A Lagrangian Perspective on the Lifecycle and Radiative Effect of Deep Convective Clouds



William Kenneth Jones
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

2024

"One of the most fascinating things about clouds
is that they are not all the same"
Robert A. Houze, Jr.

Acknowledgements

I would like to thank my supervisor, Philip Stier, for his excellent mentorship and for allowing me the freedom to explore the many aspects of cloud and storm physics. The path to writing this thesis has been far from the easiest one, and I would not have made it without such a supportive supervisor. Matt Christensen, my co-advisor and colleague, has provided support and inspiration throughout, and shared with me his interest in the study of clouds through the lens of satellite observations.

There are numerous colleagues, office–mates and mentors who have contributed to my work and helped in many ways. The *tobac* development team, including, but not limited to, Max Heikenfeld, Fabian Senf, Sean Freeman, Julia Kukulies and Kelcy Brunner have greatly helped shape the development of the approaches used within the study the properties of deep convective clouds. Sue van den Heever and Graeme Stephens have provided many insightful discussions on the behaviour of deep convection, which in particular inspired the contents of chapter 4. Martin Stengel and the ESA Cloud-CCI+ team provided the cloud retrievals that made chapter 5 possible, and helped in the development of this research.

I am very grateful to Don Grainger and Johannes Quaas for volunteering their time as my examiners, and for their considered feedback that has helped greatly improve this thesis.

I give thanks to my parents, Rosie and Stephen, for their continued guidance, love and support, even if they couldn't convince me that there were options other than Physics.

Finally, but most importantly, I would like to thank my partner Cat for her love and support during times both happy and challenging.

Funding

I acknowledge the financial support provided by the European Research Council, Horizon 2020 RECAP (grant no. 724602) and the Natural Environment Research Council Environmental Research Doctoral Training Programme to fund these studies. Furthermore, I acknowledge the financial support received from the European Space Agency through the Cloud_cci project (contract no.: 4000128637/20/INB) which funded the research carried out in chapter 5.

Statement of Originality

I hereby declare that, to the best of my knowledge, all work contained in this thesis is my own, original contribution, unless otherwise stated.

Chapter 2 is adapted from the article "A semi-Lagrangian method for detecting and tracking deep convective clouds in geostationary satellite observations" [Jones et al., 2023a]. In this article, I (William Jones) led the development of the detection and tracking framework, data analysis and validation, and wrote the paper with contributions from Matthew Christensen and Philip Stier.

Chapter 5 is adapted from the article "A Lagrangian Perspective on the Lifecycle and Cloud Radiative Effect of Deep Convective Clouds Over Africa" [Jones et al., 2023b, accepted]. In this article, I (William Jones) designed the study. Martin Stengel provided the retrieved cloud properties and radiative fluxes from Meteosat SEVIRI data. I performed the detection and tracking and the data analysis, and wrote the paper with contributions from Martin Stengel and Philip Stier.

Abstract

Deep convective clouds (DCCs) play a key role in weather, the general circulation of the atmosphere, and the energy balance of the troposphere; the feedback of DCCs in response to global warming remains one of the largest uncertainties in future climate projections. A major limit in our understanding of DCCs is a lack of observations linking convective processes to anvil cloud properties. In this thesis, the unique capabilities of geostationary satellite observations are utilised to observe DCCs over their entire extent and lifetime, and focus on changes in the structure, diurnal cycle and cloud radiative effect (CRE) of anvil clouds.

To capture the complex behaviour of DCCs, cloud tracking methods are used to measure the properties of clouds across their lifetime and perform Lagrangian analyses. A novel cloud tracking algorithm is developed, enabling tracking of the entire lifetime of DCCs, connecting developing convective cores to the subsequent anvil cloud, and improving accuracy. This algorithm is applied to observations from multiple geostationary satellites to investigate aspects of DCC behaviour: the distribution, behaviour and diurnal cycle of convective cores and anvils; the impact of convective processes on anvil structure; and the impact of the diurnal cycle on anvil radiative effect. Analysis of these data provides insight into the distribution and lifecycle of DCCs and a comparison of the differences in land, ocean, tropical and extra-tropical convection. Furthermore, by linking convective processes to anvil structure insight is gained into how changes in convection may influence anvil feedbacks, and by combining retrievals of CRE, greater insight can be gained on the CRE of individual DCCs and their contribution to the net anvil CRE.

The properties of tracked DCCs show dependence on a number of factors, including seasonal and diurnal cycles, land–sea and tropics–extra-tropics contrasts, and the intensity and organising of the convection that drives them. The diurnal cycle of convection over land results in large changes

in the average age of anvils observed at different times of day. As a result, instruments that only observe specific times may result in erroneous comparisons between the properties of DCCs over land and ocean, which is particularly impactful on observations made in the early afternoon such as those from the A-train constellation. Convective intensity and organisation are shown to have opposing effects on anvil structure. While the most intense DCCs have thin anvil fractions 0.15 higher than the least intense, the thin anvil fraction of the most organised DCCs is on average 0.12 less than isolated DCCs. In addition, these changes impact DCCs at different points in their lifecycle, with convective intensity influencing anvil structure during the growing and dissipating stages of DCCs while organisation affects the mature stage. Anvil clouds over tropical Africa are found to have a net cloud radiative effect (CRE) of anvil clouds over Africa of $-0.94 \pm 0.91\,\mathrm{Wm}^{-2}$, confirming the near zero net effect seen in previous studies. Despite the net effect being near neutral, individual DCCs exhibit wide variance in their CRE across the diurnal cycle and it is only through the balance of day- and night-time convection that the net effect is achieved. In particular, smaller, isolated DCCs show more variable CRE, may be more sensitive to changes in the diurnal, and have a greater importance to the net balance than previously thought.

Overall, understanding the interplay between convective processes and the evolution of anvils across their entire lifecycle is vital to understanding the impact of DCCs on the climate, and cloud tracking provides a powerful tool to approach this. The results of this thesis may contribute to a better understanding of the impact of convective processes on DCCs, in particular the response of anvil structure to climate change and the importance of the diurnal cycle to the cloud radiative effect (CRE) of anvil clouds. In particular, these results suggest that further investigation is warranted on the response of anvil CRE to convective processes to better understand anvil feedbacks.

Contents

List of Figures **iv**

List of Tables **ix**

List of Acronyms and Symbols **x**

1 Introduction **1**

1.1 Motivation 1

1.2 Physical properties of DCCs 5

1.2.1 Cloud microphysical properties 5

1.2.2 Cloud radiative properties 9

1.2.3 Thermodynamics of deep convection 11

1.2.4 Lifecycle and structure of DCCs 16

1.2.5 Convective organisation 21

1.3 Anvil Radiative Feedbacks 24

1.4 Detection and tracking of deep convection 27

1.5 Structure of the thesis 33

2 A Semi-Lagrangian Method for Detecting and Tracking DCCs in Geostationary Satellite Observations **35**

2.1 Introduction 35

2.2 Data 41

2.2.1 Advanced Baseline Imager 41

2.2.1.1 Selection of ABI Channels and Channel Combinations 43

2.2.2 Geostationary Lightning Mapper 48

2.2.3 Next Generation Radar 48

2.3 Theory 49

2.3.1 Observing growing convective cores 50

2.3.2 Observing anvil clouds 54

2.4 Method . . . 57
2.4.1 Estimation of cloud motion vectors using optical flow . . . 57
2.4.1.1 Evaluation of Optical Flow Methods . . . 59
2.4.2 A Semi-Lagrangian Framework for Morphological Image Processing . . . 68
2.4.3 Detection of Growing Deep Convection . . . 70
2.4.4 Detection of Anvil Clouds . . . 71
2.5 Evaluation . . . 77
2.6 Summary . . . 80

3 Linking the Properties of Deep Convective Cores and their Associated Anvil Clouds Observed over North America 82
3.1 Introduction . . . 82
3.2 Data . . . 85
3.3 Method . . . 86
3.4 Results . . . 91
3.4.1 Distributions and properties of developing convective cores . . . 91
3.4.2 Diurnal cycle of convective cores . . . 101
3.4.3 Distributions and properties of observed anvil clouds . . . 105
3.4.4 Diurnal cycle of anvils . . . 115
3.5 Summary . . . 120

4 Contrasting effects of convective intensity and organisation on the structure and lifecycle of DCCs 123
4.1 Introduction . . . 123
4.2 Data . . . 125
4.3 Method . . . 127
4.3.1 Detection of thick and thin anvils . . . 127
4.3.2 Estimation of convective intensity and organisation . . . 130
4.3.3 Lifecycle analysis . . . 131
4.3.4 Composite analysis . . . 133
4.4 Results . . . 135
4.4.1 Change in thin anvil fraction with intensity and organisation . . . 135
4.4.2 Impact of intensity and organisation on DCC properties . . . 138
4.4.3 Changes in DCC lifecycle with intensity and organisation . . . 143
4.4.4 Changes in thin anvil fraction over DCC lifecycle . . . 148
4.5 Summary . . . 150

5 The Lifecycle and Cloud Radiative Effect of Deep Convective Clouds Over Africa **155**
5.1 Introduction 155
5.2 Data 158
5.3 Method 163
5.4 Results 165
5.4.1 Spatial and temporal distributions 165
5.4.2 Anvil Cloud Properties 168
5.4.3 Anvil CRE 173
5.5 Summary 178

6 Conclusion **182**
6.1 Summary 182
6.2 Discussion & Future work 188

Bibliography **191**

List of Figures

1.1 A photograph of a DCC seen over Africa from the ISS 2
1.2 An illustration of parcel ascent on a tephigram for a sounding taken on 2018-06-20 00:00 UTC at Jacksonville, Florida 14
1.3 An illustration of the three lifecycle phases of an isolated DCC . . . 19

2.1 Observations of a cluster of DCCs over North-West Florida throughout three stages of their lifecycle 37
2.2 Cross sections of the DCC observed in figure 2.1 as they develop over time 39
2.3 ABI channels and channel differences used with the detection and tracking algorithm 44
2.4 Clear sky vertical weighting functions for the 6.2, 7.3, 10.4 and 12.4 µm ABI channels 45
2.5 ToA BT spectra for a simulated cloud representing a DCC core at heights between 1 and 15 km 51
2.6 Simulated observations of the 6.2, 7.3, 10.4 and 12.4 µm ABI channels, WVD and SWD for a DCC core of increasing height 52
2.7 Simulated observed cooling rates of the 10.4 µm ABI channels and the WVD combination for a convective core growing at a rate of $1\,\mathrm{ms}^{-1}$. 53
2.8 ToA BT spectra for a simulated anvil cloud with optical thickness between 0 and 10 55
2.9 Simulated ABI BT observations for a DCC anvil of decreasing optical depth 56
2.10 A comparison of the motion vector fields generated by different optical flow methods 61
2.11 The difference in BT between subsequent ABI observations when taking into account the cloud motion predicted by the different optical flow algorithms. 63

2.12 A comparison of the BT difference and residual vector magnitude RMSE for each of the optical flow algorithms 64
2.13 A comparison of the computational performance of each optical flow method 66
2.14 An example of the motion vectors and their residual errors calculated using the Farnebäck algorithm 67
2.15 A comparison of convolution stencils with square connectivity in Eulerian and Lagrangian frameworks. 68
2.16 Detection of growing DCC using NEXRAD column radar reflectivity, 10.4 µm BT and WVD growth rates 72
2.17 Detection of anvil cloud extent for the mature DCC cluster in 2.1 using the edge gradient method 74
2.18 Detected regions of thin anvil cloud, thick anvil cloud, and developing cores through the growing, mature and dissipating phases of the DCC lifecycle 76

3.1 The sensor zenith angle of GOES-16 ABI observations across the CONUS domain 86
3.2 Detected cores and anvils from a snapshot of GOES-16 CONUS domain observations 90
3.3 The spatial distribution of observed cores by season 91
3.4 Monthly distributions of the proportion of cores detected each month over land, sea, tropics and mid-latitudes 92
3.5 The average speed and direction of propagation of cores 93
3.6 A map showing the mean lifetime in minutes of cores within each 1×1 ° grid box 94
3.7 A map of the average maximum core cooling rate for each 1×1 ° grid box 95
3.8 Histograms showing the distributions of observed core lifetimes, maximum areas, cooling rates and BT 96
3.9 The mean observed lifetimes, maximum areas, cooling rates and BT for cores detected over land, sea, tropics and mid-latitudes 98
3.10 Average core lifetime and minimum BT with increasing core cooling rate 100
3.11 A map of the average time of day of initiation of detected cores . . . 102

3.12 The diurnal distributions of the local time of detection for cores detected over land, sea, tropics and mid-latitudes 103
3.13 The diurnal distribution of cores detected in the NGP region 104
3.14 The change in average time of core detection with distance from the coastline 104
3.15 The diurnal cycle of the maximum cooling rate of cores detected over land, sea, tropics and mid-latitudes 106
3.16 Maps showing the spatial distribution of observed anvils by season . 107
3.17 Monthly distributions of the proportion of anvils detected each month over land, sea, tropics and mid-latitudes 108
3.18 The average speed and direction of propagation of anvils 109
3.19 Maps showing the spatial changes in the averages of anvil maximum area, lifetime, BT and number of cores 109
3.20 Histograms showing the distribution of observed anvil properties . . 111
3.21 The change in the average area and lifetime of anvils with different number of cores, and their total area coverage. 113
3.22 The change in the average areas, lifetimes, BT and number of cores of anvils observed over sea, land, tropics and mid-latitudes. 114
3.23 The average time of detection of anvils 116
3.24 The diurnal distributions of the local time of detection for anvils detected over sea, land, tropics and mid-latitudes 117
3.25 A map showing the average time of observation of anvils 118
3.26 The diurnal distributions of the local time of observation for anvils detected over sea, land, tropics and mid-latitudes 119
3.27 The average age of anvils observed over sea, land, tropics and mid-latitudes throughout the diurnal cycle. 120

4.1 Simulated ABI BT observations of DCC anvils at a range of heights and optical depths for detection of thick and thin anvil detection . . 128
4.2 Definition of anvil lifecycle stages based on the method of Futyan and Del Genio [2007] 132
4.3 The average evolution of composites of cores (red), thick anvils (orange) and thin anvils (blue) over their normalised lifetimes 134
4.4 Thin anvil proportion of anvils categorised by mean anvil BT, initial core cooling rate and number of cores 136

4.5 A comparison of the change in thin anvil proportion with maximum core cooling rate and number of cores for four different regions . . . 137
4.6 The effect of initial core cooling rate on maximum anvil area, anvil lifetime, and mean and minimum BT 140
4.7 The effect of the number of cores on maximum anvil area, anvil lifetime, and mean and minimum BT 141
4.8 A comparison of the proportion of thin anvil measured at the time of maximum anvil area, and the proportional difference in lifetime, compared to the overall thin anvil proportion. 142
4.9 The average lifetime of each lifecycle stage for anvils categorised by their initial core cooling rate 144
4.10 The average fraction of anvil lifetime of each lifecycle stage for anvils categorised by initial core cooling rate 145
4.11 The average lifetime of each lifecycle stage for anvils categorised by number of cores 146
4.12 The average fraction of anvil lifetime of each lifecycle stage for anvils categorised by number of cores 147
4.13 Mean thin anvil proportion for composites of anvil area of normalised lifetime categorised by the cooling rate of their initial cores. 149
4.14 Mean thin anvil proportion for composites of anvil area of normalised lifetime categorised by the number of cores 149

5.1 Example observations from the Meteosat SEVIRI instrument at 15:00:00 UTC on 2016/6/01 159
5.2 Comparison of the sensitivities of ABI (dashed lines) and SEVIRI (solid lines) to anvil clouds of different optical thickness 160
5.3 Comparison of the relative spectral response (RSR) functions for the GOES-16 ABI and Meteosat-11 SEVIRI thermal IR channels 160
5.4 An example of the ToA CRE derived using the radiative flux model . 161
5.5 Validation of derived broadband fluxes against monthly CERES-EBAF CRE 162
5.6 Simulated sensitivity of the SEVIRI 10.8 µm BT (top) and WVD minus SWD (bottom) to anvil clouds of varying optical thickness at heights of 10, 12 and 14 km 164
5.7 An example of the cores and anvils (detected by tobac-flow, shown at 3-hour time intervals 165

5.8 Number of detected cores and average hour of core detection 167
5.9 Anvil statistics by number of associated cores for average maximum area, average lifetime, occurrence of anvils by number of cores, and percentage of total anvil coverage 169
5.10 Anvil statistics by number of cores for average anvil CTT and average minimum anvil temperature 170
5.11 The distribution of time to coldest mean anvil CTT, largest anvil area and time of anvil dissipation 171
5.12 The proportion of anvil lifetime spent in the growing, mature and dissipating phase 172
5.13 Anvil net, LW, and SW CRE, cumulative mean CRE over anvil lifetime 174
5.14 The distribution of lifetime anvil CRE for all observed anvils 176
5.15 The distributions of mean anvil SW CRE and LW CRE 177
5.16 Average anvil CRE binned by the time of detection (local time) and mean anvil CTT 178
5.17 The contribution to the net CRE for anvils with differing numbers of cores 180

List of Tables

2.1 Comparison of data from ABI to that from older geostationary instruments 42
2.2 Operational NEΔT for the ABI and SEVIRI instruments. 46
2.3 Selected options for the libRadTran simulations 50
2.4 POD, FAR and F_1-score for three different detection methods validated against observed GLM flashes 79

3.1 Validity criteria for detected cores 87
3.2 Validity criteria for detected anvils 88

5.1 SEVIRI channels and their use in the DCC tracking algorithm and cloud properties retrieval. 157

List of Acronyms and Symbols

A-train afternoon train.

ABI Advanced Baseline Imager.

AMV atmospheric motion vector.

BoA bottom-of-atmosphere.

BT brightness temperature.

CAPE convective available potential energy.

CC4CL Community Cloud Retrieval for Climate.

CCI Climate Change Initiative.

CCN cloud condensation nuclei.

CERES Clouds and the Earth's Radiant Energy System.

CIN convective inhibition.

CONUS Continental United States.

CPR Cloud Profiling Radar.

CRE cloud radiative effect.

CTH cloud top height.

CTT cloud top temperature.

DCC deep convective cloud.

DIS Deep Inverse Search.

Dual TV-L^1 Dual Total Variation Regularisation & Robust L^1 Norm.

EBAF energy balanced and filled.

ESA European Space Agency.

EUMETSAT European Organisation for the Exploitation of Meteorological Satellites.

FAR false alarm rate.

FAT fixed anvil temperature.

GCM global climate model.

GLM Geostationary Lightning Mapper.

GOES Geostationary Operational Environment Satellite.

INP ice nucleating particle.

IR infrared.

ITCZ inter-tropical convergence zone.

IWC ice water content.

LCL lifted condensation level.

LFC level of free convection.

LNB level of neutral buoyancy.

LW longwave.

LWC liquid water content.

MCMIP multi-channel cloud and moisture imagery product.

MCS mesoscale convective system.

MSG Meteosat Second Generation.

NASA National Aeronautics and Space Administration.

NEΔT noise equivalent ΔT.

NEXRAD next generation (weather) radar.

NGP Northern Great Plains.

NIR near infrared.

NOAA National Oceanic and Atmospheric Administration.

OD optical depth.

PBL planetary boundary layer.

PCA principle component analysis.

POD probability of detection.

$\mathbf{r_e}$ effective radius.

RGB red, green, blue.

RMSE root mean square error.

SEVIRI Spinning Enhanced Visible Infra-Red Imager.

SW shortwave.

SWD split window difference.

ToA top-of-atmosphere.

USA United States of America.

UTC universal co-ordinated time.

WSR-88D weather surveillance radar, 1988, Doppler.

WV water vapour.

WVD water vapour difference.

Chapter 1

Introduction

1.1 Motivation

Deep convective clouds (DCCs)—also known as cumulonimbus (*'heaped raincloud'*) or thunderstorms—are dynamical atmospheric phenomena resulting from instability in the troposphere. DCCs consist of tall, central 'cores' surrounded by a large, detrained cirrus 'anvil' (see fig. 1.1). The central cores of DCCs form around intense convective updrafts that transport water from near the surface to the upper troposphere. At the top of these updrafts, the divergence of the cloudy airmass leads to the formation of a large area of detrained anvil that extends beyond the footprint of the core. With heights exceeding 10 km, and anvils that measure hundreds or thousands of km in extent, DCCs form some of the most physically imposing objects in the atmosphere and have important impacts on weather and climate.

First identified as a unique cloud type in the late 19$^{\text{th}}$ century by Abercromby [1887] and Hildebrandsson [1887], DCCs have been a focus of scientific investigation ever since. DCCs are the source of many severe weather events [Matsudo and Salio, 2011; Houze, 2014] including heavy precipitation [Westra et al., 2014], lightning [Williams et al., 1992], hail [Punge and Kunz, 2016], flooding, tornadoes and tropical cyclones.

DCCs play an important role in weather and climate beyond extreme events. In the tropics, deep convection forms the ascending branch of the Hadley cell [Riehl and



Figure 1.1: A photograph of a DCC seen over Africa from the International Space Station on 5th February 2008. The main convective core can be seen in the centre of the image as the turbulent, textured region emerging through the wide, flat anvil that surrounds it. On the right-hand side of the image a new core can be seen developing that is about to make contact with the anvil of the original core. Image credit: NASA Image and Video Library, iss016e027426.

Malkus, 1958], and in doing so begins the transport of energy through the atmosphere from the equator to the poles. In many regions of the world, from tropical Africa to the Great Plains of North America, DCCs provide the majority of precipitation [Feng et al., 2021]. The large anvil clouds of DCCs have a mediating effect on the radiative heating of the climate, reflecting incoming sunlight and trapping outgoing longwave radiation in equal amounts [Ramanathan et al., 1989; Hartmann et al., 1992; Hartmann, 2016].

DCCs are expected to have multiple responses to climate change. Increases in atmospheric temperature and humidity increase both the intensity and frequency of heavy precipitation from deep convection [e.g. Allen and Ingram, 2002; Trenberth

et al., 2003; Held and Soden, 2006]. In addition, global warming is also expected to lead to multiple radiative feedbacks from convective anvil clouds with both positive (warming) and negative (cooling) impacts [Bony et al., 2015]. These feedbacks remain one of the largest uncertainties in future climate projections [Sherwood et al., 2020]. Understanding the behaviour, interactions and feedbacks of DCCs is therefore vital for understanding both our present-day climate and its response in a changing world [Gasparini et al., 2023].

Anthropogenic aerosols also affect DCCs through aerosol–cloud interactions such as convective invigoration [Khain et al., 2005; Rosenfeld et al., 2008], although these processes remain uncertain [Varble et al., 2023]. Aerosol impacts on DCCs are difficult to study in observations due to the small magnitude of the effects in comparison to meteorological variability and measurement uncertainty [Grabowski, 2018], as well as the difficulty in separating aerosol–cloud interactions from confounding processes [Varble, 2018]. As a result, the effects of aerosols on DCCs are not considered within the scope of this thesis.

DCCs occur over a wide range of scales, and this limits the sources of observations which can fully capture their behaviour. The processes of individual DCCs span a scale from single kilometres and minutes, to hundreds of km and multiple days for large, organised convective systems. When considering the response of the global circulation and atmospheric energy balance, it may take multiple years of observations for an equilibrium state to be observed [Muller and O'Gorman, 2011]. Geostationary weather satellites, such as the National Oceanic and Atmospheric Administration (NOAA) Geostationary Operational Environment Satellite (GOES) series and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) Meteosat series have provided imagery of DCCs over many decades for monitoring and predicting the weather. However, these instruments have not traditionally supported the accuracy required for scientific studies, and it is only with the latest generation

that the capability of this imagery for studying the behaviour of deep convection is being fully realised.

The study of DCCs is made difficult by their complex, dynamic behaviours. Many techniques used to study more slowly varying or spatially uniform atmospheric phenomena (including other cloud types, such as stratus and cumulostratus regions) cannot fully capture the complexity of DCCs. Even snapshot observations, such as those made by polar-orbiting earth observation satellites, cannot fully characterise DCC behaviour, no matter how detailed they are, due to the rapid changes in the properties of DCCs over their lifetimes. Lagrangian methods—those which follow the DCCs' motion—provide vital observations of DCC properties over their lifecycle. While particle trajectory models have been successfully utilised to study boundary layer clouds from a Lagrangian perspective [e.g Eastman and Wood, 2018; Christensen et al., 2020], these techniques are not easily applicable due to the complex wind environment of DCCs, and the low accuracy of reanalysis models and derived atmospheric motion vectors (AMVs) in these environments. Instead, image processing methods that detect and track DCCs are used to study their behaviour from a Lagrangian perspective. These methods have seen a renaissance in recent years for studying deep convection, with applications including convective resolving models, weather radars and geostationary satellite imagery. There remain several limitations in existing tracking models, and so the development of novel techniques is required to better understand the full spectrum of DCCs.

Overall, understanding the behaviour of DCCs and their response to climate change is vital both to predicting extreme weather events and climate feedbacks. This introductory chapter will explore a range of topics regarding the behaviour of DCCs and their further impacts. In particular, it will focus on the relation between deep convection and the behaviour of anvil clouds, and how novel cloud tracking methods can be used to provide new insights from observations of DCCs.

1.2 Physical properties of DCCs

The physical nature of a DCC can be divided into three parts: a cloud system, consisting of a multitude of liquid and ice water cloud particles along with several kinds of precipitation; a thermodynamic system in which moisture and energy are transformed into intense vertical motion and circulations across several scales; and a radiative system in which the large anvil reflects sunlight and absorbs longwave (LW) radiation, both heating and cooling the atmosphere within and below it. These processes are not independent and interact with each other in numerous ways. As a result, understanding each of them is important to understanding the behaviour of DCCs.

1.2.1 Cloud microphysical properties

DCCs, like any other cloud, are formed from a great number of water and ice particles suspended in the atmosphere. DCCs consist of a vertically growing core spanning from the planetary boundary layer (PBL) to the tropopause, a distance often exceeding 10 km, with a diameter of 10 km and updraught velocities of around 10 ms^{-1} [Weisman, 2015], and a surrounding anvil cloud formed due to horizontal divergence of cloud ice which has been lifted to the level of neutral buoyancy [Houze, 2014]. While the anvil cirrus of a DCC consists of ice particles, many of these particles form in the liquid phase within the core, before freezing as they are lifted vertically and then detrained horizontally. As a result, while ice–phase microphysics determines the evolution of anvil cirrus over time, liquid–phase microphysics is also important in DCCs due to the initial formation of many anvil particles by homogeneous freezing of liquid droplets formed in the convective core.

When an airmass containing water vapour is lifted it expands and cools. In doing so, the saturation vapour pressure decreases—in accordance with the Clausius-Clapeyron relation (eq. 1.1)—at a faster rate than the specific vapour pressure of

The Clausius–Clapeyron relation for saturation vapour pressure (eq. 1.1), where e_s is the saturation vapour pressure, T is the parcel temperature in Kelvin, $L_\mathrm{v}(T,p)$ is the latent heat of vaporisation of water, and R_v is the gas constant of water vapour. This cannot be solved analytically due to the temperature dependence of the latent heat of vaporisation. An accurate approximation in the form of the August–Roche–Magnus equation [Magnus, 1844] be found empirically (eq. 1.2, with coefficients from Alduchov and Eskridge [1997]). Similarly, the relation for saturation vapour pressure over ice ($e_{\mathrm{s,i}}$) is given by eq. 1.3, where $L_\mathrm{s}(T,P)$ is the latent heat of sublimation of water (equivalent to $L_\mathrm{v} + L_\mathrm{f}$, the latent heat of fusion). The empirical approximation for $e_{\mathrm{s,i}}$ is given by eq. 1.4. Note that at 0 °C, $e_{\mathrm{s,i}} = e_\mathrm{s}$, whereas the the approximations differ slightly.

$$\frac{1}{e_\mathrm{s}}\frac{\mathrm{d}e_\mathrm{s}}{\mathrm{d}T} = \frac{L_\mathrm{v}(T,p)}{R_\mathrm{v}T^2} \tag{1.1}$$

$$e_\mathrm{s}(T) \cong 6.1094\ \mathrm{kPa}\left(\frac{17.625\ \mathrm{K}^{-1}\cdot(T-273.15\ \mathrm{K})}{T-30.11\ \mathrm{K}}\right) \tag{1.2}$$

$$\frac{1}{e_{\mathrm{s,i}}}\frac{\mathrm{d}e_{\mathrm{s,i}}}{\mathrm{d}T} = \frac{L_\mathrm{s}(T,p)}{R_\mathrm{v}T^2} \tag{1.3}$$

$$e_{\mathrm{s,i}}(T) \cong 6.1121\ \mathrm{kPa}\left(\frac{22.587\ \mathrm{K}^{-1}\cdot(T-273.15\ \mathrm{K})}{T+0.71\ \mathrm{K}}\right) \tag{1.4}$$

the airmass. When the vapour pressure of the airmass exceeds the saturation vapour pressure it is considered supersaturated, and can condense to form cloud droplets. It is very difficult for water droplets to form in a perfectly clean atmosphere, since for this to happen supersaturation of several hundred percent would be required. Instead, soluble aerosol particles in the supersaturated airmass become the surface on which cloud droplets form, a process known as cloud condensation nuclei (CCN) activation [V. Hobbs, 1993]. The conditions under which aerosols can be activated as CCN are given by the Köhler curves, which combine the competing effects of Kelvin's equation and Raoult's law on the equilibrium vapour pressure above the surface of a droplet. Kelvin's equation defines how the equilibrium vapour pressure increases as the radius

of curvature decreases, therefore requiring a higher supersaturation for the activation of smaller droplets. Raoult's law regards the effect of soluble ions on the equilibrium vapour pressure: a larger amount of solute within the droplet reduces the required supersaturation. The resulting curve has a peak supersaturation requirement at a certain droplet radius. For an aerosol particle to be activated and grow into a cloud droplet it must either be larger than this radius and large enough that water can condense on it at the current supersaturation, or contain enough soluble ions that the peak supersaturation of the Köhler curve is less than the airmass supersaturation. As a result, aerosol particles with high solubility—including in particular sulphate and nitrate aerosols, volatile organic compounds and sea salt—form the majority of CCN.

Activated cloud droplets grow through two processes: condensation and coalescence. Condensation growth is most effective on small droplets, as they have the largest surface area to volume ratio, and is responsible for the growth of cloud droplets from the radius of CCN (on the order of 0.1 µm) to that of a typical cloud droplet of around 10 µm [Rogers, 1979]. This process also results in a narrowing of the cloud droplet size distribution. Coalescence growth occurs through either the merging of cloud droplets due to collision, or the collection of cloud droplets by rain droplets as the fall (accretion). It becomes effective for larger cloud droplets beginning with radii of around 15 µm and is responsible for their growth into rain droplets [Rogers, 1979]. However, growth through condensation is very slow to reach the droplet sizes required for coalescence growth, and instead turbulence and the activation of giant CCN [Feingold et al., 1999] are required for rain droplets to form. On average, only about 30% of the condensed water within a cloud will become rain droplets [Trenberth et al., 2003], and the number of raindrops constitutes a tiny fraction of the total number of droplets.

Ice clouds have important and complex microphysics. Ice cloud particles may be formed either through direct deposition of water vapour into ice, or through the freez-

ing of liquid cloud droplets. Liquid water will not, however, freeze immediately below the freezing point due to the freezing energy barrier [Heymsfield and Miloshevich, 1993]. Instead, liquid cloud droplets may exist in a super-cooled state, which may be frozen by one of two processes. Homogeneous freezing occurs at temperatures below –38 °C, allowing liquid droplets to freeze without external interactions [Koop et al., 2000]. As a result, homogeneous freezing of liquid cloud droplets tends to result in a larger number of smaller ice particles [Kärcher and Lohmann, 2002; Ickes et al., 2015]. Heterogeneous freezing occurs due to the presence of ice nucleating particles (INPs) [Kanji et al., 2017]. INPs reduce the freezing energy barrier due to having similar structures to ice crystals [Hoose and Möhler, 2012], and so allow freezing at warmer temperatures [Kärcher and Ström, 2003]. However, INPs are somewhat rare in the atmosphere [Burrows et al., 2022], and so heterogeneous freezing tends to result in fewer ice cloud particles, and, in some cases, mixed-phase clouds which consist of both ice and liquid droplets.

Ice particles may grow through deposition; the direct freezing of water vapour onto the surface of the crystal. Similarly, ice crystals may lose mass due to sublimation, although due to the low rates of sublimation at cold temperatures this process is slow at high altitudes [Boehm et al., 1999; Seeley et al., 2019a]. Ice crystals may grow through aggregation, where multiple crystals join together, or through riming, where super-cooled water droplets freeze on contact with an ice crystal [Ryan et al., 1976]. In some conditions, additional, smaller ice crystals are produced through secondary ice production processes including rime-splintering [Hallett and Mossop, 1974], droplet shattering and collision fragmentation [Field et al., 2017]. These processes are common in deep convection due to strong updrafts, turbulence and the occurrence of homogeneous freezing, and so the most intense DCCs tend to create anvils with smaller ice particles than those formed in in-situ cirrus. Ice crystals are commonly removed from the atmosphere via sedimentation, which plays an important role in

precipitation [Mülmenstädt et al., 2015].

