} \hat{\beta}_j X_j\right) - \ln\left[1+\exp\left(\hat{\beta}_0 + \sum_{j=1}^{p} \hat{\beta}_j X_j\right)\right], \tag{2.3.6}$$

where y_i is the binary predictand, based on the probability of p_i. The log-likelihood function can be maximised using a variety of methods, such as the Newton-Raphson algorithm (Hastie et al., 2009). Logistic regression can be extended from a binary problem to one that classifies the response variable with more than two classes. However, the approach is simplified in this project as the focus is on determining whether an environment supports an SCT hazard or it does not. In order to attribute SCT hazards, the two classes of non-severe storm and no storm are combined so that the binary classification problem is to classify 'severe hazard' (e.g., severe hail) versus 'non-severe storm or no storm'.

2.3.2 Support Vector Machines

Support vector machines (SVMs) are a type of machine learning method that can be applied to both classification and regression problems. This method aims to determine an optimal separating hyperplane between classes. Given a training data set x, a hyperplane can be defined by

$$\{X : f(x) = x^T\beta + \beta_0 = 0\}, \tag{2.3.7}$$

where β is a unit vector. The hyperplane that creates the largest margin, M, between the classes can be determined from the optimisation problem

$$\begin{aligned} &\max_{\beta,\ \beta_0,\ \|\beta\|=1} \ \mathrm{M} \\ &\text{subject to } y_i(x_i^T\beta + \beta_0) \geq M,\ i = 1,\cdots,N. \end{aligned} \tag{2.3.8}$$

However, it is often the case that the classes are not linearly separable. In this situation, slack variables can be introduced $\xi = (\xi_1, \xi_2, \cdots, \xi_N)$ and the constraint can be redefined as

$$y_i(x_i^T\beta + \beta_0) \geq M(1-\xi_i), \tag{2.3.9}$$

where $\xi_i \geq 0\ \forall i$ and $\sum_{i=1}^{N} \xi_i \leq K$, where K is some constant that limits the amount of overlap from data points on the wrong side of the separating hyperplane. By letting M= $1/\|\beta\|$, the

optimisation problem becomes

$$\begin{aligned} &\min_{\beta,\,\beta_0} \|\beta\| \\ &\text{subject to } y_i(x_i^T\beta + \beta_0) \geq 1 - \xi_i \;\forall i \\ &\xi_i \geq 0,\; \sum \xi_i \leq K. \end{aligned} \tag{2.3.10}$$

This can be re-written as a quadratic optimisation problem:

$$\begin{aligned} &\min_{\beta,\,\beta_0} \frac{1}{2}\|\beta\|^2 + C\sum_{i=1}^{N}\xi_i \\ &\text{subject to } \xi_i \geq 0,\; y_i(x_i^T\beta + \beta_0) \geq 1 - \xi_i \;\forall i, \end{aligned} \tag{2.3.11}$$

where C is the *cost* parameter, replacing K. To solve the problem, consider the Lagrange function

$$L(\beta, \lambda, \mu) = \frac{1}{2}\|\beta\|^2 + C\sum_{i=1}^{N}\xi_i + \sum_{i=1}^{N}\lambda_i\left[y_i\left(x_i^T\beta + \beta_0\right) - (1 - \xi_i)\right] - \sum_{i=1}^{N}\mu_i\xi_i. \tag{2.3.12}$$

By computing the partial derivatives with respect to β, β_0 and ξ_i, and setting them to zero, the following optimality conditions are obtained

$$\frac{\partial L}{\partial \beta} = \beta - \sum_{i=1}^{N}\lambda_i y_i x_i = 0, \tag{2.3.13}$$

$$\frac{\partial L}{\partial \beta_0} = -\sum_{i=1}^{N}\lambda_i y_i = 0, \tag{2.3.14}$$

$$\frac{\partial L}{\partial \xi_i} = C - \lambda_i - \mu_i = 0. \tag{2.3.15}$$

Substitution of equations (2.3.13)-(2.3.15) into (2.3.12) obtains the dual formulation

$$\begin{aligned} &\max L(\lambda) = \sum_{i=1}^{N}\lambda_i - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}\lambda_i\lambda_j y_i y_j x_i^2 x_j \\ &\text{subject to } 0 \leq \lambda_i \leq C \\ &\sum_{i=1}^{N}\lambda_i y_i = 0, \end{aligned} \tag{2.3.16}$$

and the solution for β has the form

$$\hat{\beta} = \sum_{i=1}^{N}\hat{\lambda}_i y_i x_i. \tag{2.3.17}$$

Those observations i for which $\lambda_i > 0$ are called *support vectors* as they define β. By solving the quadratic optimisation problem (2.3.16), values of λ are determined which can be used to determine β and β_0 and hence the classification function:

$$G(x) = sign\left[x^T\hat{\beta} + \hat{\beta}_0\right], \tag{2.3.18}$$

where the cost parameter C is tuned by the user.

Unfortunately, most data sets are not linearly separable, and while the cost constraint provides some flexibility in defining a separating hyperplane, more flexibility may be obtained by enlarging the feature space. This can be done using basis expansions, that map the data set into a higher dimensional space that is usually linearly separable. Often, the linear boundary in this higher dimensional space is non-linear in the original space. Let $h(x) = (h_1(x_i),\ h_2(x_i),\ \cdots,\ h_M(x_i))$, for $i = 1,\ \cdots,\ N$ and consider instead the non-linear function $f(x) = h(x)^T + \hat{\beta} + \hat{\beta}_0$. The Lagrange function of the dual formulation in (2.3.16) becomes

$$L(\lambda) = \sum_{i=1}^{N} \lambda_i - \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \lambda_i \lambda_j y_i y_j \langle h(x_i), h(x_j) \rangle, \tag{2.3.19}$$

and $f(x) = \sum_{i=1}^{N} \lambda_i y_i \langle h(x), h(x_i) \rangle + \beta_0$. Specification of $h(x)$ is not required. Instead, all that is necessary is the *kernel function*, $K(x, x') = \langle h(x), h(x') \rangle$, that computes inner products in the higher dimensional space. Popular kernel functions that are frequently considered by practitioners include

$$\begin{aligned} \text{d-Degree polynomial: } & K(x, x') = (1 + \langle x, x' \rangle)^d, \\ \text{Radial basis: } & K(x, x') = \exp\left(-\gamma \|x - x'\|^2\right), \\ \text{Hyperbolic tangent: } & K(x, x') = \tanh\left(\kappa_1 \langle x, x' \langle + \kappa_2\right). \end{aligned} \tag{2.3.20}$$

For this project, the d-Degree polynomial (Poly) and radial basis functions (RBF) are considered. The cost and γ of both kernels, and the degree of the polynomial in the Poly kernel are all parameters that are tuned. The degree of the polynomial and the γ parameter both influence how strongly the model fits to the training data, while the cost parameter relates to the cost of violating constraints in the optimisation problem. These free parameters are often selected using a grid search (Hsu et al., 2003).

2.3.3 Random Forests

Decision trees form the basis of random forests (RFs). Here, the predictor space X is partitioned into regions R. In regression trees, for each observation in region R_j, the mean of the responses for training observations in R_j is used as the prediction. The regions are selected to divide the predictor space into high-dimensional boxes, with a goal of minimising the RSS:

$$\sum_{j=1}^{J} \sum_{i \in R_j} (y_i - \hat{y}_{R_j})^2, \tag{2.3.21}$$

where $\hat{y}_{R_j}$ is the mean response for training observations in the jth box. This is done by recursive binary splitting. At each step, all predictors not yet a part of the tree, and all possible cutpoints splitting the predictor space into regions $\{X | X_j < s\}$ and $\{X | X_j \geq s\}$, are considered. The predictor X_j and cutpoint s leading to the greatest reduction in the RSS is selected. In the next step, a similar process occurs, but the splitting is over one of the two regions rather than the entire predictor space. This generally continues until a pre-specified maximum number of observations exist in each region.

For classification trees, it is assumed that each observation belongs to the most commonly occurring class of training observations in the region that it belongs. To select the regions that

partition the space, a different measure to RSS is required, since the data are qualitative. The *Gini index* is one of two commonly used metrics in growing a classification tree, given by

$$G = \sum_{k=1}^{K} \hat{p}_{mk}(1 - \hat{p}_{mk}), \tag{2.3.22}$$

where $\hat{p}_{mk}$ is the proportion of training observations in the mth region from the kth class, and K is the total number of classes. *Entropy* is the other common alternative, given by

$$D = -\sum_{k=1}^{K} \hat{p}_{mk} \log \hat{p}_{mk}. \tag{2.3.23}$$

Both the Gini index and entropy will take smaller values the more a node contains observations from a single class. Therefore, selection of the mth node is based on which splitting rule leads to the lowest Gini index or entropy. Unfortunately, decision trees created alone suffer from high variance, which leads to low prediction accuracy on a test data set.

Random forests improve on this variance problem through developing a large number of decorrelated trees. Bootstrap resampling is first applied to develop numerous training sets from the population, and a decision tree is fitted to each training set. When growing the tree, at each splitting step, randomly select m of p variables as candidates for the split variable, with $m \leq p$. Often $m = \sqrt{p}$ is used for classification problems and $m = p/3$ for regression problems. Let the set of trees grown be $\{T_n\}_{n=1}^{N}$, where N is the number of trees grown. Then, for regression, the prediction is taken as the average over all trees: $\hat{f}_{rf}(x) = \frac{1}{N}\sum_{n=1}^{N} T_n(x)$. For classification, the prediction is taken as the majority vote among the trees: $\hat{C}_{rf}(x) = majority\{\hat{C}_n(x)\}_{n=1}^{N}$.

2.4 The Weather Research and Forecasting Model

The Weather Research and Forecasting (WRF) model is a numerical model developed by the National Center for Atmospheric Research, National Centers for Environmental Prediction, Earth System Research Laboratory, United States Air Force, Naval Research Laboratory, University of Oklahoma and Federal Aviation Administration. It was developed for both research and numerical weather prediction purposes, to simulate the Earth system, including the atmosphere, on a range of spatial and temporal scales. WRF version 4.0.2 was used in the present study with the Advanced Research WRF (ARW) dynamical core to simulate severe convective storm events (Skamarock et al., 2019). The ARW solver integrates the compressible, non-hydrostatic Euler equations with a modified terrain-following hydrostatic pressure vertical coordinate. Let η denote a more traditional terrain-following hydrostatic pressure vertical coordinate be defined by

$$\eta = \frac{p_d - p_t}{p_s - p_t}, \tag{2.4.1}$$

where p_d is the hydrostatic component of dry air, and both p_s and p_t are the hydrostatic components of dry air at the surface and top, respectively. Define p_d as

$$p_d = B(\eta)(p_s - p_t) + [\eta - B(\eta)](p_0 - p_t) + p_t, \tag{2.4.2}$$

where p_0 is a reference pressure at sea-level and $B(\eta)$ is given by the third order polynomial

$$B(\eta) = c_1 + c_2\eta + c_3\eta^2 + c_4\eta^3 \tag{2.4.3}$$

subject to $B(1) = 1$, $B_\eta(1) = 1$, $B(\eta_c) = 0$ and $B_\eta(\eta_c) = 0$, where a subscript η denotes differentiation. Furthermore, let

$$c_1 = \frac{2\eta_c^2}{(1-\eta_c)^3},\ c_2 = \frac{-\eta_c(4+\eta_c+\eta_c^3)}{(1-\eta_c)^3},\ c_3 = \frac{2(1+\eta_c+\eta_c^2)}{(1-\eta_c)^3},\ c_4 = \frac{-(1+\eta_c)}{(1-\eta_c)^2}, \tag{2.4.4}$$

with η_c being the value of η where it becomes a pressure coordinate. The modified vertical coordinate metric is defined as

$$\mu_d = \frac{\partial p_d}{\partial \eta} = B_\eta(\eta)(p_s - p_t) + [1 - B_\eta(\eta)](p_0 - p_t). \tag{2.4.5}$$

The prognostic variables are the three wind components $\mathbf{v} = (u, v, w)$, contravariant vertical velocity $\omega = \dot{\eta}$, moist potential temperature $\theta_m = \theta(1 + 1.61q_v)$ (where θ is potential temperature) and moisture mixing ratios q_m=q_v, q_c, q_r, q_i, $\cdots$ for water vapour, cloud water, rain water, ice, $\cdots$. These are modified by the vertical coordinate metric as

$$\mathbf{V} = \mu_d \mathbf{v} = (U, V, W), \quad \Omega = \mu_d \omega, \quad \Theta_m = \mu_d \theta_m, \quad Q_m = \mu_d q_m. \tag{2.4.6}$$

Additionally, geopotential $\phi = gz$, where g is gravitational acceleration. However, it is not written in flux form as $\mu_d\phi$ is not conserved. The governing equations for velocity, moist potential temperature, moisture, the vertical coordinate metric and geopotential are thus:

$$\partial_t U + \nabla \cdot \mathbf{V}u + \mu_d \alpha \partial_x p + (\alpha/\alpha_d)\, \partial_\eta p \partial_x \phi = F_U, \tag{2.4.7}$$

$$\partial_t V + \nabla \cdot \mathbf{V}v + \mu_d \alpha \partial_y p + (\alpha/\alpha_d)\, \partial_\eta p \partial_y \phi = F_V, \tag{2.4.8}$$

$$\partial_t W + \nabla \cdot \mathbf{V}w - g\left[(\alpha/\alpha_d)\, \partial_\eta p - \mu_d\right] = F_W, \tag{2.4.9}$$

$$\partial_t \Theta_m + \nabla \cdot \mathbf{V}\theta_m = F_{\Theta_m}, \tag{2.4.10}$$

$$\partial_t \mu_d + \nabla \cdot \mathbf{V} = 0, \tag{2.4.11}$$

$$\partial_t \phi + \mu_d^{-1}\left[(\mathbf{V} \cdot \nabla\phi) - gW\right] = 0, \tag{2.4.12}$$

$$\partial_t Q_m + \nabla \cdot \mathbf{V}q_m = F_{Q_m}, \tag{2.4.13}$$

where α_d is the inverse density of dry air and α is the inverse density of the full parcel. Subscripts x, y and η denote differentiation, and for some variable a

$$\nabla \cdot \mathbf{V}a = \partial_x(Ua) + \partial_y(Va) + \partial_\eta(\Omega a),$$

and

$$\mathbf{V} \cdot \nabla a = U\partial_x a + V\partial_y a + \Omega\partial_\eta a.$$

The right hand side of the governing equations represent forcing terms from model physics, turbulent mixing, spherical projections and Earth's rotation. The diagnostic equation for dry hydrostatic pressure is

$$\partial_\eta \phi = -\alpha_d \mu_d, \tag{2.4.14}$$

and the diagnostic equation for full pressure is

$$p = p_0 \left(\frac{R_d \theta_m}{p_0 \alpha_d}\right)^\gamma, \tag{2.4.15}$$

where $R_d \approx 287.04$ kg^{-1} K^{-1} is the gas constant for dry air, $\gamma = 1.4$ is the ratio of the heat capacity for dry air and p_0 is a reference surface pressure (usually 10^5 Pa).

There are four different projections to the sphere that are supported in ARW. The isotropic projections (where $(\Delta x/\Delta y)_{\text{Earth}} =$ constant everywhere) — Lambert conformal, polar stereographic, and Mercator; and the anisotropic latitude-longitude projection, which can be extended to a global model. To apply these projections, map-scale factors are introduced

$$(m_x, m_y) = \frac{(\Delta x, \Delta y)}{\text{distance on Earth}} \tag{2.4.16}$$

that allow the momentum variables to be redefined as

$$U = \mu_d u/m_y,\ V = \mu_d v/m_x,\ W = \mu_d w/m_y,\ \Omega = \mu_d \omega/m_y. \tag{2.4.17}$$

These can then be substituted into the governing equations to obtain the equations projected onto the sphere. The Mercator projection is used in this project since the case study performed occurs in the mid-latitudes. This projection (along with the other two isotropic projections) defines the map-scale factors to be equal in both horizontal directions, i.e. $m = m_x = m_y$, with m only varying by latitude.

Finally, the governing equations are recast once more using perturbation variables, as this reduces truncation errors in the horizontal pressure gradient calculations and machine rounding errors in the vertical pressure gradient and buoyancy calculations. To do this, reference state variables $\bar{a}(\bar{z})$ are defined as a function of height only, and to satisfy the governing equations for an atmosphere at rest. The perturbation variables are denoted by a prime symbol, and we redefine $p = \bar{p}(\bar{z}) + p'$, $\phi = \bar{\phi}(\bar{z}) + \phi'$, $\alpha = \bar{\alpha}(\bar{z}) + \alpha'$ and $\mu_d = \bar{\mu}_d(x, y) + \mu'_d$. Further information on the map projections and recasting of the governing equations can be found in Sections 2.3 and 2.4 of Skamarock et al. (2019), respectively.

There are a number of terms in the governing equations that are meteorologically insignificant, but high frequency (acoustic) and therefore propagate fast through the domain and can introduce numerical instability in the model. This is handled for the flux-form equations using the time-splitting method described in Klemp et al. (2007). In this approach, the meteorologically significant modes are integrated using a Runge-Kutta (RK3) time integration scheme, while the acoustic modes are integrated over a smaller time step. Those acoustic modes and gravity waves that propagate in the horizontal are integrated using a forward-backward time integration scheme, while the vertically propagating acoustic modes and buoyancy oscillations are integrated with a vertically implicit scheme. Given the prognostic variables $\Phi = (U, V, W, \Theta_m, \phi', \mu'_d, Q_m)$ and the model equations defined as $\Phi_t = R(\Phi)$, the solution $\Phi(t)$ is advanced forward to $\Phi(t + \Delta t)$ using the three steps:

$$\Phi^* = \Phi^t + \frac{\Delta t}{3} R\left(\Phi^t\right), \tag{2.4.18}$$

$$\Phi^{**} = \Phi^t + \frac{\Delta t}{2} R\left(\Phi^*\right), \tag{2.4.19}$$

$$\Phi^{t+\Delta t} = \Phi^t + \Delta t R\left(\Phi^{**}\right), \tag{2.4.20}$$

where Δt is the model time step. It is important to note that this is actually not a true Runge-Kutta scheme, as it is third-order accurate for linear equations but only second-order accurate for nonlinear equations.

As discussed above, the use of a smaller time step to integrate the acoustic modes reduces the restrictions that would be imposed on the model to maintain numerical stability. To increase the accuracy of this approach, a perturbation form of the governing equations is used

with smaller acoustic time steps within the RK3 scheme. Denote the most recent RK3 predictor by the superscript t^*, representing either Φ^t, Φ^* or Φ^{**}, then the deviations of each variable from the RK3 predictor are

$$\mathbf{V}'' = \mathbf{V} - \mathbf{V}^{t^*}, \qquad \Omega'' = \Omega - \Omega^{t^*}, \qquad \Theta''_m = \Theta_m - \Theta^{t^*}_m,$$
$$\phi'' = \phi' - \phi'^{t^*}, \qquad \alpha''_d = \alpha'_d - \alpha'^{t^*}_d, \qquad \mu''_d = \mu'_d - \mu'^{t^*}_d.$$

This redefines the hydrostatic relation as

$$\alpha''_d = \frac{-1}{\mu^{t^*}_d}\left(\partial_\eta \phi'' + \alpha^{t^*}_d \mu''_d\right). \tag{2.4.21}$$

The equation of state may be linearised about t^*:

$$p'' = \frac{c_s^2}{\alpha^{t^*}_d}\left(\frac{\Theta''_m}{\Theta^{t^*}_m} - \frac{\alpha''_d}{\alpha^{t^*}_d} - \frac{\mu''_d}{\mu^{t^*}_d}\right), \tag{2.4.22}$$

where $c_s^2 = \gamma p^{t^*} \alpha^{t^*}_d$ is the square speed of sound. Additionally, given Equations 2.4.21 and 2.4.22, the vertical pressure gradient can be defined as

$$\partial_\eta p'' = \partial_\eta (C \partial_\eta \phi'') + \partial_\eta \left(\frac{c_s^2}{\alpha^{t^*}_d}\frac{\Theta''_m}{\Theta^{t^*}_m}\right), \tag{2.4.23}$$

where $C = c_s^2 / \left(\mu^{t^*}_d \alpha^{t^{*2}}_d\right)$. These variables and Equation 2.4.23 can be substituted into the governing equations to obtain the equations on the acoustic time-step, $\Delta\tau$. The discrete acoustic time-step operator is

$$\delta_\tau a = \frac{a^{\tau+\Delta\tau} - a^\tau}{\Delta\tau}, \tag{2.4.24}$$

and terms averaged in time over $\Delta\tau$ are slightly foward centred with the averaging operator

$$\bar{a}^\tau = \frac{1+\beta}{2} a^{\tau+\Delta\tau} + \frac{1-\beta}{2} a^\tau, \tag{2.4.25}$$

where β is user-specified.

To integrate over the acoustic time steps, U'' and V'' are stepped forward to $\tau + \Delta\tau$. Then μ_d is advanced to compute $\Omega''^{\tau+\Delta\tau}$, and Θ_m is advanced. W and ϕ are the next two state variables to be advanced and finally p'' and α'' are diagnosed using Equations 2.4.22 and 2.4.21, respectively. The acoustic time step $\Delta\tau$ is specified by the user, except for the first RK3 substep where $\tau = \Delta t/3$ is used to advance the solution regardless of the user-specified acoustic time step.

Arakawa C-grid staggering is used to discretise the domain, with normal velocities staggered one-half grid length from thermodynamic variables. *Mass* point variables include θ, q_i and the coordinate metric μ. The velocities u, v and w are on u points, v points and w points, respectively; and geopotential ϕ is defined on w points. The diagnostic variables, pressure p and inverse density α, are also defined at mass points. While the horizontal grid lengths Δx and Δy are constant, the vertical grid length $\Delta\eta$ is not constant. The η values may either be user-specified, or determined by the model, with the condition that $\eta = 1$ at the model surface and $\eta = 0$ at the model top. Generally, more η-levels are defined closer to the model surface, so that the boundary layer is better resolved, with vertical grid spacing increasing towards the model top.

The advection of momentum, scalars and geopotential is performed during the RK3 time integration. Given the flux divergence for a scalar q in discrete form:

$$R^{t^*}_{q_{\text{adv}}} = -m_x m_y \left[\delta_x(U\bar{q}^{x_{\text{adv}}}) + \delta_y(V\bar{q}^{y_{\text{adv}}})\right] - m_y\delta_\eta(\Omega\bar{q}^{\eta_{\text{adv}}}), \tag{2.4.26}$$

where the discrete operator is defined as

$$\delta_x(U\bar{q}^{x_{\text{adv}}}) = \Delta x^{-1}\left[(U\bar{q}^{x_{\text{adv}}})_{i+1/2} - (U\bar{q}^{x_{\text{adv}}})_{i-1/2}\right]. \tag{2.4.27}$$

The ARW offers 2^{nd} to 6^{th} order accurate spatial discretisations of the flux divergence. For the even orders:

$$2^{nd}\text{:}\quad (\bar{q}^{x_{\text{adv}}})_{i-1/2} = \frac{1}{2}(q_i + q_{i-1}), \tag{2.4.28}$$

$$4^{th}\text{:}\quad (\bar{q}^{x_{\text{adv}}})_{i-1/2} = \frac{7}{12}(q_i + q_{i-1}) - \frac{1}{12}(q_{i+1} + q_{i-2}), \tag{2.4.29}$$

$$6^{th}\text{:}\quad (\bar{q}^{x_{\text{adv}}})_{i-1/2} = \frac{37}{60}(q_i + q_{i-1}) - \frac{2}{15}(q_{i+1} + q_{i-2}) + \frac{1}{60}(q_{i+2} + q_{i-3}), \tag{2.4.30}$$

and the odd orders:

$$3^{rd}\text{:}\quad (\bar{q}^{x_{\text{adv}}})_{i-1/2} = (\bar{q}^{x_{\text{adv}}})^{4th}_{i-1/2} + \text{sign}(U) + \frac{1}{12}\left[(q_{i+1} - q_{i-2}) - 3\left(q_i - q_{i-1}\right)\right], \tag{2.4.31}$$

$$5^{th}\text{:}\quad (\bar{q}^{x_{\text{adv}}})_{i-1/2} = (\bar{q}^{x_{\text{adv}}})^{6th}_{i-1/2} - \text{sign}(U)\frac{1}{60}\left[(q_{i+2} - q_{i-3}) - 5(q_{i+1} - q_{i-2}) + 10(q_i - q_{i-1})\right]. \tag{2.4.32}$$

It can be seen that the even-ordered schemes are spatially centred, with no implicit diffusion outside of that in the RK3 time integration, while the odd-ordered schemes are upwind-biased and spatial discretisation is diffusive. It is worth noting also, that mixing ratios of moisture, chemical and tracer species should be positive definite, and there should be shape-preservation of mixing ratios (transport should not introduce new maxima/minima). This is treated in the ARW using a positive-definite renormalisation process as in Skamarock and Weisman (2009), and the shape-preserving extension of Wang et al. (2009). Advection is 5th order in the horizontal and 3rd order in the vertical for all simulations in this project.

The Courant number limits both the model and acoustic time steps. RK3 is limited by the advective courant number $\frac{u\Delta t}{\Delta x}$, and the choice of advective scheme. For three-dimensional applications, the general rule of thumb is to choose a time step so that:

$$\Delta t < 0.25\frac{\text{Cr}_{\text{theory}}}{\sqrt{3}} \cdot \frac{\Delta x}{u_{\max}},$$

where $\text{CR}_{\text{theory}}$ is the theoretical Courant number defined in Table 3.1 of Skamarock et al. (2019), and $u_{\max}$ is the expected maximum velocity in the simulation. The acoustic time step should satisfy

$$\Delta\tau < \frac{1}{2} \cdot \frac{\Delta x}{c_s}.$$

Vertically-propagating gravity waves can reflect off the upper model boundary, which is unphysical, and there are three formulations that can be used to absorb the gravity waves. In the simulations presented for this project, an implicit Rayleigh damping is applied, that damps vertical velocity in the implicit solution for vertically propagating acoustic modes. This is done by subtracting $\tau(z)\Delta\tau W''^{\tau+\Delta\tau}$ from the solution to the tridiagonal equation in the vertical

direction for $W''^{\tau+\Delta\tau}$, prior to updating the geopotential ϕ. After this, the geopotential is updated using the damped vertical velocity solution. The vertical structure of the damping layer is given by $\tau(z)$

$$\tau(z) = \begin{cases} \gamma_r \sin^2\left[\frac{\pi}{2}\left(1 - \frac{z_{\text{top}}-z}{z_d}\right)\right] & \text{for } z \geq (z_{\text{top}} - z_d) \\ 0 & \text{otherwise,} \end{cases} \tag{2.4.33}$$

where γ_r is a user-specified damping coefficient, z_{top} is the height of the model top for a particular grid column and z_d is the damping layer.

Given initial condition data, that is typically provided in GRIB format, there are pre-processing programs that provide ARW with input data that is correctly staggered, has a hydrostatically balanced reference state and perturbation fields, and metadata about the date, grid characteristics and projection. The first step creates a physical grid with specific projection type, location on Earth, number of grid points, nest locations and grid distances, and interpolates a number of static fields such as topography and land-use, to the prescribed domain. Then, the externally analysed or forecast data is decoded from GRIB format to an internal binary format that the pre-processing software can interpolate to the physical grid. The output from this pre-processing software is the three-dimensional state of the atmosphere on the user-specified model grid, at selected times. Data that is required for simulations to be performed include three-dimensional and surface temperature (K), relative humidity (%), geopotential height (m), pressure (Pa), and the horizontal components of wind speed (m s^{-1}). Two-dimensional static fields that are required include information about albedo, Coriolis parameters, terrain, vegetation/land use, a land/water mask, map scale factors, map rotation angle, soil texture category, vegetation greenness fraction, annual mean temperature and latitude/longitude. Finally, two-dimensional, time-dependent fields that must be provided are surface pressure and mean sea-level pressure (Pa), layers of soil temperature (K) and soil moisture (kg/kg), snow depth (m), skin temperature (K), sea surface temperature (K) and a sea/ice flag.

The final piece of pre-processing software takes the interpolated meteorological data, and static fields, and creates initial condition files for each nest in the simulation, and a boundary condition file for the coarse resolution outer domain. The lateral boundary condition file includes the variables u, v, θ, q_v, ϕ', and μ'_d and is used to constrain the lateral boundaries. Each variable has a valid lateral boundary value at the initial time and a tendency term to get to the next boundary time period. This is user-specified, based on the availability of data. For example, if the data provided is 3-hourly, then the first time period of the lateral boundary u would contain data for u coupled with μ_d and the inverse map factor m^x at 0 h

$$U_{0h} = \left.\frac{\bar{\mu}_d^x u}{\bar{m}^x}\right|_{0h},$$

and the tendency is given by

$$U_t = \frac{U_{3h} - U_{0h}}{3h}.$$

Note that only horizontal momentum is coupled with μ_d and the inverse map factor, while other fields are coupled with just μ_d, and μ'_d is not coupled at all. The lateral boundary fields are defined along all four sides of the horizontal grid, and the user can specify the width of the lateral boundary.

The higher the resolution of the model, the more computationally intensive it is to run. Often, there might be a region of interest within a domain that the user would prefer a higher

resolution. However, the resolution might be too high to feasibly run the simulation over the entire domain that is required. Horizontal nesting can be used to enhance the resolution of the model over a particular region of interest by introducing additional grids that are completely contained within the coarser resolution *parent* domain. Nesting can either be one-way, where information is only exchanged from the parent to the *child* domain; or it can be two-way, where the solution from the finer resolution child replaces those solutions of the parent for grid points that are contained in the child domain. The lateral boundary conditions for nested grids are interpolated from the parent domain; however, there is no relaxation zone in the nested domain. Multiple levels of nesting can be used, and each parent can contain one or more child grids, including multiple children that share the same parent. However, overlapping grids cannot be considered. The ratio of parent horizontal grid distance to child horizontal grid distance is required to be an integer (similarly for time step ratios), and a general rule of thumb is an odd integer, typically 3 or 5. This rule of thumb may be understood by considering the Arakawa-C grid staggering, where coincident points exist between the child and parent grids when an odd integer grid ratio is used. This is particularly important for masked fields. For example, it does not make sense to average the soil temperature of neighbouring locations on the child grid if the coarse grid point for which this value is fed back to is considered a water grid point.

There are a number of different physics options within ARW to describe radiation, cumulus parameterisation, the planetary boundary layer (PBL), land-surface, and microphysics. The physics options are user-specified and can have a noticeable effect on simulation results (e.g., Adlerman and Droegemeier, 2002; Dawson II et al., 2010; Labriola et al., 2019). Tendencies for potential temperature, moisture and velocity components are computed after the first RK3 step, prior to the acoustic steps, within the radiation, surface, PBL and cumulus schemes. These tendencies are fixed for the remaining RK3 steps, and microphysics is computed at the end of the RK3 steps so saturation conditions at the end of the time-step are correct.

Physics parameterisations interact with each other via model state variables and their tendencies. The radiation or PBL schemes do not need to be called at the same frequency as each other or the model time step. In particular, radiation schemes are too computationally expensive to be called every time step, so their tendencies are kept constant between calls. The radiation scheme provides temperature tendencies due to radiative flux divergence and surface downward longwave and shortwave radiation for the ground heat budget. It is called first as the fluxes are required for the land-surface model (LSM). The surface-layer scheme is called next to calculate friction velocities and exchange coefficients for the LSM and PBL schemes. The LSM also requires precipitation forcing from the microphysics (and cumulus) parameterisation(s), but this is taken from the previous time step as the microphysics scheme is called last. The LSM takes radiative and precipitation forcings, along with information from the surface-layer scheme, land state variables and land-surface properties to calculate heat and moisture fluxes over both land and sea-ice points. These fluxes are then provided to the PBL scheme as a lower boundary condition on vertical transport. The PBL scheme calculates vertical sub-grid-scale fluxes due to eddy transports within the whole atmospheric column, determines the flux profiles in the well-mixed boundary layer and stable layer, and provides tendencies of temperature, moisture and horizontal momentum. Finally, the microphysics scheme is called, which resolves water vapour, cloud and precipitation processes. It is called at the end of the time step to directly update state variables, because condensation adjustments made at this time ensure an accurate saturation balance. For the simulations presented in this project, radiation, PBL and land-surface physics packages were not included. There is the option to include cumulus parameterisation; however, this is not used as the grid resolution is high enough for deep convection to be simulated. Two microphysics schemes were considered in this project: the

Morrison microphysics scheme (Morrison et al., 2005, 2009), and the National Severe Storms Laboratory double-moment 4-ice scheme (hereafter labelled the NSSL scheme; Mansell et al., 2010).

2.4.1 Microphysics

As mentioned in Section 2.4, the physics schemes that are used in a simulation affect the results. This is particularly noticeable when varying microphysics schemes in studies of SCTs, as they describe the various cloud-scale processes that occur, such as evaporation and melting of hydrometeors. Due to the potential influence of microphysics schemes on simulation results of severe convection, both the Morrison and NSSL schemes have been considered in this project.

Morrison Microphysics
The Morrison microphysics scheme is a *double-moment* microphysics scheme, meaning it predicts two moments of the hydrometeor spectra. Generally, double-moment microphysics schemes predict the number concentration and mass mixing ratio of hydrometeor species, providing a more robust treatment of the particle size distribution. In the Morrison scheme, the mass mixing ratio q_χ and number concentration N_χ are predicted for $\chi = i$ (cloud ice), s (snow), r (rain) and g (graupel/hail), while only the mass mixing ratio is calculated for c (cloud droplets) and v (water vapour). The user is free to select whether graupel or hail is used for the heavy rimed ice category. In the simulations presented for this project, hail has been selected. The particle size distribution for the hydrometeor class χ is represented by the gamma function

$$n_\chi(D) = N_{0_\chi} D^{\mu_\chi} e^{-\lambda_\chi D}, \tag{2.4.34}$$

where N_{0_χ}, λ_χ and μ_χ are the intercept, slope and shape parameters, and D is the particle diameter. N_{0_χ} and λ_χ are obtained from the predicted N_χ and q_χ, and specified μ_χ for each species from the following equations:

$$\lambda_\chi = \left[\frac{c_\chi N_\chi \Gamma(\mu_x + d_\chi + 1)}{q_\chi \Gamma(\mu_x + 1)} \right]^{\frac{1}{d_\chi}}, \tag{2.4.35}$$

and

$$N_{0_\chi} = \frac{N_\chi \lambda_\chi^{\mu_\chi + 1}}{\Gamma(\mu_\chi + 1)}, \tag{2.4.36}$$

where Γ represents the gamma function, c_χ and d_χ are obtained from the power-law mass-diameter relationship for each hydrometeor species $m_\chi = c_\chi D_\chi^d$. For precipitation (rain, snow and graupel/hail) and cloud ice, $\mu_\chi = 0$, while for cloud droplets μ_c is a function of the predicted droplet number concentration, following Martin et al. (1994). There are upper and lower bounds on values of λ_χ, for each hydrometeor species, specified in Morrison et al. (2005), with the number concentration adjusted so λ_χ does not exceed these values.

For simplicity, all particles are assumed to be spherical, and bulk particle density for ice species is given by Reisner et al. (1998). The moment-weighted terminal fall velocities of each hydrometeor species is computed from the velocity-diameter relation

$$V_\chi(D) = \frac{\rho_{\text{sur}} a_\chi D^{b_\chi}}{\rho}, \tag{2.4.37}$$

where ρ_{sur} is air density at sea level, and a_χ and b_χ are empirically determined for each species. A range of microphysical processes are considered in the Morrison microphysics scheme including accretion, autoconversion, deposition, evaporation, ice nucleation, ice splintering, melting,

riming, sedimentation and sublimation, with details of these processes in Appendix A of Morrison et al. (2005).

NSSL Microphysics

The NSSL microphysics scheme also is a double-moment scheme (Mansell et al., 2010), adapting the microphysics scheme of Ziegler (1985) and the distribution functions of Seifert and Beheng (2006). Here, the mass mixing ratio and number concentration is predicted for cloud droplets, cloud ice, rain, snow, graupel and hail (note: it predicts both rimed ice categories, rather than one as in the Morrison microphysics scheme). The particle size distribution, in terms of diameter, is given by

$$n_\chi(D) = 3A_\chi \frac{D^{3\nu_\chi+2}}{\bar{D}^{3\nu_\chi+3}} \exp\left[-B_\chi \left(\frac{D}{\bar{D}}\right)^{3\mu_\chi}\right], \tag{2.4.38}$$

where ν_χ and μ_χ are constants based on hydrometeor species. The coefficient A is expressed as a function of the number density

$$A_\chi = \frac{\mu_\chi N_{t_\chi}}{\Gamma\left(\frac{\nu_\chi+1}{\mu_\chi}\right)} B^{(\nu_\chi+1)/\mu_\chi}, \tag{2.4.39}$$

where N_{t_χ} is the number density, and the coefficient B is expressed as

$$B_\chi = \left[\frac{\Gamma\left(\frac{\nu_\chi+1}{\mu_\chi}\right)}{\Gamma\left(\frac{\nu_\chi+2}{\mu_\chi}\right)}\right]^{-\mu_\chi}. \tag{2.4.40}$$

All species have a non-zero terminal fall speed, and separate mass- and number-weighted average terminal speeds (except for cloud droplets). This is slightly more complex than the moment-weighted terminal fall velocities in the Morrison scheme as each species has a different equation. For cloud droplets, the terminal fall speed is

$$V_{T,c} = \frac{g(\rho_w - \bar{\rho}_{\text{air}})\bar{D}_w^2}{18\bar{\rho}_{\text{air}}\nu_k}. \tag{2.4.41}$$

The mass-weighted (m) and number-weighted (n) terminal fall speeds for rain are

$$V_{T,r,m}(\bar{D}) = \frac{6\gamma}{\pi}(0.04771 + 3788.0\bar{D} - 1.105\times10^6\bar{D}^2 = 1.412\times10^8\bar{D}^3 - 6.527\times10^9\bar{D}^4), \tag{2.4.42}$$

$$V_{T,r,n}(\bar{D}) = \gamma(0.09112 + 2714.0\bar{D} - 4.872\times10^5\bar{D}^2 + 4.495\times10^7\bar{D}^3 - 1.626\times10^9\bar{D}^4). \tag{2.4.43}$$

For ice crystals, the terminal fall speeds are determined from particle volume as

$$V_{T,i,m}(\bar{v}) = \frac{a_i\Gamma(b_i+\nu_i+2)}{(\nu_i+1)^{b_i}\Gamma(\nu_i+2)}\gamma\bar{v}^{b_i}, \tag{2.4.44}$$

$$V_{T,i,n}(\bar{v}) = \frac{a_i\Gamma(b_i+\nu_i+1)}{(\nu_i+1)^{b_i}\Gamma(\nu_i+1)}\gamma\bar{v}^{b_i}. \tag{2.4.45}$$

Snow terminal fall speeds are given by

$$V_{T,s,m}(\bar{D}) = \frac{5.61743\Gamma(\nu_s+25/12)}{\pi^{1/12}(\nu_s+1)^{13/12}\Gamma(\nu_s+1)}\gamma\bar{v_s}^{1/12}, \tag{2.4.46}$$

$$V_{T,s,n}(\bar{D}) = \frac{5.61743\Gamma(\nu_s + 13/12)}{\pi^{1/12}(\nu_s + 1)^{1/12}\Gamma(\nu_s + 1)}\gamma\bar{v}_s^{1/12}. \tag{2.4.47}$$

Graupel terminal fall speeds are determined by

$$V_{T,g}(D) = \gamma a_g D^{b_g}, \tag{2.4.48}$$

where

$$a_g = \left(\frac{4\rho_g g}{3C_D\rho_{\text{air}}}\right)^{1/2}, \tag{2.4.49}$$

and $b_g = \frac{1}{2}$ (Wisner et al., 1972). The drag coefficient, C_D can be set to a constant (generally 0.8), or scaled by particle density ($0.45 - 0.8$, lower C_D for higher density). Hail fall speed is treated the same as graupel, but with C_D in a different range (typically $0.45 - 0.6$).

2.5 Other Data

2.5.1 Severe Thunderstorm Archive

The Bureau of Meteorology's Severe Thunderstorm Archive (STA; Bureau of Meteorology, 2012) contains reports of observed severe thunderstorms back to the 18th Century, specific to each hazard (hail, wind, tornado and rain). Additionally, it contains data on lightning causing injury/death, waterspouts and damaging dust devils. In this project, the focus is on SCTs producing hail, wind or tornadoes. STA data was obtained up to 10/06/2019 (DD/MM/YYYY), with 13412 raw reports of such phenomena as outlined in Table 2.2 by hazard type. However, not all events recorded actually satisfied the meteorological definition of an SCT. They likely were just reported by the public as severe due to the societal impact, which is seen in other SCT databases around the world (e.g., Shafer and Doswell, 2010). Therefore, all reports were filtered using the report comments and information on the intensity of the event observed (hail size, maximum gust/mean wind speed). If report comments and/or intensity information indicated the event did not meet severe thresholds (hail at least 2 cm in diameter, wind gusts of at least 90 km/h or a tornado), then the report was removed. Additionally, reports were removed prior to 1990 and after 2018 to correspond with the length of the reanalysis data set. Reports also were removed if no information locating the report was provided, or if the report comments suggested the event was not associated with a convective storm (e.g., severe winds associated with synoptic-scale systems during winter). The filtering of non-convective storm events would be backed up by subjective diagnosis of weather radar information, if available. If the report time was missing, then available weather radar data would be used to determine the time of the event, otherwise the report would be removed. There were 60 tornado reports that only mention the sighting of funnels clouds, with no evidence of a tornado reaching the ground (Table 2.2). Many tornado events can go unnoticed due to their often short durations and frequent occurrence in regions with low population density (Griffiths et al., 1993; Allen and Karoly, 2014). Hence, reports of funnel clouds only are kept in the tornado report data set as it is more likely for a funnel cloud to be observed than a tornado touching down, and the inclusion of funnel cloud reports only slightly increases the total reports.

2.5.2 Radar and Lightning Data

The STA only provides data of SCT events, not non-severe thunderstorms or dates where no thunderstorm was observed. In order to build a categorical model that determines when an environment is supportive of an SCT, it is necessary to have additional data of non-severe thunderstorms or no thunderstorm events. To determine likely non-severe thunderstorm and no

Table 2.2: Number of SCT reports by hazard type, for the raw STA data set and the filtered data set.

Report Type	Unfiltered Data	Filtered Data
Hail	4141	2881
Wind	8059	4813
Tornado with funnel	1212	627
Tornado excl. funnel	-	567

thunderstorm dates, lightning strike detections were obtained under a research license from the GLD360 lightning network operated by Vaisala Inc. for the three regions in Australia depicted in Figure 2.5.1. This data provided the date, time and location of lightning strikes over 2014–2018 that could have been associated with convective events. However, some of these convective events could have been severe but not recorded in the STA. Therefore, Maximum Estimated Size of Hail (MESH) data was obtained from the Australian Unified Radar Archive (Soderholm et al., 2019) for the same regions where the lightning data were obtained. Lightning reports that correspond with MESH $\geq$ 2 cm were labelled severe (for hail, as no other information was available to determine if the storm produced severe wind or a tornado), while all other lightning reports were labelled non-severe. Additionally, random dates between 2014 and 2018 were checked against the lightning data and STA to obtain reports where no thunderstorms were recorded. The data of no thunderstorm reports, and non-severe thunderstorm reports were collated into one group that was labelled non-severe. From this data, categorical models were developed on the classes of severe (as reported in the STA, and bolstered by MESH $\geq$ 2 cm reports), and non-severe.

2.5.3 Reanalysis Data Sets

Rawinsondes are often used to provide a vertical profile that represents the state of the atmosphere near to when a storm develops, or even the atmosphere for which a storm propagates into and changes morphology. These vertical profiles provide measurements including temperature, dew point temperature, height, pressure, wind speed and wind direction at different levels of the atmosphere. However, to be representative of the storm environment, they need to be near the development of the storm in both space and time (Rasmussen and Blanchard, 1998; Brooks et al., 1994; Thompson et al., 2003). Unfortunately, the location that rawinsondes are released is inconsistent across Australia, with the nearest vertical profile often hundreds of kilometres away from the SCT in question (Allen et al., 2011). Additionally, the vertical profiles are generally captured at 0000 UTC, with some stations also collecting data at 1200 UTC. Since the diurnal peak of thunderstorm occurrence is in the mid-to-late afternoon (Allen et al., 2011; Soderholm et al., 2017; Warren et al., 2020; Dowdy et al., 2020) the rawinsonde data is rarely representative of the environment in which the storm developed.

For this reason, atmospheric reanalysis data is obtained to provide vertical profiles that are used in the development of the SCT climatology. Additionally, this data is used as both the initial and boundary condition data in WRF for the simulation of the SCT case study. Reanalysis data uses earth system models (similar to WRF) and available observations (including automatic weather station, radar and satellite data) to constrain the model solutions to realistically represent the atmospheric state at the time and location of the available observations. This provides a data set that fills in missing information spatially and temporally, allowing a more detailed view of the atmosphere across the spatial domain considered. For this project, both the Bureau of Meteorology Atmospheric high-resolution Regional Reanalysis for Australia (BARRA; Su et al., 2019) and European Centre for Medium-Range Weather

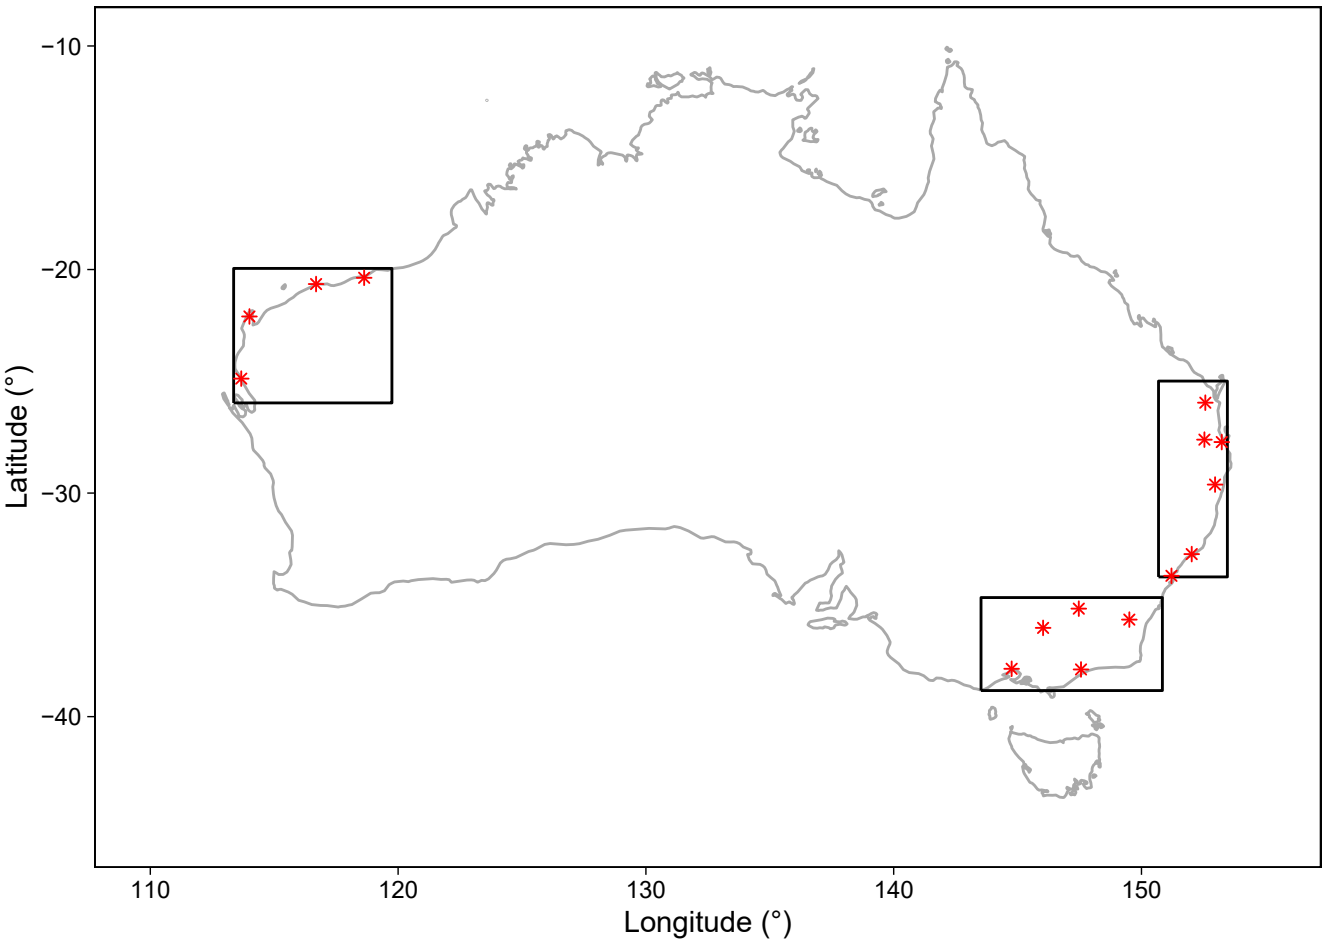

Figure 2.5.1: Map displaying the three regions where Vaisala Inc. GLD360 lightning network data was obtained from 2014–2018 (black rectangles), and BoM weather radars (red stars) within these regions to determine if detected lightning was associated with severe hail, as determined by calculated MESH $\geq$ 2 cm.

Forecasts' (ECMWF's) ERA5 reanalysis data set (Hersbach et al., 2020) are used. BARRA provides hourly data on an Australian-wide domain with horizontal resolution of 12 km, and 1.5 km sub-domains are available covering Sydney, Hobart, Adelaide and Perth from 1990-2018. Meanwhile, the ERA5 data set provides hourly data on a global grid with horizontal resolution of 30 km, covering the period 1950-present.

The BARRA data are used to determine the environmental attributes of SCT hazards across Australia, from the many thermodynamic, kinematic and combined variables used to forecast SCTs. The relationships developed for attribute selection are then applied to the BARRA dataset to develop a climatology of environments supportive of each SCT hazard. The ERA5 data set is used for initial and boundary conditions for the SCT case study with WRF. Additional simulations have been performed using the BARRA data set over the Sydney domain, with 1.5 km horizontal resolution, to compare against the results from ERA5, but have not been included in this thesis.

SCT Parameters

As discussed, a number of parameters were considered due to their application in SCT forecasting. With so many parameters available, machine learning approaches can be used to deter-

mine the parameters with greatest association to observed SCT hazards in Australia. CAPE is perhaps one of the most widely used parameters for assessing convective instability and the potential for storms to develop. However, as mentioned in Section 1.3.2, CAPE is sensitive to the parcel considered (Doswell and Rasmussen, 1994). Therefore, three different parcels have been considered in this study: the surface-based (SB), most unstable (MU), and mixed-layer (ML) parcel over 100 mb have all been used to calculate these variations of CAPE. Additionally, these parcels have been used to determine the CIN, LCL, LFC, EL, lapse rate (LapR), and LI. Additional thermodynamic indices considered include the K-index (George, 2014), TT, precipitable water (PW), and freezing level (i.e., height that air temperature reaches 0°C, in pressure). Kinematic indices considered include $0-1$ km, $0-3$ km, and $0-6$ km bulk shear; and SRH from $0-1$ km, $0-3$ km, and over the effective layer (Davies-Jones, 1984; Thompson et al., 2004). Indices that combine both thermodynamic and kinematic indices also were considered, including the severe weather threat index (SWEAT; Miller, 1972), supercell composite parameter (SCP; Thompson et al., 2002), significant tornado parameter (STP; Thompson et al., 2002, 2003), significant hail parameter (SHIP; Storm Prediction Center, 2014), severe hazards in environments with reduced buoyancy index (SHERB; Sherburn and Parker, 2014), and the energy helicity index (EHI; Thompson et al., 2003) over 0-1 km, 0-3 km, and the effective layer for each parcel type. Additionally, two-way interactions between parameters, obtained through multiplication (e.g., MUCAPE*SFC1kmSHEAR), were considered. The two-way interactions considered are described in greater detail in Chapter 8.

Bibliography

Adlerman, E. and K. Droegemeier. The Sensitivity of Numerically Simulated Cyclic Mesocyclogenesis to Variations in Model Physical and Computational Parameters. Mon. Weather Rev., 130:2671–2691, 2002.

Allen, J. and D. Karoly. A climatology of Australian severe thunderstorm environments of 1979–2011: inter-annual variability and ENSO influence. Int. J. Climatol., 34:81–97, 2014.

Allen, J., D. Karoly, and G. Mills. A severe thunderstorm climatology for Australia and associated thunderstorm environments. Australian Meteorological and Oceanographic Journal, 61:143–158, 2011.

Bamston, A., M. Chelliah, and S. Goldenber. Documentation of a highly ENSO-related sst region in the equatorial pacific: Research note. Atmos. Ocean, 35:367–383, 1997.

Brooks, H., C. Doswell, III, and J. Cooper. On the Environments of Tornadic and Nontornadic Mesocyclones. Weather Forecast., 9:606–618, 1994.

Bureau of Meteorology. Severe Storms Archive, 2012. URL `http://www.bom.gov.au/australia/stormarchive/`.

Davies-Jones, R. Streamwise Vorticity: The Origin of Updraft Rotation in Supercell Storms. J. Atmos. Sci., 59:3178–3205, 1984.

Dawson II, D., M. Xue, J. Milbrandt, and M. Yau. Comparison of Evaporation and Cold Pool Development between Single-Moment and Multimoment Bulk Microphysics Schemes in Idealized Simulations of Tornadic Thunderstorms. Mon. Weather Rev., 138:1152–1171, 2010. doi: 10.1175/2009MWR2956.1.

Doswell, I., C.A. and E. Rasmussen. The Effect of Neglecting the Virtual Temperature Correction on CAPE Calculations. Weather Forecast., 9:625–629, 1994.

Dowdy, A., J. Soderholm, J. Brook, A. Brown, and H. McGowan. Quantifying Hail and Lightning Risk Factors Using Long-Term Observations Around Australia. J. Geophys. Res. Atmos., 125:1–14, 2020. 2020JD0331101.

Enfield, D., M.-N. nez A.M., and P. Trimble. The Atlantic multidecadal oscillation and its relation to rainfall and river flows in the continental U.S. Geophys. Res. Lett., 28:2077–2080, 2001.

George, J. J. Weather Forecasting for Aeronautics. Academic press, 2014.

Griffiths, D., J. Colquhoun, K. Batt, and T. Casinader. Severe Thunderstorms in New South Wales: Climatology and Means of Assessing the Impact of Climate Change. Clim. Change, 25:369–388, 1993.

Hastie, T., R. Tibshirani, and J. Friedman. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media, 2009.

Henley, B., J. Gergis, D. Karoly, S. Power, J. Kennedy, and C. Folland. A Tripole Index for the Interdecadal Pacific Oscillation. Clim. Dyn., 45:3077–3090, 2015.

Hersbach, H., B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. The ERA5 global reanalysis. Q. J. R. Meteorol. Soc., 146:1999–2049, 2020.

Holgate, C., J. Evans, A. van Dijk, A. Pitman, and G. Di Virgilio. Australian Precipitation Recycling and Evaporative Source Regions. J. Clim., 33:8721–8735, 2020.

Hsu, C.-W., C.-C. Chang, and C.-J. Lin. A Practical Guide to Support Vector Classification. Technical report, National Taiwan University, 2003. 16pp.

Johnson, Z., Y. Chikamoto, J.-J. Luo, and T. Mochizuki. Ocean Impacts on Australian Interannual to Decadal Precipitation Variability. Climate, 6(3):61, 2018.

Klemp, J., W. Skamarock, and J. Dudhia. Conservative Split-Explicit Time Integration Methods for the Compressible Nonhydrostatic Equations. Mon. Weather Rev., 135:2897–2913, 2007.

Kucharski, F., I.-S. Kang, R. Farneti, and L. Feudale. Tropical Pacific response to 20th century Atlantic warming. Geophys. Res. Lett., 38:5, 2011.

Kucharski, F., F. Ikram, F. Molteni, R. Farneti, S. Kang, H.-H. No, M. King, G. Giuliani, and K. Mogensen. Atlantic forcing of Pacific decadal variability. Clim. Dyn., 46:2337–2351, 2016.

Kushner, P., I. Held, and T. Delworth. Southern Hemisphere Atmospheric Circulation Response to Global Warming. J. Clim., 14:2238–2249, 2001.

Labriola, J., N. Snook, M. Xue, and K. Thomas. Forecasting the 8 May 2017 Severe Hail Storm in Denver, Colorado, at a Convection-Allowing Resolution: Understanding Rimed Ice Treatments in Multimoment Microphysics Schemes and Their Effects on Hail Size Forecasts. Mon. Weather Rev., 147:3045–3068, 2019.

Lau, K.-M. and H. Weng. Climate Signal Detection Using Wavelet Transform: How to Make a Time Series Sing. Bull. Am. Meteorol. Soc., 76(12):2391–2402, 1995.

Limpasuvan, V. and D. Hartmann. Eddies and the annular modes of climate variability. Geophys. Res. Lett., 26:3133–3136, 1999.

Mansell, E., C. Ziegler, and E. Bruning. Simulated Electrification of a Small Thunderstorm with Two-Moment Bulk Microphysics. J. Atmos. Sci., 67:171–194, 2010. doi: 10.1175/2009JAS2965.1.

Mantua, N. and S. Hare. The Pacific Decadal Oscillation. J. Oceanogr., 58:35–44, 2002.

Martin, G., D. Johnson, and A. Spice. The Measurement and Parameterization of Effective Radius of Droplets in Warm Stratocumulus Clouds. J. Atmos. Sci., 51(13):1823–1842, 1994.

McBride, J. and N. Nicholls. Seasonal Relationships between Australian Rainfall and the Southern Oscillation. Mon. Weather Rev., 111:1998–2004, 1983.

Miller, R. Notes On Analysis and Severe Storm Forecasting Procedures of the Air Force Global Weather Central. Technical Report 200, Air Weather Service, United States Air Force, 1972.

Morrison, H., J. Curry, and V. Khvorostyanov. A New Double-Moment Microphysics Parameterization for Application in Cloud and Climate Models. Part I: Description. J. Atmos. Sci., 62(6):1665–1677, 2005.

Morrison, H., G. Thompson, and V. Tatarskii. Impact of Cloud Microphysics on the Development of Trailling Stratiform Precipitation in a Simulated Squall Line: Comparison of One- and Two-Moment Schemes. Mon. Weather Rev., 137:991–1007, 2009. doi: 10.1175/2008MWR2556.1.

Nicholls, N. Sea Surface Temperatures and Australian Winter Rainfall. J. Clim., 2:965–973, 1989.

Rasmussen, E. and D. Blanchard. A Baseline Climatology of Sounding-Derived Supercell and Tornado Forecast Parameters. Weather Forecast., 13:1148–1164, 1998.

Reisner, J., R. Rasmussen, and R. Bruintjes. Explicit forecasting of supercooled liquid water in winter storms using the MM5 mesoscale model. Q. J. R. Meteorol. Soc., 124(548):1071–1107, 1998.

Saji, N., B. Goswami, P. Vinayachandran, and T. Yamagata. A dipole mode in the tropical Indian Ocean. Nature, 40:360–363, 1999.

Scher, S. and G. Messori. How Global Warming Changes the Difficulty of Synoptic Weather Forecasting. Geophys. Res. Lett., 46:2931–2939, 2019.

Seifert, A. and K. Beheng. A two-moment cloud microphysics parameterization for mixed-phase clouds. Part 1: Model description. Meteorol. Atmos. Phys., 92:45–66, 2006.

Shafer, C. and I. Doswell, C.A. A Multivariate Index for Ranking and Classifying Severe Weather Outbreaks. Electronic J. Severe Storms Meteor., 5:1–39, 2010.

Sherburn, K. and M. Parker. Climatology and Ingredients of Significant Severe Convection in High-Shear, Low-CAPE Environments. Weather Forecast., 29:854–877, 2014.

Skamarock, W. and M. Weisman. The Impact of Positive-Definite Moisture Transport on NWP Precipitation Forecasts. Mon. Weather Rev., 137(1):488–494, 2009.

Skamarock, W., J. Klemp, J. Dudhia, D. Gill, Z. Liu, J. Berner, W. Wang, J. Powers, M. Duda, D. Barker, and X.-y. Huang. A Description of the Advanced Research WRF Model Version 4. Technical Report Note NCAR/TN-556+STR, National Center for Atmospheric Research, 2019.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Weckwerth, and M. Coleman. An 18-year climatology of hailstorm trends and related drivers across southeast Queensland Australia. Q. J. R. Meteorol. Soc., 143:1123–1135, 2017.

Soderholm, J., A. Protat, and C. Jakob. Australian operational weather radar dataset. electronic dataset, national computing infrastructure, 2019.

Storm Prediction Center. Significant Hail Parameter, 2014. URL `http://www.spc.noaa.gov/exper/mesoanalysis/help/help_sigh.html`.

Su, C.-H., N. Eizenberg, P. Steinle, D. Jakob, P. Fox-Hughes, C. White, S. Rennie, C. Franklin, I. Dharssi, and H. Zhu. BARRAv1.0: the Bureau of Meteorology Atmospheric high-resolution Regional Reanalysis for Australia. Geosci. Model Dev., 12:2049–2068, 2019.

Thompson, R., R. Edwards, and J. Hart. Evaluation and Interpretation of the Supercell Composite and Significant Tornado Parameters at the Storm Prediction Center. In 21st Conference on Severe Local Storms. American Meteorological Society San Antonio, Texas, 2002.

Thompson, R., R. Edwards, J. Hart, K. Elmore, and P. Markowski. Close Proximity Soundings within Supercell Environments Obtained from the Rapid Update Cycle. Weather Forecast., 18:1243–1261, 2003.

Thompson, R., R. Edwards, and C. Mead. An Update to the Supercell Composite and Significant Tornado Parameters. In Preprints, 22nd Conf. on Severe Local Storms, Hyannis, MA, Amer. Meteor. Soc. P, volume 8, 2004.

Torrence, C. and G. Compo. A Practical Guide to Wavelet Analysis. Bull. Am. Meteorol. Soc., 79(1):61–78, 1998.

Trenberth, K. Attribution of climate variations and trends to human influences and natural variability. WIREs Clim. Change, 2:925–930, 2011.

Troup, A. The 'southern oscillation'. Quart. J. Roy. Meteor. Soc., 91:490–506, 1965.

Ummenhofer, C., A. Gupta, P. Briggs, M. England, P. McIntosh, G. Meyers, M. Pook, M. Raupach, and J. Risbey. Indian and Pacific Ocean Influences on Southeast Australian Drought and Soil Moisture. J. Clim., 24:1313–1336, 2011.

Wang, H., W. Skamarock, and G. Feingold. Evaluation of Scalar Advection Schemes in the Advanced Research WRF Model Using Large-Eddy Simulations of Aerosol–Cloud Interactions. Mon. Weather Rev., 137(8):2547–2558, 2009.

Warren, R., H. Ramsay, S. Siems, M. Manton, J. Peter, A. Protat, and A. Pillalamarri. Radar-based climatology of damaging hailstorms in Brisbane and Sydney, Australia. Q. J. R. Meteorol. Soc., 146:505–530, 2020. doi: 10.1002/qj.3693.

Webster, P., A. Moore, J. Loschnigg, and R. Leben. Coupled ocean-atmosphere dynamics in the Indian Ocean during 1997-98. Nature, 40:356–360, 1999.

Weisman, M. and J. Klemp. The Dependence of Numerically Simulated Convective Storms on Vertical Wind Shear and Buoyancy. Mon. Weather Rev., 110:504–520, 1982.

Wisner, C., H. Orville, and C. Myers. A Numerical Model of a Hail-Bearing Cloud. J. Atmos. Sci., 29(6):1160–1181, 1972.

Ziegler, C. Retrieval of Thermal and Microphysical Variables in Observed Convective Storms. Part 1: Model Development and Preliminary Testing. J. Atmos. Sci., 42:1487–1509, 1985.

Chapter 3

Comparing Precipitation and Temperature Trends Between Inland and Coastal Locations

3.1 Overview

Australia is no stranger to drought, with numerous long-term droughts recorded throughout history including the Federation drought (1895–1902; Bureau of Meteorology, 2009a), World War II drought (1937–1945; Bureau of Meteorology, 2009b) and Millennium drought (1997–2009; Chiew et al., 2014). More recently, a severe drought impacted southeast Australia, significantly affecting agriculture and water availability.

Droughts have a number of stages that often occur sequentially, beginning with meteorological drought where precipitation deficits accumulate over time (Trenberth et al., 2014; Richman and Leslie, 2020). This is followed by agricultural drought, where soil moisture in the root zone impacts crops, then hydrological drought where streamflow, lake and groundwater levels are below average. Ecological drought results from soil moisture levels being low for a sufficient amount of time that the ecosystem begins to fail. Socioeconomic drought is often the final stage, where deficits in water availability are so large that there is a failure to supply crops, water and some services to the public (Richman and Leslie, 2020).

Overall, droughts have a significant and sometimes devastating impact on the human and natural environment. They can increase the likelihood and intensity of bushfires (Trenberth et al., 2014; Richman and Leslie, 2015, 2020); increase the risk of flooding and landslides when rain falls on dry soil (Richman and Leslie, 2015); cause massive ecosystem damage and loss (Bond et al., 2008); reduce agricultural yields (Cheeseman, 2016); impact tourism (Dube and Nhamo, 2020; Dube et al., 2020); increase energy demand (van Vliet et al., 2016; Ibrahim et al., 2021); and lead to various effects on human health including decreased sanitation and negative mental health (Dey et al., 2011; Stanke et al., 2013; Edwards et al., 2015; Wutich et al., 2020).

The significant effects that drought has on society was particularly evident at the beginning of my candidature. Therefore, along with my supervisors Professor Lance Leslie and Dr Shev MacNamara, it was considered important to understand how droughts might be changing in southeast Australia. In this Chapter, a case study on drought in the Hunter Valley region, north of Sydney, Australia, is presented. Time series and wavelet analysis is applied to understand changes in drought severity between a coastal site, Newcastle, and an inland site, Scone. This case study is focused on changes to wet season precipitation and temperature, where a wet

season is subjectively defined as a period of months where the observing site experiences the majority of its annual rainfall. Changes to precipitation during this period would ultimately have a significant effect on the amount of water available to a region. As lead author of this study, I collected the data, wrote the code to perform time series analysis, modified code to perform wavelet analysis, analysed the results and wrote much of the paper. Professor Lance Leslie and Dr Shev MacNamara conceived the project and assisted with analysis and writing of the paper. This paper has been published in the *ANZIAM Journal* as part of the peer-reviewed proceedings for the Computational Techniques and Applications Conference.

3.2 Introduction

Onset of the current drought (Bureau of Meteorology, 2012) has been extremely sudden, affecting much of southern and eastern Australia. This drought began soon after the record breaking Millennium Drought, which affected much of southern Australia and lasted well over ten years. A number of droughts in recent history have affected similar latitudes globally, including the Cape Town and California droughts (Richman and Leslie, 2015, 2018). These areas have observed an increase in drought frequency over the past few decades, consequent on the global warming trend. With changing precipitation patterns affecting areas of a similar latitude to southeast Australia, it is desirable to determine regions threatened by reduced annual wet season rainfall.

Mean temperatures over Australia have been increasing, particularly since the 1970s, and have been projected to continue increasing into the future Reisinger et al. (2014). An increase in temperature increases evaporation rates which can seriously affect water security. A decrease in precipitation in a region, combined with an increase in mean temperature and population, will tighten water availability, which can have vast economic and social impacts including decreased sanitation, enhanced fire risk, decreased irrigation affecting the agricultural sector and human health impacts. Little work has been conducted to assess trends in precipitation within the Hunter region. This region contains both the second largest city in New South Wales, and a range of agricultural businesses. It has also been heavily affected by the current drought (Figure 3.2.1). This statistical study compares trends in annual wet season precipitation and temperature between a coastal site (Newcastle) and an inland site (Scone), which are both located within the Hunter region. An understanding of these trends is necessary for water resourcing efforts. Advanced statistical techniques are utilised to gain a deeper understanding of how precipitation and temperature have changed over time.

3.3 Data and methodology

Monthly mean maximum temperature (TMax), mean minimum temperature (TMin) and total monthly precipitation time series were obtained for Newcastle (32.9°S, 151.8°E) and Scone (32.1°S, 150.9°E) (Figure 3.2.1) from the Climate Data Online section of the Australian Bureau of Meteorology (BoM) website (`http://www.bom.gov.au/climate/data/index.shtml`). Time series data were chosen from stations with relatively long records and few data gaps compared to other available weather station data within the one area. Nobby's Signal Station was used for Newcastle, which has a continuous record from the late 1800s. Meanwhile, multiple close sites within the township of Scone were utilised to create a long term data set. From 1957–1964, there were no temperature records made at Scone. To combat this, data from the nearby town of Jerrys Plains (32.5°S, 150.9°E) was used with adjustments made based on the difference in mean temperature between the two sites. Any other missing data points were filled

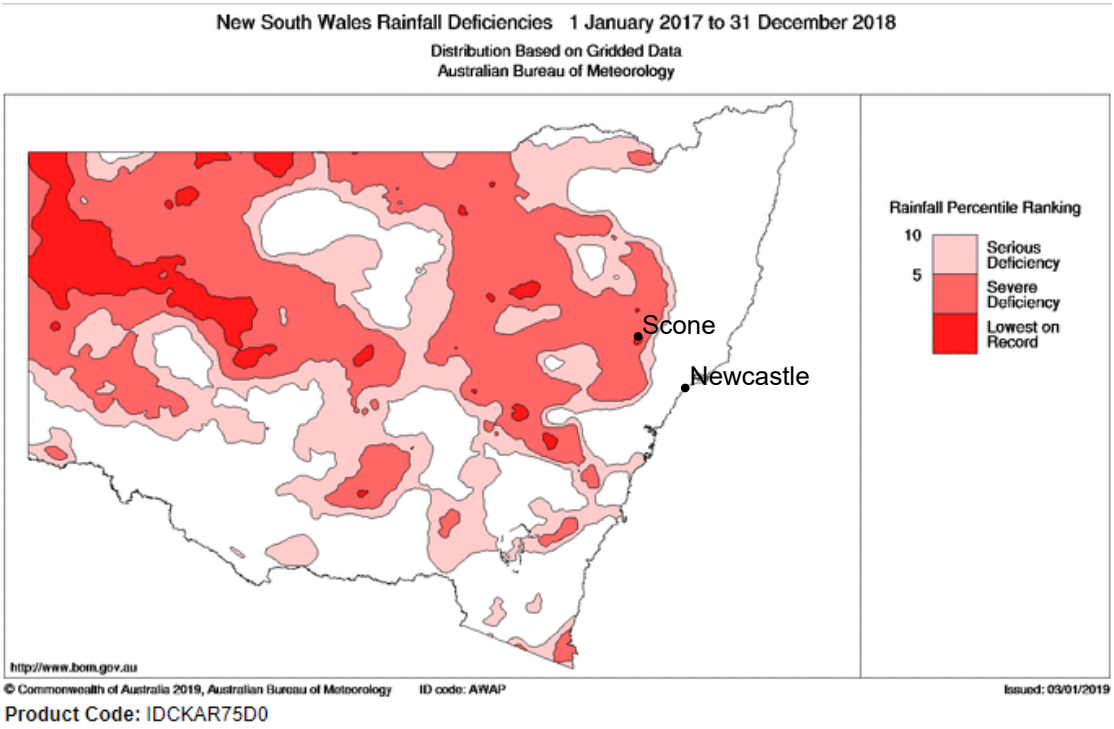

Figure 3.2.1: Drought in NSW over the past 24 months, taken from the Bureau of Meteorology (2018). Locations of Newcastle and Scone are included.

using the moving average centred on the year for which the data was missing. The precipitation time series start at 1890 as this was the earliest year that both locations recorded precipitation, while the temperature time series start at 1910 as this is when consistent temperature recording practices were introduced in Australia.

The focus for this study is changes in precipitation and temperature for both locations during their respective wet seasons. To determine the wet season, monthly averages of precipitation over the entire time series were calculated and compared. Newcastle's wet season was found to run from January through to July, while Scone's wet season runs from October through to March. The time series were first graphed with percentiles to gain an overview of any general changes to precipitation or temperature. The data was then grouped into 20-year blocks and bootstrap resampling was applied with 5000 resamples to gain a deeper picture of any trends within the data. To test if there is a statistically significant difference in the mean of the distribution underlying two 20-year periods, permutation testing was applied, which is a resampling technique that doesn't require prior knowledge of the distribution.

Wavelet analysis was performed for both sites following the approach of Torrence and Compo (1998) and using the accompanying MATLAB software. This computational technique is well suited to our data as wavelets provide information on the time evolution of the periodicity of signals within a time series, which is useful for detecting potential climate drivers, such as the El-Niño Southern Oscillation (ENSO). Unlike Fourier transforms, windowed Fourier transforms are used to analyse local frequency in a time series. However, they inadequately resolve low-frequency signals and require analysis of numerous window lengths to determine the most suitable window, which is not appropriate if there are multiple dominant frequencies in the time series Lau and Weng (1995). The wavelet transform efficiently resolves both high- and low-frequency signals by using local base functions that are stretched and translated in both

frequency and time. This study uses the Morlet wavelet, which is described by

$$\psi_0(\eta) = \pi^{-\frac{1}{4}} e^{i\omega_0\eta} e^{\frac{-\eta^2}{2}}, \tag{3.3.1}$$

where η is a non-dimensional time parameter, and ω_0 the non-dimensional frequency. Given a time series x_n, the continuous wavelet transform is defined as the convolution

$$W_n(s) = \sum_{n'=0}^{N-1} x_{n'} \psi^* \left(\frac{(n'-n)\,\delta t}{s} \right), \tag{3.3.2}$$

where s is the wavelet scale, n is a localised time index, N is the number of points in the time series and * denotes the complex conjugate. In this study, the wavelet transform is computed in Fourier space as it is more efficient than directly computing (3.3.2). The discrete Fourier transform of x_n is

$$\hat{x}_k = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i k n/N}, \tag{3.3.3}$$

where $k = 0, 1, \ldots, N-1$ is the frequency index, and the Fourier transform of the function $\psi(t/s)$ is given by $\hat{\psi}(s\omega)$. Finally, using the convolution theorem and taking the inverse Fourier transform, the wavelet transform is

$$W_n(s) = \sum_{k=0}^{N-1} \hat{x}_k \hat{\psi}^*(s\omega_k) e^{i\omega_k n \delta t}, \tag{3.3.4}$$

where

$$\omega_k = \begin{cases} \frac{2\pi k}{N\delta t} & : \quad k \leq \frac{N}{2} \\ \\ \frac{-2\pi k}{N\delta t} & : \quad k > \frac{N}{2}. \end{cases} \tag{3.3.5}$$

Equation 3.3.4 allows calculation of the continuous wavelet transform for a given wavelet scale using a routine such as the fast Fourier transform in MATLAB. The wavelet power spectrum is defined as

$$P(s) = |W_n(s)|^2, \tag{3.3.6}$$

and is used to visualise the time evolution of periodic signals in a time series. Meanwhile, the global power spectrum is equivalent to applying a Fourier transform to the time series and thereby provides a global view of the dominant amplitudes. The global power spectrum is the average over local wavelet spectra through time

$$\overline{W}^2(s) = \frac{1}{N} \sum_{n=0}^{N-1} |W_n(s)|^2. \tag{3.3.7}$$

As the time series covers only a period of approximately 120 years, the wavelet analysis conducted only considers periodicities from 0–32 years.

3.4 Results

3.4.1 Evolution of precipitation

Figure 3.4.1 presents time series of precipitation for both Newcastle and Scone's respective wet season, and box plots of the bootstrapped mean wet season precipitation over 20-year

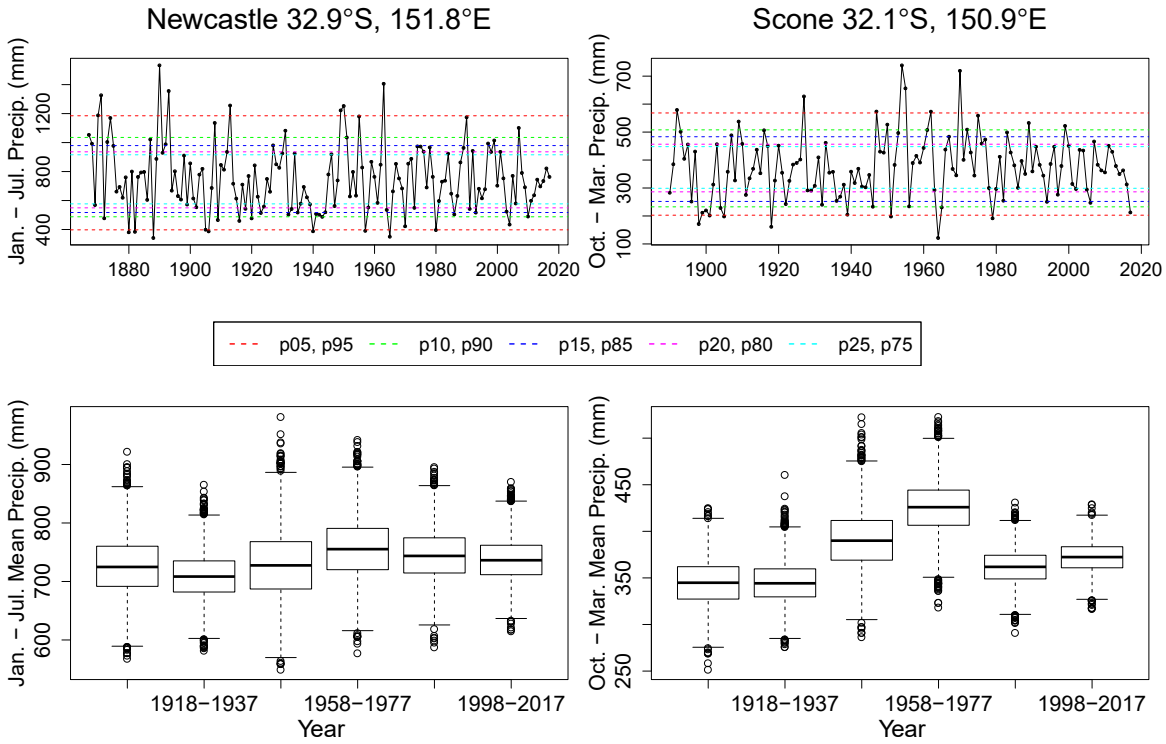

Figure 3.4.1: Time series of total wet season precipitation (top); and box plots of the bootstrapped mean wet season precipitation over 20-year periods (bottom) for Newcastle and Scone. Dashed lines indicate the 5th and 95th (bottom and top red), 10th and 90th (bottom and top green) 15th and 85th (bottom and top dark blue), 20th and 80th (bottom and top pink), and 25th and 75th percentiles (bottom and top light blue).

Table 3.1: P-values for the permutation test on the difference in the mean wet season precipitation between two 20-year periods. Tests conducted used 5000 resamples.

Periods Compared	Newcastle	Scone
1958–1977 v 1978–1997	0.840	0.072
1958–1977 v 1998–2017	0.793	0.123
1978–1997 v 1998–2017	0.900	0.694

periods. While having a consistent y-axis range is ideal, more information is gleaned by not adjusting the range on the plot as wet season precipitation is completely different at each location.