Overall, ice cloud microphysics has complex dependencies on a wide range of factors. In addition, these complex processes lead to a wide variety of shapes and sizes of ice particles, including both regular and irregular shapes [Waitz et al., 2022], which adds further variance to ice crystal interactions. The understanding of ice cloud microphysics, along with its parameterisations in climate models, is a large and important source of uncertainty in understanding clouds and future climate change [Sullivan and Voigt, 2021; Gasparini et al., 2023].

1.2.2 Cloud radiative properties

Cloud droplets interact with radiation both through the reflectance and scattering of solar visible and near infrared (NIR) radiation, and also through the absorbance and emission in the LW spectrum. While the reflection of incoming radiation has a cooling effect on the top-of-atmosphere (ToA) atmospheric energy balance, the LW effect of clouds is warming. As a result, a cloud may have a net warming or cooling effect depending on the balance of these two factors. The difference between the ToA radiative flux with clouds versus that for a clear sky is referred to as the cloud radiative effect (CRE).

The shortwave (SW) reflectance of a cloud depends upon both the microphysical properties of the cloud particles and the bulk properties of the cloud, including the number of droplets and the amount of liquid or ice. The optical properties of a cloud are generally measured by two variables: the effective radius (r_e) and the cloud optical thickness (τ). r_e is the ratio of the third and second moments of the particle size distribution, which provides a good approximation of the mean scattering radius within a volume of cloud particles as long as the radius is larger than the wavelength (which, for cloud particles that tend to reflect mostly at shorter wavelengths, is almost always the case) [Hansen and Travis, 1974]. While this is relatively straightforward

for liquid cloud droplets, for ice clouds this is complicated by their non-spherical geometries [Wyser, 1998]. τ is the negative logarithm of the transmittance through a cloud layer, and clouds with high optical thickness reflect more SW radiation. While τ varies with wavelength, it is mostly uniform throughout the visible wavelengths for cloud droplets [Hu and Stamnes, 1993] and tends to be provided without a specified wavelength. Unless stated otherwise, all values for τ provided in this thesis are taken at 550 nm. Due to the small variance of τ across the range of visible wavelengths, these values are considered directly comparable with those retrieved using geostationary satellite observations at 640 nm.

The LW effect depends upon both the temperature (and hence height) of the cloud as well as their thickness and microphysics. Clouds that are colder than the surface will emit less LW radiation than they absorb and hence have a warming effect. As the LW absorption of clouds reduces more slowly as a cloud thins than their SW reflectivity, thin clouds will tend to have a greater LW effect than SW.

Thick, low-level clouds, such as cumulus, stratocumulus and stratus have an overall cooling effect as they reflect a large proportion of incoming solar radiation but emit in the LW at a similar temperature to the surface. On the contrary, high, thin, cirrus clouds have a strong warming effect as they transmit the majority of incoming solar radiation but emit at a much lower temperature than the surface. DCC anvils, being both thick and at a high altitude, tend to have a balanced effect within 10 Wm^{-2} of neutral [Ramanathan et al., 1989; Hartmann et al., 1992; Hartmann, 2016], however their SW and LW CRE both have large magnitudes. Overall, CRE has a cooling effect on the climate, but displays a positive feedback to global warming, generally attributed to a reduction in low, oceanic cloud coverage [Bony et al., 2015]. This feedback is one of the largest uncertainties in future climate change [Sherwood et al., 2020], and recent research has found that this uncertainty may be underestimated by current global climate models (GCMs) [Hill et al., 2023]. The feedbacks of DCCs

to climate change are particularly uncertain, and recent research has highlighted our lack of understanding of the response of the optical properties of DCCs to increasing temperatures [McKim et al., 2024]. The response of DCCs to climate change will be discussed in more detail in section 1.3.

1.2.3 Thermodynamics of deep convection

Deep convection requires three ingredients to occur [Brooks et al., 2019]:. First, a build-up of conditional instability throughout the troposphere. Second, a source of moisture in the lower troposphere. Third, a buoyancy perturbation or vertical motion which lifts a parcel of moist, low-level air above the lifted condensation level (LCL) and the level of free convection (LFC).

Equation 1.5 The buoyancy equation for an air parcel perturbed from the surrounding environment. The buoyant acceleration, B, is approximately equal to gravitational acceleration times four terms, from left the right: the temperature term, the vapour term, where q_v is the mixing ratio of water vapour, the pressure perturbation term, and the hydrometeor drag, where q_H is the mixing ratio of hydrometeors. Variables with bars represent the background state.

$$B = g\left(\frac{T-\bar{T}}{\bar{T}} + 0.61\frac{q_\mathrm{v}-\bar{q}_\mathrm{v}}{\bar{q}_\mathrm{v}} - \frac{p-\bar{p}}{\bar{p}} - q_\mathrm{H}\right) \tag{1.5}$$

Conditional instability refers to a situation in which the decrease in temperature with height in an atmospheric column is greater than the moist pseudo-adiabatic lapse rate, but less than that of the dry adiabatic lapse rate. Whereas unconditional instability (where the temperature lapse rate is greater than the dry adiabatic lapse rate) is quickly stabilised through dry convection, conditional instability can continue to build until a convective cloud is formed. When water condenses, the latent heat released results in a positive buoyancy force (eq. 1.5 which accelerates the parcel upwards leading to large vertical velocities. The build-up of instability is measured as the convective available potential energy (CAPE) of an air parcel (eq. 1.7), which

represents the potential kinetic energy provided by the buoyancy of a parcel lifted above the LFC. The work required to lift the parcel to the LFC is measured as the convective inhibition (CIN) (eq. 1.8). High values of CAPE are associated with intense convection.

Equation 1.6 By ignoring the pressure, vapour and hydrometeor terms of eq. 1.5 the buoyancy equation can be written in terms of the vertical velocity, w, the height z and temperature. CAPE is calculated by integrating the right-hand side of the equation over height between the LFC and the level of neutral buoyancy (LNB) (eq. 1.7). CIN is calculated by integrating the negative buoyancy between the initial position of the parcel and the LFC (eq 1.8).

$$\frac{\mathrm{D}}{\mathrm{D}z}\left(\frac{w^2}{2}\right) = g\left(\frac{T-\bar{T}}{\bar{T}}\right) \tag{1.6}$$

$$\mathrm{CAPE} = g\int_{\mathrm{LFC}}^{\mathrm{LNB}}\left(\frac{T-\bar{T}}{\bar{T}}\right)\mathrm{d}z \tag{1.7}$$

$$\mathrm{CIN} = -g\int_{z_0}^{\mathrm{LFC}}\left(\frac{T-\bar{T}}{\bar{T}}\right)\mathrm{d}z \tag{1.8}$$

The properties of a convective airmass are commonly derived using parcel theory [Emanuel, 1983, e.g.]. This approach approximates the air parcel as undergoing adiabatic processes while it is lifted from the lower to the upper troposphere. The buoyancy of the air parcel along this trajectory is dependent on the initial height and moisture content of the parcel as well as the temperature profile of the environment it exists in. Analysis of these trajectories can be performed by plotting them on a thermodynamic diagram. An example of a tephigram, which plots the ascent of a parcel on axes of temperature and potential temperature which are skewed by 45 degrees, is shown in fig. 1.2. The tephigram has advantages over other thermodynamic diagrams commonly used in meteorology, such as the T–logP diagram, as its axes are

orthogonal to each other and therefore thermodynamic properties such as CAPE and CIN are accurately represented by their areas.

From its initial position, the parcel is assumed to undergo adiabatic cooling as it is lifted until its temperature reaches that of the dew point of the initial airmass. At this point, the LCL, the saturation of the airmass reaches 100%, condensation occurs and the cloud begins to form. Beyond this point, the temperature of the parcel is assumed to follow the moist pseudo-adiabatic lapse rate. As this lapse rate is less than that of an environment with conditional instability, as the parcel continues to rise it will reach a point where its temperature is greater than that of the surrounding environment referred to as the LFC. Above the LFC the parcel will experience positive buoyancy. As the moist adiabatic lapse rate approaches that of the dry adiabatic lapse rate at low temperature and humidity (and hence at higher altitudes), and the environmental lapse rate reduces as it approaches the tropopause layer [Fueglistaler et al., 2009], the parcel will cool below the temperature of the surrounding environment. The point at which the temperature profile of the parcel again crosses that of the environment is called the level of neutral buoyancy (LNB). Although the air parcel may continue to rise above this point due to its momentum, negative buoyancy will act to halt its ascent.

Two further properties related to convection are also shown on the tephigram. CAPE can be calculated as the area between the parcel temperature profile and the environment profile between the LFC and the LNB, showing clearly how CAPE is the total work exerted by positive buoyancy forces between these two levels. From this, CAPE can be used to calculate the theoretical maximum updraft velocity if all CAPE is converted to kinetic energy, although typical maximum updraft velocities are half of this amount due to the dilution of updraft plumes through mixing with the surrounding environment [LeMone and Zipser, 1980; Romps and Kuang, 2010]. Similarly, convective inhibition (CIN) can be calculated as the area between the

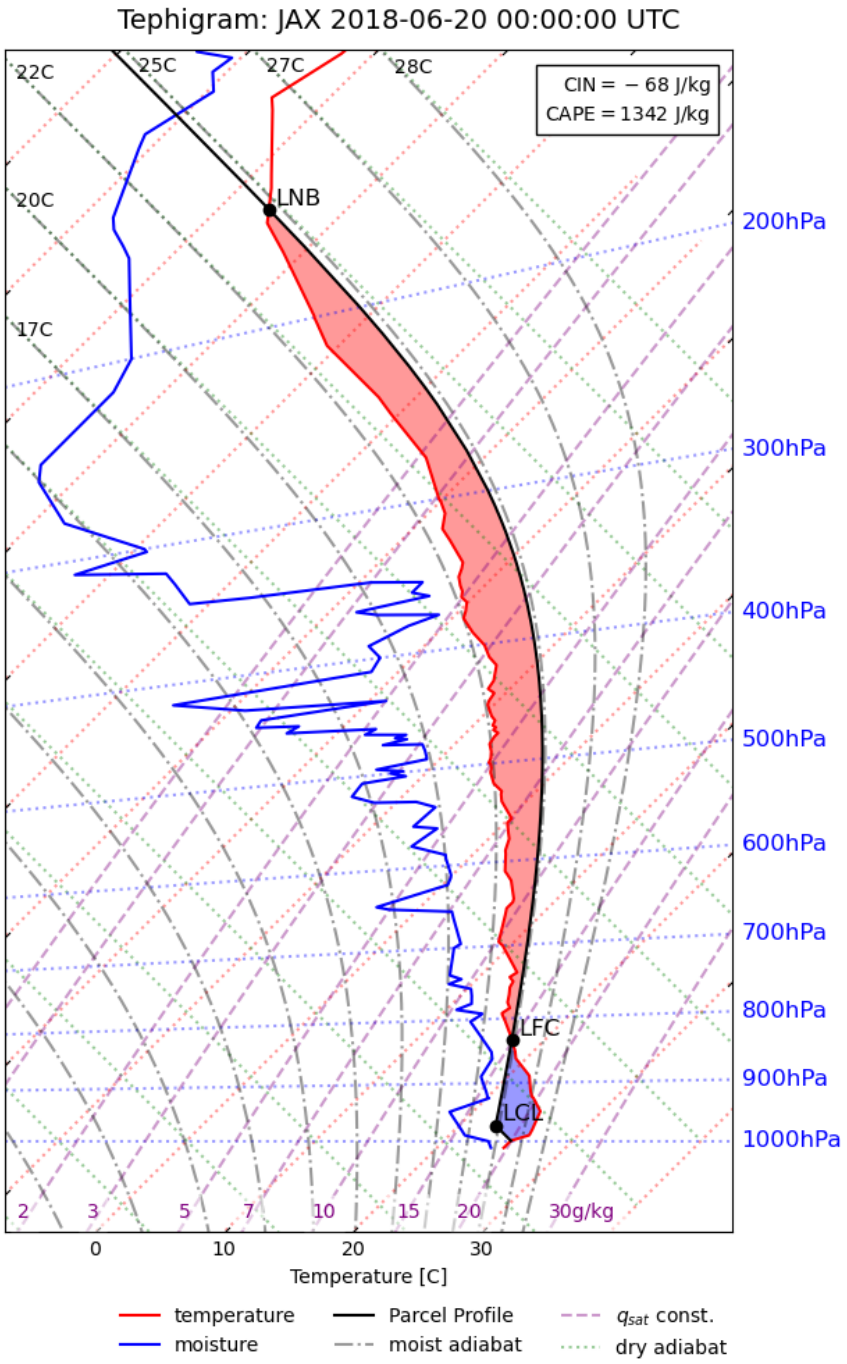


Figure 1.2: An illustration of parcel ascent on a tephigram for a sounding taken on 2018-06-20 00:00 UTC at Jacksonville, Florida. The environmental temperature profile is shown by the red line, and the dewpoint temperature by the blue line. The path of the parcel, which starts near the surface, is shown by the black line. Points marking the LCL, LFC and LNB are shown. The CAPE and CIN are shown by the red and blue areas respectively.

temperature profiles between the initial parcel location and the LFC. CIN measures the work required to lift the parcel, overcoming negative buoyancy forces, in order to reach the LFC.

There are a number of caveats with the parcel approach for deep convection that must be considered. The convective profile of a parcel depends upon its starting conditions including initial height, temperature and humidity. As a result, for even a single atmospheric column, a whole ensemble of parcels initialised at different locations within the PBL can be considered, each with different values of LCL, LFC, LNB, CAPE and CIN. To account for this, the mixed-layer CAPE is often calculated from the average for parcels initiated in the PBL [Stevens, 2005]. In addition, the properties of the most unstable parcel may also be found in the same manner, which may provide useful information for predicting convective initiation.

Secondly, it is generally assumed that above the LCL the parcel follows the moist pseudo-adiabat, which approximates that all condensed water is removed from the parcel immediately via precipitation [Emanuel, 1994]. Although this approximation may seem to agree with the high precipitation rates of DCCs, observed profiles of convective updrafts more closely match the moist adiabatic lapse rate (which assumes that all condensed water remains within the parcel) [Xu and Emanuel, 1989]. In addition, while it is often considered that MSE is conserved in convective processes this is only correct in the case of environments in hydrostatic balance [Peters and Chavas, 2021]. In deep convection, where this is not the case, it is a better approximation that MSE minus CAPE is conserved [Romps, 2015].

Finally, and arguably most importantly, is that the parcel model assumes that there is no mixing between the air parcel and the surrounding environment, a state that is referred to as undiluted. In the dynamic and turbulent environment of deep convection mixing does occur, however, and it is very rare for an undiluted state to occur [Zipser and LeMone, 1980; Romps and Kuang, 2010]. The entrainment of

dry air into convective updrafts due to mixing reduces both CAPE [Zhang, 2009] and the LNB [Masunaga and Luo, 2016] as it adjusts the parcel profile closer to the background temperature profile, reducing the buoyancy forces. Observations of DCCs have shown that while the highest cloud tops of DCCs reach or even exceed the LNB, the majority of the anvil cloud is detrained at a level substantially below the LNB [Takahashi and Luo, 2012; Takahashi et al., 2017a].

An additional factor not considered by the parcel model is the impact of circulation on instability. Low-level convergence may transport additional heat and moisture to the base of the profile, increasing the potential for convection. Wind shear—change in the speed and/or direction of wind with height—may also increase instability by transporting a colder airmass over a warmer boundary layer. Wind shear plays an important role in the behaviour of DCCs which will be explored in the following section.

1.2.4 Lifecycle and structure of DCCs

The lifecycle of deep convection begins with the development of instability throughout the troposphere. The primary driver of instability is cooling through LW emissions from water vapour, which cools the troposphere at around 2 K per day [Mapes, 2001; Jeevanjee and Fueglistaler, 2020]. SW heating also affects instability which leads to diurnal differences which show large land–sea contrasts. Over land, SW absorption warms the surface during the daytime, which sensible heating transfers to the lower troposphere. This SW heating leads to an increase in CAPE throughout the day, and, in particular, a peak in convective initiation in the afternoon due to surface heating and the resulting enthalpy flux into the PBL [Hendon and Woodberry, 1993]. On the contrary, over the ocean, SW heating has less of an effect on the surface due to the greater heat capacity of water and the ocean mixed layer. Instead, a peak in instability is seen during the night and early morning [Gray and Jacobson,

1977], which is believed to be due to the effects of SW heating on the mid and upper troposphere during the day [Wall et al., 2018]. The diurnal cycle of convection is much less pronounced over the ocean than over land [Soden, 2000]. Instability can also be driven by the relative motion of air masses. Over the great plains of North America, warm, moist air is transported northward from the Gulf of Mexico which leads to instability as it moves under colder air [Walters and Winkler, 2001].

Once sufficient instability has built up throughout the troposphere for deep convection to occur, a lifting action or buoyancy perturbation must occur to overcome CIN and initiate the formation of a DCC. This initiating process is referred to as the 'triggering' of convection. Triggering can occur through perturbations of the first three terms in the buoyancy equation (eq. 1.5). SW heating of the surface can cause an increase in surface air temperature through sensible heat fluxes, leading to dry convection and convective initiation. Surface heating also causes latent heat fluxes, increasing the enthalpy and moisture content of the boundary layer. In addition, the convergence of air masses within the boundary layer can also increase buoyancy by convergences of moisture and positive pressure perturbations. Mechanical lifting mechanisms may also be involved in the initiation of convection. Orographic lifting can occur when warm, moist air is blown towards mountains and cools as it ascends [Hodges and Thorncroft, 1997]. Gust fronts of cold, dense air may also act to lift warmer, moist air masses above them. These gust fronts may be caused by several processes, including synoptic frontal systems [Wilson and Schreiber, 1986; Jirak and Cotton, 2007], sea breeze [Tripoli and Cotton, 1979; Park et al., 2020], and cold pools caused by precipitation [Grant and van den Heever, 2016].

The occurrence of cumulus and cumulus congestus clouds (convective clouds in the low- and mid-troposphere [Johnson et al., 1999]) may also lead to deep convection. These lower-level convective clouds cause convergences of warm, moist air near the surface, building instability, and may also act to 'condition' the lower troposphere to

convection by adjusting the environmental temperature profile closer to that of an ascending air parcel and hence reducing CIN [Masunaga and L'Ecuyer, 2014; Schulz and Stevens, 2018].

Once initiated, the lifecycle of DCCs can be separated into three sections: a growing phase, where the core develops vertically (fig. 1.3 a); a mature phase in which the anvil cloud develops horizontally while convection continues within the core (fig. 1.3 b); and a dissipating phase in which the anvil cloud dissipates after convective activity ceases within the core (fig. 1.3 c) [Wall et al., 2018]. For isolated DCCs (those consisting of a single core) the overall lifecycle typically spans 1-3 hours [Chen and Houze, 1997]. However, DCCs may also form with multiple cores feeding a single anvil cloud [Roca et al., 2017], and in these cases may span areas several orders of magnitude larger [Houze, 2004], and exist for 10-20 hours or longer [Chen and Houze, 1997].

The initiation stage occurs between the point of initiation and the time at which the top of the convective core stops growing upwards. During this stage precipitation may occur in both the warm and ice phases, and the horizontal growth of the cloud is small in comparison to the vertical growth. The initiation stage can be further broken down into the periods before and after the onset of freezing. Freezing releases additional latent heat of fusion, although this may not contribute additional buoyancy as it is compensated for by the environmental lapse rate [Seeley and Romps, 2016].

Due to the large vertical velocities and high supersaturation present in deep convective updrafts, freezing is dominated by homogeneous freezing processes. Furthermore, the intense dynamics result in a high rate of secondary ice production through the collision of droplets and rime–splintering. As a result, DCCs tend to form with a large number of small ice particles. In convective updrafts with large vertical velocity, the momentum of the growing cloud may propel it above the LNB. This produces a phenomenon known as an overshooting top, where the top of the core penetrates the

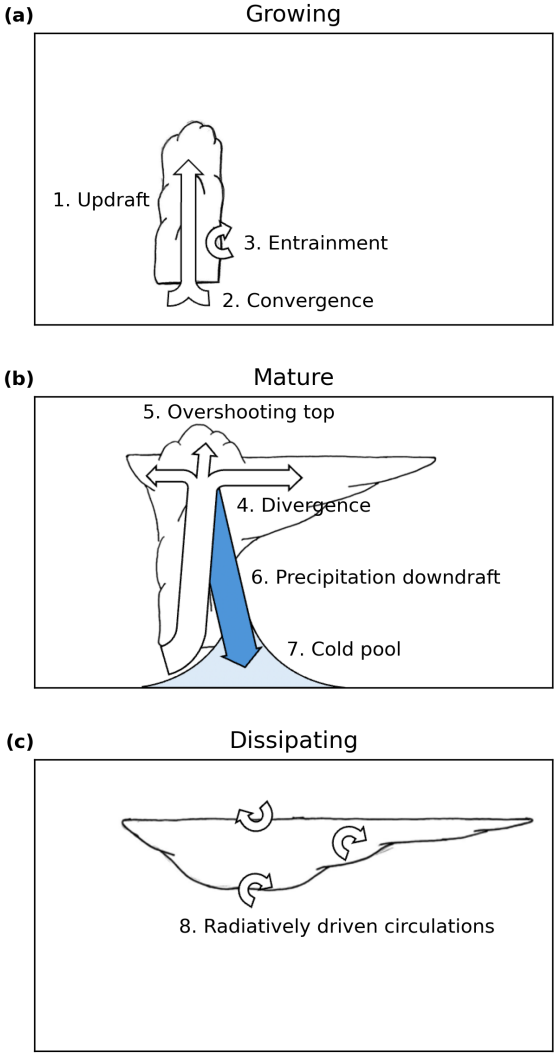


Figure 1.3: An illustration of the three lifecycle phases of an isolated DCC: the growing (a), mature (b) and dissipating (c) phases. In the growing phase, the formation of a convective updraft (1) through positive buoyancy and convergence of warm, moist air at the cloud base (2) leads to a convective core that rapidly increases in height. Turbulent mixing with the surrounding environment (3) dilutes the updraft, reducing the buoyancy and hence vertical velocity. In the mature phase, the updraft reaches its LNB resulting in divergence (4) and the formation of an anvil cloud. The momentum of the updraft within the core may result in it reaching the stratosphere, leading to an overshooting top (5). The onset of precipitation leads to downdrafts (6) due to hydrometeor drag and evaporative cooling, compensation for the downdraft. At the surface, this causes cold pools (7) which disrupt the inflow of air to the base of the storm, preventing further convective activity. In the dissipating phase there are no more convective dynamics, but radiative driven circulations (8) act to thin the thicker anvils due to entrainment of surrounding dry air, but increase the lifetime of thinner anvils.

tropopause and extends into the stratosphere despite experiencing negative buoyancy. As the cloud will continue to cool adiabatically as they rise, the tops of these updrafts can reach substantially colder temperatures than the surrounding environment [Proud and Bachmeier, 2021].

The mature stage occurs after the top of the DCC has reached the level of neutral buoyancy. The heaviest rates of convective precipitation occur during this stage, and the anvil cloud is formed [Houze, 2014]. The occurrence of heavy precipitation suppresses the convective core both through the generation of downdrafts and through the stabilisation effect of evaporating rain droplets. This process will eventually weaken and dissipate the convective core of the DCC unless the wind shear is large enough to advect the convective rainfall away from the convective core. Furthermore, the evaporation of falling precipitation leads to the divergence of cold, dense air near the surface. These cold pools may suppress further convective activity, while their gust fronts can also trigger new convective initiations.

The final stage of the DCC lifecycle occurs after the convective core has dissipated and convective rainfall has stopped, and is referred to as the dissipating phase. During this phase, the anvil cloud may continue to expand, with maximum anvil cloud extent occurring much later than maximum convective intensity [Futyan and Del Genio, 2007]. Additional, stratiform, precipitation may occur throughout the anvil cloud, but this will not be as intense as the earlier convective precipitation [Houze, 2014]. The dissipating stage may also be split into two separate periods; those before and after the end of stratiform precipitation [Wall et al., 2018].

DCCs can also be categorised spatially into three components. Firstly, the core region, in which the convective updraught and convective precipitation occur. Secondly, the anvil or cloud shield, which consists of a large area of thick cloud surrounding the core at the level of neutral buoyancy, and within which stratiform precipitation may occur. Finally, the area of cirrus outflow, where thin ice cloud extends beyond the

edge of the anvil cloud, particularly within the dissipating phase [Lilly, 1988]. This structure of an isolated DCC is also referred to as a 'hot tower' [Riehl et al., 1951].

Radiative heating of the anvil cloud base and cooling of the anvil cloud top causes two regions of instability to occur. These unstable profiles drive circulations at the cloud base and top which increase the entrainment of dry air into the anvil while also lifting the anvil higher in the troposphere. While SW heating during the daytime reduces the instability at the cloud top, the greater LW heating at the cloud base results in a greater net dissipation of anvil during the daytime [Sokol and Hartmann, 2020]. As the anvil cloud thins, these two regions of instability join into a single lapse rate. This state drives circulations within the anvil cloud which lift ice particles, reducing sedimentation and increasing particle growth [Gasparini et al., 2023]. These circulations prolong the lifetime of the thin anvil cloud [Sokol and Hartmann, 2020]. Further thinning of the cirrus anvil results in uniform radiative heating throughout the anvil, ending the circulation and leading the to complete dissipation of the anvil cloud.

1.2.5 Convective organisation

While the majority of DCCs form as isolated storms, with a single convective core that detrains an anvil cloud separated from its neighbours [Riehl et al., 1951]. In other cases, DCCs exist with multiple convective cores feeding a single anvil cloud [Zipser, 1969; Nakazawa, 1988]. At the largest end of this progression are mesoscale convective systems (MCSs), which occur when a cluster of deep convective cores form a single large area of anvil referred to as a cloud shield [Roca et al., 2017]. These MCS cloud shields can cover an area of greater than 10,000 km^2, several orders of magnitude greater than that of individual DCCs [Houze, 2004].

The convective cores that form MCSs do not group together randomly, but are instead clustered by a complex group of processes that are referred to as convective

organisation. In general, these processes tend to lead to the triggering of new convective cores that feed the same cloud shield, expanding and extending the life of the system as they do. MCSs occur in many shapes and sizes, including mesoscale convective complexes, squall lines and tropical cloud clusters [Tsakraklides and Evans, 2003] which each have unique properties. However, all MCSs display large-scale processes that initiate new convective updrafts, and increase the size and lifetime of the anvil.

The lifecycle of MCSs can be described similarly to that of isolated DCCs through the three-phase growing, mature, dissipating structure [Futyan and Del Genio, 2007]. It should be kept in mind however that when applied to MCSs, these lifecycle stages do not represent the changes in convective processes as described for isolated DCCs, but the mesoscale development and decay of the entire system. The lifetime of these systems is substantially lengthened, with typical MCSs lasting for 10 to 20 hours or longer [Chen and Houze, 1997]. In particular, there is an increase in the lengths of the initiation and mature phases compared to isolated DCCs [Wall et al., 2018] due to the continuous development of new cores throughout the active lifetime of the MCS. The total lifetime of MCSs is determined by the rate at which new cores are triggered, and observations have shown that they tend to decay once the area of convective cores falls below 10% of the total anvil area [Elsaesser et al., 2022].

A key feature that distinguishes MCSs from smaller DCCs is the development of large, mesoscale circulations that occur on larger scales than that of individual convective cores. As well as the convergence of heat and moisture at low levels due to the convective cores, this mesoscale circulation draws in additional air from the surrounding areas at the mid-levels of the atmosphere. This mid-level air does not enter the cloud shield through the convective cores, but instead condenses at higher altitudes, leading to a top-heavy latent heating profile within the cloud shield [Schumacher et al., 2004]. These organised convective cloud systems have thermodynamic impacts

on the environment to a much greater extent than isolated DCCs. The large cloud shields of the MCS s result in much larger amounts of stratiform precipitation than individual DCCs, which is distributed over a much wider area [Houze, 2014]. These organised convective cloud systems have thermodynamic impacts on the environment to a much greater extent than isolated DCCs, with idealised simulations showing that the thermodynamic interactions of organised convective systems can propagate thousands of kilometres within the troposphere [Beucler and Cronin, 2019]. Over time, the mesoscale circulation results in a moistening of the convective system and a drying of the surrounding atmosphere, creating a sharp contrast between the two regions [Bretherton et al., 2005].

MCSs are also characterised by both high rates and large volumes of precipitation, providing the majority of precipitation throughout the tropics as well as in many other regions [Feng et al., 2021]. Unlike isolated DCCs, much of the precipitation occurs not as convective rainfall, but as stratiform rain from the cloud shield [Schumacher and Houze, 2003]. As much as 70 % of the total precipitation of an MCS can come from the stratiform rain, with a higher proportion observed over the ocean than in MCSs over land. Much of this stratiform precipitation can be attributed to the mid-level, mesoscale inflow to the cloud shield as well as the top-heavy latent heating profile.

The processes through which MCSs form in the first place are still uncertain, however [Houze, 2018]. Convective organisation processes have been observed in satellite remote sensing, radiative-convective equilibrium models and cloud-resolving models [Holloway et al., 2017]. The formation and lifetime of MCSs are strongly linked to the dynamics of the surrounding environment through the convergence of moist, low-level air and the divergence of air at the top of the MCS [Houze, 2014]. Cold pools have a strong influence on the organisation of convective cores, and set the scales over which organisation occurs [Jeevanjee and Romps, 2013], as their gust fronts, transport of moisture and collisions are vital to triggering new convection

within MCSs [Feng et al., 2015]. MCSs tend to trigger additional convection around the storm, blurring the boundary between the MCS and isolated convection [Mapes, 1993].

MCSs are, in general, poorly represented in large-scale climate models due to the inability of the parameterised convection to organise between grid squares [Houze, 2018]. While convective resolving models have been proposed as a solution to this issue [Stevens et al., 2020], these models still do not represent MCSs accurately with a tendency towards too little organisation [Prein et al., 2021]. Gaining a better understanding of the properties and processes of MCS, both in observations and models, is vital therefore for improving their representation [Feng et al., 2023b]. Changes in convective organisation have important effects on the global climate, with increases in organisation correlated with domain-wide ToA cooling in satellite observations [Bony et al., 2020], and so future changes in organisation may have a large influence on the equilibrium climate sensitivity.

1.3 Anvil Radiative Feedbacks

There are a number of hypotheses regarding the CRE of tropical anvil clouds that consider whether the neutral CRE of tropical anvils is the result of a feedback mechanism. Ramanathan et al. [1989] proposed the thermostat hypothesis in which, in response to a warming environment, anvil clouds produce thicker cirrus which acts to cool the tropics through increased SW reflectance. The Iris hypothesis proposes that anvil cirrus will decrease in area, resulting in greater LW emission from the surrounding clear-sky regions. Lindzen et al. [2001] first proposed this as a result of increased precipitation efficiency, however evidence for this effect is disputed [Del Genio and Kovari, 2002; Lin et al., 2004]. Bony et al. [2016] proposed a 'stability iris' feedback, in which the established trends of increased dry static stability [Held and Soden, 2006] and a reduction in the tropical overturning circulation [Vecchi and Soden, 2007] re-

duce the detrainment of anvil cirrus. Although the anvil cloud response is generally considered to be a negative climate feedback, the predicted magnitude varies widely and it represents the greatest uncertainty among all cloud feedbacks [Sherwood et al., 2020].

On the other hand, the fixed anvil temperature (FAT) hypothesis argues that the anvil cloud top temperature (CTT) remains constant in a warming climate, and the greater difference between anvil and surface temperatures results in a positive LW feedback [Hartmann and Larson, 2002]. The basis for fixed anvil temperature (FAT) is that LW cooling of the troposphere due to water vapour becomes inefficient below 220 K [Jeevanjee and Fueglistaler, 2020], which, if relative humidity remains constant, fixes the top of the convectively active troposphere at this isotherm. While there is evidence that this is the case for the largest DCC anvils, the increase in static stability may result in a reduced positive feedback due to a 'proportionally higher' anvil temperature [Zelinka and Hartmann, 2010] which more closely matches the LW response of tropical clouds in global climate models. While satellite observations have shown a trend in anvil cloud height [Norris et al., 2016], there is not yet sufficient evidence to distinguish this from inter-annual variability [Takahashi et al., 2019]. Seeley et al. [2019b], argued that FAT is a weak constraint on anvil temperature as while the radiative tropopause temperature remains fixed, the temperature of the tropopause lapse rate inversion can vary widely. Furthermore, as anvils tend to detrain below the tropopause, [Takahashi et al., 2017a; Wang et al., 2020], anvil temperature and the tropopause temperature may only be weakly connected. Seidel and Yang [2022] however found the inclusion of CO_2 radiative heating produces anvil temperatures consistent with FAT.

While the iris and FAT feedbacks may act to cancel each other out, and hence maintain the neutral CRE of tropical anvil clouds, other important feedback mechanisms may influence this balance. Hill et al. [2023] recently showed that climate mod-

els underestimate dynamically driven cloud feedbacks including changes in convective intensity. Furthermore, convective instability is expected to scale with temperature in the same manner as the Clausius-Clapeyron relation [Seeley and Romps, 2015b; Agard and Emanuel, 2017], and some observations of tropical anvil clouds have instead suggested that warming of the surface invigorates convection [Igel et al., 2014]. This invigoration effect may result in colder anvil cloud top temperature (CTT), and hence a stronger warming feedback. Multi-decadal satellite observations have shown a cooling of upper tropospheric cloud temperature over land [Liu et al., 2023], indicating that changes in convective processes may lead to stronger cooling feedbacks.

A recent criticism of the Sherwood et al. [2020] decomposition of anvil feedbacks into the area response and the height response is that, by ignoring changes in the optical properties and structure of anvil clouds, these feedbacks are merged into the calculated iris effect, increasing its uncertainty [Raghuraman et al., 2024]. McKim et al. [2024] proposed that the anvil area feedback of $-0.40\,\mathrm{Wm^{-2}K^{-1}}$ from Sherwood et al. [2020] would require an unrealistic anvil area change of $-20\,\%\mathrm{K^{-1}}$, indicating that there are other mechanisms contributing to the measured change in anvil CRE. Recent work has proposed that the anvil cloud response reduces the area of thicker anvils with high ice water content (IWC), but not thinner cirrus, resulting in a much smaller overall feedback [McKim et al., 2024; Sokol et al., 2024].

Changes to the lifecycle and diurnal cycle of deep convection may also be an important factor, particularly when considering the SW feedback. Nowicki and Merchant [2004] used estimates of top-of-atmosphere (ToA) LW and SW radiative fluxes from Spinning Enhanced Visible Infra-Red Imager (SEVIRI) observations to estimate the diurnal cycle of anvil CRE over equatorial Africa and the equatorial Atlantic. They found that shifting the diurnal cycle of deep convection in these regions could change the CRE by $\pm 10\,\mathrm{Wm^{-2}}$, but did not track the properties of individual DCCs. Bouniol et al. [2016] compared CRE and cloud radiative heating rates to anvil cloud properties

to investigate how radiative heating affects the anvil cloud evolution. These observations were made with polar orbiting instruments however, and they highlighted the need for geostationary observations to characterise the evolution of individual anvil clouds. Subsequent research used DCC tracking methods to better characterise the lifecycle of observed anvil clouds [Bouniol et al., 2021], but as the radiative flux data was provided by polar-orbiting satellites the CRE could not be measured over the lifetime of the DCC.

1.4 Detection and tracking of deep convection

The detection and tracking of clouds has been performed since the earliest sequences of remote sensing imagery from weather radar and geostationary satellites [Menzel, 2001]. Fujita et al. [1968] compared sequences of images observed by the first geostationary weather satellite to those taken using an all-sky camera, and found that by comparing subsequent observations from the satellite one could calculate AMVs similar to those observed on the ground. This tracking of cloud position was performed by hand, and subsequently a plastic stencil 'computer' was designed to calculate cloud velocities taking into account the satellite viewing geometry [Fujita, 1969]. While these early methods compared print-outs of satellite imagery by hand, shortly after a digital computer system was developed to show sequences of images [Chang et al., 1973]. Although detection and tracking were still performed manually (albeit with the user selecting cloud positions in subsequent images using the cursor), velocity calculation was performed automatically. Wider adoption of this technology was applied to produce AMVs during field campaigns in the mid- to late-1970s [Tecson and Fujita, 1975].

There was, however, concern regarding the manual tracking of cloud velocity. This task was both time consuming and also open to subjective judgement which made uncertainties hard to estimate. Fujita et al. [1975] found large variations in AMVs

produced using these methods. Early efforts at automation applied cross-correlation techniques, previously used with weather radar, to estimate AMVs in geostationary satellite imagery [Leese et al., 1970]. Endlich et al. [1971] applied a pattern matching technique to 'brightness centres' in visible satellite imagery to estimate cloud motion. Rinehart and Garvey [1978] produced a cross-correlation algorithm for the estimation of convective cell motion in 3-D weather radar observations. While these automated methods provided more accurate motion estimates than humans, they were less capable of detecting independent cloud motions both due to errors introduced by noise, and also due to the visible imagery provided by early geostationary satellites meaning it was difficult to distinguish clouds at different altitudes.

Automatic methods for the detection and tracking of DCCs were developed for both radar reflectivity [Crane, 1979] and geostationary satellite LW infrared (IR) brightness temperature (BT) [Endlich and Wolf, 1981]. Both of these methods detected DCC features by labelling the area surrounding local extrema (maxima for radar reflectivity, minima for BT) in individual images. The addition of 11 µm IR-window BT channels to the first operational GOES and Meteosat weather satellites allowed DCCs to be distinguished from low clouds. The centroids (the location of the feature represented as a single point) of these detected features were then linked to create DCC tracks through the use of cost-minimisation algorithms with sought to find the best match between pairs of features detected at subsequent time steps. By explicitly detecting features to track, these algorithms both improved upon the weakness of previous algorithms, but also allowed the properties of tracked objects to be studied beyond their motion vectors.

The development of DCC detection and tracking algorithms continued in parallel for both radar and satellite observations. Rosenfeld [1987] and Williams and Houze [1987] both developed overlap tracking techniques for radar reflectivity and satellite BT observations respectively. Unlike the prior centroid-based tracking methods, these

overlap methods took into account the spatial extent of detected features by linking subsequent pairs of features based on those which shared the largest overlapping area. By removing the approximation of a point-like feature, the overlap techniques better handle cases of DCCs with larger, more complex shapes. The downside to this approach, however, is that, unless the spatial extent of the DCC is propagated using an estimate of the storm motion, it requires that the detected features do not move further than their diameter between subsequent observations. As a result, for small, fast-moving objects (such as convective cells), or observations that are more widely spaced, overlap methods perform poorly.

Early detection and tracking algorithms were strongly limited by the available computational power which restricted tracking to only a small number of DCCs. Furthermore, the accuracy of early algorithms was such that human verification of tracked objects was still required. The majority of algorithms were designed with only a single source of observations, which, in particular for weather radars, reduced the area over which tracking could be performed. While geostationary satellites provided larger areas of observations, the low spatial and temporal resolution of early sensors limited studies to large and long-lived MCSs only. Maddox [1980a] tracked mesoscale convective complexes over the US, and subsequent studies showed their distributions globally [Laing and Michael Fritsch, 1997].

Subsequent development produced trends in both those algorithms used to track DCCs in satellite imagery, and those using weather radar. Algorithms using geostationary satellite imagery focused on the study of MCSs, and tended to use overlap-based tracking methods [Arnaud et al., 1992; Evans and Shemo, 1996; Carvalho and Jones, 2001; Morel and Senesi, 2002]. On the other hand, tracking algorithms using radar reflectivity tended to focus on the tracking of convective cells, and favoured centroid-based tracking approaches [Dixon and Wiener, 1993; Johnson et al., 1998; Handwerker, 2002]. Across both sets of algorithms, the use of fixed thresholds rather

than extrema for locating features became favoured due to the low computational cost, ease of customisation and resilience to noise [Augustine and Howard, 1988]. While the majority of approaches used a single threshold, some algorithms began using multiple thresholds to better characterise the detected features of DCCs. Johnson et al. [1998] used a sequence of increasingly larger thresholds to better distinguish individual convective cells, and classify them according the intensity. For MCS detection, the 'detect-and-spread' approach was used by Evans and Shemo [1996] and Boer and Ramanathan [1997] to better measure the area of MCSs by first detecting the 'core' using a cold BT threshold and then detecting the surrounding cloud area using a warmer threshold. It should be noted however that Augustine and Howard [1988] argued that the area estimated using the warmer threshold is highly subjective.

Hodges [1994] developed a more general algorithm for use with multiple gridded datasets, including model data, and subsequently expanded this approach to consider detection and tracking on a sphere for use with GCMs [Hodges, 1995]. The low spatial and temporal resolution of GCMs of the time however meant that detection and tracking of individual DCCs was not possible. Other data, such as rainfall measurements and lightning flash observations, were also used to develop detection and tracking algorithms [Steinacker et al., 2000], however the vast majority of algorithms continued to use either radar reflectivity or geostationary satellite BT observations.

The development of radar detection and tracking algorithms was strongly driven by the need for nowcasting (short-term, 30- to 60-minute forecasts) of convective activity [Wilson et al., 1998]. As these algorithms are required to predict the future motion of observed DCCs, good estimates of the storm motion, including growth and decay, are required. Development of these algorithms, therefore, focused primarily on the tracking aspect rather than the detection of DCCs [Lakshmanan and Smith, 2010]. The Sydney 2000 Olympics provided both a demonstration of, and a comparison between, the capabilities of multiple tracking algorithms [Keenan et al., 2003].

The study into the performance of these algorithms found that no particular tracking approach worked best overall, and that each had strengths and weaknesses in different situations [Wilson et al., 2004]. In response, subsequent algorithms used hybrid approaches which combined multiple tracking methods to provide better performance over a range of scenarios [Lakshmanan et al., 2007; Han et al., 2009].

The second generation of geostationary weather satellites provided greater capability to track individual DCCs due to both their higher spatial and temporal resolutions and increased number of channels across the visible and IR spectrum. Roberts and Rutledge [2003] showed that these observations could be used to detect initiating DCCs up to 30 minutes before they appear in weather radar observations. Because of their use of a single BT threshold, existing methods could only track DCCs once they had matured. While this was generally acceptable for tracking MCSs, this limited their capability for tracking isolated DCCs. Mecikalski and Bedka [2006] developed an algorithm which combined visible imagery, estimates of BT cooling rate over multiple observations and AMVs to indicate convective initiation. Zinner et al. [2008] used multiple IR BT channels from the Meteosat SEVIRI imager, along with the high resolution visible channel and motion vectors derived using a pyramidal matching algorithm to detect and track developing DCCs over multiple stages from initiation to maturity.

A common theme throughout all the algorithms described so far is the separation of feature detection and tracking into separate procedures, with features detected independently at each time step. While this allows for simplification of the overall process, the independent detection of features at each time step introduces a large degree of inconsistency in the area and location of features detected over time. Furthermore, when using a fixed threshold, this also prevents the detection of features before or after they reach the threshold, limiting the detection of growing or decaying DCCs. Some recent algorithms have sought to address this by combining both feature

detection and tracking into a single process. Fiolleau and Roca [2013] applied the 'detect-and-spread' approach to a '3-D' stack of images over time. By first detecting the core region using a cold BT threshold, and then detecting the surrounding cloud volume using a warmer threshold, the algorithm is better able to detect the growing and decaying phases of DCCs without erroneously detecting warmer clouds. However, the algorithm does not consider the motion of the tracked objects in any manner, and so is only suitable for detecting MCSs whose areas are large compared to the distance moved between observations. Thomas et al. [2010] used a variational-data-assimilation model to both detect and track DCCs in geostationary satellite images. This approach not only allowed detection and tracking as a single process, but was also resilient to cases of noisy or missing data.

Improvements in climate modelling have led to new applications of cloud detection and tracking algorithms in cloud-resolving models [Plant, 2009], large eddy simulations [Dawe and Austin, 2012; Heus and Seifert, 2013] and for the evaluation of GCMs [Clark et al., 2014]. A wide range of modern algorithms have been developed specifically for general application to radar, satellite and model data [Heikenfeld et al., 2019; Ullrich and Zarzycki, 2017; Ullrich et al., 2021; Raut et al., 2021; Feng et al., 2023a]. These new methods have allowed studies into the differences in DCCs between observations and models, as well as the response of DCCs to perturbations across multiple models [Marinescu et al., 2021; Feng et al., 2023b]. Furthermore, the flexibility of these general purpose approaches has allowed their application for the detection and tracking of atmospheric phenomena beyond DCCs [Bukowski and van den Heever, 2021; Zhang et al., 2023].

A further focus of modern algorithm development has been for studying trends in the behaviour of DCCs observed in long time series of observations, with techniques optimised to these problems [Núñez Ocasio et al., 2020; Hayden et al., 2021]. Overall, these developments have allowed cloud tracking studies to move from smaller cases

to large data problems involving the properties of tens, if not hundreds of thousands of DCCs.

DCC detection and algorithms have made vast advances since the earliest approaches that supplanted human tracking. Through this development process, the requirements for a successful algorithm have become apparent. The detection process needs to accurately distinguish between DCCs and other clouds, while also detecting the largest possible extent and proportion of the DCC lifetime, while being consistent between time steps and resilient to noise. The first two aspects pose a challenge, as in general improving one of these sensitivities means worsening the other. The tracking approach needs to then accurately connect these features without mistakenly connecting or failing to track storms, taking into account the motion of each DCC as well as splits and merges. Many modern algorithm developments have only focused on addressing a few of these concerns, without taking into account developments made by other algorithms. Many of the same issues from older algorithms still exist. For example, many modern detection approaches using satellite BT only use a single 11 µm channel, despite the availability of many different channels providing additional information about DCC properties from modern imagers.

One of the key remaining challenges is the split in algorithms designed to track MCSs and those designed to track individual DCCs. The gap in capabilities between these two approaches makes it challenging, if not impossible, to study the behaviour of DCCs across a wide range of scales. Investigating these properties, therefore, will require further developments on top of those already made on DCC detection and tracking.

1.5 Structure of the thesis

Improving our knowledge of the processes influencing DCCs is vital to better understanding of their response to a warming world. Recent research has highlighted the

need to study the optical properties and structure of DCCs, and improve the connections between convective processes and anvil properties [Gasparini et al., 2023]. In addition, understanding how these properties change over the evolution of the anvil clouds has become an increasing focus of studies [Sokol and Hartmann, 2020; Wall et al., 2018; Bouniol et al., 2021]. The body of this thesis consists of four major chapters, through which we aim to answer several questions about the study and behaviour of DCCs. In chapter 2, a novel method for tracking both the developing cores and anvils of DCCs is developed, addressing some of the limitations of existing tracking methods for the study of DCCs across their entire lifecycle. This methodology is then applied to five years of geostationary satellite observations to create a dataset linking DCC properties, which is presented and analysed in chapter 3. Further insights into the impact of convective processes on anvil clouds can be gained using this dataset. In chapter 4, the response of anvil cloud structure to changes in convective intensity and organisation is studied, including changes across the observed lifetime of DCC. In chapter 5 we investigate the CRE of individual DCCs, how organisation impacts anvil CRE and how the effect of the observed population of DCCs produces the near–zero net CRE seen in previous studies.

Through exploring these topics, we aim to address gaps in the observational study of DCCs. While we cannot hope to "solve" the response of DCCs to warming, by answering these questions we can provide novel, key insights into poorly understood convective processes and their importance for climate change.

Chapter 2

A Semi-Lagrangian Method for Detecting and Tracking DCCs in Geostationary Satellite Observations

2.1 Introduction

Sequences of images from satellite instruments have been used to detect and track the motion of DCCs and tropical storms since the earliest geostationary weather satellites [Menzel, 2001]. Whereas early detection and tracking were performed by hand, numerous algorithms have been developed to perform this task automatically, and are widely used for both forecasting and research purposes (section 1.4). There is a continual effort both to improve existing algorithms and develop new methods to support these activities. However, it is important to understand the differences in observations of DCCs from satellite imagery to those of other sources, particularly radar and lightning observations.

Visible and IR imagery from modern geostationary weather satellite instruments provide unique observations of DCCs and their surrounding environment. Figure 2.1 compares observations of DCCs throughout three different stages of their lifecycle between satellite visible and IR imagery, doppler cloud radar and lightning flash observations. Composite red, green, blue (RGB) images from a combination of visible

and NIR channels aboard the Advanced Baseline Imager (ABI) show a.: small, isolated cores during the growing phase; b.: a large area of optically thick anvil during the mature phase, and c.: a large area of optically thin anvil cloud during the dissipating phase. BT imagery from the ABI 10.4 µm channel displays d.: cold, small, developing cores; e.: a large, cold anvil cloud, and f.: warmer BTs caused by thermal radiation from the surface penetrating the optically thin dissipating anvil. Lightning flash locations observed by the Geostationary Lightning Mapper (GLM) aboard GOES-16 show g.: low frequency during the growing phase; h.: high frequency during the mature phase, and i.: no lightning activity in the dissipating phase. Column mean radar reflectivity observed by next generation (weather) radar (NEXRAD) Doppler cloud radar shows high radar reflectivity in the convective cores during the j. growing and k. mature phases, and no area of high radar reflectivity during the dissipating phase. The outline of the region of BTs below 270 K observed by ABI is shown by the orange dashed contour over the GLM flash locations and NEXRAD radar reflectivity to indicate their observations relative to the anvil cloud.

These instruments are capable of observing the extent of the anvil clouds associated with DCCs over their entire lifecycle, even after convective activity has ceased (fig. 2.1 f). This is of particular importance due to the influence of anvil cloud radiative forcing on the climate, their response to temperature change [Bony et al., 2016; Hartmann, 2016; Ceppi et al., 2017; Gasparini et al., 2019] and possible feedbacks on subsequent convective activity. The newest generation of geostationary imaging satellites offers greater opportunities for the study of DCCs due to their high spatial and temporal resolution – allowing the detection and tracking of individual convective cores [Heikenfeld et al., 2019] – and also due to their high signal-to-noise ratio allowing research quality observations [Iacovazzi and Wu, 2020].

The detection and tracking of DCCs from satellite imagery remains challenging due to the inability to directly observe the convection that drives DCCs using passive

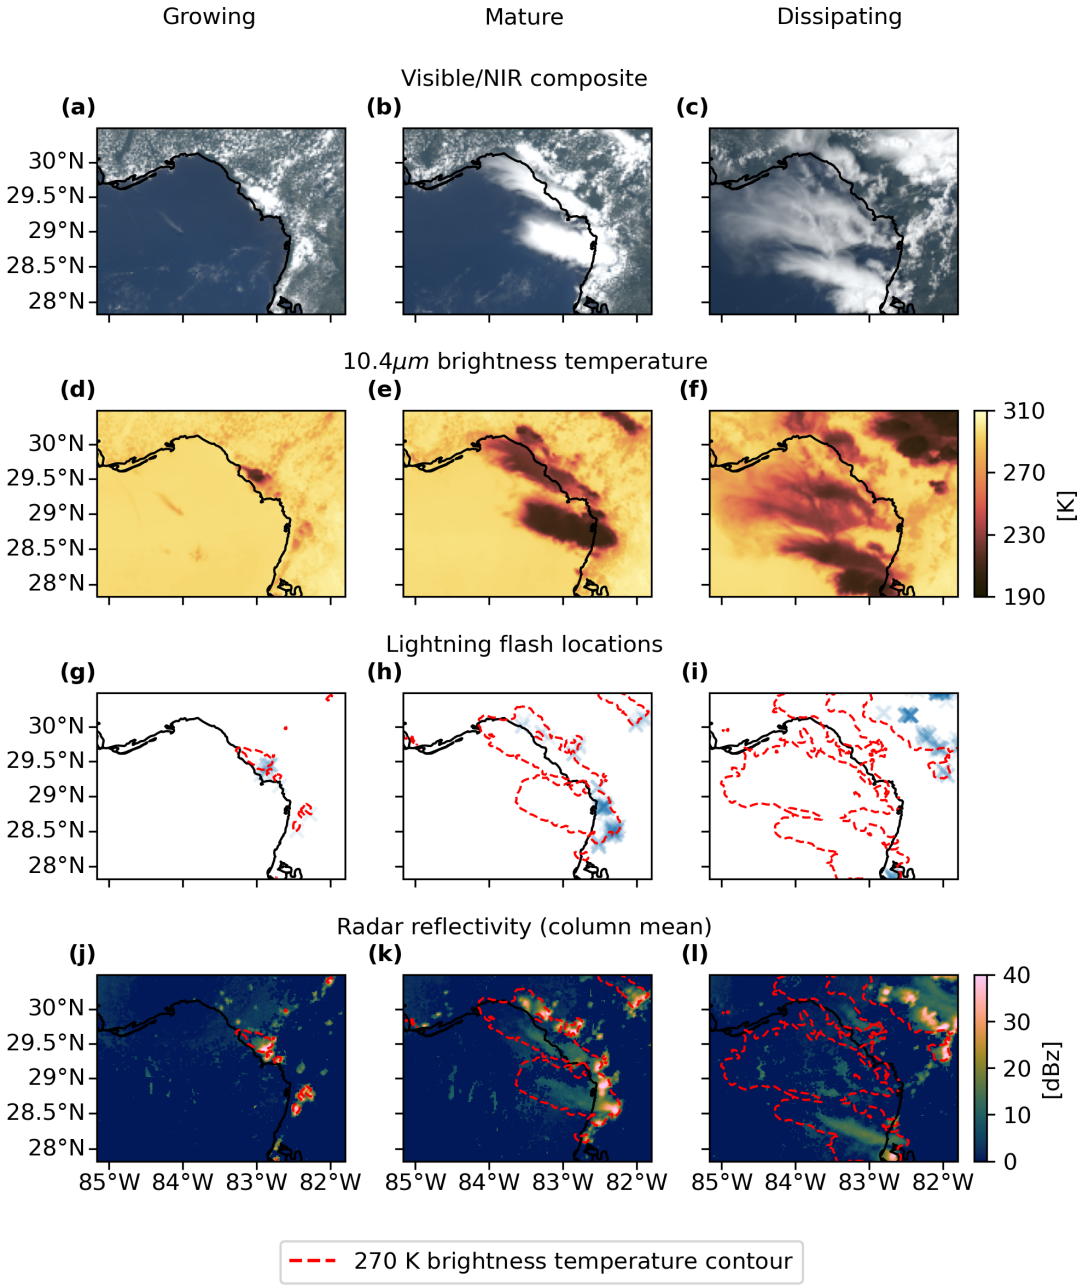


Figure 2.1: Observations of a cluster of DCCs over North-West Florida throughout three stages of their lifecycle. This cluster of DCCs occurred on the afternoon of 19th June 2018. The 'growing' column was observed at 17:00 UTC, the 'mature' column at 19:00 UTC, and the 'dissipating' column at 21:00 UTC. The four rows show observations of ABI visible/NIR composites (a,b,c), and 10.4 µm BT (d,e,f); lightning flash locations observed by GLM (g,h,i), and column mean radar reflectivity observed by NEXRAD (j,k,l). Note that, unless otherwise specified, this case study is used for all subsequent figures in this chapter.

visible and IR observations. This is unlike radar and lightning observations, which can directly observe deep convection due to the strong correlations between core updraft intensity and radar reflectivity and polarisation [Austin, 1987; Zipser and Lutz, 1994], and lightning flash occurrence [Williams et al., 1989]. Instead, a proxy for convective activity must be used to detect deep convection in visible/IR satellite imagery. The approaches used for this can generally be separated into two separate methods. Firstly, the use of thresholds on BT or other observed fields, which are capable of detecting DCC anvil clouds [e.g. Schmetz et al., 1997; Hong et al., 2005; Schröder et al., 2009]. Secondly, the detection of rapidly growing cloud tops by observing changes in the anvil cloud-top radiative cooling, or by other similar approximations of cloud growth [Zinner et al., 2008; Bedka et al., 2010].

Developing a detection method using either approach is made challenging by the dynamic nature of DCCs themselves. DCC cores typically have diameters of around 10 km, and updraft velocities on the order of $10\,\mathrm{ms}^{-1}$ [Weisman, 2015], and exist for 1-3 hours [Chen and Houze, 1997]. Large, mesoscale convective systems (consisting of multiple cores joined by a single large anvil [Roca et al., 2017]) may span areas several orders of magnitude larger than isolated DCCs [Houze, 2004], and typically exist to 10-20 hours or longer [Chen and Houze, 1997]. The life cycle of a DCC can be split into three phases: an initiation or growing phase, a mature phase and a dissipating phase after the cessation of convective activity [Wall et al., 2018]. There exists a significant difference between the diurnal cycles of deep convection over the land and over the ocean, with observed DCCs over land clustered towards the end of the day [Taylor et al., 2017].

The difficulties of detecting DCCs using various proxy approaches are demonstrated by the cross-sections of an observed DCC over time in fig. 2.2. The observed BT of the DCC anvil cloud shows wide variation over time, with the anvil cloud thinning due to dissipation after the end of convective activity which results in a warmer

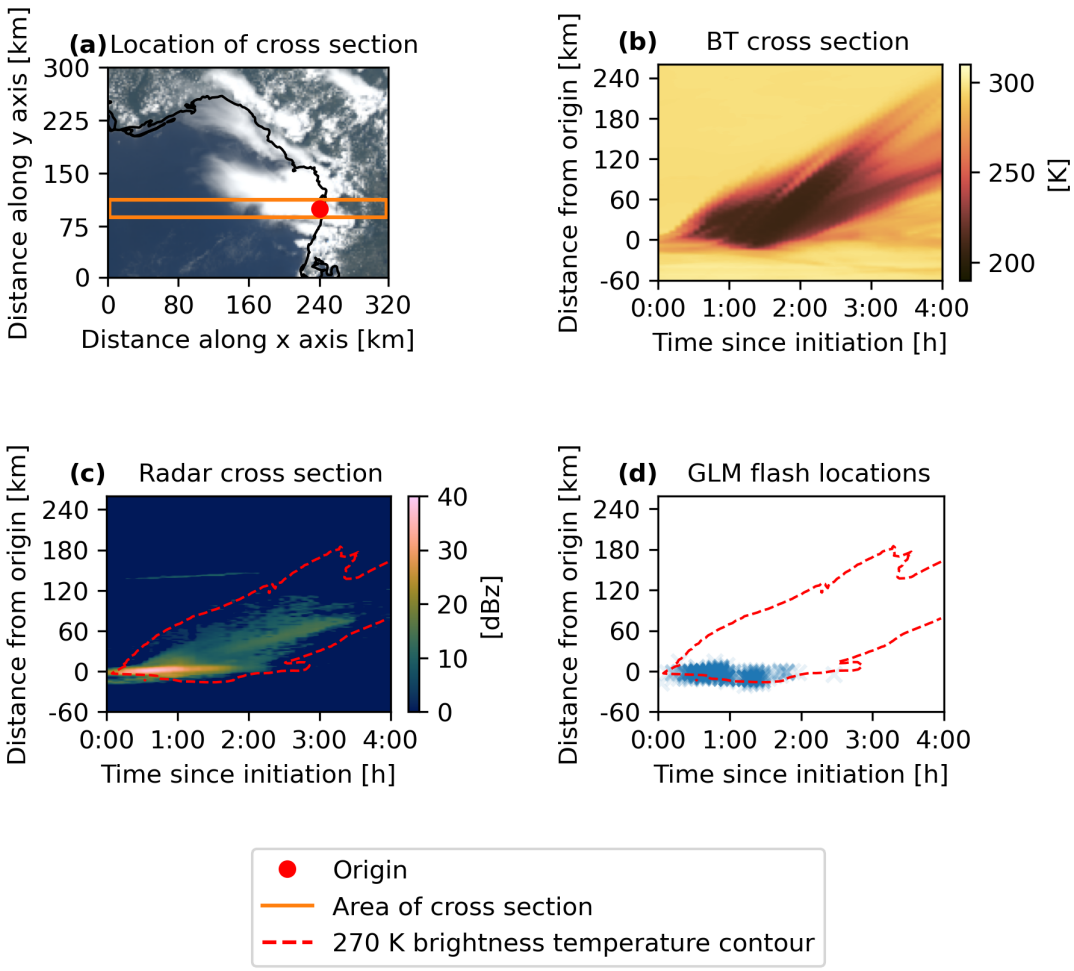


Figure 2.2: Cross sections of the DCC observed in figure 2.1 as they develop over time. a.: The location of the cross-section within the observed DCC. The mean of values is taken in the North-South axis. b.: ABI 10.8 µm BT, showing rapid cooling for the first 30 minutes, followed by an expanding region of anvil cloud that begins to thin and warm after 2-3 hours. c.: Column mean radar reflectivity, showing the presence and location of the convective core. d.: Lightning flash locations observed by GLM, which closely match the core observed by NEXRAD. Initiation occurred at 82.0 °W 28.5 °N at a time of 17:00 UTC.

BT due to greater signal from the surface. This variety of observed temperatures leads to large differences in the chosen threshold value between different algorithms [see discussion in Bennartz and Schroeder, 2012]. This choice of the threshold value is further complicated due to the overlap in observed BTs between DCC anvils and non-convective clouds [Konduru et al., 2013]. As a result, any detection method using a BT threshold must compromise between missed detection of DCCs, or false detections of non-DCC clouds.

The cooling of the cloud top is only visible for a short period during the initial phase of the DCC, before the anvil cloud top reaches the tropopause temperature after approximately 30 minutes. As a result, any method that solely relies on detecting the growth of the DCC will be unable to detect the anvil cloud after this initial growth phase has ended. While such algorithms provide accurate detection of these early phases of DCC growth [Zinner et al., 2013], they are unable to continue tracking the anvil cloud after convective activity is no longer observed.

Fiolleau and Roca [2013] identified this need to compromise on the accuracy of detecting DCCs as a problem caused by the commonly used two-step framework for detecting and tracking DCCs. In this framework, DCCs are first detected in individual images, and then linked together over time in sequences of images. As a result, the detection method chosen must be capable of detecting DCCs at each individual time step in order to track their entire lifecycle. Instead Fiolleau and Roca [2013] implemented a single-step framework for mesoscale convective systems that treats a sequence of images as a '3-D' volume (consisting of two spatial dimensions and one temporal dimension), and performs detection and tracking simultaneously by applying a watershed method over both spatial and temporal dimensions. Whereas this approach was successful for large, mesoscale systems, where the advection of the anvil is small compared to the overall anvil area, it is less capable of tracking small, rapidly moving convective cores. A semi-Lagrangian framework has been developed

which allows for single-step detection and tracking which accounts for the motion of DCCs using optical flow, improving the tracking of small DCCs.

Utilising the semi-Lagrangian framework, the best elements of both growth-based and threshold-based detection methods can be combined. It is possible to detect growing DCCs to a high degree of accuracy using methods similar to those of Zinner et al. [2008], and then extend the detected DCC over the entire anvil cloud using the 3-D watershed method of Fiolleau and Roca [2013]. This framework reduces the compromise required between the rate of missed DCCs and falsely detected DCCs, improving the overall accuracy of our detection method compared to existing approaches. Furthermore, this method allows the anvil cloud to be detected and tracked even after the region of cloud top cooling is no longer detected. Finally, the 3-D method handles the merging and splitting of intersecting DCCs by detecting all DCCs that intersect at any point during their lifetime as a merged object.

2.2 Data

Three sources of data are used throughout this chapter. Primarily, visible and IR imagery from ABI aboard the GOES-16 weather satellite is used for the detection of DCCs. Secondarily, observations from the NEXRAD weather radar network and the GLM (also aboard GOES-16) are used to assess and validate the tracking and detection method presented here.

2.2.1 Advanced Baseline Imager

The ABI is a visible and IR radiometer aboard the GOES-16 series of weather satellites [Schmit et al., 2016]. GOES-16, also known as GOES-East, is situated in a geostationary orbit at 75.2 °W above the equator, providing a field of view (or 'Earth-disc') covering most of the western hemisphere, including all of South America and most of North America. ABI has 16 channels operating in a range of spectral bands

Instrument	ABI	SEVIRI	Imager
Temporal resolution (minutes)	5	15	30
Nadir spatial resolution (km)	2	3	4 (8 for WV)
Number of IR LW window channels	3	2	2
Number of IR WV channels	3	2	1

Table 2.1: Comparison of data from ABI to that from older geostationary instruments SEVIRI aboard the second generation meteosat satellites, and the imager aboard the second generation GOES.

in the visible, NIR and thermal-IR. The majority of these channels have a resolution of 2 km at the sub-satellite point, although this reduces to approximately 3 km across most of the Continental United States (CONUS) due to the satellite viewing angle. ABI operates in a flexible scan mode, making observations of the CONUS once every 5 minutes, the full disc every 10 minutes (15 minutes prior to April 2019), and two mesoscale regions of approximately 2500 by 2500 km every minute. Additionally, it is capable of scanning the full disc every five minutes if no other scans are performed. This combination of high spatial and temporal resolution makes ABI suitable for detecting and tracking small and developing DCCs, as well as providing the spatial coverage to also track large mesoscale convective systems [Heikenfeld et al., 2019].