Over the past century, wet season precipitation in Newcastle has remained steady as seen by the lack of statistically significant differences between any two 20-year periods tested (Table 3.1). The frequency of years above the 90th percentile or below the 10th percentile has decreased, as have the tail lengths in the box plots, suggesting reduced variability in wet season precipitation for Newcastle (Figure 3.4.1). Scone's recorded wet season rainfall contrasts with Newcastle's, with a marked increase for the period 1938–1977, compared to the previous forty years. The mean wet season rainfall has since decreased to be consistent with amounts recorded in the early twentieth century. While the past two 20-year periods are similar (p-value = 0.694; Table 3.1), the difference in precipitation between 1958–1977 and 1998–2017 is statistically significant (p-value = 0.123; Table 3.1). Over the past 30 years, there has been a clear decrease in the frequency of years with recorded precipitation above the 85th percentile for Scone (Figure 3.4.1), a trend requiring continued monitoring.

Table 3.2: As in Table 3.1 but for the difference in the mean maximum temperature (TMax), and for the mean minimum temperature (TMin) during the wet season.

	TMax		TMin	
Periods Compared	Newcastle	Scone	Newcastle	Scone
1958–1977 v 1978–1997	0.001	0.162	0.015	0.787
1958–1977 v 1998–2017	0.000	0.058	0.000	0.595
1978–1997 v 1998–2017	0.002	0.661	0.007	0.783

3.4.2 Evolution of temperature

Figure 3.4.2 is the same as Figure 3.4.1, but for the mean maximum temperature (TMax) for each location's wet season. Similarly, Figure 3.4.3 provides an overview of the mean minimum temperature (TMin) for each location's wet season. If the trends in TMax and TMin are consistent, then the mean temperature should also remain roughly consistent with these trends.

Since the 1960s, a strongly increasing trend in TMax has been present in Newcastle. The last year with a record below the 5th percentile was during the 1960s, and since the mid-1980s there have been no mean maximum temperatures below the 25th percentile. The difference in TMax between the periods 1958–1977 and 1998–2017 is statistically significant (p-value = 0.000; Table 3.2). Scone's TMax shows a rather interesting pattern. There was a decrease during 1938–1977, which is most likely due to increased cloud cover consistent with the increased precipitation during this period (Figure 3.4.1). Since the 1980s, the frequency of years recording TMax below the 10th percentile has increased. Despite this, the box plots for Scone show an increasing trend emerging for TMax (Figure 3.4.2), with the difference in mean TMax between 1958–1977 and 1998–2017 being statistically significant (p-value = 0.058; Table 3.2).

The TMin for Newcastle appears to begin increasing in the early-1980s, which is slightly later than TMax. This is evident through the increasing frequency of years with TMin recorded above the 90th percentile, particularly from the 1990s, and the decrease in frequency of years with TMin recorded below the 10th percentile from the 1960s (Figure 3.4.3). Further, there is a statistically significant difference between mean TMin in 1958–1977 and 1998–2017 (p-value = 0.000; Table 3.2). These results, along with those for TMax, suggest that the mean wet season temperature for Newcastle has been increasing since at least the early-1980s. The TMin trends for Scone again show a different pattern. There is a clear increase in TMin between the years 1918–1937 and 1938–1957. However, TMin has since remained relatively constant. For example, there is no statistically significant difference between the years 1958–1977 and 1998–2017 (p-value = 0.595; Table 3.2). While TMin is remaining steady, TMax for Scone has been increasing since the late-1970s, and so the mean wet season temperature for Scone may still be increasing overall.

3.4.3 Wavelet analysis of precipitation

Figure 3.4.4 shows wavelet power spectra (3.3.6) and global power spectra (3.3.7) for wet season precipitation. Both locations show a clear two to seven year periodic signal, suggestive of an influence from the El-Niño Southern Oscillation (ENSO). The global power spectrum for Newcastle has one statistically significant peak around four years, while Scone has two peaks around both four and seven years. The seven year peak for Scone is clearly influenced by the high local power between 1940 and 1970, corresponding to the increase in precipitation occurring in this period, while the four year peak is related to a consistently appearing local signal. There also is a 16–32-year periodic signal which, while not statistically significant,

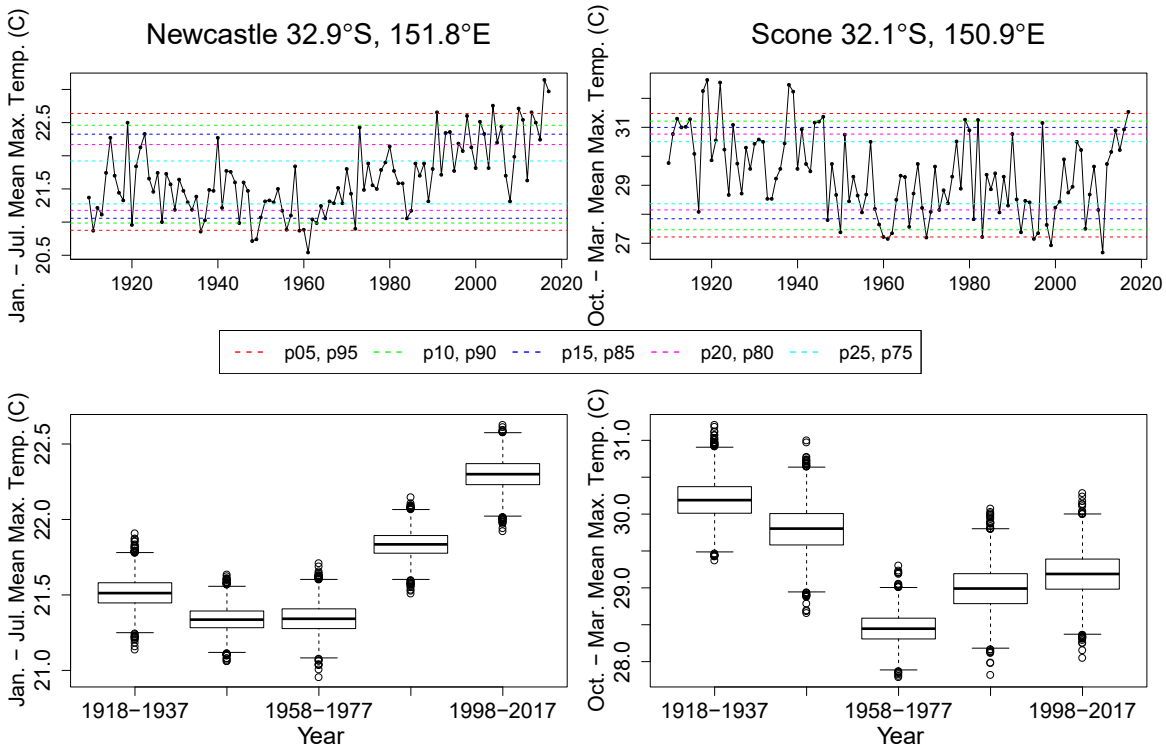

Figure 3.4.2: As in Figure 3.4.1 but for TMax during each location's respective wet season.

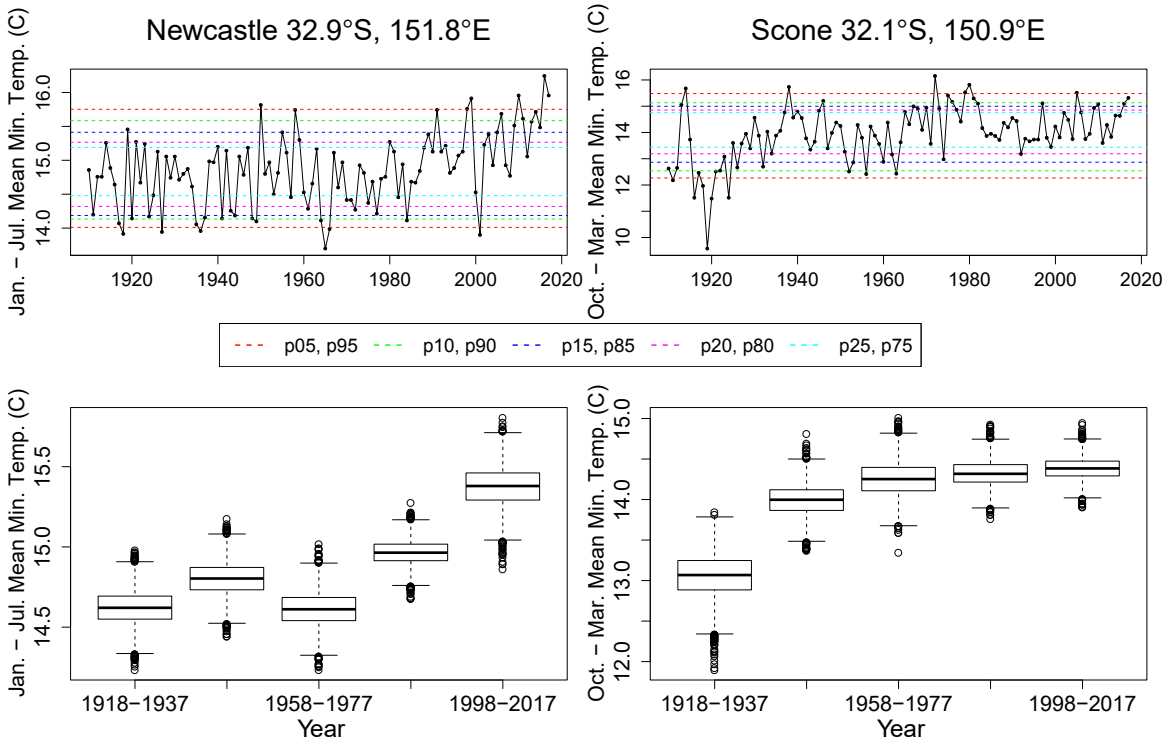

Figure 3.4.3: As in Figure 3.4.1 but for TMin during each location's respective wet season.

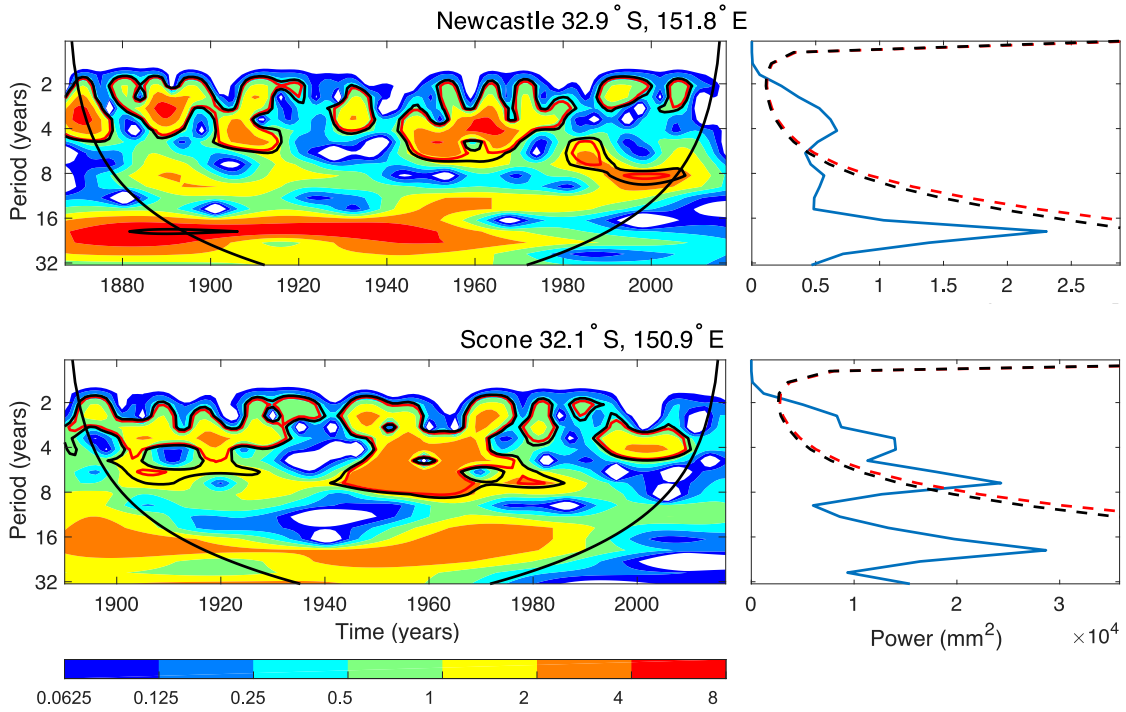

Figure 3.4.4: Wavelet analysis for total wet season precipitation for Newcastle and Scone. Low values (blue) in the wavelet power spectrum (left) indicate low variability while high values (red) indicate high variability. Peaks in the global power spectrum (right) indicate high variability. The dashed/solid red line indicates the 95% confidence level, while the dashed/solid black line indicates the 90% confidence level.

appears to have some influence on the recorded precipitation for both sites.

3.4.4 Wavelet analysis of temperature

Wavelet power spectra also were computed for wet season TMax and TMin (Figure 3.4.5). The TMax power spectra for both locations shows some power in the two to four year range, suggesting that ENSO influences TMax for both of these locations. From the 1970s, the number of years with high power in the two to four year range suggests that ENSO has become more influential since then. There is high power in the two to seven year range for Newcastle TMin, suggesting an ENSO influence. In contrast, Scone's TMin wavelet power spectra has little power in the two to seven year range, except during the 1940s–1960s and the late 1960s–late 1970s. Consequently, Scone's global power spectra has no statistically significant peaks over any periods, implying little influence on Scone's TMin.

3.5 Discussion and Conclusions

Currently, intense drought is affecting much of southern and eastern Australia. This drought followed soon after the Millennium Drought, which affected southern Australia during 1997–2009. This study analysed wet season precipitation and temperature for a coastal and inland station, Newcastle and Scone, using statistical resampling methods and wavelets.

Newcastle precipitation was stable around its long term mean, though variability has decreased since the 1960s. In contrast, Scone's precipitation has fluctuated, with an increase during 1938–1977 followed by a marked decrease during 1978–2017. Wavelet analysis revealed a two to seven

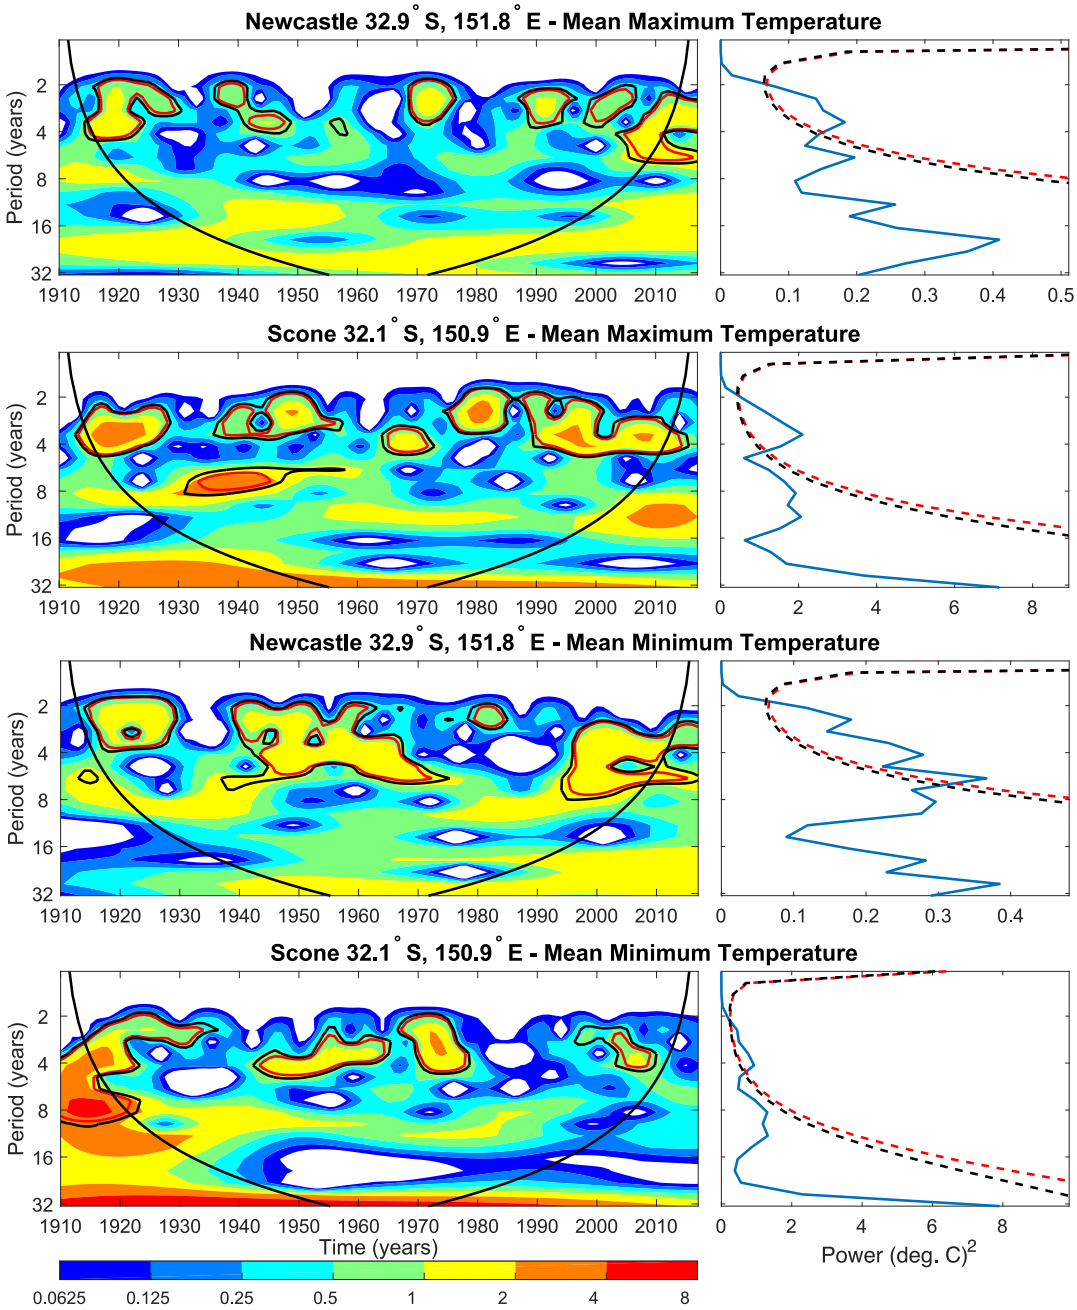

Figure 3.4.5: As in Figure 3.4.4, but for TMax (top two rows) and TMin (bottom two rows).

year signal influencing precipitation in both locations, suggesting some reliance on ENSO. A much longer (16–32-year) signal also is present in both time series. While a lack of precipitation is the main cause for drought, increasing temperatures increase evaporation and hence produce drier soils, thereby exacerbating the drought intensity such that previous rainfall deficits that may not have caused drought could now do so. Thus, an analysis of temperature trends also is necessary. In Newcastle, both the mean maximum temperature (TMax) and mean minimum temperature (TMin) have strongly increased since the mid-20th century. TMax in Scone has increased more gradually since this time, while TMin has hardly changed. This suggests that, overall, the mean temperature for both locations may be increasing, causing dry conditions to be harder felt. Wavelet analysis shows that ENSO may have an influence on TMax for both locations. Meanwhile, Newcastle's TMin is somewhat affected by ENSO, but a clear driver for TMin in Scone was not identified.

Periods of dryness are clearly having greater impact, further accentuated by a growing population, increasing strain on water supplies. The current drought is drastically affecting the Hunter Valley, while coastal areas like Newcastle are relatively unscathed (Figure 3.2.1). Further study of meteorological factors causing this disparity, and of precipitation and temperature trends over eastern Australia is necessary to inform the agriculture and water management sectors. Without further detailed study of changes in precipitation and temperature, drought affected regions will be inadequately prepared, worsening its effects.

Bibliography

Bond, N., P. Lake, and A. Arthington. The impacts of drought on freshwater ecosystems: an Australian perspective. Hydrobiologia, 600:3–16, 2008.

Bureau of Meteorology. The "federation drought", 1895–1902. https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought1.html, 2009a.

Bureau of Meteorology. The World War II droughts 1937–1945. https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought3.html, 2009b.

Bureau of Meteorology. Climate Glossary - Drought, 2012. URL http://www.bom.gov.au/climate/glossary/drought.shtml.

Bureau of Meteorology. Climate Maps - Rainfall Latest, 2018. URL http://www.bom.gov.au/jsp/awap/rain/index.jsp?colour=colour&time=latest&step=0&map=decile&period=24month&area=nat.

Cheeseman, J. Food Security in the Face of Salinity, Drought, Climate Change, and Population Growth. In Halophytes for food security in dry lands, pages 111–123. Elsevier, 2016.

Chiew, F. H. S., N. J. Potter, J. Vaze, C. Petheram, L. Zhang, J. Teng, and D. A. Post. Observed hydrologic non-stationarity in far south-eastern Australia: implications for modelling and prediction. Stoch. Environ. Res. Risk Assess., 28:3–15, 2014.

Dey, N., M. Alam, A. Sajjan, M. Bhuiyan, L. Ghose, Y. Ibaraki, and F. Karim. Assessing Environmental and Health Impact of Drought in the Northwest Bangladesh. J. Environ. Sci. Nat. Resour., 4:89–97, 2011.

Dube, K. and G. Nhamo. Vulnerability of nature-based tourism to climate variability and change: Case of Kariba resort town, Zimbabwe. J. Outdoor Recreat. Tour., 29:1–13, 2020. 100281.

Dube, K., G. Nhamo, and D. Chikodzi. Climate change-induced droughts and tourism: Impacts and responses of Western Cape province, South Africa. J. Outdoor Recreat. Tour., pages 1–10, 2020. 100319.

Edwards, B., M. Gray, and B. Hunter. The Impact of Drought on Mental Health in Rural and Regional Australia. Soc. Indic. Res., 121:177–194, 2015.

Ibrahim, N., S. Wan Alwi, Z. Abdul Manan, A. Mustaffa, and K. Kidam. Impact of Drought Phenomenon on Renewable and Non-renewable Energy Systems in the ASEAN Countries. Chem. Eng. Trans., 83:73–78, 2021.

Lau, K.-M. and H. Weng. Climate Signal Detection Using Wavelet Transform: How to Make a Time Series Sing. Bull. Am. Meteorol. Soc., 76(12):2391–2402, 1995.

Reisinger, A., R. Kitching, F. Chiew, L. Hughes, P. Newton, S. Schuster, A. Tait, and P. Whetton. Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment of the Intergovernmental Panel on Climate Change. Technical report, 2014. 1371-1438 pp.

Richman, M. and L. Leslie. The 2015–2017 Cape Town Drought: Attribution and Prediction Using Machine Learning. Procedia Comput. Sci., 140:248–257, 2018.

Richman, M. B. and L. M. Leslie. Uniqueness and Causes of the California Drought. Procedia Comput. Sci., 61:428–435, 2015.

Richman, M. B. and L. M. Leslie. Machine Learning for Attribution of Heat and Drought in Southwestern Australia. Procedia Comput. Sci., pages 3–10, 2020.

Stanke, C., M. Kerac, C. Prudhomme, J. Medlock, and V. Murray. Health Effects of Drought: a Systematic Review of the Evidence. PLOS Curr., 5, 2013.

Torrence, C. and G. Compo. A Practical Guide to Wavelet Analysis. Bull. Am. Meteorol. Soc., 79(1):61–78, 1998.

Trenberth, K., A. Dai, G. van der Schrier, P. Jones, J. Barichivich, K. Briffa, and J. Sheffield. Global warming and changes in drought. Nat. Clim. Change, 4:17–22, 2014.

van Vliet, M., J. Sheffield, D. Wiberg, and E. Wood. Impacts of recent drought and warm years on water resources and electricity supply worldwide. Environ. Res. Lett., 11:1–10, 2016. 124021.

Wutich, A., A. Brewis, and A. Tsai. Water and mental health. WIREs Water, 7:1–16, 2020. e1461.

Chapter 4

Application of Machine Learning to Attribution and Prediction of Seasonal Precipitation and Temperature Trends in Canberra, Australia

4.1 Overview

It is expected that droughts in southeast Australia will continue to intensify under a warming climate due to expansion of the Hadley cell causing rain-bearing low-pressure systems to propagate further south (Timbal and Fawcett, 2013; Timbal and Drosdowsky, 2013; Post et al., 2014), and intensification of the water cycle leading to more periods of dry weather between more intense precipitation (Trenberth et al., 2014). Improving preparedness for drought includes utilising a wide array of technologies to increase water availability, decrease water loss due to evaporation and increase efficiency of water usage (De Fraiture and Wichelns, 2010; Deo and Şahin, 2015; Bagirov et al., 2017). Additional preparedness comes from improved understanding of the climate drivers that influence water availability and how these influences might be changing, along with improving forecasting from the climate-scale down to hourly forecasting (Choudhury et al., 2019). To this end, a number of studies have begun applying statistical models to better predict precipitation and drought (e.g., Deo and Şahin, 2015; Bagirov et al., 2017). Statistical models require less computational resources to run than numerical weather prediction or climate models, and can be used to understand the complex relationships between climate drivers that influence precipitation locally.

The application of statistical models to understand climate influences on precipitation, and predict precipitation within Australia have seldom been used compared to other modelling techniques. In this Chapter, a case study is presented focusing on Canberra, Australia's Capital City. The first part of this Chapter analyses trends in precipitation and temperature utilising similar techniques to Chapter 3. However, following work published in the peer-reviewed conference proceedings for the International Congress on Modelling and Simulation (Hartigan et al., 2019) but not presented here, the trends in precipitation and temperature are considered both annually and for the four meteorological seasons. The second part of this Chapter applies linear and non-linear statistical models to determine the main climate drivers of annual precipitation and annual mean maximum temperature in Canberra, and assess the predictive performance of these statistical models. As lead author of this study, I collected the data, performed the time series and wavelet analysis, wrote the code to develop the statistical models considered, analysed the results and wrote much of the paper. Professor Lance Leslie and Dr Shev Mac-

Namara helped devise the project and assisted with analysis and writing of the paper. This paper has been published in the international peer-reviewed journal: *Climate*.

4.2 Introduction

Numerous long-term droughts have occurred in Australia throughout recorded history. They include the Federation drought (1895–1902; Bureau of Meteorology, 2009a), the World War II drought (1937–1945; Bureau of Meteorology, 2009b) and the Millennium Drought (1997–2009; Chiew et al., 2014). The current drought affecting southeast Australia began soon after the Millennium drought, with massive impacts on agriculture and water availability, including Australia's capital city, Canberra. It resulted in extreme fire weather, culminating in devastating wildfires affecting a large portion of the region in the 2019–2020 fire season. Globally, locations at regions of similar latitudes to southern Australia, such as Cape Town and California, also have recently been affected by severe droughts (e.g., Richman and Leslie, 2015, 2018). Drought is becoming increasingly frequent and more severe in these regions (Niang et al., 2014; Diffenbaugh et al., 2015) as increasing mean temperatures increase evaporation and enhance bushfire threat. Furthermore, land surface changes and an increasing population can place greater strain on water resources, exacerbating drought impacts (Narisma and Pitman, 2003; De Fraiture and Wichelns, 2010).

In addition to enhanced fire risk, drought can have significant socio-economic effects (Wilhite et al., 2014) through reduced irrigation for agriculture, serious impacts on human health including decreased sanitation, increased energy use, and greater risk of flooding as rain falls on dry soil that inhibits runoff. Canberra is located in southeast Australia and relies on nearby, catchment area rainfall for its water supply. It is situated between Sydney and Melbourne, Australia's two most populous cities, and is representative of many aspects of southeast Australian climate warming. Following the Millennium Drought, numerous programs aimed at decreasing per capita water consumption were introduced, resulting in significant reductions (Icon Water, 2019). However, it is only a temporary solution to the water availability problem as Canberra's population continues to increase rapidly (Australian Bureau of Statistics, 2018). Furthermore, annual precipitation over southeast Australia has decreased, particularly during autumn, due to the increasing strength and poleward progression of the subtropical ridge (Timbal and Fawcett, 2013; Timbal and Drosdowsky, 2013; Post et al., 2014). Projects aimed at increasing water availability include the enlargement of the Cotter dam, and development of a pipeline from the Murrumbidgee River to Googong dam (Icon Water, 2019). The pipeline supply still ultimately is controlled by rainfall, and the current drought led only to a short period of operation of the pipeline, due to low water levels (Burgess, 2018). As the Murrumbidgee River is a major contributor to the Murray-Darling basin (MDB), which produces 39% of the national food supply (Holland et al., 2015), it has limited utility during severe drought conditions.

Management of future water supply in Canberra and the wider southeast Australia region requires both a detailed assessment of any trends in important variables, and the development of accurate models to predict these variables (Deo and Şahin, 2015; Bagirov et al., 2017). This study focusses on trends in precipitation and temperature, as these are key variables in meteorological drought (Richman and Leslie, 2015). These variables can be predicted using global circulation models, statistical models, or a combination of each model (Choudhury et al., 2019). In this study, statistical models are utilised as they require less computational resources, can be used to attribute the main climate drivers associated with these variables, and have performed well in previous studies on prediction of climatic variables (e.g., Deo and Şahin, 2015; Bagirov et al., 2017; Richman and Leslie, 2018, 2020).

4.3 Materials and Methods

4.3.1 Data

Monthly mean maximum temperature (TMax), mean minimum temperature (TMin) and total monthly precipitation time series were obtained for Canberra (35.3° S, 149.1° E) from the publically available Climate Data Online provided by the Australian Bureau of Meteorology (BoM) website (`http://www.bom.gov.au/climate/data/index.shtml`). Stations with relatively long records in the area were chosen to represent Canberra. Any missing data points, which accounted for approximately 1% of the data, were filled using a moving average centred on the year for which the data was missing. The precipitation and temperature time series start at 1938 as this is the earliest period in which both data sets were available for Canberra.

4.3.2 Statistical Analysis

To examine trends in Canberra precipitation and temperature, the data are analysed both annually, and over each of the four seasons, to determine any variations in precipitation or temperature over shorter time-scales. An overview of trends present in the time series was first gained by plotting the time series and their associated percentiles. The data then were grouped into 20-year periods and bootstrap resampling was applied with 5000 resamples to gain a deeper understanding of any trends present in the data. Permutation testing then was applied with replacement to test for statistical significance between the means of the various two 20-year periods. Statistical tests of series breaks could have been performed in this study to determine the periods used (e.g., Buishand, 1982, 1984; Hubert et al., 1989); however, the aim of this study was not to identify precise breakpoints in the time series. As the climate continues to change, and natural variability can enhance or reduce changes in the climate system over different periods, the detection of breakpoints in the time series can be challenging especially when there is no abrupt change. The decision to group data into 20-year periods was made in order to have equal length periods of sufficient size to compare and identify any changes in the mean between these periods.

4.3.3 Attribute Selection

First, a wavelet analysis was performed on the entire detrended time series anomalies, following the approach of Torrence and Compo (1998). This provides an understanding of the time evolution of the periodicity of signals within a time series, which possibly is useful for detecting potential climate drivers, such as the ENSO phases, and analysing how the influence of these climate drivers might change over time. Wavelets are used as they efficiently resolve both high- and low-frequency signals (Lau and Weng, 1995). In this study, the Morlet wavelet is used.

Potential climate drivers under consideration in this study are numerous and include the Atlantic Multidecadal Oscillation (AMO), the Dipole Mode Index (DMI), the global sea surface temperature anomalies (GlobalSSTA), the global temperature anomalies (GlobalT), Niño3.4, the Pacific Decadal Oscillation (PDO), the Southern Annular Mode (SAM), and the Southern Oscillation Index (SOI), obtained from the Earth System Research Laboratory (`http://www.esrl.noaa.gov/psd/gcos_wgsp/Timeseries/`). Tasman Sea sea surface temperature anomalies (TSSST) were also used, which were obtained from the BoM (`http://www.bom.gov.au/climate/change/#tabs=Tracker&tracker=timeseries`). All data are from 1957–2017. It is worth noting that there are many other potential climate drivers that could be considered instead of those used here, such as the Interdecadal Pacific Oscillation (IPO; Henley et al., 2015) being used over the PDO. However, the climate drivers included in this study cover all

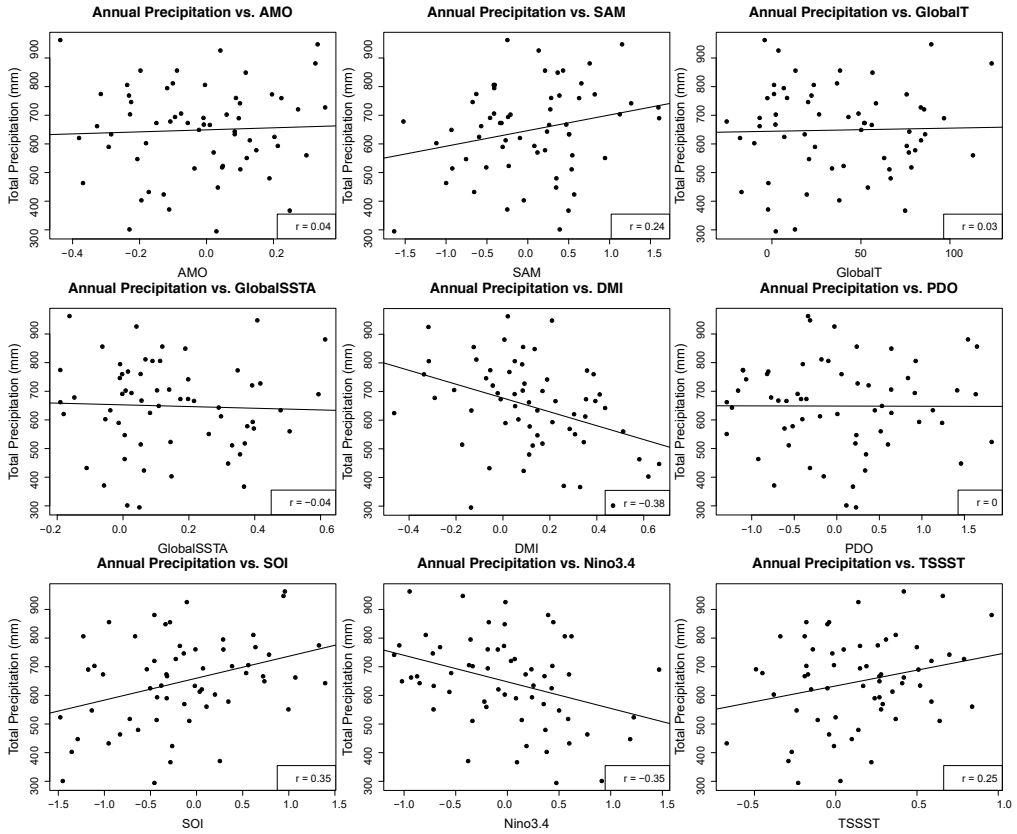

Figure 4.3.1: Scatter plot of annual precipitation in Canberra against each climate driver that serve as possible attributes, showing non-linear relationships between precipitation and the climate drivers. The thin black line is a linear fit against the data using least squares regression. Correlations of each predictor against precipitation are provided in the box to the lower right of each sub-panel, low correlations ($|r| < 0.3$) indicate little linear relationship between a predictor and precipitation.

ocean basins, and many are correlated to those that were excluded. For example the correlation between the IPO and the PDO over the period 1957–2017 is 0.699.

Figure 4.3.1 shows scatter plots of annual precipitation in Canberra against each of these potential climate attributes, with a linear fit and the correlation coefficient between the climate attributes and annual precipitation. Figure 4.3.2 shows the scatter plots for annual TMax. Many of the potential climate drivers have low correlations with annual precipitation or TMax, and exhibit non-linear relationships. This study set out to attribute the climate drivers which have greatest influence on annual precipitation and TMax. As there are non-linear relationships between the variables, non-linear statistical models are used for attribution and prediction. This study performed preliminary work on selection of attributes for annual precipitation and TMax using support vector regression (SVR), similar to the approach of Richman et al. (2015), where both polynomial (poly) and radial basis function (RBF) kernels were considered. Multiple linear regression (LR) was also performed to compare the performance of non-linear techniques with conventional linear techniques.

Two-way attribute interactions, obtained simply by multiplying one variable with another (e.g., Niño3.4*DMI), can be important in statistical models as one predictor might reinforce another. For example, a positive IOD can enhance the drying effect of an El-Niño event in Australia (Risbey et al., 2009; Cai et al., 2011). Past studies using regression and machine learning techniques have found the optimum model selected using different variable selection techniques can involve interaction terms (e.g., Ramsay et al., 2014; Richman and Leslie, 2018, 2020)). As such, this study has included all possible two-way interactions between the aforementioned climate drivers as potential attributes. Other possible relationships, such as an inverse rela-

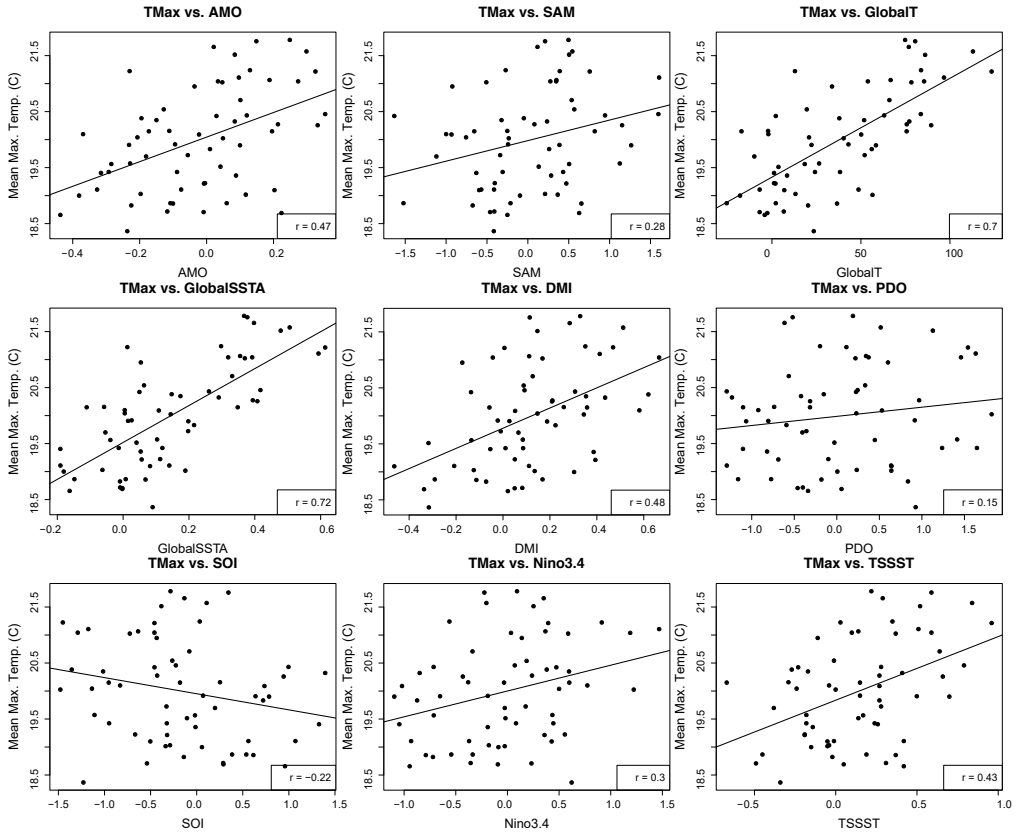

Figure 4.3.2: As in Figure 4.3.1 except for TMax. Predictors with $|r| > 0.4$ indicate moderate or strong relationships with TMax likely because they are attributes of global warming, or have been influenced by global warming.

tionship, could exist between two or more variables. However, in this study it was decided to only use the two-way multiplicative relationships as potential predictors due to the amount of testing required to choose between this or other relationships that might exist.

4.3.4 Training and Prediction

Including interaction terms, there is a total of 45 potential attributes that could serve to predict annual precipitation and TMax. However, using all attributes would result in a model that is severely overfit. This can cause large errors when the model is applied to the test data set, and increases the complexity of the model, making it very difficult to obtain a physical understanding underlying the predictions. As such, it is desirable to select only the attributes that generalize well (Prasad et al., 2017; Richman and Leslie, 2020). In this study, 80% of the data (1957–2005) was selected to train each model. Ten-fold cross-validation (Efron and Hastie, 2016) was applied to the training set using forward selection through the space of potential attributes, with a similar approach to that of Maldonado and Weber (2009).

For the RBF SVR, there are two parameters that need to be selected. The gamma value (G) is a constant in the RBF kernel function that relates to how much influence support vectors have on the decision boundary. The cost parameter (C) is the cost of violating constraints in the optimization problem. The exponent (E) of the poly kernel, and the cost parameter are free parameters that need to be selected in the poly SVR. The free parameters for both RBF and poly kernels are often selected by performing a grid search (Hsu et al., 2003), and they influence how strongly the model fits to the training data, thereby affecting the bias and variance of the model on the testing data set. During the grid search for the final model fitting step in this study, values of C tested ranged from $2^{0.5}$ to 2^3, with powers varying by 0.5; values of G ranged

Table 4.1: Percentage of folds selecting attributes for precipitation using LR, and SVR with the radial basis (RBF) and polynomial (poly) kernel functions.

Attribute	**LR**	**SVR (RBF)**	**SVR (Poly)**	**Attribute**	**LR**	**SVR (RBF)**	**SVR (Poly)**
AMO	10	80	40	DMI*TSSST	10	70	50
DMI	100	60	50	GlobalSSTA*GlobalT	0	80	40
GlobalSSTA	0	60	50	GlobalSSTA*Niño3.4	0	60	60
GlobalT	0	70	70	GlobalSSTA*PDO	20	50	50
Niño3.4	30	50	60	GlobalSSTA*SAM	10	70	50
PDO	10	60	60	GlobalSSTA*SOI	10	40	60
SAM	100	60	70	GlobalSSTA*TSSST	0	70	60
SOI	60	50	50	GlobalT*Niño3.4	30	50	50
TSSST	0	70	60	GlobalT*PDO	40	50	70
AMO*DMI	40	60	30	GlobalT*SAM	20	70	70
AMO*GlobalSSTA	0	90	70	GlobalT*SOI	10	50	50
AMO*GlobalT	0	90	60	GlobalT*TSSST	0	70	50
AMO*Niño3.4	20	60	40	Niño3.4*PDO	0	60	70
AMO*PDO	20	70	20	Niño3.4*SAM	20	80	60
AMO*SAM	70	70	30	Niño3.4*SOI	0	50	60
AMO*SOI	0	70	50	Niño3.4*TSSST	30	50	50
AMO*TSSST	0	70	50	PDO*SAM	10	80	40
DMI*GlobalSSTA	0	30	70	PDO*SOI	10	60	50
DMI*GlobalT	0	30	40	PDO*TSSST	10	70	60
DMI*Niño3.4	0	50	40	SAM*SOI	0	80	40
DMI*PDO	20	50	50	SAM*TSSST	60	80	60
DMI*SAM	40	70	30	SOI*TSSST	90	60	50
DMI*SOI	0	70	30				

from 0 to 1, varying by 0.5; and values of E tested were 1, 2 and 3.

Table 4.1 lists the percentage of folds each attribute appeared in for the models trained on annual precipitation, while Table 4.2 lists the percentage of folds for annual TMax. The SVR models used in the development of these tables did not have the C, G and E parameters tuned, instead taking default values of 1, 1/(no. of predictors) and 3 respectively, in order to decrease computation time. Those attributes that appeared in at least 5 of the 10 folds were retained for further selection using the LR and SVR methods. The most parsimonious model, with the highest correlation and lowest root mean square error (RMSE) against the observed training data was selected and used to predict both annual precipitation and TMax on the testing data set (2006–2017).

Table 4.2: Same as Table 4.1 except for TMax.

Attribute	LR	SVR (RBF)	SVR (Poly)	Attribute	LR	SVR (RBF)	SVR (Poly)
AMO	10	50	80	DMI*TSSST	0	60	70
DMI	100	70	70	GlobalSSTA*GlobalT	10	60	60
GlobalSSTA	100	40	70	GlobalSSTA*Niño3.4	0	70	60
GlobalT	70	40	70	GlobalSSTA*PDO	0	80	60
Niño3.4	60	50	60	GlobalSSTA*SAM	0	60	80
PDO	10	60	70	GlobalSSTA*SOI	0	60	60
SAM	100	70	70	GlobalSSTA*TSSST	90	60	70
SOI	10	60	60	GlobalT*Niño3.4	0	40	60
TSSST	0	30	80	GlobalT*PDO	90	80	70
AMO*DMI	0	60	70	GlobalT*SAM	0	60	90
AMO*GlobalSSTA	30	70	70	GlobalT*SOI	0	60	60
AMO*GlobalT	10	80	60	GlobalT*TSSST	0	60	70
AMO*Niño3.4	0	60	60	Niño3.4*PDO	0	60	70
AMO*PDO	20	80	60	Niño3.4*SAM	0	80	90
AMO*SAM	90	50	90	Niño3.4*SOI	0	60	60
AMO*SOI	0	70	70	Niño3.4*TSSST	0	70	60
AMO*TSSST	0	50	50	PDO*SAM	0	100	50
DMI*GlobalSSTA	0	50	50	PDO*SOI	10	50	70
DMI*GlobalT	0	50	70	PDO*TSSST	0	70	70
DMI*Niño3.4	0	80	40	SAM*SOI	0	90	90
DMI*PDO	10	50	60	SAM*TSSST	0	60	90
DMI*SAM	10	70	50	SOI*TSSST	0	60	70
DMI*SOI	70	70	60				

4.4 Results

4.4.1 Evolution of Precipitation

Time series of precipitation and box plots of bootstrapped mean precipitation over 20-year periods, from 1939–1958 through to 1999–2018, for Canberra are provided in Figure 4.4.1. The data are annual and cover the four seasons, as any seasonal changes to precipitation can potentially be important in influencing drought. The p-values from permutation testing, comparing the means for the periods 1939–1958 and 1979–1998 against the mean for 1999–2018, are shown in Table 4.3. In Table 4.3 we report multiple p-values, so the issue of multiple testing may be considered, as discussed in Chapter 15 of Efron and Hastie (2016), for example. One way to control the false discovery rate is the well-known Bonferroni correction. In our example, with N = 30 tests, we could consider instead replacing the 0.05 threshold by the threshold $0.05/N$. If we did so, then some of the results in Table 4.3 would no longer be significant. However, most of those results remain significant even with a more stringent test, and we note that a Bonferroni correction can sometimes be pessimistic.

The Canberra mean annual precipitation has remained relatively stable throughout the observational record, with little change between 20-year periods, except for the slightly wetter period 1959–1978 (Figure 4.4.1), which does not have a statistically significant difference in the mean compared to other 20-year periods (not shown). However, there appears to have been a modest decline in variability since the 2000s, with a decrease in recorded precipitation above the 75th percentile and below the 25th percentile since the 2000s. Further, there has been a reduction in the interquartile range (IQR) and whisker lengths in the 1999–2018 box plot compared to previous 20-year periods. This decline in variability of rainfall between 1939–1958 and 1999–2018 is statistically significant at the 90th percentile (p-value = 0.093). In contrast, mean precipitation during the seasons is not always as stable. During the Canberra autumn (March–May), a decline in precipitation is apparent as the frequency of years recorded below (above) the 15th (75th) percentile has increased (decreased) since the 1990s. This decline in autumn precipitation also is clear in the box plots, with a statistically significant decrease in mean precipitation between 1939–1958 and 1999–2018 (p-value = 0.026). Meanwhile, mean

Table 4.3: p-values for the permutation test on the difference in the mean wet season precipitation, mean maximum temperature or mean minimum temperature between 1939–1958 and 1999–2018; and 1979–1998 and 1999–2018. Tests conducted used 5000 resamples. Text in bold face highlight statistical significance at the 95% confidence level, text in italics highlight statistical significance at the 90% confidence level.

Years	**Precipitation**	**TMax**	**TMin**
		Annual	
1939–1958 vs 1999–2018	0.7610	**0.0000**	**0.0022**
1979–1998 vs 1999–2018	0.8620	**0.0000**	0.5210
		Autumn	
1939–1958 vs 1999–2018	**0.0262**	**0.0000**	0.5730
1979–1998 vs 1999–2018	*0.0840*	**0.0094**	0.2350
		Winter	
1939–1958 vs 1999–2018	0.8790	**0.0000**	0.5760
1979–1998 vs 1999–2018	0.3680	**0.0000**	0.8050
		Spring	
1939–1958 vs 1999–2018	0.2160	**0.0000**	**0.0002**
1979–1998 vs 1999–2018	0.6120	**0.0002**	0.3030
		Summer	
1939–1958 vs 1999–2018	**0.0466**	**0.0004**	**0.0000**
1979–1998 vs 1999–2018	*0.0978*	**0.0150**	**0.0372**

winter (June–August) precipitation has undergone little change, with no statistically significant difference between 20-year periods. Overall, this suggests preciptation over the cooler months of the year, which are vital for catchment inflows, is decreasing.

Both spring (September-November) and summer (December-February) precipitation appear to be more variable than the other half of the year, likely due to the convective nature of precipitation during this period. The frequency of years recording below the 25th percentile during spring has decreased since the late-1990s, with a corresponding increase in frequency of years recording above the 85th percentile. As a result, the box plots for spring show a very slight increase between 1939–1958 and 1999–2018. However this is not statistically significant (p-value = 0.216). The frequency of years recording below the 25th percentile in summer has also decreased since the 1990s, while the frequency of years above the 80th percentile has increased. This also is apparent in the box plots, with a potential increasing trend in summer precipitation. The difference in mean summer precipitation between 1939–1958 and 1999–2018 is statistically significant (p-value = 0.047). Further, there is relatively significant variability between 20-year periods during summer, with low mean precipitation periods followed by higher mean precipitation periods. For example, the difference in mean precipitation between 1979–1998 and 1999–2018 is statistically significant at the 10% confidence level (p-value = 0.098). The difference in mean precipitation between 1939–1958 and 1959–1978 is just outside of the 15% confidence level, with a low p-value of 0.152. Comparing the low and high mean precipitation periods against each other (e.g., 1959–1978 vs 1999–2018), neither period exhibits a statistically significant difference in mean precipitation (p-value = 0.485). This is despite the suggestion in the box plots that mean precipitation might have increased from 1959–1978 to 1999–2018 (Figure 4.4.1(b)).

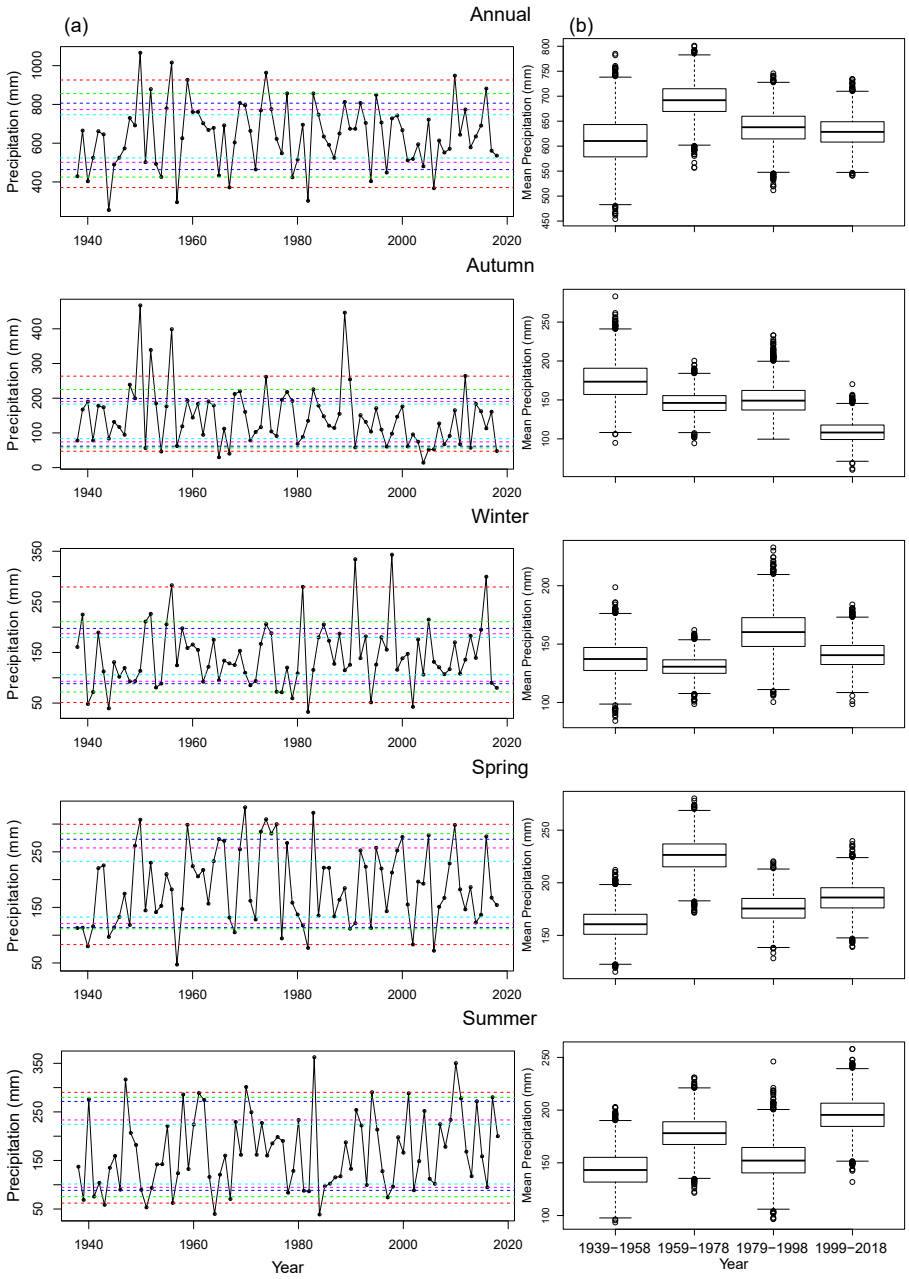

Figure 4.4.1: (a) Annual precipitation time series; and (b) box plots of the bootstrapped mean precipitation over 20-year periods for Canberra, both annually and for all four seasons. Dashed lines indicate the 5th and 95th (bottom and top red), 10th and 90th (bottom and top green) 15th and 85th (bottom and top dark blue), 20th and 80th (bottom and top pink), and 25th and 75th percentiles (bottom and top light blue).

While precipitation appears to have a potentially increasing trend for some seasons, the only season with a statistically significant trend is autumn (p-value = 0.026). As other seasons have stable or gradual increases in precipitation, annual precipitation appears to remain stable over time. Despite this, the decreasing trend in autumn precipitation is of concern because runoff over the cooler months of the year is greatest in southeast Australia (Kirono et al., 2010) and autumn precipitation is necessary for saturation of the soil, allowing runoff into the catchments from the subsequent winter and spring rainfall (Cai and Cowan, 2008; Cai et al., 2009). In addition, the inter-decadal variability of summer precipitation could result in a relatively prolonged period of low annual precipitation, which is currently being realised.

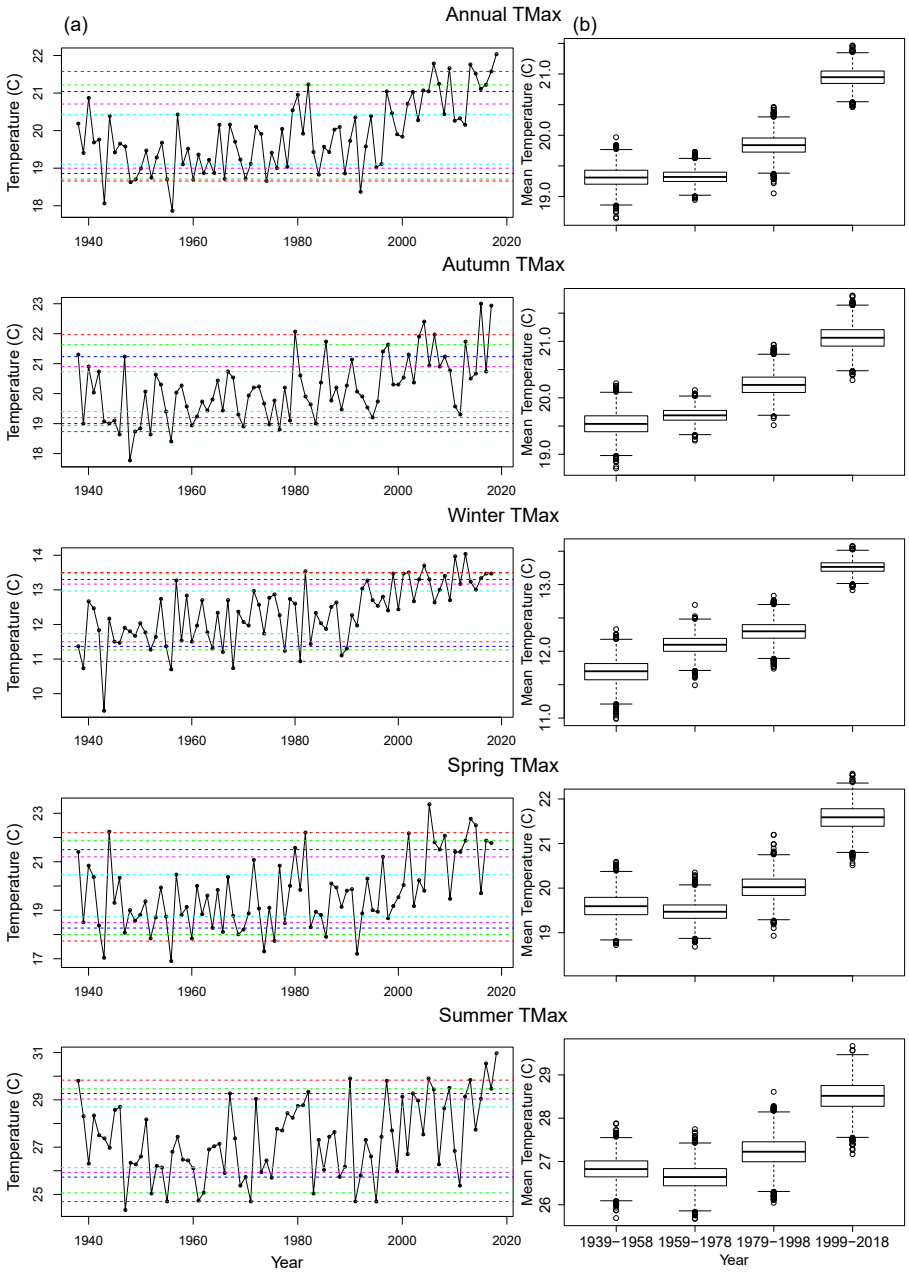

Figure 4.4.2: (a) Time series and (b) box plots as in Figure 4.4.1 except for TMax.

4.4.2 Evolution of Temperature

Temperature plays an important role in drought as it modulates the potential evapotranspiration, creating a positive feedback, as drier soil increases air temperature (Richman and Leslie, 2015). While higher temperatures can allow the atmosphere to hold more moisture, sufficient increases in temperature, particularly in regions experiencing drought, will reduce the number of clouds in the atmosphere due to drying of the soil, resulting in less precipitation (Richman and Leslie, 2018). Thus, it is necessary to monitor how temperature might also be evolving. An approximation to mean temperature changes can be made by observing changes to the mean minimum and maximum temperatures. Figure 4.4.2 provides the time series and box plot figures for TMax in Canberra, both annually and over the four seasons, whereas Figure 4.4.3 provides the overview for TMin.

Annual TMax has been increasing in Canberra, with the difference in mean TMax over the past two twenty-year periods being statistically significant (p-value = 0.000; Table 4.3). Meanwhile, TMin has had a statisically significant increase compared to the earlier part of the

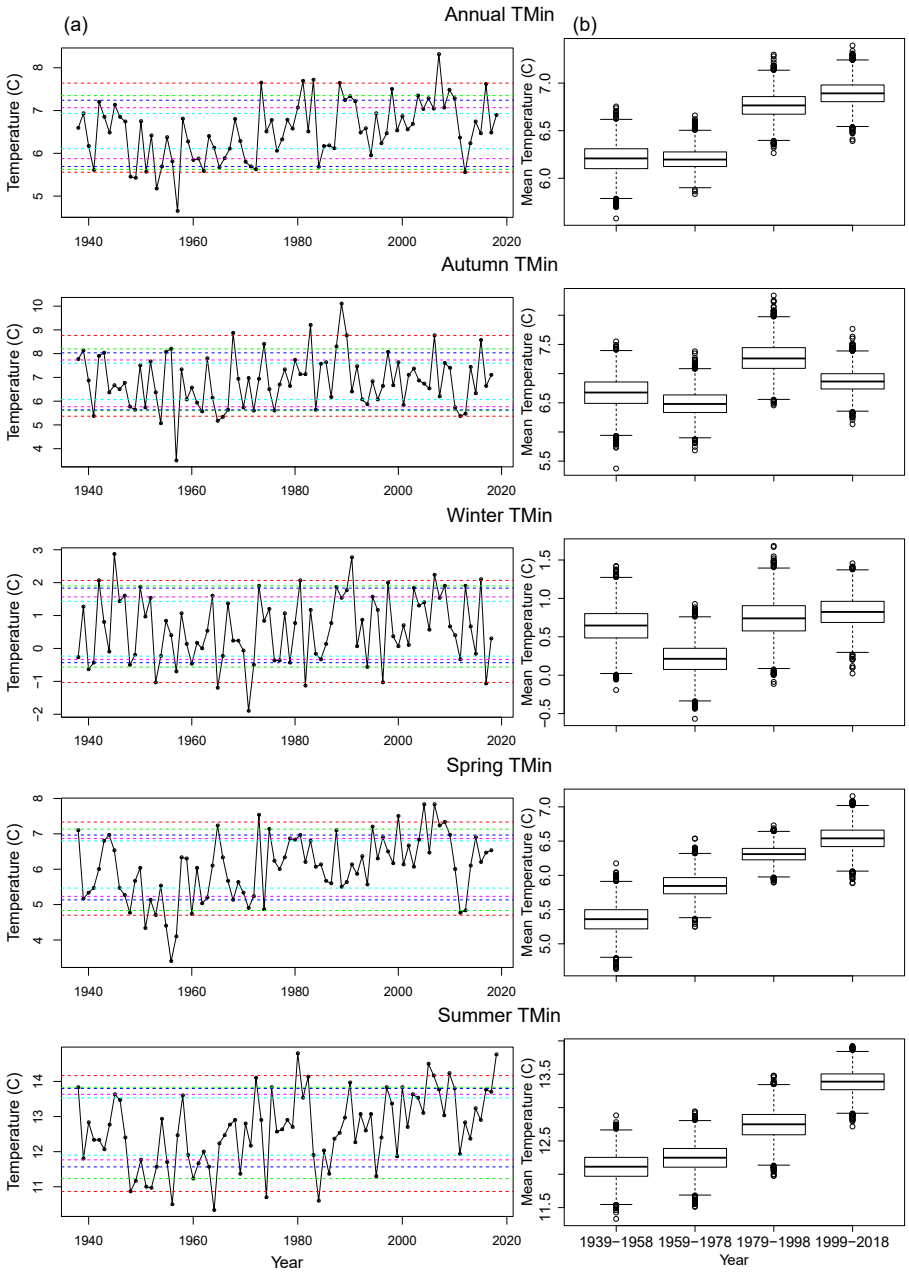

Figure 4.4.3: (a) Time series and (b) box plots as in Figure 4.4.1 except for TMin.

record (1939–1958 vs 1999–2018 p-value = 0.002). However, TMin over the last two twenty year periods has remained relatively stable with no statistically significant difference in the mean (p-value = 0.521). While TMin has increased since the beginning of the record, the increase is far less than TMax, where mean TMax has increased by approximately 1.5 °C over the length of the record (Figure 4.4.2(b)). TMax over the last twenty years stands out against the rest of the time series, with no years recorded below the median, and the majority of years above the 75th percentile occurring during this period. For TMin, the past twenty years have more years above the median, but fewer years above the 95th percentile as the twenty-year period before it, which contributes to these periods having the same means. Variability for both TMax and TMin has remained approximately the same over the observational record with little change to both IQR and the range of the box plots.

For the cooler autumn and winter months, there has been a clear increasing trend in TMax with the difference in mean TMax between 1939–1958 and 1999–2018, and between 1979–1998 and 1999–2018 both being statistically significant (p-values < 0.01; Table 4.3). The increasing

trend for autumn TMax has been relatively gradual, with the frequency of years recorded below (above) the 25th (75th) percentile decreasing (increasing) since the 1980s (Figure 4.4.2(a)). In contrast, there has been a strong, sudden increase in winter TMax over the past twenty years. Since the late-1990s, there are no years below the median and almost all years above the 75th percentile have occurred during this period. This abrupt change in winter TMax also is evident in the box plots with a statistically significant difference in mean TMax for 1999–2018 compared to 1979–1998 (p-value = 0.000). Further, there has been a strong reduction in variability of winter TMax which is marked by the reduced IQR size in the box plots (Figure 4.4.2(b)). Meanwhile, for TMin in the cooler months there has been little change in both the mean and the variance. This is most clear in the box plots, where the only period showing a statistically significant difference in mean TMin is 1959–1978 (p-value compared to 1999–2018 = 0.042; Figure 4.4.3(b)). Overall, as TMax is increasing and TMin is remaining stable, this suggests the mean temperature during the cooler months is increasing. This increase is of concern as most inflow to water storage occurs during these months, and higher temperatures can reduce these inflows and further evaporate water that is already stored (Chiew et al., 2014).

For the warmer months (September-February), TMax again has an increasing trend with statistically significant differences in the mean for both spring and summer in all periods tested (Table 4.3). This increase is relatively gradual with an increase (decrease) in frequency of years above (below) the 75th (25th) percentile since the 1980s in both spring and summer. Notably, there are no years with TMax below the 25th percentile since 2000 in spring, and a large number of years above the 90th percentile in spring and summer (Figure 4.4.2(a)). TMin observations during spring show an increasing trend, although this has tapered off in the most recent 20-year period, with the mean TMin not significantly different from that of 1979–1998 (p-value = 0.303). Since the 1980s, there have only been two years with spring TMin below the 25th percentile, both during the 2000s (Figure 4.4.3(a)). There also was an increased frequency of TMin above the 90th percentile during the late-1990s and 2000s. TMin during summer has continued to increase, with the difference between the two most recent twenty-year periods being statistically significant (p-value = 0.037). Since the 2000s, only one year was recorded just above the 25th percentile, whereas there was a greater frequency in years above the 75th percentile compared to other twenty-year periods. Overall, the mean temperature in the warmer months also appears to be increasing. Because the annual mean temperature is increasing, along with mean temperature during each season, this suggests the region will experience higher mean temperatures during all seasons, resulting in greatly reduced runoff and increased evapotranspiration (Chiew et al., 2014).

4.4.3 Wavelet Analysis of Precipitation

Wavelet power spectra and global power spectra for Canberra precipitation, both annually and over each season, are presented in Figure 4.4.4. There is a strong signal in the 2–7 year period for annual precipitation, which is evident in the global power spectra as a peak over the 3-year and 4-year periods within the 95% confidence band. This is suggestive of an influence from the El-Niño Southern Oscillation (ENSO). Although the influence appears to have weakened slightly over the last 20 years, it remains statistically significant. There are also suggestions at an interdecadal mode with a power spike over the 8–16 year period, although this is not statistically significant due to the length of the time series. ENSO appears to influence precipitation during all seasons, with high power over the 2–7 year period.

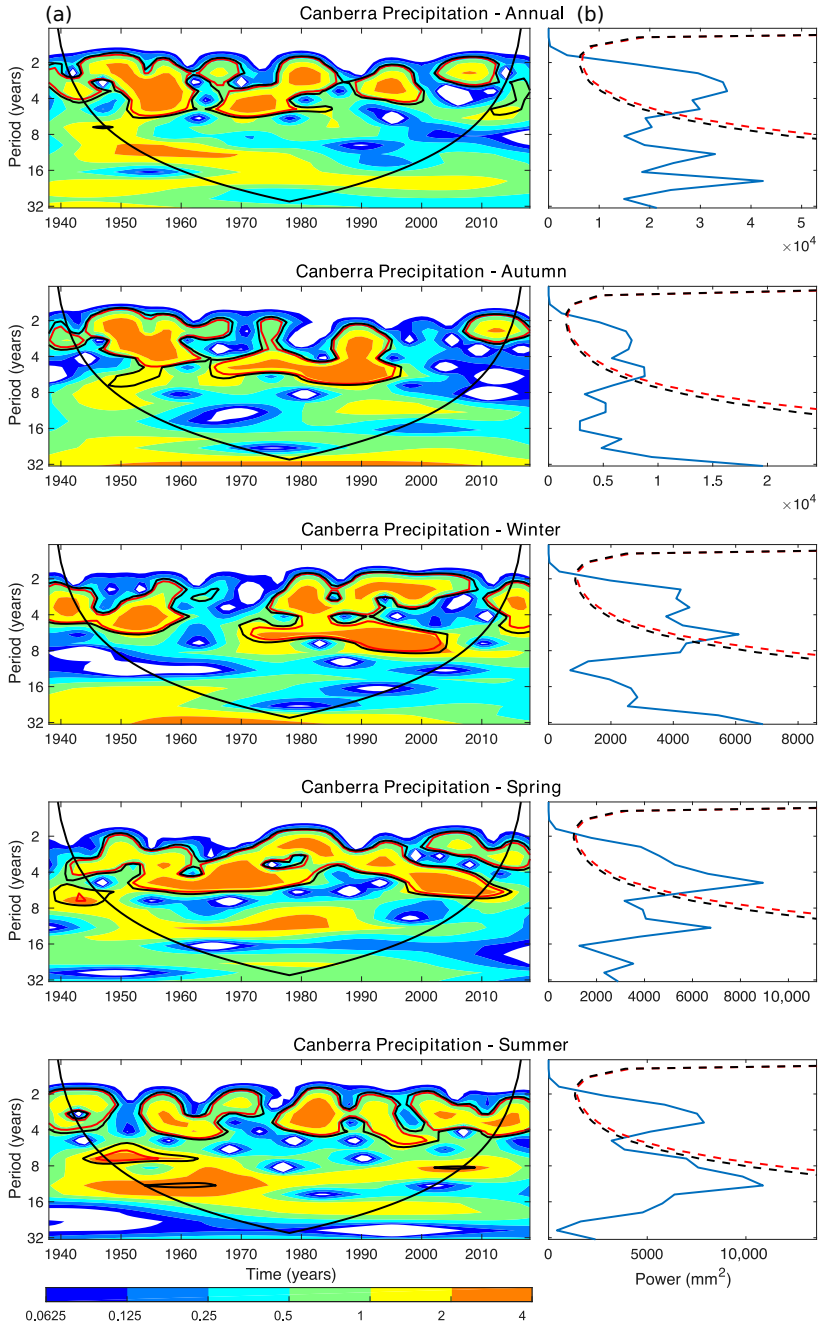

Figure 4.4.4: Wavelet analysis for total wet season precipitation for Canberra. Low values (blue) in the wavelet power spectrum (a) indicate low periodicity while high values (red) indicate high periodicity. Peaks in the global power spectrum (b) on the right panel indicate high periodicity. The dashed/solid red line indicates the 95% confidence level, while the dashed/solid black line indicates the 90% confidence level in each plot. The cone of influence is depicted in the wavelet power spectrum by the solid black cone.

Consistent with previous studies, this influence is strongest during spring, with greater total power and a strong peak over the 4-year period in the global power spectrum (Chiew et al., 1998). The ENSO influence over autumn precipitation is quite strong with high global power spectra, although it has reduced over the past 20 years which could be a consequence of the decreasing trend in autumn precipitation. There is high power over the 8–16 year period in both spring and summer, suggesting these seasons might be influenced by an interdecadal mode of variability such as the PDO or IPO, which is evident in the summer box plots (Figure 4.4.1(b)). However, this signal is mostly not statistically significant at either the 90th or 95th confidence

percentiles. Although the PDO/IPO displays an influence on precipitation and drought risk in eastern Australia (e.g., Power et al., 1999; Kiem and Franks, 2004), the decadal influence on autumn and winter precipitation in Canberra is not apparent. Given the lack of a statistically significant influence in spring and summer, this is likely why power over the 8–16 year period in annual precipitation is weak and not statistically significant.

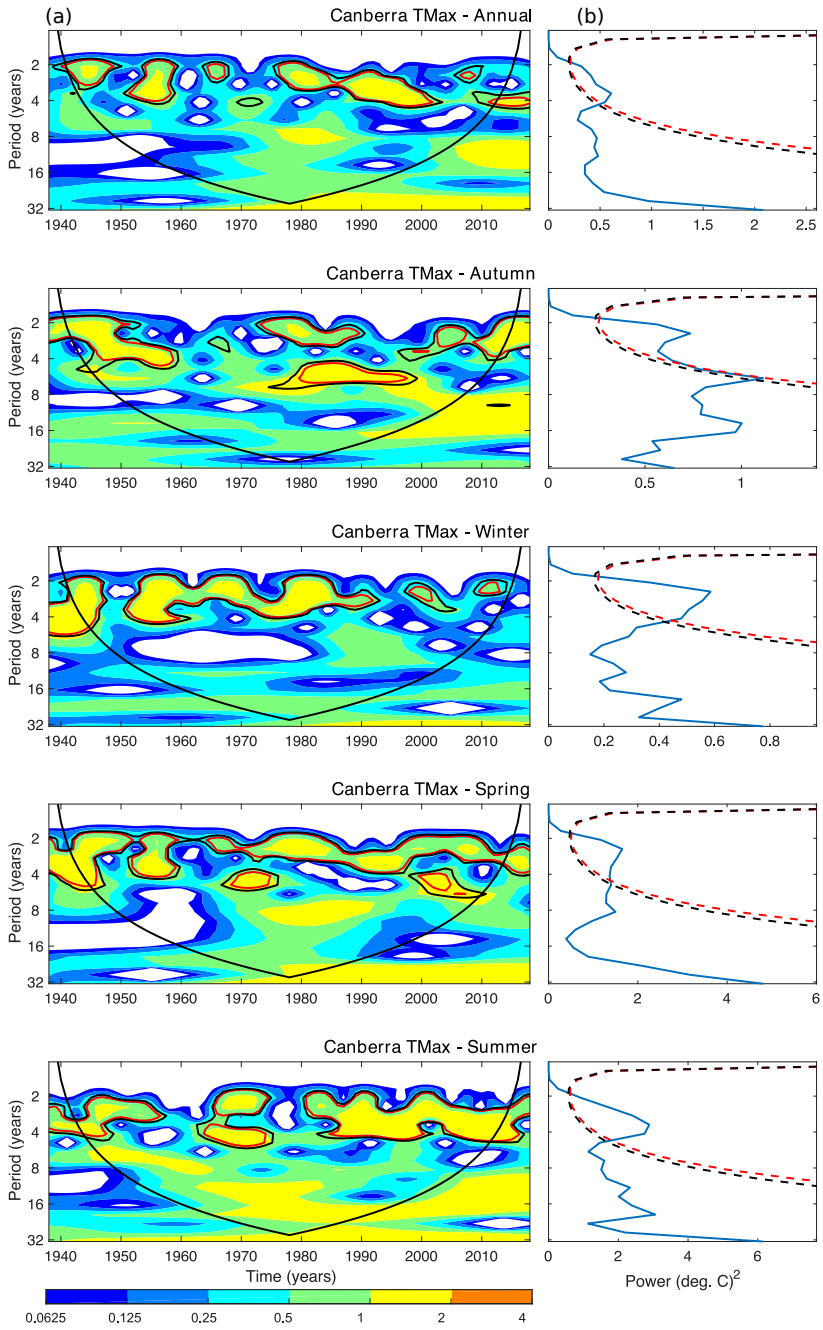

Figure 4.4.5: (a) Wavelet power spectra and (b) global power spectra for TMax, with low and high values in the power spectra indicating low and high periodicity, respectively, the cone of influence and statistical significance bands as already explained in Figure 4.4.4.

4.4.4 Wavelet Analysis of Temperature

Wavelet power spectra and global power spectra also were computed for TMax and TMin (Figures 4.4.5 and 4.4.6, respectively). There is weak, but statistically significant, power over the 2–4 year period for annual TMax that is suggestive of an ENSO influence. However, climate

drivers generally appear to have less influence on TMax than on precipitation. Occasionally, there is moderate power over the 2–7 year period during autumn, though it does not appear consistently over time. For winter, there is moderate power over the 2–7 year period until the 1990s, when this power largely disappears likely due to the sudden shift in the winter TMax time series (Figure 4.4.2). Spring and summer appear to be the most consistent seasons, with moderate power over the 2–7 year period remaining throughout the time series. Summer TMax has weaker power during the earlier part of the observational record, although the strength of this signal has increased since the 1990s, which is seen in the time series with a clearer periodic signal in summer TMax (Figure 4.4.2(a)).

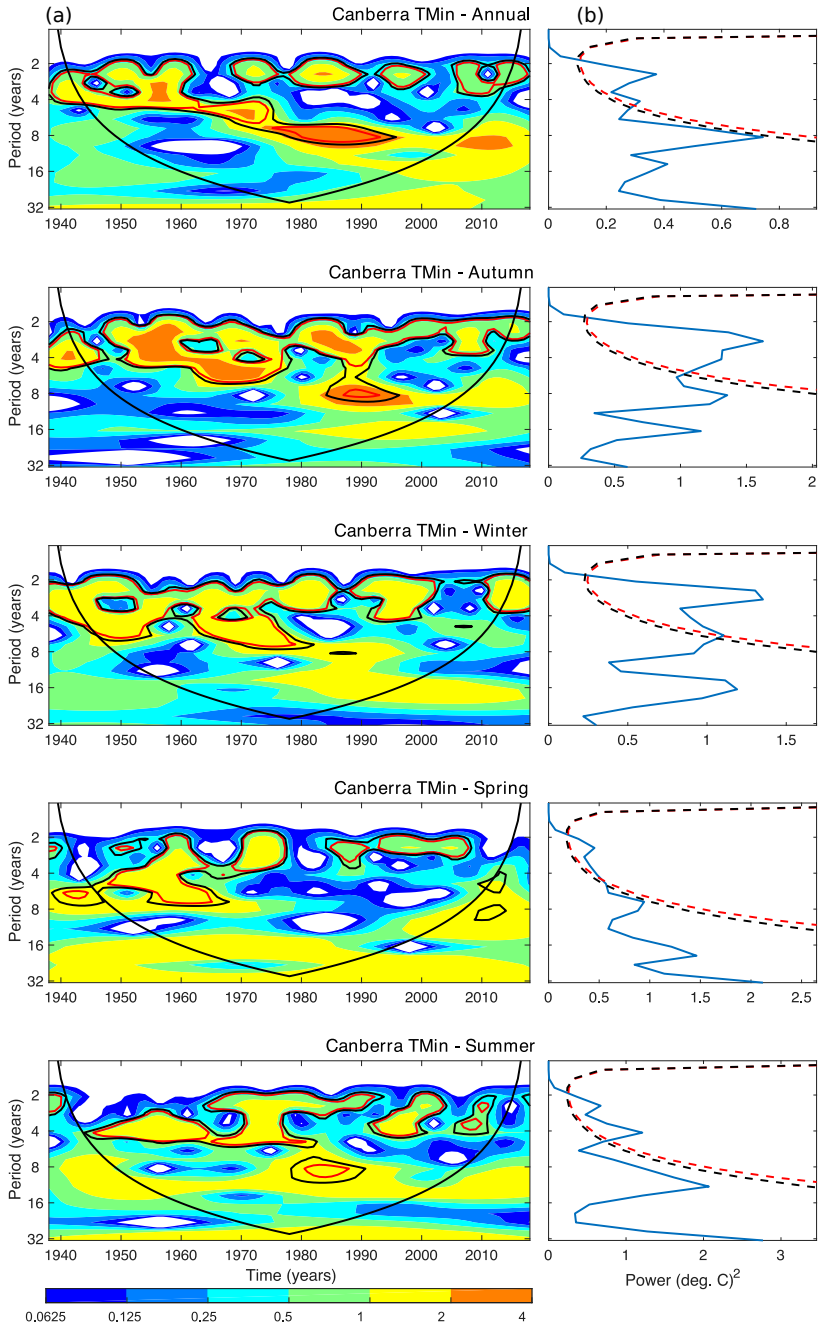

Figure 4.4.6: (a) Wavelet power spectra and (b) global power spectra for TMin, with low and high values in the power spectra indicating low and high periodicity, respectively, the cone of influence and statistical significance bands as already explained in Figure 4.4.4.