Compared to older geostationary instruments, ABI has higher spatial and temporal resolution, and more channels in both the LW IR window spectrum and the LW IR water vapour (WV) spectrum (table 2.1) [Iacovazzi and Wu, 2020]. This, combined with many of the channels being derived from those aboard the Visible Infrared Imaging Radiometer Suite, make the data from ABI more suitable for research purposes than that from older instruments [Heidinger et al., 2020]. Several artifacts are known to occur in ABI imagery [Gunshor et al., 2020]. Although the majority of these artifacts are removed using the data quality flag associated with the ABI data, in a number of cases bad detector stripes (described in section 3.2 of Gunshor et al., 2020) are not flagged in the data. These regions are detected separately, and replaced

with missing data prior to cloud tracking.

In this chapter ABI level 2 multi-channel cloud and moisture imagery product (MCMIP) is used, which provides calibrated reflectances and BTs for all ABI channels on a common grid [Schmit and Gunshor, 2020], using the 5-minute frequency imagery provided over the CONUS region. The case study shown in the figures throughout this chapter is for a subset of the CONUS scan region centred at 83.7 °W, 29.2 °N, over the time period of 17:00:00 to 21:00:00 UTC on the 19th June 2018. Validation was performed on the CONUS scan region over the entirety of 2018. All data has been sourced through the NOAA Big Data Program.

2.2.1.1 Selection of ABI Channels and Channel Combinations

In order to have an equal performance during both day and nighttime, a selection of LW IR ABI channels are used for the detection and tracking of DCCs (see fig. 2.3). These channels consist of the 10.4 µm and 12.4 µm LW (also known as the clean and dirty window channels respectively due to the presence of WV in the latter), and the upper and lower troposphere WV channels at 6.2 µm and 7.3 µm respectively. Whereas the LW window IR BT is commonly used for the detection of anvil clouds using threshold-based methods, it is not used for this purpose in this method due to the wide range of BTs observed within anvil clouds, and the variance of anvil cloud temperature with changes in tropopause temperature due to meteorology and latitude. However, the information contained within this field is used for the optical flow calculation of the cloud motion field.

Two additional combinations of channels are used to detect areas of DCC anvil. The water vapour difference (WVD) combination (fig. 2.3 c) of the upper troposphere WV channel minus the lower troposphere WV channel has been shown to provide a high detection rate for DCCs [Müller et al., 2018, 2019]. In clear sky or low cloud conditions, WVD shows the temperature difference between the upper and lower troposphere of generally around -20 K. While the 6.2 µm has an additional contribution

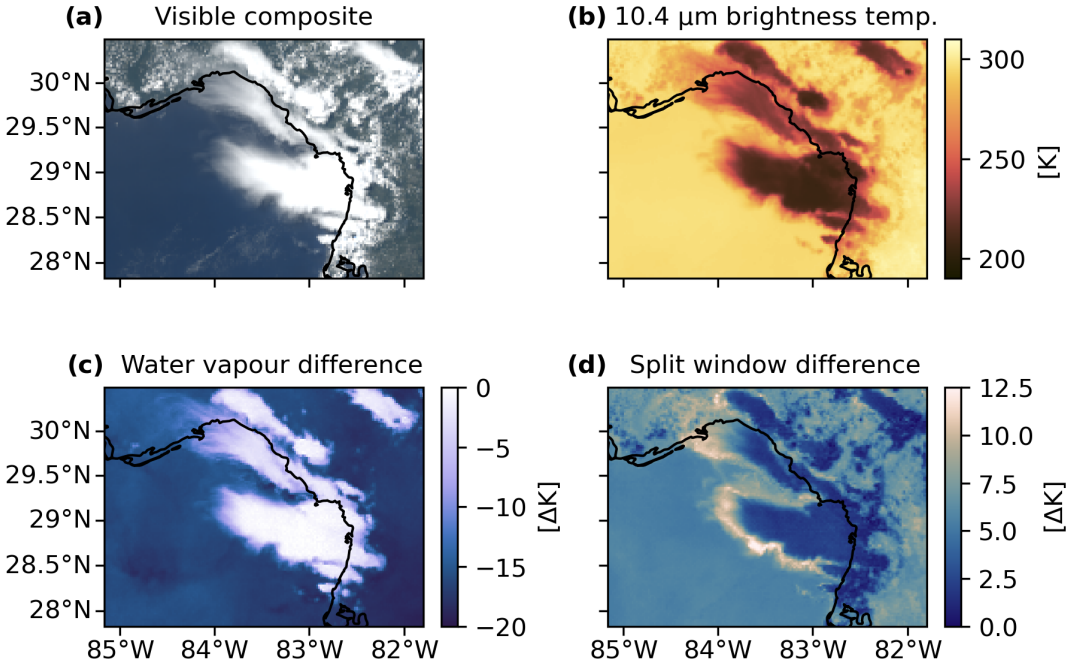


Figure 2.3: ABI channels and channel differences used with the detection and tracking algorithm. a.: A composite of visible and NIR channels. b.: The 10.4 µm BT, or clean LW window channel, which can differentiate clouds at all altitudes by their BT. c.: the WVD combination, of the 6.2 µm upper troposphere WV channel minus the 7.3 µm lower troposphere WV channel, which is strongly negative for clear sky and low cloud, but approaches positive values for thick, high clouds. d.: the SWD combination of the 10.4 µm clean LW window channel minus the 12.4 µm dirty LW window channel, which is near zero for thick clouds, around K for clear skies and approximately 10 K for thin, ice clouds.

from stratospheric WV [Schmetz et al., 1997], this does not have a significant effect on the WVD due to the small size of this absorption (fig. 2.4. Because both the WV channels are strongly absorbed by WV in the lower troposphere, the WVD field is not affected by surface and low altitude features and so provides a clear distinction between thick, high clouds and the background across a wide range of situations. Müller et al. [2019] found that a threshold of -5 K gave a high detection rate of anvil clouds. Furthermore, as the WVD values are relative to the lower stratosphere temperatures, this field is much less affected by location and meteorology than the LW IR channels. However, the WVD is still prone to the false detections of non-convective

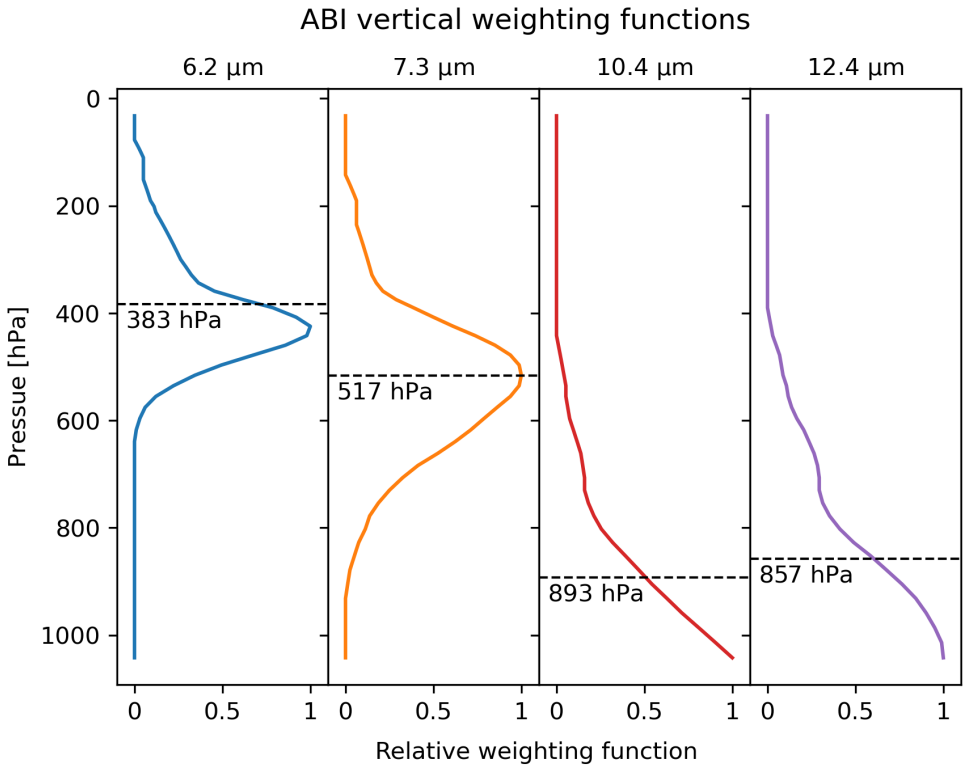


Figure 2.4: Clear sky vertical weighting functions for the 6.2, 7.3, 10.4 and 12.4 µm ABI channels calculated from rawinsonde profiles measured on 12:00:00 UTC 2020/08/10 at the Tampa RAOB station (KTBW - 72210). The dashed lines and associated pressure values show the weighted average emission height for each channel. Data from https://cimss.ssec.wisc.edu/goes-wf/plot-viewer/#/plot-viewer/plot/raob/abi16/default/20200810_1200Z/72210

clouds when using a thresholding method as it cannot directly distinguish between thick, high-altitude clouds that are associated with deep convection and those that are not.

The split window difference (SWD), consisting of the clean IR window channel minus the dirty IR window channel (fig. 2.3 d), aids in the detection and separation of optically thin anvil cloud (including cirrus outflow) from optically thick anvil due to the difference in ice particle emissivity between these two channels [Heidinger and Pavolonis, 2009]. As a result, this combination displays warm temperatures of around 10 K for thin, ice clouds, near 0 K for thick clouds, and approximately 5 K for clear

Instrument	wavelength	NEΔT [300 K]	NEΔT [220 K]
ABI (GOES–16)	6.2 µm	0.013 K	0.117 K
	7.4 µm	0.024 K	0.137 K
	10.4 µm	0.028 K	0.082 K
	12.4 µm	0.023 K	0.052 K
SEVIRI (MSG–11)	6.2 µm	0.05 K	0.45 K
	7.3 µm	0.05 K	0.29 K
	10.8 µm	0.06 K	0.17 K
	12.0 µm	0.10 K	0.24 K

Table 2.2: Operational NEΔT for the ABI and SEVIRI instruments. NEΔT is shown both at 300 K and at 220 K. ABI data is from https://www.star.nesdis.noaa.gov/GOESCal/G16_ABI_INST_CAL_daily_allmode.php; SEVIRI data is from Pili et al. [2016].

skies due to the contribution of boundary layer WV. The SWD is, however, also sensitive to low-level clouds and low-level WV concentrations, and so cannot be used alone to detect DCCs. It remains important to consider the SWD field due to the difficulty in separating anvil clouds from cirrus when using LW IR BT alone [Hong et al., 2005]. By subtracting the value of the SWD from the WVD, the sensitivity of the detection scheme to thin cirrus clouds is decreased, reducing the rate of erroneous detections. Additionally, adding the SWD field to the WVD field can enhance the appearance of cirrus, enabling the detection of thin ice clouds associated with cirrus outflow and dissipating anvils.

Figure 2.4 shows the clear sky vertical weighing functions for the 6.2, 7.3, 10.4 and 12.4 µm ABI channels. The weighting functions are calculated using a radiative transfer model with temperature and humidity profiles observed by rawinsonde observations on 12:00:00 UTC 2020/08/10 at the Tampa RAOB station (KTBW - 72210).

Measured, on-orbit instrument noise for the ABI and SEVIRI instruments are provided for each of the four channels selected for use in table 2.2. The noise equivalent ΔT (NEΔT) values are provided at a reference temperature of 300 K. However, as

NEΔT varies with temperature, it is also important to consider the noise at the colder temperatures typical of anvil clouds. NEΔT at 220 K were calculated by multiplying the NEΔT at 300 K by the ratio of the partial derivative of the Planck function at 300 K to that at 220 K. This results in a noticeably larger NEΔT at the lower temperature, in particular for the WV channels due to the difference in the gradient of the Planck function between the two temperatures. Propagating these uncertainties provides estimates of the noise in the WVD and SWD combinations of 0.180 K and 0.097 K respectively for ABI, and 0.54 K and 0.29 K for SEVIRI.

In addition to the NEΔT, there are several sources of systematic errors which may affect the observed BT and channel differences. Any bias in the ABI and SEVIRI channels are accounted for in calibration. Limb darkening at larger satellite zenith angles is harder to account for. This will have two impacts on the channel differences. Firstly, it will increase the height at which WV emissions are observed in the 6.2 and 7.4 µm channels, increasing the height at which clouds can be observed. Secondly, for the SWD combination it may reduce the difference between the two channels, as a higher zenith angle the path length through a cloud layer is longer for the same thickness of cloud, and therefore the difference in temperature will be less between the two channels.

In addition, environmental conditions may affect both the WVD and SWD combinations. Changes in humidity may affect the height at which the WVD is sensitive to clouds, with lower humidity resulting in sensitivity at lower cloud heights. In addition, the difference between the surface temperature and the cloud temperature will affect the SWD, as the temperature difference is due to a combination of cloud emission and surface emission with a ratio that varies across the two channels. As a result, if the difference between the cloud and surface temperatures is smaller, so too will be the SWD. While these biases cannot be removed from the data without retrieving the cloud properties, the choice of methods for cloud detection can limit

their impact on the results of detection and tracking.

2.2.2 Geostationary Lightning Mapper

The GLM is also mounted on GOES-16 and detects lightning flashes using an optical transient detector. The optical transient detector utilises a single, narrow-band NIR channel centred on 777 nm [Orville and Henderson, 1984] to detect momentary changes in brightness associated with lightning events at a frequency of 400 µs [Christian et al., 2003], providing a 70 % minimum efficiency of detection [Goodman et al., 2013]. GLM has the same field of view as the ABI instrument, albeit with a lower spatial resolution of 8 km at the sub-satellite point.

As lightning observations are strongly correlated with DCCs, data from GLM is used to validate the detection of DCCs using ABI. The level 2 GLM lightning cluster-filter algorithm product provides a dataset of events, groups and flashes processed from the GLM data [Peterson, 2019], and filters artifacts from the level 1 GLM data [Peterson, 2020]. From this dataset, detected flashes are extracted as evidence of DCC occurrence. These locations are then processed by mapping their frequency onto the ABI grid for validation of the algorithm.

2.2.3 Next Generation Radar

NEXRAD, also known by its technical name weather surveillance radar, 1988, Doppler (WSR-88D), is a network of weather radars operated by the National Weather Service across the USA [Crum and Alberty, 1993]. WSR-88D operates in the S-band spectrum, between 2700 and 3000 MHz. NEXRAD stations scan at a range of elevations, typically between 0.5° and 19.5° above horizontal, with a typical scan cycle taking between 4 ½ and 6 minutes, comparable to the temporal sampling of ABI over the CONUS region.

Cloud radar reflectivity is proportional to the droplet number density and the droplet radius to the sixth power, making it particularly sensitive to convective rain-

fall [Yau and Rogers, 1989]. As a result, cloud radar observations are ideal for showing the locations of convective cores [Austin, 1987; Rosenfeld et al., 1993; Zipser and Lutz, 1994], and in this chapter it is used to qualitatively assess our ability to detect developing convective cores using ABI. Level 2 NEXRAD radar reflectivity observations from multiple sites are gridded to the same resolution as ABI, and column mean reflectivity is calculated between the altitudes of 2.5 and 15 km.

2.3 Theory

To better understand how DCCs appear in GOES ABI observations, a series of experiments were performed using a radiative transfer model. These experiments were performed using libRadTran v2.0.4 [Emde et al., 2016], utilising the DISORT radiative transfer solver [Buras et al., 2011] and the REPTRAN absorbtion parameterisation [Gasteiger et al., 2014]. The experiments were run using the pyLRT wrapper for libRadTran [Gryspeerdt, 2024]. A tropical atmospheric profile was used to represent the atmospheric conditions under which deep convection typically occurs. A list of the options used to set up the radiative transfer model across all simulations is provided in table 2.3. All other options have been left as the default settings, including the treatment of relativity humidity which is set to 100 % within clouds.

In each of the simulations, a cloud layer is included in addition to these options to represent a DCC in different phases of the lifecycle. For liquid and ice cloud droplets the Hu parameterisation [Hu and Stamnes, 1993] and the Fu scheme [Fu, 1996; Fu et al., 1998] are used respectively. To show how the simulated clouds are observed by GOES-16, BT are simulated from the simulated radiances, and then integrated over the spectral response function of each of the ABI channels.

Option	Value	Description
rte_solver	disort	DISORT solver
source	thermal	thermal IR spectra
wavelength	3000–15000	3–15 µm
output_user	lambda edir eup uu	wavelength, direct, diffuse irradiance and all radiances
zout	0 TOA	surface and ToA
albedo	0.5	surface albedo
umu	-1.0, 1.0	downward & upward
sza	0	solar zenith angle
mol_abs_param	reptran fine	REPTRAN parameterisation
atmosphere_file	afglt.dat	tropical atmosphere profile
liquid microphysics	Hu	Hu and Stamnes [1993] parameterisation
ice microphysics	Fu	Fu [1996] parameterisation

Table 2.3: Selected options for the libRadTran simulations. All other parameters are left as defaults.

2.3.1 Observing growing convective cores

The first experiment aims to represent how a vertically developing DCC core would appear in ABI observations. A series of simulations were run with increasing cloud top heights between 1 and 15 km at 1 km intervals. These cloud layers are given a base height of 0 km, liquid water content (LWC) of 1,000 gm^{-2}, and droplet effective radius (r_e) of 15 µm. These values were chosen to represent typical conditions seen in a convective core. Although liquid droplets are used at all altitudes, simulations with ice cloud droplets showed negligible differences for cloud layers of this thickness.

Figure 2.5 shows BT spectra for increasing cloud height, with the spectral response functions of the ten thermal IR ABI channels plotted in the background. With such a large LWC, the BT matches the CTT apart from the absorption regions of CO_2 (4.2–4.5 µm and 13.2 µm) and ozone (9.4–10.2 µm), and the WV absorption for low-level clouds (5–8 µm).

The change in observed BT with height for the 6.2, 7.3, 10.4, and 12.4 µm BT channels along with the WVD and SWD is plotted in Fig. 2.6. Above 3 km, the 10.4

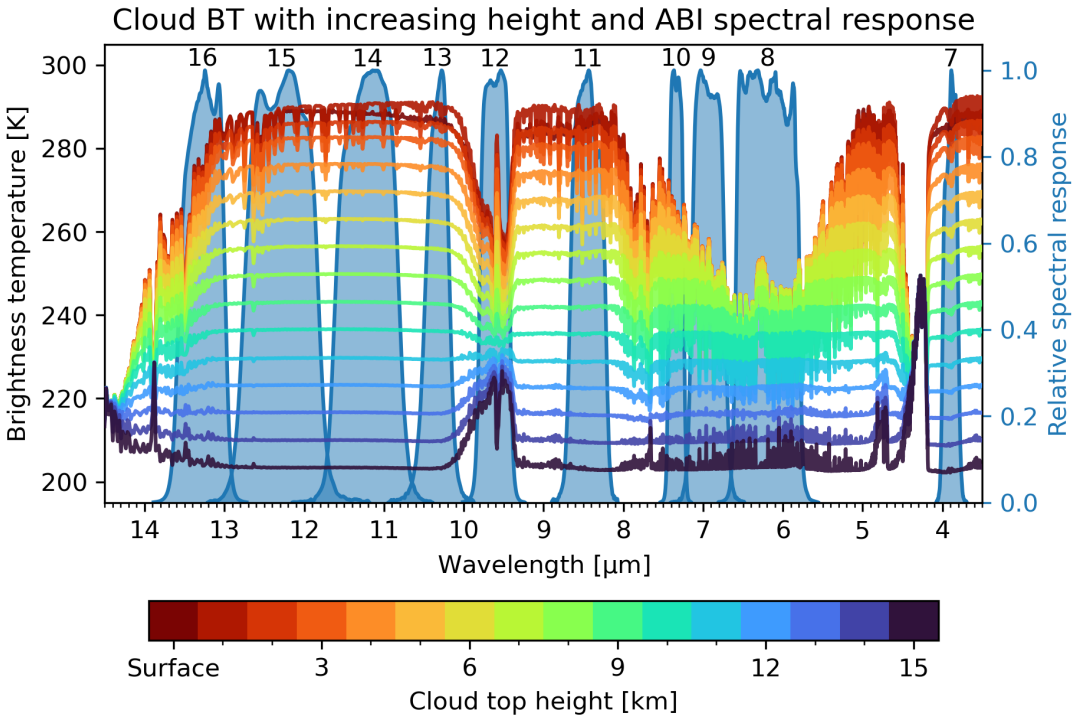


Figure 2.5: ToA BT spectra for a simulated cloud representing a DCC core at heights between 1 and 15 km, indicating that the largest difference in BT are seen in the LW window regions. The filled blue areas in the background show the relative spectral response functions for each of the LW ABI channels.

and 12.4 µm channels show a constant decrease observed BT at the moist pseudo-adiabatic lapse rate of 6 K. Due to the contribution of WV to the 6.2 and 7.3 µm channels at lower cloud top height (CTH), the WVD shows the most response to vertical development between 6 and 10 km. The SWD shows a response to clouds developing at a low altitude. It has been proposed that this can be used to detect the onset of convection in clear sky conditions before other observations, such as cloud radars [Lindsey et al., 2014, 2018].

Figure 2.7 shows the change of observed BT in units of K·minute^{-1} for a convective core rising at a rate of 1 ms^{-1}. Above 3 km, the 10.4 µm shows a consistent cooling of 0.5 K·minute^{-1}. Roberts and Rutledge [2003] found a threshold for severe convection of 8 K cooling over 15 minutes, or approximately 0.5 K·minute^{-1}, which

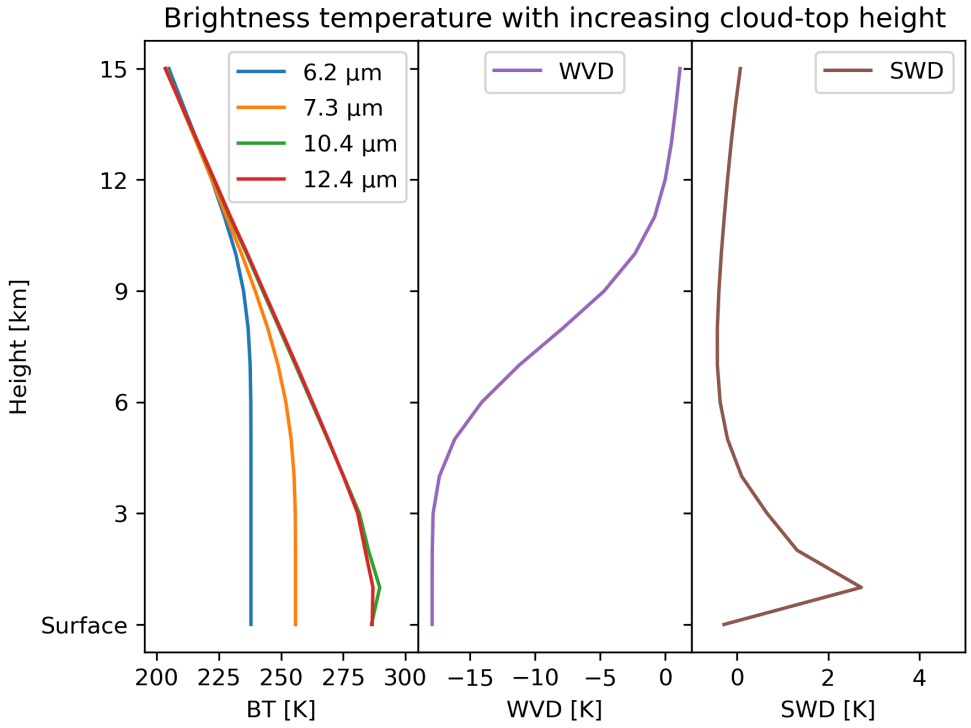


Figure 2.6: Simulated observations of the 6.2, 7.3, 10.4 and 12.4 µm ABI channels, WVD and SWD for a DCC core of increasing height.

would represent a cloud top vertical velocity of 1.25 ms^{-1} according to our simulation. The WVD shows a maximum rate of change between 7 and 8 km, with a warming rate of half the magnitude of the BT cooling rate.

Using the NEΔT values provided in table 2.2, the uncertainty due to sensor noise in the BT difference over a five–minute period is approximately 0.047 $\mathrm{K{\cdot}minute}^{-1}$, or just less than 10% of the threshold given by Roberts and Rutledge [2003]. By averaging the cooling rates over 5 pixels and a 15–minute time period the uncertainty is reduced to 0.012 $\mathrm{K{\cdot}minute}^{-1}$, or less than 2.5%.

Overall, this experiment shows that the 10.4 µm BT channel from ABI can be used to clearly identify a vertically developing core across a wide range of altitudes. On the other hand, the WVD is only sensitive to vertically developing clouds in a narrower range of altitudes, and shows a weaker rate of BT change. However, this sensitivity

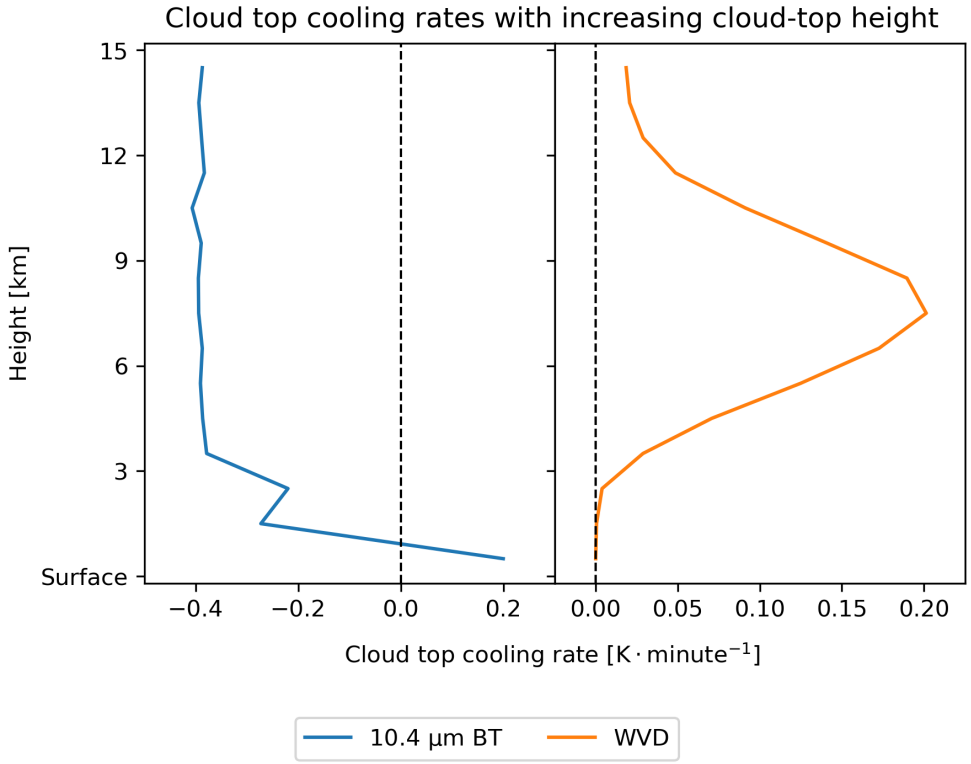


Figure 2.7: Simulated observed cooling rates of the 10.4 µm ABI channels and the WVD combination for a convective core growing at a rate of 1 ms^{-1}.

to vertical development at a mid to high altitude may help distinguish between DCCs and other, vertically developed clouds at lower altitudes, such as cumulus congestus.

It should be noted that these simulations were performed with the cloud temperature in equilibrium with the surrounding environment, this is not the case for a convective core which will be warmer than the surrounding air. Taking 6 K/km as a typical value for the moist pseudo–adiabatic lapse rate, the cooling rate for a convective cloud top rising at 1 ms^{-1} is 0.36 K · minute^{-1}. This is similar to the value shown in fig. 2.7 for the 10.4 µm BT. In addition, due to the mixing of surrounding air with the convective cloud top, it is expected that the observed temperature profile of the core will be somewhere between that of the environment profile and the moist pseudo–adiabatic lapse rate. In general, the temperature of the top of convective cores varies very little from the surrounding environment [Zipser and LeMone, 1980].

As overshooting tops also cool according to the moist adiabatic lapse rate, the BT difference method may also detect convective cores that overshoot underlying anvil clouds. While mixing with the warmer stratospheric air will reduce the cooling rate of the cloud top, the intense convection required to form overshooting tops means that they are still likely to be detected using the same thresholds. By basing further analysis of convective intensity on the most intense period of the observed cooling rate, the low bias on measurements of convective cores can be reduced. Detection of overshooting tops may also result in a bias on the number of cores associated with an anvil cloud, as more intense convection may result in detected overshooting tops whereas less intense convection may not be visible above the anvil. Previous case studies [Churchill and Houze, 1984] have shown that the majority of subsequent convection in organised DCCs tends to occur around the edge of the storm, so it is expected that the masking of developing cores by existing anvil clouds will only have a small impact on the detected number of cores. In future, research into methods for detecting convective cores under anvil clouds [Apke et al., 2018, e.g.] could avoid this bias along with enabling tracking of the full extent of convective core lifetimes.

2.3.2 Observing anvil clouds

In the second experiment, the appearance of anvil clouds of different optical depth (OD) in ABI observations is investigated. The simulations were run with an ice cloud with CTH at 14 km, cloud base height at 12 km and cloud top r_e of 20 µm, aiming to represent the typical properties of a dissipating anvil cloud [Sokol and Hartmann, 2020]. This simulation was then run for a range of OD between 0.125 and 10, with intervals of 0.125 between 0.125 and 1, 0.25 between 1 and 2, 0.5 between 2 and 5, and 1 between 5 and 10.

BT spectra for these simulations are shown in Fig. 2.8, along with the ABI spectral response functions as in Fig. 2.5. As before, absorption due to CO_2, ozone, and WV is

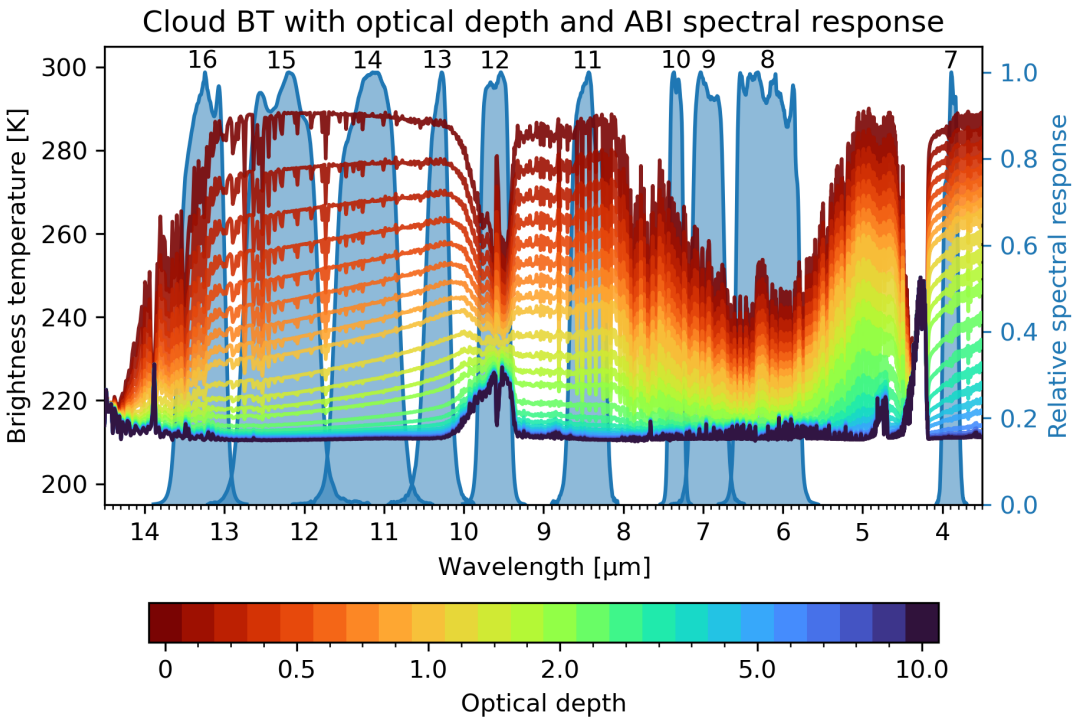


Figure 2.8: ToA BT spectra for a simulated anvil cloud with optical thickness between 0 and 10.

seen throughout the spectra. Unlike the simulations in section 2.3.1 however, there is a difference in the BT spectra across the LW window (10–13 µm) for clouds with OD between 0 and 2, with colder BT at longer wavelengths. This is due to the difference in ice emissivity across this range of wavelengths [Fu, 2015]. As the emissivity of ice particles reduces for wavelengths at 10 µm and below there is a greater contribution from the atmosphere below the cloud layer and, as a result, warmer BTs are observed. It should be noted that this difference is dependent on the size of the ice particles, with smaller r_e resulting in a larger difference in emissivity [Dubuisson et al., 2008].