TMin appears to have greater periodicity within the time series. Annual TMin has high

power over the 2–4 year period, suggestive of an ENSO influence, but weakens over the most recent 20 years. Autumn TMin has high power over the 2–7 year period, especially prior to the 1990s. Meanwhile, winter TMin has consistent moderate power over the 2–7 year period throughout the time series. TMin during spring and summer also exhibits an ENSO influence. However, in contrast to TMax, this influence begins to weaken from the 1990s for spring, and from the 2000s for summer.

4.4.5 Training and Prediction of Precipitation

Table 4.4 shows the correlation, error statistics and skill scores for the final LR and SVR models on the testing data set. Following Wilks (2011), the definition for skill is:

$$\text{Skill} = 1 - \frac{MSE_{\text{model}}}{MSE_{\text{climatology}}}, \tag{4.4.1}$$

where $\text{MSE}_{\text{model}}$ is the mean square error (MSE) of the model against the observations, and $\text{MSE}_{\text{climatology}}$ is the MSE between the mean of the climate from 1957 up to the year prior to the observation, and the year of the observation.

For precipitation prediction, the LR model that performed best is:

$$\text{precipitation} = 689.14 + 75.07 \times \text{SAM} + 70.23 \times \text{SOI} - 238.73 \times \text{DMI} + \varepsilon, \tag{4.4.2}$$

where ε represents the model error. This model achieved modest correlation, and low RMSE (Table 4.4). In contrast, the RBF SVR model utilises AMO*SOI, DMI*TSSST and SAM*SOI as predictors. The RBF SVR model for precipitation in this study has $\text{G} = 0.5$ and $\text{C} = 2^{0.5}$. The poly SVR uses only DMI and SAM as standalone predictors, with weights of -0.65 and 0.37 respectively, and parameters $\text{E} = 1$ and $\text{C} = 2$.

Table 4.4: Performance of each model on the testing data set (2006–2017). The best performing model for predicting annual precipitation and TMax is highlighted, based on low RMSE, and both high skill and correlation.

Model	Precip. LR	Precip. SVR (RBF)	Precip. SVR (Poly)	TMax LR	TMax SVR (RBF)	TMax SVR (Poly)
RMSE ($\text{mm}^2/°\text{C}^2$)	131.5	**126.0**	127.502	0.871	1.113	**0.783**
Skill	0.246	**0.308**	0.291	0.594	0.337	**0.673**
Correlation	0.516	**0.738**	0.568	0.415	0.708	**0.531**
R^2	0.266	**0.544**	0.322	0.173	0.502	**0.282**
Error Mean (mm/°C)	−21.1	**−22.6**	0.2	0.435	−1.026	**−0.595**
Error SD (mm/°C)	135.5	**129.4**	133.2	0.788	0.452	**0.531**
Error Skewness	−0.181	**−0.115**	0.133	−0.395	0.387	**0.032**
Error Kurtosis	−1.191	**−0.813**	−0.639	−1.453	−1.068	**−1.577**

In Table 4.1, there are many potential attributes for each method that could have been selected by the model. The models occasionally selected attributes that appeared in less folds than others. For example, the RBF SVR for annual precipitation selected AMO*SOI, which appeared in 70% of folds. Meanwhile, other attributes like AMO*GlobalT and AMO*GlobalSSTA, which appeared in 90% of folds, were not selected. It might seem counterintuitive to exclude these attributes, which potentially hold some physical relevance to annual precipitation in Canberra. Despite appearing in a higher percentage of folds, it is likely that these attributes are added later in the forward selection process. As the goal is to obtain a model with the least number of attributes, low RMSE and high correlation, attributes are no longer added to the model when the resulting RMSE on the training set begins to increase.

The RBF SVR model achieves slight reductions in RMSE compared with LR, but has improvements in correlation from 0.516 in the LR model to 0.738 in the SVR RBF (Table 4.4). Meanwhile, the poly SVR improved on the LR model but did not perform as well as the RBF

SVR with slightly higher RMSE and a lower correlation of 0.568 (Table 4.4). Skill scores of all models were positive, suggesting they perform better than using the mean climatology in predicting precipitation for that year. The RBF SVR performed best, with an increase of approximately 25% in skill, compared to LR (Table 4.4). However, these skill scores still are relatively low, with the RBF SVR model only achieving a skill of 0.308, suggesting that the current models for precipitation could be improved upon in future work.

Both the LR and RBF SVR tend to underpredict precipitation. Looking at model predictions against observations (not shown), all three models considerably underpredicted annual precipitation in 2010, 2012 and 2016. Annual precipitation in both 2010 and 2016 were above the 10th percentile (Figure 4.4.1(a)). For the RBF SVR, extremes in both high and low precipitation years tend to be underpredicted in the training data (not shown), although the shape of the precipitation time series is mostly achieved. Therefore, some of this prediction error is likely because the model did not train as close to precipitation extremes. Furthermore, both LR and RBF SVR models contain SOI as a predictor, which is highly correlated to ENSO and its impact on Australia (McBride and Nicholls, 1983; Chiew et al., 1998). The wavelet analysis in this study suggested ENSO influences annual precipitation (Figure 4.4.4). However, this influence has decreased over the most recent 20 years. It is likely that the reduced ENSO relationship could contribute to model errors both in the testing data set and into the future.

4.4.6 Training and Prediction of TMax

For TMax, the best performing LR model was:

$$\text{TMax} = 19.44 + 1.51 \times \text{GlobalSSTA} - 1.28 \times \text{DMI} \times \text{SOI} + 2.92 \times \text{GlobalSSTA} \times \text{TSSST} + \varepsilon. \quad (4.4.3)$$

This model obtained the lowest correlation (0.415; Table 4.4) of all models. The RBF SVR model used DMI*PDO, DMI*TSSST and GlobalT*TSSST as predictors. The free parameters used in this model were G = 0.5 and C = $2^{0.5}$. Meanwhile, the poly SVR model used DMI*SAM and GlobalSSTA with weights of -0.63 and 2.18 respectively, and free parameters E = 1 and C = $2^{0.5}$. The RBF SVR model achieved an RMSE of 1.113, which was worse than the RMSE of 0.871 for LR (Table 4.4). Meanwhile, the poly SVR recorded the lowest RMSE (0.783; Table 4.4) on the testing data set but it obtained a lower correlation of 0.531 compared to 0.708 for the RBF SVR (Table 4.4). Overall, the poly SVR appears to perform best out of these three models with an increase in skill of approximately 13% over the LR model and 100% over the RBF SVR model. Again, all models show improvement compared to the mean climatology, with greater improvements in skill compared to the precipitation models. For example, the poly SVR achieves a skill of 0.673, whereas the RBF SVR for precipitation has a skill of 0.308 (Table 4.4).

Unlike the LR model, both SVR models tended to underpredict, with the RBF SVR largely underpredicting (Table 4.4). While both models include attributes related to global warming, it is likely that the contribution has been underestimated. Considering the time series and box plots of annual TMax (Figure 4.4.2), the increase in TMax is greatest in the most recent 20-year period, which covers the entire testing data set.

4.5 Discussion

4.5.1 Drought Vulnerability

Annual precipitation is not found to be decreasing in Canberra, however the seasonal evolution of precipitation exhibits a decline in autumn precipitation consistent with previous studies

(Timbal and Fawcett, 2013; Timbal and Drosdowsky, 2013; Post et al., 2014). Summer precipitation shows a slight increasing trend, which was not statistically significant, and both winter and spring precipitation were stable. While annual precipitation is stable, this change in seasonal precipitation indicates the importance of splitting the time series into smaller subsets to observe any trends masked in the annual time series. Canberra water storages are reliant on cool season rainfall, as this is when runoff tends to be greater and there is less evapotranspiration (Kirono et al., 2010; Kiem et al., 2016). Further, summer precipitation is highly variable as it is often related to convective storms, with short bursts of heavy rain occurring rather than long periods of soaking rain. This finding suggests that catchment inflows might become less reliable in the future as the climate warms.

Annual mean maximum (TMax) and mean minimum (TMin) temperature show an increasing trend, suggesting the mean temperature in Canberra also is increasing. TMax is increasing in every season, while TMin remains relatively stable over the cooler months and is increasing during the warmer spring and summer months. Hence, the mean temperature appears to be increasing across every season. As a result, potential evapotranspiration will increase in the region, leading to reduced water storages and less rainfall runoff (Cai and Cowan, 2008; Cai et al., 2009; Kirono et al., 2010; Chiew et al., 2014). Consequently, the trends in both precipitation and temperature suggest that Canberra is likely becoming more susceptible to drought conditions.

4.5.2 Attribution and Prediction

Attribute selection on annual precipitation and TMax was carried out using LR and SVR with 10-fold cross-validation. The SVR model was developed using both polynomial (poly) and radial basis function (RBF) kernels. Both SVR models improved on precipitation prediction compared to LR, with the greatest improvements from the RBF SVR. This outperformed the LR model, with a correlation of 0.738 compared to 0.516, and skill against the mean climatology of 0.308 compared to 0.246 (Table 4.4). Predictors influencing annual precipitation in Canberra included DMI, SAM, SOI and TSSST. The wavelet analysis suggested ENSO as a potential driver of precipitation (Figure 4.4.4), which can be defined using the SOI (McBride and Nicholls, 1983). However, other selected attributes were not suggested by the wavelet analysis as their periods are shorter than a year.

Meanwhile, the SVR models for TMax showed mixed results compared to the LR model. The skill from the LR model was 0.594, but the RBF SVR had a reduced skill of 0.337 (Table 4.4). However, the poly SVR appreciably improved in skill (0.673) and correlation (0.531) compared to LR (correlation = 0.415; Table 4.4). The wavelet analysis suggested some ENSO influence on TMax, however this was relatively weak (Figure 4.4.5). The common attributes selected included DMI and measures of global warming, such as global sea surface temperature anomalies or global temperature anomalies.

The reliance of the precipitation models on SAM is likely because this climate driver relates to the position of the belt of westerly winds in the southern hemisphere and the associated position of rain-bearing low pressure systems along southern Australia (Hendon et al., 2007; Risbey et al., 2009). The DMI is recognised to have a strong influence on precipitation especially for southeast Australia as it affects how much moisture is transported into precipitation-bearing systems (Risbey et al., 2009; Ummenhofer et al., 2009). Furthermore, the DMI has been shown to influence decadal variability of precipitation in southeast Australia (Ummenhofer et al., 2011). This influence is reflected in both the LR and poly SVR models as the DMI term has the greatest weight out of all predictors. The Tasman Sea is adjacent to eastern Australia, so the two-way interaction between DMI and TSSST used in the RBF SVR suggests that there are concurrent effects of the Tasman Sea and Indian Ocean on moisture advection into

precipitation systems that appears to improve prediction. The SOI affects where precipitation is most favourable in the Pacific region. The LR and poly SVR models have either weak or no contribution from SOI, however the RBF SVR includes the two-way interaction between SOI and AMO. This suggests a teleconnection between the Pacific and Atlantic basins that may influence eastern Australia rainfall more than ENSO alone, which has been shown by Choudhury et al. (2017) and Johnson et al. (2018). Interestingly, despite the known effect of co-occurring ENSO and DMI events of the same sign on precipitation extremes in southeast Australia (Ummenhofer et al., 2011), the two-way interaction term between DMI and either SOI or Niño3.4 was not selected by any model. This could be a result of the models not training close to precipitation extremes. As all models utilise SAM and DMI in some form, these climate drivers appear to have the greatest influence on precipitation in Canberra.

As DMI influences moisture in the region (Risbey et al., 2009; Ummenhofer et al., 2009), it might influence TMax through modulating how much cloud cover is present, thereby affecting incoming and outgoing solar radiation. Similarly, due to changing the location of low-pressure systems (Hendon et al., 2007; Risbey et al., 2009), SAM also can have an effect on cloud cover, and therefore TMax in the region, which is likely why it was selected by the better performing poly SVR. Further, as there is a strong warming signal in the time series (Figures 4.4.2 and 4.4.3), inclusion of GlobalSSTA and GlobalT, which are attributes of global warming, would likely improve prediction of TMax. In fact, the global warming signals contribute close to 80% of the poly SVR weights and similarly make 80% of the predictor contribution in LR.

The results highlight that relatively accurate models of precipitation and TMax can be developed utilising machine learning techniques. However, only forward selection was applied to select the models in this study. Other model selection techniques, such as a backward search (Maldonado and Weber, 2009) or correlation-based feature selection (Hall, 1999), could be applied in the future to improve confidence in the generalizability of the results. This especially is the case for predicting annual precipitation, as the skill in the models developed in this study, compared with climatology, were relatively low (skill for LR is 0.246, for SVR RBF is 0.308; Table 4.4). This reduced skill could be partly due to the weakening relationship of ENSO on annual precipitation (Figure 4.4.4(a)). Further insight might also be gained by considering other climate drivers as potential predictors such as the Atlantic Niño or trans-basin variability index, which have also displayed influences on precipitation variability in Australia (Choudhury et al., 2017; Johnson et al., 2018). Despite including two-way interaction terms and predictors that are well-correlated with other climate drivers not included in this study, replacement with these different climate drivers could result in improved predictive skill over the models presented in this study. Lagged time series of climate predictors could also be used as predictor variables in future work, as they have been shown to be of benefit in past drought modelling studies (e.g., Abbot and Marohasy, 2017; Prasad et al., 2017). Furthermore, because precipitation patterns vary across the seasons, improved predictive skill might be obtained by developing models for each season (Choudhury et al., 2019).

4.6 Conclusions

In recent years heavy rainfall events, notably in 2010, 2012 and 2016, have provided relief to drought affected southeast Australia. However, it is well-known that Australia is frequently affected by drought, and for southeast Australia it is expected to become increasingly common as global warming continues (Post et al., 2014). Numerous initiatives aimed to improve water availability in Canberra were introduced following the Millennium drought (Icon Water, 2019) but, as the most recent drought has shown, they do not render the city immune to future drought. This study found that precipitation and temperature are evolving in a way that can decrease catchment inflows and increase the likelihood of drought in Canberra. Greater

demand on water resources is expected into the future as the population continues to increase, making it vital to continue research in how factors influencing drought vary over time. This also necessitates the study of improving predictive capabilities of drought.

Multiple linear regression (LR) and support vector regression (SVR) was used to model annual precipitation and mean maximum temperature (TMax), as this can highlight the most important climate drivers and prove useful to management of future water resources (Deo and Şahin, 2015; Bagirov et al., 2017). Annual precipitation appeared most influenced by DMI, SAM, SOI and Tasman Sea SST anomalies. Meanwhile, DMI and measures of global warming influenced annual TMax. The SVR models notably improved on prediction of annual precipitation and TMax compared to LR. Overall, the results of this study highlight how beneficial continued research in machine learning model development can be for improving the prediction of, and for increasing our understanding of, the underlying dynamics of atmospheric variables and weather patterns.

Bibliography

Abbot, J. and J. Marohasy. Skilful rainfall forecasts from artificial neural networks with long duration series and single-month optimization. Atmos. Res., 197:289–299, 2017.

Australian Bureau of Statistics. Australian Demographic Statistics. Technical report, 2018. 60 pp.

Bagirov, A., A. Mahmood, and A. Barton. Prediction of monthly rainfall in Victoria, Australia: Clusterwise linear regression approach. Atmos. Res., 188:20–29, 2017.

Buishand, T. Some methods for testing the homogeneity of rainfall records. J. Hydrol., 58: 11–27, 1982.

Buishand, T. Tests for detecting a shift in the mean of hydrological time series. J. Hydrol., 73: 51–69, 1984.

Bureau of Meteorology. The "federation drought", 1895–1902. https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought1.html, 2009a.

Bureau of Meteorology. The World War II droughts 1937–1945. https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought3.html, 2009b.

Burgess, K. Canberra's emergency water pipeline makes short-lived debut. The Canberra Times, 2018. URL https://www.canberratimes.com.au/politics/act/canberra-s-emergency-water-pipeline-makes-short-lived-debut-20181112-p50fk5.html.

Cai, W. and T. Cowan. Dynamics of late autumn rainfall reduction over southeastern Australia. Geophys. Res. Lett., 35:L09708, 2008.

Cai, W., T. Cowan, P. Briggs, and M. Raupach. Rising temperature depletes soil moisture and exacerbates severe drought conditions across southeast Australia. Geophys. Res. Lett., (21): L21709, 2009.

Cai, W., P. van Rensch, T. Cowan, and H. Hendon. Teleconnection Pathways of ENSO and the IOD and the Mechanisms for Impacts on Australian Rainfall. J. Clim., 24:3910–3923, 2011.

Chiew, F., T. Piechota, and T. McMahon. El-Niño/Southern Oscillation and Australian rainfall, streamflow and drought: Links and potential for forecasting. J. Hydrol., 204:138–149, 1998.

Chiew, F. H. S., N. J. Potter, J. Vaze, C. Petheram, L. Zhang, J. Teng, and D. A. Post. Observed hydrologic non-stationarity in far south-eastern Australia: implications for modelling and prediction. Stoch. Environ. Res. Risk Assess., 28:3–15, 2014.

Choudhury, D., A. Gupta, A. Sharma, A. Taschetto, R. Mehrotra, and B. Sivakumar. Impacts of the tropical trans-basin variability on Australian rainfall. Clim. Dyn., 49:1617–1629, 2017.

Choudhury, D., R. Mehrotra, A. Sharma, A. Gupta, and B. Sivakumar. Effectiveness of CMIP5 Decadal Experiments for Interannual Rainfall Prediction Over Australia. Water Resour. Res., 55:7400–7418, 2019.

De Fraiture, C. and D. Wichelns. Satisfying future water demands for agriculture. Agr. Water Manag., 97:502–511, 2010.

Deo, R. and M. Şahin. Application of the Artificial Neural Network model for prediction of monthly Standardized Precipitation and Evapotranspiration Index using hydrometeorological parameters and climate indices in eastern Australia. Atmos. Res., 161:65–81, 2015.

Diffenbaugh, N. S., D. L. Swain, and D. Touma. Anthropogenic warming has increased drought risk in California. Proc. Natl. Acad. Sci. U.S.A., 112:3931–3936, 2015.

Efron, B. and T. Hastie. Computer Age Statistical Inference. Cambridge University Press, 2016.

Hall, M. A. Correlation-based Feature Selection for Machine Learning. 1999.

Hartigan, J., S. MacNamara, and L. Leslie. Trends in precipitation and temperature in Canberra. In Elsawah, S. (ed.) MODSIM2019, 23rd International Congress on Modelling and Simulation. Modelling and Simulation Society of Australia and New Zealand, 2019.

Hendon, H. H., D. W. J. Thompson, and M. C. Wheeler. Australian Rainfall and Surface Temperature Variations Associated with the Southern Hemisphere Annular Mode. J. Clim., 20:2452–2467, 2007.

Henley, B., J. Gergis, D. Karoly, S. Power, J. Kennedy, and C. Folland. A tripole index for the interdecadal Pacific oscillation. Clim. Dyn., 45:3077–3090, 2015.

Holland, J., G. Luck, and C. Finlayson. Threats to food production and water quality in the Murray-Darling Basin of Australia. Ecosyst. Serv., 12:55–70, 2015.

Hsu, C.-W., C.-C. Chang, and C.-J. Lin. A Practical Guide to Support Vector Classification. Technical report, National Taiwan University, 2003. 16pp.

Hubert, P., J. Carbonnel, and A. Chaouche. Segmentation des séries hydrométéorologiques—application à des séries de précipitations et de débits de l'afrique de l'ouest. J. Hydrol., 110:349–367, 1989.

Icon Water. Water for Generations, 2019. URL `https://www.iconwater.com.au/Water-education/Water-and-sewerage-system/Water-for-Generations.aspx`.

Johnson, Z., Y. Chikamoto, J.-J. Luo, and T. Mochizuki. Ocean Impacts on Australian Interannual to Decadal Precipitation Variability. Climate, 6(3):61, 2018.

Kiem, A. and S. Franks. Multi-decadal variability of drought risk, eastern Australia. Hydrol. Process., 18:2039–2050, 2004.

Kiem, A. S., F. Johnson, S. Westra, A. van Dijk, J. P. Evans, A. O'Donnell, A. Rouillard, C. Barr, J. Tyler, M. Thyer, D. Jakob, F. Woldemeskel, B. Sivakumar, and R. Mehrotra. Natural hazards in Australia: droughts. Clim. Change, 139:37–54, 2016.

Kirono, D., F. Chiew, and D. Kent. Identification of best predictors for forecasting seasonal rainfall and runoff in Australia. Hydrol. Process., 24:1237–1247, 2010.

Lau, K.-M. and H. Weng. Climate Signal Detection Using Wavelet Transform: How to Make a Time Series Sing. Bull. Am. Meteorol. Soc., 76(12):2391–2402, 1995.

Maldonado, S. and R. Weber. A wrapper method for feature selection using Support Vector Machines. Inf. Sci., 179:2208–2217, 2009.

McBride, J. and N. Nicholls. Seasonal Relationships between Australian Rainfall and the Southern Oscillation. Mon. Weather Rev., 111:1998–2004, 1983.

Narisma, G. and A. Pitman. The Impact of 200 Years of Land Cover Change on the Australian Near-Surface Climate. J. Hydrometeorol., 4:424–436, 2003.

Niang, I., O. C. Ruppel, M. A. Abdrabo, A. Essel, C. Lennard, J. Padgham, and P. Urquhart. 2014: Africa. In: Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, 2014. 1199–1265 pp.

Post, D., B. Timbal, F. Chiew, H. Hendon, H. Nguyen, and R. Moran. Decrease in southeastern Australian water availability linked to ongoing Hadley cell expansion. Earth's Future, 2:231–238, 2014.

Power, S., T. Casey, C. Folland, A. Colman, and V. Mehta. Inter-decadal modulation of the impact of ENSO on Australia. Clim. Dyn., 15:319–324, 1999.

Prasad, R., R. Deo, Y. Li, and T. Maraseni. Input selection and performance optimization of ANN-based streamflow forecasts in the drought-prone Murray Darling Basin region using IIS and MODWT algorithm. Atmos. Res., 197:42–63, 2017.

Ramsay, H., M. Richman, and L. Leslie. Seasonal Tropical Cyclone Predictions Using Optimized Combinations of ENSO Regions: Application to the Coral Sea Basin. J. Clim., 27:8527–8542, 2014.

Richman, M. and L. Leslie. The 2015–2017 Cape Town Drought: Attribution and Prediction Using Machine Learning. Procedia Comput. Sci., 140:248–257, 2018.

Richman, M., L. Leslie, T. Trafalis, and H. Mansouri. Data Selection Using Support Vector Regression. Adv. Atmos. Sci., 32:277–286, 2015.

Richman, M. B. and L. M. Leslie. Uniqueness and Causes of the California Drought. Procedia Comput. Sci., 61:428–435, 2015.

Richman, M. B. and L. M. Leslie. Machine Learning for Attribution of Heat and Drought in Southwestern Australia. Procedia Comput. Sci., pages 3–10, 2020.

Risbey, J., M. Pook, P. McIntosh, M. Wheeler, and H. Hendon. On the Remote Drivers of Rainfall Variability in Australia. Mon. Weather Rev., 137:3233–3253, 2009.

Timbal, B. and W. Drosdowsky. The relationship between the decline of Southeastern Australian rainfall and the strengthening of the subtropical ridge. Int. J. Climatol., 33:1021–1034, 2013.

Timbal, B. and R. Fawcett. A Historical Perspective on Southeastern Australian Rainfall since 1865 Using the Instrumental Record. J. Clim., 26:1112–1129, 2013.

Torrence, C. and G. Compo. A Practical Guide to Wavelet Analysis. Bull. Am. Meteorol. Soc., 79(1):61–78, 1998.

Trenberth, K., A. Dai, G. van der Schrier, P. Jones, J. Barichivich, K. Briffa, and J. Sheffield. Global warming and changes in drought. Nat. Clim. Change, 4:17–22, 2014.

Ummenhofer, C., A. Gupta, P. Briggs, M. England, P. McIntosh, G. Meyers, M. Pook, M. Raupach, and J. Risbey. Indian and Pacific Ocean Influences on Southeast Australian Drought and Soil Moisture. J. Clim., 24:1313–1336, 2011.

Ummenhofer, C. C., M. H. England, P. C. McIntosh, G. A. Meyers, M. J. Pook, J. S. Risbey, A. S. Gupta, and A. S. Taschetto. What causes southeast Australia's worst droughts? Geophys. Res. Lett., 36, 2009.

Wilhite, D., M. Sivakumar, and R. Pulwarty. Managing drought risk in a changing climate: The role of national drought policy. Weather Clim. Extremes, 3:4–13, 2014.

Wilks, D. S. Statistical Methods in the Atmospheric Sciences, volume 100. Academic Press, 2011.

Chapter 5

Attribution and Prediction of Precipitation and Temperature Trends within the Sydney Catchment Using Machine Learning

5.1 Overview

The ongoing and intense nature of the 2017–2019 drought resulted in rapidly decreasing water storages and the lowest inflows ever recorded for the Sydney Catchment Area (SCA), leading to tight water restrictions for Australia's most populous city (Sydney Water, 2020). This drought, which affected much of Australia, culminated in the devastating Black Summer bushfires during the extended 2019-2020 fire season, with more than 24 million hectares burned, 3100 homes destroyed and 33 lives lost directly from the fires (Commonwealth of Australia, 2020; Davey and Sarre, 2020; Jalaludin and Morgan, 2021). There were additional impacts from these fires, including a large number of days where southeast Australia was shrouded in smoke, causing an estimated 4456–4753 hospitilisations and 417–429 deaths due to smoke inhalation (Beggs and Zhang, 2020; Johnston et al., 2021; Linnenluecke and Marrone, 2021; Simmons et al., 2021). Overall, this drought, and the bushfires and flooding that followed, are a clear indication of the influence that global warming is playing on Australia's climate system and the resulting extreme weather.

In order to be better prepared for future droughts, it is important to understand the strength of any changes affecting the climate system, and develop tools that can skillfully predict future weather and climate, with the ability to improve such models as more data becomes available. Following the extreme impacts the most recent drought had on the Sydney region, this Chapter presents a case study analysing trends in precipitation and temperature across the SCA. Linear and non-linear statistical models are applied to attribute climate drivers and predict these trends. This case study increases the statistical models and attribute selection techniques considered and applies them to annual, autumn and winter precipitation across a broad area rather than a single automatic weather station. Additionally, the SCA region was broken into two sub-regions to determine if there are any significant differences between trends or climate attributes. As lead author of this study, I collected the data, performed the time series and wavelet analysis, developed the statistical models, analysed the results and wrote much of the paper. Professor Lance Leslie, Dr Shev MacNamara and Dr Milton Speer helped devise the project and assisted with analysis and writing of the paper. This paper has been published in the international peer-reviewed journal: *Climate*.

5.2 Introduction

Droughts often are undetected for a period of time before they are identified (Kiem et al., 2016). Typically, they begin as precipitation deficits accumulated over time (Richman and Leslie, 2020). There are numerous non-linear effects and positive feedbacks that influence drought (Mukherjee et al., 2018), such as temperature changes, land surface changes, increases in potential evapotranspiration, higher wind speeds and reduced humidity. Furthermore, its impacts on the human and natural environment are widespread, with enhanced bushfire threat; increased flooding risk when rain falls on dry soil; reduced agricultural yields (Cheeseman, 2016); various effects on human health, such as decreased sanitation and negative mental health (Stanke et al., 2013); and greater energy use.

Many such effects have been observed across southeast Australia during the most recent drought, with destructive wildfires that affected much of the region, followed by flooding rain from east coast low events (Dowdy et al., 2019). During 2019, the Sydney catchment area (SCA; Figure 5.2.1) recorded rapidly decreasing water stores and its lowest inflows on record, leading to tight water restrictions in Australia's most populous city, Sydney (Sydney Water, 2020). Southeast Australia has experienced the effects of many historical long-term droughts, such as the Federation drought (1895–1902; Bureau of Meteorology, 2009a), the World War II drought (1937–1945; Bureau of Meteorology, 2009b) and the Millennium Drought (1997–2009; Chiew et al., 2014). Regions with similar latitudes also have recently experienced severe drought, including California (Richman and Leslie, 2015) and Cape Town (Richman and Leslie, 2018). It is expected that drought will be more frequent and severe in these regions with the global warming trend (Niang et al., 2014; Diffenbaugh et al., 2015). Likewise, precipitation has been decreasing along the east coast of Australia as a result of the poleward expansion of the Hadley cell (Speer et al., 2011; Timbal and Fawcett, 2013; Timbal and Drosdowsky, 2013; Post et al., 2014).

The population of Sydney continues to increase (Australian Bureau of Statistics, 2019), which will place further strain on water resources. Therefore, it is important to analyze any potential trends in variables that impact water availability, including precipitation and temperature. Studies utilising non-linear statistical models in the atmospheric sciences are relatively limited; however, they can help improve understanding of the complex relationships between climate indices and atmospheric variables. Additionally, development of accurate models on a range of time scales will help inform water management (Choudhury et al., 2019), thereby improving water security. This study will focus on identifying trends in precipitation and temperature across the SCA. Statistical models are then used to attribute specific climate drivers and predict annual precipitation, and precipitation over autumn and winter, as these techniques have performed well in previous studies predicting climatic variables (e.g., Deo and Şahin, 2015; Bagirov et al., 2017; Richman and Leslie, 2018, 2020) but are seldom applied in Australia, and to the author's knowledge, have not been performed on precipitation in the SCA.

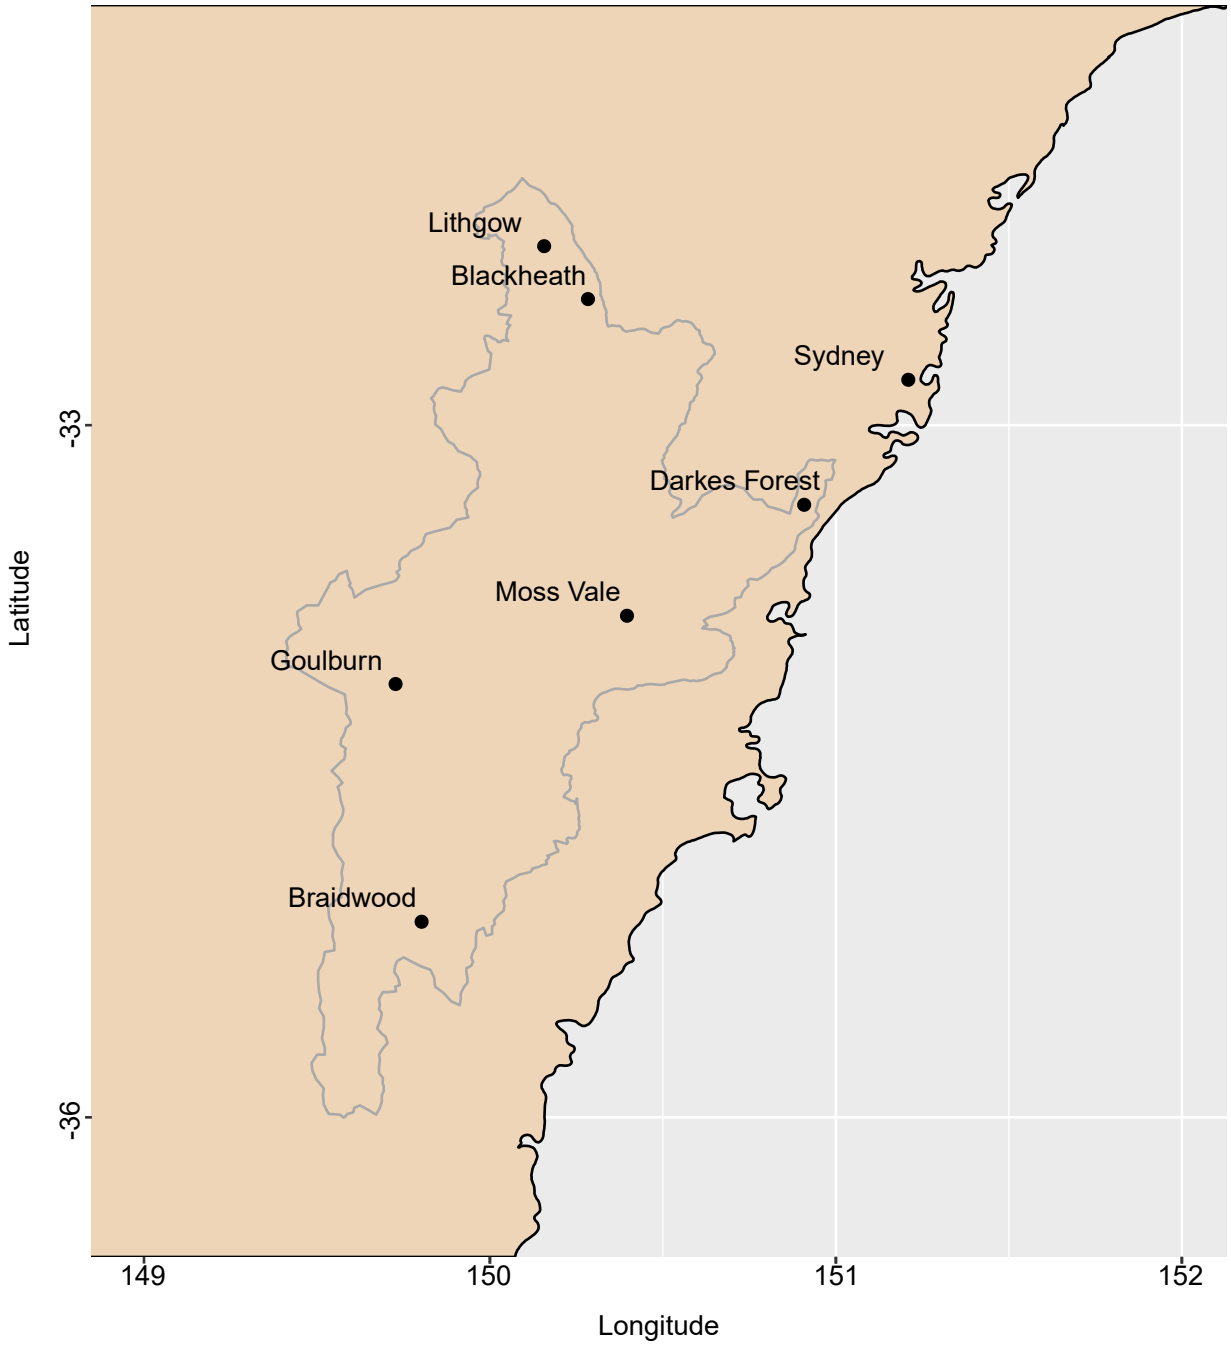

Figure 5.2.1: Location map of the Sydney Catchment Area (outline in dark grey): the six sites where data were obtained and the Sydney central business district (black dots).

5.3 Data and Methodology

5.3.1 Data

Monthly precipitation data were obtained for Blackheath, Braidwood, Darkes Forest, Goulburn, Lithgow and Moss Vale from the publicly available Climate Data Online provided by the Australian Bureau of Meteorology (BoM) website (`http://www.bom.gov.au/climate/data/index.shtml`). Similarly, monthly mean maximum temperature (TMax) and mean minimum temperature (TMin) data were obtained for Goulburn, Lithgow and Moss Vale. Sites were selected within the greater SCA (Figure 5.2.1) to include relatively continuous, long-term records. Any missing data points, which accounted for approximately 8% of the data, were filled using the running mean centred on the year of the missing record. The time series data covered the

period from 1957 to 2019. The precipitation data were summed to obtain a proxy to total precipitation falling across the entire catchment, the northern region of the catchment (containing Lithgow, Blackheath and Darkes Forest) and the southern region of the catchment (containing Moss Vale, Goulburn and Braidwood). The mean of the three temperature time series data was taken to obtain a mean of TMax and TMin across the SCA. Temperature data were not further split into northern and southern catchment regions, as there are only three stations, with Lithgow being the only station in the northern region of the SCA.

5.3.2 Statistical Analysis

To analyse trends present in the time series, the data were first plotted along with their percentiles, both annually and over the four seasons. These data were then grouped into two 31-year periods (1958–1988 and 1989–2019) covering the length of the time series. Bootstrap resampling was applied with replacement, using 5000 resamples, to improve understanding of any trends present in the mean and variance of the time series data. Two-sided permutation testing was then applied to evaluate statistical significance of any differences between the two time periods. This technique is similar to bootstraping; however, it builds the sampling distribution by permuting the observed data without replacement. The difference between the two distributions can then be tested, where the null hypothesis is that there is no difference between the two distributions. In this study, a significance level of $\alpha = 0.1$ was considered to represent a statistically significant result.

5.3.3 Attribute Selection and Prediction

Initially, wavelet analysis was applied on the entire detrended time series anomalies following the approach of Torrence and Compo (1998). The local wavelet power spectra display how periodic signals within the time series evolve over time, allowing for the detection of potential climate drivers, such as the El-Niño Southern Oscillation (ENSO), and analysis of how their influence might change over time. Wavelets were used as they efficiently resolve both high- and low-frequency signals by stretching and translating local base functions in both space and time (Lau and Weng, 1995). The Morlet wavelet was used in this study, which is described by

$$\psi_0(\eta) = \pi^{-1/4} e^{i\omega_0 \eta} e^{-\eta^2/2}, \tag{5.3.1}$$

where η is a non-dimensional time parameter and ω_0 is the non-dimensional frequency. The continuous wavelet for a time series, x_n, is given by

$$W_n(s) = \sum_{n'=0}^{N-1} x_{n'} \psi^* \left(\frac{(n'-n)\delta t}{s} \right), \tag{5.3.2}$$

where n is a localised time index, s is the wavelet scale, N is the number of points in the time series and $*$ is the complex conjugate. For computational efficiency, the wavelet transform is computed in Fourier space using the discrete Fourier transform of x_n:

$$\hat{x}_k = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i k n/N}, \tag{5.3.3}$$

where $k = 0,\ 1, \cdots, N-1$ is the frequency index. The Fourier transform of $\psi(t/s)$ is given by $\hat{\psi}(s\omega)$. Using the convolution theorem, and taking the Fourier transform, the wavelet transform in Fourier space is given by

$$W_n(s) = \sum_{k=0}^{N-1} \hat{x}_k \hat{\psi}^*(s\omega_k) e^{-\omega_k n \delta t}, \tag{5.3.4}$$

where

$$\omega_k = \begin{cases} \frac{2\pi k}{N\delta t} & : \quad k \leq \frac{N}{2} \\ \\ \frac{-2\pi k}{N\delta t} & : \quad k > \frac{N}{2}. \end{cases} \tag{5.3.5}$$

Of particular interest is the local wavelet power spectrum, defined by

$$P(s) = |W_n(s)|^2 , \tag{5.3.6}$$

which displays the time evolution of periodic signals in a time series. Additionally, the global power spectrum, which is equivalent to applying a Fourier transform on the time series, provides an overview of the dominant amplitudes within the time series, and is given by

$$\overline{W}^2(s) = \frac{1}{N} \sum_{n=0}^{N-1} |W_n(s)|^2 . \tag{5.3.7}$$

A range of climate drivers was considered to select important attributes for prediction of annual precipitation, and precipitation during autumn (March–May) and winter (June–August). The complete list of predictors considered included: the Atlantic Multidecadal Oscillation (AMO), the Dipole Mode Index (DMI), the global sea surface temperature anomalies (GlobalSSTA) the global temperature anomalies (GlobalT), Niño3.4, the Tripole Index for the Interdecadal Pacific Oscillation (TPI), the Southern Annular Mode (SAM) and the Southern Oscillation Index (SOI), obtained from the Earth System Research Laboratory (`http://www.esrl.noaa.gov/psd/gcos_wgsp/Timeseries/`). Tasman Sea surface temperature anomalies (TSSST) were obtained from the BoM (`http://www.bom.gov.au/climate/change/?ref=ftr#tabs=Tracker&tracker=timeseries`) and considered as an additional potential climate driver. All data are from 1957–2019.

Two-way interaction terms between the above predictors also were considered as potential predictors, obtained by multiplying one variable with another (e.g., AMO*DMI). These predictors can potentially be influential in statistical models, as one predictor might reinforce another, and they have been selected in past statistical modelling studies (e.g., Richman and Leslie, 2018, 2020; Hartigan et al., 2020). There are additional climate drivers that could have been considered in this study; however, many of these climate drivers are correlated with those already considered, and all ocean basins have been covered by those considered here. There also are possible alternative relationships between climate drivers, such as an inverse relationship, that have not been considered in this study, as it would require a significant amount of testing to choose between all possible relationships.

Figure 5.3.1 displays scatter plots of annual precipitation across the SCA, with a linear fit and the correlation coefficient between the climate drivers and annual precipitation. Some of the climate drivers exhibit moderate linear relationships with annual SCA precipitation (e.g., Niño3.4 has a correlation of -0.38, Figure 5.3.1), which suggests multiple linear regression (LR) might perform well in predicting annual precipitation. However, not all relationships between the climate drivers are strong, and some appear to be non-linear. As such, this study develops both linear and non-linear statistical models for annual precipitation. The non-linear models considered are support vector regression (SVR) (Vapnik, 1995), with both polynomial (Poly) and radial basis function (RBF) kernels, and random forests (RF) (Breiman, 2001; Hastie et al., 2009). Precipitation during autumn is important as autumn rainfall saturates the soil, allowing runoff to occur in the following winter and spring rainfall (Cai and Cowan, 2008; Cai et al., 2009). Furthermore, runoff is typically greatest in this region of Australia during the cooler winter months (Kirono et al., 2010). Therefore, this study also develops models of autumn and winter precipitation to investigate which climate drivers have greater influence on precipitation during these seasons, and whether improvements in predictability could be achieved.

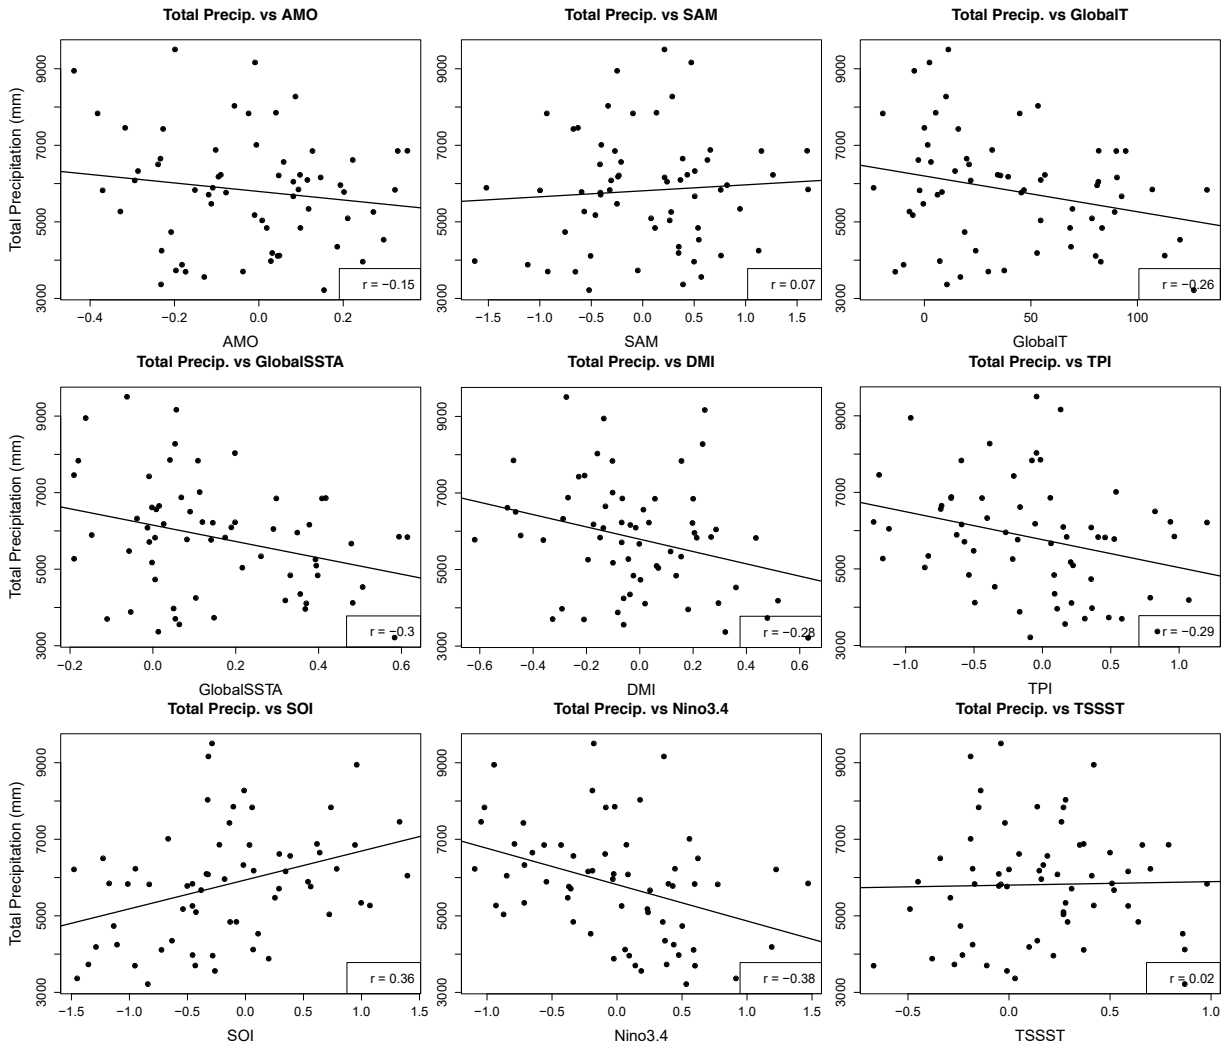

Figure 5.3.1: Scatter plot of annual precipitation across the Sydney Catchment Area (SCA) against each climate driver that serves as a possible attribute. The thin black line shows a linear fit against the data using least squares regression. Correlations of each predictor against precipitation are provided in the box to the lower right of each sub-panel; low correlations ($|r| < 0.3$) indicate a weak linear relationship between a predictor and precipitation.

Although the number of climate drivers has been reduced through selecting a subset to ensure that all ocean basins were covered, when including the two-way interaction terms there were 45 potential attributes (Table 5.1) considered in this study to predict annual, autumn and winter precipitation. Using all of these attributes would result in an overfit model that can reduce physical understanding underlying the predictions made and lead to large errors when applying the model to the test data set. This study used 85% of the data (1957–2010) to train each model. To select attributes that generalise well to the data, ten-fold cross validation was applied to the training set using forward selection through the space of potential attributes, with a similar approach to Maldonado and Weber (2009).

Forward selection does not necessarily result in the desired outcome of an optimal combination of predictors that leads to a parsimonious model with low error on the test data set. In this study, an additional selection technique was applied where models were trained individually on precipitation falling in either the southern catchment area or the northern catchment area using forward selection. Two models for precipitation over the entire SCA were then developed using only the variables selected in the models for the southern catchment and northern catchment areas, separately. Each of these models were compared against the model developed using

forward selection over the entire SCA.

For both the RBF SVR and Poly SVR kernels, the cost parameter (C) needs to be selected, which specifies the cost of violating constraints in the optimization problem. In addition, for RBF SVR the gamma value (G) needs to be selected, and the exponent (E) needs to be selected for the Poly kernel. These free parameters influence how strongly the model fits to the training data and often are selected using a grid search (Hsu et al., 2003). In the final model fitting step of this study, values of C from $2^{0.5}$ to 2^3 were tested with powers varying by 0.5; values of G ranged from 0 to 1, varying by 0.25; and values of E tested were 1, 2 and 3. For the random forests, the number of trees in each forest was tuned with values tested ranging from 50 to 800, varying by 50.

Table 5.1 lists the percentage of folds each attribute appeared in for the models trained on annual, autumn and winter precipitation. The SVR models used in the development of these tables did not have C, G or E parameters tuned, in order to decrease computation time. Instead, default values of 1, 1/(number of predictors) and 3 were used, respectively. Attributes that appeared in at least 5 of 10 folds were retained for selection using LR, SVR and RF methods. The model that obtained the highest correlation and skill, and lowest root mean square error (RMSE) against the observed training data was selected and used to predict annual, autumn and winter precipitation on the testing data set (2011–2019), where RMSE is defined according to Wilks (2011):

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{k=1}^{n}(y_k - o_k)^2}, \tag{5.3.8}$$

where n is the number of paired forecast (y_k) and observed (o_k) values. The Pearson correlation coefficient is used in this study, defined as:

$$r_{yo} = \frac{\text{Cov}(y, o)}{s_y s_o}, \tag{5.3.9}$$

where Cov(y, o) is the covariance of forecasts (y) against observations (o), and s represents the standard deviation. Skill is defined as:

$$\text{Skill} = 1 - \frac{MSE_{\text{model}}}{MSE_{\text{climatology}}}, \tag{5.3.10}$$

where MSE_{model} is the mean square error (MSE, the square of RMSE) of the model against observations, and $MSE_{\text{climatology}}$ is the MSE between the mean of the climate from 1957 up to the year prior to the observation, and the year of the observation.

Table 5.1: Percentage of folds selecting attributes for precipitation in linear regression (LR); support vector regression (SVR) with the radial basis (RBF) and polynomial (Poly) kernel functions; and random forests (RF).

Attribute	LR	SVR (RBF)	SVR (Poly)	RF	Attribute	LR	SVR (RBF)	SVR (Poly)	RF
AMO	30	100	80	40	DMI*TSSST	90	80	60	20
DMI	0	70	80	30	GlobalSSTA*GlobalT	0	60	80	10
GlobalSSTA	0	80	80	30	GlobalSSTA*Niño3.4	10	60	80	40
GlobalT	80	40	90	20	GlobalSSTA*TPI	20	60	80	40
Niño3.4	60	30	90	20	GlobalSSTA*SAM	0	90	70	30
TPI	10	60	80	30	GlobalSSTA*SOI	10	60	90	30
SAM	20	80	90	20	GlobalSSTA*TSSST	0	70	60	30
SOI	60	60	80	40	GlobalT*Niño3.4	30	40	90	30
TSSST	10	30	90	10	GlobalT*TPI	0	60	90	20
AMO*DMI	60	80	70	40	GlobalT*SAM	0	100	80	50
AMO*GlobalSSTA	10	90	80	40	GlobalT*SOI	10	50	80	30
AMO*GlobalT	10	80	70	40	GlobalT*TSSST	10	70	50	50
AMO*Niño3.4	20	50	70	30	Niño3.4*TPI	0	40	70	30
AMO*TPI	0	70	50	30	Niño3.4*SAM	40	80	70	30
AMO*SAM	30	80	90	40	Niño3.4*SOI	0	40	60	30
AMO*SOI	10	80	90	30	Niño3.4*TSSST	0	40	70	30
AMO*TSSST	0	50	60	20	TPI*SAM	30	70	80	30
DMI*GlobalSSTA	10	10	80	70	TPI*SOI	10	50	70	30
DMI*GlobalT	10	30	80	50	TPI*TSSST	40	50	70	20
DMI*Niño3.4	50	40	90	40	SAM*SOI	0	80	60	40
DMI*TPI	40	60	70	40	SAM*TSSST	20	100	90	50
DMI*SAM	30	80	80	30	SOI*TSSST	10	40	60	40
DMI*SOI	20	70	80	60					

5.4 Results

5.4.1 Evolution of Precipitation

Time series of annual precipitation and box plots of bootstrapped mean annual precipitation over the 31-year periods 1958–1988 and 1989–2019, for the SCA, are shown in Figure 5.4.1. Data presented in Figure 5.4.1 are annual (a,b) and cover the four seasons (c–j). In addition, p-values from permutation testing comparing the means and variances for the period 1958–1988 against 1989–2019 are presented in Table 5.2.

A significant decline in annual precipitation is apparent across the catchment, with a reduction in years recording above the 75th percentile since the 1990s, with only two years recording above the 75th percentile since the 2000s. Additionally, there is an increase in the frequency of years recording below the 25th percentile (Figure 5.4.1a). The most recent dry period is pronounced, with annual precipitation recording below the 25th percentile for the last three years, and 2019 in particular recording the lowest total annual precipitation over the observation period. This decline in annual precipitation is apparent in the 31-year box plots, exhibiting a statistically significant decrease in precipitation from a mean of approximately 6300 mm in 1958–1988 to a mean of 5400 mm during 1989–2019 (p-value = 0.0258; Figure 5.4.1b and Table 5.2). This decrease in annual mean precipitation also is accompanied by a statistically significant reduction in variance, as evidenced by the reduced lengths of the interquartile range and whiskers on the box plots (p-value = 0.028; Table 5.2).

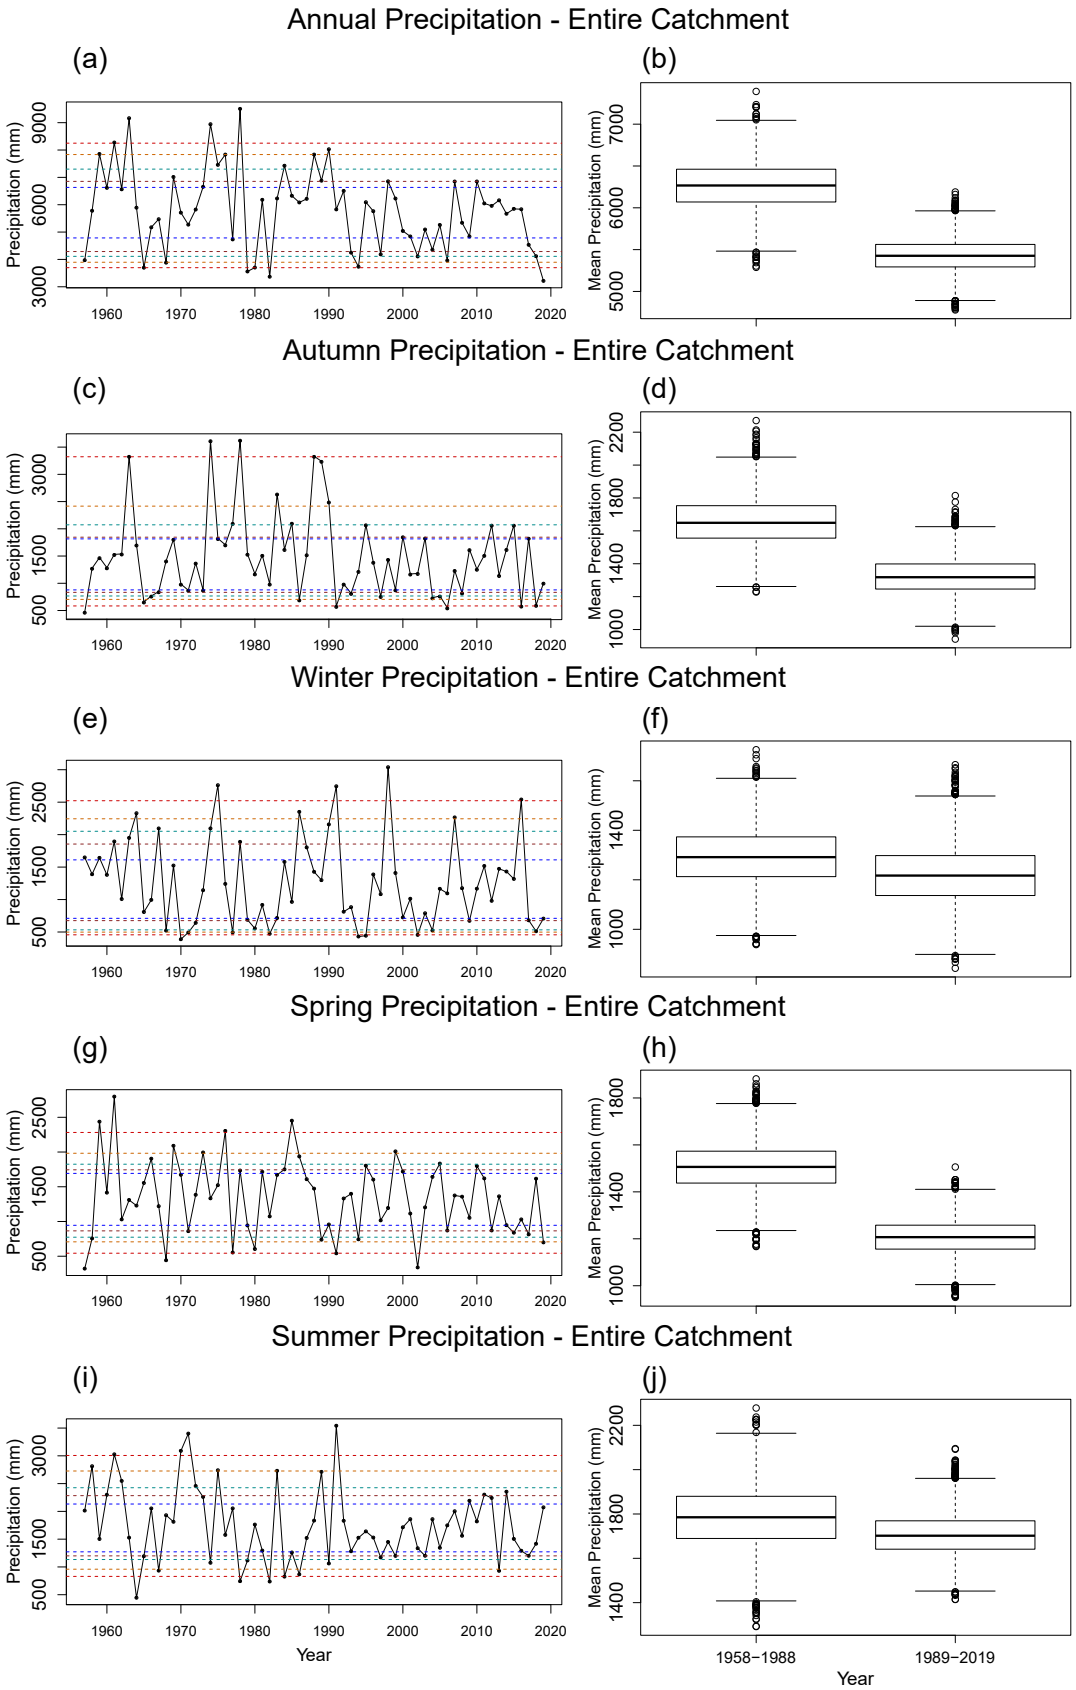

Figure 5.4.1: (left panels, **a**–**i**) Annual precipitation time series summed across all six sites. Dashed lines indicate the 5th and 95th (bottom and top red), 10th and 90th (bottom and top orange), 15th and 85th (bottom and top green), 20th and 80th (bottom and top brown) and 25th and 75th percentiles (bottom and top dark blue); and (right panels, **b**–**j**) box plots of the bootstrapped mean annual precipitation over 31-year periods for the SCA.

Table 5.2: p-values for the permutation test on the difference between the mean and variance of precipitation, TMax or TMin between 1958–1988 and 1989–2019, annually and across all four seasons. Tests conducted used 5000 resamples. Text in bold face highlights statistical significance at the 95% confidence level; text in italics highlights statistical significance at the 90% confidence level.

Years	**Precipitation**	**TMax**	**TMin**
		Annual	
Mean	**0.0258**	**0.0036**	**0.0000**
Variance	**0.0280**	0.2570	0.2300
		Autumn	
Mean	*0.0778*	0.3600	0.2470
Variance	0.2180	*0.0914*	0.8190
		Winter	
Mean	0.6850	**0.0226**	**0.0004**
Variance	0.8430	0.9290	0.7860
		Spring	
Mean	**0.0228**	**0.0070**	**0.0002**
Variance	*0.0802*	0.5840	0.2680
		Summer	
Mean	0.6390	**0.0440**	**0.0000**
Variance	**0.0438**	**0.0408**	0.3980

The decreasing precipitation trend also is apparent for autumn (March–May) precipitation, with a statistically significant difference in mean precipitation between the two 31-year periods, greater than the 90% confidence level (p-value = 0.0778; Table 5.2 and Figure 5.4.1d). This reduction is evident in the time series, with no years recorded above the 90th percentile since 1990 and an increase in years recording below the 25th percentile since 1990 (Figure 5.4.1c). There appears to be a reduction in precipitation variance, with reduced interquartile range in the box plots; however, this is not statistically significant with a p-value = 0.218 (Table 5.2). In contrast, winter (June–August) precipitation appears relatively stable. There is a decrease in the number of years recording above the 75th percentile since the 1990s, although all but one of the years with winter precipitation above the 95th percentile occured during this period (Figure 5.4.1e). There also is little change to the frequency of years recording below the 25th percentile. This corresponds to only a slight decrease in winter precipitation in the 31-year period 1989–2019 that is not statistically significant (p-value = 0.685; Table 5.2 and Figure 5.4.1f).

There is another statistically significant decrease in spring (September–November) precipitation (p-value = 0.0228; Table 5.2 and Figure 5.4.1h). This is consistent with a reduction in years above the 75th percentile, and particularly years above the 90th percentile, from the 1990s (Figure 5.4.1g). There has also been an increased frequency of years below the 25th percentile since this time; however, years below the 10th percentile have reduced in frequency. This has resulted in a statistically significant reduction in variance of spring precipitation (p-value = 0.0802; Table 5.2 and Figure 5.4.1h). Meanwhile, there has been little change to mean summer precipitation (Figure 5.4.1j). However, the extremes of summer precipitation occurred less frequently from the 1990s with only one year recorded above the 95th percentile and one year below the 10th percentile in the last 30 years (Figure 5.4.1a). As such, the variance of summer precipitation is much reduced, with little variation around the mean (p-value of difference in variance between 1958–1988 and 1989–2019 = 0.0438; Table 5.2).

When splitting precipitation into the northern and southern catchment regions, it is apparent that the precipitation trends observed over the entire SCA are occurring in both regions

(not shown). However, the trends in the northern catchment are not as substantial. The overall trend of decreasing annual precipitation (p-value = 0.0258; Table 5.2) over the SCA will increase pressure on water resources, especially as the population continues to increase. Although winter precipitation remains stable, the statistically significant reduction in autumn precipitation (p-value = 0.0778) is particularly of concern, as autumn precipitation is necessary to saturate the soil for winter and spring runoff (Cai and Cowan, 2008; Cai et al., 2009). The cooler months of the year are those when runoff is greatest in southeast Australia (Kirono et al., 2010), and the decreasing trends in both autumn and spring precipitation suggest a trend to much reduced runoff in the future.

5.4.2 Evolution of Temperature

Figure 5.4.2 provides the time series and box plots of TMax annually and over the four seasons, while Figure 5.4.3 provides this overview for TMin. Understanding changes in temperature is important, as temperature influences potential evapotranspiration (Cai and Cowan, 2008; Cai et al., 2009; Kirono et al., 2010; Chiew et al., 2014), and there is a non-linear relationship between temperature and precipitation with higher mean temperatures reducing precipitation during drought (Richman and Leslie, 2018).

Annual TMax displays a clear increasing trend with a reduction in years below the 25th percentile since the 2000s and a corresponding increase in years above the 75th percentile (Figure 5.4.2a). This trend is statistically significant, with a p-value = 0.0036 (Table 5.2, Figure 5.4.2b). The last three years stand out, with TMax recorded above the 90th percentile in 2017, and the 95th percentile in 2018 and 2019 (Figure 5.4.2a). In particular, 2019 recorded the warmest catchment-average TMax over the period 1957–2019 of approximately 20.6 °C, which is 1.7 °C above the mean for the period of approximately 18.9 °C. This increasing trend in TMax occurs across all seasons except autumn, where the increase in mean is not statistically significant (p-value = 0.36; Table 5.2 and Figure 5.4.2d). However, variance of mean TMax appears to have increased in autumn with a greater interquartile range in the 1989–2019 box plot. This change in autumn variance is statistically significant at the 90% confidence level (p-value = 0.0914; Table 5.2). There are similar, statistically significant, changes to variance of TMax in summer as displayed in the box plots (p-value = 0.0408; Table 5.2 and Figure 5.4.2j).

There is a noticeable increasing trend in annual TMin, with all years above the 90th percentile occurring since the late 1990s and a reduction in years recorded below the 25th percentile (Figure 5.4.3a). The change in annual TMin is statistically significant (p-value = 0.0000; Table 5.2), which is highlighted in the box plot figures where there is no overlap between the two 31-year periods (Figure 5.4.3b). Similar to TMax, the increase in TMin is statistically significant for all seasons except autumn (autumn p-value = 0.2470, while p-value for all other seasons < 0.0004; Table 5.2 and Figure 5.4.3d). There are hints that autumn TMin might be starting to increase, with a reduction in the number of years recording below the 25th percentile since the 1990s, except for the late-2000s and early-2010s, and an upward shift in the 1989–2019 box plot with no statistically significant change in variance (Figure 5.4.3c).

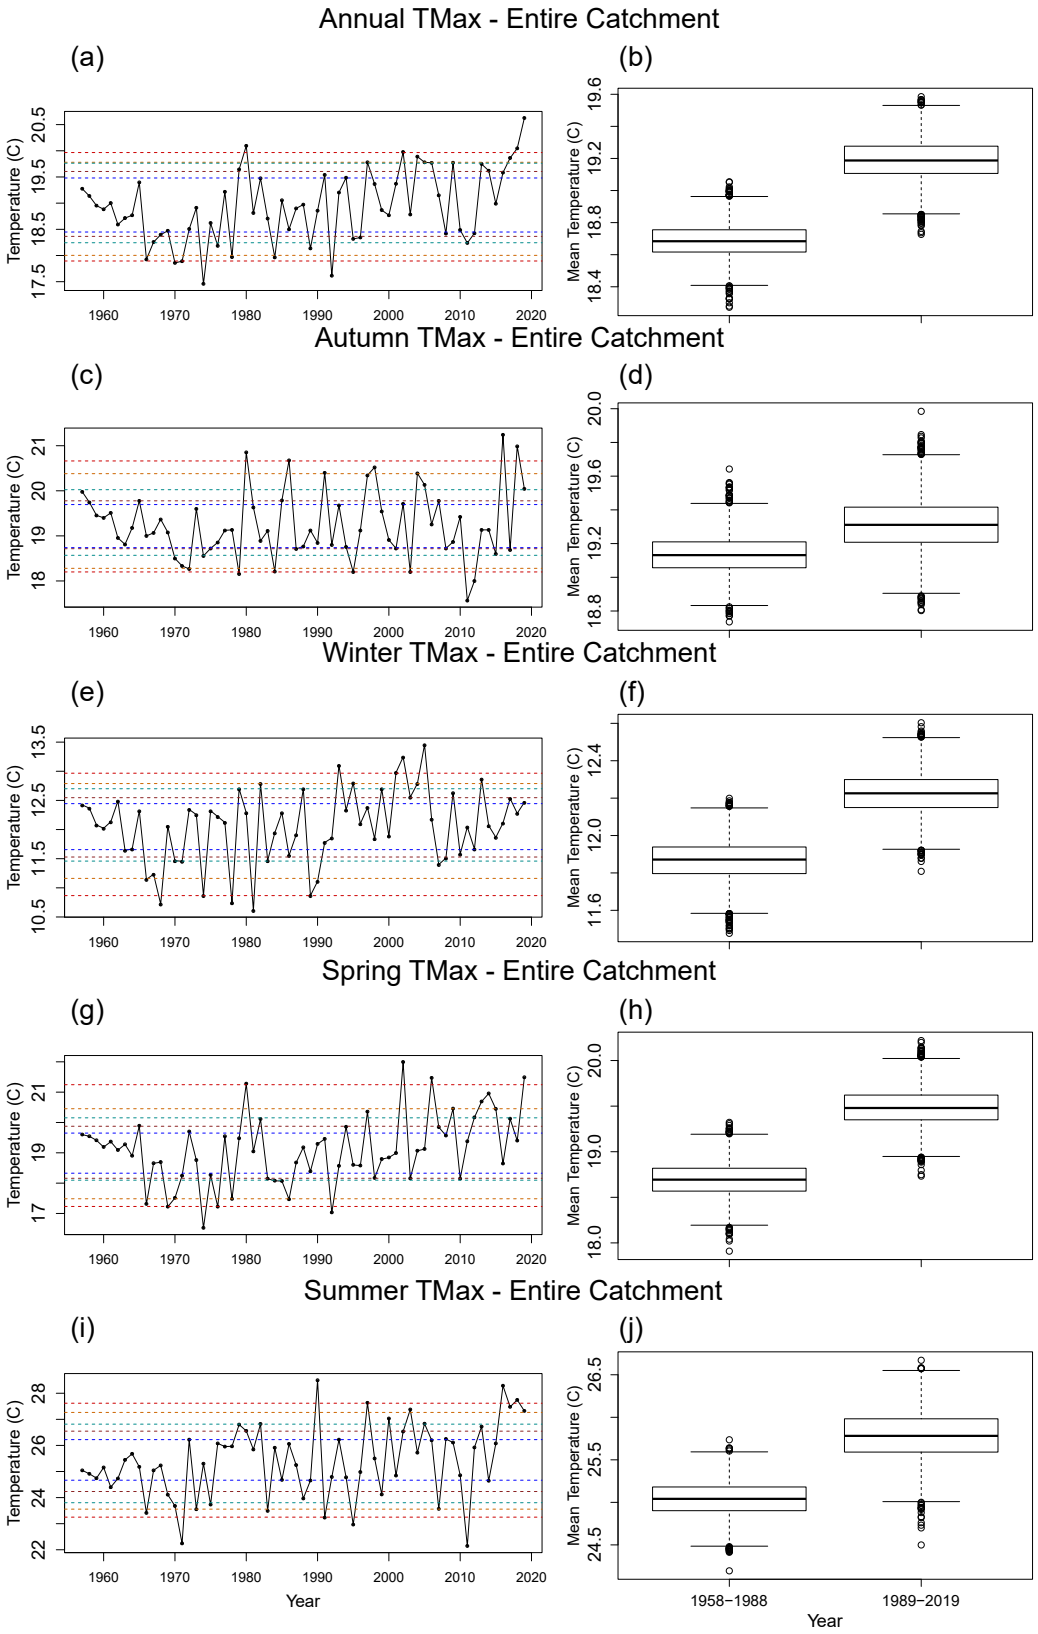

Figure 5.4.2: (left panels, **a**–**i**) Annual mean TMax time series across the SCA; and (right panels, **b**–**j**) box plots of the bootstrapped mean TMax over 31-year periods for the SCA, as in Figure 5.4.1 but for mean TMax.

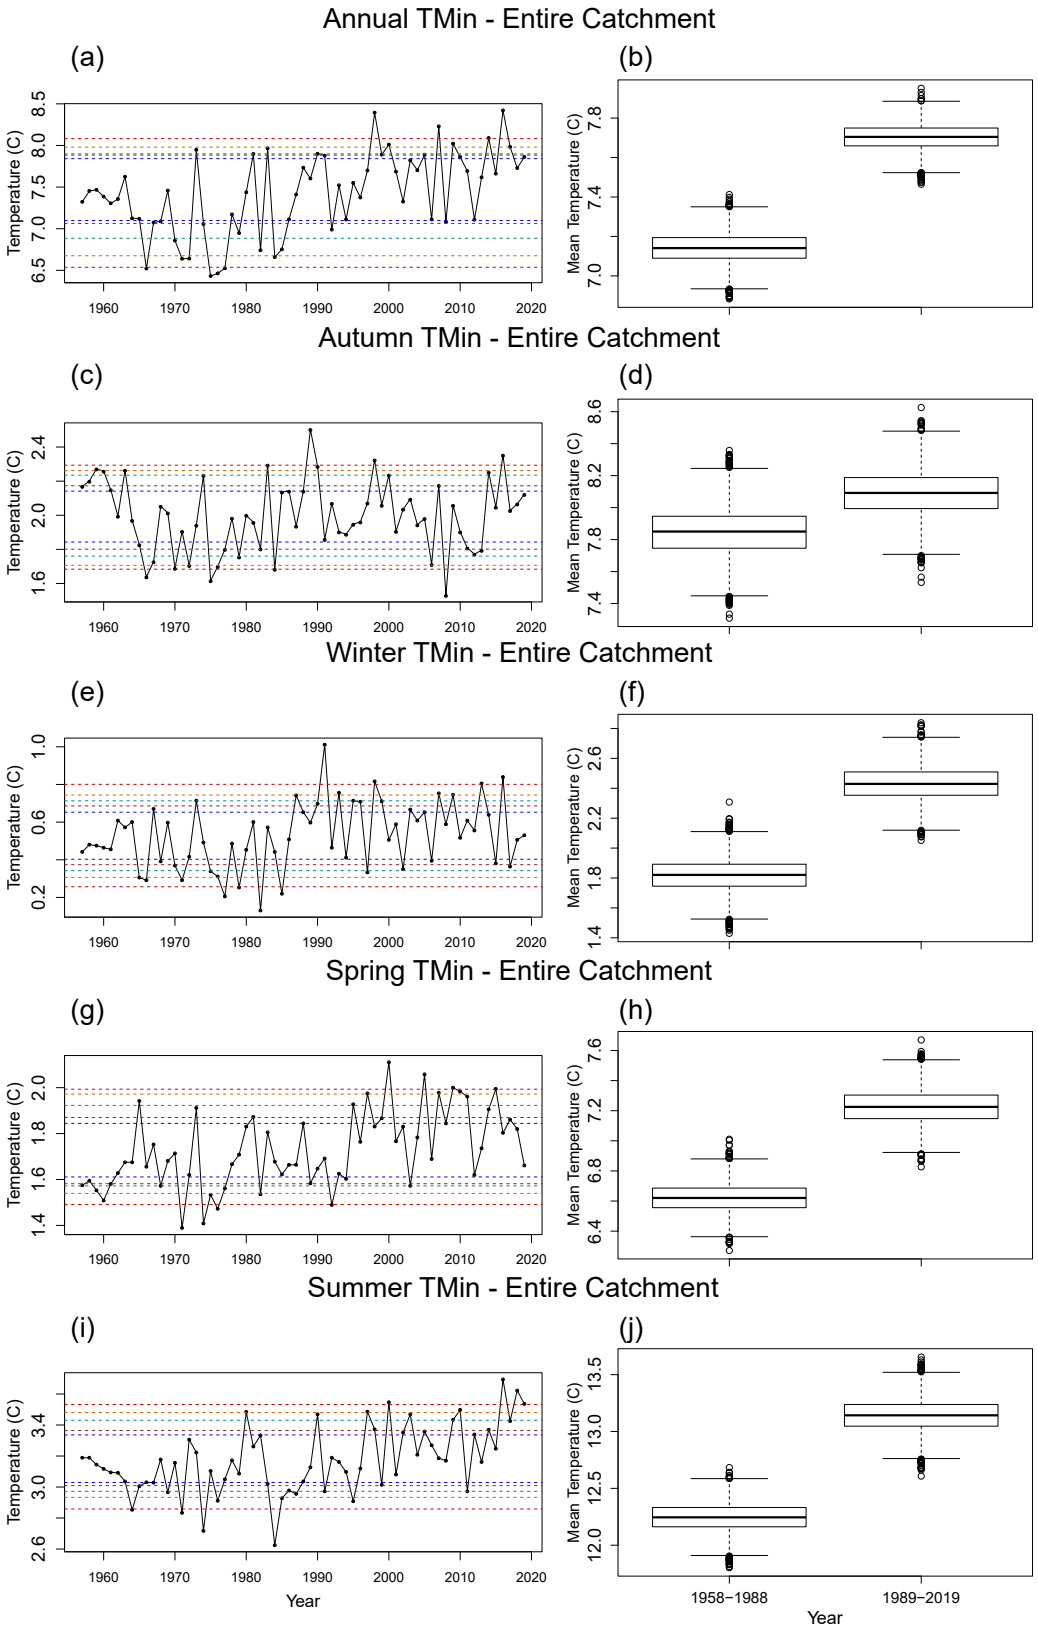

Figure 5.4.3: (left panels, **a**–**i**) Annual mean TMin time series across the SCA; and (right panels, **b**–**j**) box plots of the bootstrapped mean TMin over 31-year periods for the SCA, as in Figure 5.4.1 but for mean TMin.

Since both TMax and TMin are increasing across all seasons except autumn, it appears that mean temperatures are increasing across the SCA. This suggests further reductions in runoff and increased evapotranspiration even without the reduction in annual precipitation. Although little change in autumn mean temperature suggests evapotranspiration will remain

much the same during this season, the reduced precipitation during autumn will influence runoff later in the year which will be further compounded by increases in mean temperature.

5.4.3 Wavelet Analysis of Precipitation

Global and local wavelet power spectra for precipitation in the SCA both annually and over the four seasons are provided in Figure 5.4.4. Annual precipitation displays statistically significant high global power over the 2–4 year period, which is suggestive of an ENSO influence on annual precipitation (Figure 5.4.4b). The ENSO signal is apparent to the late-1980s; however, it has weakened in the most recent 30 years. There also is indication of an interdecadal influence on annual precipitation with high global power over the 8–16 year period, which is partially in the 90% confidence band. This is reflected in the local wavelet power spectra where there is statistically significant power over the 8–16 year period between the 1960s and 1990s, and weak power from the 1990s onwards.

The ENSO influence on precipitation is apparent in all seasons globally, with statistically significant power over the 2–4 year period (Figure 5.4.4d–j). This influence is maintained throughout the time series for winter and spring (Figure 5.4.4e,g). However, the influence has greatly weakened from the 2000s in both autumn and summer precipitation (Figure 5.4.4c,i). The 8–16 year periodic signal is strongest in winter, with high power between 1970–2010 causing a statistically significant peak in the global power spectra (Figure 5.4.4e,f). There also is some high power over this period in the other seasons, although this is generally not statistically significant except between the late-1970s and early-1990s in spring (Figure 5.4.4g), and the 1960s to early-1970s in summer, the latter of which is outside the cone of influence (Figure 5.4.4i).

5.4.4 Wavelet Analysis of Temperature

Local and global wavelet power spectra were also computed for TMax and TMin and are presented in Figures 5.4.5 and 5.4.6, respectively. There is statistically significant power over the 2–4 year period for TMax, which remains throughout the time series (Figure 5.4.5). This is visible throughout all seasons and displays little temporal change, although the signal only appears in autumn TMax from the late-1970s (Figure 5.4.5c). Additionally, the power for TMax in summer increases from the 1980s (Figure 5.4.5i).

Similarly, there is statistically significant power over the 2–4 year period for annual TMin, although this appears weaker than TMax over the global power spectrum (Figures 5.4.5b and 5.4.6b). This also is noticeable in the local power spectra, with less continuity of statistically significant periodicities throughout the time series (Figure 5.4.6a). Annual TMin also displays a statistically significant signal in the 8–12 year period between the late-1970s and mid-1990s, which is suggestive of some interdecadal influence on annual TMin. This is nearly strong enough to reach the 90% confidence level in the global power spectrum (Figure 5.4.6b). High power over the 2–4 year period occurs over all four seasons (Figure 5.4.6d–j), although it is non-existent during summer before the 1970s (Figure 5.4.6i). The interdecadal signal in annual TMin appears mostly influenced by a statistically significant periodicity during summer over a similar time period (Figure 5.4.6a,i).

Overall, it appears that both TMax and TMin are influenced by ENSO although the ENSO influence on TMin is weaker. This influence generally holds throughout all seasons and has only small temporal changes, mainly during autumn and summer.

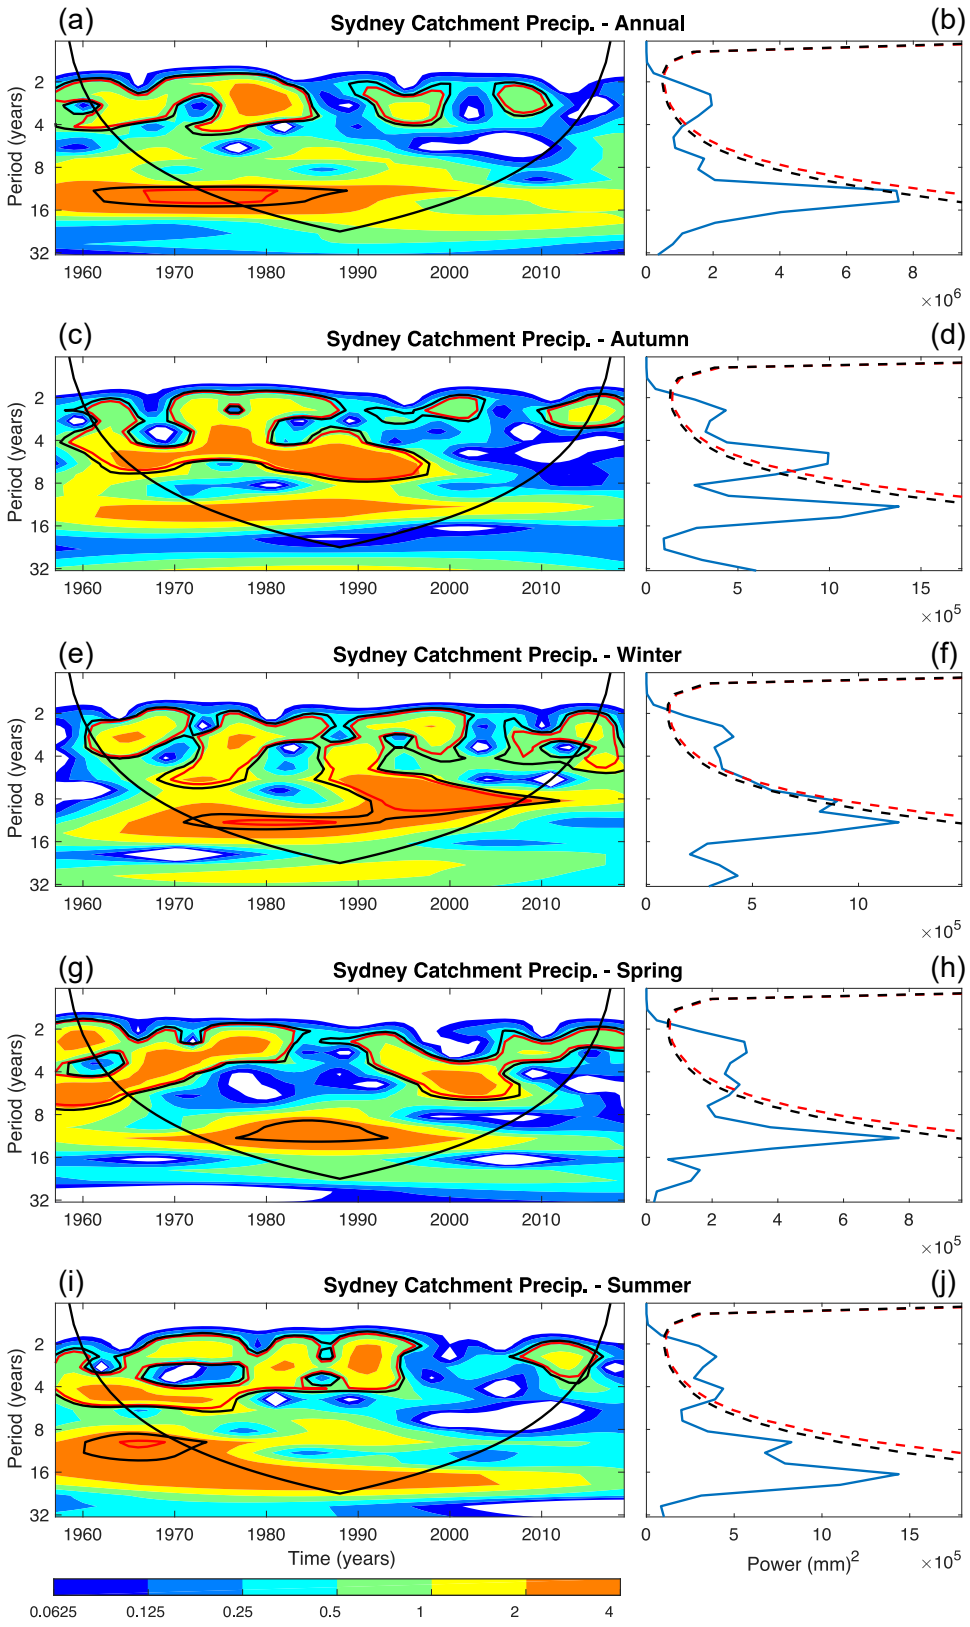

Figure 5.4.4: Wavelet analysis for total precipitation across the SCA both annually (**a**–**b**) and over the four seasons (**c**–**j**). Low values (blue) in the wavelet power spectrum (left panels, **a**–**i**) indicate low periodicity, while high values (orange) indicate high periodicity. Horizontal peaks in the global power spectrum (right panel, **b**–**j**) indicate high periodicity over the time series. The dashed/solid red (black) line indicates the 95% (90%) confidence level in each plot. The cone of influence is depicted in the wavelet power spectrum by the solid black cone.

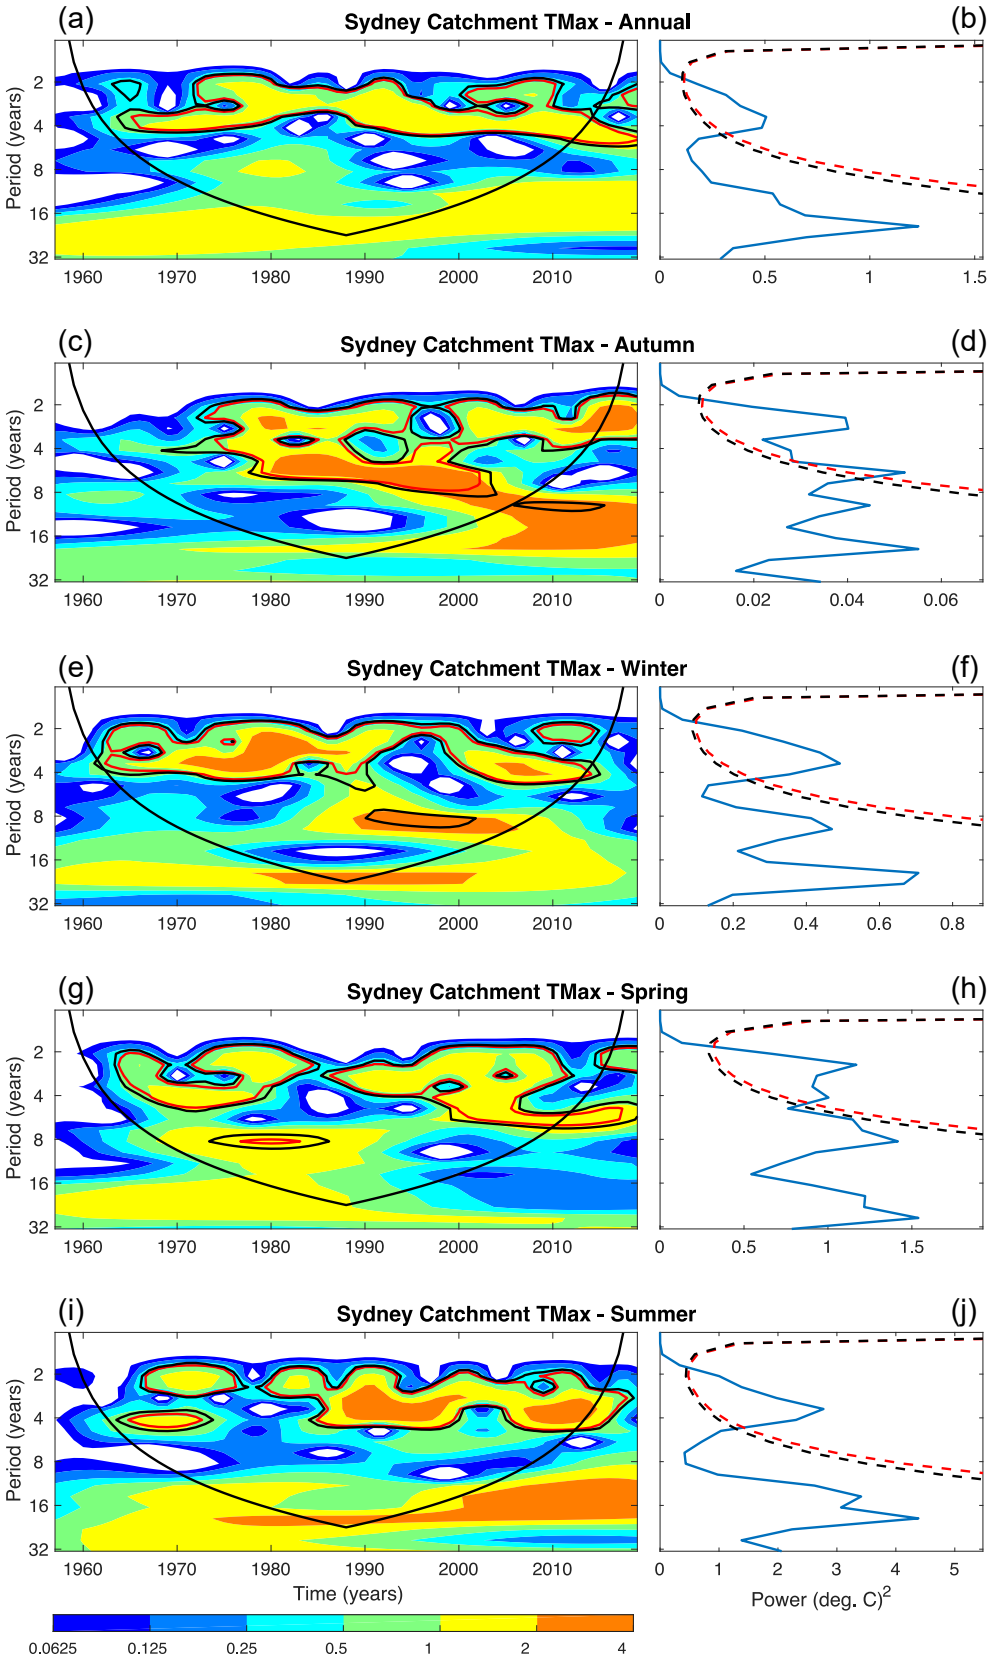

Figure 5.4.5: Wavelet power spectra (left panel, **a**–**i**) and global power spectra (right panel, **b**–**j**) for TMax, with low and high values in the power spectra indicating low and high periodicity, respectively; the cone of influence and statistical significance bands as already explained in Figure 5.4.4.

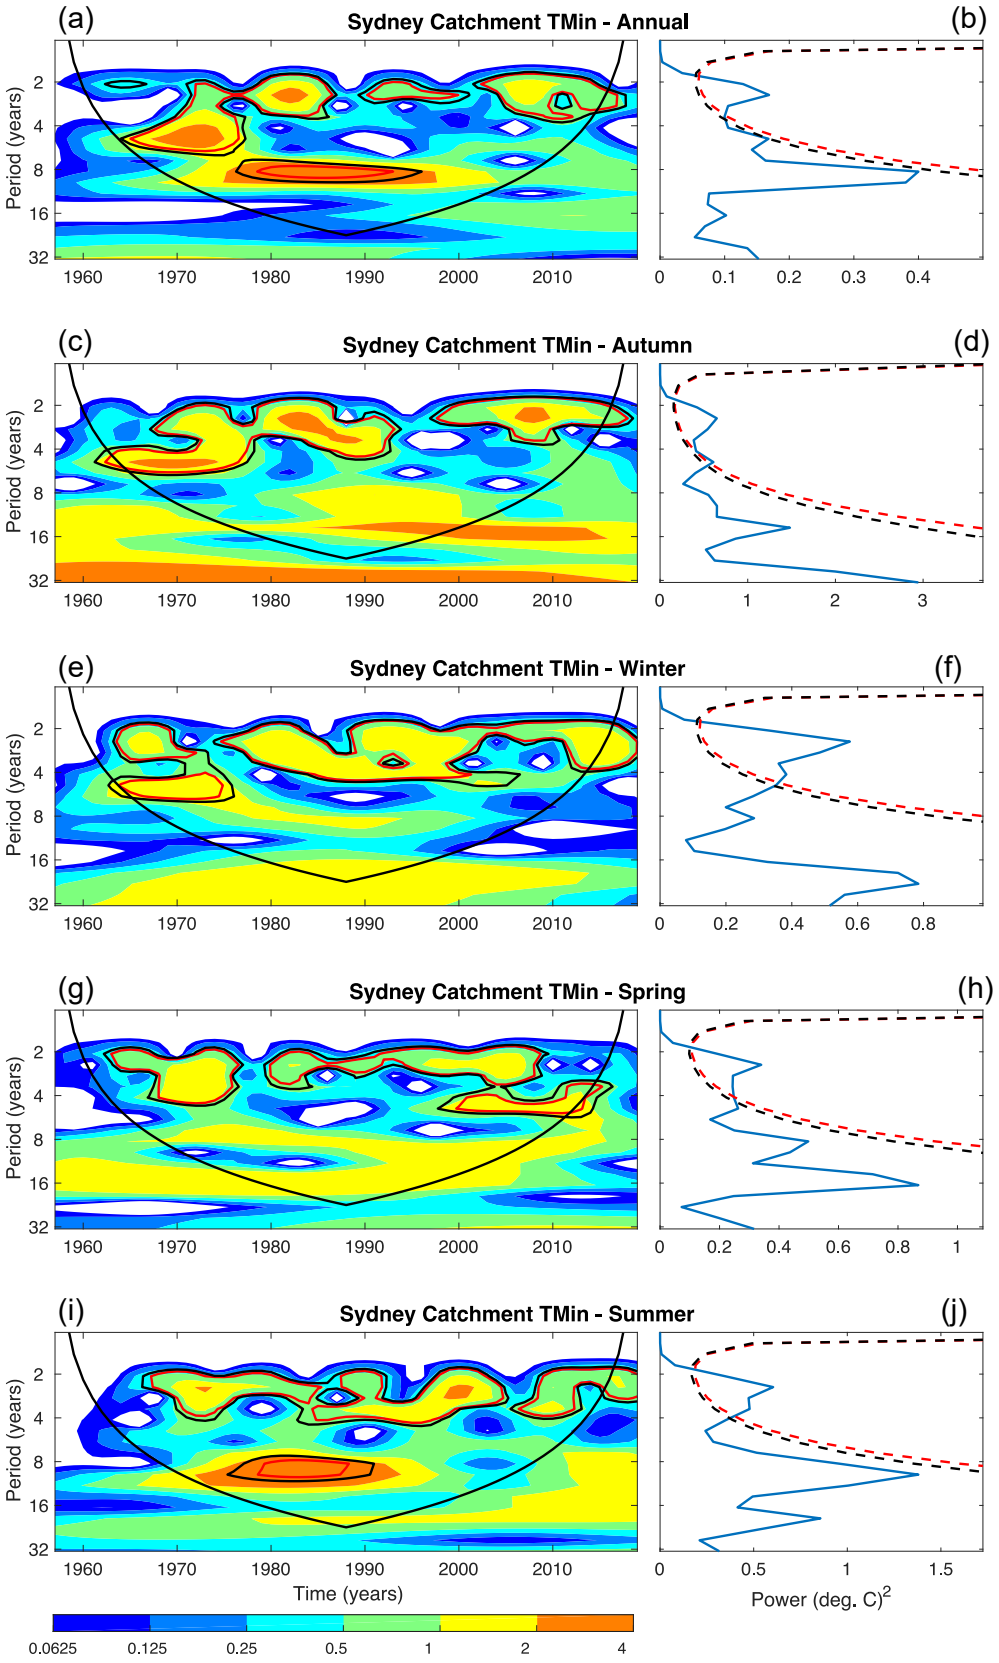

Figure 5.4.6: Wavelet power spectra (left panel, **a**–**i**) and global power spectra (right panel, **b**–**j**) for TMin, with low and high values in the power spectra indicating low and high periodicity, respectively; the cone of influence and statistical significance bands as already explained in Figure 5.4.4.

5.4.5 Attribution and Prediction of Annual Precipitation

For prediction of annual precipitation, the LR model that performed best is described by the equation:

$$\text{Precipitation} = 5887.4 - 941.8 \times \text{Niño3.4} - 5339.1 \times \text{DMI} \times \text{TSSST} + \varepsilon, \tag{5.4.1}$$

where ε represents the model error (some standard diagnostic plots, such as Q-Q and scale-location plots, were checked with no significant pattern displayed by the residuals to suggest the LR model could be significantly improved). The Poly SVR model used similar predictors to the LR, choosing Niño3.4, DMI*TSSST and Niño3.4*SAM, with free parameters C = 2, G = 0.25 and E = 1. Meanwhile, the RBF SVR model used a combination of GlobalT*Niño3.4 and SAM*TSSST as predictors, with C = $2^{1.5}$ and G = 0.25. The RF model used the predictors DMI*GlobalSSTA, DMI*GlobalT and DMI*SOI, with 150 trees. Shapiro–Wilks tests were applied to the residuals to test for normality, with p-values ranging from 0.21 to 0.25 for the four models. Therefore, it can be concluded that the residuals are normally distributed.

Table 5.3 displays the correlation, error statistics and skill scores for the final models on the testing data. The LR model achieved a high correlation of 0.815 (Table 5.3) but also had a relatively high RMSE of 764.971. Meanwhile, there was a slight decrease in correlation for the RBF SVR model to 0.761, but an improvement in RMSE by approximately 15% to 651.478. The best performing model is the Poly SVR, with an RMSE of 613.704 and high correlation of 0.921 (Table 5.3). Furthermore, the skill score of 0.721 of the Poly SVR suggests great improvement in results compared to the mean climatology. In contrast, the worst performing model was the RF model, which obtained an RMSE of 780.455 and correlation of 0.622 (Table 5.3). Although the RF model underperforms compared to others, the skill score of all models is positive, and at least 0.55 (for the RF model; Table 5.3), which shows that all models greatly improve on predictions compared to the mean climatology.

Table 5.3: Performance of each of the annual precipitation models on the testing data set (2011–2019). The best performing model for predicting annual precipitation is highlighted, determined by low RMSE, and both high correlation and skill.

Model	**LR**	**SVR (RBF)**	**SVR (Poly)**	**RF**
RMSE (mm^2)	764.971	651.478	**613.704**	780.455
Skill	0.567	0.686	**0.721**	0.550
Correlation	0.815	0.761	**0.921**	0.622
R^2	0.663	0.579	**0.848**	0.387
Residuals Mean (mm)	258.617	−22.463	**310.671**	98.182
Residuals SD (mm)	763.600	690.585	**561.365**	821.221
Residuals Skewness	1.365	0.084	**0.566**	0.283
Residuals Kurtosis	0.705	−1.314	**−1.158**	−1.729

Notably, both Poly SVR and LR models utilize the same predictors, except the Poly SVR has the added contribution of Niño3.4*SAM. Additionally, the kernel that was selected for the Poly SVR, after a grid search was performed, is essentially the linear kernel (i.e., E = 1). Looking at model predictions against observations, the climatology actually predicts annual precipitation fairly well until 2017 (not shown). This is unsurprising given annual precipitation does not appear to vary greatly over this part of the testing period, and remains approximately halfway between the 25th and 75th percentiles (Figure 5.4.1a). The greatest contribution to RMSE for the LR model is an underprediction in 2015 (not shown). The Poly SVR model improves on this, although it still underpredicts precipitation in 2015. This underprediction appears to be the result of a strongly positive Niño3.4 and SAM, which does not occur at

any other point in the testing period (not shown). Both models closely predict the observed pattern of annual precipitation, including the significantly reduced precipitation over 2017–2019. However, the Poly SVR underpredicts 2019 precipitation more than the LR model. Both RBF SVR and RF models predict the shape of annual precipitation fairly well, including the reduced precipitation during 2017–2019, but they tended to overpredict precipitation over this period, which is not beneficial in preparing for drought.

Total precipitation over the northern and southern regions of the catchment also was modelled using the same approach as modelling precipitation over the entire SCA. Although these models performed relatively well e.g., the Poly SVR model on southern catchment precipitation obtained a correlation of 0.904 and skill of 0.66 (not shown), they did not collectively perform better than models of precipitation over the entire SCA. The best performing predictor sets from both northern and southern catchment models, however, were used in separate models of entire catchment precipitation to test if they performed better than models selected through forward selection on annual precipitation over the entire SCA. Interestingly, all models presented here, except the LR model, were formed using predictors that were part of either the northern or southern catchment precipitation models as these performed better than those selected using the traditional forward selection technique (not shown).

5.4.6 Attribution and Prediction of Autumn Precipitation

The best performing LR model for autumn precipitation is:

$$\text{Precipitation} = 1516.96 + 5121.52 \times \text{AMO} \times \text{DMI} - 4477.40 \times \text{AMO} \times \text{TSSST} + \varepsilon. \quad (5.4.2)$$

Meanwhile, the RBF SVR used AMO*SOI, TPI*TSSST and Niño3.4*SAM as predictors, with $C = 2^{1.5}$ and $G = 0.25$. The Poly SVR used predictors AMO*GlobalT, AMO*TSSST, GlobalSSTA*TPI and SAM, with $C = 2^{0.5}$, $G = 0.25$ and $E = 1$. The RF model used AMO*SOI, SOI*TPI and TPI as predictors, with 50 trees.

The LR model performs best for autumn precipitation, with an RMSE of 418.036 and a correlation of 0.679 (Table 5.4). The Poly SVR has an increase in RMSE of approximately 23% from the LR model, and achieves a much lower correlation of 0.493. There are further increases in RMSE for the RBF SVR and RF models. This results in both models having negative skill compared to the mean climatology of autumn precipitation (skill = -0.133 for RBF SVR, -0.226 for RF; Table 5.4). When comparing predictions against observations, all models largely tend to predict the shape of precipitation fairly well (not shown); however, the RF model often overpredicts precipitation while the other three models usually underpredict precipitation (Table 5.4). Of the observations in the testing set, only 2016 and 2018 recorded extremes in autumn precipitation. In these cases, both years were below the 90th percentile (Figure 5.4.1a). Of the models presented here, only the LR model predicts close to the 2016 observation, albeit underpredicting (not shown). Meanwhile both SVR models, particularly the Poly SVR model, predict close to the 2018 record.

Table 5.4: Performance of each of the autumn precipitation models on the testing data set (2011–2019), similar to Table 5.3.

Model	LR	SVR (RBF)	SVR (Poly)	RF
RMSE (mm^2)	**418.036**	597.255	515.242	621.334
Skill	**0.445**	−0.133	0.157	−0.226
Correlation	**0.679**	0.116	0.493	0.094
R^2	**0.461**	0.013	0.243	0.009
Residuals Mean (mm)	**118.211**	138.001	162.255	−112.422
Residuals SD (mm)	**425.297**	616.342	518.692	648.147
Residuals Skewness	**−0.545**	−0.360	−0.044	0.469
Residuals Kurtosis	**−0.856**	−0.864	−1.148	−1.414

5.4.7 Attribution and Prediction of Winter Precipitation

The best LR model for winter precipitation found through forward selection is:

$$\text{Precipitation} = 1310.47 + 1008.38 \times \text{SOI} + 9843.28 \times \text{AMO} \times \text{SAM} - 8858.67 \times \text{DMI} \times \text{TPI} + 4059.01 \times \text{SAM} \times \text{TPI} + \varepsilon. \quad (5.4.3)$$

The RBF SVR uses only AMO*GlobalSSTA and TPI as predictors, with C = $2^{0.5}$ and G = 0.25. The Poly SVR uses the predictors GlobalT*SOI, TPI*TSSST and TSSST, with C = $2^{0.5}$, G = 0.5 and E = 1. The RF model uses AMO*SAM, DMI*TPI and Niño3.4*TPI, with 650 trees.

Here, the best performing model is the Poly SVR, which achieves an RMSE of 439.781 and a correlation of 0.703 (Table 5.5). The RBF SVR sees a decrease in correlation to 0.521, and increase in RMSE to 502.886. The LR and RF models perform fairly similarly, with LR having higher RMSE (533.319 compared to 522.228 for the RF model; Table 5.5) but also higher correlation of the predictions against observations (0.514 compared to 0.461 for the RF model). The LR model does not appear to capture the shape of precipitation over the testing period (2011–2019) very well (not shown). In contrast, the other models tend to capture the shape fairly well, especially the Poly SVR. These models were more inclined to underpredict precipitation whereas the LR model often overpredicted. The underprediction in the Poly SVR appears to mostly be related to the years 2015 and 2016. Precipitation during winter in 2016 was recorded above the 95th percentile, one of four years during the entire time series that recorded above the 95th percentile. The other three years were in the training data set; however, the model only appears to be trained relatively well to winter precipitation in 1998 (although still underpredicting, not shown).

Table 5.5: Performance of each of the winter precipitation models on the testing data set (2011–2019), similar to Table 5.3.

Model	LR	SVR (RBF)	SVR (Poly)	RF
RMSE (mm^2)	533.319	502.886	**439.781**	522.228
Skill	0.179	0.270	**0.442**	0.213
Correlation	0.514	0.521	**0.703**	0.461
R^2	0.264	0.272	**0.494**	0.212
Residuals Mean (mm)	−147.827	58.190	**145.452**	67.985
Residuals SD (mm)	543.506	529.808	**440.207**	549.193
Residuals Skewness	0.585	−0.217	**−0.205**	1.009
Residuals Kurtosis	−0.745	−1.356	**−1.202**	−0.444

5.5 Discussion

5.5.1 Drought Vulnerability

Global annual precipitation is expected to increase by approximately 7% per 1 °C of warming due to changes in the atmosphere's water-holding capacity (Trenberth, 2011). A slight increasing trend in global precipitation has been observed; however, this is related largely to ENSO variability (Trenberth et al., 2014; Adler et al., 2017). Despite the observed and expected increase in global precipitation, annual precipitation is decreasing in the SCA, due to a statistically significant decline in autumn and spring precipitation (p-values = 0.0778 and 0.0228, respectively; Table 5.2). There also is reduced variance in annual precipitation, due to decreasing variance in spring and summer. The decline in mean precipitation has been observed over southeast Australia and eastern Australia in previous studies, and is suggested to continue as a consequence of global warming causing the continued expansion of the Hadley cell, and reducing the frequency of rain-bearing systems reaching southern Australia (Speer et al., 2011; Timbal and Fawcett, 2013; Timbal and Drosdowsky, 2013; Post et al., 2014). Autumn precipitation is important for water catchment inflows over southeast Australia, because it saturates the soil for cool season runoff and evapotranspiration is lower during this period (Cai and Cowan, 2008; Cai et al., 2009). The decline in precipitation is further compounded by reductions in spring precipitation, which would result in notable reductions in runoff for the SCA.

Concomitant with the global warming trend (Stocker et al., 2014), there is a statistically significant increase in mean TMax and TMin annually and over all seasons except autumn (e.g., annual p-values = 0.0036 for TMax and 0.0000 for TMin; Table 5.2). This general increase in mean TMax and TMin suggests mean temperatures across the catchment are also increasing. Increasing temperatures lead to increased evapotranspiration, which reduces water storages. Combined with the clear reduction in annual precipitation, it is likely that there will be future droughts that are potentially more intense than those that have been experienced in the past. Furthermore, future water resources for the Sydney region will be strained.

5.5.2 Attribution and Prediction

Although non-linear statistical models have proved valuable for the prediction of variables such as precipitation, and attribution of these variables with climate drivers (Deo and Şahin, 2015; Bagirov et al., 2017; Richman and Leslie, 2018, 2020; Hartigan et al., 2020), their use in Australian hydrological problems has been limited. This is especially the case for Australia's most populous city, Sydney, where little has been published examining the applicability of non-linear models to predict and attribute annual and seasonal precipitation over the SCA. The most recent drought that affected the Sydney region, and resulted in tight water restrictions, highlights the importance of pursuing research into improving predictability of precipitation on various time scales for improved water management. In this study, multiple LR, SVR and RF were used with 10-fold cross-validation for attribution and prediction of annual precipitation, and precipitation during autumn and winter in the SCA. Two separate SVR models were developed, one used the radial basis function kernel (RBF) while the other used the polynomial kernel (Poly). The two seasons were selected to create precipitation models as runoff in this region of Australia typically is greatest during the cooler months of the year (Kirono et al., 2010). The northern region (Lithgow, Blackheath and Darkes Forest) typically records greater precipitation than the southern region (Moss Vale, Goulburn and Braidwood), suggestive of potential differences in the drivers of precipitation in these regions. Therefore, the catchment precipitation data also was split in order to develop models of precipitation for these regions. These models did not collectively perform better than models of the precipitation over the entire SCA, so they were not presented here. However, the predictors of the best performing models for the northern

and southern catchments were also considered for all statistical models of precipitation over the SCA. This resulted in better models when compared to those selected through forward selection with 10-fold cross-validation on catchment-wide precipitation for three of the four annual precipitation models, and one of four winter precipitation models.

Overall, the Poly SVR performed best for prediction of annual and winter precipitation over the testing period (2011–2019), while the LR model performed best on autumn precipitation. Interestingly, when the Poly SVR performed best, the grid search for an optimal exponent yielded a kernel with degree one. This does not necessarily suggest the LR model will perform similarly, for example, although the LR model for annual precipitation performs fairly well (RMSE = 764.971 and correlation = 0.815; Table 5.3), it achieves a skill score that is approximately 21% lower than that of the Poly SVR. This is despite sharing all but one predictor in common with the Poly SVR. The results presented in this paper, where an SVR with a linear kernel or the LR model perform best on the data, reflect the moderately strong linear relationships between the predictors and annual precipitation shown in Figure 5.3.1. Although nonlinear statistical models perform better than LR in some circumstances (e.g., Hartigan et al., 2020), and even perform somewhat better here, it is important to survey the performance of various statistical models as the results are dependent on location-specific features.

For annual precipitation in the SCA, there are many common attributes. In fact, both the LR model and Poly SVR model chose DMI*TSSST and Niño3.4 as predictors, while the Poly SVR also included the interaction term Niño3.4*SAM. Niño3.4 and SOI are both indices related to ENSO. This attribute is included at least once in all models, highlighting the known strong influence of ENSO on precipitation in eastern Australia (McBride and Nicholls, 1983), which also was identified by the wavelet analysis (Figure 5.4.4). Additional influences on annual precipitation appear to be the DMI and a measure of global warming, typically the TSSST.

Attributes in common among the models predicting autumn precipitation include AMO, TPI and TSSST. Meanwhile, the winter precipitation models have AMO, TPI and predictors related to ENSO. It is notable that measures of global warming, including SST anomalies or global temperature anomalies, are included in the Poly SVR for winter precipitation, despite no clear trend in either the mean or variance of winter precipitation (Figure 5.4.1). The influence of ENSO on autumn and winter precipitation was identified by the wavelet analysis (Figure 5.4.4c–f), although the exclusion of predictors related to ENSO in the autumn precipitation models is likely due to the reduced influence of ENSO from the 1990s. Additionally, the wavelet analysis displays high periodicity over the 8–16 year period, with statistical significance for winter precipitation. This is likely why TPI was a common attribute among the autumn and winter precipitation models, and potentially why the AMO was selected, which has been shown to exert a decadal influence over ENSO (Kucharski et al., 2011, 2016; Johnson et al., 2018). The commonly selected attributes for precipitation during autumn and winter differ somewhat from those selected for annual precipitation, with the inclusion of AMO and TPI, and exclusion of DMI. This highlights how seasonal precipitation can be influenced by different climate drivers of annual precipitation.

Autumn and winter precipitation individually contribute approximately 25% and 21%, respectively, to the annual mean precipitation. Despite the selection of different climate drivers that might be more applicable for the individual seasons, the models for individual seasons do not perform better overall than the best performing annual precipitation model. For example, the RMSE for the best performing models in either autumn or winter is more than 50% of the RMSE for the best performing model of annual precipitation. This is likely because the variance between seasons is fairly large compared to annual precipitation, especially for winter, with the timing and location of east coast low events influential on the resulting precipitation in a season. Rather than being highly accurate models of seasonal precipitation, the models for individual seasons appear more suitable as indicators of whether the season is likely to be

dry or wet, relative to the climatology (not shown).

ENSO influences the location of convection in the Pacific Ocean, with El-Niño events responsible for decreased convection over tropical Australia and a reduction of precipitation over eastern Australia (Ummenhofer et al., 2011; Bureau of Meteorology, 2012). The DMI also influences precipitation in eastern Australia, with large-scale moisture advection from the Indian Ocean feeding into precipitation bearing systems, such as northwest cloud bands (Risbey et al., 2009; Ummenhofer et al., 2009). During a negative DMI event, there is greater precipitation over eastern Australia, whereas a positive DMI is associated with cooler waters off the northwest of Australia and reduced precipitation in eastern Australia from reduced moisture advection. SST anomalies are considered a measure of global warming that is more reliable than the global temperature anomalies that include land surfaces due to reduced variability compared to the near-surface atmosphere (Choudhury et al., 2019). As there is a decreasing trend in annual precipitation, it is not surprising that some form of SST anomalies are included as a predictor in each model, and SST anomalies might be favoured over global temperature anomalies because the ocean responds slower than land to global temperature change (Hansen et al., 2010). The selection of TSSST by the LR and Poly SVR models is likely because it is adjacent to the SCA, therefore it influences local moisture advection and precipitation on a smaller scale. The inclusion of the interaction between DMI and TSSST highlights a compound effect of the Indian Ocean and Tasman Sea on precipitation systems. The AMO and TPI were selected by numerous models on autumn and winter precipitation, generally as a two-way interaction with another variable. Previous studies have found Atlantic Ocean variability can influence tropical Pacific climate variability, by modulating the global Walker circulation (Kucharski et al., 2011, 2016). This has been found to influence Australian precipitation variability by altering sea-level pressure anomalies in the region (Johnson et al., 2018). The TPI (or interdecadal Pacific oscillation) influences rainfall over eastern Australia through varying the impact of ENSO on precipitation (Power et al., 1999; Kiem and Franks, 2004; Verdon et al., 2004), likely through changes in the position of the South Pacific Convergence Zone (Folland et al., 2002). The relationships between precipitation and the various climate drivers selected by the models are complex and future work is required to explore these relationships in detail, particularly for the effect of different climate drivers on autumn and winter precipitation in the region studied.

In this study, there was a general improvement in skill compared with climatology in modelling precipitation using either non-linear or linear statistical methods, particularly for annual precipitation where the skill score for the best performing model is 0.721 (Table 5.3). The importance of considering a range of statistical models, including linear models, is clear in this study as linear models sometimes outperform the non-linear techniques, or a linear kernel outperforms the non-linear RBF kernel in SVR. Neither the RBF SVR or RF models performed best in modelling precipitation, which raises the question of whether it is worthwhile considering these models in the future. Although they did not perform the best, they still performed comparatively well except in modelling autumn precipitation. Additionally, they selected attributes of precipitation that were similar to those selected by the better performing LR or Poly SVR models, suggesting that underlying physical relationships are being captured by the models rather than random relationships between the variables. Through forward selection, all models included at least one two-way interaction term, indicating the need to consider not only singular indices but also potential relationships between indices including multiplicative and possibly inverse relationships.

It was expected that splitting the precipitation time series into seasons and creating models for autumn and winter precipitation would improve predictability. However, this did not occur with the RMSE of these models being more than half that of the worst performing annual precipitation model. Future work should investigate how to improve this, as precipitation over the cooler months of the year are highly important for catchment inflows. In this study, only

forward selection was applied to determine the predictors used in the models. However, this does not necessarily yield an optimal model from the subset of predictors (Hastie et al., 2009). This was clear after some optimal models for entire catchment precipitation were selected from the predictors used to model precipitation in either the northern or southern regions of the catchment. Therefore, one approach could be to employ a wider range of selection techniques, such as correlation-based feature selection (Hall, 1999). Another approach could be to investigate the use of lagged time series of climate predictors, which proved beneficial in previous drought modelling studies (e.g., Prasad et al., 2017). Finally, as found in Section 5.4.1, precipitation in the SCA is non-stationary. It is expected that relationships with climate variables will change as global warming continues (Benoit et al., 2020). Despite the high performance of models in this study, the predictors that were useful here might not be in the future. Therefore, there is a need to reassess the important attributes to SCA precipitation and develop updated models that are relevant to future climate regimes.

5.6 Conclusions

This study found that annual precipitation over the Sydney catchment area (SCA) is decreasing in the 31-year period 1989–2019, due to a decline in precipitation during both autumn and spring from the 1990s. This decrease in precipitation is expected to continue with the global warming trend. In addition, mean temperature over the SCA displayed an increasing trend over all seasons except autumn. This suggests the SCA is going to be increasingly vulnerable to drought conditions in the future. Further strain will be placed on water supply as the population continues to increase (Australian Bureau of Statistics, 2019). Therefore, it is necessary to continue monitoring trends in variables associated with water supply, and improve water supply through infrastructure investment and better seasonal forecasting.

In this study, multiple linear regression (LR), support vector regression (SVR) and random forests (RF) were used to model annual, autumn and winter precipitation. The models highlighted the importance of DMI, ENSO and Tasman Sea SST anomalies on annual precipitation. Different attributes were found to be associated with autumn precipitation (AMO, TPI and TSSST) and winter precipitation (AMO, TPI and ENSO). However, the models of autumn and winter precipitation did not perform as well as annual precipitation, with model errors individually more than 50% greater than those for annual precipitation. The models of autumn and winter precipitation appeared to be better used as indicators whether that season would likely be wetter or drier than average.

Overall, the models performed very well, with clear links between climate indices and annual, autumn and winter precipitation, despite the observed trends in annual and autumn precipitation. By performing studies using linear and non-linear statistical modelling, among other methods, the complex relationships between climate indices and atmospheric variables can be better understood. This can improve the application of climate driver indices in seasonal forecasts of atmospheric variables such as precipitation and temperature, which are performed by national weather agencies including the Australian Bureau of Meteorology, UK Met Office and National Weather Service. The best performing models in this study, depending on season, were either LR or the Poly SVR with an exponent of one. This highlights how linear models compare with non-linear models, and the importance of testing a wide-range of modelling techniques when developing a statistical model, rather than assuming that one model will perform better than another.

Bibliography

Adler, R., G. Gu, M. Sapiano, J.-J. Wang, and G. Huffman. Global Precipitation: Means, Variations and Trends During the Satellite Era (1979–2014). Surv. Geophys., 38:679–699, 2017.

Australian Bureau of Statistics. Australian Demographic Statistics. Technical report, 2019. Technical Reportl Cat. No. 3101.0, 44 pp.

Bagirov, A., A. Mahmood, and A. Barton. Prediction of monthly rainfall in Victoria, Australia: Clusterwise linear regression approach. Atmos. Res., 188:20–29, 2017.

Beggs, P. and Y. Zhang. The 2020 special reports of the MJA-Lancet Countdown on health and climate change: lessons learnt from Australia's '"Black Summer"' (summary. Med. J. Aust., 213:1–4, 2020.

Benoit, L., M. Vrac, and G. Mariethoz. Nonstationary stochastic rain type generation: accounting for climate drivers. Hydrol. Earth Syst. Sci., 24:2841–2854, 2020.

Breiman, L. Random forests. Mach. Learn., 45:5–32, 2001.

Bureau of Meteorology. The "federation drought", 1895–1902. `https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought1.html`, 2009a.

Bureau of Meteorology. The World War II droughts 1937–1945. `https://webarchive.nla.gov.au/awa/20090330051442/http://pandora.nla.gov.au/pan/96122/20090317-1643/www.bom.gov.au/lam/climate/levelthree/c20thc/drought3.html`, 2009b.

Bureau of Meteorology. Climate glossary. `http://www.bom.gov.au/climate/glossary/elnino.shtml`, 2012.

Cai, W. and T. Cowan. Dynamics of late autumn rainfall reduction over southeastern Australia. Geophys. Res. Lett., 35:L09708, 2008.

Cai, W., T. Cowan, P. Briggs, and M. Raupach. Rising temperature depletes soil moisture and exacerbates severe drought conditions across southeast Australia. Geophys. Res. Lett., (21): L21709, 2009.

Cheeseman, J. Food Security in the Face of Salinity, Drought, Climate Change, and Population Growth. In Halophytes for food security in dry lands, pages 111–123. Elsevier, 2016.

Chiew, F. H. S., N. J. Potter, J. Vaze, C. Petheram, L. Zhang, J. Teng, and D. A. Post. Observed hydrologic non-stationarity in far south-eastern Australia: implications for modelling and prediction. Stoch. Environ. Res. Risk Assess., 28:3–15, 2014.

Choudhury, D., R. Mehrotra, A. Sharma, A. Gupta, and B. Sivakumar. Effectiveness of CMIP5 Decadal Experiments for Interannual Rainfall Prediction Over Australia. Water Resour. Res., 55:7400–7418, 2019.

Commonwealth of Australia. Royal Commission into National Disaster Arrangements Report, 2020. URL `https://naturaldisaster.royalcommission.gov.au/system/files/2020-11/Royal%20Commission%20into%20National%20Natural%20Disaster%20Arrangements%20-%20Report%20%20%5Baccessible%5D.pdf`.

Davey, S. and A. Sarre. Editorial: the 2019/20 Black Summer bushfires. Aust. For., 83:47–51, 2020.

Deo, R. and M. Şahin. Application of the Artificial Neural Network model for prediction of monthly Standardized Precipitation and Evapotranspiration Index using hydrometeorological parameters and climate indices in eastern Australia. Atmos. Res., 161:65–81, 2015.

Diffenbaugh, N. S., D. L. Swain, and D. Touma. Anthropogenic warming has increased drought risk in California. Proc. Natl. Acad. Sci. U.S.A., 112:3931–3936, 2015.

Dowdy, A., A. Pepler, A. Di Luca, L. Cavicchia, G. Mills, J. Evans, S. Louis, K. McInnes, and K. Walsh. Review of Australian east coast low pressure systems and associated extremes. Clim. Dyn., 53:4887–4910, 2019.

Folland, C., J. Renwick, M. Salinger, and A. Mullan. Relative influences of the Interecadal Pacific Oscillation and ENSO on the South Pacific Convergence Zone. Geophys. Res. Lett., 29:21–1, 2002.

Hall, M. A. Correlation-based Feature Selection for Machine Learning. 1999.

Hansen, J., R. Ruedy, M. Sato, and K. Lo. Global Surface Temperature Change. Rev. Geophys., 48, 2010.

Hartigan, J., S. MacNamara, and L. Leslie. Application of Machine Learning to Attribution and Prediction of Seasonal Precipitation and Temperature Trends in Canberra, Australia. Climate, 8:76, 2020.

Hastie, T., R. Tibshirani, and J. Friedman. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media, 2009.

Hsu, C.-W., C.-C. Chang, and C.-J. Lin. A Practical Guide to Support Vector Classification. Technical report, National Taiwan University, 2003. 16pp.

Jalaludin, B. and G. Morgan. What does climate change have to do with bushfires? Aust. Health Rev., 45:4–6, 2021.

Johnson, Z., Y. Chikamoto, J.-J. Luo, and T. Mochizuki. Ocean Impacts on Australian Interannual to Decadal Precipitation Variability. Climate, 6(3):61, 2018.

Johnston, F., N. Borchers-Arriagada, G. Morgan, B. Jalaludin, A. Palmer, G. Williamson, and D. Bowman. Unprecedented health costs of smoke-related $PM_{2.5}$ from the 2019–20 Australian megafires. Nat. Sustain., 4:42–47, 2021.

Kiem, A. and S. Franks. Multi-decadal variability of drought risk, eastern Australia. Hydrol. Process., 18:2039–2050, 2004.

Kiem, A. S., F. Johnson, S. Westra, A. van Dijk, J. P. Evans, A. O'Donnell, A. Rouillard, C. Barr, J. Tyler, M. Thyer, D. Jakob, F. Woldemeskel, B. Sivakumar, and R. Mehrotra. Natural hazards in Australia: droughts. Clim. Change, 139:37–54, 2016.

Kirono, D., F. Chiew, and D. Kent. Identification of best predictors for forecasting seasonal rainfall and runoff in Australia. Hydrol. Process., 24:1237–1247, 2010.

Kucharski, F., I.-S. Kang, R. Farneti, and L. Feudale. Tropical Pacific response to 20th century Atlantic warming. Geophys. Res. Lett., 38:5, 2011.

Kucharski, F., F. Ikram, F. Molteni, R. Farneti, S. Kang, H.-H. No, M. King, G. Giuliani, and K. Mogensen. Atlantic forcing of Pacific decadal variability. Clim. Dyn., 46:2337–2351, 2016.

Lau, K.-M. and H. Weng. Climate Signal Detection Using Wavelet Transform: How to Make a Time Series Sing. Bull. Am. Meteorol. Soc., 76(12):2391–2402, 1995.

Linnenluecke, M. and M. Marrone. Air pollution, human health and climate change: newspaper coverage of Australian bushfires. Environ. Res. Lett., 16:1–11, 2021. 125003.

Maldonado, S. and R. Weber. A wrapper method for feature selection using Support Vector Machines. Inf. Sci., 179:2208–2217, 2009.

McBride, J. and N. Nicholls. Seasonal Relationships between Australian Rainfall and the Southern Oscillation. Mon. Weather Rev., 111:1998–2004, 1983.

Mukherjee, S., A. Mishra, and K. Trenberth. Climate Change and Drought: a Perspective on Drought Indices. Curr. Clim. Change Rep., 4:145–163, 2018.

Niang, I., O. C. Ruppel, M. A. Abdrabo, A. Essel, C. Lennard, J. Padgham, and P. Urquhart. 2014: Africa. In: Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, 2014. 1199–1265 pp.

Post, D., B. Timbal, F. Chiew, H. Hendon, H. Nguyen, and R. Moran. Decrease in southeastern Australian water availability linked to ongoing Hadley cell expansion. Earth's Future, 2:231–238, 2014.

Power, S., T. Casey, C. Folland, A. Colman, and V. Mehta. Inter-decadal modulation of the impact of ENSO on Australia. Clim. Dyn., 15:319–324, 1999.

Prasad, R., R. Deo, Y. Li, and T. Maraseni. Input selection and performance optimization of ANN-based streamflow forecasts in the drought-prone Murray Darling Basin region using IIS and MODWT algorithm. Atmos. Res., 197:42–63, 2017.

Richman, M. and L. Leslie. The 2015–2017 Cape Town Drought: Attribution and Prediction Using Machine Learning. Procedia Comput. Sci., 140:248–257, 2018.

Richman, M. B. and L. M. Leslie. Uniqueness and Causes of the California Drought. Procedia Comput. Sci., 61:428–435, 2015.

Richman, M. B. and L. M. Leslie. Machine Learning for Attribution of Heat and Drought in Southwestern Australia. Procedia Comput. Sci., pages 3–10, 2020.

Risbey, J., M. Pook, P. McIntosh, M. Wheeler, and H. Hendon. On the Remote Drivers of Rainfall Variability in Australia. Mon. Weather Rev., 137:3233–3253, 2009.

Simmons, J., C. Paton-Walsh, A. Mouat, J. Kaiser, R. Humphries, M. Keywood, A. Sutresna, D. Griffith, J. Naylor, and J. Ramirez-Gamboa. The gas and aerosol phase composition of smoke plumes from the 2019–2020 Black Summer bushfires and potential implications for human health. pages 1–29, 2021. Preprint.

Speer, M., L. Leslie, and A. Fierro. Australian east coast rainfall decline related to large scale climate drivers. Clim. Dyn., 36:1419–1429, 2011.

Stanke, C., M. Kerac, C. Prudhomme, J. Medlock, and V. Murray. Health Effects of Drought: a Systematic Review of the Evidence. PLOS Curr., 5, 2013.

Stocker, T., D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley. Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, 2014.

Sydney Water. Sydney's drought: Aquabumps shows just how dry it got. https://www.sydneywater.com.au/SW/about-us/our-publications/Media/sydney-s-drought--aquabumps-shows-just-how-dry-it-got/index.htm, 2020.

Timbal, B. and W. Drosdowsky. The relationship between the decline of Southeastern Australian rainfall and the strengthening of the subtropical ridge. Int. J. Climatol., 33:1021–1034, 2013.

Timbal, B. and R. Fawcett. A Historical Perspective on Southeastern Australian Rainfall since 1865 Using the Instrumental Record. J. Clim., 26:1112–1129, 2013.

Torrence, C. and G. Compo. A Practical Guide to Wavelet Analysis. Bull. Am. Meteorol. Soc., 79(1):61–78, 1998.

Trenberth, K. Changes in precipitation with climate change. Clim. Res., 47:123–138, 2011.

Trenberth, K., A. Dai, G. van der Schrier, P. Jones, J. Barichivich, K. Briffa, and J. Sheffield. Global warming and changes in drought. Nat. Clim. Change, 4:17–22, 2014.

Ummenhofer, C., A. Gupta, P. Briggs, M. England, P. McIntosh, G. Meyers, M. Pook, M. Raupach, and J. Risbey. Indian and Pacific Ocean Influences on Southeast Australian Drought and Soil Moisture. J. Clim., 24:1313–1336, 2011.

Ummenhofer, C. C., M. H. England, P. C. McIntosh, G. A. Meyers, M. J. Pook, J. S. Risbey, A. S. Gupta, and A. S. Taschetto. What causes southeast Australia's worst droughts? Geophys. Res. Lett., 36, 2009.

Vapnik, V. The Nature of Statistical Learning Theory. Springer, 1995.

Verdon, D., A. Wyatt, A. Kiem, and S. Franks. Multidecadal variability of rainfall and streamflow: Eastern Australia. Water Resour. Res., 40, 2004. W10201.

Wilks, D. S. Statistical Methods in the Atmospheric Sciences, volume 100. Academic Press, 2011.

Chapter 6

Impact of accelerated global warming on rainfall and temperature trends in Australia's northern Murray-Darling Basin using statistical analysis and machine learning

6.1 Overview

Australia's Murray-Darling Basin (MDB) region covers a large area of southeast Australia, consisting of a network of river systems including Australia's longest river, the Darling River. The MDB supports a large agricultural sector that produces over 40% of the national food supply. However, water availability in the MDB has faced significant pressure as a result of long-lasting droughts including the Millennium drought, and mismanagement of resources partly due to its spatial extent covering 4 Australian states and one territory, resulting in policies that vary across districts. Water flow within the rivers that constitute the MDB is vital for the environment, with the lack of flow causing algal blooms and mass fish deaths in the system. Unfortunately, it is expected that water availability within the MDB will continue to decrease as a result of global warming, with the ecosystem surrounding the MDB likely to collapse without urgent intervention (Bergstrom et al., 2021).

Following from the work in previous Chapters, along with an article focused on the southern MDB that I co-authored (Speer et al., 2021), this Chapter aims to better understand how drought is changing in the northern MDB, and the large-scale climate influences on rainfall in the region, including the influences of global warming. As second author to this study, I obtained the data for analysis, produced the time series, box plots and wavelets, and developed the machine learning and linear regression models. I also helped with analysis and preparation of the manuscript, alongside Dr Milton Speer, who conceived this study with Professor Lance Leslie, and led this study. This manuscript has been submitted for review in the international, peer-reviewed journal Big Earth Data.

6.2 Introduction

6.2.1 Australia's Murray-Darling Basin

Droughts and increasingly long dry spells are well-known characteristics, over the last two decades, of Australia's most important agricultural growing area, the Murray-Darling Basin (MDB; Holland et al., 2015), located in southeast Australia (Figure 6.2.1 inset). The northern MDB (Figure 6.2.1), hereafter NMDB, has a variable and intermittent rainfall pattern, with long dry periods and droughts interspersed with intense rainfall events. The NMDB occupies a large geographical area situated in the northern subtropical latitudes (between 25°S and 34°S) of Australia. It is more influenced by tropical climate drivers such as the El-Niño Southern Oscillation phenomenon (ENSO; McBride and Nicholls, 1983) than the southern MDB (between 34° and 38°S), which is more influenced by mid-latitude weather systems in the cool-season months of April-September (Risbey et al., 2009; Pook et al., 2014), so the two areas have different climate regimes.

Historically, as will be discussed in Section 6.2.2, NMDB floods have occurred as a result of winter/spring rains. However, floods also can occur in summer, particularly in La-Niña years, for example, in 2010–2012 (Bureau of Meteorology, 2012) and recently in 2020–2022. The NMDB summer flood rainfall typically results from tropical low pressure systems, including landfalling tropical cyclones (`https://www.mdba.gov.au/issues-facing-basin/flooding-murray-darling-basin`). Globally, each decade since 1980 has been warmer than the previous decade with 2010–2019 being around 0.2°C warmer than 2000-2009 (Bureau of Meteorology and CSIRO, 2020), indicating an acceleration in global warming in recent decades. In addition, mean annual near-surface air temperature since 1910 in Australia has increased by nearly 1.5°C, relative to the 1910–1950 mean, which is approximately 30% greater than the global average (CSIRO, 2020).

Since the early 1990s climate change research has revealed a reduction of about 12% in April-October precipitation over southeast Australia (CSIRO, 2020). Notably, there is an increase in summer precipitation(CSIRO, 2020) and a decrease in autumn (March-May) precipitation (Cai and Cowan, 2013). For the southern MDB, Speer et al. (2021) found a reduction in precipitation in late autumn (April-May) since the 1990s.

Figure 6.2.1: **Map of northern and southern Murray-Darling Basin in southeast Australia** The MDB lies within subtropical latitudes (25°S – 38°S) of the Australian continent. Observation stations used for precipitation, TMax and TMin averaging that represent the NMDB are marked and indicated in a legend. (Source: Murray-Darling Basin Authority, G.P.O. Box 1801, Canberra City, ACT 2601 Australia. `https://www.mdba.gov.au/sites/default/files/pubs/Murray-Darling_Basin_Boundary.pdf`) Reproduced with some place name deletions and insertions via license: Creative Commons Attribution-Non Commercial-NoDerivatives4.0 International Public License (CC BY-NC-ND 4.0)

6.2.2 The NMDB

Rainfall in the NMDB is generated by weather systems on various time and space scales. They range from short duration and small space scale thunderstorms to larger scale systems, likely resulting from climate drivers such as the ENSO, the southern annular mode (SAM), the Dipole Mode Index of the Indian Ocean Dipole (IOD), and the Interdecadal Pacific Oscillation (IPO), or its Tripole Index (TPI), on weekly to seasonal and decadal time scales (`http://www.bom.gov.au/climate/about/australian-climate-influences.shtml`). These modes drive the high variability and unreliability in seasonal and annual rainfall amounts and hence in stream inflows that eventually reach the Darling River. However, contributions to the precipitation reduction in this area of southeast Australia in late autumn (April-May) from the climate drivers described above using techniques such as machine learning have not yet been assessed.

The Darling River connects the entire river network of the NMDB with the southern MDB (Figure 6.2.1). Notably, most years do not deliver the annual mean water contribution from the Darling River. The annual mean contribution masks the extreme fluctuations that occur in the Darling River flows which, in turn, are driven by the highly variable NMDB annual catchment area rainfall. Nevertheless, farming, and other local communities rely on river water for

human consumption and for irrigation of agricultural crops, underlining the social and economic importance of river water for the NMDB. There is a wide variety of crops requiring irrigation, the most important being cotton, of which the NMDB is Australia's largest producer (http://www.murrayriver.com.au/about-the-murray/murray-darling-basin/#agriculture).

In addition, key NMDB environmental regions require an adequate water supply to sustain important NMDB wetlands, and to maintain the many river water holes in which fish species have adapted to survive long interruptions of water flow (Sheldon, 2017). Periods of no flows in the Barwon-Darling River system have been a regular feature, occurring 45 times between 1885 and 1960. However, there has been an increased frequency and duration of no flow periods in the Barwon-Darling since 2000 and a decreased frequency and duration of low flow pulses, particularly downstream of Bourke (Sheldon, 2017). Mass fish deaths in the Lower Darling River in 2019 were caused by the decrease in low flows, due mainly to poor river regulation combined with recent drought and higher temperatures (https://australiainstitute.org.au/report/a-fish-kill-qanda/). Small flushing flows improve water quality along the river by reducing salinity levels, breaking up stratification in pools, mitigating algal blooms, and are essential for ensuring the diversity and abundance of aquatic life such as prawns, mussels and aquatic insects which form a vital food resource for fish and waterbirds (Sheldon, 2017). Nevertheless, while acknowledging that there are known management issues affecting river levels, 94% of the gauges in the northern MDB show a declining trend in streamflow since records began in 1970, and in the headwaters of the NMDB catchment there are statistically significant declining trends (CSIRO, 2020).

Rainfall extremes have continued after the Millennium Drought (1997–2009). Following a wetter than average 2016, most of the southern MDB received decile 2-3 or decile 1 rainfall over the 4-year period from January 2017 to December 2020. During the same period, much of the NMDB fared even worse, with almost all receiving either decile 1 rainfall or its driest 4 years on record (Fig 6.2.2). In 2021, even after significant rainfall finally replenished the NMDB total water storage in late summer and mid-winter, it was still only at 82% capacity by the end of the cool season in late September 2021 (https://www.mdba.gov.au/water-management/murray-darling-reports-data/water-in-storages). Moreover, given that river inflows have almost ceased, the levels of NMDB dams, which are mostly in the state of New South Wales (NSW), are expected to drop sharply because of their small storage capacities, as shown by similar, historical characteristics in the graphs of their storage volumes (https://realtimedata.waternsw.com.au/).

Because the NMDB climate is so variable, it also experiences flood events. While flood events do not occur regularly in the NMDB, there were two major flood events during the recent 2010–2012 and 2020–2022 periods. These two major floods occurred as a result of rainfall in the upper part of the catchment where the observation stations are located (Figure 6.2.1). Consequently, with the decrease in MDB flooding events (Holland et al., 2015), and the drying out of river waterholes in the last two decades, the fish deaths are compounded by an increased spread of blue-green algae from the reduced river flows. Therefore, it is vital to analyze potential trends in variables that impact water availability, such as decreased NMDB catchment precipitation, and increased evaporation rates due to higher temperatures resulting from global warming. For example, rainfall was very low across the NMDB during January–February 2019, typically the wetter months of the year, resulting in 2018–2019 being the driest since 2002–2003. This led to area-averaged actual evapotranspiration and soil moisture deciles being among the lowest on record as shown by Figure 6.2.3a and Figure 6.2.3b, respectively. Importantly, Stephens

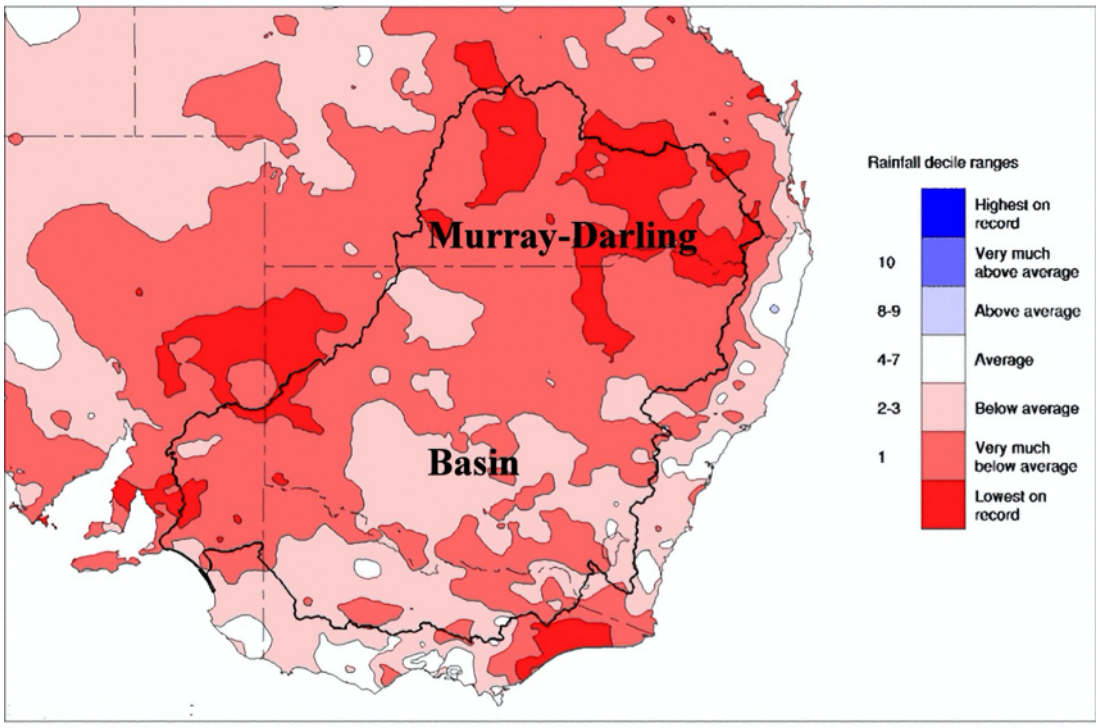

Figure 6.2.2: **Murray-Darling Basin rainfall deciles** Rainfall deciles for the 48 months January 2017 to December 2020 in southeast Australia focusing on the MDB defined by the area within the solid black line. Note the lowest on record in the north of the basin and the very much below or below average rainfall in the rest of the basin. (Reproduced with permission under Creative Commons Attribution Licence 3.0 from the Australian Bureau of Meteorology). Available at : `http://www.bom.gov.au/climate/maps/rainfall/?variable=rainfall&map=decile&period=48month&region=md&year=2020&month=12&day=31`

et al. (2018), in a study covering 1975–2016, found that prior to 1994 wind speed dominated evaporation rates across southern Australia, but after 1994 temperature became dominant. It is notable that evapotranspiration in the NMDB during 2018–2019 was even greater than in the southern MDB (Figs. 6.2.3a, b). Moreover, the water storage in the NMDB dams is about 1/3 of the capacity of the storage of southern MDB dams. Hence, when the NMDB dams fill in the warm season (October-March), as they did in 2020/2021 and 2021/2022, irrigation usage and high evaporation rates can rapidly reduce dam levels. It also indicates that an excessively high proportion of the NMDB inflow reaching the Darling River is the direct result of rainfall inflows rather than releases from dams. Because most tributaries in the NMDB are unregulated, historically this is why most farms have found it necessary to construct on-farm dams for water storage (`https://www.mdba.gov.au/sites/default/files/pubs/Final-Report-Independent-Panel-fish-deaths-lower%20Darling_4.pdf`). The future development of more accurate seasonal and annual prediction models on a range of time scales, would benefit from climate drivers identified in this study, which we investigate using machine learning techniques. Hence, valuable guidance could be provided for state and local water authorities, thereby assisting future water allocation decisions.

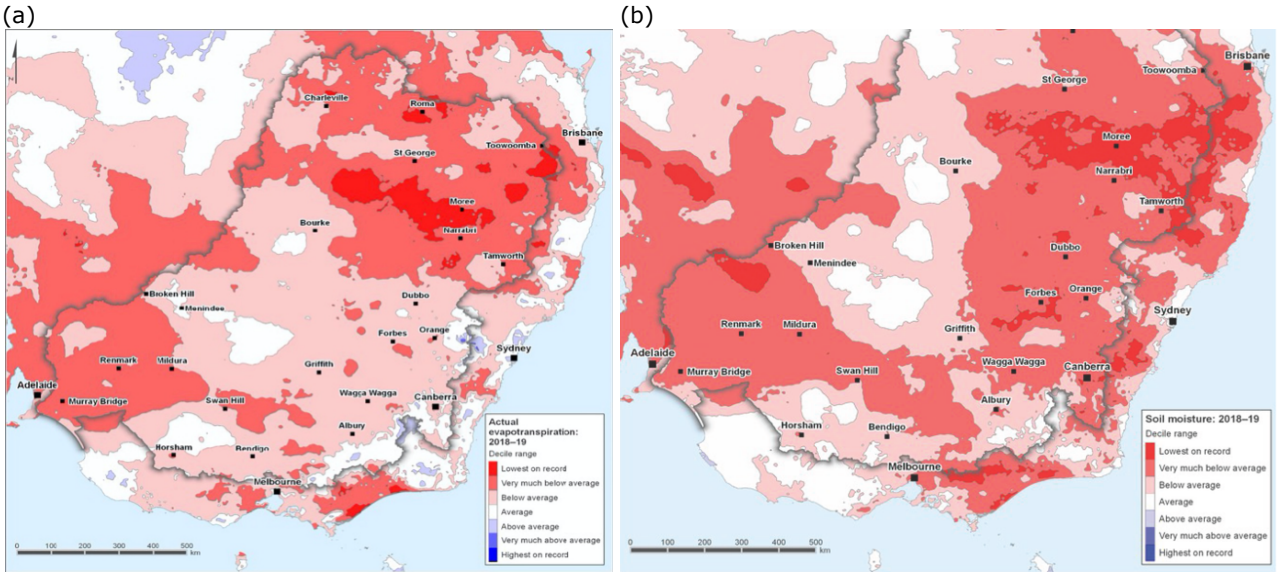

Figure 6.2.3: **Annual deciles of actual evapotranspiration and soil moisture 2018–2019** Map of southeast Australia showing for the MDB region deciles during the 2018–2019 year for, (a) annual area-averaged actual evapotranspiration. Note the decile area of lowest on record in the NMDB, and (b) annual area-averaged soil moisture. Note the decile area of the lowest on record in the NMDB. (Reproduced with permission under Creative Commons Attribution Licence 3.0 from the Australian Bureau of Meteorology. Available at: `http://www.bom.gov.au/water/nwa/2019/mdb/climateandwater/climateandwater.shtml`).

In summary, the major aims of this study are as follows. First, is to identify trends in precipitation and temperature from sections of the vast NMDB catchment area that determine water availability in the Darling River, and from the numerous river systems and tributaries that feed into it. Next is a detailed attribution analysis of climate drivers affecting annual and seasonal rainfall using machine learning techniques. Such an approach has not been attempted in the region previously. Here, it is performed by ranking the numerous possible attributes in importance, and by using a range of linear and non-linear machine learning regression techniques, described in Sections 6.3.3 and 6.4.4, fitted to the observed precipitation amounts, for the period 1965–2018. The attributes are assessed both individually and in combinations, with the five most effective attributes chosen on their percentage appearance in 10-fold cross-validation (to avoid overfitting) of annual precipitation, precipitation in late autumn (April-May), in the remaining cool season (JJAS), and in the warm season (October-March), to highlight warm season and cool season influences. The regression techniques chosen have proved to be successful in previous studies by the authors, in identifying attributes in previous climate change studies (e.g., Hartigan et al., 2020a,b). In addition, a wavelet analysis was applied to investigate the contributions from a range of climate drivers such as ENSO.

6.3 Data and Methodology

6.3.1 Data

Monthly precipitation data were obtained from the Bureau of Meteorology's (BoM's) homogeneous climate record through its climate change site network. (`http://www.bom.gov.au/climate/change/?ref=ftr#tabs=Tracker&tracker=site-networks`). These data have undergone complex quality control to address inconsistencies and errors. There are 8 precipitation stations comprising Augathella, Cunnamulla, Normandy, Miles, Surat, Bellata, Bingara and Curlewis. Additionally, monthly mean maximum temperature (TMax) and mean

minimum temperature (TMin) data were obtained for Charleville, Miles, Thargomindah, St. George. Moree, Inverell, Walgett and Gunnedah from the Bureau of Meteorology's climate change site network. Sites were selected to cover a broad area of the NMDB (Figure 6.2.1) and to contain a long, continuous record beginning in 1910 in a similar manner to an earlier MDB study (Murray-Darling Basin Authority, 2018). The precipitation data were summed to obtain a proxy to the total precipitation falling across the entire NMDB catchment area, while the mean of the temperature time series data was calculated to obtain the mean of TMax and TMin across the NMDB. The available time series data cover the period from 1910–2019.

6.3.2 Statistical analysis

To analyze trends in the data, the time series were plotted with their percentiles displayed in four panels: annually, late autumn (April-May), the four remaining cool season months (JJAS), and the warm season (October-March). The data were then grouped into four 27-year periods (1911–1937, 1938–1964, 1965–1991 and 1992–2018), covering almost all of the time series. Additionally, the decision on the groupings was influenced by the fact that the last three 27-year periods mentioned above coincide with the main time span from 1950 when global warming became prominent (Pachauri and Meyer, 2014), then from the 1970s when global warming began to accelerate (Morice et al., 2012; National Centers for Environmental Information, 2020), and particularly from the early to mid-1990s (Bureau of Meteorology and CSIRO, 2020). Change point analysis of the data is not applicable because the acceleration in global warming occurred over several years in the early to mid-1990s. If years other than 1991 had instead been chosen to demarcate the last two periods, the findings would have remained almost unchanged. Bootstrap resampling with replacement was applied, with 5000 resamples, to the mean and variance of the time series data to identify potential trends between the time periods. Two-sided permutation testing was applied to test for statistical significance of these trends. A significance level of $\alpha < 0.1$ was considered to represent a marginally statistically significant result, and $\alpha < 0.01$ represents a highly significant result.

6.3.3 Attribute selection

Numerous climate drivers were considered as potential attributes for annual, April-May, JJAS, and October-March precipitation. The drivers assessed were the Atlantic Multidecadal Oscillation (AMO), the Indian Ocean Dipole Mode Index (DMI), global sea surface temperature anomalies (GlobalSSTA), global temperature anomalies (GlobalT), Niño3.4, TPI, the Southern Annular Mode (SAM), the Southern Oscillation Index (SOI), and the Tasman Sea surface temperature anomalies (TSSST). Time series of these climate drivers were obtained from the Earth System Research Laboratory (`http://www.esrl.noaa.gov/psd/gcos_wgsp/Timeseries/`), except for TSSST which was obtained from the BoM (`http://www.bom.gov.au/climate/change/?ref=ftr#tabs=Tracker&tracker=timeseries`). Two-way interaction terms between the above predictors, where one variable is multiplied by another (e.g., AMO*IOD) were also considered potential predictors. These two-way interaction terms are considered as possible attributes as one attribute might reinforce another. Additional relationships between possible attributes, such as additive and quotient relationships, although assessed by Richman and Leslie (2020) for a single station, were not considered in this study, because some of these are highly correlated with other attributes considered here, and all ocean basins already are represented. As relationships between climate drivers and precipitation can vary between weak and strong, linear and non-linear, and combinations, this study develops both linear and non-linear statistical models of annual, April-May, JJAS and October-March precipitation, for a range of climate drivers. The non-linear models considered are support vector

regression (SVR) (Vapnik, 1995), with either the polynomial (Poly) or radial basis function (RBF) kernels, and random forests (RF) (Breiman, 2001; Hastie et al., 2009). When the two-way interaction combinations are included, there are 45 potential attributes considered for the prediction of precipitation. If all attributes were used, the models would be subject to overfitting, which reduces physical understanding and is likely to lead to large errors if the models are applied to test data set in prediction mode. To select attributes that generalize well from the set of 45 attributes, ten-fold cross-validation were applied to the data set, using both forward and backward selection through the space of potential attributes as in Maldonado and Weber (2009). The comprehensive, detailed set of results from all attributes in all models are detailed in Table 6.2.

6.4 Results and Discussion

First, precipitation and temperature time series are discussed followed by the P-values and box-whisker plots for precipitation, TMax and TMin. Then wavelet analysis of precipitation, TMax and TMin is discussed to highlight the periodicity in the climate drivers. Finally, attribute selection of climate drivers is described for each time period.

6.4.1 Precipitation and Temperature Time Series in the Northern Murray-Darling Basin

Precipitation

April-May is the only period of the year exhibiting long-term changes, with an apparent decrease in both the mean and variance from the 1990s (Figure 6.4.1b). The annual time series displays no long term trend (Figure 6.4.1a), which also is the case for JJAS and October-March. However, the steep decrease in the annual times series after 2011, to well below the 5th percentile in 2019, has contributions from all of the periods.

TMax

The annual times series of TMax shows an increasing trend from the early 1990s (Figure 6.4.1c). That is due to the strong contributions from JJAS in particular, but also from April-May. However, there is a steep increase approaching the 100th percentile in the period October-March (Figure 6.4.1d), which mirrors the steep decrease in 2019 precipitation.

TMin

The annual TMin time series shows an almost linear increase since the 1960s with a steep increase and much less variation since 2014 (Figure 6.4.1e). The increase is due mostly to similar increases in JJAS and October-March. The steep increase and lower variability are present also in JJAS and October-March. However, after an initial steep increase in April-May from the early 1950s (Figure 6.4.1f), there appears to be no clear trend after that, aside from a steep increase and lower variability from 2014.

6.4.2 P-values and box-whisker plots for precipitation, TMax and TMin

The p-values in Table 6.1 allow comparisons of the four 27-year intervals. Statistically significant differences (i.e., p-value < 0.1) are shown, in bold italics, for means and variances of

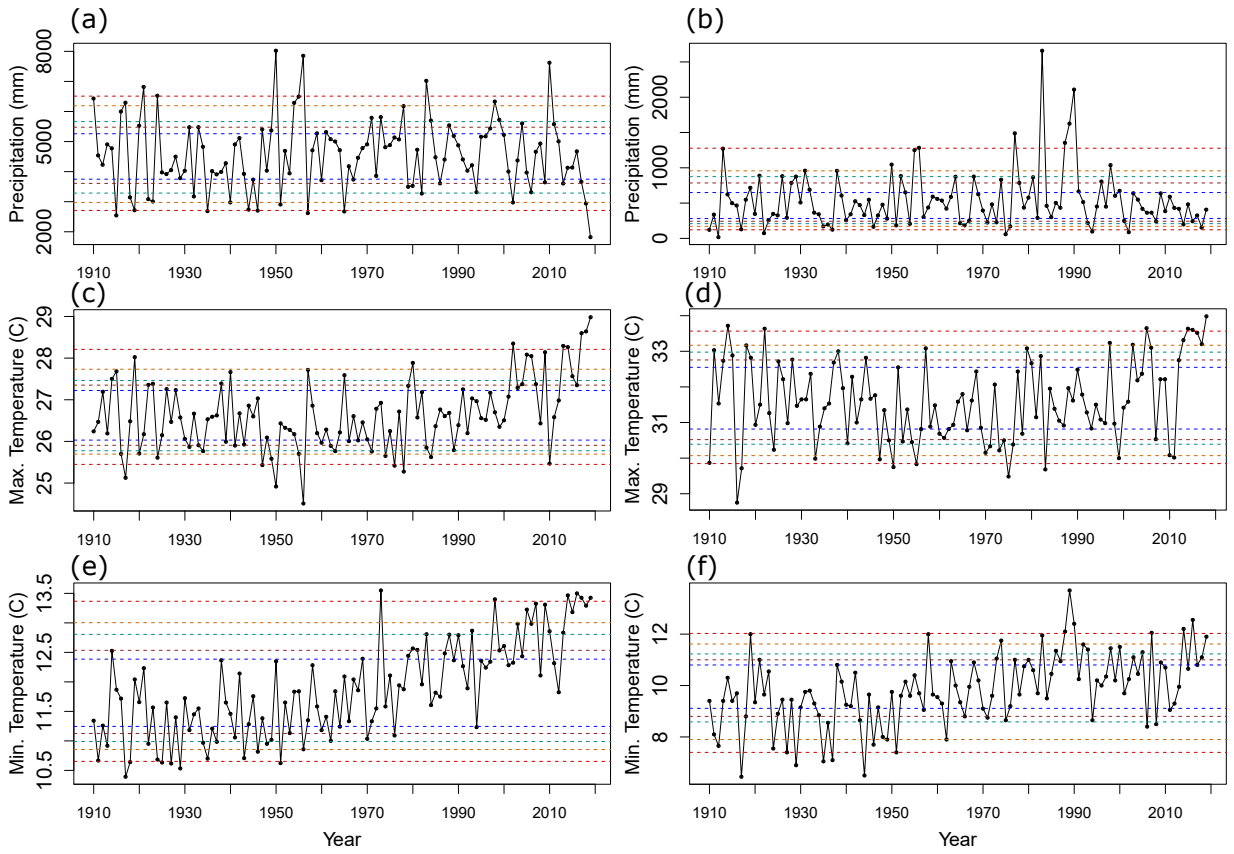

Figure 6.4.1: **Precipitation, TMax & TMin time series in the NMDB** Precipitation time series in the NMDB for (a) Annual, (b) April-May. Dashed lines indicate percentiles 5th and 95th (red); 10th and 90th (orange); 15th and 85th (light green); 20th and 80th (brown); and, 25th and 75th (dark blue). Note the apparent decrease in mean and reduction in values greater than the 75th percentile for April-May since the 1990s and the decrease from 2011 to well below the 5th percentile in 2019 for the annual time series; the time series of TMax in the NMDB for (c) Annual, (d) October-March. Note the steep increase approaching the 100th percentile in both October-March and annual time series, which mirrors the steep decrease in 2019 annual precipitation (Figure 6.4.1b); the time series of TMin in the NMDB for (e) Annual, (f) April-May. Note the almost linear increase in the annual time series since the 1960s and the steep increase with much less variation since 2014, while there is no clear trend in April-May apart from one steep increase after the 1950s and another from 2014.

precipitation, TMax and TMin. There are comparisons involving the four intervals 1911–1937, 1938–1964, 1965–1991 and 1992–2018.

Precipitation

The most important precipitation finding is the significant decrease in the April-May mean and highly significant decrease in variance, between the intervals 1965–1991 and 1992–2018, with a p-value of 0.0318 and 0.0002 (Table 6.1). It is noteworthy that the significant decrease in the mean is an abrupt reversal of the increasing trend in the mean in the periods prior to 1992–2018. The decrease in variance between 1992–2018 and previous intervals, including 1965–1991, also is highly significant, with a p-value of 0.0002, as shown in Table 6.1, and is very apparent in the box-whisker plots (Figures 6.2.2a, b). Possible climate influences are discussed in the following wavelets Section.

Table 6.1: **P-values from permutation testing differences in interval means and variances** P-values from permutation testing differences in interval means and variances for April–May, JJAS, October–March and annual precipitation, TMax and TMin, based on area averages of observing stations in the northeast part of the NMDB. Marginally significant values (p-value ≤ 0.10) are in bold italics. Note that the p-value for each variance test is calculated after one sample has had bias correction in the mean. Key points to note are the significant and highly significant p-values (p-value < 0.05 and p-value < 0.01, respectively) for the April–May mean and variance precipitation decreases from 1965–1991 to 1992–2018; and the highly significant increases in mean TMin, and in mean TMax for most of the periods.