Figure 2.9 shows simulated ABI observations for the 6.2, 7.3, 10.4, and 12.4 µm BT channels, the WVD and the SWD for a range of OD between 0.1 and 10. At high OD, all channels show a cold BT close to that of the CTT, however, for OD below 3, the increased contribution from the atmosphere below results in warmer

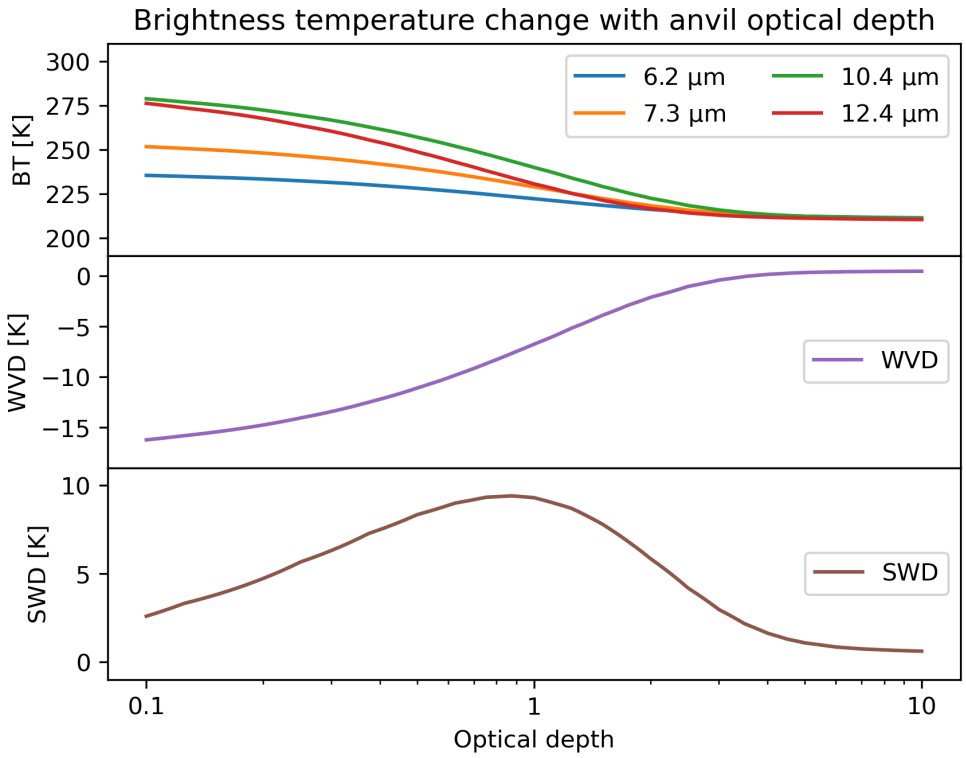


Figure 2.9: Simulated ABI BT observations for a DCC anvil of decreasing optical depth for the 6.2, 7.3, 10.4, and 12.4 µm BT channels, the WVD and the SWD.

BT. The WVD similarly shows a more negative BT difference at OD below 3, with a continual decrease for smaller OD. The SWD however shows an increasingly positive BT difference for OD below 5, peaking at 10 K around 1 OD. Below this OD the SWD decreases towards smaller positive values. As a result, the SWD has the potential to detect thin anvil cirrus at much lower OD than BT channels or WVD, down to values around 0.25 optical depth (OD).

While the BT, WVD and SWD are plotted in fig. 2.9 as a function of OD only, other factors may influence the observations. Changes in cloud height will affect the observed BT, as higher anvils will have lower temperatures and vice versa. The WVD should be unaffected by cloud height as long as the cloud is above heights at which WV contributes to the vertical weighting functions (fig. 2.4), and above 12 km there is little change with height (fig. 2.7). Changes in tropospheric humidity will affect

the WVD as this will change the height of the weighting function response and hence the height at which the WVD is sensitive to clouds. As the majority of DCCs are expected to occur in relatively humid environments it is expected that this effect will be minimal. However, in particularly cold, dry conditions this may result in lower level clouds being erroneously detected as anvils. Finally, the SWD has some sensitivity to r_e, as larger particles show a smaller difference in emissivity between 10 and 12 µm. The effect of this is small, and the peak SWD remains present at the same optical depth.

2.4 Method

In this section, a novel method for detecting and tracking both the growing cores and anvils clouds of DCCs, consisting of the following steps:

1. Ingest of LW IR BT fields from geostationary satellite imagery, including calculation of WVD and SWD fields from IR WV and LW IR window channels.

2. Calculation of optical flow vectors to be used as an *a priori* estimate of cloud motion for use in the semi-lagrangian framework.

3. Detection of growing DCC cores using cloud top cooling rate.

4. Detection of thick and thin anvil clouds associated with detected cores using a semi-lagrangian "3D" watershedding method.

5. Grouping of cores into multi-core systems, calculation of statistics and validation using lightning observations.

2.4.1 Estimation of cloud motion vectors using optical flow

The retrieval of AMVs has been performed since the earliest geostationary satellite observations [Menzel, 2001]. AMVs provide information about the motion of clouds

in the atmosphere, including DCCs [Bedka and Mecikalski, 2005], and are routinely generated for the majority of operational geostationary earth observation satellites, including GOES-16 [Daniels et al., 2016]. However, although AMVs may provide useful information about the motion of DCCs, the non-geostrophic nature of wind fields in these conditions may result in the AMVs being calculated inaccurately or rejected by quality control checks [Bedka and Mecikalski, 2005].

Optical flow algorithms are a family of algorithms used to estimate the apparent motion of objects observed in a series of images [Aggarwal and Nandhakumar, 1988]. A wide range of optical flow algorithms exist, and these have been successfully applied to many computer vision applications. It should be noted that optical flow does not necessarily represent the physical motion of an object, but is instead an estimation of the relative motion between an object and the observer and additionally any change in the apparent object (including growing, shrinking or other warping of the object).

Optical flow algorithms have been previously shown to be accurate for the prediction of AMVs using geostationary satellite images [Wu et al., 2016], as long as the observations are sufficiently frequent such that the motion of unique features between images is less than the length scale at which neighbouring features can be resolved [Bresky and Daniels, 2006]. Heikenfeld et al. [2019] found that at imaging frequencies of less than 5 minutes the motion of DCC cores was less than the spacing between neighbouring cores in the majority of cases, indicating that the frequency of the ABI CONUS scan region is suitable for calculating optical flow vectors of DCCs. The use of optical flow has several advantages over traditional AMVs for the retrieval of DCC motion vectors: optical flow can be calculated quickly using only two subsequent images and no *a priori* information, aiding in near real-time applications; and also have no requirement for geostrophic balance. Optical flow algorithms are routinely used in the nowcasting of convective precipitation, and can be used to provide accurate predictions of DCC with an hour of lead time using either radar or satellite observations

[e.g. Bowler et al., 2004; Bechini and Chandrasekar, 2017; Woo and Wong, 2017]. Optical flow, and similar motion vector techniques, have also been successfully applied to both the detection of developing deep convection [Zinner et al., 2008; Zhang et al., 2014] and tracking detected deep convective features [Senf et al., 2018a] separately.

It should be noted that optical flow is used to estimate the apparent motion of the cloud field between subsequent images, with the aim of using these vectors to map the locations of DCCs from one step to the next, instead of calculating actual AMVs corresponding to winds. This approach avoids a number of challenges with the use of optical flow for calculating AMVs including the estimation of the height of estimated flow vectors and the detection of diverging or converging motion vector fields in situations of growing and dissipating clouds respectively. In the latter case, both divergence and convergence within the optical flow vector field are intentionally included to map both the location and shape of observed clouds between time steps.

2.4.1.1 Evaluation of Optical Flow Methods

As different optical flow algorithms may provide better accuracy in different situations [Baker et al., 2011], several algorithms are compared to assess their suitability for tracking the motions of DCCs. Six different optical flow algorithms, implemented in the OpenCV library [Bradski, 2000], are considered, which are listed below:

- Farnebäck [Farnebäck, 2003]
- Deep Inverse Search (DIS) [Kroeger et al., 2016]
- DeepFlow [Weinzaepfel et al., 2013]
- Dual Total Variation Regularisation & Robust L^1 Norm (Dual TV-L^1) [Zach et al., 2007; Pérez et al., 2013]
- Sparse to Dense [Bouguet, 1999]

- principle component analysis (PCA)-Flow [Wulff and Black, 2015]

The first four of these algorithms are dense optical flow methods, which calculate motion vectors for every pixel in the image. The final two calculate sparse optical flow—motion vectors between detected objects in each image, rather than individual pixels—and then map these to all pixels based on a region segmentation approach. The sparse to dense and PCA-flow algorithms use the Lucas-Kanade method [Lucas and Kanade, 1981] and the PCA method respectively to calculate the sparse motion vectors.

Each algorithm was applied to a sequence of images from the 10.4 µm channel, with the algorithms estimating motion vectors between each consecutive pair of images. As the methods require 8-bit images, a linear normalisation was applied to the BT images between each pair's minimum and maximum values. To improve the accuracy of the detected motion vectors, a variational refinement process [Brox et al., 2004] is applied to the calculated flow vectors. This process minimises the energy functional of the pixel value, pixel gradient, and local smoothness to reduce the angular error of the detected flow vectors. Finally, vectors are calculated both forwards and backward in time between each pair of images, and these two sets of vectors are interpolated to each other's locations and averaged to ensure that they are equal and opposite.

Figure 2.10 shows the ABI BT at 19:00:00 UTC (fig. 2.10 a), along with the optical flow vectors calculated by each of the algorithms with the subsequent time step. For each set of optical flow vectors, the average vector over each 5×5-pixel area is plotted by the red arrows, and the vector velocity magnitude is plotted in the background. The two sparse methods (fig. 2.10 f,g) show a closer match of the areas where optical flow vectors to the edge of the anvil cloud than the dense methods (fig. 2.10 b–e), which have a tendency to smooth the detected vectors over a large area. Of all the methods, the Farnebäck algorithm tends to produce larger magnitudes, but correctly detects zero flow in areas without clouds, unlike most of the other algorithms.

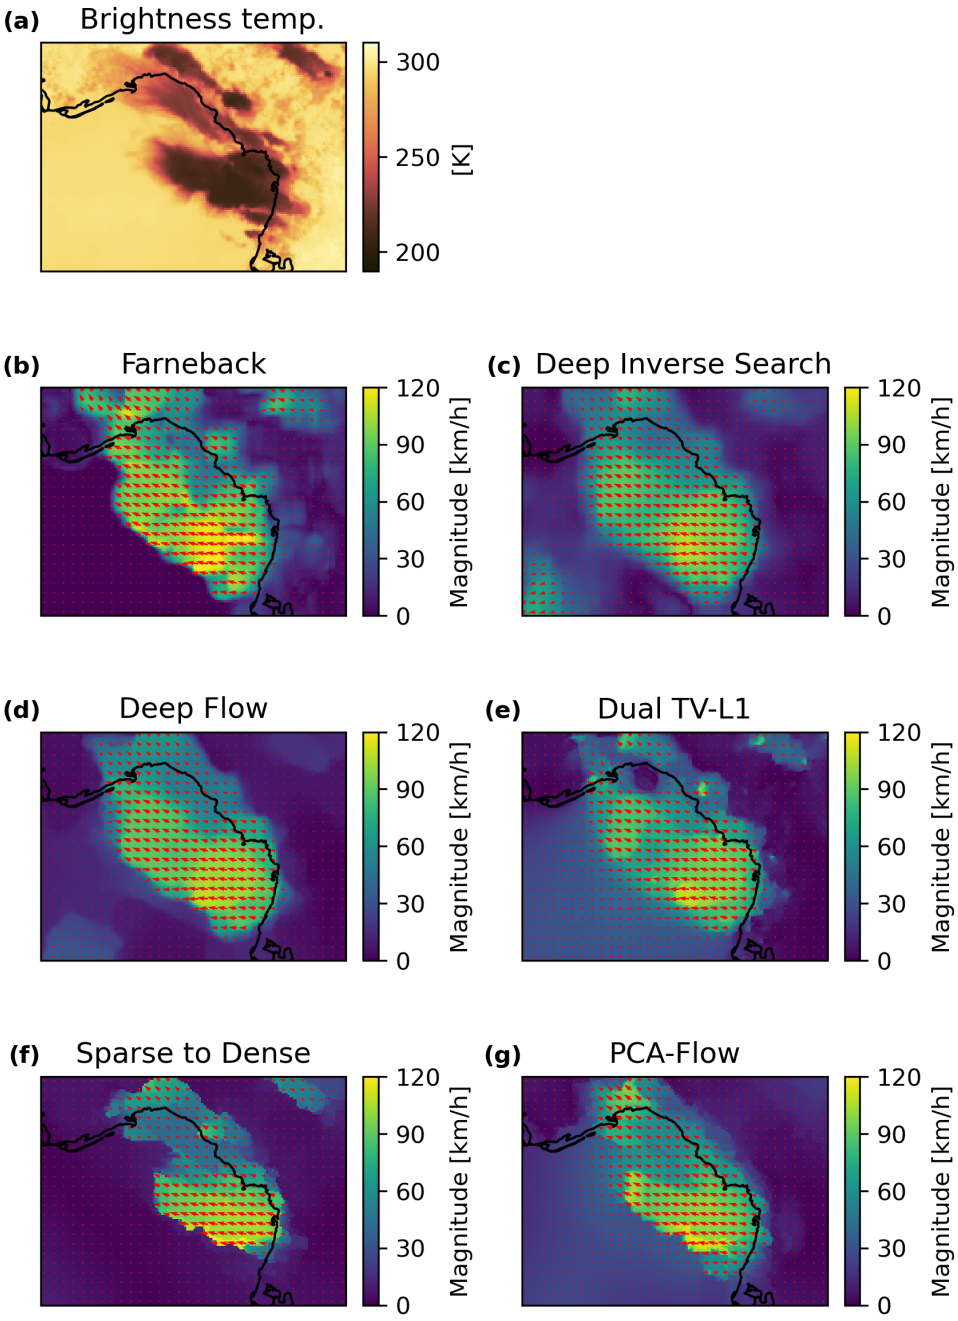


Figure 2.10: A comparison of the motion vector fields generated by the (b) Farnebäck, (c) DIS, (d) DeepFlow, (e) Dual TV-L[1], (f) Sparse to Dense, and (g) PCA-Flow. Optical flow vectors shown are estimated between the 10.4 µm channel at 19:00:00 UTC (a) and the subsequent image. Red arrows (b–g) show the motion vectors averaged over each 5×5-pixel area is plotted by the red arrows, and the background shows the magnitude of the velocity vectors.

Figure 2.11 shows the difference in BT between subsequent observations when taking into account the cloud motion predicted by the different optical flow algorithms. The plots are again shown for the time step at 19:00:00 UTC, with the 10.4 µm BT at this time step shown again for reference in fig. 2.11 a. Figure 2.11 b shows the Eulerian difference in BT; the difference in the BT observed at the same pixel locations on subsequent time steps. Large cooling rates are seen of the left (downwind) side of the anvil cloud, not because the CTT is cooling, but because of the expansion of the anvil cloud into previously cloud-free areas. Separating these observations of cooling BT caused by the motion and expansion of the anvil cloud from those due to the vertical growth of the DCC is the main objective of the semi-Lagrangian approach.

The dense optical flow methods (fig. 2.11 c–f) show good performance in removing the anomalous cooling observations seen downwind of the anvil cloud. The Farnebäck, DIS and DeepFlow algorithms (fig. 2.11 c–e) all show a warming tendency towards the edge of the anvil cloud. Although this may be explained by an overestimation of the cloud motion by these algorithms, this may also be a real observation of apparent warming due to thinning of the detraining anvil cloud. The Dual TV-L^1 algorithm (fig. 2.11 f) shows generally smaller differences than the other algorithms. This is likely due to the algorithm accounting for changes in brightness between observations, whereas the other algorithms assume that the brightness of tracked objects remains constant between observations.

The two sparse optical flow algorithms (fig. 2.11 g,h) both show worse performance, with the sparse to dense method performing particularly poorly. Both algorithms show cooling downwind of the anvil, likely because the region-based methods used to map the sparse motion vectors to the whole image do not detect this region well.

Figure 2.12 a shows the root mean square error (RMSE) of the BT differences shown in fig. 2.11 for each algorithm. As seen in fig. 2.11, the Dual TV-L^1 algorithm has the smallest differences, the other dense algorithms perform similarly, and the

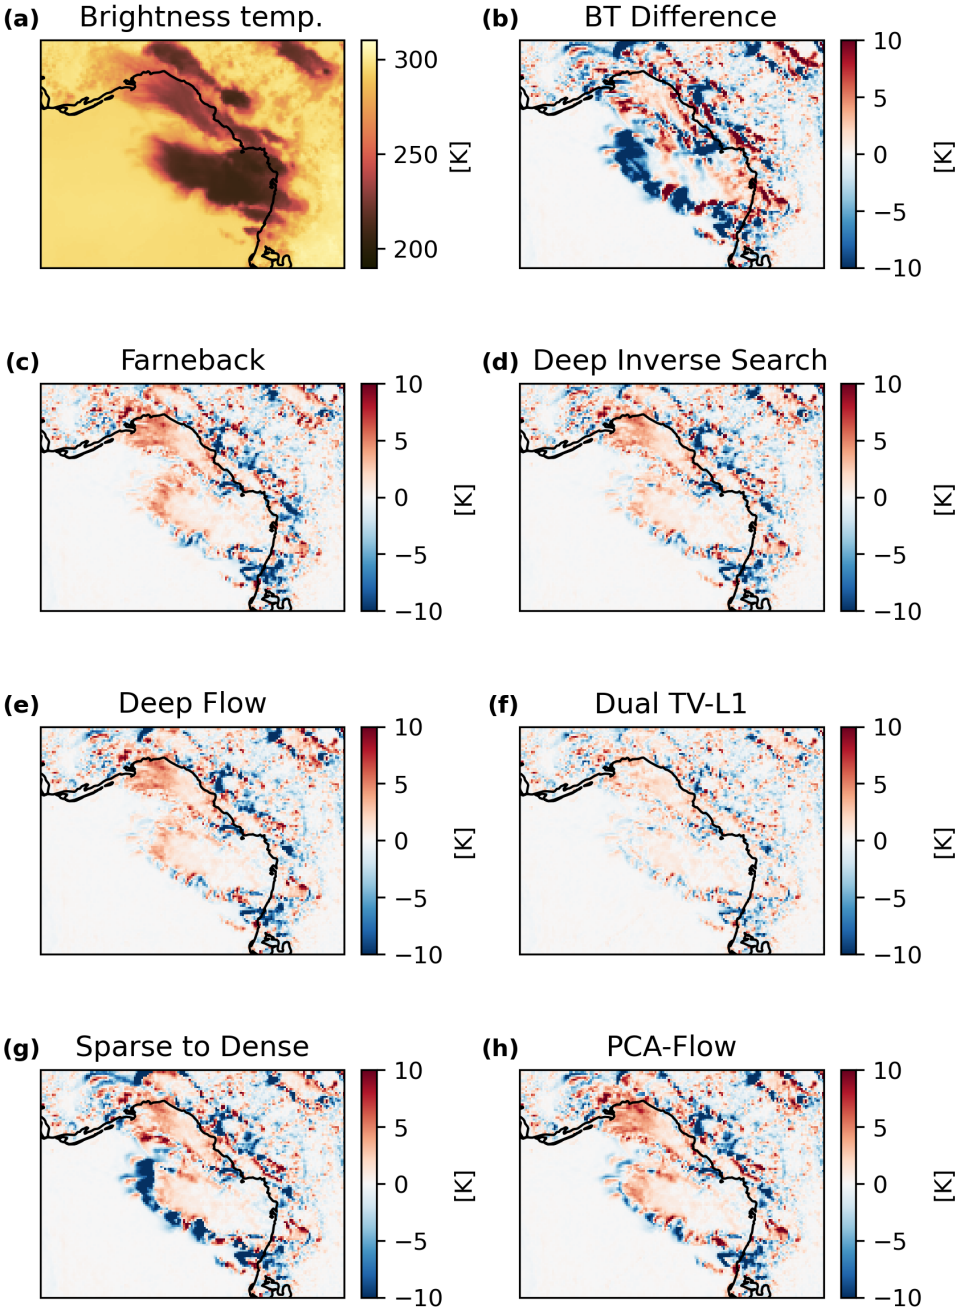


Figure 2.11: The difference in BT between subsequent ABI observations when taking into account the cloud motion predicted by the different optical flow algorithms. The plots are again shown for the time step at 19:00:00 UTC, with the 10.4 µm BT at this time step shown for reference (a). b: the difference in BT without accounting for cloud motion; c–h: the Lagrangian difference in BT using motion vectors predicted by each of the optical flow algorithms.

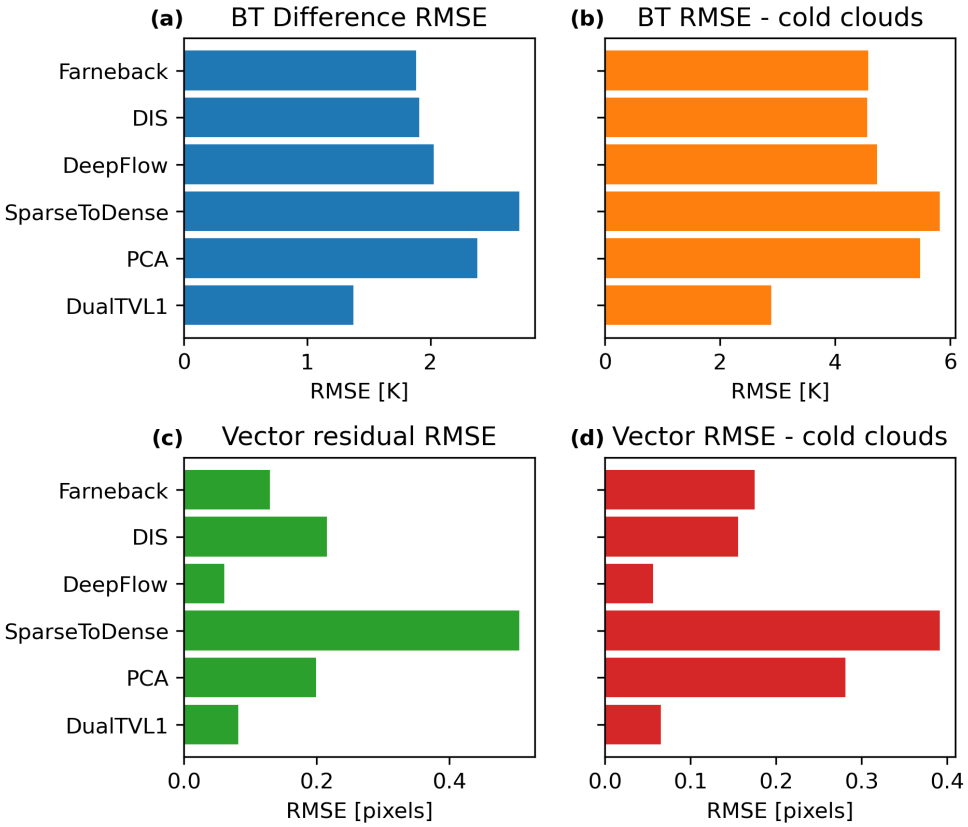


Figure 2.12: A comparison of the RMSE for each of the optical flow algorithms. a: BT difference RMSE for all pixels and (b) for pixels with BT $< 273\,\mathrm{K}$. c: RMSE of the residual vector magnitude for all pixels and (d) for pixels with BT $< 273\,\mathrm{K}$.

PCA-Flow and sparse to dense algorithms perform the worst. Some differences in the BT should be expected due to actual changes in the CTT. Figure 2.12 b shows the RMSE of the BT differences for pixels with BT less than 273 K, representing cold cloud observations. While the RMSE shown here for all algorithms is large compared to the sensor noise previously calculated for the BT difference in section 2.3.1 (0.47 K), it should be noted that some of the RMSE here is due to contributions from changes in the height of the anvil cloud. As a result, the actual error in the BT difference over time due to horizontal cloud motions not captured by the optical flow algorithm is likely much smaller, on a similar magnitude to the NEΔT.

To estimate the uncertainty in the calculated motion vectors, the residual motion

vectors are calculated for each algorithm. The residual motion is calculated by first warping each image to the locations predicted by the optical flow vectors, in the same manner used to calculate the BT differences in fig. 2.11. Then the optical flow algorithms are used to calculate a second set of residual vectors between each image and the subsequent warped image. In theory, if the initial motion vectors were calculated perfectly, warping the image by the optical flow vectors should result in the subsequent cloud in the same position as the prior image, and so the second set of motion vectors should calculate as zero. The magnitude of the residual motion vectors is therefore considered as the uncertainty in the initial set of motion vectors.

Figure 2.12 c,d show the RMSE of the residual vector magnitudes for all pixels and pixels less than 273 K respectively. The DeepFlow and Dual TV-L^1 algorithms again show the best accuracy, and sparse to dense the worst. The Farnebäck algorithm shows good accuracy across all pixels, as seen in the ability to correctly detect zero motion in cloudless skies in fig. 2.10. However, when considering only cold pixels the Farnebäck algorithm performs slightly worse than DIS.

The results shown in fig. 2.12 suggest that the Dual TV-L^1 algorithm is the best choice for estimating the motion vectors in satellite observations of DCC. However, due to the large size and long sequences of images produced by ABI, it is also important to consider the computational cost of each algorithm. Figure 2.13 shows the run time in seconds for each algorithm applied to a sequence of 48 300×300 pixel ABI BT images. Both the DeepFlow and Dual TV-L^1 algorithms show significantly longer run times than any of the other algorithms. As a result, the Farnebäck algorithm has been selected for the tracking method as it offers the best balance of accuracy and computational performance for running on large datasets.

An example of the motion vectors calculated by the Farnebäck algorithm when applied to the 10.4 µm BT field is shown in fig. 2.14. Figure 2.14 c shows the residual flow vectors, which are small compared to the detected cloud motion vectors (fig. 2.14 b).

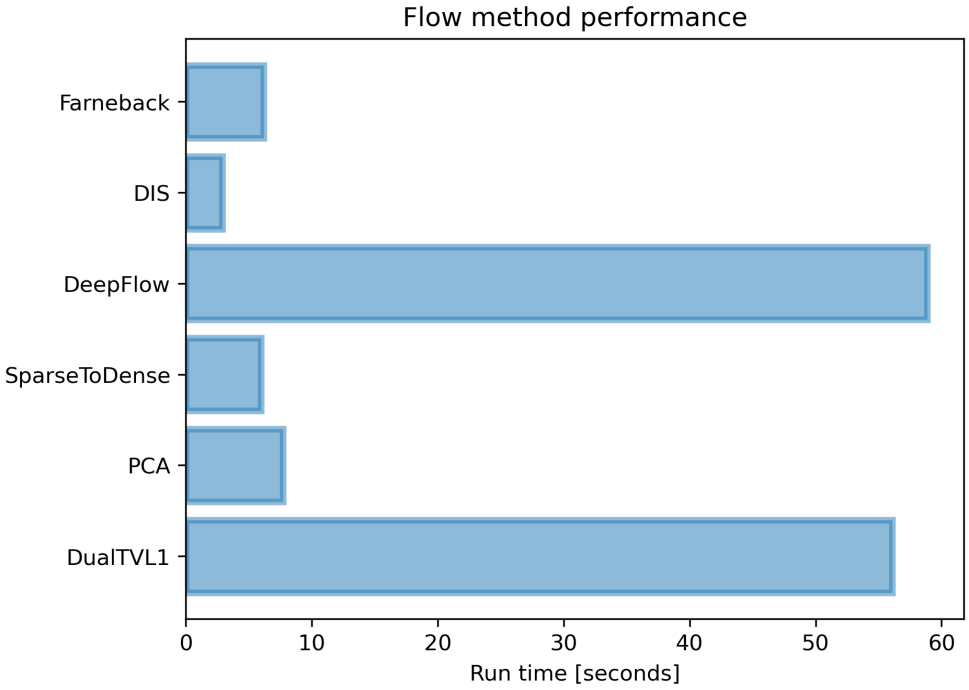


Figure 2.13: A comparison of the computational performance of each optical flow method. Run-time was measured for a sequence of 48 300×300 pixel ABI 10.4 µm BT images.

The main areas in which residual flow vectors can be seen are the regions around the edge of anvil clouds and in the centre of thick anvil clouds where there are no clear features to track. Figure 2.14 d shows the distribution of residual vector errors proportional to the detected cloud motion vectors. The majority of the residual errors are close to zero, and very few exceed 0.5. As, in the majority of applications for the semi-Lagrangian framework, only the motion rounded to the nearest pixel is of interest, these errors are considered acceptable.

Two potential sources of uncertainty are the assumptions made by the Farnebäck algorithm that the feature being tracked remains the same size and intensity in subsequent images. For optical flow tracking using BT images of growing DCCs, neither of these assumptions are true. However, steps have been taken to reduce the impact of both these sources of uncertainty on the tracking algorithm. For the first of these

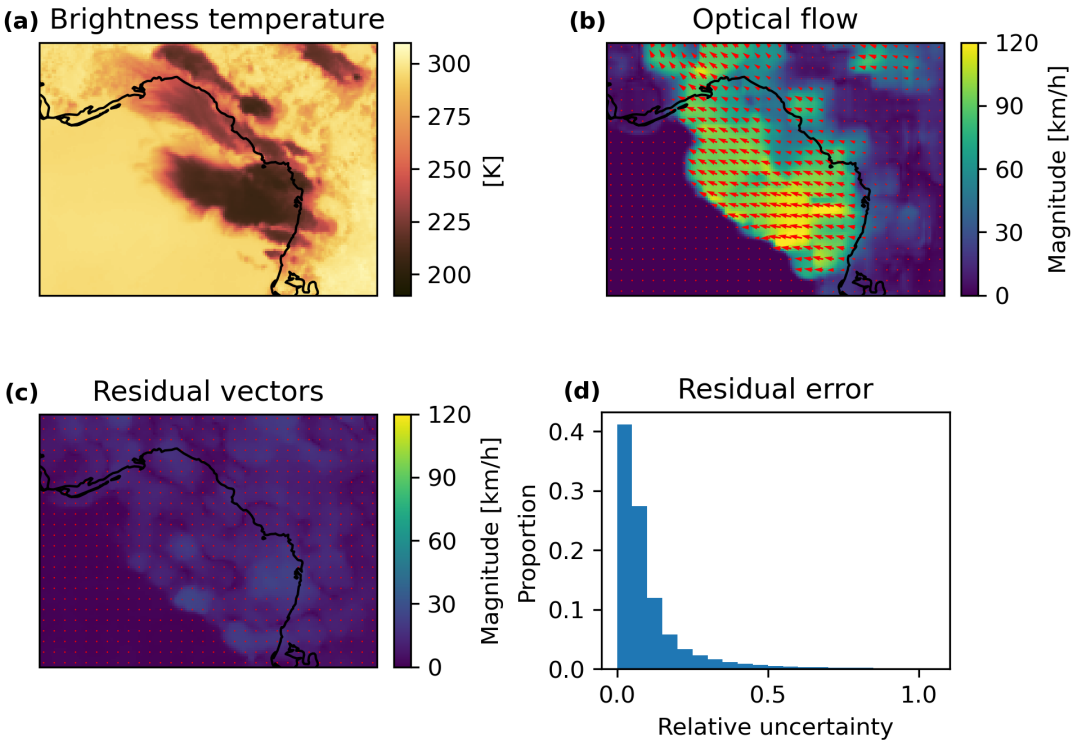


Figure 2.14: An example of the motion vectors and their residual errors calculated using the Farnebäck algorithm. a: the 10.4 µm BT field at 19:00:00 UTC for which the motion vectors are shown. b: the calculated motion vectors. c: the residual motion vectors, indicating the uncertainty in the optical flow. d: the magnitude of the residual motion vectors as a proportion of the magnitude of the optical flow vectors for cold pixels (less than 273 K)

assumptions, it is found that in the case of small, fast-moving DCCs—where the accuracy of the optical flow vectors is most important—the changes in the size of the DCC is small compared to the overall motion, and so the uncertainty introduced is small. Comparably, for large DCCs where the changes in size may be large compared to the motion, the uncertainty has less impact on the tracking algorithm. In the worst-case scenario, the estimated optical flow field will be zero, in which case the "3D" detection and tracking algorithm works in the same manner as that of Fiolleau and Roca [2013], which is suitable for use on these larger DCCs.

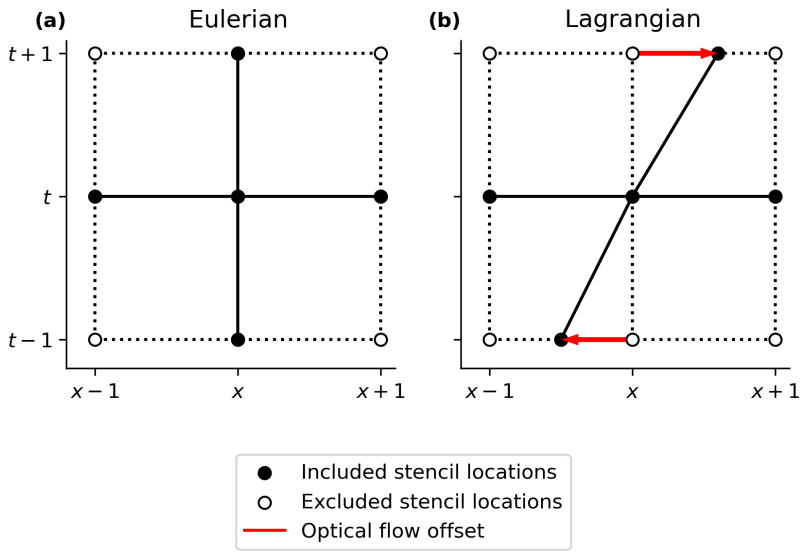


Figure 2.15: A comparison of convolution stencils with square connectivity in Eulerian (a) and Lagrangian (b) frameworks. In the Lagrangian framework, the points at prior and subsequent time steps are offset by the calculated optical flow field.

2.4.2 A Semi-Lagrangian Framework for Morphological Image Processing

Morphological image operations analyse images using their geometrical and structural properties. Core to many morphological algorithms, from simple filters to complex neural networks [Kalchbrenner et al., 2014], is the kernel, or convolution method. A convolution method performs operations on the pixels of an image by applying a convolution stencil to the pixel and its neighbours. In a conventional convolution scheme, such as that used in the methods of Fiolleau and Roca [2013], the convolution stencil acts on adjacent pixels in both time and space (see fig. 2.15 a). In this Eulerian framework, different locations in time are considered in the same manner as those in the spatial dimensions. However, previous analysis of DCCs has shown that the motion of convective cores between images can be similar to the spacing of cores and their size [Heikenfeld et al., 2019]. As a result, it is important to include the effects of advection when comparing images across time steps.

To perform morphological operations which take into account this advection, a

novel Lagrangian convolution method has been developed. For spatial operations, the Lagrangian stencil operates identically to that of a classical convolution method. However, when sampling points at prior or subsequent time steps, the locations of the stencil points are offset by the relevant optical flow vectors (fig. 2.15 b). Values at the offset stencil locations are interpolated, providing a Lagrangian reference frame for changes in the observations over time. Applying the convolution stencil to every pixel in a sequence of images provides a semi-Lagrangian framework for morphological operations, combining the Lagrangian reference frame for evaluating changes over time while maintaining the regular grid of the images.

New implementations of several existing image processing operations have also been developed within the Lagrangian convolution framework, including:

- Sobel edge detection [Sobel, 2014]
- Watershed segmentation using the connected-components method [Bieniek and Moga, 2000]
- Labelling of connected components [Hoshen and Kopelman, 1976]

These operations are used in this method to detect the full extent of the anvil cloud associated with the DCC, to perform detection continuously across multiple time periods while accounting for the motion of the DCC, and to identify individual DCCs and DCC clusters across multiple time periods respectively.

The Sobel method detects edges in an image using the magnitude of the local gradient at each pixel. Edge detection enables the segmentation of an image into separate regions without pre-defined thresholds (such as in BT) to separate them.

Watershed algorithms are a method of image segmentation that equates an image to a topographical map, with elevation according to the value of the pixel. Each pixel is then descended towards its local minima until it reaches a predefined marker region. The method takes its name from the geographical feature of the same name, which

refers to the separation between adjacent drainage basins. Although this physical interpretation of the algorithm applies to two-dimensional images, the method can be applied to arrays with any number of dimensions, such as the method used by Fiolleau and Roca [2013] which applied watershedding to a three-dimensional field.

Labelling algorithms assign unique identifiers to each segmented region provided by either the edge detection or watershed algorithms.

2.4.3 Detection of Growing Deep Convection

Growing deep convective cores are detected in a similar manner to that used by Zinner et al. [2008]. The growth rates are calculated using the finite difference of the 10.4 µm BT and WVD fields in the Lagrangian perspective. Combining both the 10.4 µm BT and the WVD field provides the best observations for detecting growing deep convective cores. The BT field shows cooling cloud tops throughout the troposphere, while the WVD isolates growth in the mid- to upper-troposphere, removing spurious observations of growth due to boundary layer convection and cloud formation.

A region of growing deep convection is classified as an area showing continuous cooling of greater than 0.5 K per minute in the 10.4 µm BT field. This growth rate must be detected over a minimum 15-minute period, covering an area of at least 3 by 3 pixels (approximately 7.5 by 7.5 km) at each time step. Each region of detected growth is given a unique label, the average BT of the detected cloud at each time step is calculated, and any label which fails to meet the cooling rate thresholds of 8 K over a 15-minute period (from [Roberts and Rutledge, 2003; Hartung et al., 2013]) are removed. Finally, each label is checked to ensure that the growth region ends with a WVD field with a value of greater than -5, indicating that the growing core reaches a high altitude [Müller et al., 2018]. While threshold detection methods are sensitive to noise, the uncertainty in the BT cooling rates due to sensor noise, combined with the uncertainty in the optical flow vectors, is not expected to exceed 5% of the threshold.

A simple threshold method applied to the BT cooling rate is, therefore, suitable for the detection of developing deep convection.

Figure 2.16 shows a comparison between the detected core cooling rates in ABI imagery and the corresponding column radar reflectivity measured by NEXRAD for the case of newly developing convection (fig. 2.16 a,c,e) and mature DCC (fig. 2.16 b,d,f) Both 10.4 µm BT and WVD growth rates show growth detected in the same locations as the high column radar reflectivity from NEXRAD. As predicted by the results in section 2.3.1, the observed growth rates in the WVD are about half of those in the 10.4 µm BT.

Both 10.4 µm and WVD growth rates in the growing phase (fig. 2.16 c,e) show good agreement with column radar reflectivity (fig. 2.16 a). However, during the mature stage of the DCC, discrepancies develop between the observed growth rates (fig.2.16 d,f) and the radar reflectivity (fig. 2.16 b) due to the development of the anvil cloud blocking satellite observations of the core underneath.

2.4.4 Detection of Anvil Clouds

The region of anvil cloud associated with the growing convective clouds detected in section 2.4.3 is detected and tracked using an edge-based watershed segmentation approach. The edge-based approach to cloud detection avoids the use of a fixed threshold for anvil temperature, and so can detect a more accurate extent of the anvil cloud [Dim and Takamura, 2013]. An upper threshold for the WVD field of -5 K is defined, as used by Müller et al. [2018], and a lower threshold of -12.5 K, which represents definite non-anvil cloud. Because the presence of thin cirrus outflow from the anvil clouds can make it difficult to determine the extent of the anvil cloud, the SWD field is used as described in section 2.1 to either remove or include the region of cirrus outflow in the detected anvil region. To detect the thick anvil cloud, the SWD field is subtracted from the WVD field (fig. 2.17 a). In this case, the

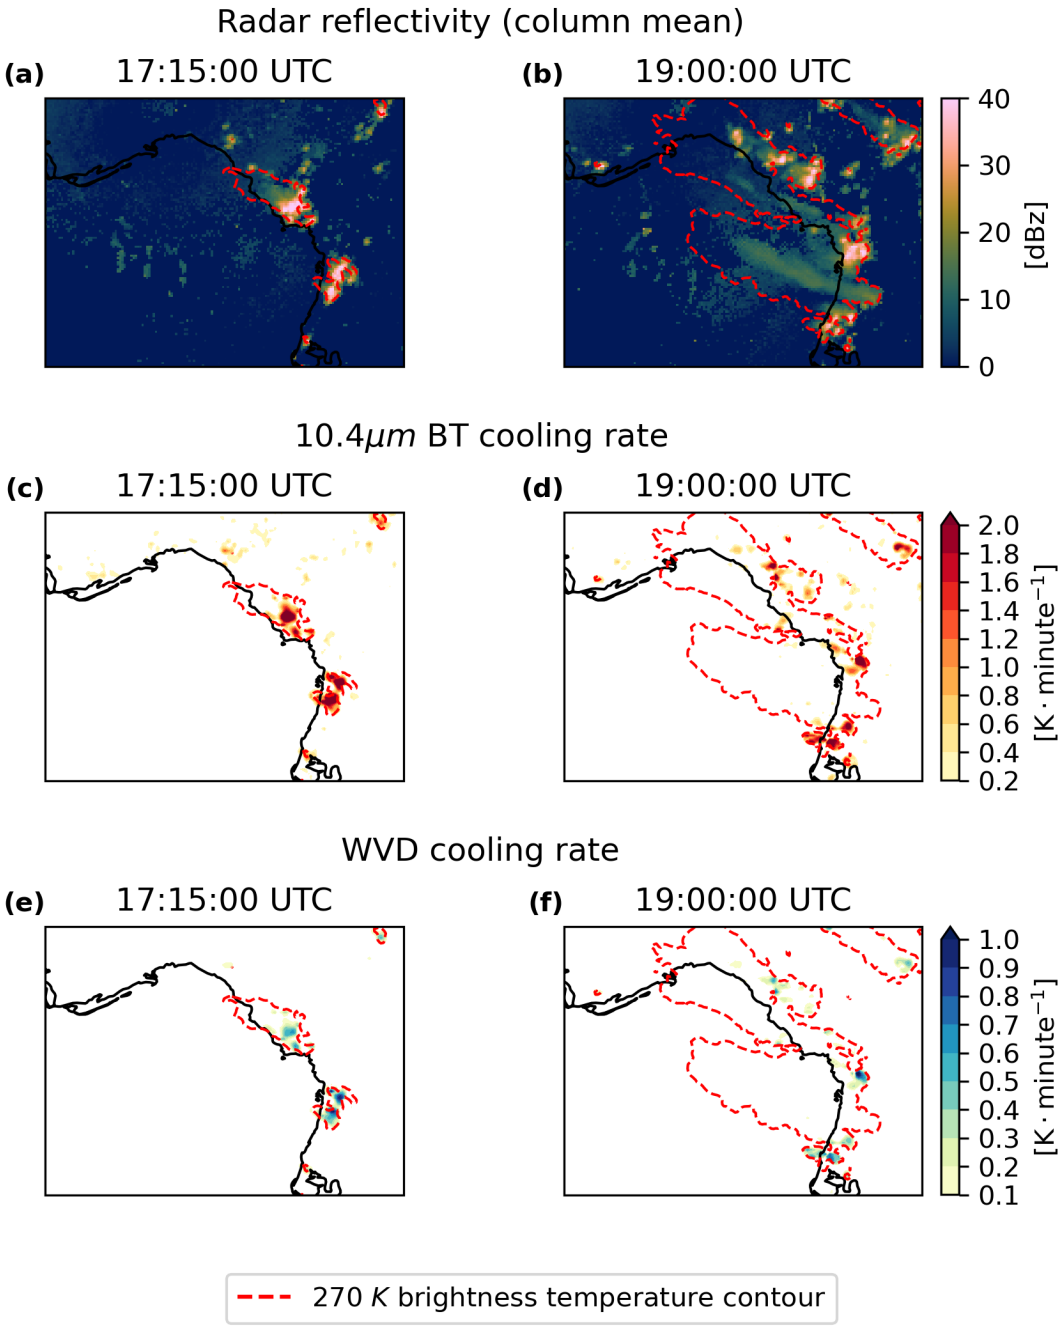


Figure 2.16: Detection of growing DCC regions for the DCC cluster in figure 2.1. a,b: NEXRAD column mean radar reflectivity remapped to the ABI grid. c,d: cooling rate of the 10.4 µm BT field. e,f: warming rate of the WVD field.

upper and lower thresholds remain the same as the SWD field is approximately 0 K for thick, high clouds, and so has no effect on the temperature of these features. For detecting the thin anvil region, the SWD field is added to the water vapour difference (WVD), and the value of both thresholds is increased by 5 K to 0 K and -7.5 K respectively (fig. 2.17 c). This change is made to account for the effect of low-level WV on the SWD field which gives a background value of approximately 5 K. Between these two thresholds there is a region in which the extent of the anvil cloud is uncertain. By applying a Sobel filter to detect the local gradient magnitude of the combined WVD/SWD field [Sobel, 2014], the outer extent of the anvil cloud is detected within this region where the greatest magnitude in the detected edges is found (see fig. 2.17 b,d).

The uncertainty due to sensor noise is small relative to the range between the two thresholds, at 2.7 % for the combined WVD and SWD fields. While this low noise may indicate that a simple threshold approach may be suitable, such an approach would be susceptible to the systematic biases in the WVD and SWD observations identified in section 2.2.1.1. Combining the edge detection approach with the two-threshold approach provides robustness to systematic bias, since as long as the maximum gradient of the cloud edge remains between the two thresholds the same cloud extent should be detected. This allows systematic biases on the order of several K without significantly affecting the results, allowing better comparison of DCCs observed in different environments.

When applied to the detected edges of the anvil clouds using the Sobel filter, with the growth regions detected previously as markers, the watershed method allows us to detect those anvil regions associated with detected regions of growing DCCs, while avoiding the detection of non-convective regions of high, cold clouds. Furthermore, due to the application of the watershed algorithm to both the spatial and temporal dimensions of the sequence of images through the semi-Lagrangian framework, the

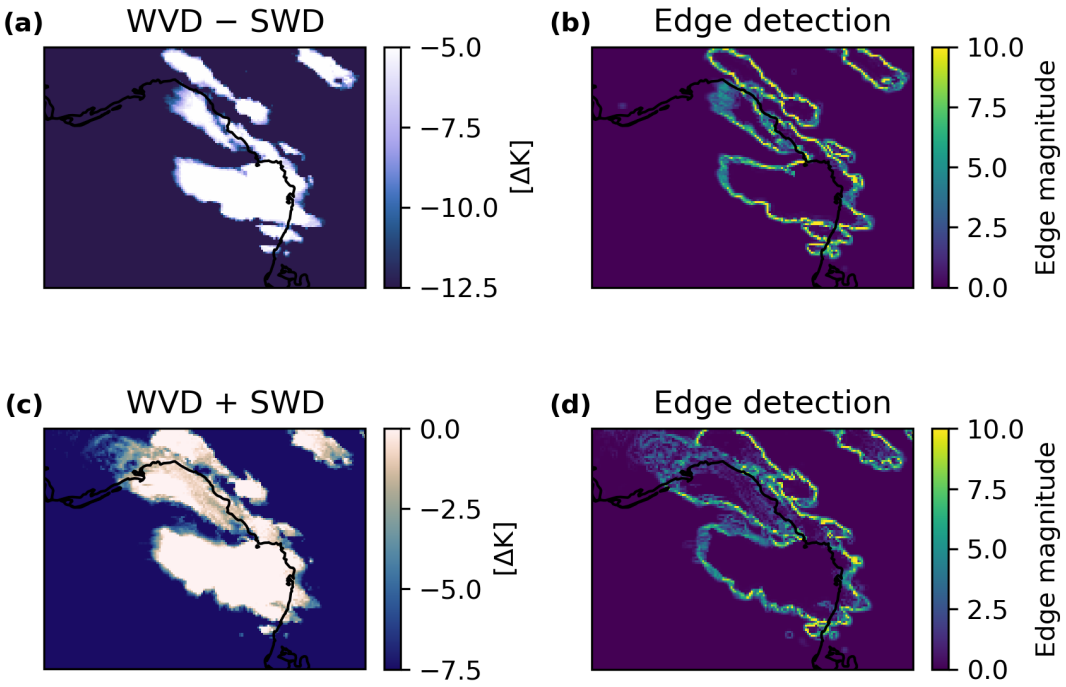


Figure 2.17: Detection of anvil cloud extent for the mature DCC cluster in 2.1 using the edge gradient method. a.: The combined field of the WVD minus the SWD, to isolate the thick anvil, between the upper and lower thresholds of -5 and -12.5 K respectively. b.: the detected edge gradient magnitude of the field between these thresholds, which is used to detect the outer extent of the thick anvil cloud. c.: the combined field of WVD plus the SWD, to enhance the thin anvil, and d.: the calculated edge magnitude of this field.

associated anvil clouds can be detected and tracked after the growth of the DCC is no longer observed.

Figure 2.18 shows an example of the results of detecting and tracking DCC cores (outlined in red) and the associated anvils (outlined in orange and blue for the thick and thin anvil regions respectively). During the growing phase (fig. 2.18 a,b) a number of developing cores are seen, as well as the initial growth of primarily thick anvil cloud. Figure 2.18 c–e shows the mature phase of the DCC, where the thick anvil expands and the thin anvil cirrus starts to detrain towards the upper-left side of the figure. Due to the anvil cloud reaching its maximum height, the initial developing cores seen in fig. 2.18 a,b are no longer detected as the cloud-top cooling cannot be

observed. Instead, new cores are detected developing on the edges of the anvil cloud. In fig. 2.18 f–h, the detected anvil cloud begins to dissipate. The thick anvil area decreases while the thin anvil area continues to increase as the anvil dissipates and detrains. At this point in the lifetime of the tracked DCC new growing cores are no longer observed, however the "3D" approach allows the continued tracking of the anvil cloud until it dissipates.

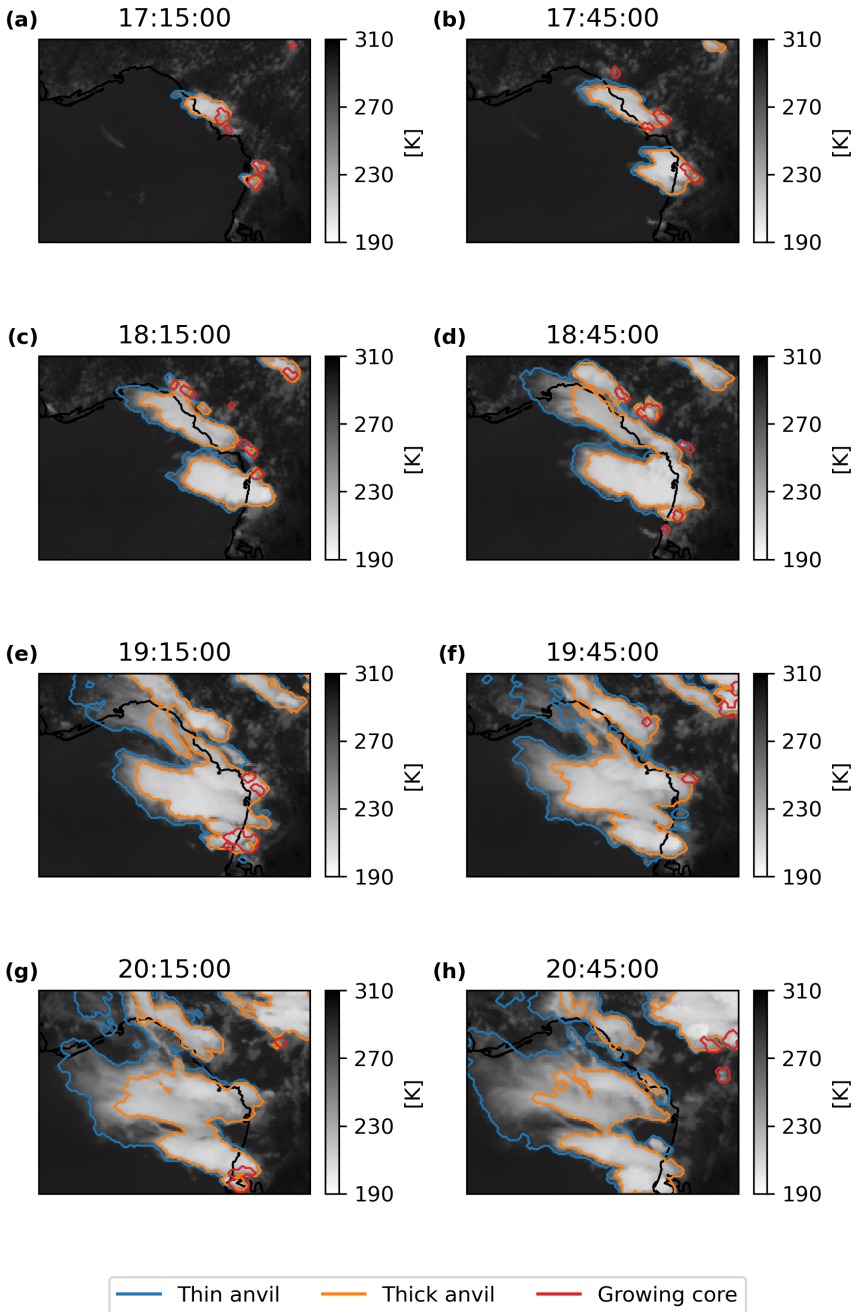


Figure 2.18: Detected regions of thin anvil cloud (blue), thick anvil cloud (orange), and developing cores (red) overlaid on the GOES-16 ABI 10.4 µm BT field for the DCC cluster from figure 2.1. The three stages of the DCC lifecycle are shown; the growth phase (a,b), the mature phase (c–e), and the dissipating phase (f–g).

2.5 Evaluation

The effectiveness of the semi-Lagrangian framework for the detection of DCCs is evaluated by analysing the proximity of detected anvil cloud regions to lightning flash detection from GLM. Lightning observations are frequently used to validate detection methods for deep convection [e.g. Zinner et al., 2013; Müller et al., 2019] due to the strong correlation between deep convective updraughts and lightning activity. Although GLM is not capable of detecting all lightning events (approximately 70% of lightning events are detected) [Peterson, 2020], the high frequency of lightning flashes per DCC mean that these observations provide a suitable ground truth for validation. It should be noted that lightning observations are only suitable for validating the detection of the thick anvil region, as lightning does not occur in the cirrus outflow. As a result, validation of the detection of the thin anvil region would require the use of other data such as cloud profiling radar or lidar observations, and is not considered further in this paper.

Here, the same validation method as that used by Müller et al. [2019] to evaluate the semi-Lagrangian framework for the detection of DCCs. Detection events are classified into three categories:

- Correct detection (CD), when the algorithm detects a DCC that is collocated with one or more lightning observations
- False detection (FD), when the algorithm detects a DCC but no lightning flash is observed
- Missed detection (ND), when the algorithm does not detect a DCC but a lightning flash is observed

Using these three categories of events two measures of accuracy can be defined for the detection of DCCs. The probability of detection (POD) is defined as the number

of correct detections divided by the total number of correct and missed detections. This provides a measure of how likely the algorithm is to detect a DCC that exists in the ground truth. The false alarm rate (FAR) is defined as the number of false detections divided by the total number of correct and false detections. This provides a measure of how likely a DCC detected by the algorithm is not present in the ground truth. The F_1-score is calculated as the harmonic mean of the POD and the recall (where the recall in is $1 - \text{FAR}$), and provides an overall measure of accuracy between 0 and 1.

When evaluating whether detected DCC regions and lightning observations were collocated, Müller et al. [2019] considered events within 32 km and 15 minutes to be collocated. This margin of uncertainty was separated into two components, half from the physical separation between observed lightning strikes, and the remaining half from uncertainty in the collocation and geolocation of the satellite and lightning observations. For a typical ABI pixel length over the CONUS of 2-2.5 km, this margin of error translates into 15 pixels in the ABI view. The distance between a GLM lightning flash and detected cloud region is defined as the distance between the flash and the nearest ABI pixel within that region, with GLM flashes that fall within a detected DCC given a distance of 0. When considering that the resolution of GLM is a factor of four less than that of ABI, the same justification for the margin of error used by Müller et al. [2019] is also applicable to collocated observations from ABI and GLM.

Validation was performed using GOES-16 ABI data from the CONUS scan region for the entirety of 2018, which was processed using the method described in this chapter. In total validation was performed for 319 days of ABI data, the remaining 46 days being excluded due to missing observations from either the ABI (1 day) or GLM (45 days) instruments aboard GOES-16. Detection and tracking of DCCs was performed over the CONUS scan region of ABI for daily periods. By performing

Detection Method	n	POD	FAR	F_1-score
Growth-based	598,038	0.4017	0.2136	0.5318
WVD threshold	678,854	0.9727	0.6457	0.5194
Semi-Lagrangian	145,969	0.9837	0.1611	0.9056

Table 2.4: POD, FAR and F_1-score for three different detection methods validated against observed GLM flashes (n=116,671,289). Growth-based refers to the detection of growing DCCs using the method described in section 2.4.3 WVD threshold uses the threshold method developed by Müller et al. [2018]. Semi-Lagrangian refers to the detection of anvil clouds connected to growing cores using the edge-based watershedding method described in section 2.4.4.

validation over both a large region, including a range of both land and ocean domains, and a full year time period, any bias in the validation associated with the variability of the accuracy of the method with location and season is avoided.

Results of the validation of the detected anvil region, as well as those for the detection of growing deep convection and the WVD filter, are shown in table 2.4. The regions of growing DCCs detected using the method described in section 2.4.3 shows low scores for both the FAR and POD metrics. While the detection of growing DCCs shows a low FAR of 0.21, the short time frame in which growth can be observed leads to a high rate of missed detections of lightning flashes, which results in a POD of 0.40. This is not necessarily because the algorithm fails to detect cores, but because many lighting observations occur during the mature phase of convection (see fig. 2.1 h, 2.2 d), and these lightning flashes are not detected as the core is only observed during the growing phase.

For comparison, the accuracy of detecting anvils only by a fixed threshold of the WVD without detecting growing cores is also evaluated, as used by Müller et al. [2018]. Compared to the detection of growing DCC regions, the WVD threshold shows a much higher POD of 0.97, but also has a high FAR of 0.64, repeating the findings of Müller et al. [2019] which show that although the WVD threshold method is capable of detecting the majority of DCCs, it is incapable of distinguishing between

anvil clouds and other thick, high altitude clouds. Furthermore, the WVD threshold detection detects a larger number of objects (n=678,854) compared to either of the other detection methods, further indicating that a large number of non-convective clouds are detected using the threshold method on its own. Note that both the core detection and WVD threshold approach have similar F_1-scores (0.53 and 0.51 respectively) as both prioritise one measure of accuracy over the other.

Finally, the anvil regions detected using a combination of the detected growth regions and the WVD field using the semi-Lagrangian framework described in section 2.4.4 are validated. The novel method has a high POD of 0.98 similar to that of the WVD threshold, while also maintaining much of the low FAR of the detection of growing DCCs (FAR=0.16). As a result, the semi-Lagrangian method displays a high overall F_1-score of 0.91. This result highlights the capability of the semi-Lagrangian detection framework to use growth-based detection methods to substantially reduce the compromise between POD and FAR error rates by combining multiple methods for the detection of DCCs.

2.6 Summary

Algorithms for the detection and tracking of DCCs perform a vital role in both forecasting and research applications. Sequences from geostationary satellites provide unique observations of DCC anvil clouds over their entire lifecycle. However, the traditional framework used by such algorithms requires a compromise between the rates of false and missed detections due to the overlap in signature from convective and non-convective clouds [Konduru et al., 2013]. Whereas novel methods have approached this problem for the detection of large, mesoscale convective systems [Fiolleau and Roca, 2013], such approaches do not take advantage of the capability of the latest generation of geostationary imaging satellites to detect individual deep convective cores.

By developing and implementing a novel semi-Lagrangian framework for the detection and tracking of DCCs, the detection of growing DCC cores [Zinner et al., 2008] and DCC anvils [Müller et al., 2018] to detect and track DCCs over their entire lifecycles can be combined. The novel methods developed here for the semi-Lagrangian computer vision framework, along with implementations of multiple image processing operations commonly used for object detection, allow the accurate detection and tracking of moving objects utilising both spatial and temporal information. These methods may have impacts on applications of computer vision beyond the detection and tracking of DCCs. Furthermore, the novel framework can achieve higher levels of accuracy without compromising on the number of DCCs detected, as with previous algorithms [Müller et al., 2019].

Using this novel methodology, both small, isolated DCCs and large, mesoscale convective systems can be detected and tracked with a high degree of accuracy, high spatial and temporal resolution and across large domains such as the CONUS. The data provided about the behaviour of DCCs over their entire lifetime will allow new research into vital topics such as the response of deep convection and climate change, and the interactions and feedbacks between DCCs and large-scale atmospheric thermodynamics [Varble, 2018].

Chapter 3

Linking the Properties of Deep Convective Cores and their Associated Anvil Clouds Observed over North America

3.1 Introduction

Understanding the relationships between the properties of deep convective cores and anvils is vital to understanding the behaviour of DCCs in both the present and future climate. Our ability to study these relationships is limited by a lack of datasets connecting convective processes to anvil properties over the entire DCC lifetime [Gasparini et al., 2023]. The capabilities of the newest generation of geostationary satellite instruments provide opportunities to address this issue. Data from the GOES-16 ABI instrument's CONUS domain provide many advantages for investigating convective processes. The higher spatial and temporal resolutions and number of channels allow us to move beyond the traditional tracking of large, cold cloud shields to instead track DCCs from the initial development of the core to the final dissipation of the anvil across scales spanning isolated DCCs to large MCSs. Furthermore, the CONUS domain contains a wide variety of regimes which impact the properties of DCCs, including land, ocean, tropics and mid-latitudes.

Deep convective storms play an important role in both the weather and climate

of North America. The continent experiences tropical, subtropical and extra-tropical convection across a range of modes, including isolated DCCs, MCSs and supercell convection [Brooks et al., 2019]. The North American Monsoon, which transports warm, moist air from the south-east Pacific and the Gulf of Mexico into the continent, strongly influences the seasonal cycle of these convective events [Adams and Comrie, 1997; Higgins and Shi, 2001]. DCCs are responsible for a wide range of extreme weather events, including heavy rainfall and flooding, hail, derechos, lightning and tornadoes [Westra et al., 2014; Houze, 2014; Williams et al., 1992; Bruning and MacGorman, 2013; Punge and Kunz, 2016; Matsudo and Salio, 2011]. Additionally, DCCs—in particular MCSs—provide the majority of precipitation across many regions of North America [Feng et al., 2019; Li et al., 2021]. As a result, a wide array of observational networks have been deployed to study deep convection over North America, including satellite observations and cloud radar [Brooks et al., 2019].

The United States of America (USA) East of the Rocky Mountains experiences a wide variety of convection. The Rocky Mountains block the Westerly zonal winds except at higher altitudes. At lower altitudes, instead, there is a Southerly low-level jet transporting warm, moist air from the Gulf of Mexico. The combination of these two air masses provides both a high shear environment and a high atmospheric lapse rate and hence high instability which provides the conditions for intense convection—including both MCSs and supercells—to initiate over the Great Plains and Midwest regions of the USA [Coniglio et al., 2010; Song et al., 2019]. These MCSs propagate eastward and provide the majority of precipitation across these regions [Feng et al., 2019]. In the southeastern USA the lapse rates, and hence instability, tend to be lower, but the water vapour mixing ratio is higher which produces a tendency towards more frequent but less intense convection [Brooks et al., 2007].

Mexico also experiences frequent MCSs, which produce heavy rainfall and risks from flooding [Douglas et al., 1993]. Topographic interactions play an important role

in the development of these systems. Warm, moist air from the Eastern Pacific and Gulf of Mexico is lifted and converges as it meets the mountainous terrain of Mexico, driving the development of MCSs [Farfán and Zehnder, 1994]. Unlike in the USA, there is typically less shear in this environment, and so there is a lower tendency for the production of supercell convection (and the associated risks) except in the North of Mexico [Weiss and Zeitler, 2008].

The impacts of extreme weather across North America drive interest in the research of convective systems and their behaviour. Furthermore, global warming is expected to drive an increase in several factors affecting convection, including CAPE [Seeley and Romps, 2015b], with a corresponding increase in the intensity of deep convective storms [Trapp et al., 2007; Seeley and Romps, 2015a]. However, many models generally do a poor job of representing convective storms, particularly MCSs, over much of North America [Pinto et al., 2015]. Although convective resolving models have improved this [Stevens et al., 2020], there are still shortcomings in their representation of convective cloud processes which limit their capabilities [Jeevanjee, 2017; Prein et al., 2021]. As a result, studying the distribution of convective systems and their properties remains an important task for understanding these processes, improving models and forecasting the weather [Brooks et al., 2019].