Period	***Observation***	***Mean*** & ***Variance***	***1911–1937 vs 1965–1991***	***1938–1964 vs 1965–1991***	***1911-1937 vs 1992–2018***	***1965–1991 vs 1992–2018***
April-May	Precip.	Mean	***0.1***	0.289	0.48	***0.0318***
		Variance	***0.0168***	***0.0154***	***0.0946***	***0.0002***
	TMax	Mean	0.194	***0.0542***	***0.0002***	***0.009***
		Variance	0.934	0.599	0.677	0.73
	TMin	Mean	***0.0000***	***0.0016***	***0.0000***	0.847
		Variance	0.726	0.897	0.318	0.584
JJAS	Precip.	Mean	0.593	0.788	0.388	0.679
		Variance	0.23	0.815	0.581	0.585
	TMax	Mean	0.689	0.577	***0.0000***	***0.0000***
		Variance	0.602	0.474	0.569	0.935
	TMin	Mean	***0.0230***	***0.0142***	***0.0000***	***0.0022***
		Variance	0.913	0.989	0.86	0.795
Oct–Mar	Precip.	Mean	0.225	0.705	***0.0424***	0.404
		Variance	0.664	***0.1***	0.573	0.909
	TMax	Mean	0.157	0.698	0.373	***0.0194***
		Variance	0.251	0.718	0.864	***0.0626***
	TMin	Mean	***0.0004***	***0.0048***	***0.0000***	***0.0000***
		Variance	0.759	0.94	0.556	0.211
Annual	Precip.	Mean	0.295	0.933	0.498	0.688
		Variance	0.22	0.12	0.584	0.644
	TMax	Mean	0.628	0.304	***0.0008***	***0.0000***
		Variance	0.619	0.614	0.42	0.207
	TMin	Mean	***0.0000***	***0.0002***	***0.0000***	***0.0000***
		Variance	0.859	0.416	0.825	0.97

TMax

Mean TMax values are highly significant throughout the year between the intervals 1965–1991 and 1992–2018, with a p-value of 0.0000 (Table 6.1) and confine clearly the accelerated global warming over the last 50 years (Morice et al., 2012; National Centers for Environmental Information, 2020). The annual mean box-whisker plot illustrates this significance (Figure 6.2.2c).

TMin

There also is high significance in mean TMin values between all periods with all intervals and 1992–2018, apart from the notable exception of April-May between 1965–1991 and 1992–2018, where the p-value is 0.847 (Table 6.1), as shown in the box-whisker plots (Figures 6.2.2d, e). The most likely explanation is that clearer nighttime and early morning skies due to the highly significant decrease in precipitation, and possibly lighter winds, from the 1990s, has increased radiational cooling relative to 1965–1991. The very high significance in the TMin increase from 1938–1964 to 1965–1991 for all periods (Table 6.1) corresponds approximately to 1950, when the importance of global warming is documented (Pachauri and Meyer, 2014).

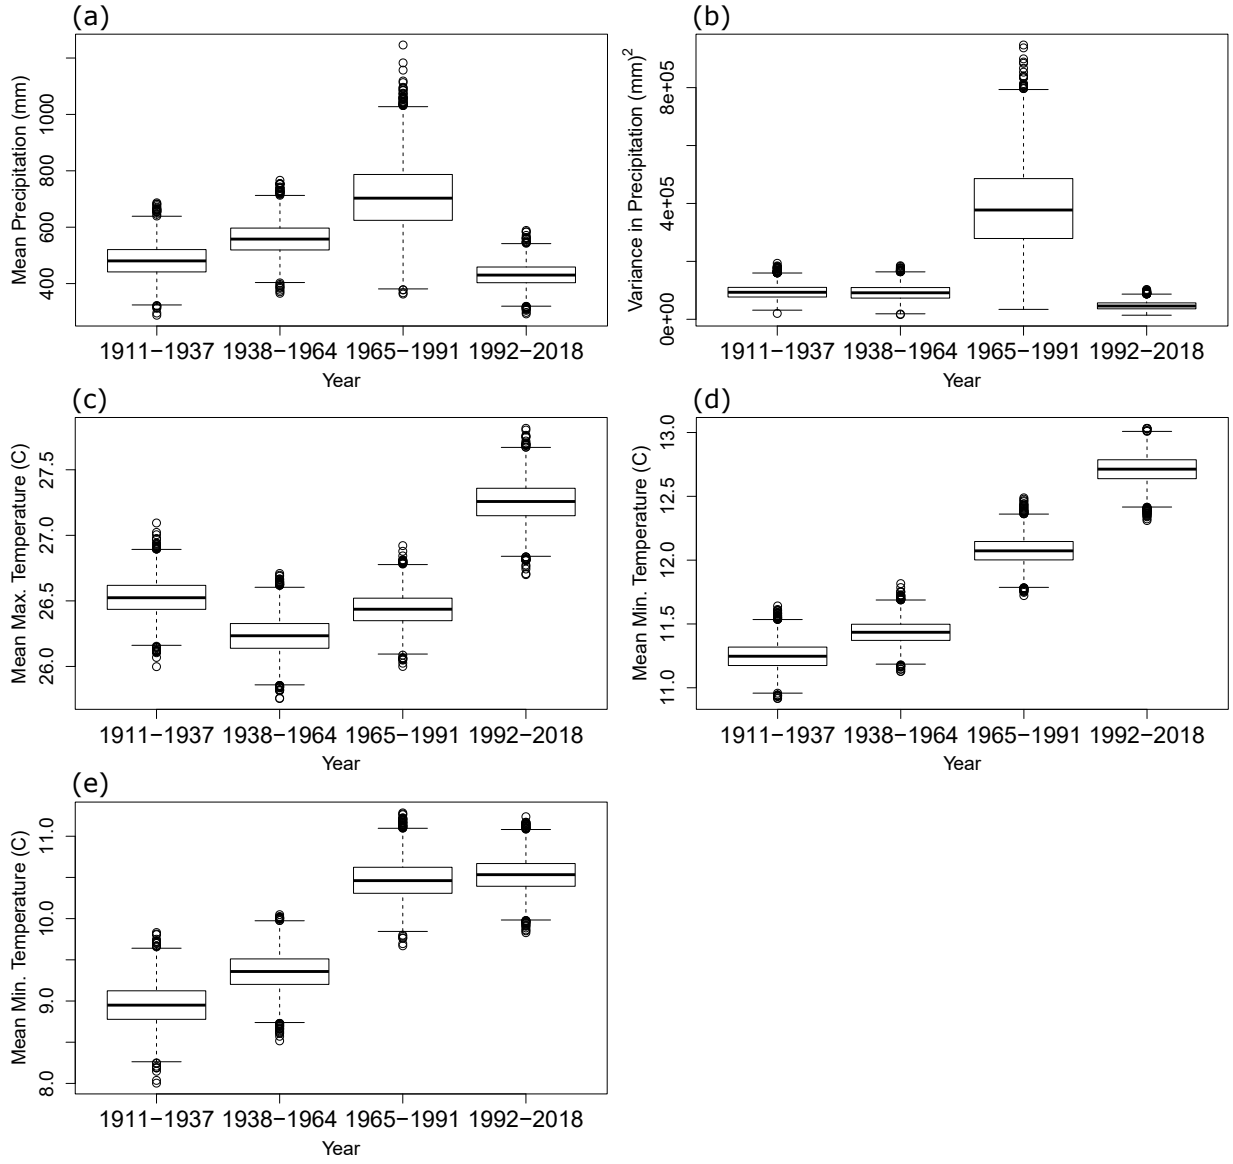

Figure 6.4.2: **Box and whisker plots of NMDB precipitation, TMax & TMin** Box and whisker plots of NMDB for (a) mean April-May precipitation, and (b) variance; (c) mean annual TMax; (d) mean annual TMin; (e) mean April-May TMin.

6.4.3 Wavelet analysis of temperature and precipitation 1911–2018

It is noted that oscillations in atmospheric and oceanic phenomena such as ENSO are expected to appear in wavelets of their time series. However, many climate drivers, such as the SAM, IPO, IOD, and the Atlantic Multidecadal Oscillation (AMO), are modes that operate on variable time scales. As such, they are not as well-represented in wavelet analyses as ENSO. Following Torrence and Compo (1998), wavelet analysis was applied to detrended anomalies of the precipitation and temperature time series. Wavelet analysis produces both a local and a global power spectrum. The local wavelet power spectrum reveals the temporal evolution of those periodic signals located in the time series, allowing the detection of oscillatory climate drivers such as ENSO, and shows their temporal evolution. The global power spectrum also is generated, and essentially is a Fourier transform, displaying the major amplitudes present in the time series. Here, the Morlet wavelet is chosen as the mother wavelet. The following annual, April-May, JJAS and October-March wavelets for the NMDB are interpreted for precipitation. The wavelets for TMax and TMin were very similar to precipitation, emphasizing the 2–8-year ENSO periodicity. Hence, they are not shown here.

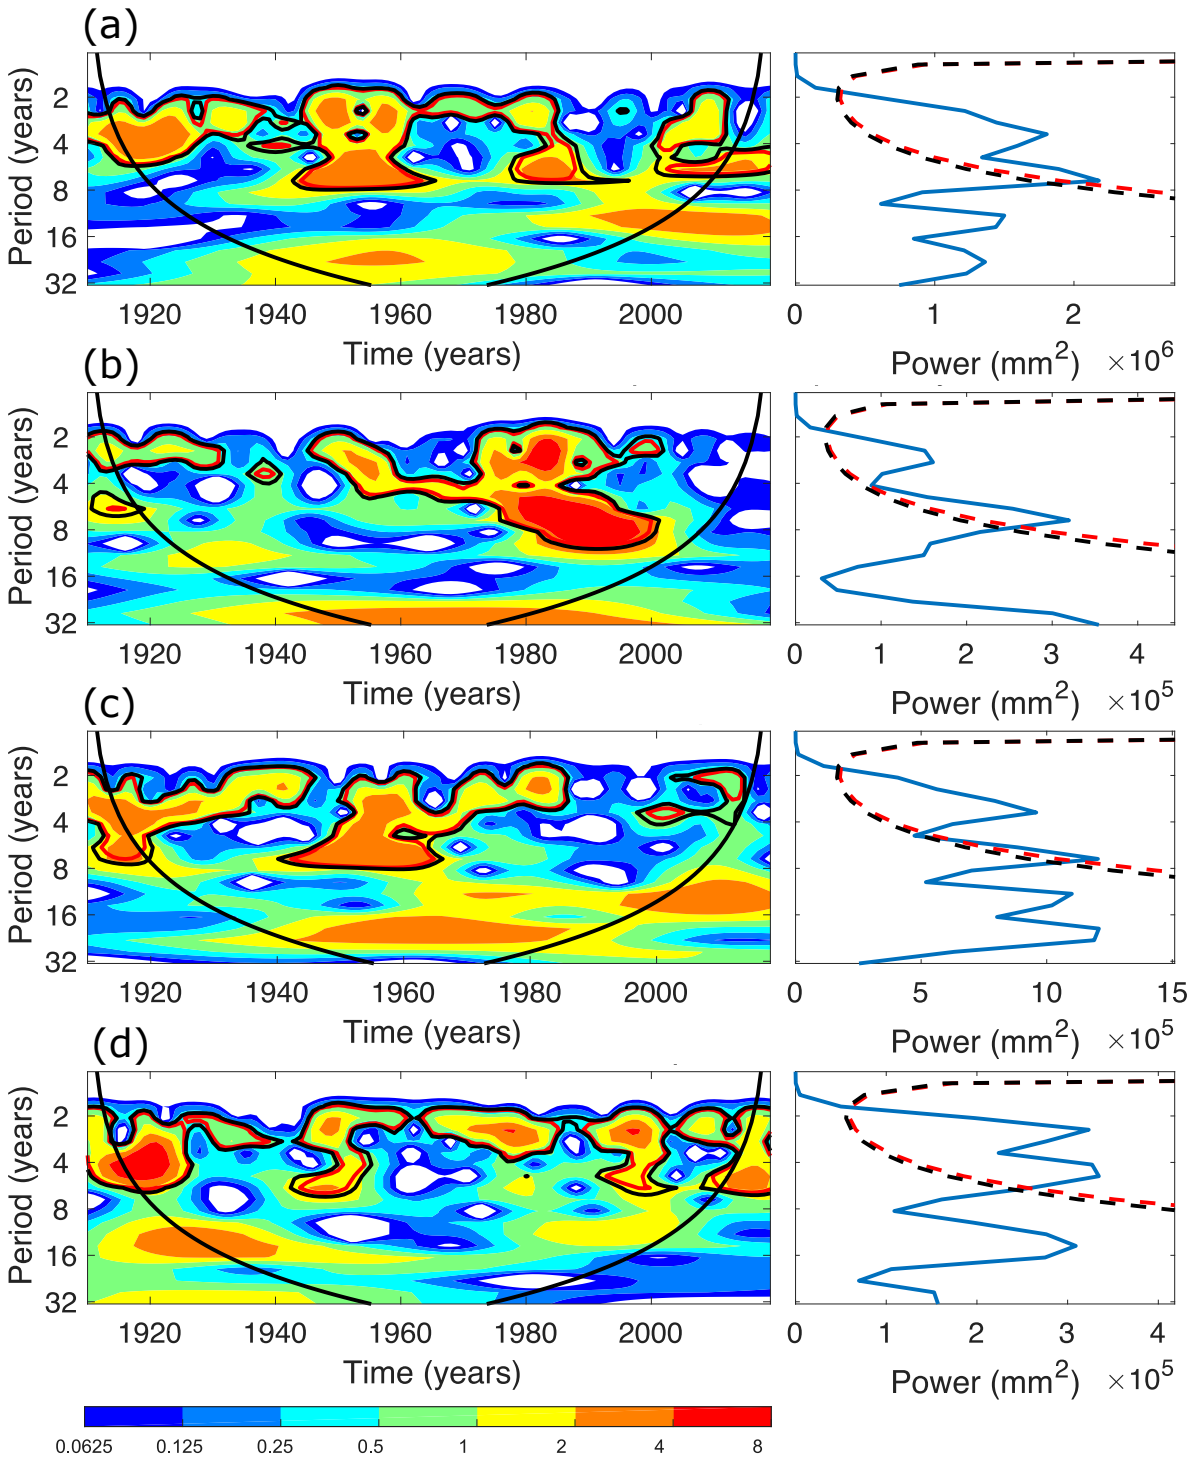

Figure 6.4.3: **Wavelets for precipitation** Wavelets representing NMDB precipitation for the periods (a) annual, (b) April-May, (c) JJAS, and (d) October-March. The dashed black and red lines are the 90th and 95th confidence percentiles, respectively.

Precipitation

The annual wavelet plot shows ENSO significance at both the 90th and 95th percentiles in the 2-8 year period (Figure 6.4.3a), dominated by October-March (Figure 6.4.3b) and JJAS (Figure 6.4.3c), and a significant period from 1950–2000 in April-May (Figure 6.4.3d). There is a notable non-significant, but possible, IPO periodicity in the 15-–30 year time frame. It is noteworthy that although it is not part of this study, the recent La-Niña of late 2020 and early 2021 provided increased precipitation in the northern MDB, highlighting the ongoing impact

Table 6.2: **Major precipitation attributes identified for each time period** The five major precipitation attributes identified, for each time period. They had the highest percentages of appearances in the 10-fold, cross-validation of the machine learning schemes, applied to the 1965–2018 observed precipitation data set.

Annual	April-May	JJAS	Oct-Mar
GlobalSSTA*TSSST	SOI	SOI	SOI
AMO*SAM	GlobalSSTA*TSSST	Niño3.4	DMI*SAM
AMO*Niño3.4	TPI	SAM	SAM
TPI	GlobalT	TPI	GlobalT*GlobalSSTA
DMI*Niño3.4	DMI	DMI*TSSST	SOI*TPI

of ENSO (Bureau of Meteorology and CSIRO, 2020).

TMax

The wavelets for TMax (not shown) also indicate that the ENSO periodicity is the dominant influence in all four periods. Interestingly, its significance in approximately 12-year periodicity starts from the 2000s, outside the cone of significance in October-March, and appears annually.

TMin

As for TMax, ENSO is the only periodicity detected in the TMin wavelets (not shown) and is apparent in all four time periods.

6.4.4 Attribute Selection

As described in Section 6.3.3, below, the possible attributes considered in this study are the AMO, DMI, GlobalSSTA, GlobalT, the Niño3.4, the TPI, the SAM, the SOI, and the TSSST. Two-way interaction terms between these predictors also were considered (e.g., AMO*DMI). The total number of potential attributes was 45. The precipitation data covers the period 1965–2018 and 10-fold cross-validation was applied to limit overfitting with the above techniques. Those attributes that appeared in at least 50% of the folds across the eight training methods are considered likely attributes of NMDB precipitation. However, because many attributes appeared in at least 50% of folds, only the five attributes with the highest percentages above 50% were selected here as the key attributes of NMDB precipitation.

For annual precipitation, these attributes were GlobalSSTA*TSSST, AMO*SAM, AMO*Niño3.4, TPI and DMI*Niño3.4 (Table 6.2). Two-way interactions were selected more often than either attribute in isolation, highlighting the importance of considering two-way interactions as potential attributes. For example, GlobalSSTA*TSSST occurs in 60% of folds, whereas GlobalSSTA occurs only in 46.25% and TSSST in 55%. For April-May precipitation, the key attributes were SOI, GlobalSSTA*TSSST, TPI, GlobalT and DMI. For JJAS precipitation those selected were SOI, Niño3.4, SAM, TPI and DMI*TSSST. Finally, the key drivers of October-March precipitation were SOI, DMI*SAM, SAM, GlobalT*GlobalSSTA and SOI*TPI (Table 6.2).

Notably, for each precipitation period, the selected attributes are related to global warming (GlobalT, GlobalSSTA, and TSSST), highlighting the profound influence that global warming now has on precipitation in the NMDB. This is particularly relevant for April-May precipitation, where it was shown in Table 6.1 that there is a statistically significant decrease between the periods 1965–1991 and 1992–2018, as two of the climate drivers are direct attributes of

accelerated global warming during 1992–2018. Additionally, Table 6.2 reveals that drivers related to the atmospheric-based SOI, the ENSO-based tropical Pacific Ocean indices, the TPI and DMI are known to influence precipitation across all precipitation groupings. Finally, AMO appears to play a role in annual precipitation in the NMDB by modulating the influence of SAM and Niño3.4.

6.5 Conclusions

Australia's northern Murray-Darling Basin (NMDB), which occupies a vast area of subtropical eastern Australia, is found to be experiencing the impacts of accelerated global warming on its catchment area rainfall seasonality and effectiveness. Observations of precipitation, mean maximum temperature (TMax) and mean minimum temperature (TMin) from 1910–2018 were divided into four consecutive 27-year periods, 1911–1937, 1938–1964, 1965–1991, and 1992–2018. They revealed that mean April-May precipitation over the NMDB, for the period 1992–2018, decreased significantly (p-value = 0.0318; Table 6.1) from that for the period 1965–1991. April-May are the late Autumn season months that moisten the catchment area prior to the cool season rainfall events that flow into the NMDB river systems. However, there is no compensating increasing trend in the remaining cool-season months of the year (JJAS) to offset the drying in April-May. In the warm season months (October-March), La-Niña phases have produced record rainfall in parts of southeast Australia in the years 2010-2012 and very recently in 2020–2022. Nevertheless, significant increases in mean TMax and TMin through JJAS and October-March further disrupt the hydrological cycle and reduce the rainfall effectiveness as the plant evapotranspiration increased in the period 1992–2018 due to rising values of TMax. Hence, this combination of decreased April-May precipitation and associated higher daytime temperatures decreases run-off into streams that feed into the Darling River and its tributaries, thereby reducing NMDB water availability, frequently with drastic consequences that are expected to worsen further in the future, as global warming continues.

The precipitation wavelets highlight the importance in 2–8 year periodicity of ENSO for the annual, JJAS (late cool-season) and October-March (warm season) periods. The climate drivers selected as likely attributes for the observed annual and seasonal precipitation patterns varied between the four different periods. For precipitation, the main attributes contributing to the trends in rainfall seasonality variability patterns were identified by using the time series in 10-fold machine learning models. It was found that the key attributes differed between the different periods. However, they had in common an emphasis on the role of global warming both individually and in combination with known local and large-scale climate drivers, notably ENSO, SAM, IOD, TPI, and Tasman Sea SSTs, because four of which, ENSO, IOD,TPI and Tasman Sea SSTs, being ocean-based time series, are so highly correlated with the recent quasi-linearly increasing trend in the Global SST time series.

Further work, beyond attribute identification, is planned, with the aim of developing sub-seasonal and seasonal outlooks, and annual prediction schemes, likely by combining high resolution climate model projections with machine leaning models. This hybrid approach is expected to be necessary as the traditional machine learning approach of using training and test data sets can rapidly lose predictive skill, particularly if there is a high level of non-stationarity in the training (observational) data set, and especially if this non-stationarity continues into the test data set.

Bibliography

Bergstrom, D., B. Wienecke, J. van den Hoff, L. Hughes, D. Lindenmayer, T. Ainsworth, C. Baker, L. Bland, D. Bowman, S. Brooks, J. Canadell, A. Constable, K. Dafforn, M. Depledge, C. Dickson, N. Duke, K. Helmstedt, A. Holz, C. Johnson, M. McGeoch, J. Melbourne-Thomas, R. Morgain, E. Nicholson, S. Prober, B. Raymond, E. Ritchie, S. Robinson, K. Ruthrof, S. Setterfield, C. Sgrò, J. Stark, T. Travers, R. Trebilco, D. Ward, G. Wardle, K. Williams, P. Zylstra, and J. Shaw. Combating ecosystem collapse from the tropics to the Antarctic. Glob. Change Biol., 27:1692–1703, 2021.

Breiman, L. Random forests. Mach. Learn., 45:5–32, 2001.

Bureau of Meteorology. Record-breaking la-Niña events. `http://www.bom.gov.au/climate/enso/history/La-Nina-2010-12.pdf`, 2012.

Bureau of Meteorology and CSIRO. State of the Climate 2020. `http://www.bom.gov.au/state-of-the-climate/documents/State-of-the-Climate-2020.pdf`, 2020.

Cai, W. and T. Cowan. Southeast Australia Autumn Rainfall Reduction: A Climate-Change-Induced Poleward Shift of Ocean-Atmosphere Circulation. J. Clim., 26:189–205, 2013.

CSIRO. Science solutions for the Murray-Darling Basin: Managing today's resources for the future. `https://publications.csiro.au/rpr/download?pid=csiro:EP206574&dsid=DS1`, 2020.

Hartigan, J., S. MacNamara, and L. Leslie. Application of Machine Learning to Attribution and Prediction of Seasonal Precipitation and Temperature Trends in Canberra, Australia. Climate, 8:76, 2020a.

Hartigan, J., S. MacNamara, L. Leslie, and M. Speer. Attribution and Prediction of Precipitation and Temperature Trends within the Sydney Catchment Using Machine Learning. Climate, 8:120, 2020b.

Hastie, T., R. Tibshirani, and J. Friedman. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media, 2009.

Holland, J., G. Luck, and C. Finlayson. Threats to food production and water quality in the Murray-Darling Basin of Australia. Ecosyst. Serv., 12:55–70, 2015.

Maldonado, S. and R. Weber. A wrapper method for feature selection using Support Vector Machines. Inf. Sci., 179:2208–2217, 2009.

McBride, J. and N. Nicholls. Seasonal Relationships between Australian Rainfall and the Southern Oscillation. Mon. Weather Rev., 111:1998–2004, 1983.

Morice, C., J. Kennedy, N. Rayner, and P. Jones. Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: the HadCRUT4 dataset. J. Geophys. Res., 117(D08101), 2012.

Murray-Darling Basin Authority. Hydrologic assessment of flow changes in the northern Basin. Technical Report 35/18, 2018.

National Centers for Environmental Information. State of the Climate: Global Climate Report for 2019. Technical report, 2020. https://www.ncdc.noaa.gov/sotc/global/201913/supplemental/page-3.

Pachauri, R. and L. Meyer. 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, 2014.

Pook, M., J. Risbey, and P. McIntosh. A comparative synoptic climatology of cool-season rainfall in major grain-growing regions of southern Australia. Theor. Appl. Climatol., 117: 521–533, 2014.

Richman, M. B. and L. M. Leslie. Machine Learning for Attribution of Heat and Drought in Southwestern Australia. Procedia Comput. Sci., pages 3–10, 2020.

Risbey, J., M. Pook, P. McIntosh, C. Ummenhofer, and G. Meyers. Characteristics and variability of synoptic features associated with cool season rainfall in southeastern Australia. Int. J. Climatol., 29:1595–1613, 2009.

Sheldon, F. Characterising the ecological effect of changes in the 'low-flow hydrology' of the barwon-darling river. Technical report, Australian Rivers Institute, 2017. Advice to the Commonwealth Environmental Water Holder Office. https://www.awe.gov.au/sites/default/files/documents/characterising-eco-effects-changes-low-flow-barwon-darling.pdf.

Speer, M., L. Leslie, S. MacNamara, and J. Hartigan. From the 1990s climate change has decreased cool season catchment precipitation and reduced river heights in Australia's southern Murray-Darling Basin. Sci. Rep., 11(16136), 2021.

Stephens, C., T. McVicar, F. Johnson, and L. Marshall. Revisiting pan evaporation trends in Australia a decade on. Geophys. Res. Lett., 45:11164–11172, 2018.

Torrence, C. and G. Compo. A Practical Guide to Wavelet Analysis. Bull. Am. Meteorol. Soc., 79(1):61–78, 1998.

Vapnik, V. The Nature of Statistical Learning Theory. Springer, 1995.

Chapter 7

High resolution simulations of a tornadic storm affecting Sydney

7.1 Overview

Severe convective thunderstorms (SCTs) occur over a relatively small spatiotemporal scale, yet can result in the most costly natural disasters in Australia with many events exceeding AU$1billion in damage, including the Sydney hailstorms of 1999 and 2018, and the Brisbane hailstorm of 2014 (Buckley et al., 2001, 2010; Soderholm et al., 2017; Insurance Council of Australia, 2019). Due to their small scale, there is often only a short lead time on warning the community for specific SCTs, although significant efforts have gone into understanding the large-scale conditions to better anticipate severe weather on the day. Despite the understandings of the large-scale conditions, there are SCTs that occur at unexpected times, and with potentially unexpected effects. On the morning of 16 December 2015, a severe thunderstorm developed off the coast of New South Wales and propagated over Kurnell, a suburb in Sydney, producing a tornado that caused significant damage, impacting Sydney's desalination plant, causing it to be offline for approximately three years. The tornado also came close to Sydney airport. Thankfully, the storm did not result in any casualties despite occurring in Australia's most populous city, but it produced widespread damage with insured losses of AU$206 million making it the most costly natural disaster of the Australian 2015/2016 summer. This storm was unusual as it developed during the morning over the ocean, rather than inland during the afternoon, which is more typical for severe storms impacting the Sydney region (Allen and Allen, 2016; Warren et al., 2020).

Considering the costly impact SCTs have on Australia, research is relatively limited, with a moderate number of studies aimed at understanding the climatological occurrence of storms, and some case studies being performed. However, there have only been a very small number of numerical modelling case studies focussed on Australia, and to my knowledge, there have been no numerical modelling studies of tornado-producing SCTs in Australia. The following Chapter presents a case study using the Weather Research and Forecasting (WRF) model to simulate the tornadic storm in Kurnell at high resolution. As lead author of this study, I conceived the study, collected the data and determined the model setup (including resolution and physics packages), ran the WRF simulations, analysed the simulation output and wrote much of the paper. Professor Lance Leslie, Dr Shev MacNamara and Dr Milton Speer assisted with determining the model setup, analysis and writing. This paper has been published in the *ANZIAM Journal* as part of the peer-reviewed proceedings for the Computational Techniques and Applications Conference.

7.2 Introduction

At approximately 10:30 AM local time (UTC+11 hours), on 16 December 2015, an EF2 rated tornado (Storm Prediction Center, 2014) impacted the Sydney suburb of Kurnell. The storm and associated tornado was the most costly natural disaster of the Australian 2015/2016 summer, with insured losses of AU$206 million (Insurance Council of Australia, 2016). The Sydney desalination plant was heavily impacted and required major repairs that were not completed until the end of 2018, months before needing to be used during the intense drought in 2019. Furthermore, the tornado was located only 4 km southeast of Sydney airport, Australia's busiest airport.

The storm itself was unusual as it developed during the morning and over the ocean, rather than the usual development inland during the afternoon, associated with maximum daytime heating and Convective Available Potential Energy (CAPE); this afternoon or evening storm development is more typical of severe thunderstorms in the Sydney region (Allen and Allen, 2016; Warren et al., 2020). Additionally, observational evidence suggests that this storm potentially produced multiple tornadoes over its lifetime (Richter et al., 2016).

Numerical simulations of severe convective storms can provide detailed information not available from routine weather observations. The output from these simulations can be used to better understand how storms developed and intensified, and how specific storm-related hazards were supported by the atmosphere. In this study, high resolution simulations are performed using the weather research and forecasting (WRF) model (Powers et al., 2017) as an initial investigation into the characteristics of the storm.

7.3 Data and methodology

The WRF model was configured on a double-nested domain centred on Sydney. The outer nest is 1050×1050 km^2 with 1.5 km horizontal resolution, while the inner nest is 450.5×500 km^2 with 0.5 km resolution. The model top is set at 50 hPa, with 33 vertical levels stretching from a resolution of approximately 52 m at the model surface to 1093 m at the model top. The Governing equations of the model are the non-hydrostatic Euler equations:

$$\partial_t U + \nabla \cdot \mathbf{V}u + \mu_d \alpha \partial_x p + (\alpha/\alpha_d)\, \partial_\eta p \partial_x \phi = F_U \tag{7.3.1}$$

$$\partial_t V + \nabla \cdot \mathbf{V}v + \mu_d \alpha \partial_y p + (\alpha/\alpha_d)\, \partial_\eta p \partial_y \phi = F_V \tag{7.3.2}$$

$$\partial_t W + \nabla \cdot \mathbf{V}w - g\left[(\alpha/\alpha_d)\, \partial_\eta p - \mu_d\right] = F_W \tag{7.3.3}$$

$$\partial_t \Theta_m + \nabla \cdot \mathbf{V}\theta_m = F_{\Theta_m} \tag{7.3.4}$$

$$\partial_t \phi + \mu_d^{-1}\left[(\mathbf{V} \cdot \nabla\phi) - gW\right] = 0 \tag{7.3.5}$$

$$\partial_t Q_m + \nabla \cdot \mathbf{V}q_m = F_{Q_m} \tag{7.3.6}$$

where $\mathbf{v} = (u, v, w)$ represents the zonal, meridional and vertical components of wind, θ_m is moist potential temperature, $\phi = gz$ is the geopotential, g is gravitational acceleration, p is the full pressure, q_m represents moisture mixing ratios, μ_d is the vertical coordinate metric, α_d is the inverse density of dry air, α is the inverse density of the full parcel and η is a terrain-following hydrostatic pressure vertical coordinate. Furthermore, $\mathbf{V} = \mu_d \mathbf{v} = (U, V, W)$, $\Theta_m = \mu_d \theta_m$ and $Q_m = \mu_d q_m$. The right hand side of the governing equations represents forcing terms from model physics, turbulent mixing, spherical projections and Earth's rotation. Finally, the diagnostic equation for hydrostatic pressure is

$$\partial_\eta \phi = -\alpha_d \mu_d. \tag{7.3.7}$$

The model is solved with an integration time step of 3 s, using a time-splitting third-order Runge-Kutta scheme, where acoustic modes are integrated over a shorter time step. For more details on the model equations, and how they are solved, the reader is referred to the WRF technical notes (Skamarock et al., 2019).

Microphysical parameterisation describes various cloud-scale processes such as evaporation and melting. These processes can be represented differently, affecting the resulting storm morphology (Dawson II et al., 2010). As such, this study has performed simulations using the Morrison (Morrison et al., 2005, 2009) and the NSSL 2-moment 4-ice (Mansell et al., 2010) microphysics schemes, which are represented in the same way on both domains. The Morrison scheme predicts the mass of cloud, rain, ice, snow, and hail, and the number concentration of all these particles except cloud drops. Whereas the NSSL scheme predicts the mass and number concentration of cloud, rain, ice, snow, graupel, and hail. These schemes utilise different equations to describe the mass, number concentration and terminal fall velocity of particles which could significantly affect the simulation results. In order to simplify the results, the simulations presented here did not include longwave or shortwave radiation, boundary layer physics or surface physics.

The initial and boundary conditions used in this study are the European Centre for Medium-Range Weather Forecast's (ECMWF) ERA5 reanalysis data set (Hersbach et al., 2020), which was originally on a 30 km grid, and has been interpolated to the 1.5 km outer domain. The simulations are set to begin on 15/12/2015 at 0000 UTC, and end on 16/12/2015 at 0900 UTC. The inner domain data was output every 5 min, while the outer domain was output every 15 min.

7.4 Event overview

Over the course of the morning, a high pressure system was located off the southeast coast of Australia, directing moist surface easterly winds over New South Wales (NSW; Figure 7.4.1a). Meanwhile, a 500 hPa trough and associated cold pool extended from the southeast coast inland towards the northwest of the continent (Figure 7.4.1b). This resulted in a broad region of unstable air along the NSW coast, with highly unstable air concentrated over the Sydney region and adjacent ocean (CAPE $\geq 1000\,\mathrm{J\ kg^{-1}}$, TT ≥ 50 K; Figure 7.4.1 c, d). The marine boundary layer was capped in the early morning; however, this cap was removed during the morning likely due to moistening and cooling of the atmosphere above the marine boundary layer (Richter et al., 2016).

Prior to development of the Kurnell tornadic storm, there was some weak convection over the Great Dividing Range extending into the Sydney and central coast regions (not shown). A cluster of convection developed off the Illawarra coast at approximately 1700 UTC 2015-12-15 (not shown). The Kurnell tornadic storm developed from the merging of two cells near Kiama at 2100 UTC (not shown). The storm tracked north, adjacent to the coast, as a supercell thunderstorm. It then crossed the coastline and tracked through Kurnell, producing an EF2 rated tornado (Figure 7.4.2). After passing through Kurnell and the eastern suburbs of Sydney, the storm moved offshore in a northeasterly direction and dissipated.

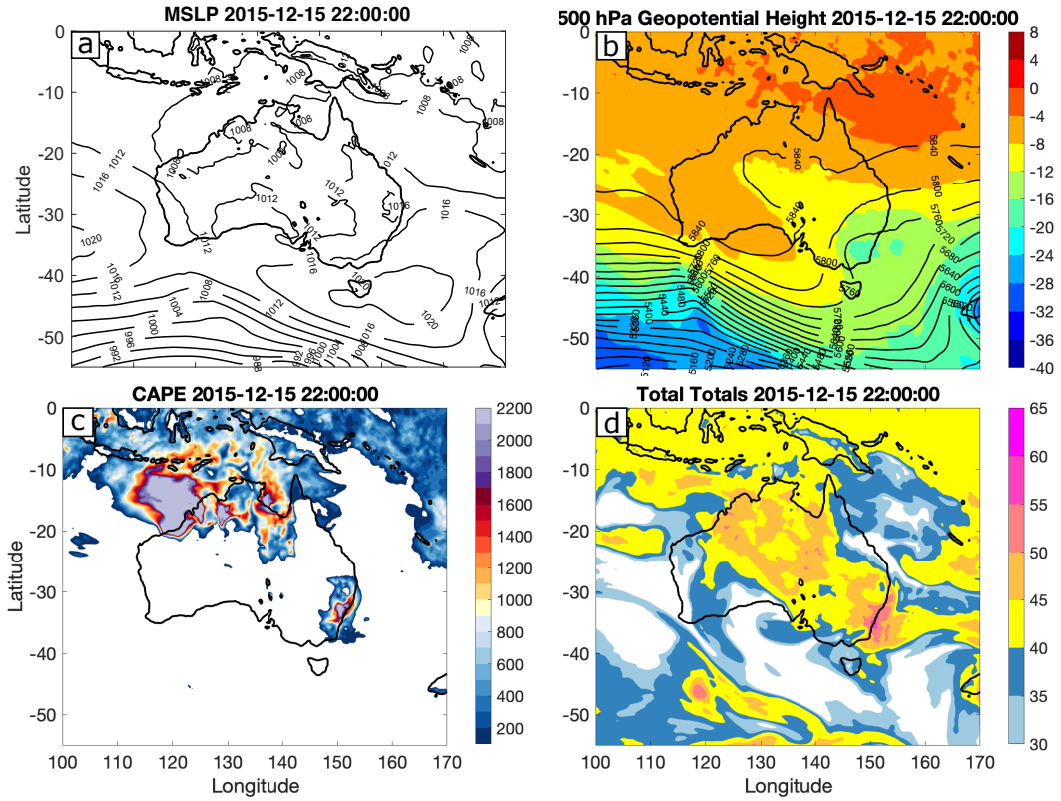

Figure 7.4.1: Atmospheric conditions from ERA5 reanalysis data, at the time of development of the Kurnell tornadic storm. Variables shown are (a) mean sea level pressure (hPa); (b) 500 hPa geopotential height (m) represented as thin black contours, and temperature (°C) represented by filled contours; (c) convective available potential energy (CAPE; J kg^{-1}); and (d) total totals index (TT; K, Miller, 1972).

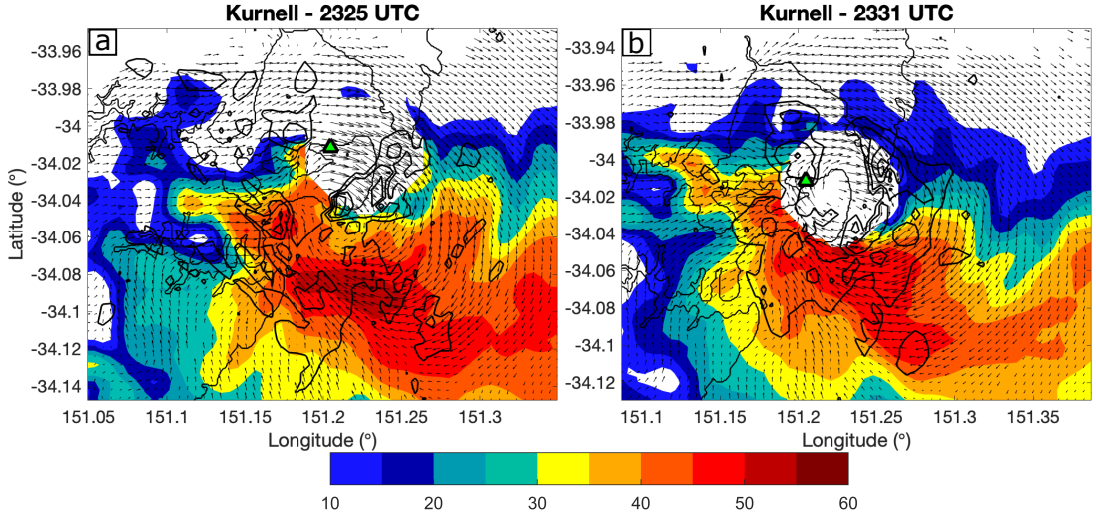

Figure 7.4.2: Filled contours display radar reflectivity (dBZ), thin black contours are vertical velocity (every 10 m s^{-1} starting at 5 m s^{-1}), and vectors show wind speed and direction, from the Kurnell radar interpolated to 2.5 km above sea level, at (a) 2325 UTC and (b) 2331 UTC. The location of Kurnell is provided by a solid green triangle. White regions are where no reflectivity is observed by the radar. The white circle in the figure is due to the storm propagating over the Kurnell radar, where sampling does not take place. However, vectors still exist in this circle as the 3D wind field is retrieved using multiple radars that sample this region.

7.5 Simulation results

Similar to observations, both simulations display convection forming off the Illawarra coast at 1600 UTC, with additional weak convection occurring inland of the Sydney and central coast region (not shown). Both microphysics schemes display a merging of two cells near Kiama, resulting in one strong storm cell propagating over the ocean adjacent to the coast, although this occurs approximately one hour earlier than the observed cell. Additionally, the storm produced in the Morrison scheme displays a quasi-linear convective system (QLCS) morphology, rather than supercell morphology (not shown). The Morrison scheme produces a larger and stronger cold pool (not shown). Investigation into the causes of this is beyond the scope of the present article but could be due to greater numbers of small rain particles being simulated in the Morrison scheme, leading to greater evaporative cooling, or it could be due to smaller hail being simulated leading to greater cooling through melting. This larger and stronger cold pool in the Morrison scheme could lead to environmental interactions that favour a QLCS mode. Due to the difference in the simulated convective mode under the Morrison scheme, the results herein will focus on the simulation that used NSSL microphysics.

The NSSL storm enters the inner domain at approximately 2115 UTC, with one main updraft that is located immediately behind the leading edge of the cold pool (not shown). This storm begins displaying mid-level updraft rotation at 2150 UTC (not shown), followed by the development of a hook echo at the lowest model level at 2215 UTC (Figure 7.5.1a), which is a common radar characteristic of supercell convection, and an area where tornadogenesis is most likely to occur. Additionally, there is a clearly defined rear flank gust front intersecting with the forward flank gust front at the hook echo region, with strong inflow directed into this area (Figure 7.5.1). This low-level mesocyclone initially draws in warmer environmental air from ahead of the gust front (Figure 7.5.2a). However, strong outflow on the rear flank appears to push the rear flank gust front away from the low-level mesocyclone (Figure 7.5.2b). This likely caused the demise of the first vortex at 2330 UTC, due to ingestion of cooler, more stable air.

The simulated storm is in the vicinity of Kurnell at 2240 UTC (Figure 7.5.1b), and continues propagating north through Sydney. Another low-level mesocyclone develops over the ocean at 2240 UTC, becoming clearly defined at 2245 UTC (Figures 7.5.1b and 7.5.2c). This mesocyclone progresses along the rear flank gust front towards the forward flank gust front, dragging warm inflow air into its circulation at 2300 UTC before dissipating (Figure 7.5.2d). It is possible that interactions of the low-level mesocyclone with the forward flank gust front led to its demise. After the storm passes through Sydney, it begins to propagate towards the east-northeast and change morphology into a QLCS. The storm is supported by convergence with the cold pool boundary against the northeasterly flow in the marine layer for some time until dissipating (not shown).

During the course of the event, there were numerous reports of wind gusts over 100 km h^{-1} and hail at most 6 cm in diameter across Sydney (Bureau of Meteorology, 2012). Figure 7.5.3 displays swath plots of maximum wind speed at 10 m, and maximum hail diameter at the surface, between 2210 and 2320 UTC from the NSSL simulation. The surface wind swath displays a large region of wind speeds over 25 m s^{-1}, with some areas above 30 m s^{-1} (Figure 7.5.3b). Furthermore, the maximum hail diameter swath indicates a relatively narrow region of hail between 3 and 4 cm in diameter (Figure 7.5.3b). Although the severe surface winds and hail all occur offshore in this simulated storm, rather than in parts of Sydney, the simulation indicates the environment was supportive of storms that could produce numerous severe weather hazards.

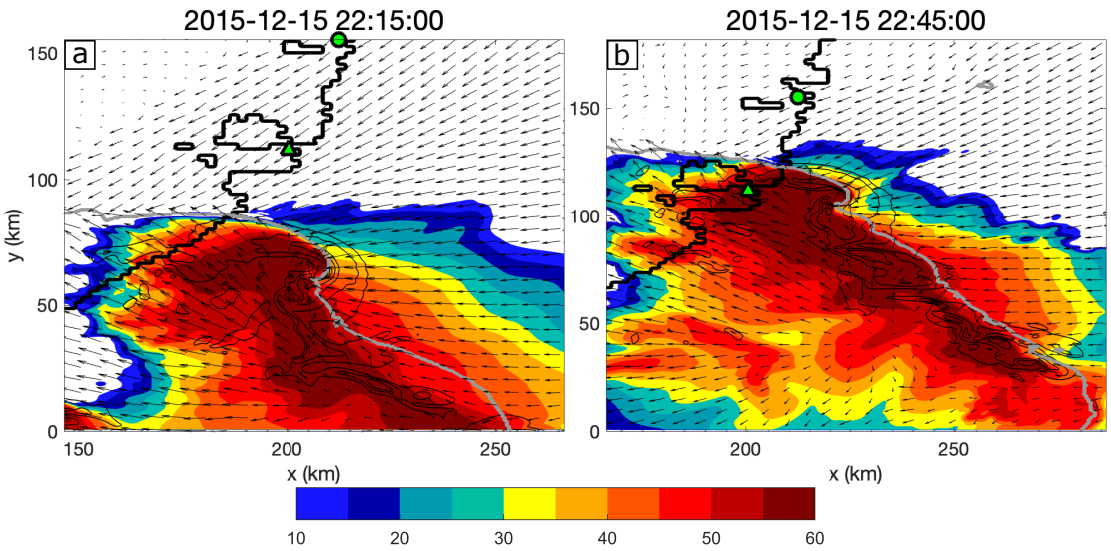

Figure 7.5.1: Variables shown are lowest model level reflectivity (filled contours; dBZ), wind speed and direction (vectors), and cold pool outline (thick grey contour representing $\theta' < -1$ K), and mid-level vertical velocity (eta-level = 17, approximately 5 km above sea level, thin black contours; every 10 m s^{-1} beginning at 5 m s^{-1}), for the NSSL simulation. The location of Sydney is provided by a solid green circle, and the location of Kurnell is provided by a solid green triangle.

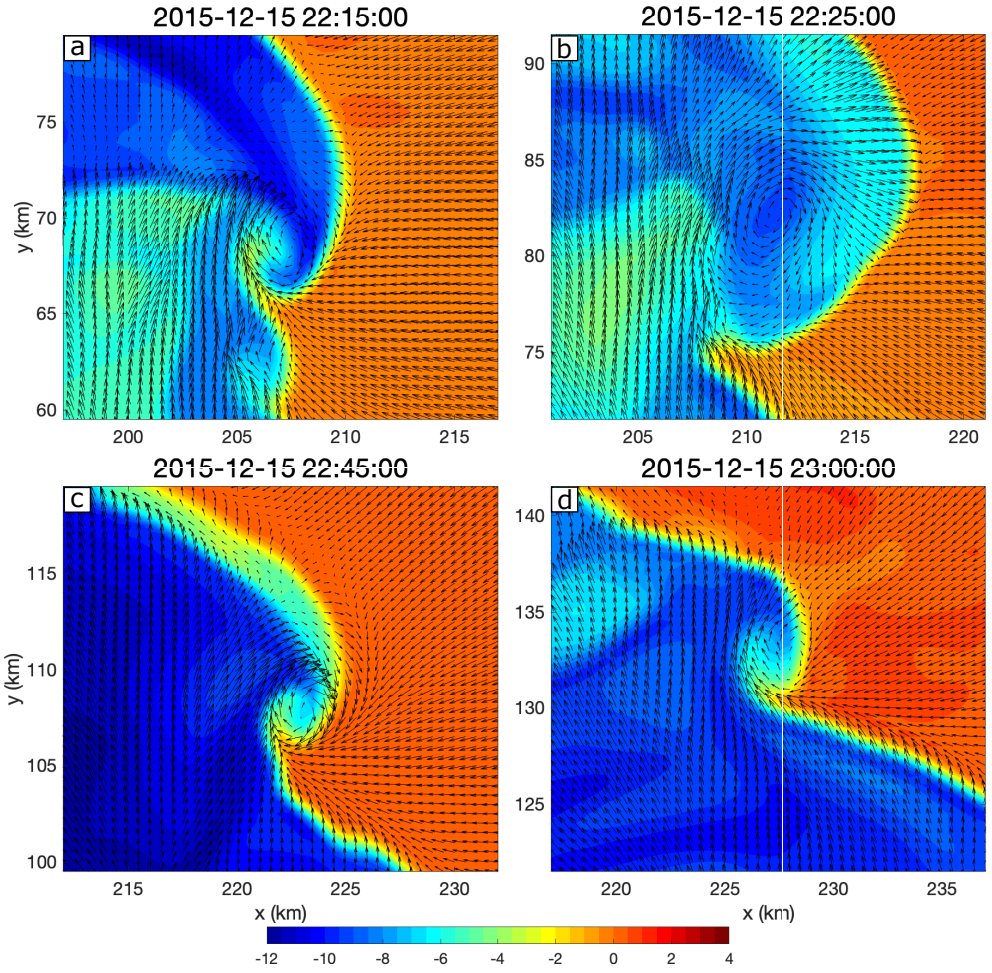

Figure 7.5.2: Filled contours display moist potential temperature perturbations (K) at the lowest model level, while vectors display lowest model level wind speed (m s^{-1}) and direction. All figures are from the NSSL simulation. This figure is focused on the region of rotation, so the coastline is not included but is immediately to the left of the domain shown.

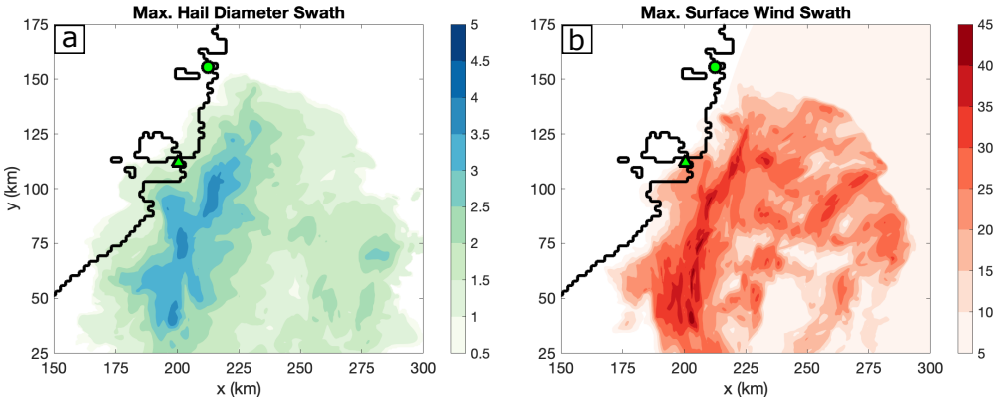

Figure 7.5.3: Simulated (a) maximum surface hail diameter (cm) and (b) maximum 10 m wind speed (m s^{-1}), between 2210 and 2320 UTC for the NSSL simulation. The location of Sydney is provided by a solid green circle, and the location of Kurnell is provided by a solid green triangle.

7.6 Conclusions and future work

In this study, two simulations of the tornadic storm that impacted Kurnell, Sydney, were performed using the WRF model with initial and boundary conditions from the ECMWF's ERA5 reanalysis data set. Although the simulations did not include longwave or shortwave radiation, boundary layer or surface physics, they both accurately depict the development of a storm at a similar time and location, and through similar mechanisms as the observed Kurnell storm.

Additionally, both simulations display this storm propagating north along the coast before impacting the Kurnell and Sydney region, and propagating back out to sea where it dissipated. This is despite the storm simulated under the Morrison microphysics scheme displaying a different morphology to that of the Kurnell and NSSL storms. Storm trajectories typically are dependent on the storm morphology due to the way the storm interacts with the ambient environment and maintains its structure. The tendency for both simulated storms to follow similar development patterns and trajectories despite different microphysics schemes was very promising. This similarity in the simulations indicates how strongly supportive the environment was for a severe convection event to develop near Sydney. However, the different storm morphology simulated in the Morrison microphysics scheme also highlights the importance of the common practice of running simulations using more than one microphysical scheme, in both research and operational settings, to capture a range of possible storm morphologies.

The storm simulated under the NSSL scheme displayed supercell characteristics, with both mid-level and low-level rotation evident. This rotation was cyclical, which was observed by Doppler radar on the day (Richter et al., 2016), providing further evidence that the storm could have produced numerous tornadoes over the ocean prior to impacting Kurnell. The NSSL storm also produced large hail and severe surface winds, although they were simulated over the ocean rather than in the Kurnell and greater Sydney region as observed. This difference is possibly due to the exclusion of surface and boundary layer physics in the model, although the results presented here are still remarkably accurate. Future work will focus on producing higher resolution simulations using the entire model physics. This will allow detailed vorticity and parcel trajectory analysis in order to better understand the contributing factors to tornadogenesis. Additionally, the Bureau of Meteorology's BARRA reanalysis data set, which is higher resolution than ERA5, will be used for initial and boundary conditions in future simulations. The results will be compared against those from ERA5, as the higher resolution would allow for higher resolution simulations to be performed at a lower computational cost.

Bibliography

Allen, J. and E. Allen. A review of severe thunderstorms in Australia. Atmos. Res., 178: 347–366, 2016.

Buckley, B., L. Leslie, and Y. Wang. The Sydney Hailstorm of April 14, 1999: Synoptic description and numerical simulation. Meteorol. Atmospheric Phys., 76:167–182, 2001.

Buckley, B., W. Sullivan, P. Chan, and M. Leplastrier. Two Record Breaking Australian Hailstorms: Storm Environments, Damage Characteristics and Rarity. In 25th Conference on Severe Local Storms, pages 11–14. American Meteorological Society Denver, Colorado, 2010.

Bureau of Meteorology. Severe Storms Archive, 2012. URL `http://www.bom.gov.au/australia/stormarchive/`.

Dawson II, D., M. Xue, J. Milbrandt, and M. Yau. Comparison of Evaporation and Cold Pool Development between Single-Moment and Multimoment Bulk Microphysics Schemes in Idealized Simulations of Tornadic Thunderstorms. Mon. Weather Rev., 138:1152–1171, 2010. doi: 10.1175/2009MWR2956.1.

Hersbach, H., B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. The ERA5 global reanalysis. Q. J. R. Meteorol. Soc., 146:1999–2049, 2020.

Insurance Council of Australia. Victorian bushfire losses push summer catastrophe bill past $550m, 2016. URL `https://www.insurancecouncil.com.au/media-centre#archiveMar-2016`.

Insurance Council of Australia. Insurance bill for summer catastrophes passes 2.2 billion, 2019.

Mansell, E., C. Ziegler, and E. Bruning. Simulated Electrification of a Small Thunderstorm with Two-Moment Bulk Microphysics. J. Atmos. Sci., 67:171–194, 2010. doi: 10.1175/2009JAS2965.1.

Miller, R. Notes On Analysis and Severe Storm Forecasting Procedures of the Air Force Global Weather Central. Technical Report 200, Air Weather Service, United States Air Force, 1972.

Morrison, H., J. Curry, and V. Khvorostyanov. A New Double-Moment Microphysics Parameterization for Application in Cloud and Climate Models. Part I: Description. J. Atmos. Sci., 62(6):1665–1677, 2005.

Morrison, H., G. Thompson, and V. Tatarskii. Impact of Cloud Microphysics on the Development of Trailling Stratiform Precipitation in a Simulated Squall Line: Comparison of One- and Two-Moment Schemes. Mon. Weather Rev., 137:991–1007, 2009. doi: 10.1175/2008MWR2556.1.

Powers, J., J. Klemp, W. Skamarock, C. Davis, J. Dudhia, D. Gill, J. Coen, D. Gochis, R. Ahmadov, S. Peckham, G. Grell, J. Michalakes, S. Trahan, S. Benjamin, C. Alexander, G. Dimego, W. Wang, C. Schwartz, G. Romine, Z. Liu, C. Snyder, F. Chen, M. Barlage, W. Yu, and M. Duda. The Weather Research and Forecasting Model: Overview, System Efforts, and Future Directions. Bull. Am. Meteorol. Soc., 98:1717–1737, 2017. doi: 10.1175/BAMS-D-15-00308.1.

Richter, H., A. Protat, J. Taylor, and J. Soderholm. Doppler radar and storm environment observations of a maritime tornadic supercell in Sydney, Australia. In Preprints, 28th Conf. on Severe Local Storms, Portland, OR, Amer. Meteor. Soc. P, 2016.

Skamarock, W., J. Klemp, J. Dudhia, D. Gill, Z. Liu, J. Berner, W. Wang, J. Powers, M. Duda, D. Barker, and X.-y. Huang. A Description of the Advanced Research WRF Model Version 4. Technical Report Note NCAR/TN-556+STR, National Center for Atmospheric Research, 2019.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Wedd, and T. Weckwerth. Diurnal Preconditioning of Subtropical Coastal Convective Storm Environments. Mon. Weather Rev., 145:3839–3859, 2017.

Storm Prediction Center. The Enhanced Fujita Scale (EF Scale), 2014. URL `https://www.spc.noaa.gov/efscale/`.

Warren, R., H. Ramsay, S. Siems, M. Manton, J. Peter, A. Protat, and A. Pillalamarri. Radar-based climatology of damaging hailstorms in Brisbane and Sydney, Australia. Q. J. R. Meteorol. Soc., 146:505–530, 2020. doi: 10.1002/qj.3693.

Chapter 8

An Australian Severe Convective Thunderstorm Hazard Climatology

8.1 Overview

Severe convective thunderstorms (SCTs) are devastating both economically and to human life across Australia, with one quarter of all natural hazard-related losses in Australia resulting from SCTs (Middelmann, 2007; Crompton and McAneney, 2008). However, much of the Australian understanding of SCTs comes from research conducted mostly in the United States of America (USA), which generally focuses on damaging supercell thunderstorms. This research has proven helpful in shaping the Australian understanding of SCTs, but it is not sufficient for accurate and informed forecasting of SCTs in Australia. Many of the environments in which they form in Australia differ from that experienced in the USA. The environments also vary across the country, which has resulted in a range of unexpected impacts and losses (Allen and Allen, 2016b).

Climatologies describe the average atmospheric conditions at some specified location, often throughout the year. Climatologies of SCTs often are developed to improve understanding of their frequency and the environments that support them, and ultimately better anticipate SCTs and their associated hazards before they form. In Australia, SCT climatologies have been developed with a focus on general SCTs rather than the individual hazards they produce, except for localised studies in major cities such as Brisbane (e.g., Soderholm et al., 2017a) or Sydney (e.g., Schuster et al., 2005). Therefore, little is known on how the formative environments of SCTs vary in Australia for each individual hazard. Additionally, there is difficulty in developing climatologies due to the limited spatial and temporal extent of observations (Schuster et al., 2005; Allen et al., 2011). Climatological studies are very much focused on developing a statistical understanding of the environments supporting a specific form of severe weather. In this Chapter, a case study is presented on the development of an SCT hazard climatology; that is, individual climatological understandings of severe convective storms that produce one of either severe wind, large hail, or tornadoes, across Australia. These climatologies are developed by applying logistic regression and machine learning techniques to draw relationships from a data set that is relatively small. This work assesses the potential benefits of applying such techniques to develop climatologies and improve understanding of SCTs and their associated hazards in Australia, for potentially more detailed studies of this kind in the future. The work presented in this Chapter has not yet been developed into a paper for the peer-reviewed literature.

8.2 Introduction

SCT climatologies provide an understanding of the spatiotemporal distribution of SCTs and conditions supportive of them (e.g., Doswell, 1980; Brooks et al., 2003; Allen et al., 2011). The impact of climate drivers such as the El-Niño Southern Oscillation (ENSO) may also be studied, which assists seasonal prediction of severe convective weather (Yeo, 2005; Lu et al., 2015; Childs et al., 2018). The key to SCT climatologies are observations of a severe thunderstorm, which is defined by the Australian Bureau of Meteorology (2017) as convective storms that produce at least one of: hail at least 2 cm in diameter, wind gusts of at least 90 km h^{-1}, a tornado, or heavy rain conducive to flash flooding. However, most weather organisations do not consider heavy rainfall as a characteristic of severe thunderstorms as non-convective weather can also be conducive to flash flooding. Furthermore, heavy rainfall events can result from *quasi-stationary convective systems*, which form as convection continually initiates over one location and are often not tied to other severe phenomena (Chappell, 1986). Australian SCT reports are maintained in the publicly available severe thunderstorm archive (STA; Bureau of Meteorology, 2012), which contains reports from 1795 to present. Population density and the small scale nature of severe convective weather leads to a clustering of reports around the cities and east coast of Australia, and a sparse distribution elsewhere (Doswell, 1980; Allen et al., 2011). This inconsistent data collection is an issue worldwide, leading to gaps in our understanding of SCTs (Brooks et al., 2003).

Climatologies of thunderstorm frequency place an upper bound on SCT activity, finding maxima during the spring and summer across northern Australia and from southeast Queensland across central and eastern NSW into northeastern Victoria (Kuleshov et al., 2002; Dowdy and Kuleshov, 2014). Many Australian SCT climatology studies have focused on one region (e.g., Matthews and Geerts, 1995; Schuster et al., 2005; Yeo, 2005). While regional analysis is important (Brooks et al., 2003; Brooks, 2009), a comprehensive understanding of the climatology across Australia is necessary. Recent work utilised *pseudo-proximity soundings* derived from atmospheric reanalysis data (Allen et al., 2011; Allen and Karoly, 2014). From these data, a covariate relationship (Brown and Murphy, 1996) was developed relating SCT occurrence to the product of CAPE and 0–6 km bulk shear (S06). The estimated frequency of SCT environments resulted in a maximum across eastern and northern Australia during spring and summer (Allen et al., 2011; Allen and Karoly, 2014). The discriminant relationship of Allen et al. (2011) was limited by the short temporal sample size of the data that it was developed on, and the assumptions that the discriminant holds over all climate regions across Australia and has equal importance for each SCT hazard.

Remote sensing techniques have improved accuracy of the climatological record in regions where SCT observations are sparse (e.g., Lukach et al., 2017; Soderholm et al., 2017a; Shikhov and Chernokulsky, 2018). This includes the use of radar data (e.g., Lukach et al., 2017; Soderholm et al., 2017a), lightning flash counts (e.g., Schultz et al., 2011), and satellite observations (e.g., Cecil and Blankenship, 2012; Ferraro et al., 2015; Bedka et al., 2018; Shikhov and Chernokulsky, 2018). Applying filtering techniques to satellite observations, a region from north of Brisbane down to Sydney and the Goldfields-Esperance region were found to be the most hail-prone areas of Australia (Bedka et al., 2018). Methods of convection initiation are often not considered in climatologies. However, Soderholm et al. (2017a) utilised radar and surface observations to identify mechanisms of severe hailstorm initiation and development in southeast Queensland. Additional work using Australian radar data for understanding the hail climatology has suggested the Brisbane region experiences 26 damaging hail days per annum, and Sydney experiences 32 damaging hail days, although there was a greater amount of missing radar data

at Brisbane that might influence the results (Warren et al., 2020). Furthermore, for 10 radar sites in Australia, hail events were found to be characterised by environments with greater wind shear, while lightning events had a higher freezing level (Dowdy et al., 2020). Despite these advances in research, there is a lack of coverage of remote-sensing instruments over the Australian continent and a short temporal record that limits the applicability of remote-sensing techniques to Australian SCT climatologies (Walsh et al., 2016). However, these techniques are expected to become highly useful in the future as the length of the record increases, spatial coverage improves and more advanced remote-sensing technology is employed (Doswell, 2001; Walsh et al., 2016).

Despite the widespread belief that Australian tornadoes are infrequent and not as intense as those in the USA (Minor et al., 1980; Allen and Allen, 2016b), numerous events have occurred throughout history (e.g., Sills et al., 2004; Allen, 2012; Taylor et al., 2016; Bureau of Meteorology, 2016). The perceived lack of threat from tornadoes in Australia has been fueled by underreporting and the low number of strong tornadoes hitting a major city (Concannon et al., 2000; Antonescu et al., 2017). However, the potential for a tornado to be observed will increase alongside an increasing population, as will the economic and societal impact of tornadic events (Antonescu et al., 2017). The climatology of Australian tornadoes has had little development, limited by the small observational record (Allen and Allen, 2016b). Early work expanded the observational record, with varying claims on the annual frequency of tornadoes (Evesson, 1970; Minor et al., 1980; Geerts and Noke-Raico, 1995). A more recent expansion of the STA suggests that Australia experiences an annual average of 60 tornadoes, with these events forming in a variety of environments not limited to typical SCT environments (Allen and Allen, 2016a).

Australian tornadoes have a bi-modal seasonal cycle with maxima in early spring/summer, and in winter (Griffiths et al., 1993; Geerts and Noke-Raico, 1995; Hanstrum et al., 2002). Approximately half of all Australian tornadoes are *cool-season tornadoes* (Hanstrum et al., 2002). The majority of cool-season tornadoes in Australia occur in southwest Western Australia (WA) and in South Australia (SA), with another peak in activity on the NSW coast (Hanstrum et al., 2002). Cool-season tornadoes are associated with intense extratropical cyclones, with environments containing low surface-based instability and high low-level VWS (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009). Little research into Australian warm-season tornadoes has been performed, with global climatology studies finding a high number of warm-season tornado environments over eastern Australia (Brooks et al., 2003; Tippett et al., 2015). This is only a foundation for further work in characterising warm-season tornado environments as the environments conducive to tornadoes in Australia are not necessarily all sampled in the USA, from which the discriminant in Brooks et al. (2003) was founded (Brooks, 2009).

SCTs significantly impact society and the surrounding environment and require a local understanding to improve forecasting and emergency management of events (Brooks, 2009; Antonescu et al., 2017). However, the majority of research towards SCTs has been conducted in, and focused on, the Great Plains of the USA. Much of the Australian SCT climatology work has been directed to SCTs as they are generally defined (e.g., Allen et al., 2011; Allen and Karoly, 2014), or regional hail studies (e.g., Yeo, 2005; Soderholm et al., 2017a). More recent work has assessed the hail climatology across Australia (Bedka et al., 2018) and for numerous radar sites across parts of Australia (Dowdy et al., 2020), and the Australian wind climatology including changes in historical and future environments (Brown and Dowdy, 2021a,b). Although historical changes in the spatial frequency of environments supportive of wind SCTs was assessed in Brown and Dowdy (2021b), understanding of the baseline spatial frequencies was not provided. There have only been small studies in understanding the tornadic SCT climatology, with

the environments and climate drivers of cool season-tornadoes fairly well classified (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009), but warm-season tornadoes poorly understood. Furthermore, the climatology for each convective hazard across Australia has not been assessed using the same methodology for each hazard.

Until more recently, much of the work has focused on the discriminant relationship of Brooks et al. (2003), considering combinations of CAPE and S06. While these two variables are important for SCT development, there are many other variables that influence deep convection and its associated hazards that have not been assessed, particularly for Australian SCTs. Logistic regression and machine learning has proven useful in past studies for determining skillful relationships between a broad array of atmospheric variables and SCTs (e.g., Mercer et al., 2009; Shafer and Doswell, 2010; Brown and Dowdy, 2021b). In this study, logistic regression and machine learning are applied to determine key attributes of the individual SCT hazards across Australia. The models used are compared amongst each other and to the discriminant relationship in Allen et al. (2011) for predictive ability. Applicability of statistical modelling for better understanding the SCT hazard climatology is assessed. Additionally, any trends within the reanalysis data set are identified.

8.3 Data and Methodology

8.3.1 Data

The Bureau of Meteorology's STA (Bureau of Meteorology, 2012) has been used as the main base of reports of SCTs around Australia. There are 13412 raw reports of hail, wind and tornadoes (Table 2.2), but reports were filtered using the comments and information on intensity of the event observed to include only events that satisfied the severe thresholds, and to ensure the report was likely convective of nature. Within the tornado reports, there were 60 reports of funnel clouds without a confirmed tornado touching down. However, many tornado events can go unnoticed due to their short durations and frequent occurrence in regions with low population density (Griffiths et al., 1993; Allen and Karoly, 2014). As a result, reports of only funnel clouds have been kept as they still indicate a large-scale environment that might support tornado development, and could have been associated with a tornado. Additionally, reports were only considered from 1990 through to 2018, to coincide with the length of the Bureau of Meteorology high-resolution Regional Reanalysis for Australia (BARRA; Su et al., 2019) data set that was used to characterise the near-storm environment. In total, 8321 reports of SCTs from the STA database were used (Table 2.2).

In order to develop a classification model of SCT hazard environments, additional data is required on likely non-severe thunderstorm days. Lightning data was supplied by Vaisala Inc. using their GLD360 lightning network (Said, 2017). This network covers the entire globe and was upgraded in 2016, with a detection efficiency of cloud-to-ground lightning over Australia of approximately 70% (Ron Holle & Ryan Said, personal communication). Due to licensing, this data was obtained for three regions in Australia, outlined in Figure 2.5.1, for the period 2014–2018. This data was filtered to determine possible dates, times and locations of SCT occurrences, considering all of the SCT reports from the STA.

Additionally, Maximum Estimated Size of Hail (MESH; Witt et al., 1998) data from the Australian Unified Radar Archive (Soderholm et al., 2019) was obtained over the 2014–2018 period using the modified 75th percentile fit in Eq. 15 of Murillo and Homeyer (2019), for radars within the three lightning detection regions Figure 2.5.1. This data was used in unison with

the lightning data to determine if a storm was likely severe (MESH $\geq$ 2 cm) or non-severe. From this, a list of severe, and likely non-severe storm reports was obtained. Finally, data was obtained for random dates between 2014 and 2018, and checked against the lightning data to obtain reports where no thunderstorms were recorded. This resulted in a data set of non-severe and no storm locations for reanalysis data to be obtained. There are some limiting assumptions from use of the above data, including:

1. If MESH $<$ 2 cm and there has been no STA report, then a recorded storm is not severe, even though a tornado or severe wind event could have occurred. Without availability of Doppler radar data and nearby automatic weather station sites, this cannot be accurately assessed, so it was assumed such storms are not severe.

2. The resulting non-severe storm data from the regions in Figure 2.5.1 is representative of non-severe storms across Australia, and potentially representative of no storm environments. Lightning data was selected from regions covering a relatively broad range of climates in Australia, while also containing a high number of radar sites, to determine from those observations which storms were likely severe or non-severe.

3. If MESH $\geq$ 2 cm, then a hail storm has definitely occurred. Although MESH has a relatively high degree of accuracy in categorising hail storms, it is not 100% accurate as it is a statistical fit on hail data.

4. Although some SCTs can produce all three hazards, it was not considered relevant to separate the data further into subgroups e.g., hail-only, hail and wind, hail and tornado, etc.

Despite these assumptions, the available data can still be valuable in assessing the SCT hazard climatology, and determining the applicability of linear and non-linear statistical modelling techniques for future study using data with fewer limitations.

In order to understand the environment surrounding an SCT, rawinsondes are often used to provide a vertical profile of the state of the atmosphere in which a storm develops. These profiles provide measurements of temperature, dew point temperature, height, pressure, and wind speed and direction at different levels of the atmosphere. However, there much debate on what is representative of the near-storm environment, as rawinsondes are often recorded only a few times a day and can be more than a hundred kilometres from the storm (Rasmussen and Blanchard, 1998; Brooks et al., 1994a; Thompson et al., 2003). This is a significant problem in Australia, with rawinsondes released only at some airports around the country and generally recorded at 0000 UTC, with few stations obtaining data at 1200 UTC. Given the diurnal peak of thunderstorm occurrence is mid-afternoon (local time), and the sites are often hundreds of kilometres away, the rawinsonde data is rarely representative of the environment in which the storm developed.

In order to better represent SCT environments, pseudo-proximity soundings are often used, developed from atmospheric reanalysis data sets (Lee, 2002). These data sets use earth system models and available observations (including automatic weather station, weather radar and satellite data) to constrain the model solutions, in order to realistically represent the atmospheric state at the time and location of the available observations. This data fills in the missing information between observations, both spatially and temporally, to form a consistently gridded data set of the historical atmosphere. From these data sets, the vertical profiles can be created that represent the atmospheric state, similarly to a rawinsonde. In this study

the BARRA data set is used to create pseudo-proximity soundings that characterise the near-storm environment. BARRA provides hourly data on an Australian-wide domain with a 12 km horizontal resolution from 1990-February 2019. There are 1.5 km horizontal resolution sub-domains over Sydney, Hobart, Adelaide and Perth, so the distance of the report from the nearest grid point was not different for reports outside these sub-domains, compared to those inside. Additionally, the data used in this study ends in 2018, in order for 29 complete years to be available.

8.3.2 Methodology

A large number of potential attributes were considered for selection of the important atmospheric indicators for each individual SCT hazard, among those indicators commonly considered in SCT forecasting. These include the CAPE, CIN, LCL, LFC, EL, lapse rate (LapR), and LI. As CAPE is sensitive to the parcel considered (Doswell and Rasmussen, 1994), three different parcels were considered in the formulation of the preceding parameters – the surface-based (SB), most-unstable (MU), and mixed-layer (ML) parcel over 100 mb. Additional thermodynamic indices considered include the K-index (George, 2014), TT (Miller, 1972), precipitable water (PW), and freezing level (FrzLevPres, the 0° height, in mb). Kinematic indices included 0–1 km, 0–3 km, and 0–6 km bulk shear (named SFCXkmSHEAR, where X represents the top of the layer it was calculated over, in km); and SRH from 0–1 km, 0–3 km, and over the effective inflow layer (given the labels SRH1km, SRH3km, and EffSRH, respectively). Indices that combine thermodynamic and kinematic properties included the severe weather threat index (SWEAT; Miller, 1972), supercell composite parameter (SCP; Thompson et al., 2002), significant tornado parameter (STP; Thompson et al., 2002, 2003), significant hail parameter (SHIP; Storm Prediction Center, 2014), severe hazards in environments with reduced buoyancy index (SHERB; Sherburn and Parker, 2014), and the energy helicity index (EHI; Thompson et al., 2003) over 0–1 km, 0–3 km, and the effective layer for each parcel type (given the label YEHIZkm, where Y is the parcel type and Z is the layer considered, similar to SRH). Finally, the following two-way interactions were considered: CAPE*SFC1kmSHEAR, CAPE*SFC3kmSHEAR, CAPE*SFC6kmSHEAR, CAPE*FrzLevPres, and CAPE*LCL; for each parcel type.

In this study, a classification problem is considered rather than a regression problem, to determine which of the above attributes are most important for each individual hazard. The statistical models considered are logistic regression (LogR; Hastie et al., 2009), support vector machines (SVM; Vapnik, 1995) with both polynomial (Poly) and radial basis function (RBF) kernels, and random forests (RF; Breiman, 2001; Hastie et al., 2009). Due to the data being unbalanced between severe and non-severe reports, a random sample is taken without replacement from the class that contains more data (either the severe or non-severe reports), in order for it to match the size of the other class. This resulted in a total of 2602 hail reports, 420 tornado reports, and 2163 wind reports for each of the severe and non-severe classes.

In total, there were 62 potential attributes considered in this study for each hazard. The data was split into 80% training data, with the remaining 20% used for testing. In order to select attributes that generalise well to the data, ten-fold cross validation (CV) was applied to the training data set using forward and backward selection through the space of potential attributes. This was first applied in order to determine whether the SB-, MU- or ML-parcel should be used. Indices based on the parcel that was selected most frequently among the folds were kept for further consideration, along with the indices that do not have a formulation that varies with parcel type. The 10-fold CV approach is then applied again to determine the attributes that generalise well to the data across the models.

Once these attributes are determined, a grid search is applied as in Hsu et al. (2003) to determine the free parameters of each model. For LogR, this is the cost parameter (C); for RBF and Poly SVM kernels there is the C parameter and gamma value (G), and the exponent (E) needs to be selected for the Poly SVM; for RFs, the number of trees is the free parameter. For the final model fitting step, the best performing model that retains between 2 and 5 attributes is selected for each model type and fitting procedure with values of C varying from $2^{0.5}$ to 2^3 with powers varying by 0.5; values of G ranging from 0 to 1, varying by 0.25; and values of E being either 1, 2 or 3. For the RFs, the number of trees in each forest was tuned with values ranging from 50 to 1000, varying by 50. The final models for each statistical method are then applied to the test data, where contingency statistics are calculated, and receiver operating characteristic (ROC) curves and precision-recall (PR) curves are created to assess the performance of each model on the testing data. The details of the contingency statistics, ROC and PR curves will be described in further detail below. For each SCT hazard, the model that performs best on the testing data is then applied to the BARRA data set to diagnose the frequency of environments supportive of that SCT hazard. To do this, pseudo-proximity soundings are created at every second grid point (in the horizontal) and every third hour in order to calculate the necessary environmental predictors for that classification model. It was decided to assess environmental frequency for every second grid point and third hour to reduce the amount of computation required, while still being frequent enough to pick up the majority of environments.

Following calculation of the spatial frequency of environments, trends in the BARRA data set were then assessed using similar methods to Chapters 3–6, but extended to a geospatial data set. The data at each grid point was broken into two periods from 1990–2003 and 2004–2018. Bootstrap resampling was applied with 5000 resamples to calculate the mean frequency of environments supportive of each SCT hazard at each grid point, over both 14-year periods. The differences in the mean at each grid point are then calculated from which a visualisation of the change in frequency of supportive SCT hazard environments can be assessed across Australia. Permutation testing was then applied, with 5000 resamples, to test for significance in any trends within the mean of each SCT hazard.

8.3.3 Classification Model Selection

In order to assess classification model performance, and determine the best performing model, contingency statistics, ROC and PR curves are considered. Although a "best performing model" is selected for all three SCT hazards in this Chapter, the classification of such a model can be highly subjective. It is dependent on how misclassification errors are valued by the model developer. In some circumstances, misclassification errors such as incorrect "yes" forecasts (i.e., false alarms) can prove costly. For this Chapter, the trade-off between increasing the number of correct "yes" forecasts and false alarms is considered, with higher skill scores (defined below) valued over reducing the number of false alarms. However, placing different levels of importance on these measures can result in a different decision of what constitutes the "best performing model".

Here, the possible outcomes for each model are that an SCT hazard occurs, or does not occur, and is described in Table 8.1. One of the simplest measures of classification accuracy is the *proportion correct* (or *hit rate*), defined in Wilks (2011) as:

$$\text{Proportion Correct} = \frac{a+d}{n}. \tag{8.3.1}$$

A proportion correct of 1 is most desirable, as this means all forecasts have been correctly made. However, the proportion correct means both correct SCT hazard and no SCT hazard forecasts have equal weighting, which is not desirable when the SCT hazard forecast is a rare occurrence, and provides no information on the relative errors. The *probability of detection* (POD) is another measure considering the ratio of correct SCT hazard forecasts to the total number of SCT hazard events:

$$\text{POD} = \frac{a}{a+c}. \tag{8.3.2}$$

Equivalently, the *probability of false detection* (PoFD) is the ratio of incorrect no SCT hazard forecasts to the total no SCT hazard events observed:

$$\text{PoFD} = \frac{b}{b+d}. \tag{8.3.3}$$

A high POD and low PoFD are desirable as this suggests a large number of SCT hazards have been correctly forecast, and few no SCT hazard events have been incorrectly forecast. An alternative statistic frequently used to assess incorrect forecasts is the *false alarm ratio* (FAR):

$$\text{FAR} = \frac{b}{a+b}, \tag{8.3.4}$$

where a number close to 0 is most desirable. The *critical success index* (CSI) weights the number of correct SCT hazard forecasts by the total number of correct SCT hazard forecasts and incorrect forecasts:

$$\text{CSI} = \frac{a}{a+b+c}. \tag{8.3.5}$$

The CSI is particularly useful when the number of SCT hazards are infrequent, with the ideal score of 1 indicating no incorrect forecasts have been made. There are also a large number of skill scores that have been developed to help summarise classification model performance. The *Heidke skill score* (HSS) provides a measure based on the proportion correct that assesses how well a model performs compared to obtaining the correct forecast by chance (Wilks, 2011), and is defined as:

$$\text{HSS} = \frac{2(ad-bc)}{(a+c)(c+d)+(a+b)(b+d)}. \tag{8.3.6}$$

A HSS close to 1 is desirable, whereas a value close to 0 indicates that many correct forecasts might be due to chance. The *Gilbert skill score* (GSS) extends on the CSI, using it as the measure of a random forecast. Given $a_{ref} = (a+b)(a+c)/n$, the GSS is defined by:

$$\text{GSS} = \frac{a-a_{ref}}{a-a_{ref}+b+c}, \tag{8.3.7}$$

and often is used when the number of correct forecasts are rare. Although, if they are extremely rare, then the GSS can be misleading and approach 0 even under skillful forecasts (Wilks, 2011). The *odds ratio skill score* uses the odds ratio to measure forecast skill, and is given by:

$$\text{Odds ratio skill score} = \frac{ad-bc}{ad+bc}, \tag{8.3.8}$$

where essentially random forecasts obtain a score of 0. A perfect forecast will exhibit a skill score of 1; however, this can also be obtained for imperfect forecasts where either b or c are zero. Finally, the *bias ratio* is a measure that compares the average forecast to observation, provided by:

$$\text{Bias ratio} = \frac{a+b}{a+c}. \tag{8.3.9}$$

Table 8.1: Example 2×2 contingency table outlining the possible outcomes from a simple yes/no classification model. Here, a represents a true positive, b a false positive, c a false negative and d a true negative.

		Observed	
		SCT Hazard	No SCT Hazard
Forecast	SCT Hazard	a	b
	No SCT Hazard	c	d

An unbiased forecast achieves a bias ratio of one, while values above one indicate the model is overforecasting and values below one indicate it is underforecasting, relative to observations. There are additional, graphical measures that demonstrate the performance of classification models. The *receiver operating characteristic* (ROC) curve, assesses the false positive rate against the true positive rate of a model as the probability threshold determining a classification is varied (Mason, 1982; Wilks, 2011). Changing the decision threshold results in different contingency tables from which these rates can be calculated and plotted. An ideal ROC curve maximises the area under the curve, which is equivalent to the Mann-Whitney U-statistic as shown by Mason and Graham (2002), and approaches the top-left of the diagram. On the other hand, a ROC curve following close to the one-to-one line, indicates a model that is not much better than a random forecast.

The *precision-recall* (PR) curve is another graphical assessment of model performance as the decision threshold is varied. PR curves are plots of *precision* (given by $a/(a+b)$) against *recall* (the POD) and consider the relative trade-off between these two quantities as a decision threshold is varied. Better performing classification models will have PR curves approaching the top right corner of the diagram, indicating both high precision and high POD, and therefore a low number of incorrect classifications. The *average precision* (AP) summarises the PR curve by calculating a weighted average of the precision at a particular decision threshold, with the change in recall from the previous decision threshold used as the weight. Higher values of AP indicate a better performing model.

8.4 Results

8.4.1 Comparison of Pseudo-proximity Soundings to Observed Soundings

As reanalysis data is derived from an earth system model, it is important to compare the reanalysis-derived pseudo-proximity soundings against observed rawinsonde vertical profiles. In order to compare these soundings, MLCAPE, MUCAPE, SBCAPE, SFC1kmSHEAR, SFC3kmSHEAR and SFC6kmSHEAR are compared for all available rawinsondes at 0000 UTC around Australia between 1990–2018 (inclusive). Details on the sites considered are provided in Table 8.2. Only rawinsondes recorded at 0000 UTC are considered because this is generally closer in time to the diurnal peak of convection, and the time most commonly sampled across sites (Allen et al., 2011; Soderholm et al., 2017a; Warren et al., 2020; Dowdy et al., 2020). Rawinsonde data was obtained from the University of Wyoming upper-air sounding archive (`http://weather.uwyo.edu/upperair/sounding.html`) for vertical profiles with MUCAPE larger than 100 J kg^{-1}. As a result, there were 22855 vertical profiles compared against the nearest pseudo-proximity sounding in latitude and longitude in the BARRA data set.

Figure 8.4.1 displays scatterplots comparing bulk shear and CAPE for the pseudo-proximity

Table 8.2: Information for each rawinsonde site in Australia that was used to compare against pseudo-proximity soundings derived from the BARRA reanalysis data set.

Location	Latitude (°S)	Longitude (°E)	State	Coastal Proximity
Adelaide	34.95	138.53	SA	Coastal
Albany	34.93	117.80	WA	Coastal
Alice Springs	23.80	133.88	NT	Inland
Brisbane	27.38	153.13	QLD	Coastal
Broome	17.95	122.23	WA	Coastal
Charleville	26.41	146.26	QLD	Inland
Darwin	12.42	130.89	NT	Coastal
Geraldton	28.80	114.70	WA	Coastal
Hobart	42.83	147.50	TAS	Coastal
Learmonth	22.24	114.10	WA	Coastal
Meekatharra	26.61	118.55	WA	Inland
Melbourne	37.66	144.85	VIC	Coastal
Moree	29.48	149.83	NSW	Inland
Mt Isa	20.68	139.48	QLD	Inland
Perth	31.93	115.96	WA	Coastal
Sydney	33.95	151.17	NSW	Coastal
Williamtown	32.8	151.83	NSW	Coastal
Woomera	31.15	136.81	SA	Inland

soundings in the BARRA data set (along the y-axis) against the observed soundings (x-axis). Ideally, points should lie along the thin black line, which represents a one-to-one correspondence between the two data sets. For bulk shear, there is a relatively narrow spread about the line, but the data points tend to lie slightly below the line indicating bulk shear is often underestimated. This is particularly the case for 0–1 km bulk shear, with values becoming more evenly spread as the depth that bulk shear is calculated over increases. Reanalysis data sets often have a bias in estimating wind speeds, and there is a known bias of underestimating winds in high wind conditions and overestimating them in low wind conditions for BARRA (Su et al., 2019). The underestimation of wind speeds is possibly due to unstable atmospheric conditions, which has been found in other studies of reanalysis data sets (Rose and Apt, 2016). As the depth of the layer that bulk shear is calculated over increases, the underestimation of wind speeds at lower model levels are likely being diffused by the more accurate estimation of wind speeds higher in the atmosphere, causing the data to fall closer to the one-to-one line.

Meanwhile, there is generally a large spread of CAPE values, with BARRA consistently overestimating. This is most evident for SBCAPE and MUCAPE, with MLCAPE appearing closer to the one-to-one line, although still lying above it. The improvement noticed in MLCAPE is likely due to it being calculated from the mean of parcels within the 100 mb nearest to the surface, rather than relying on just one parcel value. Differing values of CAPE could potentially be due to the reanalysis data misrepresenting some sub-grid scale processes, or the unstable air masses entering the grid point at a different time than in the physical world.

Figures 8.4.2 & 8.4.3 display these same scatterplots for coastal and inland rawinsonde sites, respectively. It is evident that bulk wind shear is underestimated regardless of proximity to the coast, although 0–1 km bulk shear appears more frequently underestimated inland. On the other hand, the overestimation of CAPE appears to be more significant for coastal sites. While inland sites tend to more accurately represent CAPE, there is still a notable spread rather

than being located close to the one-to-one line. This misrepresentation at coastal sites is likely due to the more dynamic and frequently changing environment along the coast as compared to inland.

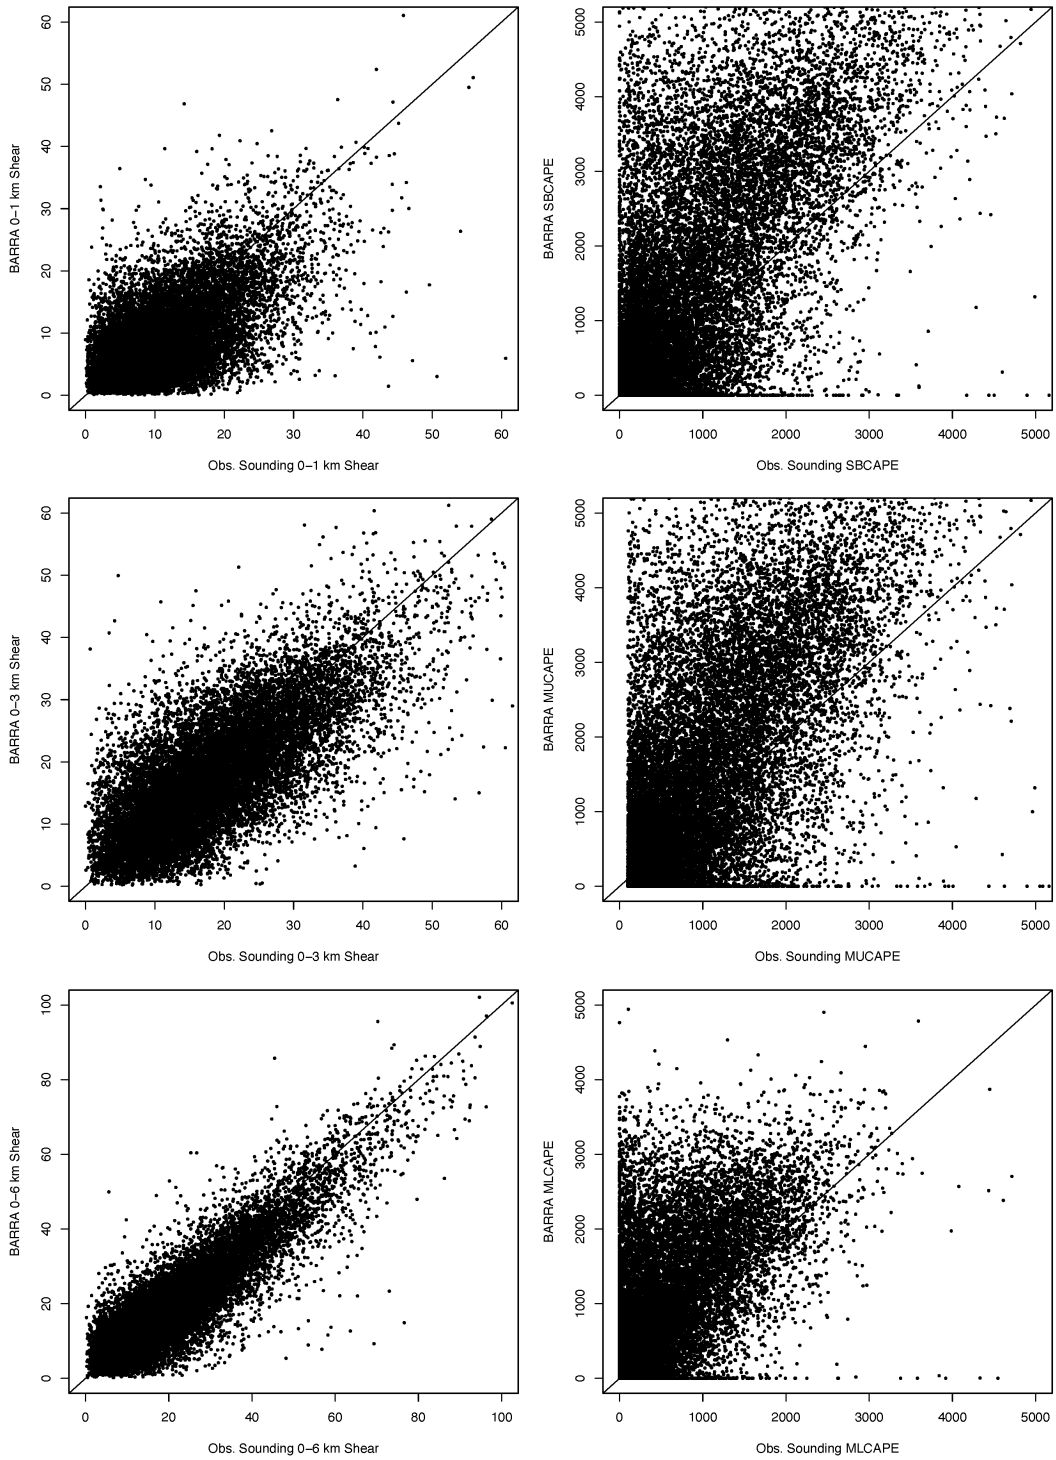

Figure 8.4.1: Scatterplots of bulk shear computed over different levels (left column) and CAPE computed for different parcel types (right column), comparing vertical profiles from the BARRA data set against observed rawinsonde profiles. The thin black line represents the line $y = x$, where points ideally should lie.

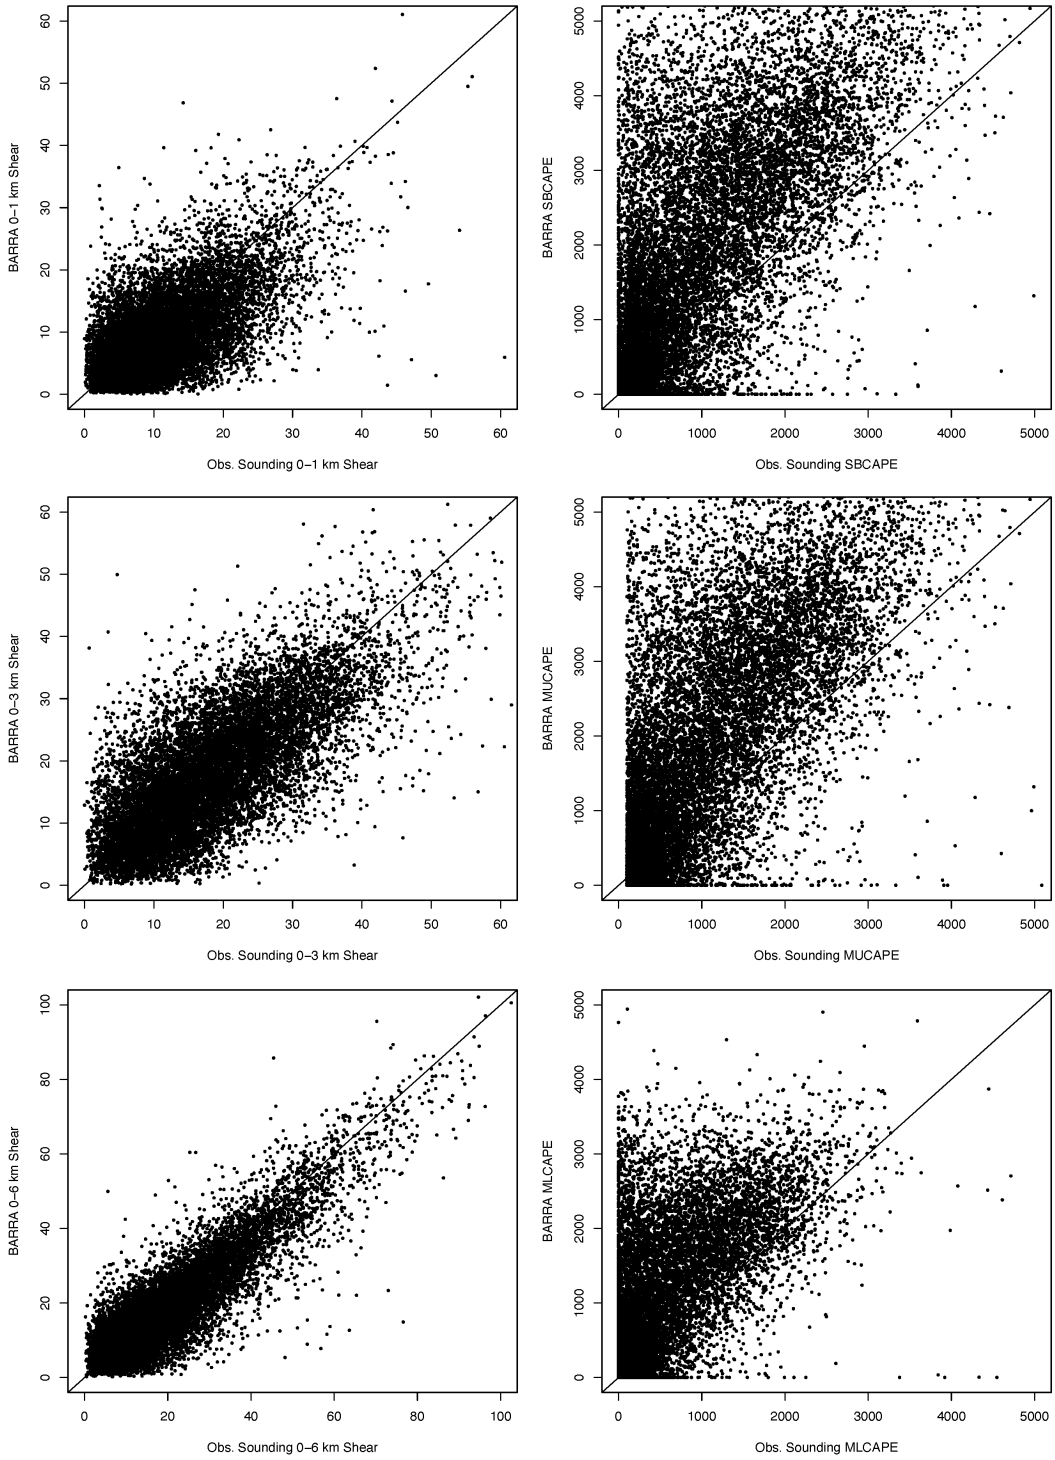

Figure 8.4.2: As in Figure 8.4.1, but for the coastal rawinsonde sites only.

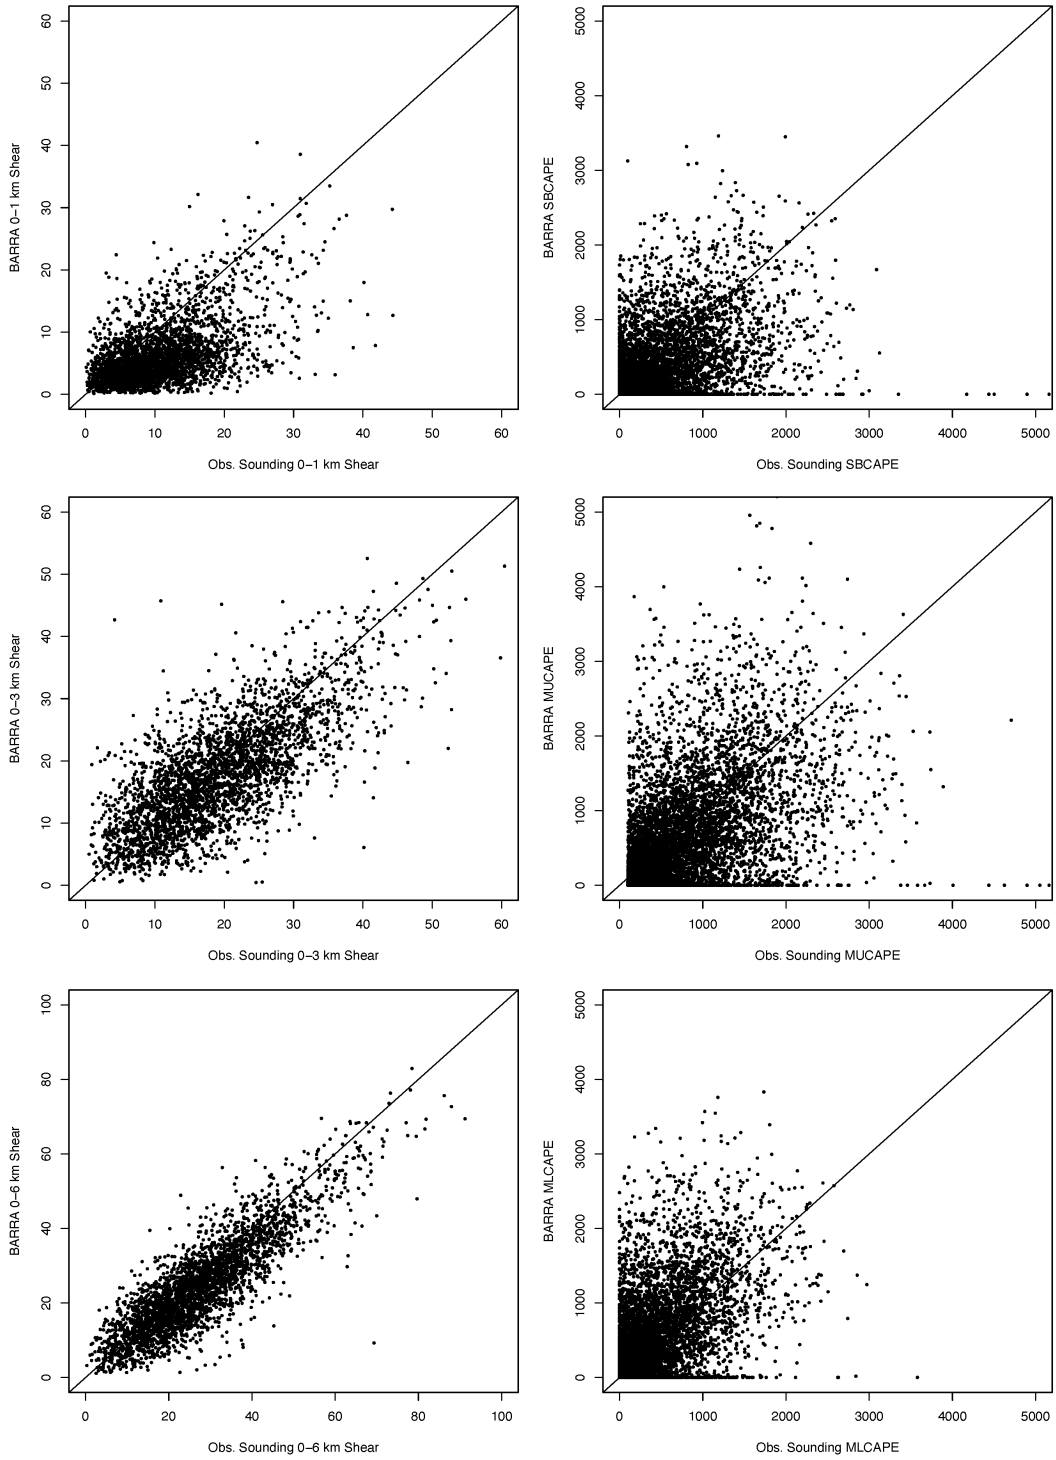

Figure 8.4.3: As in Figure 8.4.1, but for the inland rawinsonde sites only.

8.4.2 STA Climatology

Prior to fitting statistical models of SCT hazards and developing a climatology from the BARRA reanalysis data, analysis of the spatial distribution of each hazard from the STA report database was performed. Figure 8.4.4 displays the binned density of reports for each SCT hazard within the STA data across Australia. This has been performed on the filtered STA data from 1990–2018 (Table 2.2). There are some reports located in the ocean, likely due to incorrect data entry in the STA. As in previous work, reports appear centred around major population regions, especially the densely populated east coast (Allen et al., 2011). The frequency of hail reports

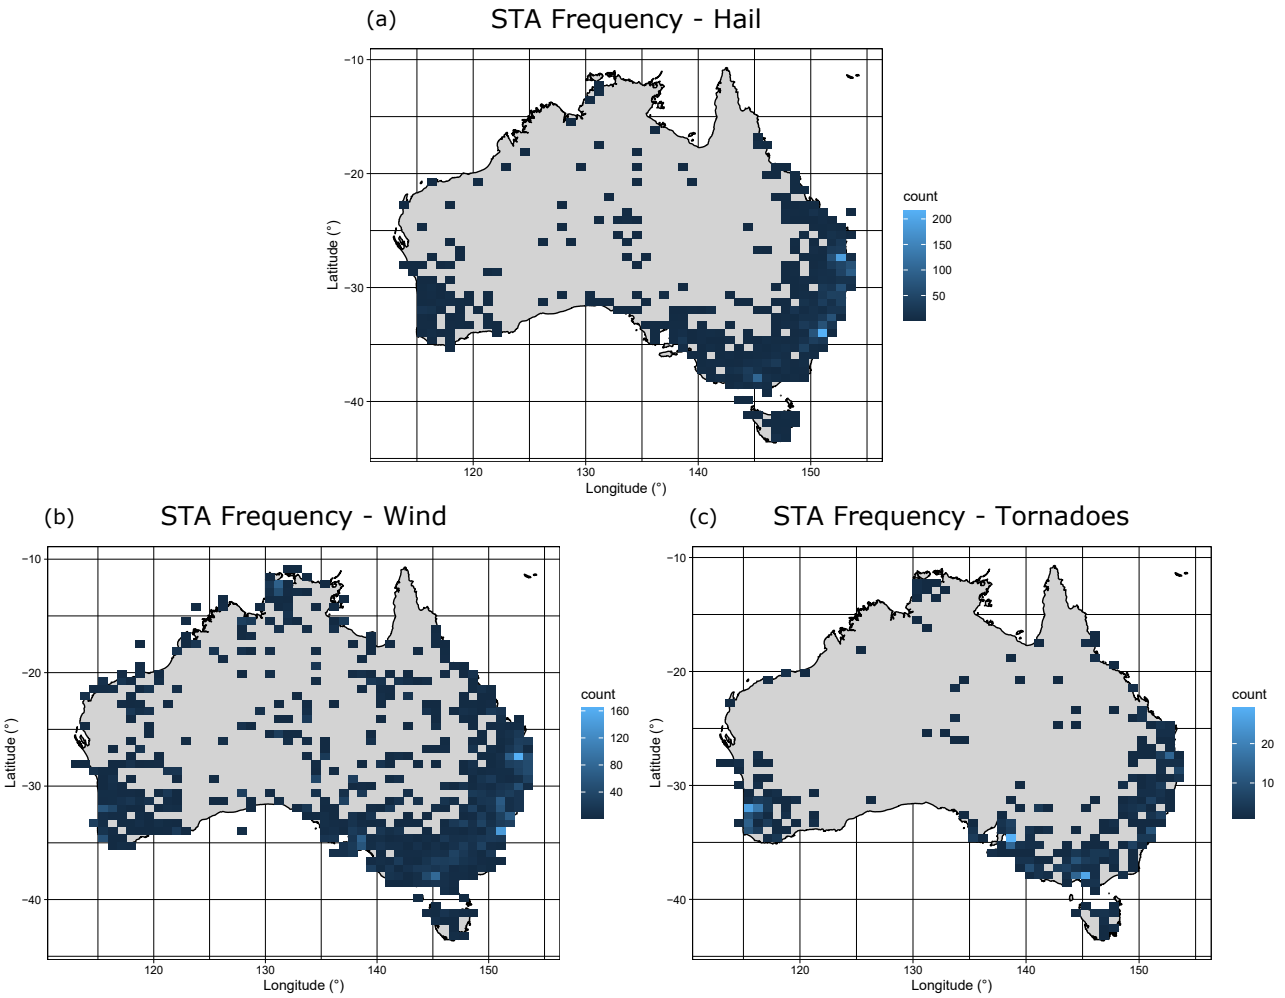

Figure 8.4.4: Binned density of total reports in the STA from 1990–2018 after manual filtering of reports, for (a) hail, (b) wind, and (c) tornadoes.

is greatest along the east coast, with the highest frequency around the Sydney metropolitan region, and southeast Queensland (Figure 8.4.4a). There also is a high frequency around Melbourne, while there are lower frequencies in Hobart, the area surrounding Perth and Adelaide. The Northern Territory has a low frequency of hail reports, partially due to the lower population density there, although the number of reports around the state capital of Darwin also is low. This lower frequency of reports likely is due to Darwin being in the tropics, where the prevailing convective mode likely is not as supportive of hailstone development and growth. Additionally, the average freezing level would be higher, so there would be greater melting of any hailstones that develop as they fall to the surface.

Wind reports appear to be least affected by the sparse population density away from the coast, with reports widely distributed across Australia compared to the other hazard types (Figure 8.4.4b), although the frequency away from the coast still is much reduced. As expected, there is a high frequency of reports along the east coast, with a peak in frequency over the Sydney Metropolitan region and southeast Queensland, similar to the distribution of hail. There also is a higher frequency of reports surrounding each state capital city, apart from Hobart. Meanwhile, tornado reports are most affected by the reduced population density around central and northwest Australia (Figure 8.4.4c). Many tornadic storms can be missed given tornadoes occur over very small spatial and temporal scales. As a result, if they have developed in an area with low population density, they can easily end up unreported. Although there are

less reports, the distribution displays important differences compared to the other two SCT hazards. Namely, the peak in frequency of tornado reports occurs in Melbourne, Adelaide, and around southwest WA, particularly around Perth. This is likely a result of the high frequency of cool-season tornadoes that affect southern Australia, with almost 50% of the total number of Australian tornadoes estimated to occur during the cool-season, and the majority of these cool-season tornadoes affecting southwest Australia, and SA (Hanstrum et al., 2002). There is a relatively high frequency of reports also around the Sydney region and southeast Queensland. Interestingly, there is a higher number of reports around Darwin compared to the frequency of hail reports, potentially due to tornadoes spawned by tropical cyclones (Figures 8.4.4a & c).

Figures 8.4.5–8.4.7 display the frequency of STA reports during each month, for hail, wind, and tornadoes, respectively. The frequency of hail reports is lowest during the cooler months of the year from May–August (Figure 8.4.5). Hail reports begin to increase in frequency over the east coast during August and September. Following this is an increase in frequency over central and WA during October and November, while there is a peak in frequency over southeast Queensland during November. The frequency of reports over the central NSW coast (including Sydney) appears to peak during December. Similarly, the greatest frequency of reports in the region surrounding Adelaide occurs during December. The frequency of hail reports over eastern Australia decreases from January through to May. Meanwhile frequency over WA remains relatively stable until May, where reporting appears focused over the southwest corner during the cooler months.

Over much of Australia, the number of severe wind reports begins to increase in frequency from August (Figure 8.4.6). The greatest frequency of reports over eastern Australia occurs during December and January, while the peak in frequency over northern Australia occurs in November. There appears to be a peak in reports over Victoria during January. Across Australia, there is a decline in severe wind reports from February until May. However, there still appears to be a number of reports during the cooler months in southern Australia, particularly around Adelaide, southwest WA, parts of Victoria and Tasmania.

The bi-modal peak in Australian tornado frequency is apparent in the monthly spatial distribution plots (Figure 8.4.7). The greatest frequency in reports is apparent over eastern Australia over October–January, with a reduction in reports outside these months. However, there still appears to be a number of reports during the cooler months of the year, possibly due to tornadoes associated with east coast low pressure systems (Hanstrum et al., 2002). Meanwhile, southern Australia experiences a peak in frequency during the cooler months of the year, particularly during June for southwest Australia and the Adelaide region. During the cool-season, there are a number of strong cold frontal systems that have been shown to produce conditions that are favourable for tornado development (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009). However, there are still some reports over southern Australia during spring (Sep–Nov) and summer (Dec–Feb).

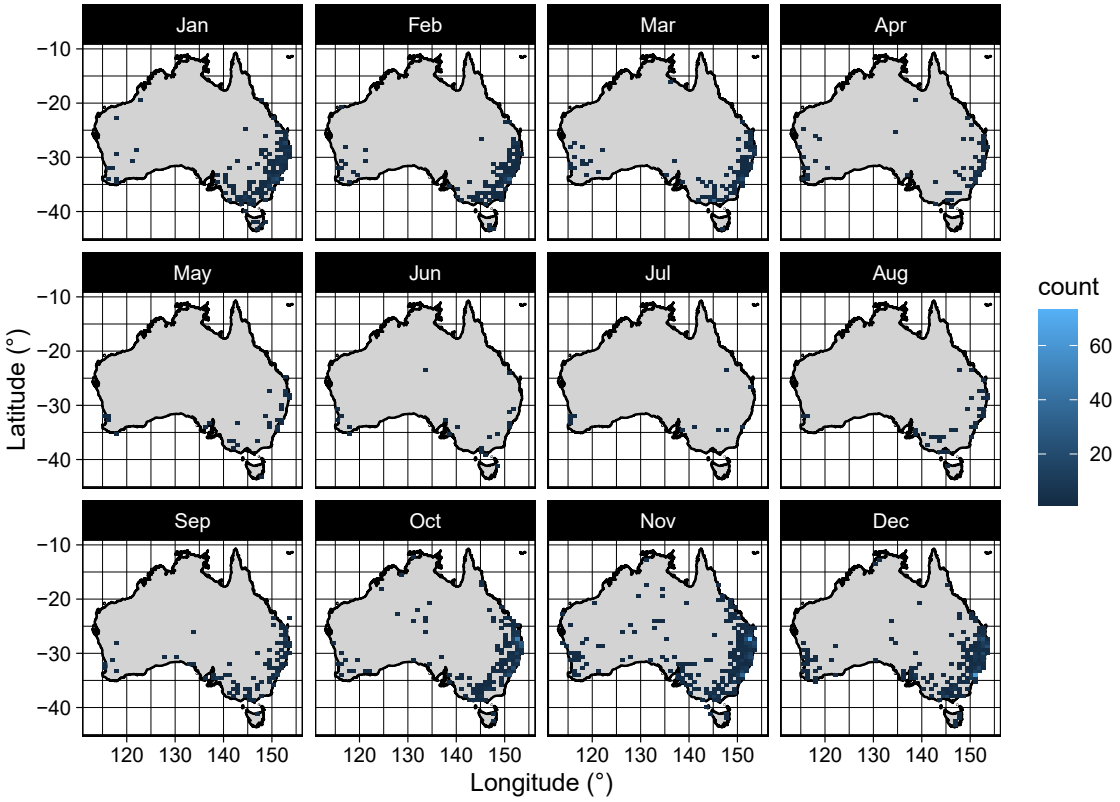

Figure 8.4.5: Binned density of all hail reports in the STA for each month from 1990–2018.

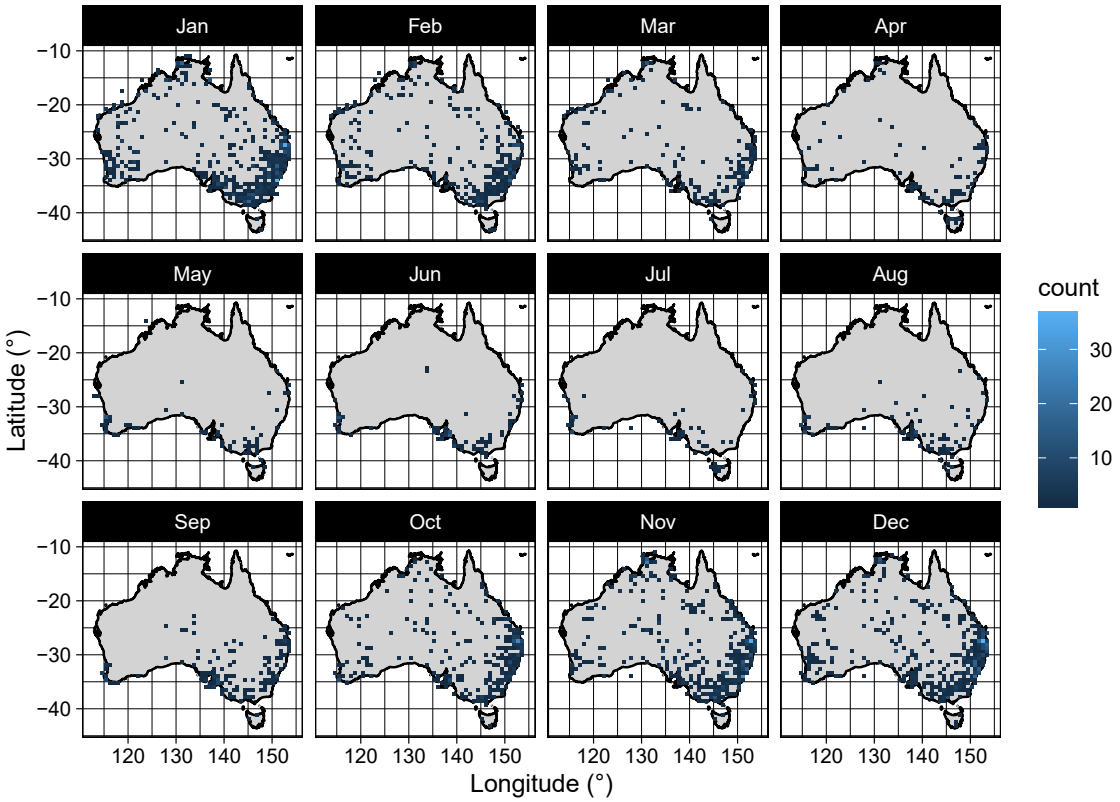

Figure 8.4.6: As in Figure 8.4.5 but for wind reports.

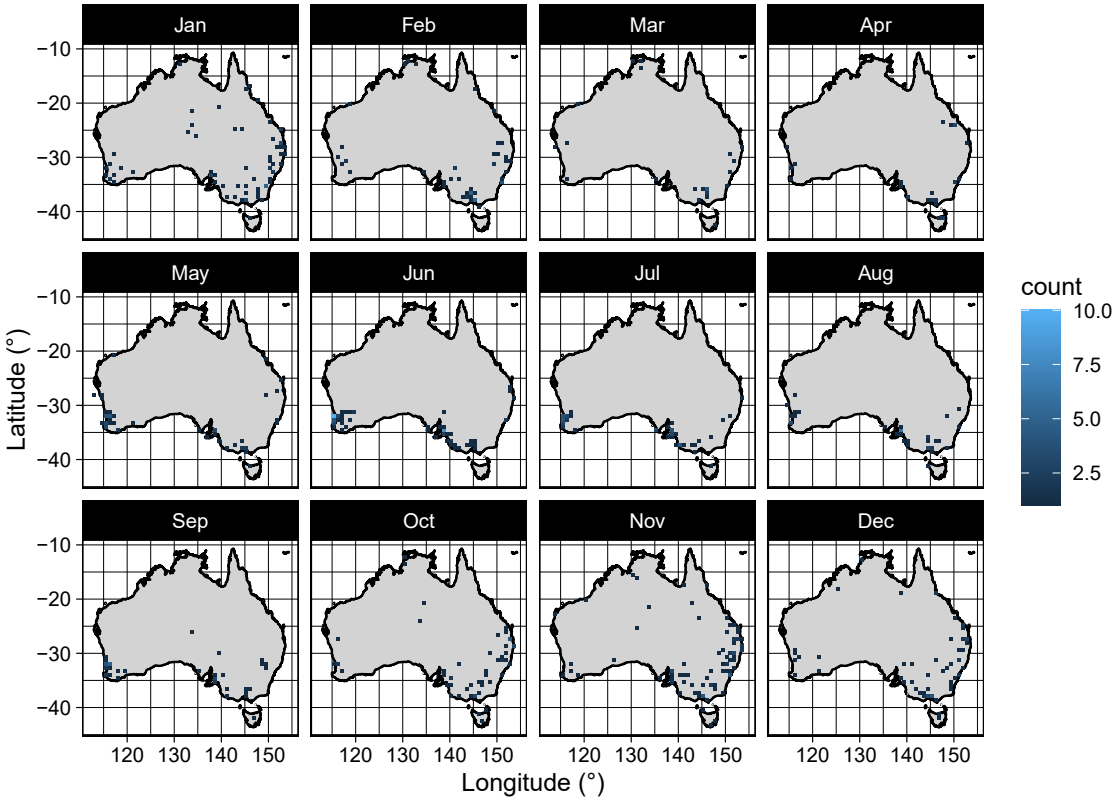

Figure 8.4.7: As in Figure 8.4.5 but for tornado reports.

8.4.3 Model Fitting

Table 8.3 lists the percentage of folds that each parcel formulation appeared in for those parameters that are calculated by considering an air parcel, across all model selection methods for hail, tornadic and wind SCTs. There is a clear consensus across the statistical methods and hazard types over selecting the ML parcel with highest priority, followed by the SB parcel, then the MU parcel. This might be in part explained by the scatterplots in Figures 8.4.1–8.4.3, where the ML parcel is most representative of observed instability in the atmosphere, and likely better discriminates SCT hazard environments in BARRA. A similar behaviour appears in Brown and Dowdy (2021a), with the four LogR models developed to predict severe convective wind environments only including indices with the ML-parcel definition. As a result, indices calculated using the ML parcel are retained for 10-fold CV to be applied to better determine which indices generalise best across each model and selection type. This reduces the number of possible attributes considered from 62 down to 30.

Table 8.3: Percentage of folds that indices calculated using the surface-based (SB), most-unstable (MU) and mixed-layer (ML) parcel formulations appeared in.

Parcel Type	Hail	Tornadic	Wind
SB Parcel	23.13	22.03	20.39
MU Parcel	18.67	16.72	15.39
ML Parcel	28.66	26.00	24.25

Tables 8.4–8.6 display the percentage of folds selecting each index for hail, tornadic and wind SCTs. At this point in the selection process, parameters were not tuned, with C, G and E parameters taking the default values 1, 1/(number of features × training variance), and 3, respectively, and the number of trees in the RF model set to 100. Indices that appeared in at

Table 8.4: Percentage of folds selecting attributes for hail SCTs using LogR, SVM with the radial basis (RBF) and polynomial (Poly) kernel functions, and random forests (RF) using either forward (F) or backward (B) selection techniques to search through the space of attributes.

Index	LogR F	SVM RBF F	SVM Poly F	RF F	LogR B	SVM RBF B	SVM Poly B	RF B
MLCAPE	60	80	50	100	40	10	10	70
MLCIN	60	30	70	100	70	80	90	50
MLLI	70	80	50	0	30	30	10	40
MLLapR	30	10	100	100	20	30	0	70
MLLCL	30	0	100	0	50	80	100	40
MLLFC	40	20	90	100	80	90	0	90
MLEL	80	60	90	0	100	100	100	50
K-Index	30	60	60	0	60	60	80	20
TT	40	30	90	0	90	30	70	0
PW	20	0	60	0	10	60	70	70
SWEAT	0	10	30	0	30	70	60	80
SFC1kmSHEAR	20	30	20	0	60	70	50	50
SFC3kmSHEAR	30	0	10	0	60	30	60	20
SFC6kmSHEAR	30	80	10	0	30	70	60	70
SRH1km	80	30	10	0	50	30	40	80
SRH3km	40	20	10	0	40	40	50	60
EffSRH	70	30	50	80	90	90	90	40
SCP	80	90	20	80	20	20	10	50
STP	70	70	60	90	100	60	100	30
SHIP	40	90	30	70	20	30	10	70
SHERB	10	0	40	0	40	70	30	80
FrzLevPres	100	100	100	0	100	100	100	90
MLEHI1km	60	50	10	100	30	70	10	10
MLEHI3km	70	50	10	100	30	10	80	50
MLEHIEff	100	100	20	100	10	30	70	40
MLCAPE*SFC1kmSHEAR	60	70	20	90	10	60	40	20
MLCAPE*SFC3kmSHEAR	50	80	50	100	50	10	10	20
MLCAPE*SFC6kmSHEAR	30	70	100	90	50	10	40	40
MLCAPE*FrzLevPres	50	100	80	100	70	20	20	50
MLCAPE*MLLCL	50	60	60	100	60	40	40	50

least 5 of 10 folds for 50% of the modelling techniques were retained for further selection using the LogR, SVM and RF methods.

Table 8.5: As in Table 8.4 but for tornadic SCTs.

Index	LogR F	SVM RBF F	SVM Poly F	RF F	LogR B	SVM RBF B	SVM Poly B	RF B
MLCAPE	60	50	20	20	0	40	10	50
MLCIN	40	30	90	100	80	60	90	90
MLLI	50	60	80	40	50	0	30	50
MLLapR	90	90	0	80	30	70	0	90
MLLCL	30	50	20	80	30	70	30	50
MLLFC	30	30	10	30	40	40	0	30
MLEL	100	40	100	100	100	90	100	20
K-Index	10	50	70	70	40	50	70	40
TT	20	60	60	60	20	0	40	80
PW	20	20	80	30	20	10	20	30
SWEAT	100	70	50	50	70	90	70	50
SFC1kmSHEAR	50	80	100	80	90	90	90	60
SFC3kmSHEAR	40	60	80	50	90	50	50	60
SFC6kmSHEAR	70	20	90	50	90	90	70	50
SRH1km	50	40	40	90	100	50	70	80
SRH3km	60	10	20	50	30	0	30	70
EffSRH	40	40	30	50	70	20	60	20
SCP	60	90	40	60	0	10	10	60
STP	30	40	10	30	20	10	50	60
SHIP	60	80	30	20	70	20	20	40
SHERB	20	80	60	40	80	70	70	20
FrzLevPres	100	50	100	90	90	90	100	80
MLEHI1km	60	70	30	40	50	40	50	30
MLEHI3km	60	60	60	50	80	50	60	30
MLEHIEff	70	60	70	30	0	70	40	50
MLCAPE*SFC1kmSHEAR	50	60	70	0	0	90	50	40
MLCAPE*SFC3kmSHEAR	30	10	30	10	40	50	20	40
MLCAPE*SFC6kmSHEAR	40	60	40	50	30	10	60	50
MLCAPE*FrzLevPres	20	0	10	40	20	90	70	30
MLCAPE*MLLCL	40	40	10	10	70	50	70	50

Table 8.6: As in Table 8.4 but for wind SCTs.

Index	LogR F	SVM RBF F	SVM Poly F	RF F	LogR B	SVM RBF B	SVM Poly B	RF B
MLCAPE	20	20	40	60	10	30	50	30
MLCIN	50	60	90	100	40	100	100	90
MLLI	100	80	90	30	50	30	70	80
MLLapR	30	30	0	30	100	40	0	20
MLLCL	10	100	100	100	100	100	100	100
MLLFC	0	30	0	100	90	40	0	80
MLEL	80	100	100	100	0	90	90	90
K-Index	60	60	100	40	0	30	90	70
TT	90	90	100	0	100	60	80	10
PW	40	70	70	30	90	90	90	70
SWEAT	10	90	90	70	100	100	90	90
SFC1kmSHEAR	90	80	80	30	30	90	100	70
SFC3kmSHEAR	90	30	10	30	40	30	0	30
SFC6kmSHEAR	0	30	40	90	70	50	50	40
SRH1km	100	10	50	100	100	70	30	70
SRH3km	70	40	10	30	70	70	30	20
EffSRH	60	20	30	40	40	60	60	30
SCP	20	20	10	100	40	20	10	40
STP	50	20	60	100	20	0	10	80
SHIP	90	0	30	30	20	20	10	10
SHERB	100	70	80	60	100	80	90	70
FrzLevPres	10	100	100	100	100	100	100	70
MLEHI1km	50	40	40	10	10	10	20	10
MLEHI3km	50	20	10	10	20	10	0	30
MLEHIEff	20	40	10	10	0	0	20	0
MLCAPE*SFC1kmSHEAR	60	30	30	20	50	10	20	50
MLCAPE*SFC3kmSHEAR	50	50	50	10	50	40	60	30
MLCAPE*SFC6kmSHEAR	60	50	60	30	50	60	40	50
MLCAPE*FrzLevPres	20	100	20	20	10	40	60	30
MLCAPE*MLLCL	20	20	0	20	0	30	30	40

Models were then trained for each method with parameter tuning using a grid search, and at most 5 attributes retained. The final model for each method was selected to have a high mean cross-validated score (defined by each method in Python's Scitkit Learn package, e.g., the RF uses mean accuracy - the fraction of correct labels) on the training data and minimal attributes. Contingency statistics were calculated on the test data set for each model developed and are presented in Tables 8.7, 8.9, and 8.11 for the hail, tornadic, and wind SCT models, respectively. Additionally both ROC and PR curves were created to assess model performance and are presented in Figures 8.4.8–8.4.13. The attributes and optimal tuning parameters used for each model are presented in Tables 8.8, 8.10, and 8.12 for the hail, tornadic, and wind SCT models, respectively.. The final model selected to diagnose the spatial frequency of environments supportive of each hazard in BARRA was the model that performed best across the contingency statistics, ROC and PR curves.

For hail SCTs, the model that performed best in contingency statistics was the SVM RBF developed using backward selection. This model has MLCIN, MLEL, SFC6kmSHEAR and FrzLevPres as attributes, with C = $2^{0.5}$ and G = 0.25. The model was among those with the lowest FAR (0.1966), and achieved the highest CSI (0.6906), HSS (0.6276), and GSS (0.4573; Table 8.7). The ROC curve for this model appears closest to the top left corner compared to other models, although the differences are rather minimal, even for the LogR models that both have slightly higher area under the ROC curve (Figure 8.4.8). The PR curve for the SVM RBF B model appears better than most models except the LogR models. Although the LogR models appear to perform marginally better than the SVM RBF B model on the ROC and PR curves, they have a higher FAR rate as a result of their higher POD. It is considered more desirable to have a lower FAR in order to reduce the number of incorrectly identified environments that are supportive of hail storms. Therefore, the SVM RBF B model was selected for analysis of the hail environment climatology.

Table 8.7: Contingency statistics on the testing data set for the best performing models on hail SCTs. The best performance among the models for each contingency statistic is highlighted in bold face.

Statistic	**LogR F**	**SVM RBF F**	**SVM Poly F**	**RF F**	**LogR B**	**SVM RBF B**	**SVM Poly B**	**RF B**
Scikit-Learn Accuracy	0.7901	0.7756	0.7685	0.7713	**0.7908**	0.7896	0.7860	0.7735
Proportion Correct	0.7994	0.8052	0.7639	0.7488	0.8052	**0.8138**	0.8052	0.7821
POD	0.8522	0.8081	**0.8868**	0.8234	0.8407	0.8311	0.8253	0.8177
FAR	0.2292	**0.1966**	0.2881	0.2629	0.2151	**0.1966**	0.2066	0.2366
PoFD	0.2536	**0.1977**	0.3589	0.2937	0.2303	0.2035	0.2150	0.2534
CSI	0.6799	0.6747	0.6525	0.6365	0.6833	**0.6906**	0.6793	0.6524
HSS	0.5988	0.6104	0.5278	0.5298	0.6104	**0.6276**	0.6104	0.5643
GSS	0.4274	0.4392	0.3586	0.3603	0.4392	**0.4573**	0.4392	0.3930
Bias Ratio	1.1056	**1.0058**	1.2457	1.1171	1.0710	1.0345	1.0403	1.0710
Odds Ratio Skill Score	0.8888	0.8894	0.8665	0.8363	0.8927	**0.9013**	0.8904	0.8593

Table 8.8: Features and tuning parameters selected for the best performing models presented in Table 8.7.

Model	Features	Parameters
LogR F	MLCAPE, MLEL, K-Index, FrzLevPres, MLCAPE*SFC3kmSHEAR	C = $2^{0.5}$
SVM RBF F	FrzLevPres, MLCAPE*FrzLevPres	C = $2^{2.5}$, G = 0.75
SVM Poly F	MLLCL, MLLFC, MLEL, PW, EffSRH	C = 4, E = 3, G = 0.25
RF F	MLLFC, MLEHI1km, MLCAPE*SFC1kmSHEAR, MLCAPE*SFC3kmSHEAR	ntree = 350
LogR B	MLEL, K-Index, EffSRH, FrzLevPres, MLCAPE*SFC6kmSHEAR	C = $2^{0.5}$
SVM RBF B	MLCIN, MLEL, SFC6kmSHEAR, FrzLevPres	C = $2^{0.5}$, G = 0.25
SVM Poly B	MLEL, SFC6kmSHEAR, EffSRH, FrzLevPres, MLCAPE*MLLCL	C = $2^{0.5}$, E = 3, G = 0.25
RF B	PW, FrzLevPres, MLEHI3km, MLCAPE*SFC3kmSHEAR	ntree = 500

Table 8.9: As in Table 8.7 but for tornadic SCTs.

Statistic	**LogR F**	**SVM RBF F**	**SVM Poly F**	**RF F**	**LogR B**	**SVM RBF B**	**SVM Poly B**	**RF B**
Scitkit-Learn Accuracy	0.8992	0.8992	0.9022	0.8933	**0.9066**	0.8963	0.8889	0.8993
Proportion Correct	0.8810	**0.8988**	**0.8988**	**0.8988**	**0.8898**	0.8869	0.8869	0.8869
POD	0.9167	0.9405	0.9167	0.9167	0.9286	0.9286	**0.9643**	0.9048
FAR	0.1444	0.1319	**0.1149**	**0.1149**	0.1236	0.1429	0.1649	0.1264
PoFD	0.1548	0.1429	**0.1190**	**0.1190**	0.1310	0.1548	0.1905	0.1310
CSI	0.7938	**0.8229**	0.8191	0.8191	0.8211	0.8041	0.81	0.8
HSS	0.7619	**0.7976**	**0.7976**	**0.7976**	**0.7976**	0.7738	0.7738	0.7738
GSS	0.6154	**0.6634**	**0.6634**	**0.6634**	**0.6634**	0.6311	0.6311	0.6311
Bias Ratio	1.0714	1.0833	**1.0357**	**1.0357**	1.0595	1.0833	1.1548	**1.0357**
Odds Ratio Skill Score	0.9673	0.9791	0.9757	0.9757	0.9771	0.9722	**0.9827**	0.9688

Table 8.10: Features and tuning parameters selected for the best performing models presented in Table 8.9.

Model	Features	Parameters
LogR F	MLLapR, MLEL, SWEAT, SRH1km, FrzLevPres	C = $2^{0.5}$
SVM RBF F	MLLapR, MLEL, SWEAT, SHERB, FrzLevPres	C = 2, G = 0.25
SVM Poly F	MLEL, SWEAT, SHERB, FrzLevPres	C = $2^{0.5}$, E = 1, G = 0.75
RF F	MLLCL, MLEL, SWEAT, SHERB, FrzLevPres	ntree = 750
LogR B	MLEL, K-Index, SWEAT, SRH1km, FrzLevPres	C = 2
SVM RBF B	MLCIN, MLEL, SWEAT, SRH1km, FrzLevPres	C = $2^{0.5}$, G = 0.25
SVM Poly B	MLCIN, MLEL, SFC1kmSHEAR, FrzLevPres	C = 2, E = 1, G = 0.25
RF B	MLEL, TT, SRH1km, FrzLevPres, MLEHI1km	ntree = 950

Table 8.11: As in Table 8.7 but for wind SCTs.

Statistic	LogR F	SVM RBF F	SVM Poly F	RF F	LogR B	SVM RBF B	SVM Poly B	RF B
Scitkit-Learn Accuracy	0.8596	0.8920	0.8787	0.8914	0.8452	0.8778	0.8845	**0.8995**
Proportion Correct	**0.8834**	0.8776	0.8695	**0.8834**	0.8453	**0.8834**	0.8707	0.8741
POD	**0.9238**	0.8984	0.8707	0.9145	0.8545	0.9099	0.8614	0.8822
FAR	0.1453	0.1375	0.1313	0.1391	0.1610	0.1360	**0.1224**	0.1318
PoFD	0.1570	0.1432	0.1316	0.1478	0.1640	0.1432	**0.1201**	0.1339
CSI	**0.7984**	0.7859	0.7694	0.7968	0.7341	0.7960	0.7691	0.7780
HSS	**0.7667**	0.7552	0.7390	**0.7667**	0.6905	**0.7667**	0.7413	0.7483
GSS	**0.6217**	0.6067	0.5861	**0.6217**	0.5273	**0.6217**	0.5890	0.5978
Bias Ratio	1.0808	1.0416	**1.0023**	1.0624	1.0185	1.0531	0.9815	1.0162
Odds Ratio Skill Score	**0.9697**	0.9629	0.9560	0.9681	0.9354	0.9675	0.9570	0.9595

Table 8.12: Features and tuning parameters selected for the best performing models presented in Table 8.11.

Model	Features	Parameters
LogR F	MLLI, MLEL, TT, SFC1kmSHEAR, SHERB	C = $2^{0.5}$
SVM RBF F	MLLI, MLLCL, SWEAT, SHERB, FrzLevPres	C = $2^{0.5}$, G = 0.25
SVM Poly F	MLEL, TT, SHERB, FrzLevPres, MLCAPE*SFC3kmSHEAR	C = $2^{0.5}$, E = 3, G = 0.75
RF F	MLLCL, MLEL, SFC1kmSHEAR, STP, MLCAPE*SFC6kmSHEAR	ntree = 500
LogR B	MLLI, MLLCL, SWEAT	C = $2^{0.5}$
SVM RBF B	MLLI, MLLCL, SHERB, FrzLevPres	C = $2^{2.5}$, G = 0.25
SVM Poly B	MLLI, MLLCL, SWEAT, FrzLevPres, MLCAPE*SFC3kmSHEAR	C = 2, E = 3, G = 0.5
RF B	MLLCL, MLEL, PW, SWEAT, SHERB	ntree = 300

The most frequently selected index across the hail SCT models is FrzLevPres. The height of the freezing level significantly affects hail growth. A lower freezing level reduces the amount of melting that can occur before a hail stone reaches the ground (Dessens et al., 2015; Bedka et al., 2018), while allowing hail growth to occur sooner within the storm updraft, although this can also affect the amount of supercooled water in strong updraft regions, further modulating the hail size distribution (Prein and Holland, 2018). MLEL was the second most selected feature and was likely selected as a higher equilibrium level allows a thunderstorm to develop over a significant vertical depth, which is a common feature of severe convection (Bedka et al., 2018). It also would increase the time that hydrometeors reside above the freezing level, and could relate to the depth of the optimal hail growth layer (Allen et al., 2020). Furthermore, for a given quantity of CAPE, the height of the EL partially controls the depth of the unstable layer, with a lower EL leading to CAPE concentrated over a smaller region and more intense updrafts that can hold larger hail stones before they fall out of a storm (Blanchard, 1998; McCaul and Weisman, 2001). Additional parameters that were selected frequently include EffSRH, and a two-way interaction between MLCAPE and wind shear (especially MLCAPE*SFC3kmSHEAR). The amount of SRH in the lower atmosphere affects the ability for storms to develop rotation in their updraft, potentially leading to supercell development. Supercell storms are prolific hail producers (e.g., Buckley et al., 2001; Soderholm et al., 2017b) due to their ability to continually cycle hail stones through the updraft before it falls to the ground. The MLCAPE and wind shear interaction term has a similar relationship, with increasing wind shear leading to more organised storms, and significant shear resulting in supercell-type convection, while MLCAPE relates to the amount of instability available for storms to form in the presence of a lifting mechanism (Brooks et al., 2003; Allen et al., 2011).

The tornadic SCT models appear to perform better in general than the hail SCT models. However, this could in part be due to the reduced number of samples that the models were tested on. Both the SVM Poly F and RF F models performed extremely similarly on the contingency statistics (Table 8.9), with the SVM Poly F model performing slightly better on Scikit-Learn accuracy (0.9022 vs 0.8933). Additionally, the SVM Poly F model appears to

remain closer to the top left corner in the ROC curve (Figure 8.4.10) and the top right corner in the PR curve (Figure 8.4.11), despite having the same area under the curve as the RF F model. As a result, the SVM Poly F model was selected for analysis of the tornadic environment climatology. The model performs rather well with a high POD (0.9167; Table 8.9) and low FAR (0.1149) resulting in a relatively high CSI (0.8191) although this is beaten by both the SVM RBF F (CSI = 0.8229) and LogR B (CSI = 0.8211) models. Despite this, the SVM Poly F model has similarly high HSS (0.7976) and GSS (0.6634), and a low bias ratio (1.0357). Given the models all perform rather similarly, it was decided to also apply the LogR B model to the BARRA data set, to assess differences in the diagnosed tornadic SCT environment climatology. This model was selected as it achieved the highest Scikit-Learn accuracy, had similar or better statistics results to other models, and included SRH1km as a feature which is strongly related to tornadic storms overseas (Edwards and Thompson, 2000; Thompson et al., 2002, 2003; Rasmussen, 2003).

The SVM Poly F model included MLEL, SWEAT, SHERB and FrzLevPres as predictors, with C = $2^{0.5}$, E = 1, and G = 0.75. Interestingly, the RF F model included these and MLLCL as predictors. Meanwhile, the LogR B model included MLEL, K-Index, SWEAT, SRH1km and FrzLevPres. Overall, the most commonly selected features among the models were FrzLevPres, MLEL, SWEAT and SRH1km. The SWEAT index includes measures of low-level moisture, instability, low- and mid-level wind speeds, and the change in wind direction between these two levels (generally indicative of warm air advection; Miller, 1972). Low-level moisture is often characterised by the LCL height, which has been shown to discriminate between tornadic and non-tornadic storms (Rasmussen and Blanchard, 1998; Craven and Brooks, 2004). Low-level moisture can modulate the alignment of the storm mesocyclone with the near-surface circulation within the cold pool, and change buoyancy within the cold pool, both affecting how well vorticity near the surface is stretched and intensified to tornado strength (Markowski and Richardson, 2014; Brown and Nowotarski, 2019). Wind speed and direction changes over height affect the amount of environmental vorticity available for a storm to develop rotation. Furthermore, the low-level wind profile influences the relationship between low-level moisture and circulation alignment, while mid-level environmental wind can affect the location of storm outflow and resulting characteristics of the low-level mesocyclone (Brooks et al., 1994b; Brown and Nowotarski, 2019). Additionally, SRH1km is related particularly to tornadic storms as sources of low-level vorticity can be tilted by thunderstorm updrafts and stretched, causing amplification of vorticity and development of a low-level mesocyclone that is supportive of tornado development (Edwards and Thompson, 2000; Thompson et al., 2002, 2003; Rasmussen, 2003). In fact, Coffer et al. (2019) have shown calculation of SRH with progressively shallower layers to increase forecast skill, with SRH over 0–500 m being effective in discriminating between supercells that produce significant tornadoes, weak tornadoes or no tornadoes. Given the freezing level height has an impact on microphysical attributes (Dessens et al., 2015; Bedka et al., 2018), and can influence resulting outflow characteristics (Proctor, 1989; Richter et al., 2014), it likely has an effect on low-level vorticity and the ability for tornadoes to develop. The thunderstorm updraft is critical for stretching low-level vorticity to tornado strength. Since MLEL can affect updraft strength through varying the concentration of CAPE over storm depth (Blanchard, 1998; McCaul and Weisman, 2001), it is possible that this has some effect on tornado development beyond being useful as an indicator of potential SCTs (Bedka et al., 2018). Finally, low-topped supercells during winter are a common convective mode leading to cool-season tornadoes, that account for a large number of Australian tornadoes (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009), and might be indicated by the EL height, in combination with other environmental attributes.

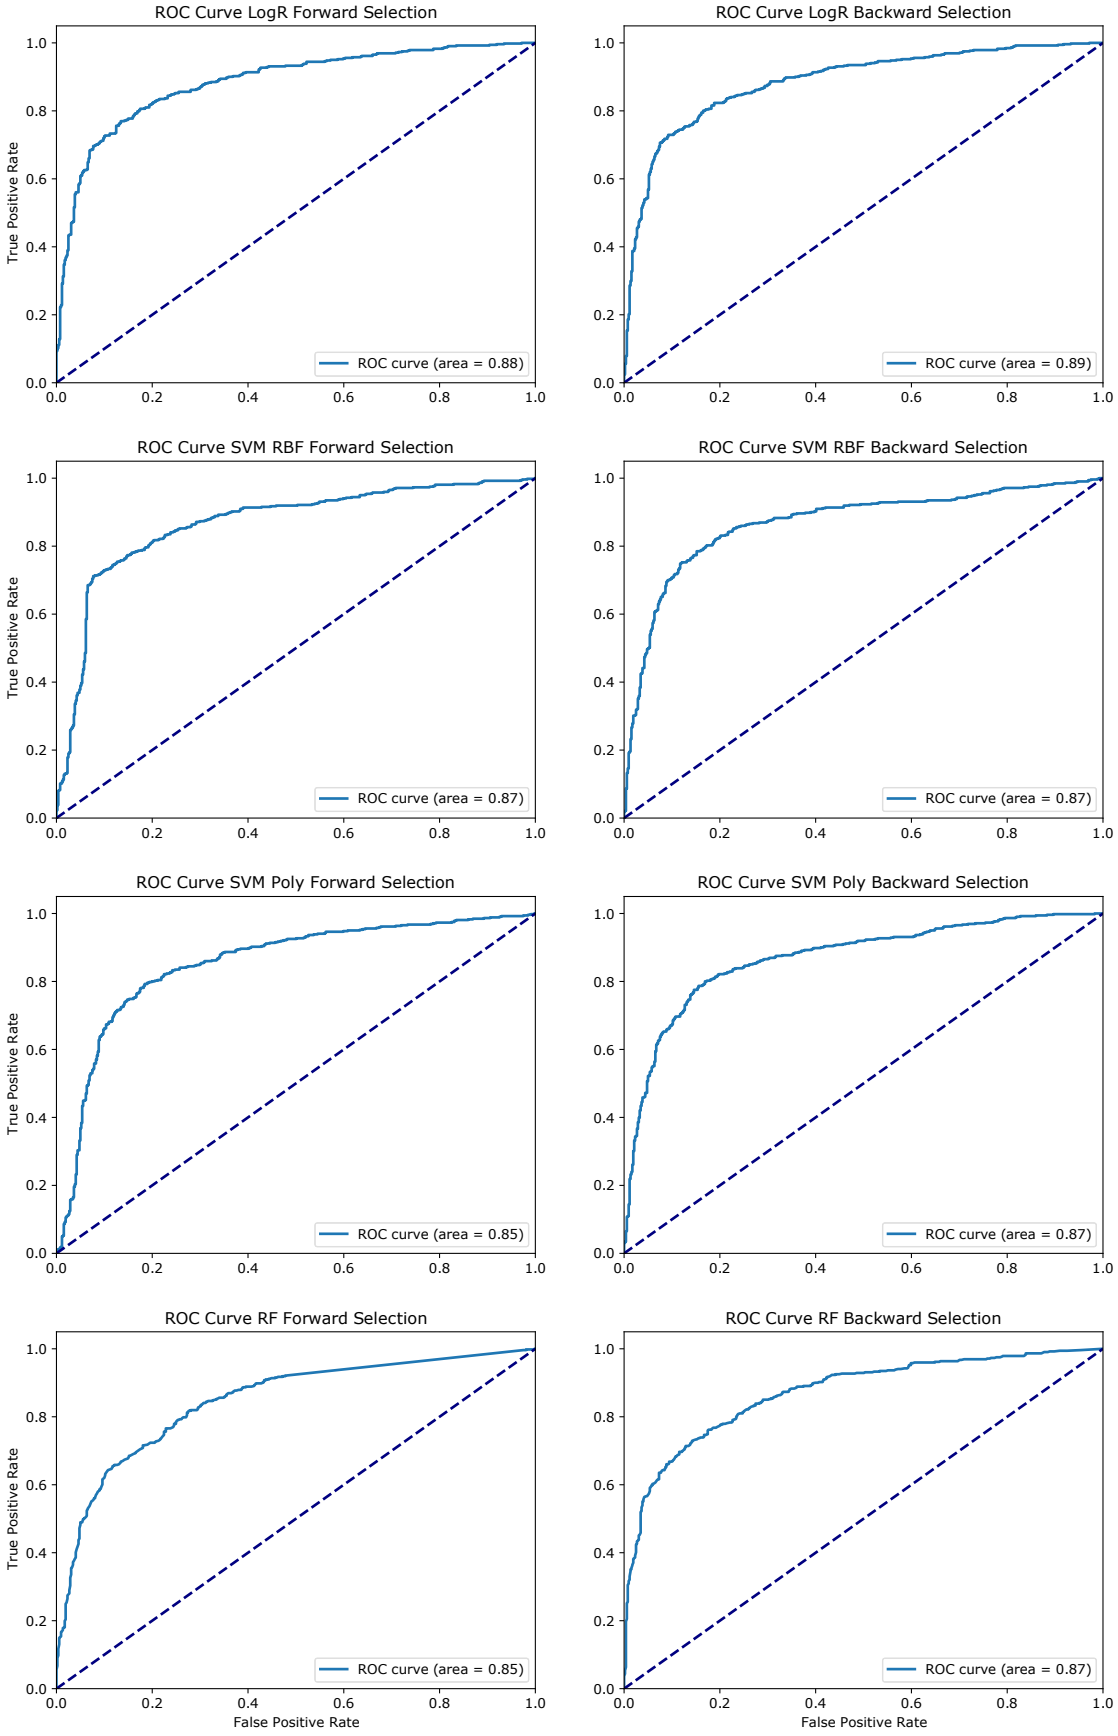

Figure 8.4.8: Receiver Operating Characteristic (ROC) curves for each statistical model for hail SCTs described in Table 8.8, with those developed through forward selection on the left and those through backward selection on the right.

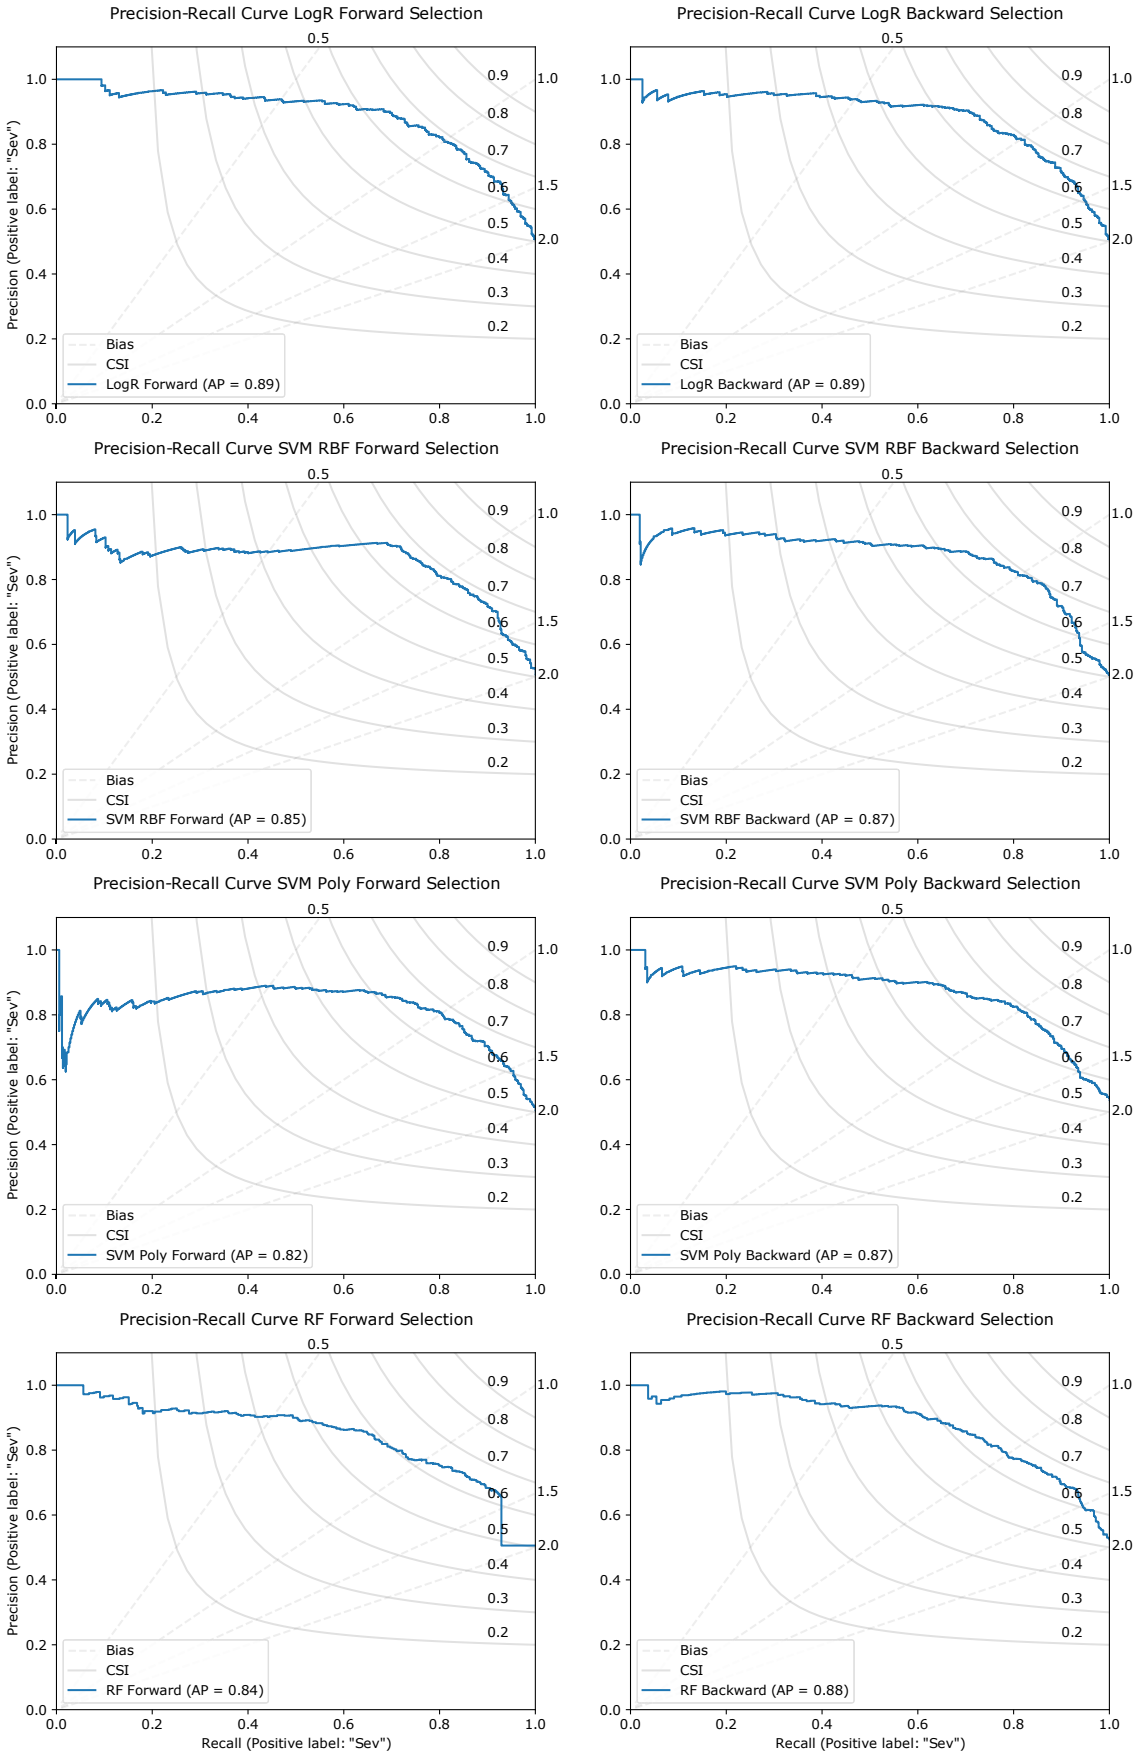

Figure 8.4.9: Precision-Recall (PR) curves (solid blue lines) for each statistical model for hail SCTs described in Table 8.8, with those developed through forward selection on the left and those through backward selection on the right. Solid, dark grey lines depict lines of constant CSI, while dashed light grey lines depict lines of constant bias.

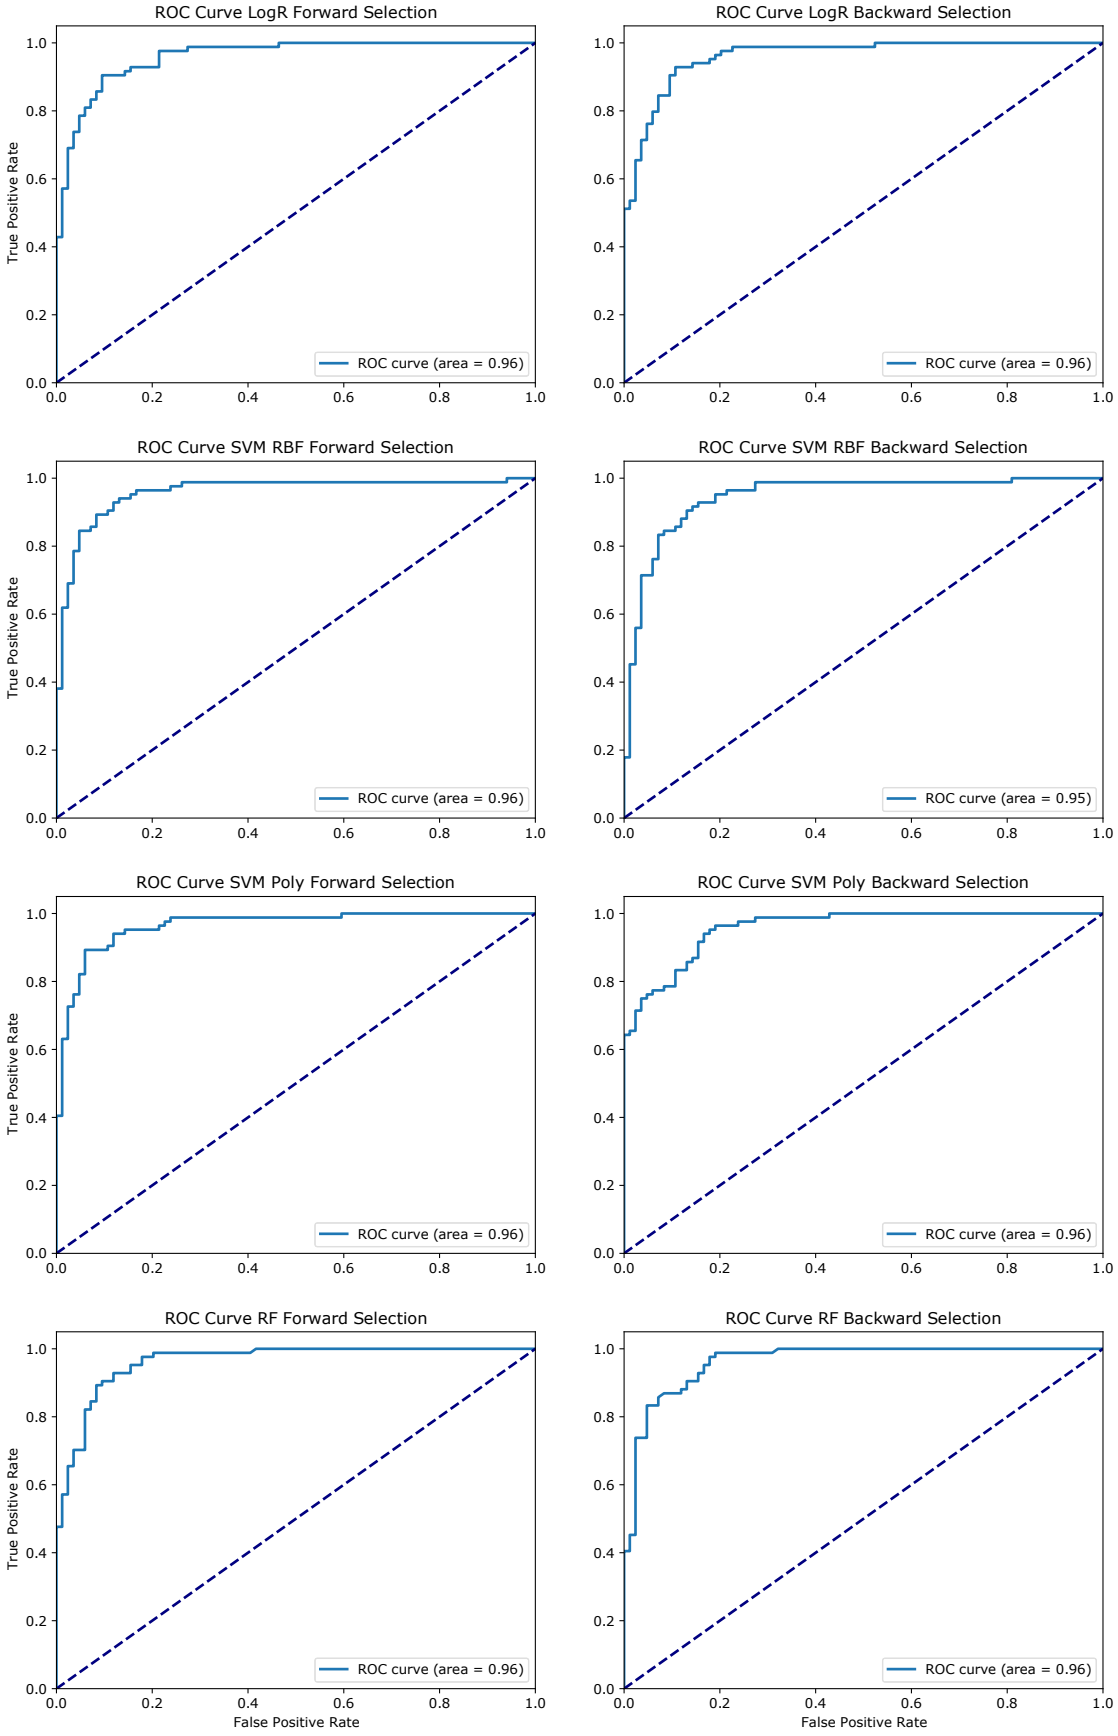

Figure 8.4.10: ROC curves for each statistical model for tornadic SCTs described in Table 8.10, as outlined in Figure 8.4.8.