Geostationary satellite observations have provided a key tool for studying the behaviour of DCCs over North America due to the large spatial coverage of their observations. Early studies focused on the behaviour of mesoscale convective complexes; large, elliptical MCSs which propagate eastward from the Rocky Mountains [Maddox, 1980b; Augustine and Howard, 1988, 1991]. Tsakraklides and Evans [2003] found that linear organised convective systems (such as squall lines) showed different lifecycles to those of mesoscale convective complexes over North America, indicating that the structure of convection plays an important role in their behaviour. More recently, Feng et al. [2019] and Li et al. [2021] used a combination of satellite and radar data

to track both isolated and mesoscale convection, and provide better characterisation of their spatial properties. However, the low temporal resolution of 1 hour limits the ability to the properties of individual cores, and the reliance on ground-based radar restricts the study to DCCs which occur over land over the USA.

By leveraging the full temporal and spatial resolution of GOES ABI observations over the CONUS, a 5-year dataset of observed DCCs has been created, tracking both cores and anvils. Analysis of this dataset can show new discoveries about the lifecycle of DCCs, and how their properties correspond to those of their cores.

3.2 Data

To detect and track DCCs across North America, ABI MCMIP data from the 6.2, 7.3, 10.4 and 12.4 µm channels observed in the CONUS region is used, as described in section 2.2.1. These observations span the full extent of the ABI CONUS domain of 2,500 pixels E–W by 1,500 pixels N–S. This domain covers a region of around 60–120 °W in longitude and 15–50 °N in latitude, covering an area of approximately 5,700 by 3,900 km. Figure 3.1 shows the satellite zenith angle of ABI observations across the CONUS domain. As the viewing angle increases the accurate detection and tracking of DCCs becomes more difficult. This is due to the confounding between vertical and horizontal motion, and also due to the area of each pixel increasing with the zenith angle. The large sensor zenith angles in the North-West of the domain may introduce large errors in the detection and tracking algorithm, and so for the rest of this chapter only those DCCs detected east of 110 °W and south of 45 °N are included in the analysis.

Five full years of data—from the start of 2018 to the end of 2022—are used to produce a dataset of detected DCCs and their properties. This period spans all complete years of operational data from GOES-16 ABI. To improve to continuity of observations a gap-filling procedure is implemented. If time gaps between observations of

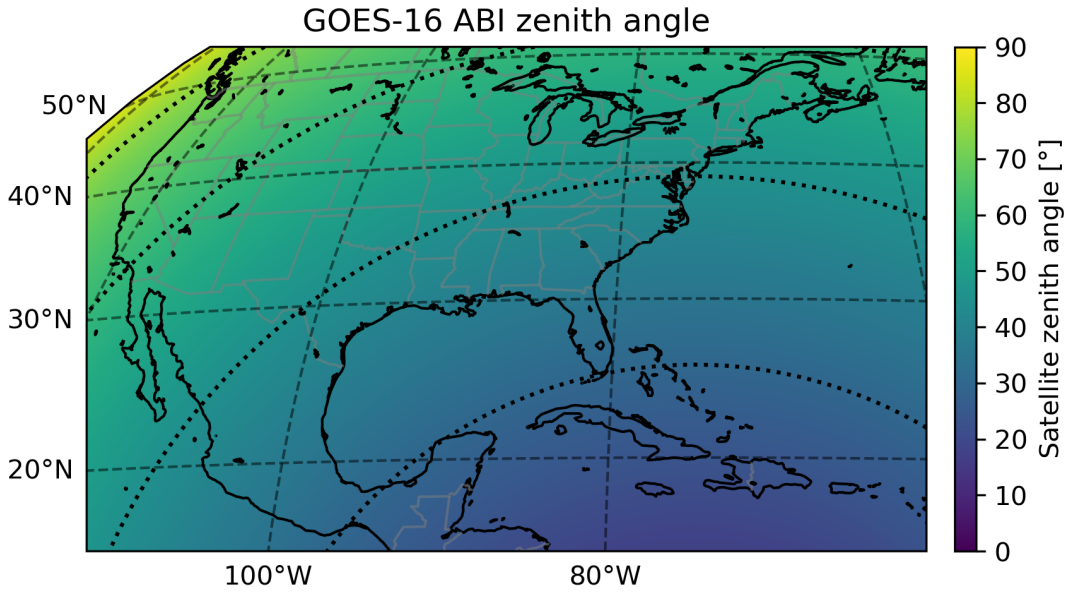


Figure 3.1: The sensor zenith angle of GOES-16 ABI observations across the CONUS domain. Dotted arcs are shown for each 15 °interval of zenith angle.

greater than 15 minutes are present, observations from the full-disc ABI scan are used to fill these gaps. Full disc imagery is typically available every 10 or 15 minutes depending on the operating mode. This gap filling is particularly important during the periods in which ABI uses its mode 4 scan pattern, in which no CONUS domain scans are made, but the full disc is scanned every 5 minutes. Using the full-disc observations allows us to maintain temporal sampling throughout these periods.

3.3 Method

Detection and tracking of convective cores and anvil clouds are performed using the *tobac-flow* method [Jones et al., 2023a] described in section 2.4. Initial detection of DCCs is performed separately over 24-hour periods spanning from 12:00:00 UTC (approximately 6 am local time over North America) to the same time the next day. This 24-hour period was dictated due to performance constraints, as the large domain combined with the high spatial and temporal resolution of ABI data results in a large

memory requirement. The start time corresponds with the minima of convective activity over land, and so was chosen to minimise the number of DCCs missed at the start and end of the detection period. Each period is extended by six ABI observations at each end to ensure at least one hour of overlap between successive days.

To track long-lived DCCs that last beyond one day, a linking algorithm is used to combine DCCs observed across multiple days. The linking algorithm combines DCCs detected at the same locations within the overlap period of two daily detection files. Splitting and merging of objects is taken into account, so a single object which splits into two, or two objects which merge into one in the subsequent file are all considered a single, tracked object. The linking algorithm is applied separately to each month of data for performance reasons.

After linking, a processing step is applied to calculate the properties of detected cores and anvils at each step of their lifecycles. Finally, core and anvil step properties are aggregated over each month of observation, and overall core and anvil properties are calculated. During this final step quality flagging is performed to isolate detected features that fail one or more quality checks. The quality criteria are split into two groups. The first set of criteria—for core or anvil removal—removes features that cannot be verified as correctly tracked DCCs due to bad data or a failure to meet the basic requirements for tracking described in chapter 2. This may happen because there are large time gaps in the dataset, or missing data due to artifacts in the ABI

Core removed if:	Core invalid if:
Initial BT – final BT $< 8\,\mathrm{K}$	Intersects edge of domain
Lifetime < 15 minutes	Intersects start of domain
Time gaps > 15 minutes	Intersects end of domain
Maximum area $> 10{,}000\,\mathrm{km}^2$	Adjacent to bad ABI data
Any NaN values in core properties	

Table 3.1: Validity criteria for detected cores. Cores which flag any of the removal criteria are removed in their entirety from the dataset. Those which flag any of the invalid criteria are retained, but removed from subsequent analysis.

Anvil removed if:	Anvil invalid if:
No associated cores	Intersects edge of domain
Lifetime < 15 minutes	Intersects start of domain
Time gaps > 15 minutes	Intersects end of domain
Maximum area < maximum core area	Adjacent to bad ABI data
Anvil detected before initial core	Associated with invalid cores
Anvil dissipated before final core ends	Maximum area reached before end of initial core
Any NaN values in anvil properties	

Table 3.2: Validity criteria for detected anvil. Anvils which flag any of the removal criteria are removed in their entirety from the dataset. Those which flag any of the invalid criteria are retained, but removed from subsequent analysis. The majority of anvils removed are due to have no associated cores, or because the anvil was observed before any developing cores.

observations. Detected features which flag any of these criteria are removed from the aggregated dataset in their entirety. This step in particular removes anvils which have no associated cores, or are not detected as initiating with a developing core.

The second set of criteria is used to identify detected cores or anvils that have not been observed over their entire extent or lifecycle. Cores and anvils which flag any of these criteria are still included within the aggregated properties dataset, but are removed from the analysis DCC properties throughout this chapter. Quality criteria for cores are listed in table 3.1, and those for anvils in table 3.2.

The complete processing pathway is outlined by the following steps:

1. Detection of cores and anvils in ABI observations over each 24-hour period.
2. Linking of overlapping objects detected in subsequent 24-hour periods over each month.
3. Calculation of core and anvil step properties.
4. Quality criteria applied to core and anvils, properties aggregated over each month

Two datasets are produced. The first consists of daily core and anvil spatial maps, with step properties, produced by processing step 3. The second, consisting of aggregated monthly core and anvil properties, is produced by step 4.

Figure 3.2 shows an example of detected cores and anvils over the CONUS region, against backgrounds of composite visible imagery (fig. 3.2 a) and 10.4 µm BT (fig. 3.2 b). Cores and anvils which are removed from the aggregated dataset are outlined with dotted lines, and those which are considered invalid are shown with dashed outlines. The large, organised convective system centred at 95 °W, 27 °N has been removed as it is intersected by a scan-line artifact later in its lifetime. The DCCs observed along the eastern edge of the domain have been marked invalid as, while they are considered true detections of DCCs, they intersect the edge of the domain.

Over the five-year observing period a total of 3,877,130 cores are detected, of which 3,615,533 are considered valid, and 1,643,030 are linked with an anvil cloud. A total of 648,345 anvils are detected, of which 391,050 are valid, and these valid anvils contain 792,522 cores. The disparity in the number of valid cores and the number of cores contained with valid anvils is due to additional filtering applied to the anvils. As the larger and longer-lived anvils are more likely to intersect the edges of the domain, they are more likely to be marked invalid than the cores. In these cases, the cores themselves are valid for analysis, but the anvil itself is excluded from analysis as its entire extent and lifetime are not captured. The exclusion of anvils that intersect the edges of the domain introduces a bias towards removing larger, long-lived systems near the edge of the domain, which should be considered when assessing the properties of the observed DCCs.

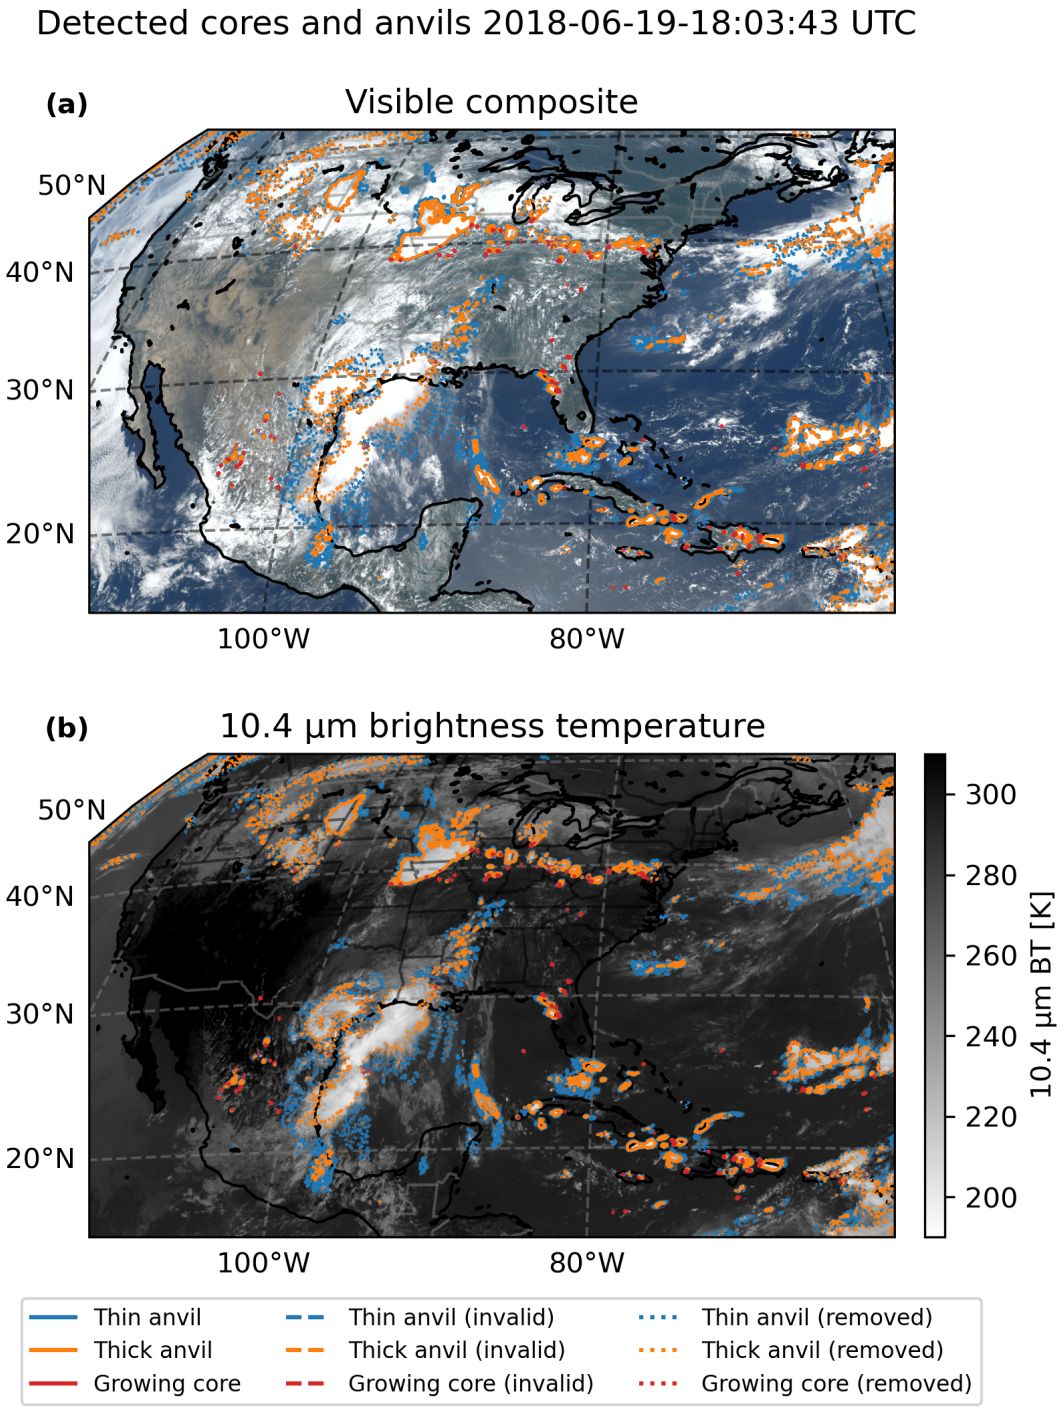


Figure 3.2: Detected cores and anvils from a snapshot of GOES-16 CONUS domain observations. Cores and anvils removed from the aggregated dataset are shown with dotted outlines. Those which are flagged as invalid are shown with dashed outlines. Detected features are shown against (a) visible composite imagery and (b) 10.4 µm BT

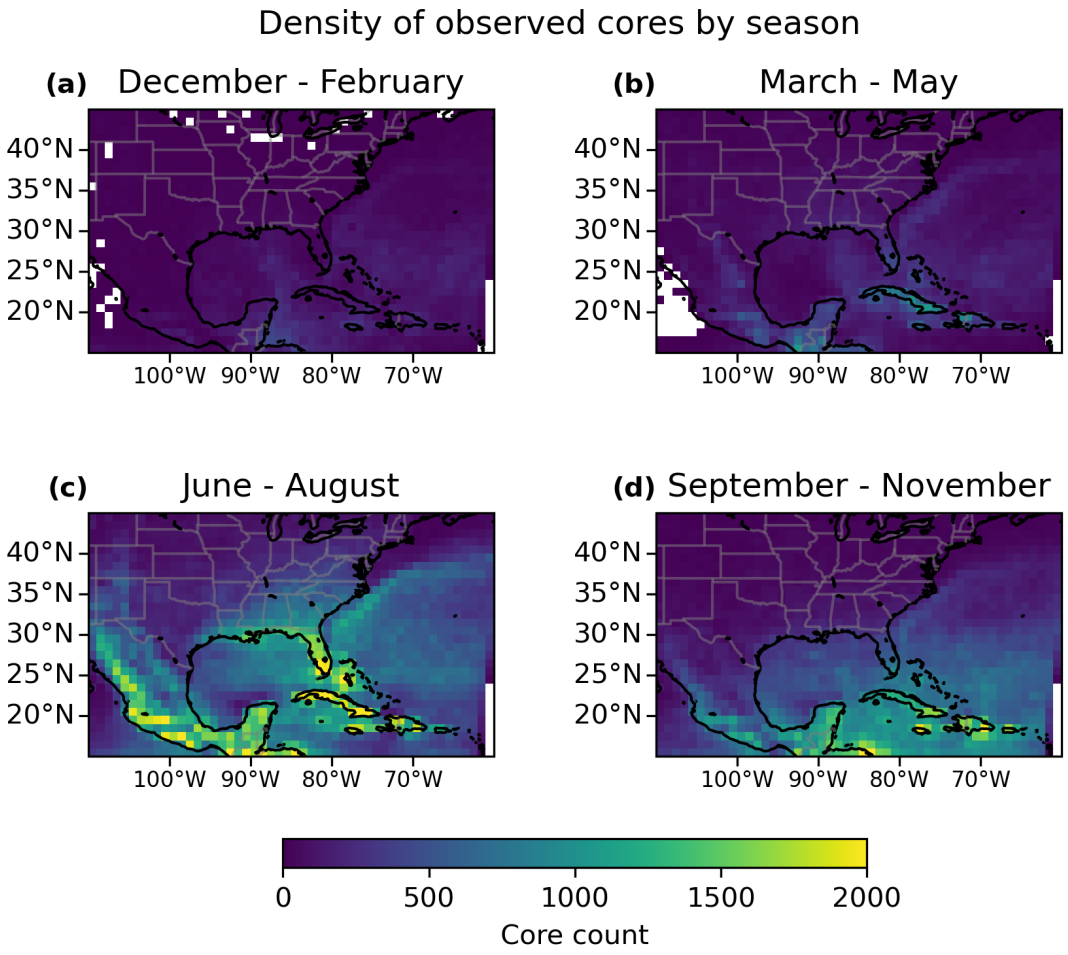


Figure 3.3: The spatial distribution of observed cores, broken down by season and accumulated into 1×1° grid boxes of latitude and longitude. The density of observed cores is greatest during summer (c) and smallest during winter (a).

3.4 Results

3.4.1 Distributions and properties of developing convective cores

To begin, the distribution of detected DCC cores throughout the dataset is investigated. Figure 3.3 shows the distribution of observed cores over North America separated by season. Large variations in the spatial distribution of cores can be observed across the different seasons. Overall, in winter and spring (fig. 3.3 a,b) there

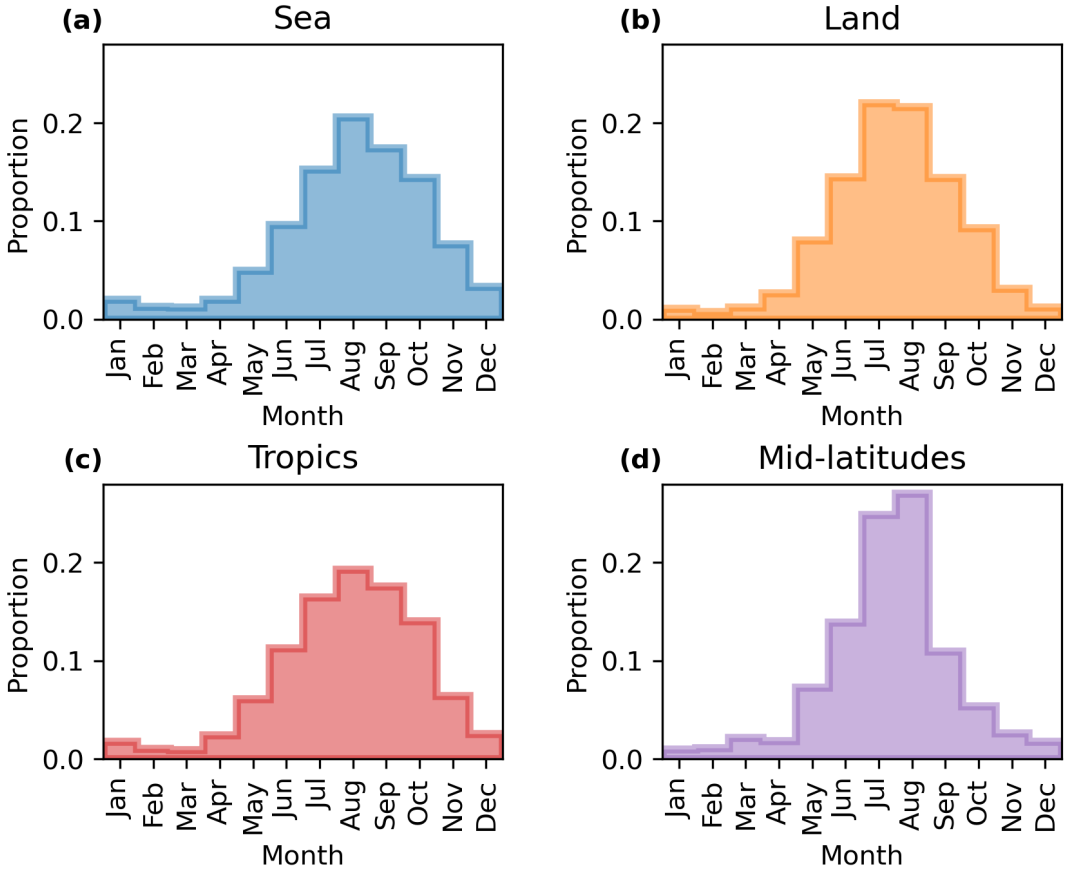


Figure 3.4: Monthly distributions of the proportion of cores detected each month over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N).

are lower rates of observed DCCs than in the summer and autumn (fig. 3.3 c,d).

In winter (fig. 3.3 a) the majority of convection observed occurs over the ocean, particularly in the areas of the Gulf of Mexico and West Atlantic associated with warm currents. In spring (fig. 3.3 b), a similar pattern shows over the ocean, and there is also an increase in convection detected over land in the Caribbean, Mexico and the central and southern USA. In summer (fig. 3.3 c) there is a large increase in convection over land and ocean, with the highest rates of convection of any season. There are, in particular, high rates of convection over Mexico, the Caribbean, the southern USA (Florida in particular) and adjacent ocean regions. In autumn (fig. 3.3 d) there is a large reduction in the number of convective cores detected over land regions. The

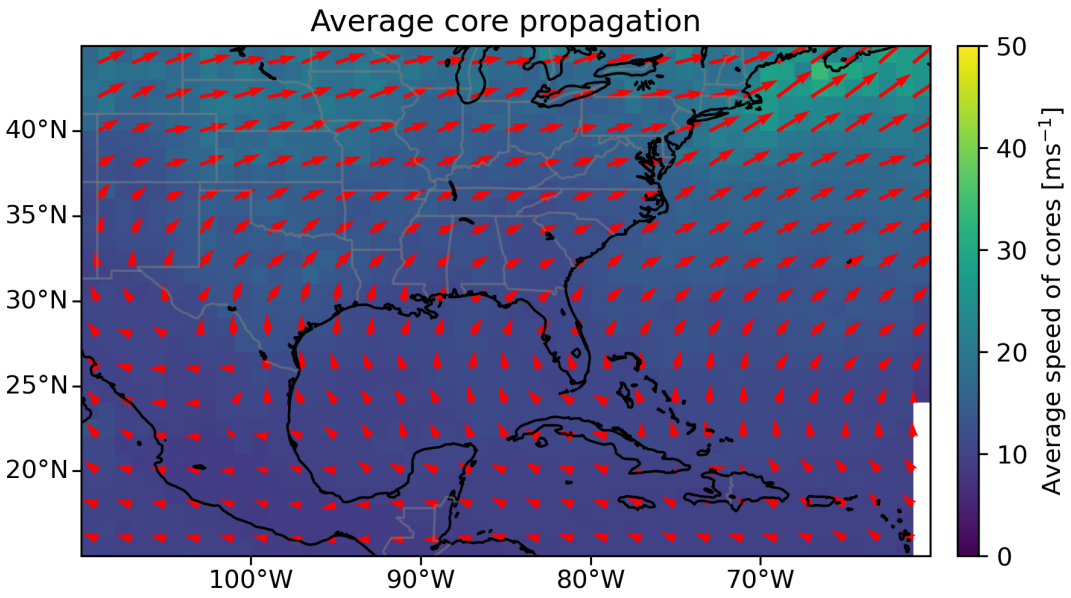


Figure 3.5: The average speed and direction of propagation of cores observed within each 1×1° grid box. The colouring shows the average speed of propagation, and the red arrows show the average direction for each 2×2° grid box

number of detections over the ocean remains high, however, indicating a possible lag in the seasonal cycle of convection over oceans compared to that over land.

Figure 3.4 shows the proportion of cores detected in each month over the annual cycle for land and sea regions. Both land and sea regions show a peak of convective activity in the summer and a low in the winter months. However, as suggested by fig. 3.3, there is a time lag of about 1 month between the annual cycle of convection over land and that over the ocean, likely due to the time lag in ocean heating. Comparing the tropics and mid-latitudes, the monthly distribution of convection is much more sharply focused on July–August in the mid-latitudes compared to the broader distribution seen in the tropics.

Figure 3.5 shows the average speed of propagation for cores observed in each 1-degree grid box, with the average direction of propagation shown by arrows for each 2-degree grid box. There is a clear change in the direction of propagation from

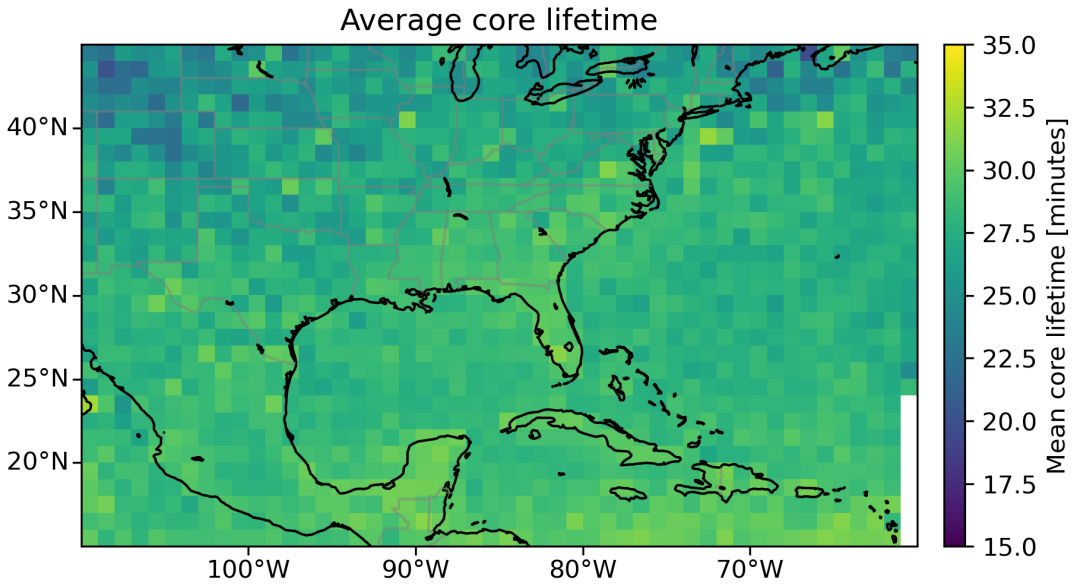


Figure 3.6: A map showing the mean lifetime in minutes of cores within each 1×1 ° grid box.

an easterly motion in the tropics (below 25 °N) to a south-westerly motion in the mid-latitudes.

Figure 3.6 displays the mean lifetime of observed cores over each 1-degree box of latitude and longitude. The lifetime is defined as the period of time over which each core is detected, which represents the time in which to core is developing vertically. Convection will continue in the core after this point, but the motion of the cloud top will instead be a horizontal divergence of the anvil. The average lifetimes appear mostly uniform across the domain, with a small reduction with increasing latitude. This reduction in the lifetime of the developing core at higher latitudes is likely due to the lower tropopause height restricting the vertical development of the cores. Similarly, there is a slight reduction in average lifetime over the more mountainous, inland regions of Mexico compared to the adjacent coastal regions. This again may be linked to the reduction in the depth of the convective cores, although in this case due to an increased surface height rather than a lower tropopause.

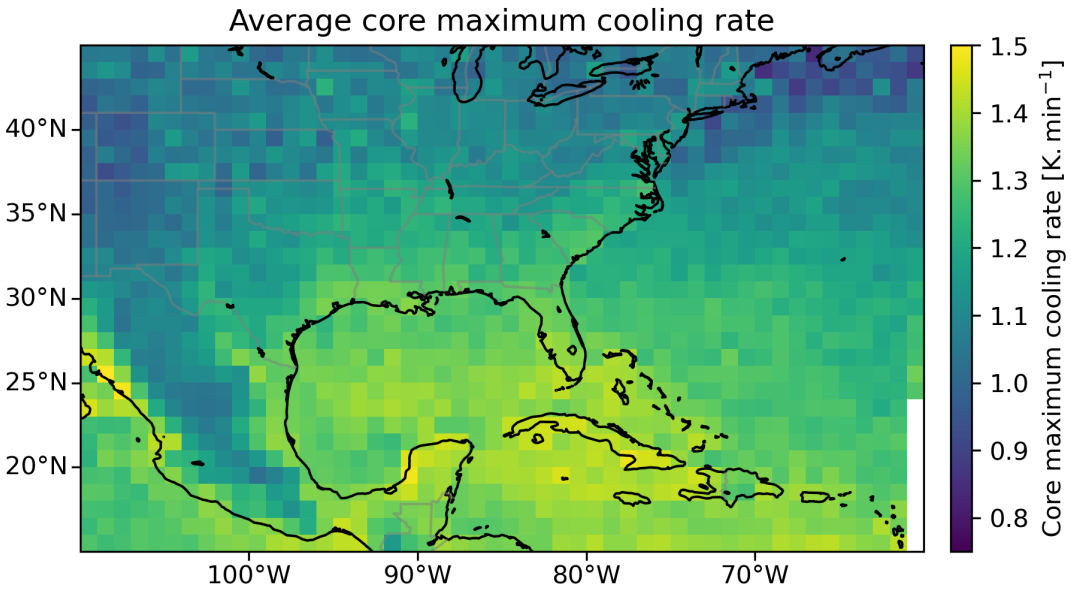


Figure 3.7: A map of the average maximum core cooling rate for each 1×1 ° grid box.

Figure 3.7 shows the average cooling rate of observed cores. The contrasts here are more pronounced than those shown in fig. 3.6. There is a reduction in the average cooling rate with latitude, likely due to the reduction in solar heating and hence lower CAPE. Orography also has a factor, with lower cooling rates observed over the mountains of Mexico and the North American Rockies. The largest average core cooling rates are observed in coastal regions, indicating potential land–sea interactions.