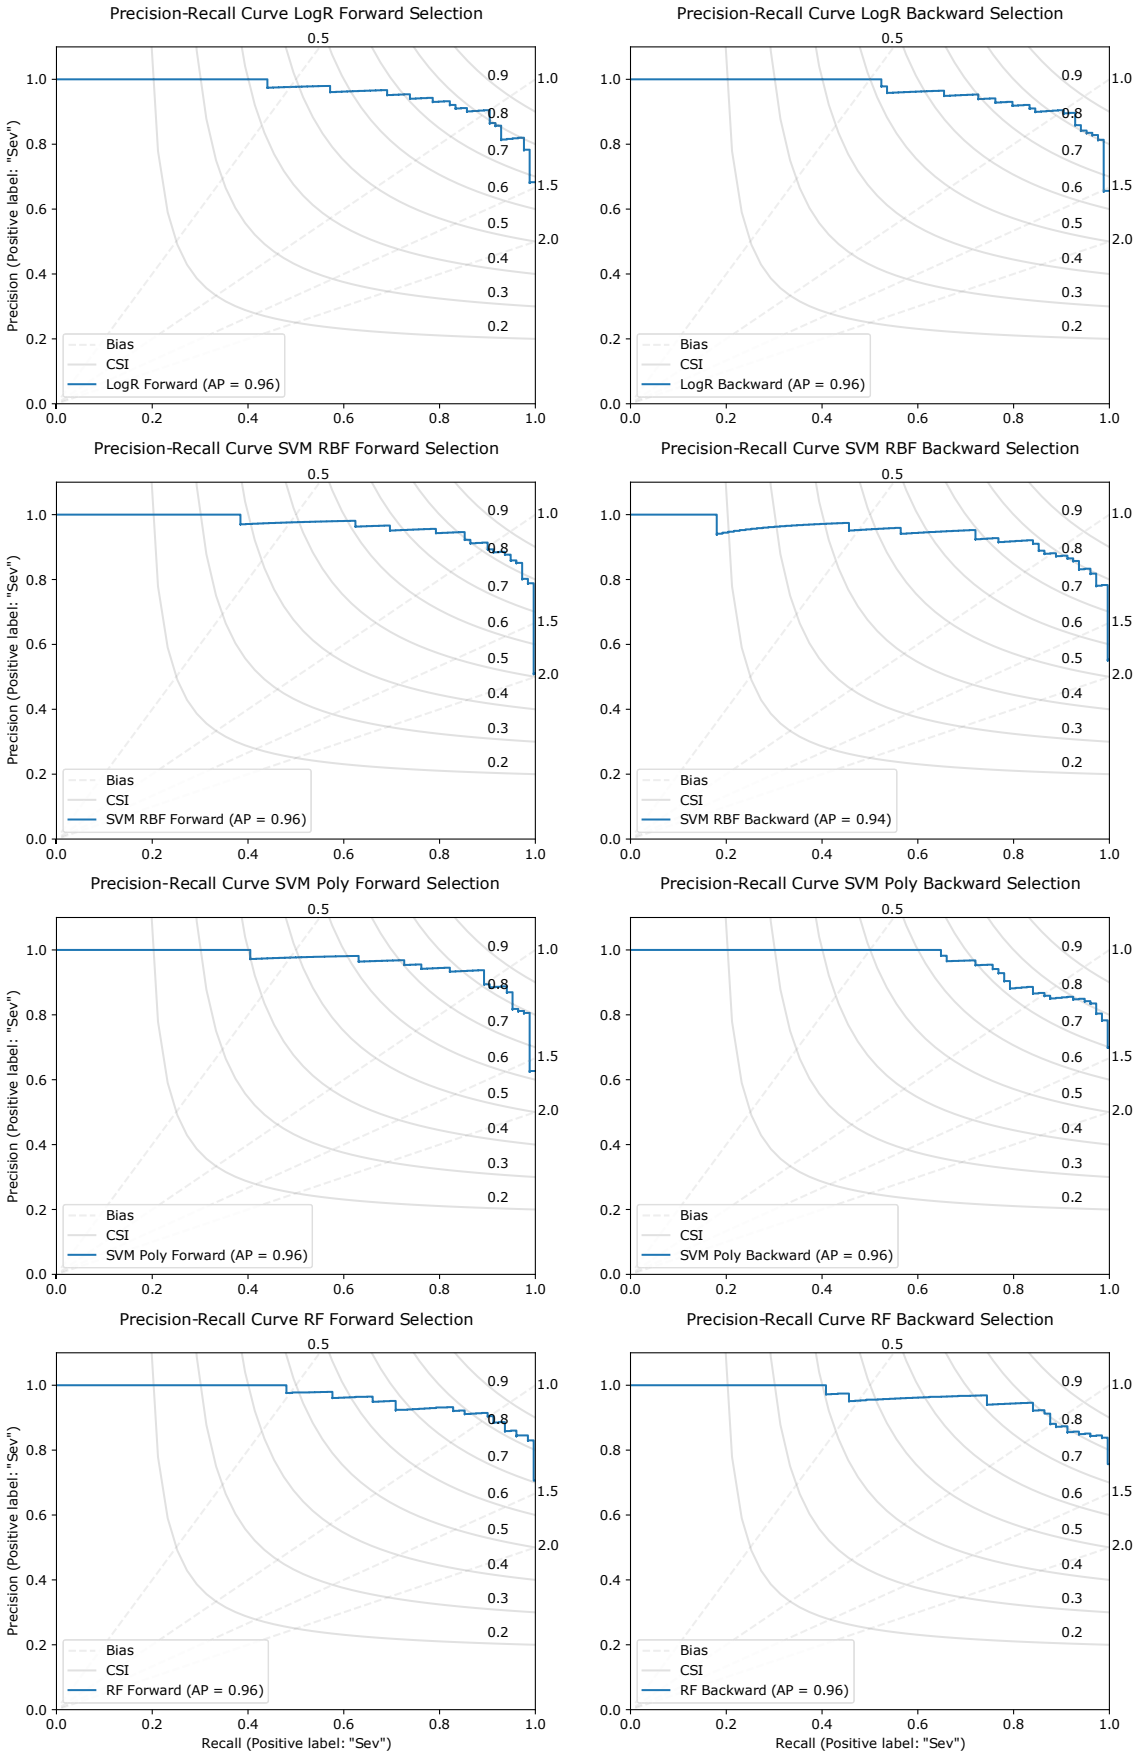

Figure 8.4.11: PR curves for each statistical model for tornadic SCTs described in Table 8.10, as outlined in Figure 8.4.9.

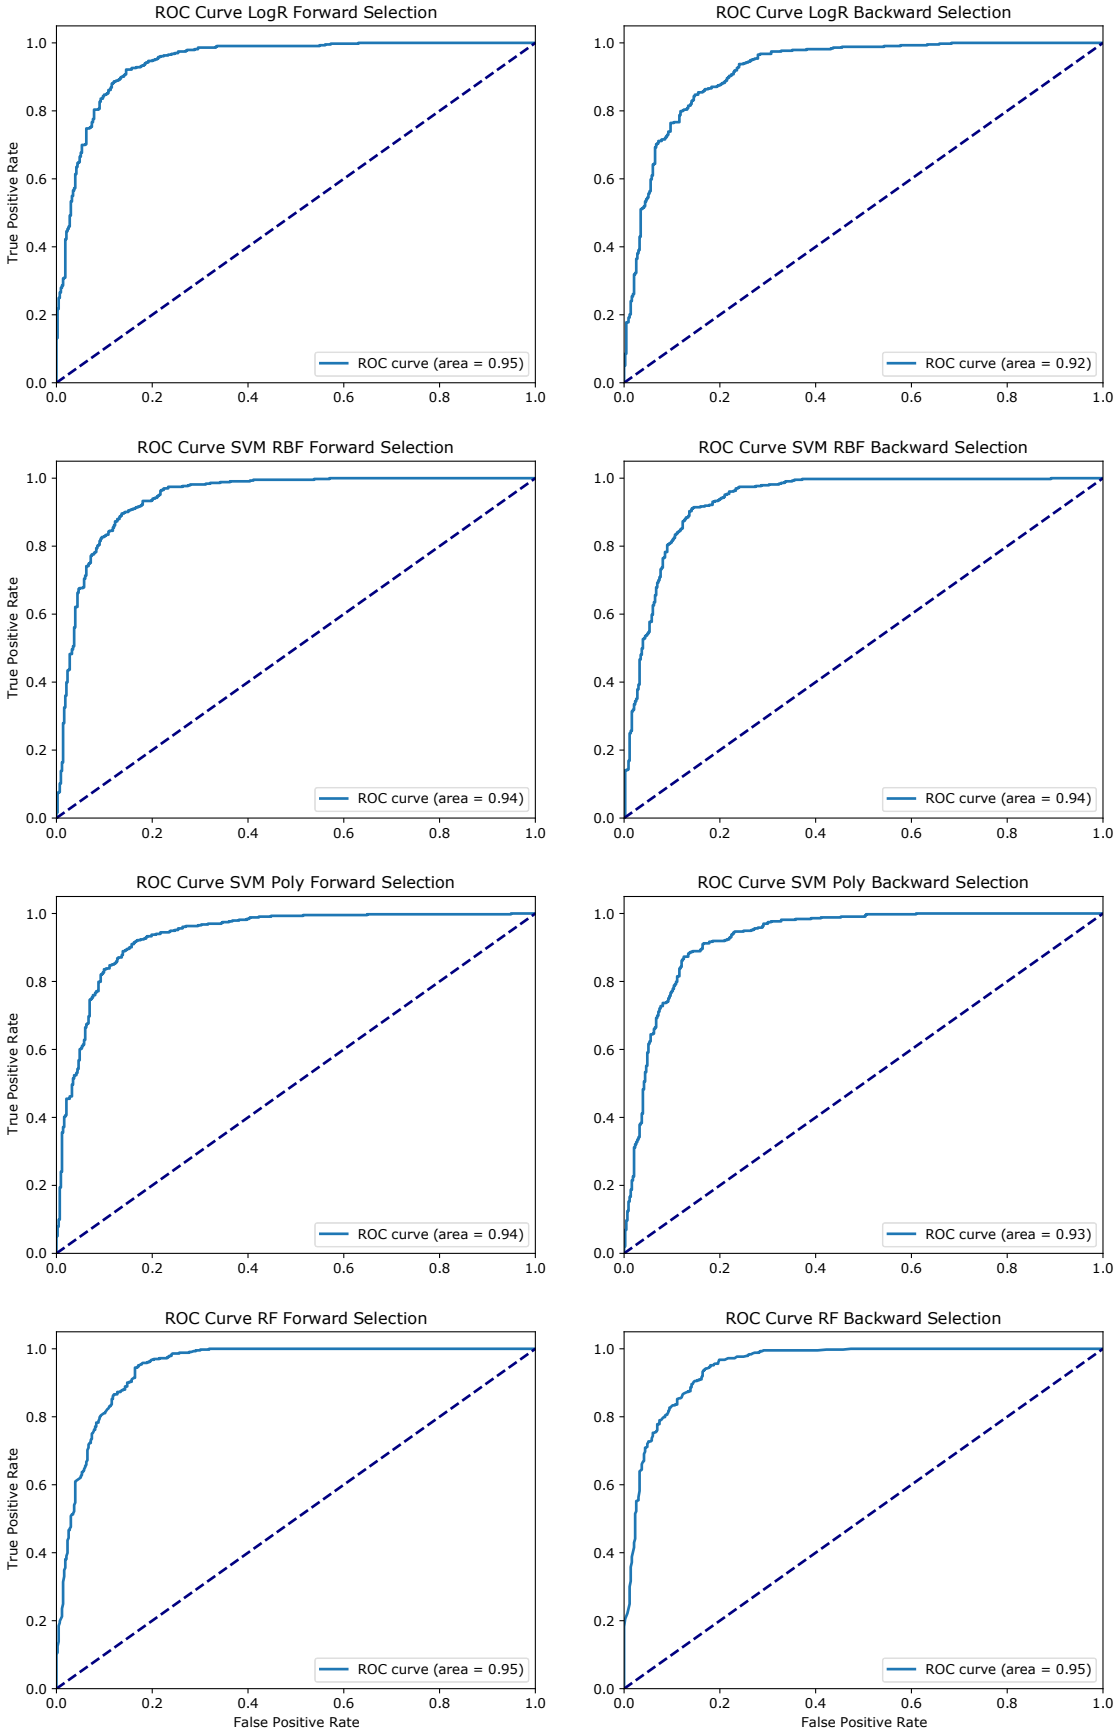

Figure 8.4.12: ROC curves for each statistical model for wind SCTs described in Table 8.12, as outlined in Figure 8.4.8.

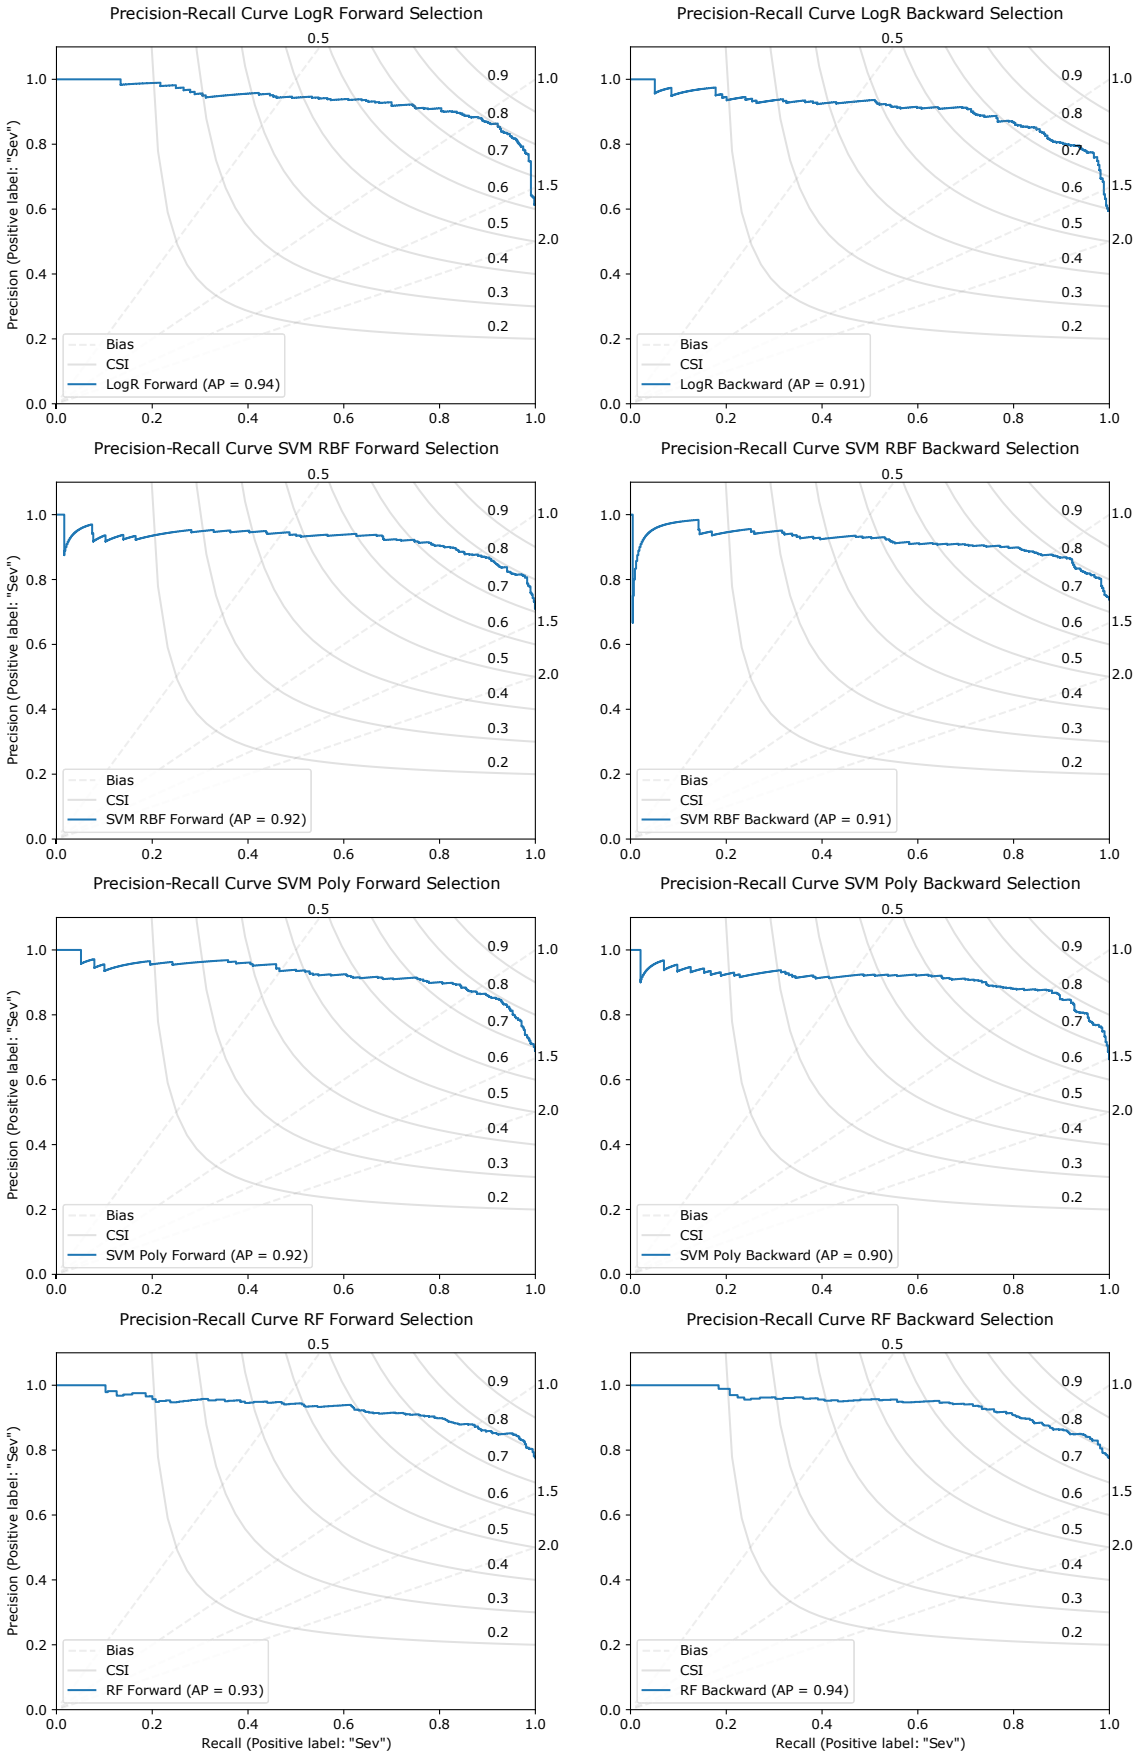

Figure 8.4.13: PR curves for each statistical model for wind SCTs described in Table 8.12, as outlined in Figure 8.4.9.

The best performing model for wind SCT environments is the LogR F model. Although

it achieves the second highest FAR and PoFD of all models (0.1453 and 0.1570 respectively; Table 8.11), it obtains the highest POD (0.9238), CSI (0.7984), HSS (0.7667) and GSS (0.6217). This model also appears to have the greatest area under the ROC curve, along with both RF models (Figure 8.4.12), and shares the greatest average precision of the PR curves with the RF B model (Figure 8.4.13). The RF B model appears to perform better for decision thresholds that achieve extremely high recall, with CSI above 0.7, compared to the LogR F model which has CSI reduce to 0.6 for very high recall. However, given the generally better performance of the LogR F model over the RF B model in the contingency statistics (Table 8.11), it was selected for the diagnosis of the wind SCT environment climatology over Australia.

The LogR F model was selected for application to the wind SCT environment climatology, with predictors MLLI, MLEL, TT, SFC1kmSHEAR and SHERB, and C= $2^{0.5}$. The most commonly selected predictors for wind SCT models were MLLI, MLLCL, SWEAT, SHERB, FrzLevPres and MLEL. MLLI likely was selected as it is an indicator of the likelihood of thunderstorm development, and severe convective wind storms appear to form in a wide variety of environments compared to other convective hazards (Taszarek et al., 2020). As mentioned earlier, MLLCL relates to moisture in the lower atmosphere. Environments of reduced low-level moisture (characterised by a higher LCL) lead to greater evaporative cooling and dry microbursts (Wakimoto, 1985; McCaul and Cohen, 2002). SWEAT was developed to distinguish environments that can produce storms that become sufficiently organised to cause severe weather (Miller, 1972). Other wind SCT climatologies have identified high-shear low-CAPE environments as part of the parameter space for severe wind producing storms (Taszarek et al., 2020), which is likely why SHERB was selected as it was developed to identify the potential of damaging winds or tornadoes in high-shear low-CAPE environments (Sherburn and Parker, 2014). Freezing level pressure also is the height of the melting level, with melting of hail contributing negative buoyancy, and dry air at the melting level shown to increase downdraft acceleration (Proctor, 1989; Richter et al., 2014). MLEL was a commonly selected parameter in the logistic regression models of Brown and Dowdy (2021a) for modelling Australian severe convective winds.

Table 8.13 provides some key contingency statistics for the discriminant relationship in Allen et al. (2011) (their Equation 7) as applied to the testing data in this project for each SCT hazard type. For hail SCTs, the discriminant relationship has an improvement of POD (0.8592, compared to 0.8138 for the SVM RBF B model; Tables 8.7 & 8.13). However, this is at the expense of an increased FAR (0.3090; Table 8.13), causing a lower CSI and HSS (0.6207 and 0.5777, respectively) in the discriminant model. For tornadic SCTs, the discriminant model has lower POD compared to all 8 models developed here (0.8636, Table 8.13) and a significantly higher FAR that is almost triple that of the SVM Poly F model (0.3214 compared to 0.1190 for the SVM Poly F model; Table 8.13). The resulting low CSI and HSS (0.6129 and 0.5714, respectively; Table 8.13) suggests that tornadic SCTs in Australia are dependent on more than CAPE and 0–6 km bulk shear. Finally, the discriminant relationship performs worse than the LogR F model that was selected for wind SCTs. It has a lower POD (0.8901 compared to 0.9238 for the LogR F model; Table 8.13), and a higher FAR (0.2702 compared to 0.1453; Table 8.13), resulting in significantly lower CSI and HSS. Overall, it appears that the discriminant relationship described in Allen et al. (2011) performs generally worse than the statistical models developed for individual hazards here, although it still provides skillful separation of non-severe and severe cases in the test data set. It is possible that this reduced performance is partially due to the short data set that the discriminant relationship was developed over, spanning from 2003–2010, with only 1550 reports (Allen et al., 2011). However, it also highlights how there is much to be gained by considering hazards individually, rather than collectively, and a wide

Table 8.13: Contingency statistics from application of the discriminant relationship described in (Allen et al., 2011, , their Equation 7), applied to each of the test data sets used for hail, tornado and wind SCTs.

Statistic	Hail	Tornado	Wind
POD	0.8592	0.8636	0.8901
FAR	0.3090	0.3214	0.2702
CSI	0.6207	0.6129	0.6695
HSS	0.5777	0.5714	0.6397

range of potential parameters to classify SCT hazard environments.

8.4.4 BARRA Climatology

Figure 8.4.14 provides filled contours of the frequency of environments (in days per year) that are supportive of each SCT hazard using the SVM RBF B model for hail (Figure 8.4.14a), the LogR F model for wind (Figure 8.4.14b), the SVM Poly F model for tornadoes (Figure 8.4.14c), and the LogR B model for tornadoes (Figure 8.4.14d, as a comparison of the potential value from including SRH1km as a predictor). Interestingly, while all models perform well on the test data, they appear to identify a significantly high frequency of environments supportive of SCTs, with some regions experiencing over 200 environments on average each year. This is particularly apparent for the region in northern Australia. Although northern Australia experiences a number of thunderstorms during the monsoon season, and some of these are severe, the high frequency of environments is perculiar. Some of this might be due to the models being trained and tested on data that is more concentrated around southern and eastern Australia. However, Allen and Karoly (2014) found a relatively high frequency of severe and significant severe environments across northern Australia in the ERA-Interim data set during the warm season (September–April), which also appears in the global climatology of Brooks et al. (2003), although the frequency is much reduced compared to here. This high frequency of severe environments could be due to significant instability in the region, which is particularly apparent over northwest Australia (Allen and Karoly, 2014). More realistic results could be achieved by constructing these models for sub-regions of Australia, or by making the sampling across Australia more uniform. Despite the potential overrepresentation of supportive SCT environments in parts of Australia, the results are discussed in the remainder of the Chapter as they still provide valuable insight into the Australian SCT hazard climatology.

Hail

There is a large number of environments supportive of hail SCTs in southeastern Australia, following the Great Dividing Range, and in northwestern Australia (Figure 8.4.14a). There are additional, weaker, peaks in hail SCT environments in southwestern Australia and across Tasmania, particularly over the eastern side. In southeast Australia, there is a maximum in the frequency of supportive environments from the central coast of New South Wales (NSW) through to northeast NSW and southeast Queensland. There is a second maximum in the far southeast of Australia. Although there likely is an overrepresentation of environments in northern Australia, given it resides in the tropics, with less VWS available for storm organisation and lifting mechanisms for storms to develop (Allen and Karoly, 2014), the distribution of favourable environments is heterogeneous. In particular, the region surrounding the Gulf of Carpentaria (particularly the Cape York Peninsula) experiences the lowest frequency of environments, while the northwest coast and the region surrounding Darwin has the highest frequency of environments. This could be due to instability being concentrated over this region

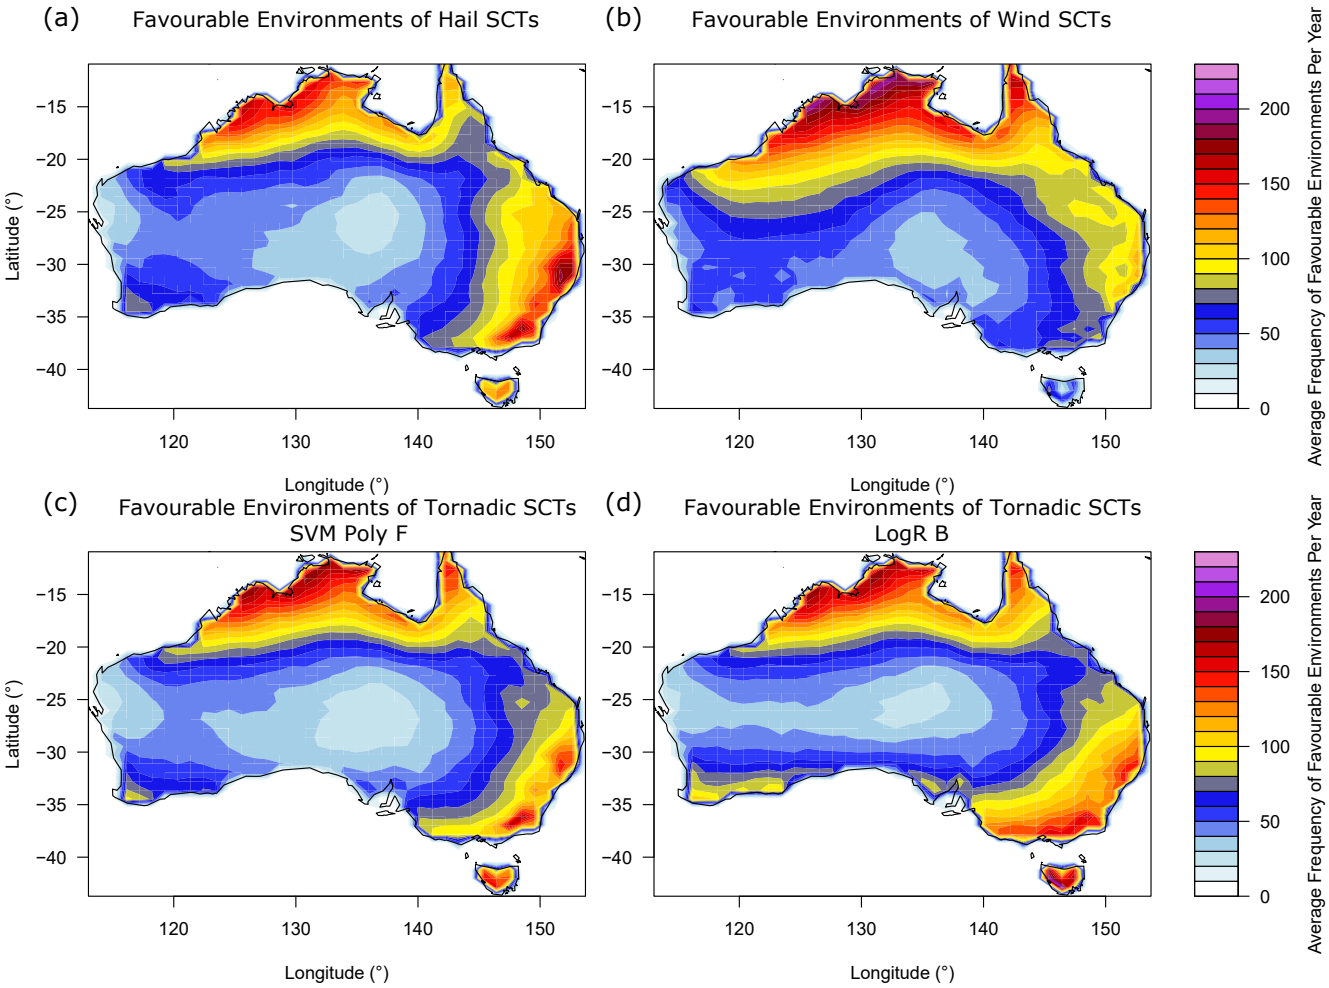

Figure 8.4.14: Average frequency of environments supportive of SCT hazards (in days per year) for (a) hail, (b) wind, (c) tornadoes as diagnosed by the SVR Poly F model, and (d) tornadoes as diagnosed by the LogR B model.

of northern Australia (Allen and Karoly, 2014). Much of central Australia has a low number of environments supportive of hail storms likely due to low moisture levels, although this also extends to coastal regions including the central west coast and the south coast of Australia. Whereas, for the regions of maximum hail environment frequency, the greatest frequency is found just inland from the coast, and degrades as the distance from the coast increases. The lower frequency of environments along the immediate coast might be due to the maritime air mass causing cooler surface air temperatures and a stronger capping inversion, with the air mass warming (and subsequent reduction of capping and increased instability) as it mixes with warmer inland air (Soderholm et al., 2017a).

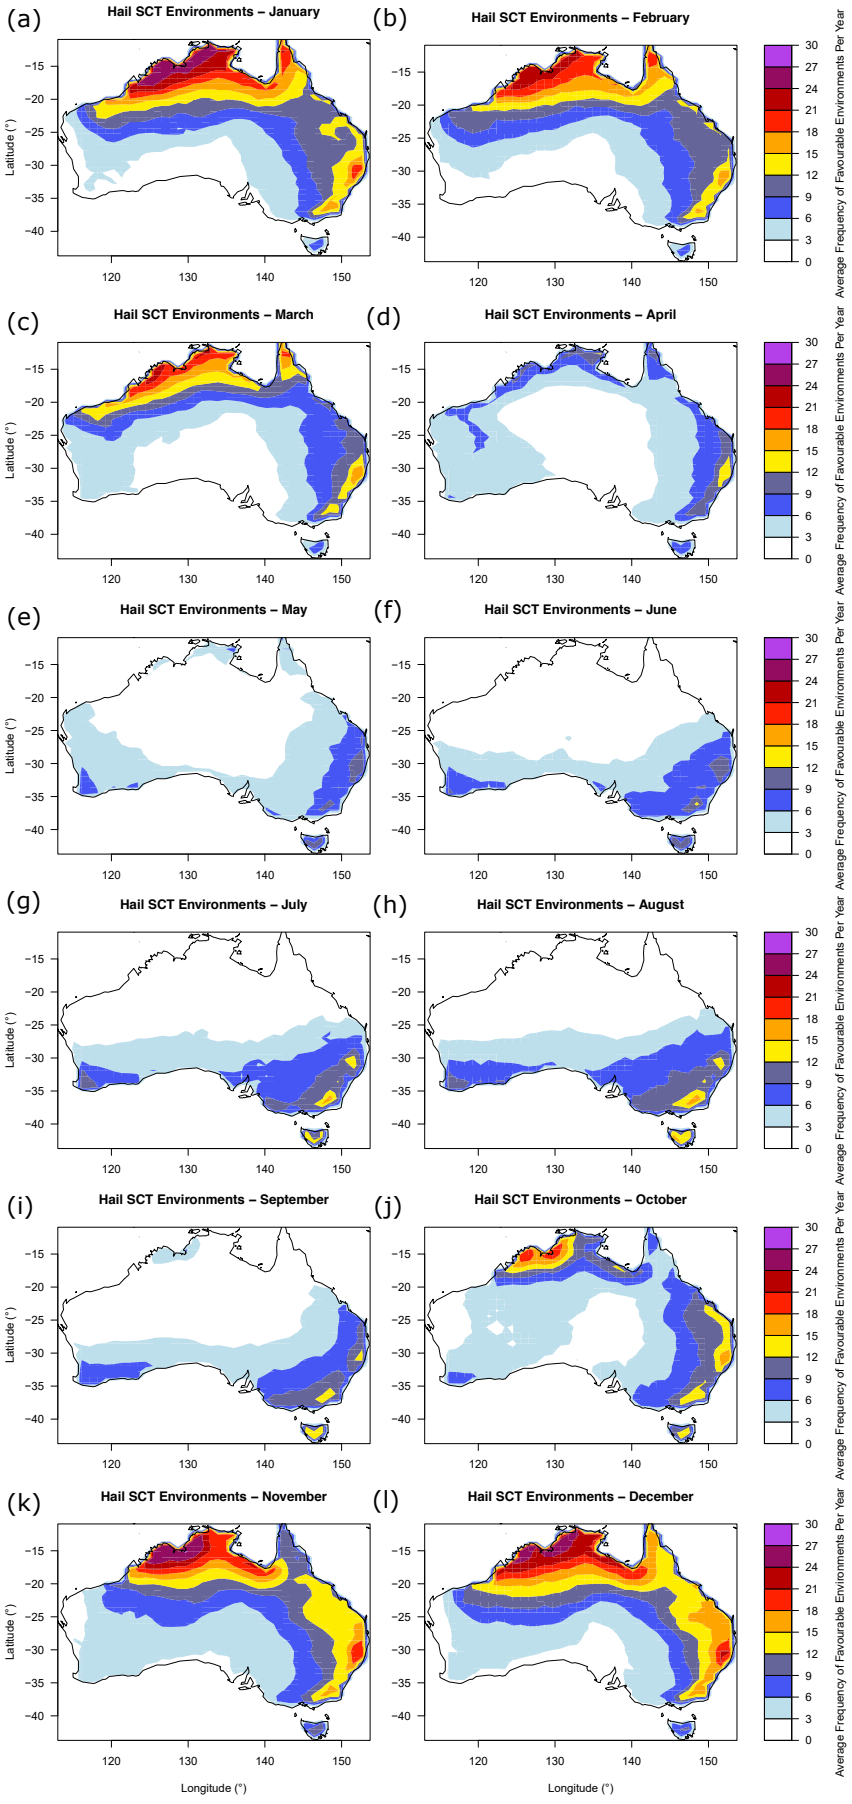

Figure 8.4.15: 1990–2018 average frequency of environments supportive of hail SCTs per year in the BARRA reanalysis data set.

Figure 8.4.15 provides an overview of the average frequency of environments supportive of hail SCTs over 1990–2018 across each month of the year. The majority of supportive environments in northern Australia occur over the wet season, with an increase in frequency appearing first around northwest Australia and some coastal regions surrounding the Gulf of Carpentaria in October, followed by a gradual increase in frequency around the Cape York Peninsula to a peak during January (Figure 8.4.15j-c). This is expected as northern Australia is within the tropics, with the wet season typically beginning after October and the dry season spanning approximately April-late September (Lo et al., 2007; Pope et al., 2009). For eastern Australia and southeastern Australia, the frequency of supportive environments is greatest over the warmer months of the year (September-March; Figure 8.4.15i-c), with the peak frequency beginning in November, and decreasing from January. However, for southeast Australia, there are still 6–18 days each month during the cooler months (April-August) where the environment is supportive of hail SCTs (Figure 8.4.15d-h). For Tasmania, the greatest frequency in supportive environments occurs from May-November (Figure 8.4.15e-k), with the peak frequency occurring during August. Meanwhile, southwest Australia experiences a low frequency of supportive environments during the warmer months of the year, with the greatest frequency occurring over the cooler months of June-September (Figure 8.4.15f-i).

Tornado

Considering the better performing SVM Poly F model in Figure 8.4.14c, the greatest frequency of environments supportive of tornadic SCTs occurs over northern Australia. This could be partly associated with tornadoes forming in the outer rain bands of tropical cyclones (Edwards, 2008, 2012). However, considering the frequency of tropical cyclones that impact the region compared to the number of environments, the resulting frequency still appears extremely high. This high frequency of favourable environments across northern Australia for all SCT hazards is likely due to significant instability in the region such that the statistical models, that are mostly trained on data south of the tropics (e.g., Figure 8.4.4), still classify the environment as severe (Allen and Karoly, 2014). Meanwhile, there is a high frequency of environments across southeast Australia, although the frequency is lower than that for hail SCT environments (Figure 8.4.14a). There is a peak in environments in the far southeast corner of the continent, and again over the central coast of NSW towards northeast NSW, although this does not extend far into southeast Queensland (Figure 8.4.14c). There also is a large number of environments over Tasmania, with the peak similar to the central NSW coast peak, and the number of environments supportive of tornadic SCTs in Tasmania generally larger than the number of environments supportive of hail SCTs (Figure 8.4.14a). Additionally, the greatest frequency of environments occurs across the western side of Tasmania. There appears to be a moderate frequency of environments supportive of tornadic SCTs over southwest Australia, particularly close to the coast. Meanwhile, the lowest frequency of environments supportive of tornadic SCTs occurs over central Australia, towards both the south coast and central west coast of Australia.

There are many similarities between the diagnosed tornadic SCT environments in the SVM Poly F model (Figure 8.4.14c) and the LogR B models (Figure 8.4.14d). However, there are a few notable differences. There is a generally higher frequency of supportive environments in southern Australia with the LogR B model (Figure 8.4.14d), including a higher portion along the SA coastline, and a moderate frequency in the region surrounding Adelaide. There also is a greater frequency of environments in southeast Australia, with much of Victoria experiencing a high frequency of supportive environments, extending through to northeast NSW. The area of supportive environments in southeast Australia also extends further inland from the coast

than for the SVM Poly F model (Figure 8.4.14c). The frequency of supportive environments in Tasmania has increased significantly, with the state experiencing some of the greatest number of environments across Australia. Interestingly, there has been little change in the spatial distribution of environments across northern Australia, highlighting that low-level SRH plays little role in the determination of tornadic environments there.

Figure 8.4.16 provides an overview of the spatial distribution of environments supportive of tornadic storms throughout the year for the SVM Poly F model, while Figure 8.4.17 provides this overview for the LogR B model. Focusing on the SVM Poly F model, the wet season again coincides with the greatest frequency of supportive environments over northern Australia, with a peak in environments during December and January across much of the region (Figure 8.4.16j-d). Similarly, the greatest frequency of supportive environments for eastern Australia, and part of southeast Australia, occurs during the warmer months of October through to February (Figure 8.4.16j-b). However, there also is a moderate number of supportive environments around the central east coast during the cooler winter months (June-August, Figure 8.4.16f-h), likely due to tornadoes spawned by thunderstorms embedded within East Coast Low pressure systems (Hanstrum et al., 2002). Meanwhile for areas of southern Australia, including the southwest, Victoria and Tasmania, tornadic SCT environments increase in frequency over the cooler months of the year, peaking in July and August (Figure 8.4.16e-i), consistent with the high number of cool-season tornadoes that affect southern Australia during this time of year Hanstrum et al. (2002); Mills (2004); Kounkou et al. (2009). This also is the peak occurrence time for the far southeast of Australia, although there is a substantial number of environments during the warmer months of the year too. The peak in Tasmania appears to wane slower than for southwest Australia and the western area of Victoria, with a moderate frequency of environments occurring from September-December (Figure 8.4.16i-l). This is likely due to its poleward location resulting in more cold fronts and low pressure systems passing through the area as they progress towards the South Pole in the spring and summer months.

Turning to the LogR B model, the patterns for northern Australia are very similar to the SVM Poly F model, highlighting the lack of influence SRH1km appears to play on tornadic environments in this region, for the models presented (Figure 8.4.17). However, there are differences in both eastern and southern Australia. The greatest frequency in environments for eastern Australia is again over the warmer months, with a peak during November and December (Figure 8.4.16j-b). The spatial distribution appears to extend slightly further inland from the east coast for the LogR B model compared to the SVM Poly F model (Figure 8.4.16j-b). The greatest contrast in the monthly distribution of tornadic SCT environments is for southern Australia. The LogR B model detects an increase in the frequency of supportive environments beginning in Tasmania in April (Figure 8.4.17d), likely as strong cold frontal systems begin to progress north. The increased frequency of supportive environments then reaches the southern mainland of Australia during May, and continues to increase in Tasmania and southern Australia, reaching the peak during July and August like the SVM Poly F model (Figure 8.4.17e-h). The peak frequency appears to extend much further north, and also encompasses parts of the SA coastline that were not detected in the SVM Poly F model (Figure 8.4.16g-h). This extension further north also includes a greater frequency of environments over the central NSW coastline (Figure 8.4.17g-h).

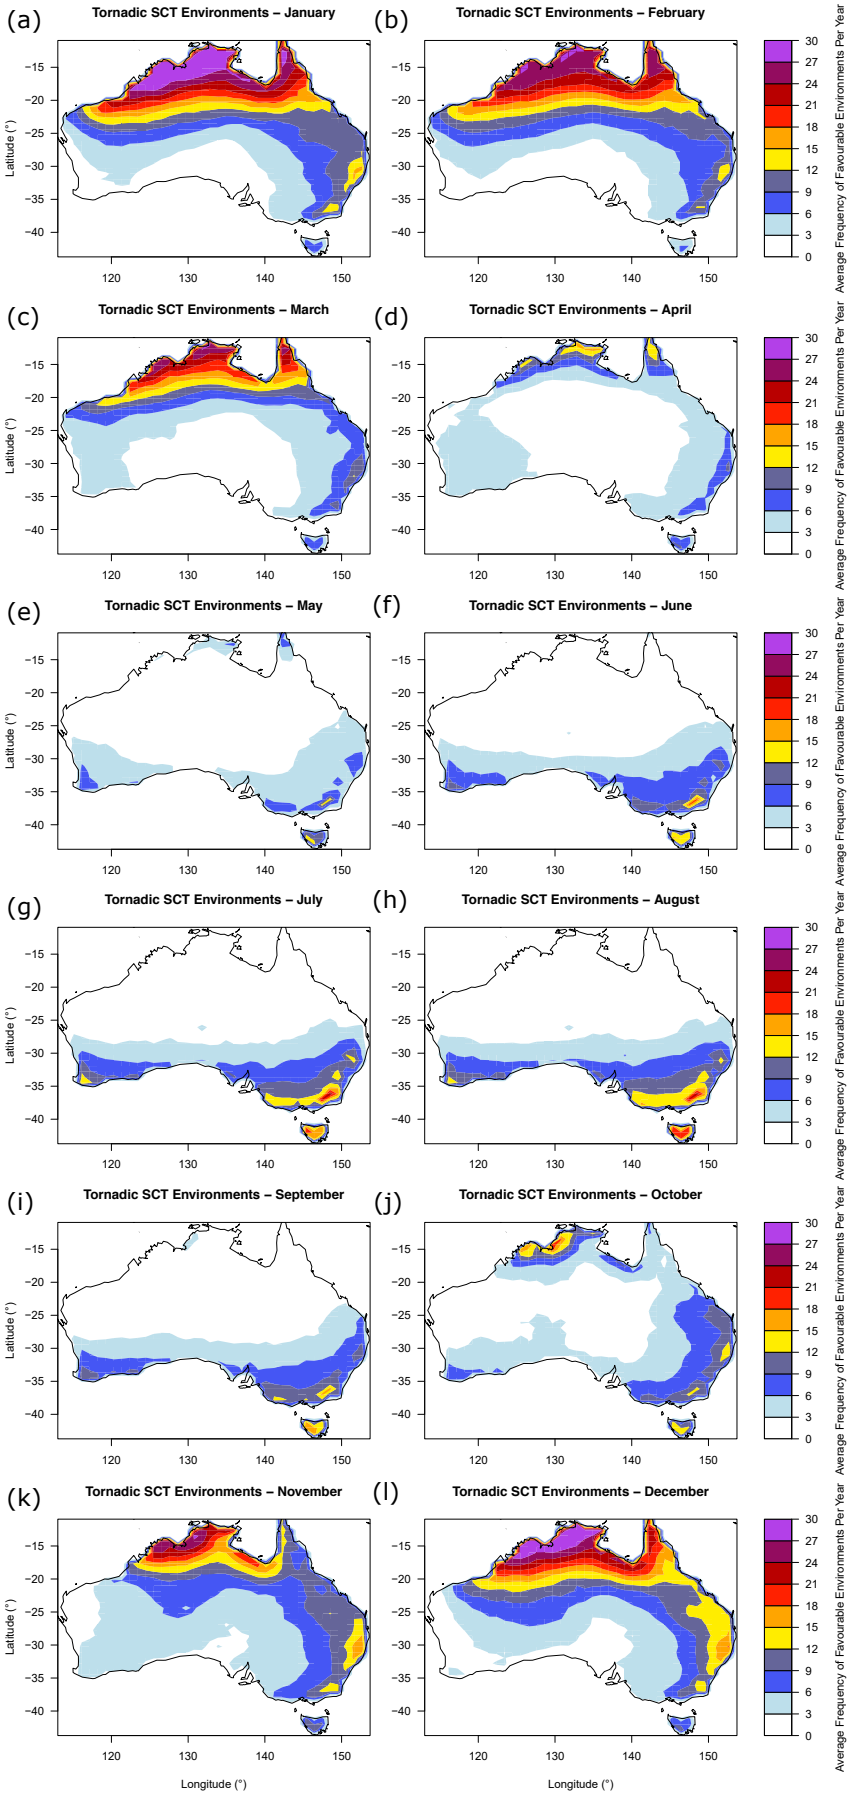

Figure 8.4.16: As in Figure 8.4.15 but for tornadic SCTs using the SVM Poly F model.

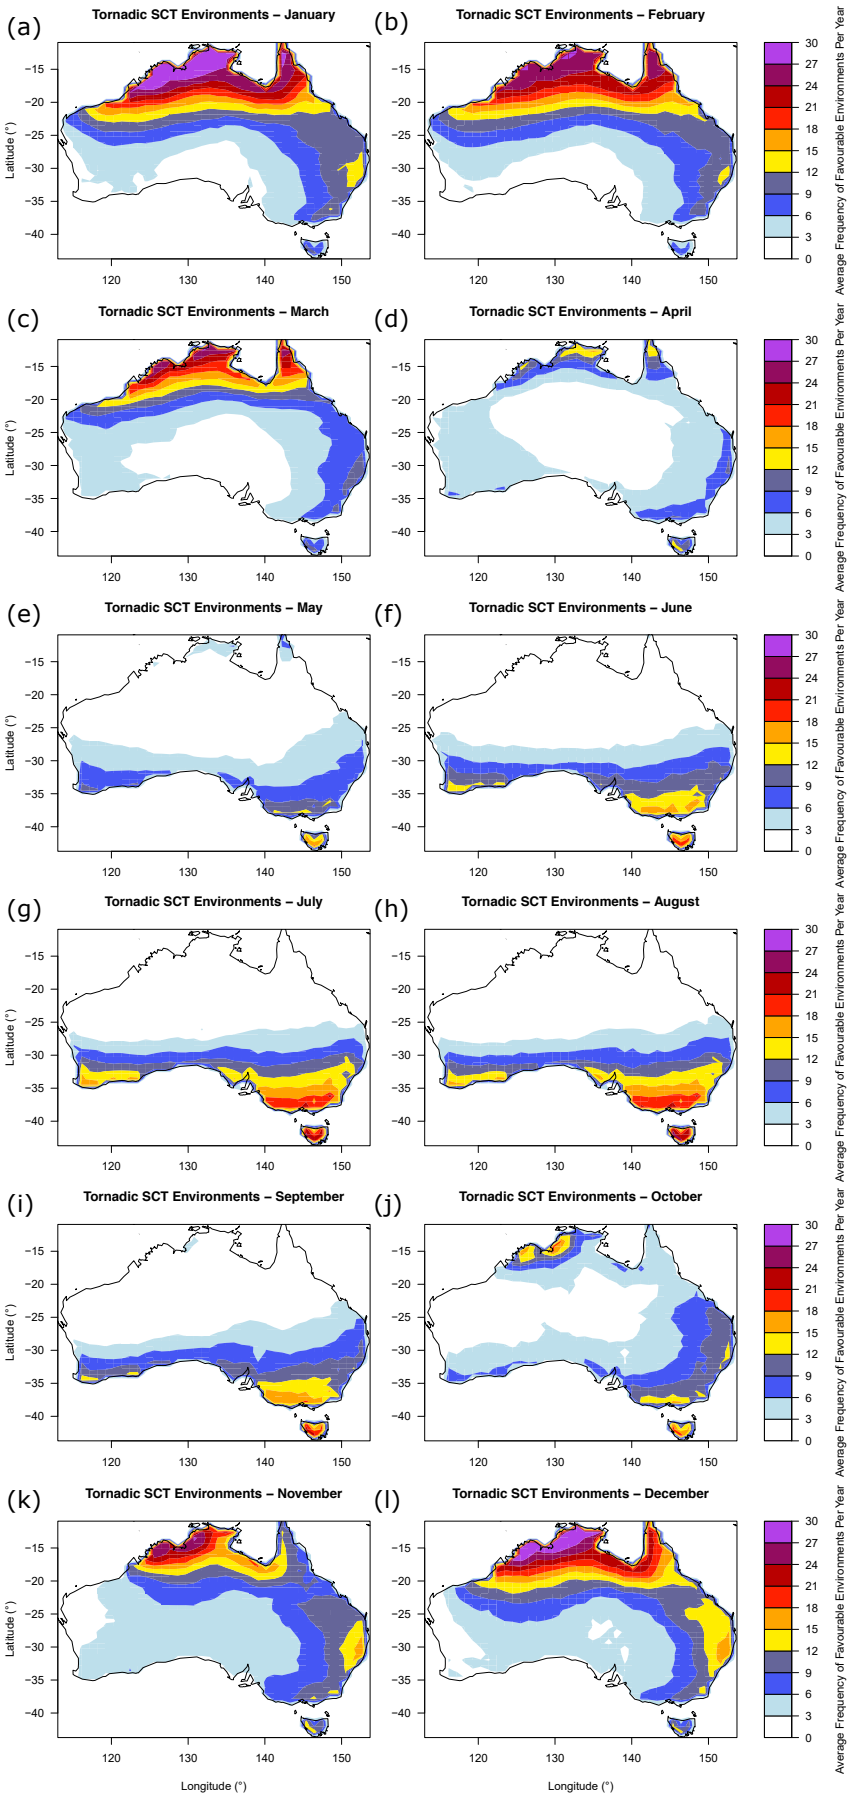

Figure 8.4.17: As in Figure 8.4.15 but for tornadic SCTs using the LogR B model.

Wind

The spatial distribution of environments supportive of wind SCTs appears more widespread over the country, with large parts of central Australia even experiencing a moderate frequency of environments (Figure 8.4.14b). Wind environments appear with the greatest frequency in northern Australia, with a peak in environments along northwest coastal regions, and the peak extending farther inland than the hail and tornadic SCT environments. There is a high frequency of environments over eastern Australia, particularly from the central coast of NSW into southeast and central Queensland. In contrast to both hail and tornadic SCT environments, the frequency of wind SCT environments in southeast and southwest Australia is lower, with only a moderate number of environments occurring in these regions (Figure 8.4.14a, c-d).

Figure 8.4.18 displays the monthly frequency of environments supportive of wind SCTs across Australia. Overall, it appears that the majority of environments occur during the warmer months from October-March, coinciding with the general peak in thunderstorm activity (Dowdy and Kuleshov, 2014). In northern Australia, the high frequency of supportive environments again coincides with the wet season, beginning with a moderate frequency in October over northwest Australia, and extending across the region to a peak frequency in January (Figure 8.4.18k-d). Much of northwest Australia experiences an average of 27–30 days/month where the environment is supportive, from November through to January, before beginning to decrease. For eastern Australia and areas of southeast Australia, the peak in environment frequency occurs over December and January (Figure 8.4.18a,l). There is an extremely low frequency of environments during the cooler months from May-September, with mostly parts of southern Australia experiencing some supportive environments (Figure 8.4.18e-i).

8.4.5 Unbalanced Data Case Study

Although the data set used in this study was balanced to have the same number of severe and non-severe reports, the initial version of Python code to perform the case study had accidentally sampled the data so it was again unbalanced. In this Section, I take a brief look at the results from training the statistical models on an unbalanced data set with reference to the balanced data set. This was performed for both hail and wind SCT environments. Table 8.14 displays the percentage of folds selecting each parcel formulation when developing the models on unbalanced data. Here, the results for hail appear similar to the balanced models (Table 8.3), with the ML parcel being selected most frequently, followed by the SB parcel and finally the MU parcel. Meanwhile, for wind SCTs, the SB parcel is selected most frequently, followed by the ML parcel, then the MU parcel. This is indicative of how different the results likely could end up being from training with unbalanced data.

Table 8.14: Percentage of folds that indices calculated using the surface-based (SB), most-unstable (MU) and mixed-layer (ML) parcel formulations appeared in, when developing the statistical models on an unbalanced data set.

Parcel Type	Hail	Wind
SB Parcel	21.89	25.16
MU Parcel	21.17	18.67
ML Parcel	28.17	21.83

Tables 8.15 and 8.16 provide the contingency statistics, and features and free parameter results, respectively, of the final model fitting step for hail SCT environments on the unbalanced data set. Tables 8.17 and 8.18 provide this for the wind SCT environments on the unbalanced

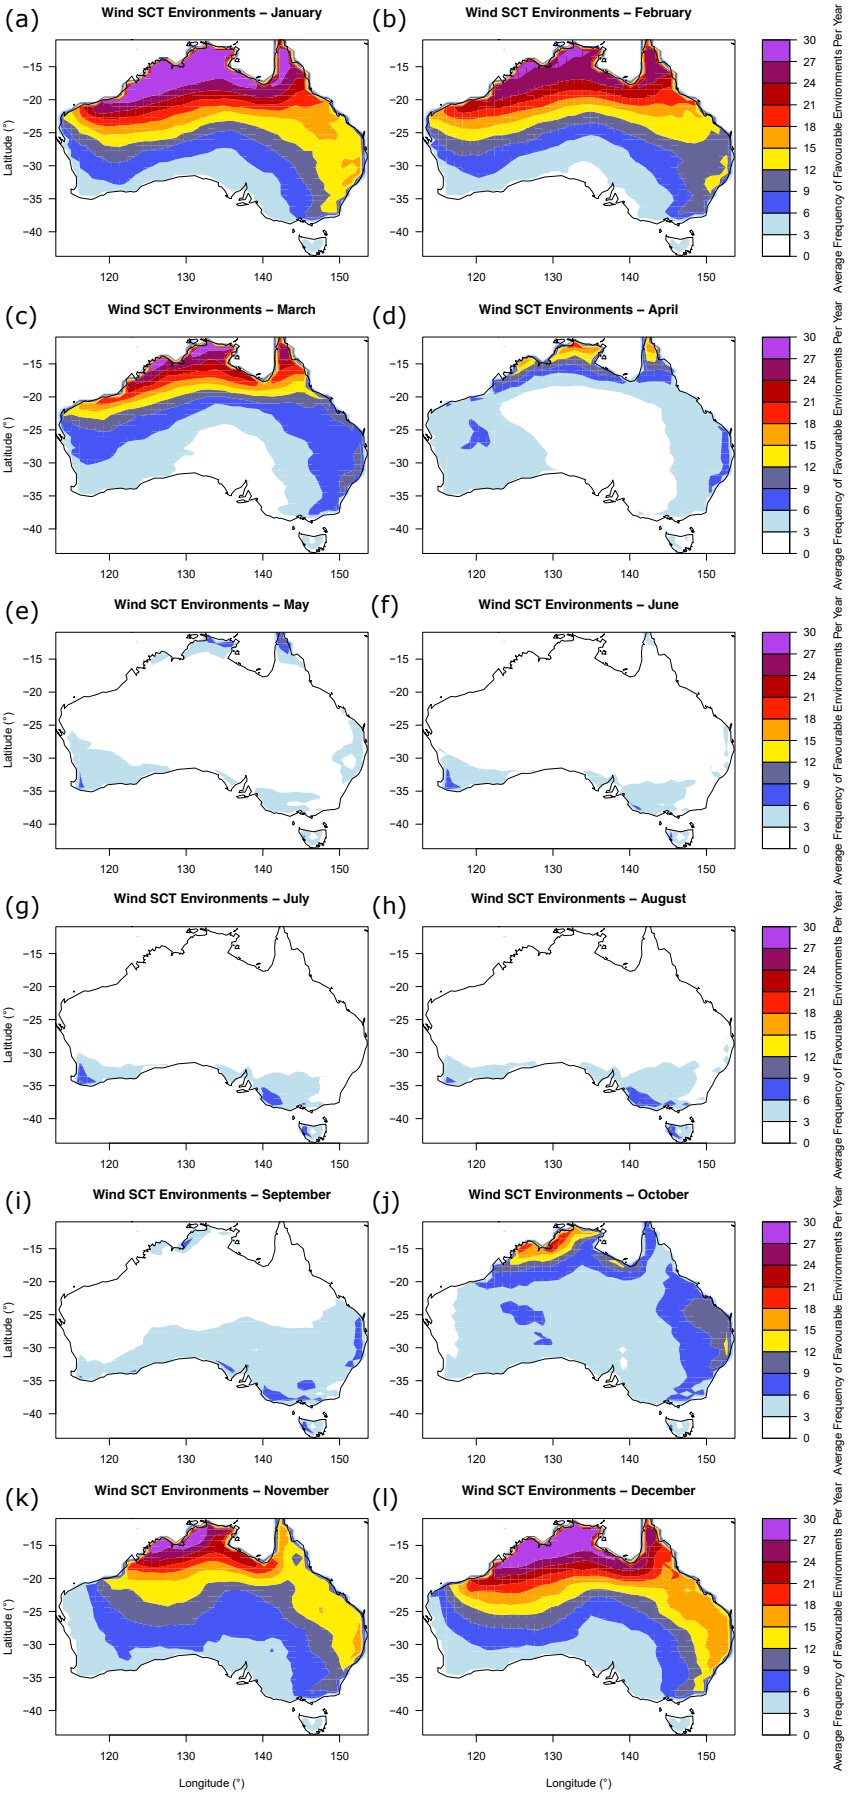

Figure 8.4.18: As in Figure 8.4.15 but for wind SCTs.

Table 8.15: As in Table 8.7 but for hail SCTs developed with unbalanced data.

Statistic	**LogR F**	**SVM RBF F**	**SVM Poly F**	**RF F**	**LogR B**	**SVM RBF B**	**SVM Poly B**	**RF B**
Scitkit-Learn Accuracy	0.7979	**0.7996**	0.7836	0.7821	0.7939	0.7891	0.7879	0.7848
Proportion Correct	0.7512	0.7493	0.7157	0.7233	0.7474	0.7541	0.7089	**0.7666**
POD	0.6449	0.6123	**0.7174**	0.5978	0.6486	0.5543	0.6812	0.6304
FAR	0.4749	0.4768	0.5240	0.5175	0.4812	0.4650	0.5335	**0.4476**
PoFD	0.2105	0.2013	0.2850	0.2314	0.2170	**0.1739**	0.2810	0.1843
CSI	0.4073	0.3930	0.4008	0.3642	0.4050	0.3741	0.3829	**0.4173**
HSS	0.4050	0.3898	0.3721	0.3404	0.3996	0.3761	0.3481	**0.4268**
GSS	0.2539	0.2421	0.2286	0.2051	0.2497	0.2316	0.2113	**0.2713**
Bias Ratio	1.2283	1.1703	1.5072	1.2391	1.2500	**1.0362**	1.4601	1.1413
Odds Ratio Skill Score	0.7441	0.7248	0.7286	0.6632	0.7389	0.7106	0.6906	**0.7661**

Table 8.16: Features and tuning parameters selected for the best performing models presented in Table 8.15.

Model	Features	Parameters
LogR F	MLLFC, K-Index, PW, FrzLevPres, MLCAPE*FrzLevPres	C = $2^{0.5}$
SVM RBF F	MLLFC, SHIP, FrzLevPres, MLEHI3km, MLCAPE*FrzLevPres	C = 2, G = 0.25
SVM Poly F	MLLFC, K-Index, EffSRH, FrzLevPres, MLCAPE*SFC6kmSHEAR	C = 2, E = 3, G = 0.75
RF F	MLLFC, STP, MLEHI3km, MLEHIEff, MLCAPE*SFC1kmSHEAR	ntree = 700
LogR B	MLLFC, K-Index, FrzLevPres, MLCAPE*FrzLevPres	C = $2^{0.5}$
SVM RBF B	FrzLevPres, MLCAPE*MLLCL	C = $2^{0.5}$, G = 0.5
SVM Poly B	MLLFC, EffSRH, FrzLevPres, MLCAPE*SFC6kmSHEAR, MLCAPE*FrzLevPres	C = 2, E = 3, G = 0.75
RF B	MLCAPE, PW, FrzLevPres	ntree = 250

Table 8.17: As in Table 8.7 but for wind SCTs developed with unbalanced data.

Statistic	**LogR F**	**SVM RBF F**	**SVM Poly F**	**RF F**	**LogR B**	**SVM RBF B**	**SVM Poly B**	**RF B**
Scitkit-Learn Accuracy	0.8518	0.8602	0.8527	**0.8836**	0.8518	0.8660	0.8527	0.8741
Proportion Correct	0.7192	0.6960	0.6751	**0.7838**	0.7192	0.7156	0.6751	0.7711
POD	1	1	1	**0.9623**	1	1	1	1
FAR	0.8215	0.8322	0.8413	**0.7839**	0.8215	0.8227	0.8413	0.7888
PoFD	0.3004	0.3239	0.3461	**0.2278**	0.3004	0.3030	0.3461	0.2438
CSI	0.1785	0.1677	0.1587	**0.2143**	0.1785	0.1773	0.1587	0.2112
HSS	0.2220	0.2037	0.1880	**0.2810**	0.2220	0.2199	0.1880	0.27537
GSS	0.1248	0.1134	0.1038	**0.1635**	0.1248	0.1236	0.1038	0.1597
Bias Ratio	5.6038	5.9623	6.3019	**4.4528**	5.6038	5.6415	6.3019	4.7358
Odds Ratio Skill Score	1	1	1	**0.9771**	1	1	1	1

Table 8.18: Features and tuning parameters selected for the best performing models presented in Table 8.17.

Model	Features	Parameters
LogR F	SBEL, SWEAT, FrzLevPres	C = $2^{0.5}$
SVM RBF F	SBLCL, SBEL, SWEAT, FrzLevPres	C = 4, G = 0.5
SVM Poly F	SBLCL, SBEL, SWEAT, FrzLevPres	C = $2^{0.5}$, E = 1, G = 0.25
RF F	SBLCL, SBEL, SWEAT, SFC1kmShear, SHIP	ntree = 500
LogR B	SBEL, SWEAT, FrzLevPres	C = $2^{0.5}$
SVM RBF B	SBCIN, SBLCL, SBEL, SWEAT, FrzLevPres	C = $2^{2.5}$, G = 0.75
SVM Poly B	SBLCL, SBEL, SWEAT, FrzLevPres	C = $2^{0.5}$, E = 1, G = 0.25
RF B	SBEL, TT, PW, SWEAT, EffSRH	ntree = 900

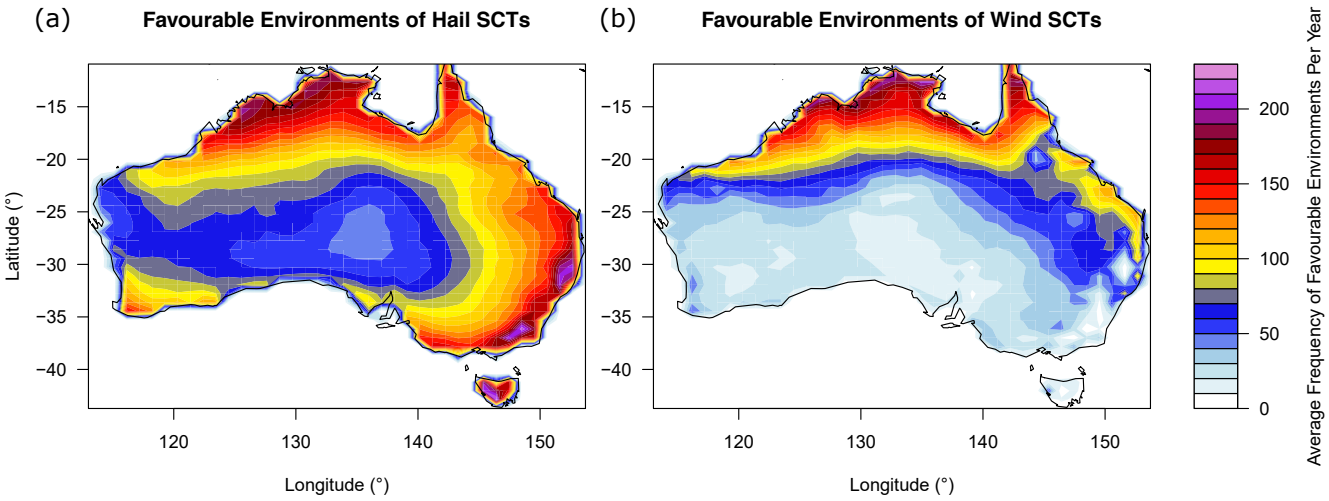

Figure 8.4.19: Average frequency of environments supportive of SCT hazards for (a) hail, and (b) wind with from the best performing models developed using the unbalanced data set.

data set. For hail SCT environments, the models fit on the unbalanced data perform worse overall with significantly lower POD, higher FAR and lower CSI and GSS. For example, the best performing models for each of these statistics obtain a POD = 0.7174 (SVM Poly F), FAR = 0.4476 (RF B), CSI = 0.4173 (RF B), and HSS = 0.4268 (RF B; Table 8.15), whereas the best performing models for each of these statistics on the balanced data obtain a POD = 0.8868 (SVM Poly F), FAR = 0.1966 (SVM RBF F and SVM RBF B), CSI = 0.6906 (SVM RBF B), and HSS = 0.6276 (SVM RBF B; Table 8.7). For the unbalanced models, the most frequently selected features were MLLFC, FrzLevPres, MLCAPE*FrzLevPres and the K-Index (Table 8.16). Of these features, only FrzLevPres was also selected frequently among the models developed on a balanced data set (Table 8.8). The remaining features that were frequently selected in the unbalanced models were not often selected in the balanced models, and vice versa. The best performing model on the unbalanced data is the RF B model, whereas for the balanced data the best performing model is the SVM RBF F model.

There is an interesting feature in the unbalanced wind SCT models, with both POD and the odds ratio skill score equal to one for all but the RF F model (Table 8.17). While a POD equal to one is desirable, which is indicative of no false negatives, this is only when it is accompanied with a low FAR. This is not the case for the unbalanced models, with the best performing FAR still very high at 0.7838. The models display some skill, but it is low, with a CSI of 0.2143 and HSS of 0.2810 for the best performing RF F model. Additionally, many of the unbalanced wind SCT models selected a similar combination of features (Table 8.18), with SBEL, SBLCL, SWEAT and FrzLevPres all selected most frequently among the models. Unlike the hail SCT models, there are common features (despite a different parcel calculation used) between the models on balanced data and those on unbalanced data, with all the features selected in the unbalanced data also frequently selected in the balanced data models. However, the frequency of selection by the balanced models is reduced, allowing both MLLI and SHERB to be among the frequently selected features there, while these features are not frequently selected by the models developed on unbalanced training data.

Figure 8.4.19 provides the spatial distribution of average annual environments favourable for hail and wind SCTs from the best performing models on the unbalanced data. For the hail SCTs, there appears to be a significant increase in the frequency of environments across Australia compared to the balanced hail SCT model (Figures 8.4.14a and 8.4.19a). The peak

frequency regions in far southeast Australia and the central/northeast coast of NSW are still depicted in the unbalanced hail model, however the frequency is much greater. Additionally, there is a significant increase in frequency in Tasmania, with a peak over the west of Tasmania, compared to the peak over the east in the balanced model (Figure 8.4.14a). There also is an increase in frequency of hail SCT environments for southwest Australia, and for inland Australia extending to the SA coast compared to the best performing balanced model (Figure 8.4.19a). The increased frequency captured in the unbalanced model is likely a result of the much increased FAR, despite a decrease in the POD (Table 8.15).

For the wind SCT environments, there is little difference between the balanced and unbalanced model across northern Australia, except for a higher frequency of environments over the Cape York Peninsula (Figure 8.4.14b and 8.4.19b). However, the rest of Australia experiences a significant reduction in favourable wind environments with a low frequency of environments over eastern Australia, especially south of Queensland, and extremely low frequency over the remainder of the continent (Figure 8.4.19b).

8.4.6 Trends in the BARRA Climatology

The BARRA data set runs from 1990 through to the end of February 2019. This provides 29 complete years of data in which potential trends in environments supportive of SCT hazards could be determined. Figures 8.4.20-8.4.23 provide mean differences (from bootstrap resampling) in the frequency of supportive environments per month at each grid point between the periods 1990–2003 and 2004–2018, and significance contours from permutation testing of these differences (indicating where the p-value < 0.1). The differences are calculated annually and for each of the four meteorological seasons. It should be noted that for summer, the year groups of comparison are 1990–2003 and 2004–2017, as the early-2019 BARRA data was not analysed in this study. While examining these figures, it is important to keep in mind that statistical significance could be in part related to false positives. This would more likely be the case for locations where the significance contours take only a few grid cells, and are not part of a broader region of significant change. Given the relatively short length of the time series, changes in the frequency of environments could be related to changes in decadal climate drivers such as the PDO, or they could be related to climate change. However, longer time series would be required to determine to what extent the changes presented here are due to climate variability or change.

Hail

Trends in the annual frequency of hail environments are particularly apparent over the east Australian coastline covering much of NSW up to far southeast Queensland, where there is a statistically significant increase by up to 1–2 environments per month (Figure 8.4.20a). There also is an increasing trend in environments over the far northwest Australian coastline, with a slightly higher magnitude compared to eastern Australia of just over 2 environments per month. Meanwhile, much of the tropical north displays a decrease in the frequency of environments. There are statistically significant reductions of up to 4 environments per month in parts of the Northern Territory, and up to 2 environments per month in the Cape York Peninsula. The remainder of the country experiences relatively small changes, often not reaching statistical significance, with some indication of a potential decrease in frequency of environments over southwestern Australia, and an increase in environments for the majority of inland eastern Australia.

Much of the changes that occur during autumn do not display statistical significance, apart

from a decrease in environments over much of WA (Figure 8.4.20b). This is particularly the case for the central west Australian coastline, extending inland, where there is a decrease of up to 4 environments per month. Much of central Australia experiences an increase of up to 2 environments per month, although this largely is not statistically significant. There appears to be a similar increase in environments along the east coast of Australia, with little statistical significance. Meanwhile, during winter there appears to be a reduction of just over 2 environments per month across much of WA, particularly the southwest and southern coastline (Figure 8.4.20c).

Northwest Australia experiences an increase in favourable environments of up to 6 environments per month, to statistical significance, during spring (Figure 8.4.20d). This region of increased environmental frequency extends south towards central parts of WA, with an increase of 2–4 environments per month. Similarly, eastern Australia displays statistically significant increases of 2–4 environments per month from the southern NSW coastline, towards northeast NSW, and inland to the Great Dividing Range. There also is an increase in 1–2 environments for inland areas of Queensland, although this mostly does not display statistical significance. Parts of northern Australia display some decrease in the frequency of environments, with little statistical significance. Meanwhile much of southern Australia displays a slight decrease in the frequency of supportive environments or little change. However, there appears to be a statistically significant reduction of up to 2 environments per month in Tasmania.

During summer, there is a large region of northern Australia that displays a reduction of generally between 4 and 6 environments per month (Figure 8.4.20e). This region includes the Cape York Peninsula and extends to part of central Australia. The far northwest coastline exhibits a slight increase of up to 2 environments per month, although this generally is not statistically significant. Parts of WA and the SA coastline also display increases of up to 2 environments per month, with some statistical significance. Meanwhile, much of eastern Australia shows an increase of 1–3 environments per month, particularly around the southeast, where there is statistical significance. There also is an increase in the frequency of environments in Tasmania, in contrast to the decrease exhibited during spring (Figure 8.4.20d).

Tornado

Annually, the trends in frequency of environments for both the SVM Poly F and LogR B models appear relatively consistent (Figure 8.4.21a and 8.4.22a). There is an increase in the frequency of environments in the northwest of Australia particularly near the coast, that appears statistically significant. This increase in frequency varies across the region, with some areas experiencing an increase of 1–2 environments per month, and others experiencing an increase of 2–4 environments per month (particularly for the SVM Poly F model; Figure 8.4.21a). Along the east coast of Australia, from the southeast of NSW and far northeast Victoria through to northeast NSW and southeast Queensland, there appears to be a statistically significant increase of approximately 1–2 environments per month. The increase in environments covers much of the Queensland coast, though is largely not statistically significant there, and extends to parts of western Queensland and the eastern Northern Territory, and the coastline surrounding part of the Gulf of Carpentaria. There is some statistical significance in this region, with an increase of 2–3 environments per month in parts. Along much of southern Australia, there is a weak decrease that is not statistically significant, except for the southwest of Australia, where the decrease of up to 2 environments per month is statistically significant.

Autumn with the SVM Poly F model displays an increase in environments across much of

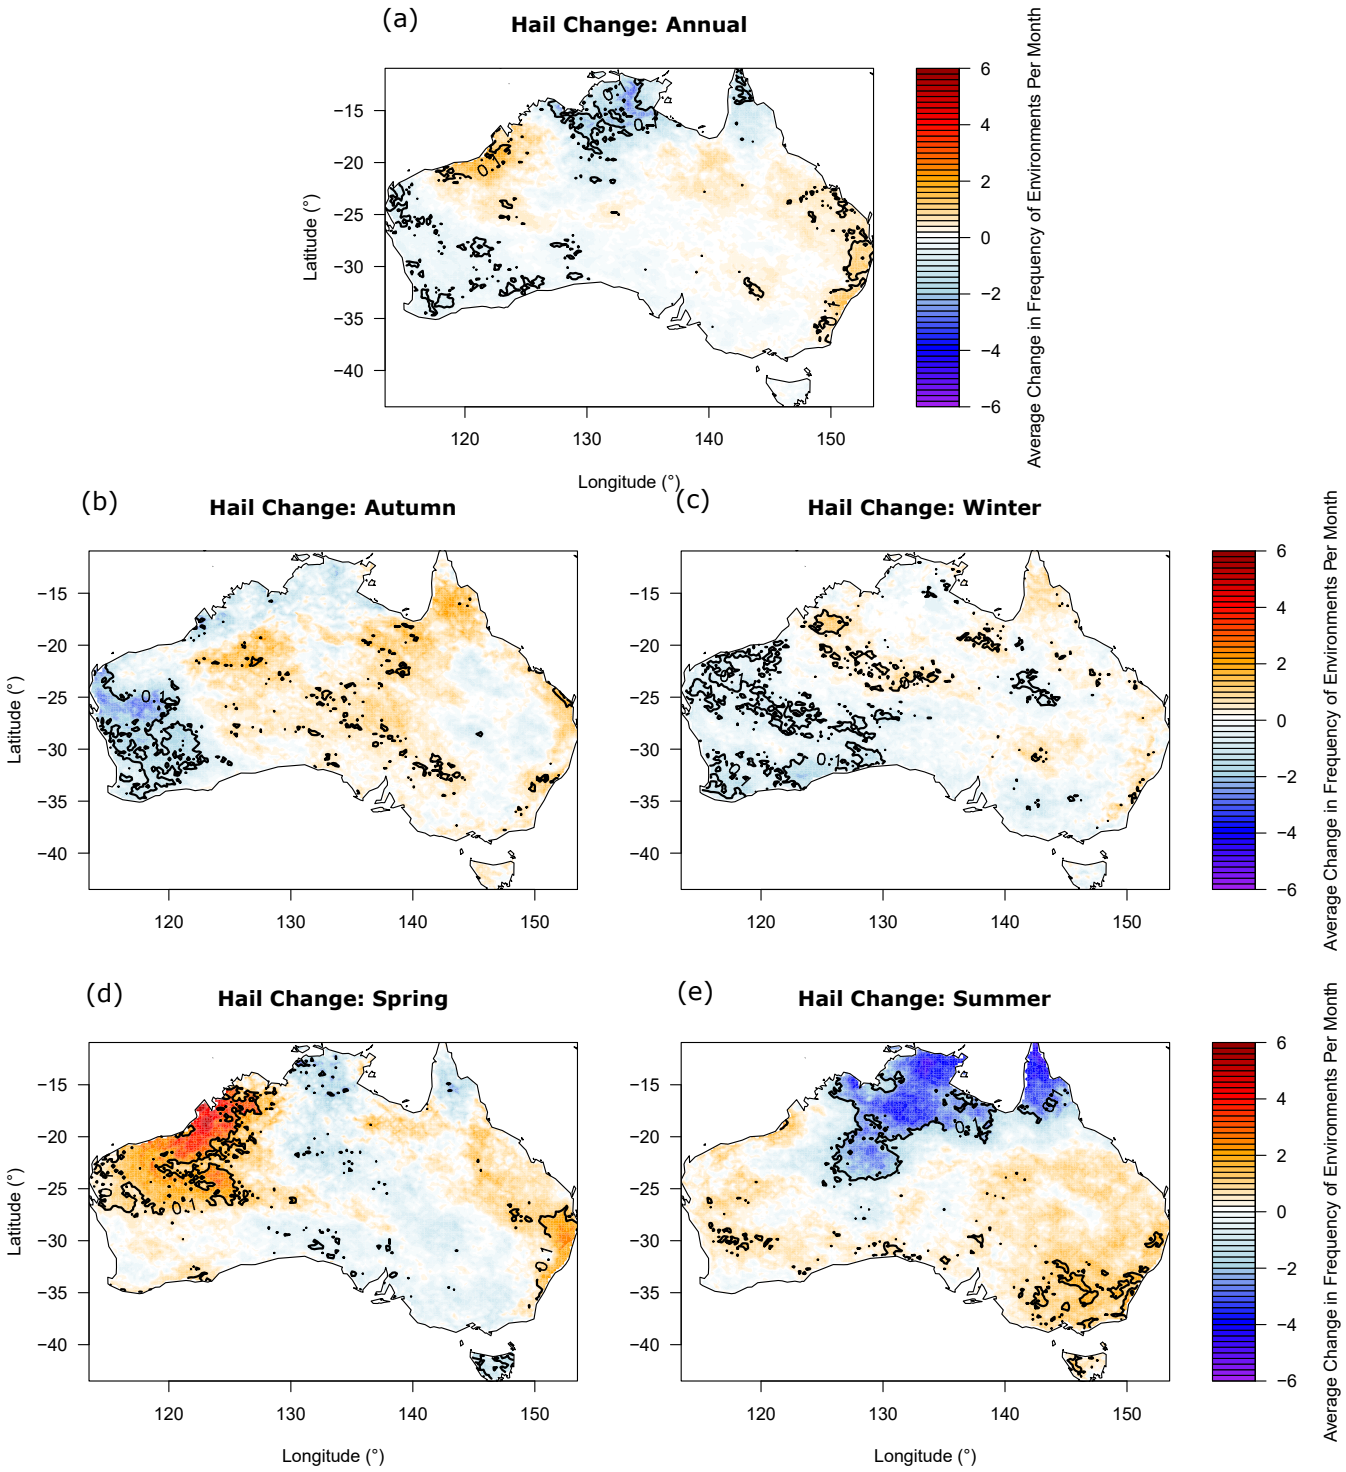

Figure 8.4.20: Filled contours display average change in frequency of environments supportive of hail SCTs between 1990–2003 and 2004–2018 using bootstrap resampling. Change is represented as a monthly average number of days (a) annually, and over (b) autumn, (c) winter, (d) spring, and (e) summer. The thick black contour displays statistically significant changes from permutation testing where the p-value < 0.1.

central Australia, with up to 4–6 more environments per month in some areas of northern central Australia, to statistical significance (Figure 8.4.21b). The LogR B model shows similar changes in the frequency of environments supportive of tornadic SCTs, but with a lower magnitude (Figure 8.4.22b). In particular, the LogR B model exhibits small areas with increases of 4–6 environments per month contained within a broader region where there is a change in frequency of about 2–4 environments per month. Meanwhile, WA displays a statistically significant decrease from the coast, extending inland. This decrease appears to vary from approximately 2 environments per month in the south, and up to 3 environments per month in the north for the SVM Poly F model (Figure 8.4.21b). Meanwhile, the decrease in environments over WA appears somewhat weaker in the LogR B model, particularly in the north of the significance region (Figure 8.4.22b). There also is a statistically significant increase in the frequency of environments classified with the LogR B model over southern central Australia that was not detected in the SVM Poly F model (Figure 8.4.21b), where a large region displays a statistically significant increase of up to 2 environments per month (Figure 8.4.22b). Finally, there appears to be an increase of up to 2 environments per month in parts of Victoria, particulary around the southwest of the region with statistical significance in the LogR B model, while there is little change in the SVM Poly F model (Figure 8.4.21b & 8.4.22b). The increase in environments favourable to tornadic SCTs extends along parts of the east coast of Australia, with no statistical significance (Figure 8.4.22b).

During winter, there is little change over much of Australia for environments classified using the SVM Poly F model (Figure 8.4.21c). There is a slight increase of 1–2 tornadic SCT environments over parts of the central east coast of Australia and adjacent inland, that displays some statistical significance. Over the west of Australia, there largely appears to be a decrease in tornadic environments, with a region of statistical significance over the southwest displaying a decrease of approximately 3 environments per month at its peak. The LogR B model displays similar results, with little change over much of the country, but a much larger region exhibits a statistically significant decrease in favourable environments over WA (Figure 8.4.22c). This decrease in favourable environments again is greatest over the southwest with a decrease of approximately 3–4 environments over large portions of the region. The reduced frequency of environments over southwest Australia in the cooler months of the year, particularly during winter, is possibly due to the reduction in strong cold frontal systems passing through the region, which has also caused an observed decrease in precipitation (Hope et al., 2006; Raut et al., 2014; Richman and Leslie, 2020).

Meanwhile during spring, there is a large area of northwestern Australia, including the interior, that displays a statistically significant increase in tornadic SCT environments for the SVM Poly F model (Figure 8.4.21d). The increase in environments is greatest over the far northwest coast, extending slightly inland, with up to 6 more environments per month. There also appears to be an increase of 1–4 environments over part of the central east coast of NSW, extending into northeast NSW. On the other hand, parts of Victoria and SA exhibit a statistically significant decrease in frequency of favourable environments with a reduction of up to 2 environments per month. This reduction in environments extends to Tasmania, with the north coast displaying a statistically significant decrease. These changes are largely consistent in the LogR B model (Figure 8.4.22d). The change in frequency of environments appears larger over northwestern Australia when the LogR B model is used to classify environments supportive of tornadic SCTs. Meanwhile, the change along the central east coast of NSW has reduced slightly in extent and magnitude, with an increase of only 1–2 environments per month. Additionally, the decreasing trend in environments seen in parts of Victoria and SA appear to extend somewhat further inland, with a slight reduction in environments apparent in central

Australia. The region of statistical significance in Tasmania also has shifted from the north coast to the west coast.

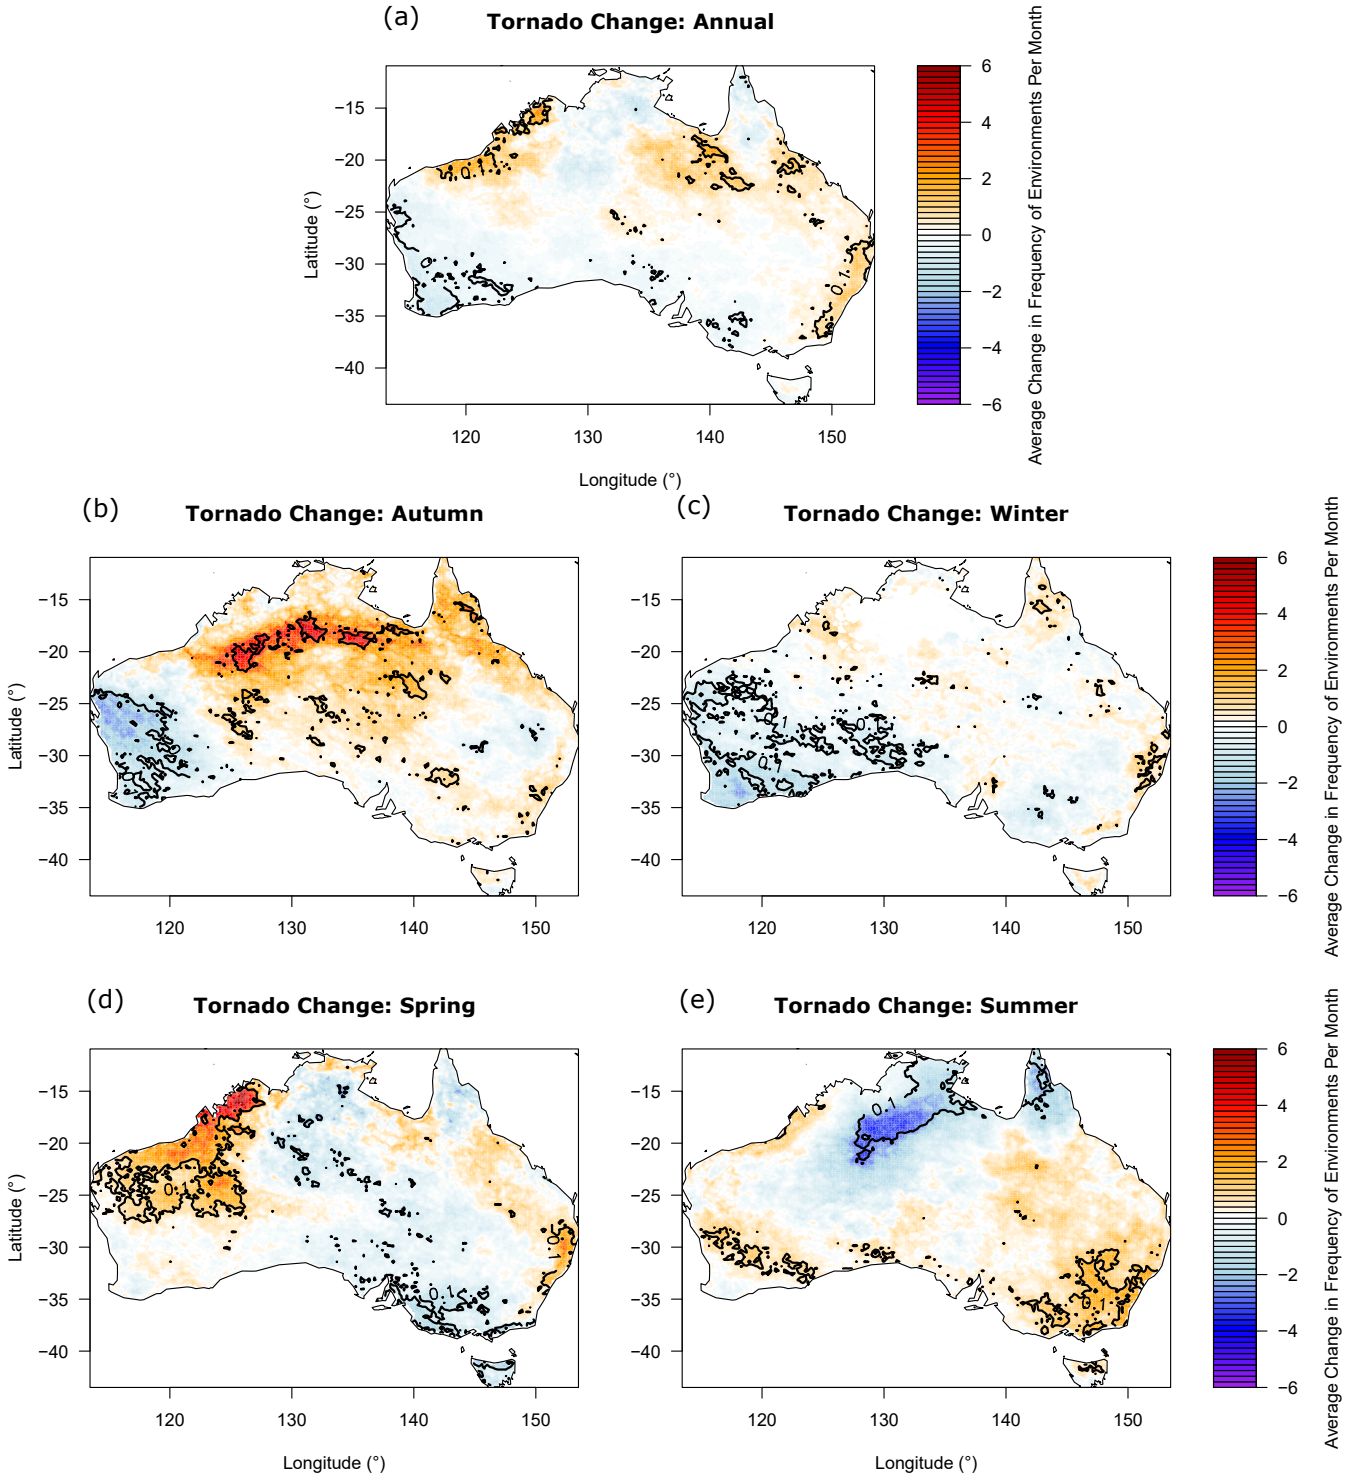

Figure 8.4.21: As in Figure 8.4.20 but for environments supportive of tornadic SCTs using the SVM Poly F model.

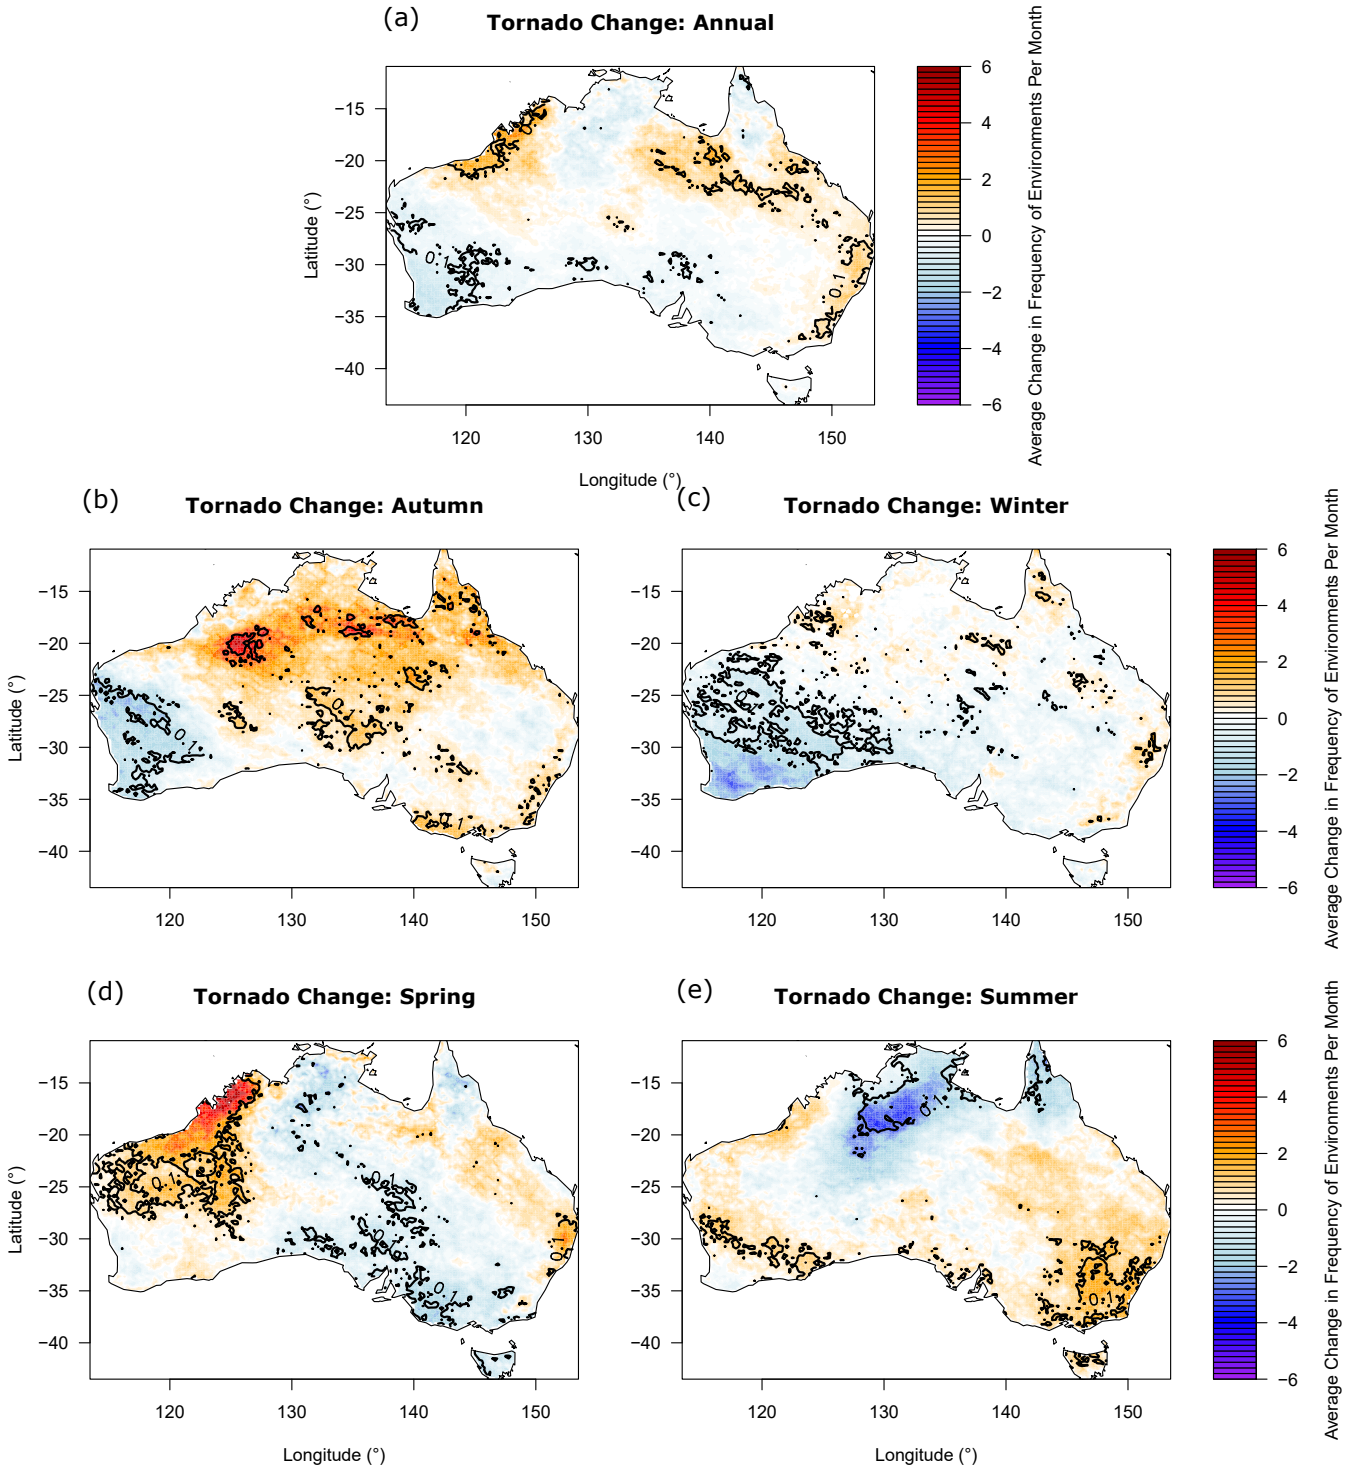

Figure 8.4.22: As in Figure 8.4.20 but for environments supportive of tornadic SCTs using the LogR B model.

Much of eastern Australia displays an increase in the frequency of environments supportive of tornadic SCTs during summer, with a statistically significant increase over southeastern Australia of approximately 2 environments per month for the SVM Poly F model (Figure 8.4.21e). This increasing trend also appears over northern Tasmania. Meanwhile, there is an increasing trend apparent in southwest Australia, except for the far southwestern corner, of between 1–2 environments per month. In contrast, northern Australia displays a clear decreasing trend, with a statistically significant decrease for much of the Northern Territory, particularly for inland regions. In the region around the west of the Northern Territory, the decrease in environments is in the range of 3–5 environments per month. Additionally, there is a decrease in

supportive environments over the Cape York Peninsula, particularly along the west coast, with a statistically significant decrease of 2–3 environments per month. For the LogR B model, the change in environments is qualitatively similar (Figure 8.4.22e). However, the decrease over the Cape York Peninsula is weaker, though still statistically significant (Figure 8.4.22e). The change in environments over Tasmania also appears different, with an increase in frequency of 1–2 supportive environments per month across much of the state rather than just the north of the state, though the statistically significant trend remains only over the north of Tasmania.

Wind

Much of Australia displays an increase in the annual frequency of environments supportive of wind SCTs between the two periods (Figure 8.4.23a). Of greatest note is the increase of 2–4 environments per month over part of northwest Australia, extending inland in a south-southeasterly direction to part of the southern coastline. However, the change in frequency along the southern coastline is much smaller, with an increase of up to 1 environment per month. Southeast Australia also displays a statistically significant increase, with inland regions and the coastline experiencing an increase of 1–2 environments per month. Additionally, there are increases of 1–2 environments per month for inland Queensland towards the coast of the Gulf of Carpentaria, and the central east Queensland coast.

There again appears to largely be an increase in the frequency of environments during autumn (Figure 8.4.23b). This increase is greatest in a region over central northern Australia, and part of the Gulf of Carpentaria coastline, in a somewhat similar pattern as the change in tornadic SCT environments as diagnosed by the SVM Poly F model (Figure 8.4.21b). There also appears to be a statistically significant increase of 1–2 environments over a region surrounding northwest Victoria, southwest NSW and far east SA (Figure 8.4.23b). Similarly, parts of the east NSW coastline, far southeast corner of Australia and Tasmania display statistically significant increases of 1–2 environments per month during autumn. Meanwhile, there is a statistically significant reduction in environments over much of WA, with 2–4 less environments for part of the inland and central west coast.

During winter, there is some indication of a reduced frequency of environments, by up to 2 environments per month, in parts of WA, particularly around the southern coast (Figure 8.4.23d). Additionally, there appears to be a region around northeast NSW and southeast Queensland that experiences a slight increase of 1–2 environments per month. However, overall there are few changes in the frequency of environments supportive of wind SCTs during winter.

Spring displays an increase in the frequency of environments over a large portion of WA with statistical significance, excluding the far southwest corner (Figure 8.4.23d). In the area around the northwest coast, there is an increase of 4–5 environments per month, whereas the central west coast and south coast exhibit an increase of 1–2 environments per month. Additionally, parts of inland WA display an increase of 3–4 environments per month. Much of central Australia appears to experience a slight decrease in the frequency of environments, though this mostly is not statistically significant. There also appears to be a statistically significant increase of 1–3 supportive environments per month along the southeast Australian coastline. Much of Queensland displays an increase of 1–2 environments per month, excluding the north into the Cape York Peninsula, although this mostly is not statistically significant.

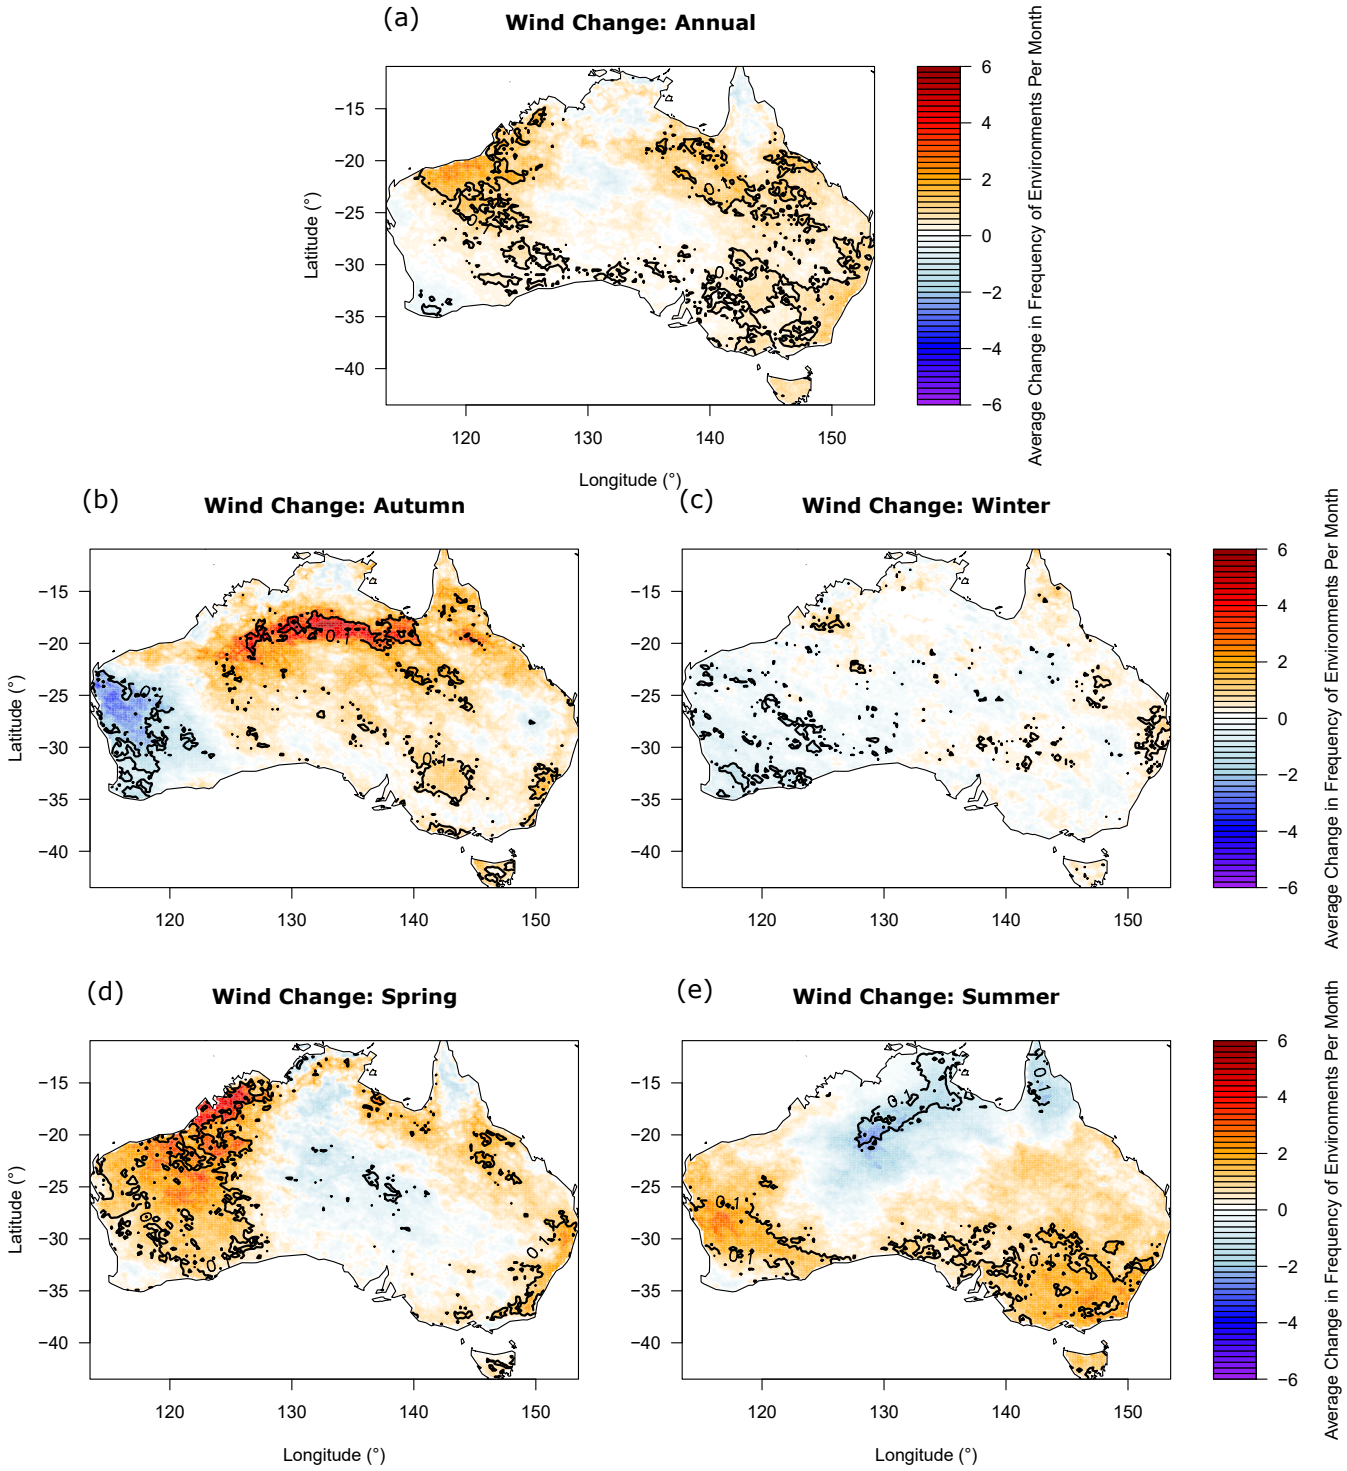

Figure 8.4.23: As in Figure 8.4.20 but for environments supportive of wind SCTs.

During summer, there is an increase of 1–3 environments per month for part of southwest Australia, excluding the far southwest coast (Figure 8.4.23e). There also is an increase of 1–2 environments per month extending to the central west coast, although this is not statistically significant. Similarly, much of southern Australia experiences an increase of 1–3 environments per month, with the greatest increase in frequency apparent over southeastern Australia. There also is an increase of 1–2 environments per month over much of Tasmania. There is a broader area displaying an increase of 1–2 environments per month over eastern Australia, although this is not statistically significant for more northern parts of this region. In contrast, the tropics display a decrease in the frequency of supportive environments during summer. This decrease is statistically significant for much of the Cape York Peninsula, with a reduction of 1–2

environments per month. Meanwhile, a region from the eastern Northern Territory, extending inland to the southwest experiences a statistically significant reduction of 1–3 environments per month.

8.5 Discussion and Conclusions

This study assessed the climatology of Australian SCT hazards through a combined data set comprising STA reports, lightning data, and weather radar data used to determine severe, likely non-severe and no storm events around Australia. Pseudo-proximity soundings were developed nearest to these reports using the BoM's BARRA reanalysis data set to obtain a set of 62 potential indicators of SCTs and their associated hazards. Statistical models were then developed, with the best model for each SCT hazard (based on contingency statistics and analysis of ROC and PR curves) applied to the BARRA reanalysis data set from 1990–2018 to determine the average frequency of environments that support each convective hazard across Australia.

The ability of BARRA pseudo-proximity soundings to reproduce the environment was assessed by comparing these soundings with atmospheric radiosondes from the University of Wyoming archive. CAPE calculated using BARRA data tended to be overestimated, particularly in coastal regions, as compared to observed soundings. However, CAPE calculated using the mixed-layer parcel appeared to be the most closely aligned to the radiosondes, possibly due to the averaging reducing variability in the calculation. This is likely why the mixed-layer parcel was most frequently selected, among the three parcel calculations tested, by the statistical models developed for each SCT hazard.

As previous studies have found both in Australia and overseas, reports of SCTs in the STA were most clustered around population centres, especially the east coast of Australia (Brooks et al., 2003; Allen et al., 2011). Despite this, there are still some notable features that are likely physical. Firstly, eastern Australia reports the greatest frequency of hail and wind SCTs, which is expected based on population density and other climatology studies (Brooks et al., 2003; Allen and Karoly, 2014; Allen and Allen, 2016b; Bedka et al., 2018). There generally is a high frequency of supportive environments due to instability associated with low pressure troughs developing on the Great Dividing Range, moisture and instability fuelled by the warm East Australian Current, and the passage of upper level systems that produce favourable wind shear profiles (Leslie et al., 2008; Allen and Karoly, 2014). The peak in frequency of hail and wind SCT reports occurs during late spring and early summer. Meanwhile, southern Australia experiences its peak in tornado occurrence during the cooler months of the year which has been shown in past studies (Hanstrum et al., 2002; Kounkou et al., 2009), whereas eastern Australia experiences a peak in tornado occurrence during the warm season (September-March; Evesson, 1970).

Overall, the statistical models developed for each hazard performed rather well with a high POD, low FAR and relatively high skill scores. This was the case across all models and attribute selection methods considered, not just the best performing model. The models developed also performed particularly well compared to the discriminant relationship in Allen et al. (2011), which is mostly hampered by the higher FAR on the test data set used in this study. There was relatively good agreement across the models in the attributes selected for each hazard. For hail SCTs, the most common attributes included the height of the freezing level and EL, effective SRH and a two-way interaction between MLCAPE and wind shear. Freezing level height influences hail size, by affecting when hail growth can begin within a thunderstorm updraft, and modulating the amount of melting that occurs before a hail stone reaches the

ground (Dessens et al., 2015; Bedka et al., 2018; Prein and Holland, 2018). Severe convection commonly develops over a significant vertical depth, which can be approximated by the EL height (Bedka et al., 2018). The EL could also relate to the depth of the optimal hail growth layer (Allen et al., 2020). However, a lower EL can mean that CAPE is concentrated over a smaller region (Blanchard, 1998; McCaul and Weisman, 2001), leading to more intense updrafts for the same amount of CAPE (given all else is equal). This would allow hail to grow larger before the terminal fall velocity exceeds the updraft velocity. The interaction between MLCAPE and wind shear is generally associated with storm severity (Brooks et al., 2003; Brooks, 2009; Allen et al., 2011) and effective SRH is associated with the likelihood of supercell convection that commonly produces large hail (Rasmussen and Blanchard, 1998; Thompson et al., 2002; Rasmussen, 2003; Thompson et al., 2003, 2004a,b).

The favoured attributes for tornadic SCTs include the freezing level height, MLEL, SWEAT and SRH between the surface and 1 km. SWEAT is an index that aims to predict severe thunderstorms and includes both wind shear and instability in its formulation, so likely has some skill in predicting the storm modes that are commonly associated with tornadic storms (Miller, 1972). SWEAT also includes low-level moisture, which has been shown to discriminate tornadic SCTs (Rasmussen and Blanchard, 1998; Craven and Brooks, 2004) and influence the efficiency of stretching near-surface vorticity to tornado strength (Markowski and Richardson, 2014; Brown and Nowotarski, 2019). On the other hand, SRH between the surface and 1 km has been shown to be highly influential on tornadic storms and the intensity of the resulting tornado (Thompson et al., 2002; Rasmussen, 2003; Thompson et al., 2003), with shallower depths potentially providing greater forecast skill (Coffer et al., 2019). The height of the equilibrium level might have been selected as low-topped supercells during winter are one of the common convective modes causing cool-season tornadoes that frequently impact southern Australia (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009). Alternatively, the EL has an effect on updraft strength (Blanchard, 1998; McCaul and Weisman, 2001), which might affect tornado development by influencing the efficiency of stretching low-level vorticity to tornado strength. Prior research has not identified the freezing level as particularly influential on the production of tornadic SCTs. However, the freezing level height could have been selected here due to the impact on microphysical structure of tornadic storms and the resulting outflow influencing low-level vorticity and tornado development (Proctor, 1989; Richter et al., 2014; Dessens et al., 2015; Bedka et al., 2018).

The environmental attributes frequently selected for wind SCTs included MLLI, MLLCL, SWEAT, SHERB, FrzLevPres and MLEL. MLLI is a simple measure of instability that can be used to predict the likelihood of thunderstorms. Meanwhile MLLCL is closely related to the amount of moisture in the lower atmosphere, with higher LCLs associated with lower moisture near the surface that can lead to greater evaporative cooling and the occurrence of microbursts (Wakimoto, 1985; McCaul and Cohen, 2002). As described in the preceding paragraph, SWEAT was developed to predict severe thunderstorms, so this parameter likely predicts environments where storms become organised and produce damaging winds (Miller, 1972). SHERB was developed to detect severe thunderstorm potential in low-CAPE high-shear environments, where a large number of severe wind producing storms occur (Taszarek et al., 2020), and this parameter has shown particular skill in detecting damaging wind storm environments (Sherburn and Parker, 2014). The height of both the freezing level and equilibrium level also were commonly selected. The height of the freezing level has been identified in the past as a potential contributor to downdraft speeds with melting of hail below this level increasing negative buoyancy (Richter et al., 2014). Therefore, a higher level would increase the amount of melting before hail reaches the ground, further increasing negative buoyancy. Additionally, dry

air at the melting level has been shown to increase downdraft speeds (Proctor, 1989; Richter et al., 2014), potentially due to evaporation of particles that have undergone transition from hail to precipitation. The height of the equilibrium level was frequently selected in the LogR models of Brown and Dowdy (2021a) for severe convective wind storms in Australia. It has an influence on how instability is distributed over the depth of a storm, which can influence storm mode, strength of the updraft and the resulting downdraft (Rasmussen and Blanchard, 1998; McCaul and Weisman, 2001). It is interesting that both MLEL and FrzLevPres were selected across all three hazards, and might indicate more important relationships between these two variables and general severe convective weather potential that are worthwhile investigating.

Programming bugs in the early stages of this project resulted in models for favourable hail and wind SCT environments that were developed on unbalanced data. Given the data that was available, an assessment on the importance of developing models with balanced data, where possible, was performed. For the hail SCTs, the POD and FAR were worse, both contribuing to low skill scores. However, the wind SCT models had near perfect POD, likely due to the lack of false negatives. The wind SCT models had very high FAR, causing extremely low skill. Although the resulting models appear to reproduce some aspects of the environmental frequency similar to the balanced models, their high FAR resulted in what appears to be a significant overestimation of the frequency of environments supportive of wind and hail SCTs across Australia.

For the best performing models developed on balanced classification data sets, the overall frequency of favourable environments for each SCT hazard within the BARRA data set appears rather high. The high frequency of occurrence is particularly notable in northern Australia, where the frequency of environments is approaching the number of days that the wet season spans. High frequencies of severe environments have been noted in previous climatologies (Brooks et al., 2003; Allen and Karoly, 2014). In reality, northern Australia experiences lower wind shear, making it less likely to be supportive of severe convection. However, the models selected include an instability component, which often is very high during the northern Australia wet season, particularly over northwest Australia (Allen and Karoly, 2014). Given the models are mostly trained on data over southeastern Australia, it is likely that they are overestimating the frequency of supportive environments in northern Australia due to the very high instability there. Considering this alongside the results when the hail and wind SCT models were developed on unbalanced classification data highlights an important factor for statistical modelling on geospatial data. Given the influence of climate variability and geomorphology differences over a large area, there likely is improved skill when developing statistical models on data that is not only balanced within the classification set, but also balanced geospatially. This additionally highlights the importance of developing regional models, so that regional differences in favourable environments are captured. Furthermore, in the absence of truly spatially balanced data, development of models focused on smaller regions likely will improve skill due to capturing the environmental influences that are important for those particular regions. Although the number of reports will be lower in some regions compared to others, there likely is enough information contained in the reports to obtain physically relevant relationships, improve understanding of the significance of particular variables on severe local storms and better describe the frequency of supportive environments.

Considering the annual environment distribution, there is a high frequency of environments supportive of hail SCTs over eastern Australia, particularly along the central coast of NSW into southeast Queensland, and the far southeast corner, over a similar area identified by Bedka et al. (2018). The peak frequency of tornadic environments covers a similar area over eastern Australia as the hail environments, likely as both hail and tornadic SCTs favour environments

supportive of supercells, that are commonly associated with large hail and tornadoes. Meanwhile, wind SCTs display a relatively high frequency of environments over part of eastern Australia. There is a high frequency of environments over northern Australia for all hazards. This likely is overestimated due to the spatial imabalance in reports, except perhaps for wind SCTs where there are a considerable number of reports in northern Australia from which the model was trained. This higher number of wind reports in the Northern Territory is likely in part due to the occurrence of squall lines that have been associated with the sea-breeze front (Wissmeier et al., 2010). Additionally, due to being in the tropics, northern Australia experiences environments with high levels of moisture, that likely support wet microbursts due to increased precipitation loading within storms (Wakimoto, 2001; Coniglio et al., 2011), without the requirement of high wind shear for the development of more organised convection capable of producing other severe hazards. Finally, there is a relatively high frequency of tornadic SCT environments over southwest Australia. Although the model including SRH1km performed slightly worse on the contingency statistics than the SVM Poly F model, it displays a greater occurrence of environments supportive of tornadic SCTs across southern Australia and appears more representative of cool-season tornadoes that occur in this region (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009).

The frequency of supportive environments throughout the year in the BARRA reanalysis appear similar to what is observed within the STA. There is a high frequency of environments supportive of hail SCTs over eastern Australia from November through to January, with a peak in December. There also appears to be a moderate frequency of environments in southeast and southwest Australia during winter, likely associated with strong low pressure systems including east coast lows. The frequency of tornadic SCT environments displays a peak over southern Australia during winter, and a peak over eastern Australia during spring and summer. These peaks in tornadic environments become more prominent when considering the LogR B model, which includes SRH1km. There is a peak in environments supportive of wind SCTs over eastern Australia during December and January, while northern Australia experiences the peak during January. Additionally, there appears to be a low frequency of environments during the cooler months over southern Australia.

Given the BARRA data set is approximately 28 years long, the data was split into two periods so bootstrap resampling and permutation testing could be applied to visualise changes in the frequency of SCT environments, and determine statistical significance of any changes, respectively. There appears to be an increase in the frequency of hail SCTs over the east coast of Australia and northwest Australia, particularly during spring, while there is a reduced frequency of environments over northern Australia during summer. There also is an increase in the frequency of environments supportive of tornadic SCTs in eastern Australia and northwest Australia during spring, and southeast Australia during summer. Meanwhile, there is a reduction in the frequency of tornadic environments over southwest Australia during autumn and winter, likely due to a reduction in the frequency of cold fronts passing through southern Australia during the cool season (Hope et al., 2006; Raut et al., 2014). These cold fronts have been associated with cool-season tornadoes in southern Australia (Hanstrum et al., 2002; Mills, 2004; Kounkou et al., 2009), and their progression towards the poles with global warming suggests that the frequency of cool-season tornadoes in southern Australia also might decrease as the warming trend continues. Additional analysis of future climate simulations have suggested a reduction in the risk of cool-season tornadoes due to increasing stability in the lower troposphere, and increasing SAM (which is negatively correlated with cool-season tornado risk; Kounkou et al., 2007, 2009). There also appears to be an increase in the frequency of environments supportive of wind SCTs over much of WA and eastern Ausralia during spring, and part

of southwest and southeast Australia during summer. These changes over WA during spring and summer are also reported in the severe wind climatology of Brown and Dowdy (2021b), while the changes over eastern Australia contrast with their results. However, different time periods were analysed, with Brown and Dowdy (2021b) comparing 1979–1998 to 1999–2018 in the ERA5 reanalysis.

Although the changes elucidated here are likely physically relevant, the length of the BARRA data set is not sufficient to determine if these changes are due to interannual variability, or if climate change also is having an influence. There is an expected increase in the frequency of severe hail storms for part of eastern Australia (Leslie et al., 2008). Meanwhile, the frequency of cool-season tornadoes appears likely to decrease (Kounkou et al., 2007), and there is uncertainty in the future change of severe convective wind environments (Brown and Dowdy, 2021b). More work is required to better understand what factors have influenced the changes found in this project. ERA5 is a longer reanalysis data set, extending back to 1950 (Hersbach et al., 2020). A similar climatology could be developed by applying the statistical models presented here to the ERA5 data, with the results compared to those here. Then wavelet analysis and statistical models can be developed to assess the relevant climate drivers associated with each hazard, similar to the methods presented in Chapters 4–6. Finally, the statistical models developed here, or those developed over smaller regions, could be applied to global circulation model (GCM) or regional circulation model (RCM) output, to determine how climate change might influence the frequency of environments that favour each SCT hazard.

Bibliography

Allen, J. Supercell Storms: Melbourne's white Christmas 2011. Bulletin of the Australian Meteorological and Oceanographic Society, 25:47–51, 2012.

Allen, J. and E. Allen. Australian Tornadoes: Climatology 1795-2014 Compared to a 'Record' 2013. 22nd Conference on Applied Climatology. American Meteorological Society, New Orleans, Louisiana, 2016a.

Allen, J. and E. Allen. A review of severe thunderstorms in Australia. Atmos. Res., 178: 347–366, 2016b.

Allen, J. and D. Karoly. A climatology of Australian severe thunderstorm environments of 1979–2011: inter-annual variability and ENSO influence. Int. J. Climatol., 34:81–97, 2014.

Allen, J., D. Karoly, and G. Mills. A severe thunderstorm climatology for Australia and associated thunderstorm environments. Australian Meteorological and Oceanographic Journal, 61:143–158, 2011.

Allen, J., I. Giammanco, M. Kumjian, H. Punge, Q. Zhang, P. Groenemeijer, M. Kunz, and K. Ortega. Understanding Hail in the Earth System. Rev. Geophys., 58:1–49, 2020. e2019RG000665.

Antonescu, B., D. Schultz, A. Holzer, and P. Groenemeijer. Tornadoes in Europe: An Underestimated Threat. Bull. Am. Meteorol. Soc., 98:713–728, 2017.

Bedka, K., J. Allen, H. Punge, M. Kunz, and D. Simanovic. A Long-Term Overshooting Cloud-Top Detection Database over Australia Derived from MTSAT Japanese Advanced Meteorological Imager Observations. J. Appl. Meteorol. Climatol., 57:937–951, 2018.

Blanchard, D. Assessing the Vertical Distribution of Convective Available Potential Energy. Weather Forecast., 13:870–877, 1998.

Breiman, L. Random forests. Mach. Learn., 45:5–32, 2001.

Brooks, H. Proximity soundings for severe convection for Europe and the United States from reanalysis data. Atmos. Res., 93:546–553, 2009.

Brooks, H., C. Doswell, III, and J. Cooper. On the Environments of Tornadic and Nontornadic Mesocyclones. Weather Forecast., 9:606–618, 1994a.

Brooks, H., I. Doswell, C.A., and R. Wilhelmson. The Role of Midtropospheric Winds in the Evolution and Maintenance of Low-Level Mesocyclones. Mon. Weather Rev., 122:126–136, 1994b.

Brooks, H., J. Lee, and J. Craven. The spatial distribution of severe thunderstorm and tornado environments from global reanalysis data. Atmos. Res., 67–68:73–94, 2003.

Brown, A. and A. Dowdy. Severe convection-related winds in Australia and their associated environments. J. South. Hemisphere Earth Syst. Sci., 71:30–52, 2021a.

Brown, A. and A. Dowdy. Severe Convective Wind Environments and Future Projected Changes in Australia. J. Geophys. Res. Atmos., 126:1–17, 2021b. e2021JD034633.

Brown, B. and A. Murphy. Verification of aircraft icing forecasts: The use of standard measures and meteorological covariates. In Preprints, 13th Conf. on Probability and Statistics in the Atmospheric Sciences, San Francisco, CA, Amer. Meteor. Soc, pages 251–252, 1996.

Brown, M. and C. Nowotarski. The Influence of Lifting Condensation Level on Low-Level Outflow and Rotation in Simulated Supercell Thunderstorms. J. Atmos. Sci., 76:1349–1372, 2019.

Buckley, B., L. Leslie, and Y. Wang. The Sydney Hailstorm of April 14, 1999: Synoptic description and numerical simulation. Meteorol. Atmospheric Phys., 76:167–182, 2001.

Bureau of Meteorology. Severe Storms Archive, 2012. URL `http://www.bom.gov.au/australia/stormarchive/`.

Bureau of Meteorology. Severe thunderstorm and tornado outbreak South Australia 28 september 2016. Technical report, Bureau of Meteorology, 2016.

Bureau of Meteorology. Severe Thunderstorms, 2017. URL `http://www.bom.gov.au/weather-services/severe-weather-knowledge-centre/severethunder.shtml`.

Cecil, D. and C. Blankenship. Towards a Global Climatology of Hailstorms as Estimated by Satellite Passive Microwave Imagers. J. Clim., 25:687–703, 2012.

Chappell, C. Quasi-Stationary Convective Events. In Mesoscale meteorology and forecasting, pages 289–310. Springer, 1986.

Childs, S., R. Schumacher, and J. Allen. Cold-Season Tornadoes: Climatological and Meteorological Insights. Weather Forecast., 33:671–691, 2018.

Coffer, B., M. Parker, R. Thompson, B. Smith, and R. Jewell. Using Near-Ground Storm Relative Helicity in Supercell Tornado Forecasting. Weather Forecast., 34:1417–1435, 2019.

Concannon, P., H. Brooks, and C. Doswell, III. Climatological Risk of Strong and Violent Tornadoes in the United States. In Preprints, 2nd Symp. on Environmental Applications, Long Beach, CA, Amer. Meteor. Soc, pages 212–219, 2000.

Coniglio, M., S. Corfidi, and J. Kain. Environment and Early Evolution of the 8 May 2009 Derecho-Producing Convective System. Mon. Weather Rev., 139:1083–1102, 2011.

Craven, J. and H. Brooks. Baseline Climatology of Sounding Derived Parameters Associated with Deep Moist Convection. Natl. Weather Dig., 28:13–24, 2004.

Crompton, R. and K. McAneney. Normalised Australian insured losses from meteorological hazards: 1967–2006. Environ. Sci. Policy, 11:371–378, 2008.

Dessens, J., C. Berthet, and J. Sanchez. Change in hailstone size distributions with an increase in the melting level height. Atmos. Res., 158–159:245–253, 2015.

Doswell, C. Severe Convective Storms—An Overview. In Severe Convective Storms, volume 50, pages 1–26. Springer, 2001.

Doswell, C., III. Synoptic-Scale Environments Associated with High Plains Severe Thunderstorms. Bull. Am. Meteorol. Soc., 61:1388–1400, 1980.

Doswell, I., C.A. and E. Rasmussen. The Effect of Neglecting the Virtual Temperature Correction on CAPE Calculations. Weather Forecast., 9:625–629, 1994.

Dowdy, A. and Y. Kuleshov. Climatology of lightning activity in Australia: spatial and seasonal variability. Australian Meteorological and Oceanographic Journal, 64:103–108, 2014.

Dowdy, A., J. Soderholm, J. Brook, A. Brown, and H. McGowan. Quantifying Hail and Lightning Risk Factors Using Long-Term Observations Around Australia. J. Geophys. Res. Atmos., 125:1–14, 2020. 2020JD0331101.

Edwards, R. Tropical Cyclone Tornadoes - A Research and Forecasting Overview. Part 1: Climatologies, distribution and Forecast Concepts. In 24th Conference on Severe Local Storms, pages 1–10. American Meteorological Society, Savannah, Georgia, 2008.

Edwards, R. Tropical Cyclone Tornadoes: A Review of Knowledge in Research and Prediction. Electronic J. Severe Storms Meteor., 7:1–33, 2012.

Edwards, R. and R. Thompson. RUC-2 Supercell Proximity Soundings, Part II: An Independent Assessment of Supercell Forecast Parameters. In Preprints, 20th Conf. on Severe Local Storms, Orlando, FL, Amer. Meteor. Soc, pages 435–438, 2000.

Evesson, D. Tornado Occurrences in New South Wales. Australian Meteorological Magazine, 17:143–165, 1970.

Ferraro, R., J. Beauchamp, D. Cecil, and G. Heymsfield. A prototype hail detection algorithm and hail climatology developed with the advanced microwave sounding unit (AMSU). Atmos. Res., 163:24–35, 2015.

Geerts, B. and M. Noke-Raico. Tornadoes in Australia: Do We Really Know? Bulletin of the Australian Meteorological and Oceanographic Society, 8:46–51, 1995.

George, J. J. Weather Forecasting for Aeronautics. Academic press, 2014.

Griffiths, D., J. Colquhoun, K. Batt, and T. Casinader. Severe Thunderstorms in New South Wales: Climatology and Means of Assessing the Impact of Climate Change. Clim. Change, 25:369–388, 1993.

Hanstrum, B., G. Mills, A. Watson, J. Monteverdi, and I. Doswell, C.A. The Cool-Season Tornadoes of California and Southern Australia. Weather Forecast., 17:705–722, 2002.

Hastie, T., R. Tibshirani, and J. Friedman. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media, 2009.

Hersbach, H., B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. The ERA5 global reanalysis. Q. J. R. Meteorol. Soc., 146:1999–2049, 2020.

Hope, P., W. Drosdowsky, and N. Nicholls. Shifts in the synoptic systems influencing southwest Western Australia. Clim. Dyn., 26:751–764, 2006.

Hsu, C.-W., C.-C. Chang, and C.-J. Lin. A Practical Guide to Support Vector Classification. Technical report, National Taiwan University, 2003. 16pp.

Kounkou, R., G. Mills, and B. Timbal. The impact of anthropogenic climate change on the risk of cool-season tornado occurrences. Technical Report 129, Bureau of Meteorology, 2007.

Kounkou, R., G. Mills, and B. Timbal. A reanalysis climatology of cool-season tornado environments over southern Australia. Int. J. Climatol., 29:2079–2090, 2009.

Kuleshov, Y., G. de Hoedt, W. Wright, and A. Brewster. Thunderstorm distribution and frequency in Australia. Australian Meteorological Magazine, 51:145–154, 2002.

Lee, J. Tornado proximity soundings from the NCEP/NCAR reanalysis data, 2002.

Leslie, L., M. Leplastrier, and B. Buckley. Estimating future trends in severe hailstorms over the Sydney Basin: A climate modelling study. Atmos. Res., 87:37–51, 2008.

Lo, F., M. Wheeler, and A. Donald. Probabilistic Forecasts of the Onset of the North Australian Wet Season. Mon. Weather Rev., 135:3506–3520, 2007.

Lu, M., M. Tippett, and U. Lall. Changes in the seasonality of tornado and favorable genesis conditions in the central United States. Geophys. Res. Lett., 42:4224–4231, 2015.

Lukach, M., L. Foresti, O. Giot, and L. Delobbe. Estimating the occurrence and severity of hail based on 10 years of observations from weather radar in Belgium. Meteorol. Appl., 24: 250–259, 2017.

Markowski, P. and Y. Richardson. The Influence of Environmental Low-Level Shear and Cold Pools on Tornadogenesis: Insights from Idealized Simulations. J. Atmos. Sci., 71:243–275, 2014.

Mason, I. A model for assessment of weather forecasts. Aust. Met. Mag., 30:291–303, 1982.

Mason, S. and N. Graham. Areas beneath the relative operating characteristics (ROC and relative operating levels (ROL) curves: Statistical significance and interpretation. Q. J. R. Meteorol. Soc., 128:2145–2166, 2002.

Matthews, C. and B. Geerts. Characteristic thunderstorm distribution in the Sydney area. Australian Meteorological Magazine, 44:127–138, 1995.

McCaul, J., E.W. and C. Cohen. The Impact on Simulated Storm Structure and Intensity of Variations in the Mixed Layer and Moist Layer Depths. Mon. Weather Rev., 130:1722–1748, 2002.

McCaul, J., E.W. and M. Weisman. The Sensitivity of Simulated Supercell Structure and Intensity to Variations in the Shapes of Environmental Buoyancy and Shear Profiles. Mon. Weather Rev., 129:664–687, 2001.

Mercer, A., C. Shafer, C. Doswelll, III, L. Leslie, and M. Richman. Objective Classification of Tornadic and Nontornadic Severe Weather Outbreaks. Mon. Weather Rev., 137:4355–4368, 2009.

Middelmann, M. Natural Hazards in Australia: Identifying Risk Analysis Requirements. Technical report, Geoscience Australia, 2007.

Miller, R. Notes On Analysis and Severe Storm Forecasting Procedures of the Air Force Global Weather Central. Technical Report 200, Air Weather Service, United States Air Force, 1972.

Mills, G. Verification of operational cool-season tornado threat-area forecasts from mesoscale NWP and a probabilistic forecast product. Australian Meteorological Magazine, 53:269–277, 2004.

Minor, J., R. Peterson, and R. Lourensz. Characteristics of Australian Tornadoes. Australian Meteorological Magazine, 28:57–77, 1980.

Murillo, E. and C. Homeyer. Severe Hail Fall and Hailstorm Detection Using Remote Sensing Observations. J. Appl. Meteorol. Climatol., 58:947–970, 2019.

Pope, M., C. Jakob, and M. Reeder. Regimes of the North Australian Wet Season. J. Clim., 22:6699–6715, 2009.

Prein, A. and G. Holland. Global estimates of damaging hail hazard. Weather Clim. Extremes, 22:10–23, 2018.

Proctor, F. Numerical Simulations of an Isolated Microburst. Part II: Sensitivity Experiments. J. Atmos. Sci., 46:2143–2165, 1989.

Rasmussen, E. Refined Supercell and Tornado Forecast Parameters. Weather Forecast., 18: 530–535, 2003.

Rasmussen, E. and D. Blanchard. A Baseline Climatology of Sounding-Derived Supercell and Tornado Forecast Parameters. Weather Forecast., 13:1148–1164, 1998.

Raut, B., C. Jakob, and M. Reeder. Rainfall Changes over Southwestern Australia and Their Relationship to the Southern Annular Mode and ENSO. J. Clim., 27:5801–5814, 2014.

Richman, M. B. and L. M. Leslie. Machine Learning for Attribution of Heat and Drought in Southwestern Australia. Procedia Comput. Sci., pages 3–10, 2020.

Richter, H., J. Peter, and S. Collis. Analysis of a Destructive Wind Storm on 16 november 2008 in Brisbane, Australia. Mon. Weather Rev., 142:3038–3060, 2014.

Rose, S. and J. Apt. Quantifying sources of uncertainty in reanalysis derived wind speed. Renew. Energy, 94:157–165, 2016.

Said, R. Towards a global lightning locating system. Weather, 72:36–40, 2017.

Schultz, C., W. Petersen, and L. Carey. Lightning and Severe Weather: A Comparison between Total and Cloud-to-Ground Lightning Trends. Weather Forecast., 26:744–755, 2011.

Schuster, S., R. Blong, and M. Speer. A Hail Climatology of the Greater Sydney Area and New South Wales, Australia. Int. J. Climatol., 25:1633–1650, 2005.

Shafer, C. and I. Doswell, C.A. A Multivariate Index for Ranking and Classifying Severe Weather Outbreaks. Electronic J. Severe Storms Meteor., 5:1–39, 2010.

Sherburn, K. and M. Parker. Climatology and Ingredients of Significant Severe Convection in High-Shear, Low-CAPE Environments. Weather Forecast., 29:854–877, 2014.

Shikhov, A. and A. Chernokulsky. A satellite-derived climatology of unreported tornadoes in forested regions of northeast Europe. Remote Sens. Environ., 204:553–567, 2018.

Sills, D., J. Wilson, P. Joe, D. Burgess, R. Webb, and N. Fox. The 3 November Tornadic Event during Sydney 2000: Storm Evolution and the Role of Low-Level Boundaries. Weather Forecast., 19:22–42, 2004.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Weckwerth, and M. Coleman. An 18-year climatology of hailstorm trends and related drivers across southeast Queensland Australia. Q. J. R. Meteorol. Soc., 143:1123–1135, 2017a.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Wedd, and T. Weckwerth. Diurnal Preconditioning of Subtropical Coastal Convective Storm Environments. Mon. Weather Rev., 145:3839–3859, 2017b.

Soderholm, J., A. Protat, and C. Jakob. Australian operational weather radar dataset. electronic dataset, national computing infrastructure, 2019.

Storm Prediction Center. Significant Hail Parameter, 2014. URL `http://www.spc.noaa.gov/exper/mesoanalysis/help/help_sigh.html`.

Su, C.-H., N. Eizenberg, P. Steinle, D. Jakob, P. Fox-Hughes, C. White, S. Rennie, C. Franklin, I. Dharssi, and H. Zhu. BARRAv1.0: the Bureau of Meteorology Atmospheric high-resolution Regional Reanalysis for Australia. Geosci. Model Dev., 12:2049–2068, 2019.

Taszarek, M., J. Allen, T. Púčik, K. Hoogewind, and H. Brooks. Severe Convective Storms across Europe and the United States. Part II: ERA5 Environments Associated with Lightning, Large Hail, Severe Wind, and Tornadoes. J. Clim., 33:10263–10286, 2020.

Taylor, J., D. Sgarbossa, and H. Richter. The Use of Convective Parameters by the Australian Bureau of Meteorology Extreme Weather Desk in Forecasting the 16 December 2015 Tornadic Supercell Impacting Kurnell, Sydney. In Preprints, 28th Conf. on Severe Local Storms, Portland, OR, Amer. Meteor. Soc. P, 2016.

Thompson, R., R. Edwards, and J. Hart. Evaluation and Interpretation of the Supercell Composite and Significant Tornado Parameters at the Storm Prediction Center. In 21st Conference on Severe Local Storms. American Meteorological Society San Antonio, Texas, 2002.

Thompson, R., R. Edwards, J. Hart, K. Elmore, and P. Markowski. Close Proximity Soundings within Supercell Environments Obtained from the Rapid Update Cycle. Weather Forecast., 18:1243–1261, 2003.

Thompson, R., R. Edwards, and C. Mead. An Update to the Supercell Composite and Significant Tornado Parameters. In Preprints, 22nd Conf. on Severe Local Storms, Hyannis, MA, Amer. Meteor. Soc. P, volume 8, 2004a.

Thompson, R., R. Edwards, and C. Mead. Effective Storm-Relative Helicity in Supercell Thunderstorm Environments. In Preprints, 22nd Conf. on Severe Local Storms, Hyannis, MA, Amer. Meteor. Soc. P, volume 8, 2004b.

Tippett, M., J. Allen, V. Gensini, and H. Brooks. Climate and Hazardous Convective Weather. Curr. Clim. Change Rep., 1:60–73, 2015.

Vapnik, V. The Nature of Statistical Learning Theory. Springer, 1995.

Wakimoto, R. Forecasting Dry Microburst Activity over the High Plains. Mon. Weather Rev., 113:1131–1143, 1985.

Wakimoto, R. Convectively Driven High Wind Events. In Severe Convective Storms, volume 50, pages 255–298. Springer, 2001.

Walsh, K., C. White, K. McInnes, J. Holmes, S. Schuster, H. Richter, J. Evans, A. Di Luca, and R. Warren. Natural hazards in Australia: storms, wind and hail. Clim. Change, 139: 55–67, 2016.

Warren, R., H. Ramsay, S. Siems, M. Manton, J. Peter, A. Protat, and A. Pillalamarri. Radar-based climatology of damaging hailstorms in Brisbane and Sydney, Australia. Q. J. R. Meteorol. Soc., 146:505–530, 2020. doi: 10.1002/qj.3693.

Wilks, D. S. Statistical Methods in the Atmospheric Sciences, volume 100. Academic Press, 2011.

Wissmeier, U., R. Smith, and R. Goler. The formation of a multicell thunderstorm behind a sea-breeze front. Q. J. R. Meteorol. Soc., 136:2176–2188, 2010.

Witt, A., M. Eilts, G. Stumpf, J. Johnson, E. Mitchell, and K. Thomas. An Enhanced Hail Detection Algorithm for the WSR-88D. Weather Forecast., 13:286–303, 1998.

Yeo, C. Severe thunderstorms in the Brisbane region and a relationship to the El niño Southern Oscillation. Australian Meteorological Magazine, 54:197–202, 2005.

Chapter 9

Conclusions and Future Work

Severe weather affects Australia on various scales from the relatively small spatiotemporal scale of thunderstorms, to larger east coast lows and riverine flooding events, to multiyear droughts that can affect large areas of the country. This thesis has investigated the use of statistical and numerical modelling to improve understanding of drought and SCTs in Australia, and trend analysis to assess how these severe weather hazards have changed throughout the observational record.

Given the large population in southeast Australia and previous major droughts affecting this region, a significant portion of this thesis has been devoted to better understanding drought in southeast Australia. Trend analysis techniques were first applied to assess how precipitation and temperature is evolving between coastal and inland areas of the Hunter Valley Region, north of Sydney, during their respective wet seasons. It was found that mean precipitation over the coast was stable, although variance decreased. Meanwhile, inland precipitation increased for some time before reducing to levels observed in the earlier record. Additionally, temperatures are increasing in both areas, which will increase potential evapotranspiration and likely exacerbate future droughts. ENSO was found to be a persistent influence on wet season precipitation and TMax, while an interdecadal influence also appeared to have an effect on precipitation.

Although wet season precipitation (which is subjectively identified as the months during the year recording the majority of total precipitation at a given location) is important, if the wet season is of sufficient length then important trends in the time series might be missed. This was found by considering work in Hartigan et al. (2019), which is not included in this thesis, alongside the work presented in Chapter 4. In Hartigan et al. (2019), the Canberra wet season was defined as the period from September through to February. However, much of southern Australian water systems benefit more from precipitation during the cooler months of the year, approximately from April through to September (Cai et al., 2009; Kirono et al., 2010). There was little change in precipitation during Canberra's wet season, although an important interdecadal influence was recognised (Hartigan et al., 2019). In contrast, Chapter 4 finds a decreasing trend in autumn precipitation, and statistically significant increases to mean temperature in Canberra. This suggests that changes to the characteristics of precipitation, and the efficiency of precipitation transferring to water systems, are occurring in the region and Canberra is likely more vulnerable to drought. Additionally, this work has highlighted the importance of considering a range of time series groupings in order to better understand how important influences on drought or other severe weather might be changing throughout the historical record.

Important trends covering a broad region can also be understood by aggregating weather station data across the region of interest. The trend analysis techniques developed in Chapters 3 and 4 were applied to data aggregated over the Sydney Catchment Area (SCA) to assess trends in annual and seasonal precipitation and mean temperature. Median annual precipitation over the SCA was found to be decreasing over the catchment by approximately 900 mm over a 31-year period. This change in annual precipitation is the result of decreases in autumn and spring precipitation. Additionally, mean temperatures are increasing across all seasons except autumn. Similar trends in precipitation and temperature were found when splitting the catchment into two smaller regions, suggesting that there is an overall drying trend across the SCA. Given this drying trend particularly affects precipitation during the cooler months of the year, it is likely that valuable runoff into water storage systems will be reduced.

The MDB is a major agriculture region in Australia, producing up to 40% of the national food supply. This region is frequently affected by drought, and trend analysis suggests that precipitation and resulting runoff characteristics are changing in both the southern MDB (Speer et al., 2021) and northern MDB. Overall, the trend analyses in this thesis highlight that much of southeast Australia is becoming increasingly vulnerable to drought, with changes in cool-season precipitation and increases in both mean maximum and mean minimum temperatures likely affecting the amount of water discharge into river systems, and increasing potential evapotranspiration.

These changes are not unexpected, given the projected expansion of the Hadley cell and poleward progression of mid-latitude storm tracks with global warming (Timbal and Drosdowsky, 2013; Timbal and Fawcett, 2013; Post et al., 2014). However, it is important to monitor observational data for such trends, and assess whether the currently observed trends are associated at least partially with global warming. By applying statistical modelling, including linear regression and machine learning techniques, it was found that the precipitation trends in the SCA and NMDB are being affected by global warming. Other climate drivers that have previously been associated with precipitation in southeast Australia also were found to influence precipitation in these particular regions, and Canberra. Given the current trends and the association of global warming that was found with these trends, it is likely that southeast Australia, including major agricultural and metropolitan regions, will face greater water security threats, and potentially food production issues, in the future.

The work presented here demonstrates the effectiveness of statistical modelling in determining key climate drivers associated with precipitation and temperature, and the influence of global warming on these variables, with many skillful models developed compared to climatology. However, as many of the time series are non-stationary, these models will have to be periodically updated, and additional statistical techniques should be investigated for their ability to better model non-stationary data. The current skill of these models suggests that they could be applied in a seasonal forecasting framework to improve the prediction of precipitation on this time scale, but these models were developed on time series that were not lagged. Therefore, any seasonal forecasting using these models would require inputs from numerical models predicting the future states of climate drivers. As such, a future research direction would be to investigate the performance of lagged-time series as predictors of precipitation, and other variables. If statistical models with lagged predictors prove skillful, they could be used in conjunction with other modelling to improve seasonal forecasts of precipitation and drought in Australia and other regions of the world.

Climatologies of each SCT hazard were developed using reports from the STA. This high-

lighted the previously noted spatial concentration of reports by hazard to populous regions of Australia (Allen et al., 2011), with wind SCT reports displaying the greatest distribution around Australia, although still having a high concentration in the populated southeast. The STA reports were used in conjunction with lightning data and radar-estimated maximum hail size of storms in three regions over Australia, including the southeast and northwest of the country. From this data, a categorical data set of severe/non-severe storm reports was developed, to which pseudo-proximity soundings were obtained from the BoM's BARRA reanalysis data set. These pseudo-proximity soundings were used to calculate a range of indices that have been used to forecast SCTs and their associated hazards by weather services across the world. Indices that rely on selecting a specific air parcel to lift were calculated for surface-based, most-unstable, and mixed-layer parcels. When comparing pseudo-proximity soundings from BARRA with observed soundings, it was found that CAPE was frequently overestimated, mainly due to calculations at coastal locations, with MLCAPE from the pseudo-proximity soundings agreeing best with observed MLCAPE. The statistical modelling methodology was first applied to all air parcel calculations to determine the calculation that would best fit the models, for each hazard. The mixed-layer parcel was found to be the most frequently selected air parcel type amongst all hazards, and was kept for further model selection.

Following this, statistical models were developed for each hazard. Interestingly, many of the models developed, including those that did not perform the best on the testing data set, tended to perform better than the linear discriminant developed in Allen et al. (2011). This is mostly due to higher FAR from the discriminant model. The best performing statistical model on the test data set for each inidividual hazard was then applied to three-hourly BARRA data across Australia to diagnose whether the environment was supportive of that SCT hazard at that grid point and hour. From this, the frequency of days supportive of each SCT hazard at each grid point was calculated and an understanding of the spatial distribution of environments supportive of each SCT hazard, and how this distribution evolves throughout the year, was developed. There was a considerably high frequency of environments considered supportive of hail and tornadic SCTs over northern Australia considering its tropical location, which would lead to a higher freezing level and lower wind shear that would be less supportive of these hazards. This could be the result of high instability in the region outweighing the negative contribution of other environmental attributes. The spatial distribution of reports show that these models were developed mostly on data from southeast Australia, so the relationships to environmental parameters are likely more suited on southeast Australian severe storm environments, and there could be an overestimation of the frequency of environments in northern Australia. There also appears to be a rather high frequency of hail and tornadic SCT environments over southeastern Australia. The moderate to high frequency of tornadic SCT environments extends over much of southern Australia when diagnosed by the model that includes SRH1km as a predictor. For wind SCTs, there is a peak over northern Australia that could be more physical, due to the better spread of reports across Australia, including this region. The atmosphere in northern Australia has high levels of moisture due to being in the tropics, so likely is frequently supportive of wet microbursts in thunderstorms that occur during the wet season. Eastern Australia displays a moderate to high frequency of environments supportive of wind SCTs, although this is much lower than over the tropical north of the country. Overall, the distribution of SCT environments throughout the year agrees well with the distribution of STA reports, with a peak in hail SCT environment frequency during late spring/early summer over much of the continent and a peak in wind SCT environment frequency during summer. There is a bi-modal peak in tornadic SCT environments with southern Australia displaying its greatest frequency during the cooler winter months, and the remainder of Australia displaying a high frequency of environments during late spring/early summer.

There were additional models for hail and wind SCTs trained on imbalanced categorical data (i.e., more severe reports than non-severe, or vice versa) due to a coding error. The resulting models were considered against the models developed on balanced data to test the effect of applying this methodology to imbalanced data. Skill scores for these models were much lower than for the models developed on the balanced data, with the hail SCT models experiencing worse POD and FAR, while the wind SCT models had near perfect POD (due to no false negatives) but extremely high FAR. There were some similarities in the spatial distribution of environments from the unbalanced models compared to the balanced models. However, the magnitude of environments was generally greater, likely due to the higher FAR in the unbalanced models and resulting overestimation of environmental frequency. This highlights the importance of balancing categorical data sets before developing statistical models. There might be additional importance in balancing data when considering the high frequency of SCT hazard environments over much of northern Australia in the spatial distribution plots. The statistical models developed were trained on data mostly from southeastern Australia, given the high number of reports there, particularly for hail and tornadic SCTs. The environments in southeast Australia that support SCTs likely differ to those in the tropics, given the significant difference in latitude, and other effects including geomorphology. The overestimation of SCT environment frequency particularly over northern Australia could possibly be mitigated by balancing the number of reports across different latitude bands over Australia, prior to training a generalised model for Australia. An alternative avenue of research would be to split Australia into smaller regions and train models on the data within those regions. The former would possibly be a better approach to developing generalised models for Australian SCT hazards, whereas the latter would allow for highly specialised models that depict meaningful relationships with environmental parameters specific to that region.

The BARRA data set was considered long enough to assess changes in the frequency of SCT hazard environments, although it is not of sufficient length to determine if climate change has impacted any change in frequency. The trend analysis techniques applied over Chapters 3–6 were extended to allow consideration of spatially gridded data. There were increases in the frequency of environments supportive of hail SCTs over the east coast and northwest of Australia during spring, with a reduced frequency of environments over northern Australia during summer. Similarly, there was an increase in the frequency of tornadic SCT environments over eastern Australia and northwest Australia during spring, and an increase over southeast Australia during summer, while southwest Australia experienced a reduced frequency of environments. There also was an increased frequency of environments supportive of wind SCTs over Western Australia and eastern Australia during spring, and parts of southwest and southeast Australia during summer. The ECMWF's ERA5 reanalysis data set extends further back in time than BARRA, although is coarser resolution. It would be valuable to apply the same models developed here to the ERA5 data set to assess if there are differences in the diagnosed frequency of environments across Australia, and to assess how the frequency of environments might be changing over a longer time series.

The SCT climatology presented in this thesis is a first step to understanding SCT hazards by employing a range of statistical models. However, there is additional work that needs to be done. Namely, much of the understanding of how climate drivers influence the frequency of SCT hazards in a given year have focused particularly on the influence of ENSO, with this influence varying by region (Leigh and Kuhnel, 2001; Yeo, 2005; Allen and Karoly, 2014; Soderholm et al., 2017). The exception being our understanding of the climate driver influences on cool-season tornadoes over southern Australia, where SAM has been found to modulate fre-

quency through affecting the location of the mid-latitude westerlies (Kounkou et al., 2009). However, there are many additional climate drivers that likely have an effect on other SCT hazards and warm-season tornadoes, including the DMI, SAM and TPI. Wavelet analysis and statistical modelling techniques like those applied in Chapters 3–6 could be applied to diagnosed environments from the ERA5 data set to improve our understanding of which climate drivers influence the frequency of environments supportive of each SCT hazard. Although, considering the influence by region appears to be highly variable (Leigh and Kuhnel, 2001; Yeo, 2005; Allen and Karoly, 2014; Soderholm et al., 2017), such analysis likely will need to be performed on smaller scales than Australia as a whole.

This thesis has demonstrated the ability of statistical modelling to improve understanding of southeast Australian drought and Australian SCT hazards. These are just two of the severe weather hazards that affect Australia. It is important to understand severe weather in a context particular to the regions affected by those hazards, rather than relying purely on research from overseas. Additionally, it is important to apply a range of analysis techniques to understand severe weather, rather than assuming that one technique applied to one type of severe weather or region will perform best everywhere. Therefore, future work should also consider applying similar techniques to those used in this thesis, along with other statistical methods, to better understand the range of severe weather hazards that Australia is exposed to. As more studies of this nature are performed, greater agreement across techniques will improve confidence on the relationships that are found to exist, and will have greater relevance to those forecasting severe weather and planning for its impact.

GCM and RCM projections typically have rather coarse horizontal and temporal resolution, due to restricted computational resources. Considering that most severe weather, particularly SCTs, are small in either spatial or temporal scale, there is additional benefit in using statistical models developed on historical data to GCM and RCM projections, given their current inability to resolve most severe weather. There is a limited number of studies focused on how Australian severe storms might change in the future, and many GCM/RCM projections have not yet been assessed. As such, it would be highly beneficial to apply models like those produced here to future climate projections to better understand how the frequency of environments supportive of each hazard might change.

While a powerful tool, statistical modelling is only one method to understand severe weather. Chapter 7 presents a numerical modelling case study of a tornadic storm that impacted the suburb of Kurnell in Sydney, Australia. In this study, the ERA5 reanalysis data set was used for initial and boundary conditions, and the performance of both the Morrison and NSSL 2-moment 4-ice microphysics schemes was assessed. Both simulations produced a severe convective storm during the morning, highlighting that the synoptic-scale conditions were supportive of severe convection. The simulation with NSSL microphysics produced a storm that displayed similar characteristics to the observed storm. This includes the development from cell mergers over the ocean, following a similar path to the observed storm, developing updraft rotation and other supercell characteristics, and transitioning to a linear mode following impacting Sydney. However, both simulations produced storms with much of the severe weather occurring over the ocean, narrowly missing Sydney. From this study, it was noted that convection that occurred inland prior to the storm developing appeared to produce an outflow boundary that was in the Sydney region when both simulated storms approached Kurnell. This outflow boundary could have acted as a source of horizontal vorticity, from which a tornado could develop, particularly if interacting with a supercell thunderstorm. Future work will include more analysis of the simulation data presented in this thesis, and analysis of additional simulations performed using

the full model physics and a triply nested domain with higher resolution, and comparison of those results with simulations performed using the BARRA reanalysis data set.

Overall, this thesis has highlighted the importance of researching severe weather not just using globally relevant relationships, but also focusing on regional influences for locations that are exposed to a particular form of severe weather. By employing a broad range of statistical and numerical modelling techniques, a better appreciation of the risk of that severe weather hazard can be developed. Through developing this regional understanding of severe weather hazards, communities can be better prepared for the impact of severe weather. Additionally, historical observations provide a valuable tool in understanding how severe weather affected communities in the past, and how the effects might be changing at present. Continued assessment of severe weather and the variables associated with it, for trends in the time series, along with assessing trends in GCMs and RCMs is highly necessary to help communities adapt to the changing climate, and is vital as greenhouse gas emmissions continue now and into the future.

Bibliography

Allen, J. and D. Karoly. A climatology of Australian severe thunderstorm environments of 1979–2011: inter-annual variability and ENSO influence. Int. J. Climatol., 34:81–97, 2014.

Allen, J., D. Karoly, and G. Mills. A severe thunderstorm climatology for Australia and associated thunderstorm environments. Australian Meteorological and Oceanographic Journal, 61:143–158, 2011.

Cai, W., T. Cowan, P. Briggs, and M. Raupach. Rising temperature depletes soil moisture and exacerbates severe drought conditions across southeast Australia. Geophys. Res. Lett., (21): L21709, 2009.

Hartigan, J., S. MacNamara, and L. Leslie. Trends in precipitation and temperature in Canberra. In Elsawah, S. (ed.) MODSIM2019, 23rd International Congress on Modelling and Simulation. Modelling and Simulation Society of Australia and New Zealand, 2019.

Kirono, D., F. Chiew, and D. Kent. Identification of best predictors for forecasting seasonal rainfall and runoff in Australia. Hydrol. Process., 24:1237–1247, 2010.

Kounkou, R., G. Mills, and B. Timbal. A reanalysis climatology of cool-season tornado environments over southern Australia. Int. J. Climatol., 29:2079–2090, 2009.

Leigh, R. and I. Kuhnel. Hailstorm Loss Modelling and Risk Assessment in the Sydney Region, Australia. Nat. Hazards, 24:171–185, 2001.

Post, D., B. Timbal, F. Chiew, H. Hendon, H. Nguyen, and R. Moran. Decrease in southeastern Australian water availability linked to ongoing Hadley cell expansion. Earth's Future, 2:231–238, 2014.

Soderholm, J., H. McGowan, H. Richter, K. Walsh, T. Weckwerth, and M. Coleman. An 18-year climatology of hailstorm trends and related drivers across southeast Queensland Australia. Q. J. R. Meteorol. Soc., 143:1123–1135, 2017.

Speer, M., L. Leslie, S. MacNamara, and J. Hartigan. From the 1990s climate change has decreased cool season catchment precipitation and reduced river heights in Australia's southern Murray-Darling Basin. Sci. Rep., 11(16136), 2021.

Timbal, B. and W. Drosdowsky. The relationship between the decline of Southeastern Australian rainfall and the strengthening of the subtropical ridge. Int. J. Climatol., 33:1021–1034, 2013.

Timbal, B. and R. Fawcett. A Historical Perspective on Southeastern Australian Rainfall since 1865 Using the Instrumental Record. J. Clim., 26:1112–1129, 2013.

Yeo, C. Severe thunderstorms in the Brisbane region and a relationship to the El niño Southern Oscillation. Australian Meteorological Magazine, 54:197–202, 2005.