Figure 3.8 shows the distributions of lifetime (fig. 3.8 a), maximum area (fig. 3.8 b), cooling rate (fig. 3.8 c), and minimum BT (fig. 3.8 d). There is a peak in the observed lifetime of cores between 15 and 20 minutes, with a large tail extending beyond 60 minutes in some cases. Note that due to the requirement of a minimum of three consecutive observations, cores lasting less than 10 minutes will not be detected, truncating the distribution. In fig. 3.8 b the maximum core area peaks around 150–200 km^2, representing cores approximately 12–14 km in diameter. Figure 3.8 c shows that the core cooling rate peaks at values of 1–1.25 Kminute^{-1}, representing vertical cloud top velocities of approximately 2.5–3 ms^{-1}. The distribution is slightly truncated

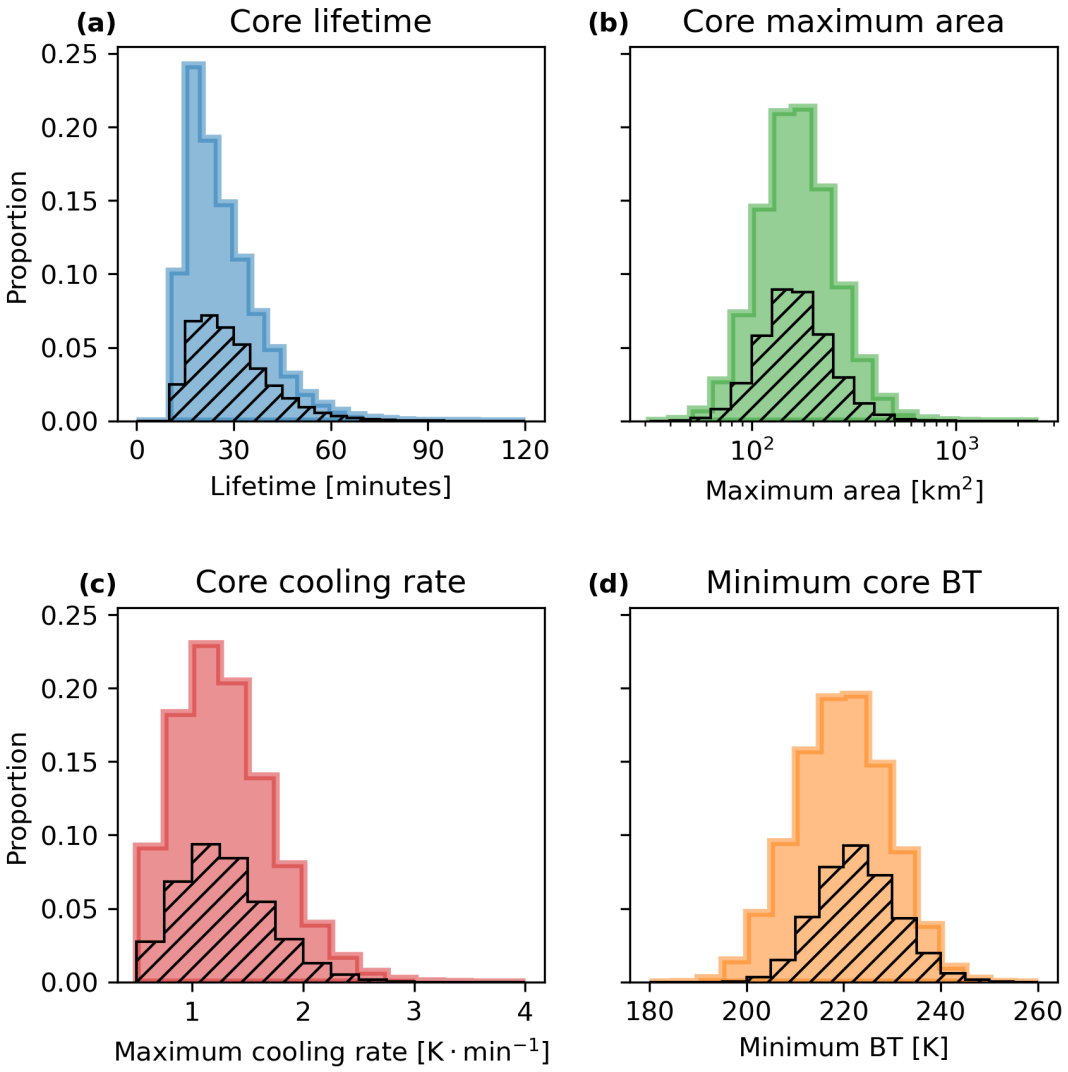


Figure 3.8: Histograms showing the distributions of observed (a) core lifetime, (b) maximum area, (c) maximum cooling rate, (d) and minimum BT. The hatched areas show the proportion of each distribution consisting of the initial detected cores within each tracked DCC.

by the requirement for detected cores to have a cooling rate of at least 0.5 Kminute^{-1}.

The minimum core BT distribution (fig. 3.8 d) peaks at around 220 K, supporting the theory that the radiative tropopause sets the height of the convectively active layer of the atmosphere. Many cores reach temperatures warmer than 220 K, indicating either a lower LNB or dilution of the convective updraft leading to reduced buoyancy. In addition, a large number of cores reach BT colder than 220 K, which, for the coldest cases, likely indicates the occurrence of overshooting tops.

The hatched areas in fig. 3.8 show the proportion of each distribution consisting of the initial cores of tracked DCCs. As their development will not be masked by an overlaying anvil cloud, these cores are expected to provide the most complete measurements of developing core properties. For the lifetime, maximum area and cooling rate these show close similarities to the shape of the distribution for all cores, indicating that masking due to anvils does not significantly affect the ability to observe the core properties. For the minimum core BT, the initial cores tend to have warmer BT than later cores, indicating an increased presence of cold cores and overshooting tops in organised convection.

Figure 3.9 shows how the mean core lifetime, maximum area, cooling rate and BT vary between different regions. There is little difference in the mean core lifetime between regions, echoing the results seen in fig. 3.6. Maximum core area is larger over land than sea, and larger in the mid-latitudes than in the tropics. There is little difference in the cooling rate between sea and land, but on average cores in the tropics cool $\sim$0.2 Kminute^{-1} faster than those in the mid-latitudes. There is a small difference in the minimum BT of cores observed over land and ocean, with the former $\sim$1.5 K colder than the latter. A more noticeable difference is seen with latitude, with cores in the tropics having, on average, minimum BT $\sim$3 K colder than those in the mid-latitudes.

Overall, fig. 3.9 show differences in core intensity between land and ocean, and

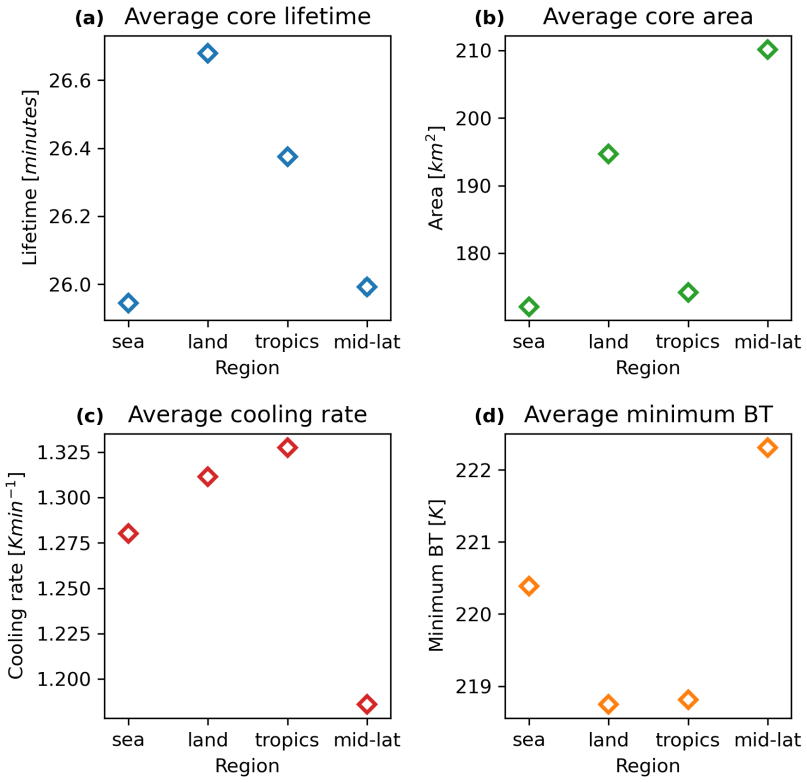


Figure 3.9: The mean observed (a) lifetime, (b) maximum area, (c) maximum cooling rate, (d) and minimum BT for cores detected over land, sea, tropics and mid-latitudes. Error bars are not shown as the standard error of the mean is too small to be visible.

between the tropics and mid-latitudes, with these differences most prominent in the minimum BT. Due to the large number of tracked cores, the standard error of the mean for each point is very small, and so from a simple perspective these differences appear significant. However, the impact of systematic bias on these properties must also be considered, in particular the impact of the sensor zenith angle as discussed in sections 2.2.1.1 and 3.2, which may impact the detection accuracy and measured properties of tracked DCCs. As more of the ocean observations are in the tropics, and more of the land DCCs observed in the mid-latitudes, covariance should be expected between the average properties of DCCs in these regions. In fig. 3.9 b larger average areas are seen for both land over ocean and mid-latitudes over the tropics, indicating that the differences may be the result of the expected systematic bias in area caused

by sensor zenith angle. However, for the other three properties, while convective cores over land have longer lifetimes, higher cooling rates and colder minimum BT than those over the ocean, cores in the tropics also have the same indicators of more intense convection when compared to the mid-latitudes. This opposes the expected covariance between these regions, and indicates that the differences seen are likely to be real and significant, and that the systematic error is smaller than the differences seen between land and ocean and between tropics and mid-latitudes. These differences may result from changes in different convective processes as the difference in cooling rate between tropics and mid-latitudes is as large when comparing sea and land.

Figure 3.10 a compares how the average lifetime of detected cores changes with their maximum observed cooling rate. The observed core cooling rates range from $0.5\,\mathrm{K{\cdot}minute^{-1}}$ to $4\,\mathrm{K{\cdot}minute^{-1}}$, representing cloud top growth rates between $1.25\,\mathrm{ms^{-1}}$ to $10\,\mathrm{ms^{-1}}$. For low values of core cooling rate, there is a similar, positive linear relationship between cooling rate and lifetime between all regions, indicating that more intense convection leads to longer periods of cores being observed. However, beyond cooling rates of around $1.5\,\mathrm{Kminute^{-1}}$ there is a peak in the core intensity–height relationship, with larger cooling rates leading to shorter lifetimes. This may be explained that the cores with faster cooling rates reach their level of neutral buoyancy faster, and therefore show cooling cloud tops for a shorter time period. The inflexion in the core lifetime may also explain why there are no apparent differences in average lifetime between different regions, despite differences in other convective properties.

Figure 3.10 b shows how the minimum BT changes with the maximum cooling rate of the core. Unlike fig. 3.10 a, there is a continuous decrease in BT with core cooling rate, indicating that faster cooling rates tend to result in cores with colder CTT and hence higher CTH. Figure 3.10 b also provides context for the results seen in fig. 3.10 a. If it is assumed that cores are detected around the freezing level (273 K), then the overall temperature change over the core lifetime ranges from $\sim$45 K for the

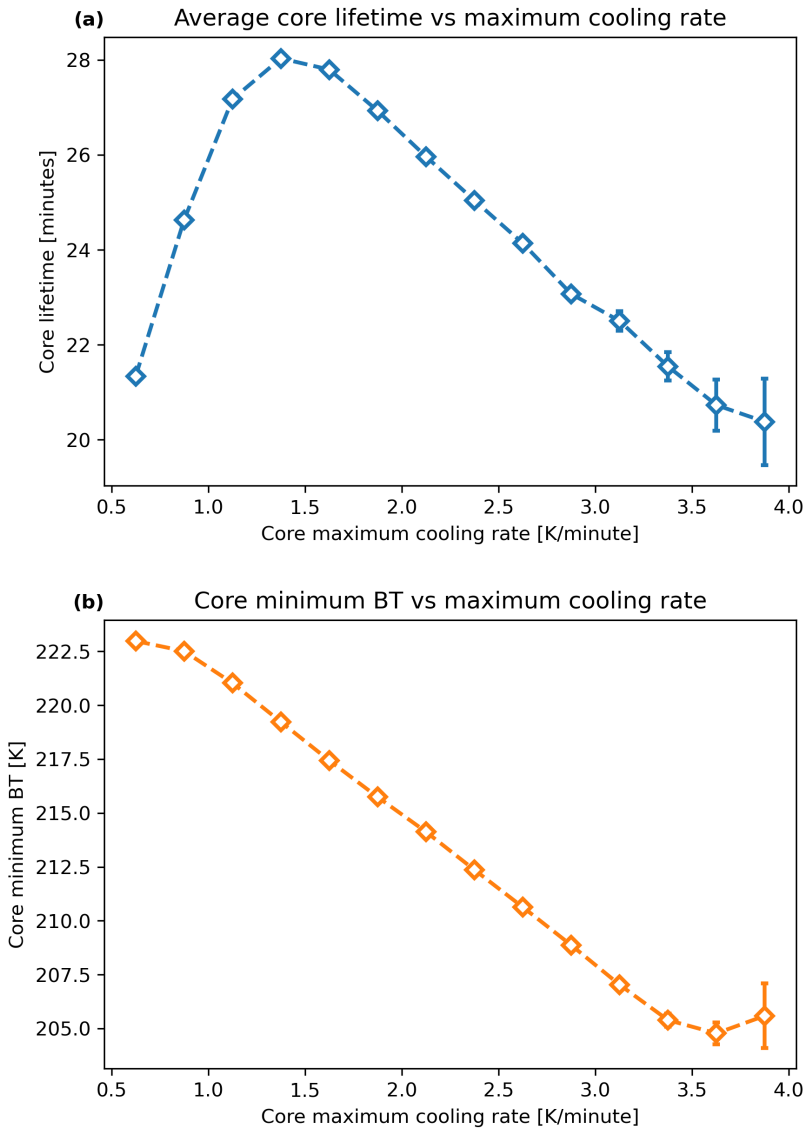


Figure 3.10: Average core lifetime (a) and minimum BT (b) with increasing core cooling rate. Error bars show the standard error of the mean, and are large enough to the visible only for the largest cooling rates.

least intense cores to ∼65 K for the most intense cores. Although this continues to increase with the cooling rate, the proportional change is less than that of the core cooling rate itself. As the core lifetime can be approximated as the core BT change divided by the cooling rate, the larger proportional change in cooling rate will result in a shorter lifetime for more rapidly cooling cores.

By multiplying the average core lifetime by the vertical velocity estimation from the cooling rate, an estimate for the convective layer height be can found. For the longest lived cores with cooling rates around 1.5 K·minute^{-1} this results in an estimated convective layer height of 6.3 km. Using the moist pseudo-adiabatic lapse rate a total temperature difference across the observed lifetime of these cores can be calculated as 38 K., resulting in an average initial detection temperature of approximately 255 K. For the most intense observed cores, with cooling rates greater than 3 K·minute^{-1} (7.5 ms^{-1}), the convective layer height is estimated as 11 km, resulting in an initial temperature of 270 K. While in both cases the estimate for the initial temperature of observations is above the freezing level, this still agrees with the explanation for the inflexion in the growing core lifetime proposed in the previous paragraph. Beyond a certain intensity, the lifetime of the growing phase of a convective core is constrained more by the convectively active layer height than the updraft intensity.

3.4.2 Diurnal cycle of convective cores

Figure 3.11, shows the mean time of detection for cores detected within each 1-degree grid square. The most notable feature in fig. 3.11 is the strong land-sea contrast, with the majority of land regions showing convective activity occurring during the afternoon, and the majority of ocean regions showing activity before midday. In addition, there are a few major features of the detected time of initiation across both land and sea regions. In coastal regions in the Gulf of Mexico and the Caribbean Sea,

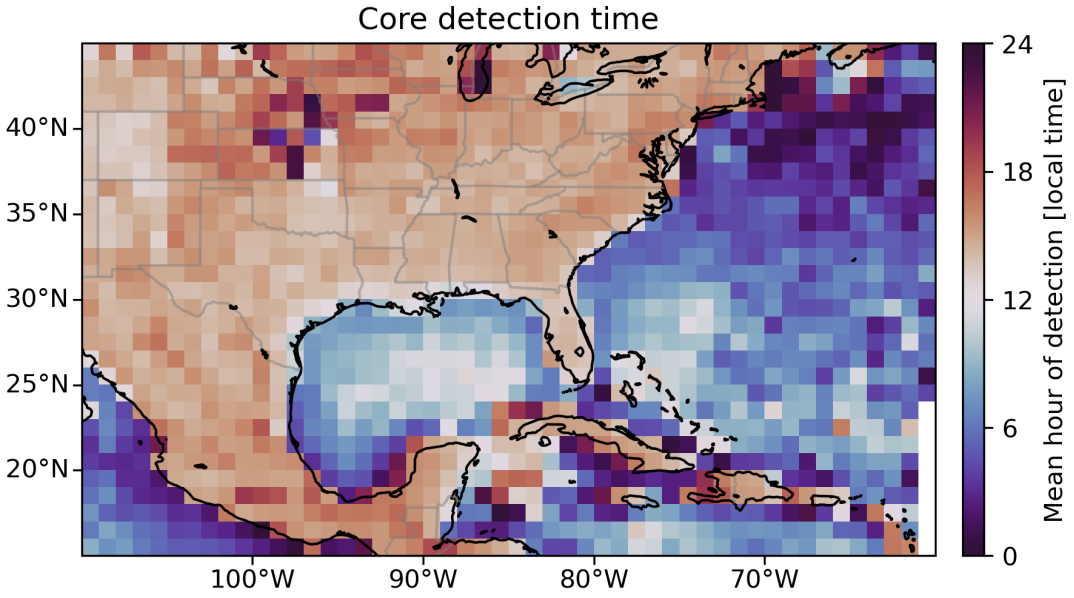


Figure 3.11: The average time of day of initiation of cores observed within each 1×1° grid box. The time of initiation is calculated as the local solar time based on longitude, and the mean is calculated as the circular mean over a 24-hour period to account for the cyclical aspect of the hour of day.

earlier initiation times are seen closer to coastlines, while regions further from land have average times of detection closer to midday. Over land there is also a later time of initiation over the Northern Great Plains (90–100 °W, 37–47 °N).

Figure 3.12 shows the distribution of core detections across the diurnal cycle for land, sea, tropics and mid-latitude regions. Over land there is a sharp peak in convective activity initiating in the mid-afternoon between 2 and 3 pm, with a tail extending into the night-time. Over the sea, the distribution of core detections is much more uniform across the diurnal cycle. There is still a peak observed in the early hours of the morning (3–6 am), and a low in the evening (6 pm), but the difference between these is much less pronounced than that over land. Both the tropics and mid-latitudes have peaks in convection at the same time as that seen for all land regions in fig. 3.12 b. The peak for mid-latitudes is slightly broader, with more convection occurring later in the day.

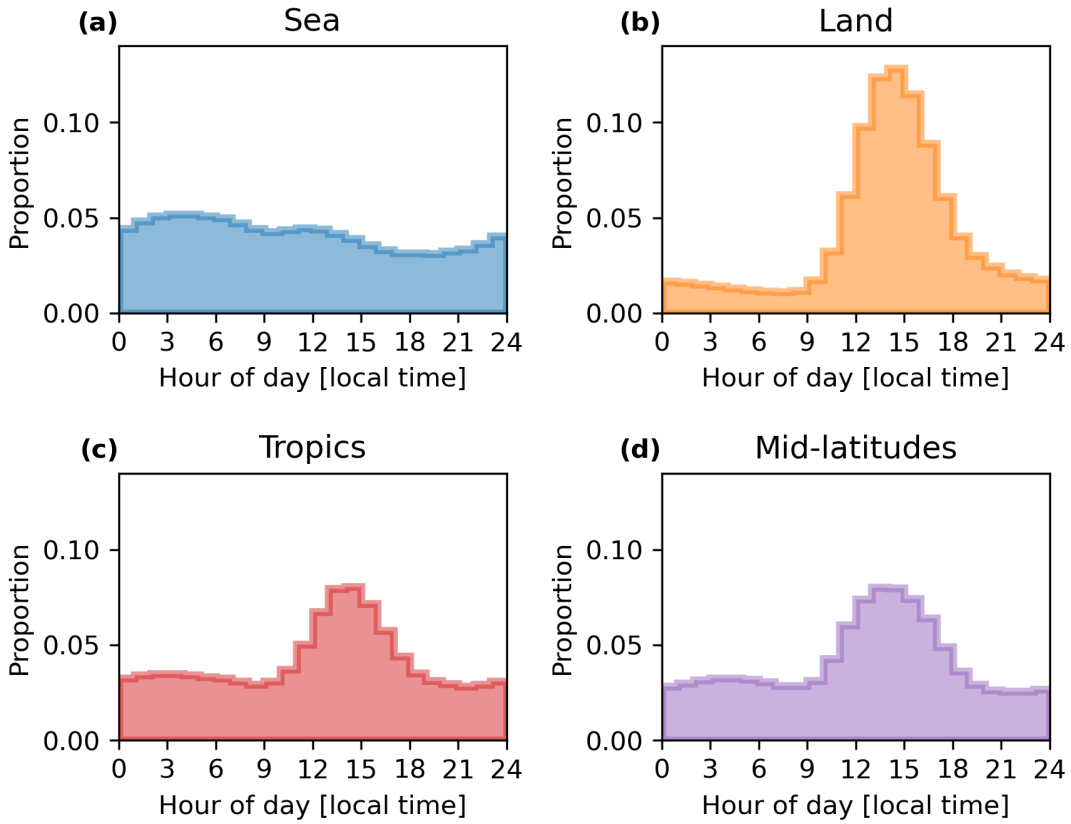


Figure 3.12: The diurnal distributions of the local time of detection for cores, binned by hour detected over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N).

Figure 3.11 showed a noticeably later average time of detection over the Northern Great Plains (NGP) region. Figure 3.13 shows how the diurnal distribution of core detection time varies within this area, which is defined as 37–45 °N, 90–100 °W. In contrast to the distribution seen over all land regions in fig. 3.12 b, the NGP region shows a later peak of convective activity between 3 and 4 pm, as well as much higher rates of convection observed during the nighttime and into the early morning until around 7 am. Previous studies have found a similar bimodal distribution in convective precipitation over the NGP Li et al. [2021].

In fig. 3.14 the effect of the distance to the coast on core detection time is examined over the Gulf Coast region (22.5–32.5 °N, 82.5–100 °W). Negative distances indicate

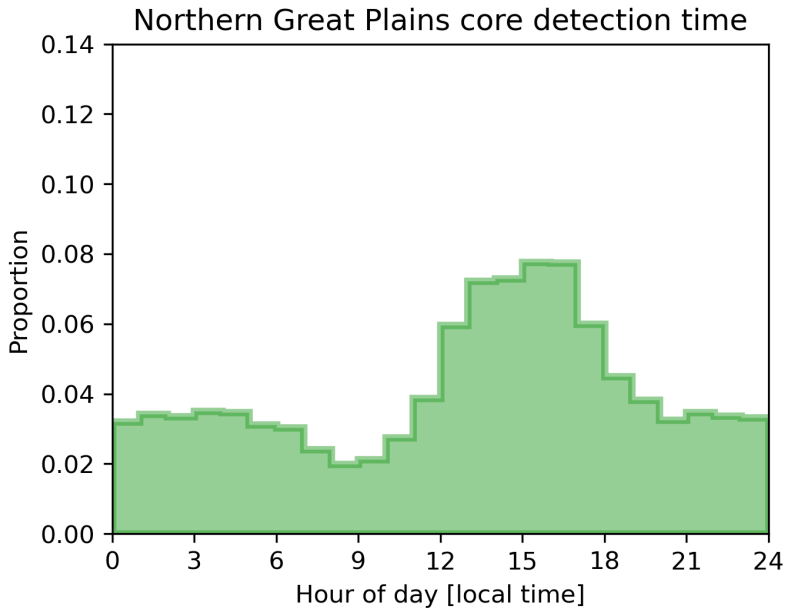


Figure 3.13: The diurnal distribution of cores detected in the NGP region (37–45 °N, 90–100 °W).

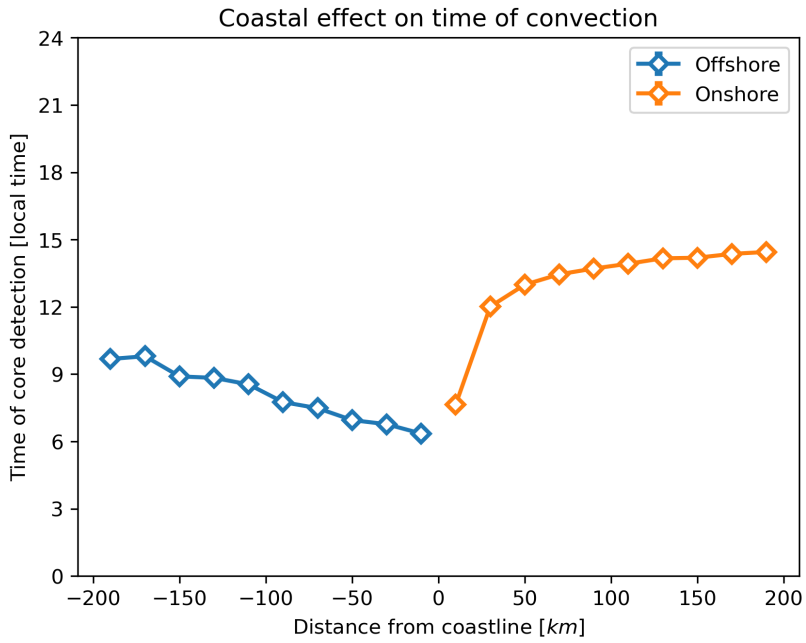


Figure 3.14: The change in average time of core detection with distance from the coastline. Negative distances along the x axis show cores detected further to sea, and positive distances show cores detected further inland. Error bars are not shown as the standard error of the mean is too small to be visible.

locations further offshore, and positive distances those locations further inland. Figure 3.11 showed a change in the average time of detection of cores in locations closer to the coast, and that effect is shown again here. For cores detected over the sea, there is a linear decrease in the average time of detection as the distance to the coast increases. For cores detected over land, there is a sharp decrease in the time of detection very close to the coast, although further than 50 km from the coast this becomes a shallower linear gradient. The linear change of core detection time with distance from the coastline may be linked to offshore and onshore breezes triggering convection. Cores over land see a reduction in the variance of the time of detection between 50 and 150 km from the coast, which may also indicate that an external forcing from sea breeze is triggering convection in these areas, constraining the time of initiation of convection.

Figure 3.15 shows how the mean maximum cooling rate of cores changes across the diurnal cycle. While over the sea there is little change in cooling rate, over land there is a noticeable increase in the cooling rate throughout the daylight hours. This leads to a peak at around 4 pm, before the cooling rate falls back to its nighttime levels which are lower than that seen over ocean. Both tropics and mid-latitudes show similar diurnal cycles to land, albeit with a difference of 0.2 $K minute^{-1}$ across the entire diurnal cycle.

3.4.3 Distributions and properties of observed anvil clouds

In this section, the properties of anvil clouds in the dataset are examined. Anvils are detected and tracked independently from cores. Although anvils must be associated with observed cores at the start of their lifetime, tracking of anvils continues beyond the extent of the observed core lifetime. This also allows the detection of anvils that are associated with multiple cores, providing insight into the effects of convective organisation on anvil properties.

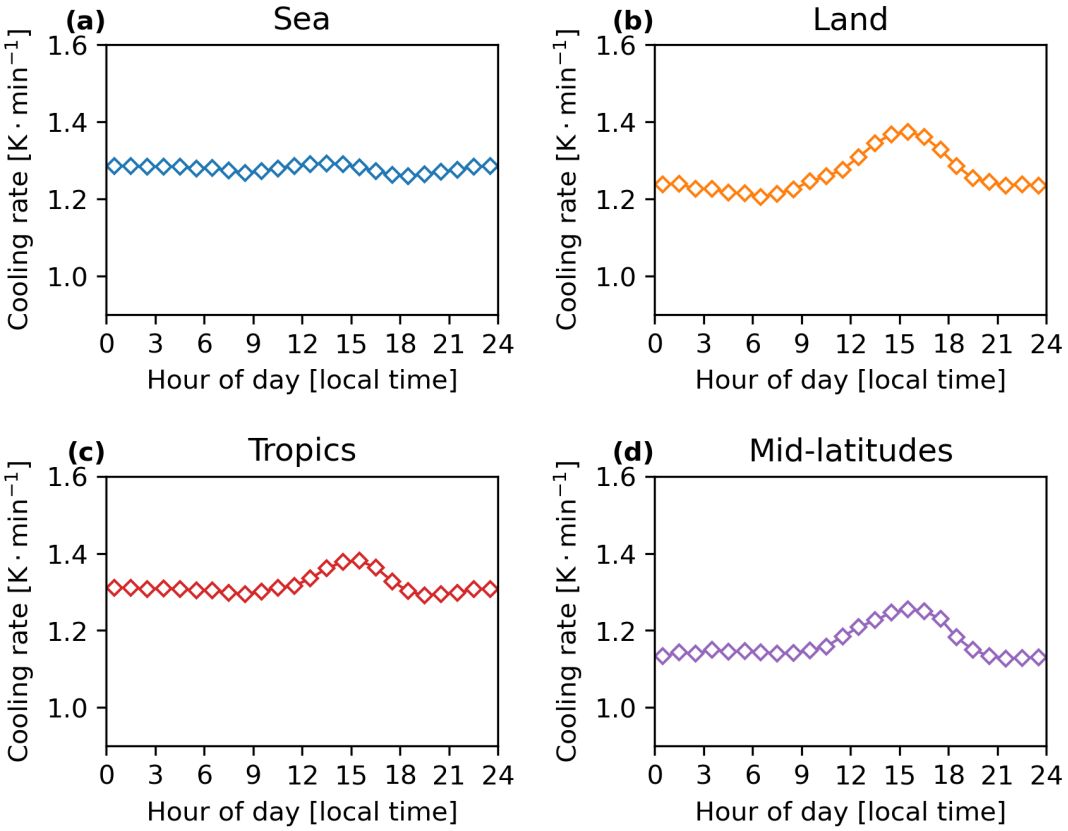


Figure 3.15: The diurnal cycle of the maximum cooling rate of cores detected over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N). Each point shows the mean of the maximum cooling rate of cores detected during that hour. Error bars are not shown as the standard error of the mean is too small to be visible.

Figure 3.16 shows the counts of anvils for each 2×2 °grid box, separated by season. There is a similar seasonal cycle and distribution to fig. 3.3. In winter (fig. 3.16 a) and spring (fig. 3.16 b) there are low rates of convection, with the majority of convection observed over warm ocean regions. In summer, fig. 3.16 c shows the highest rates of anvil detections, with a large increase in the observations of anvils over land. In spring (fig. 3.16 d), the number of anvils observed over the ocean remains high, but that over land reduces.

Figure 3.17 shows how the annual distribution of anvil detections changes by month across different regions. While similar to the annual distribution of cores

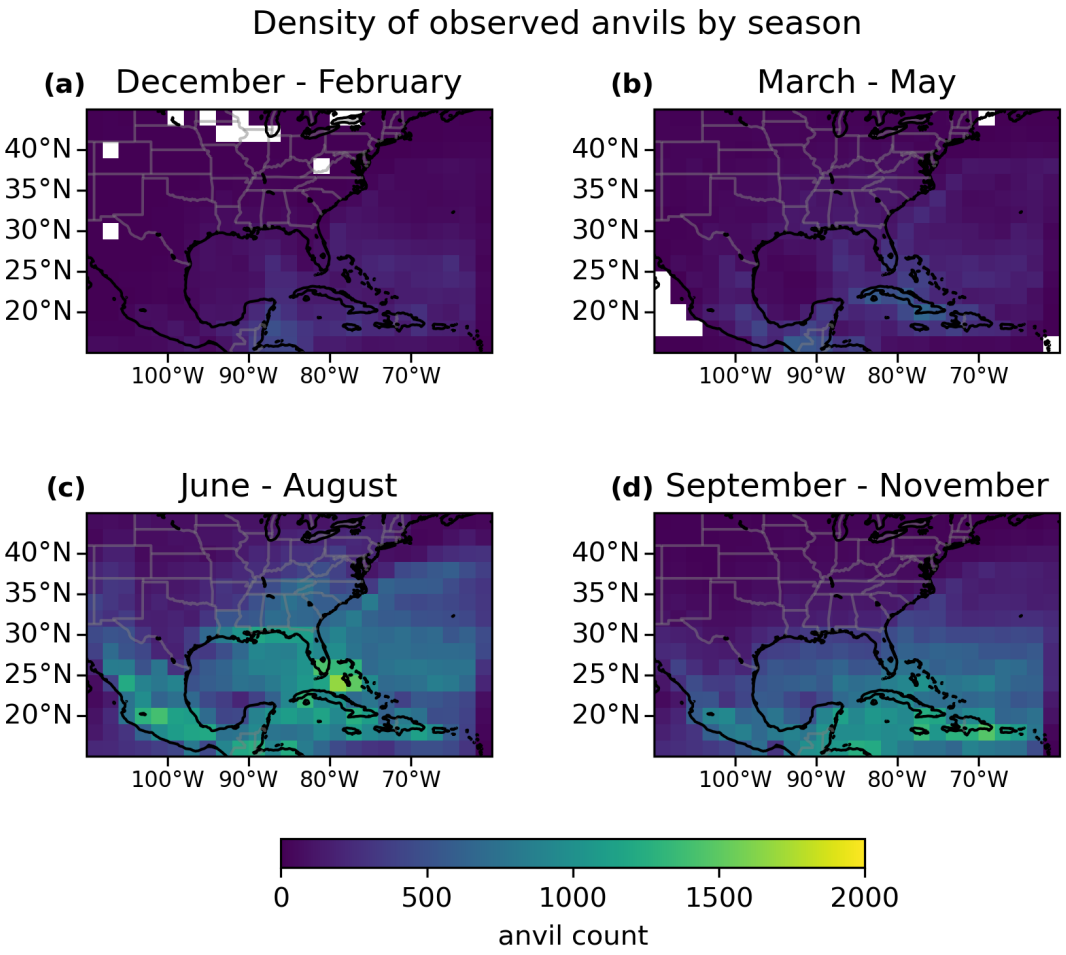


Figure 3.16: Maps showing the spatial distribution of observed anvils for (a) winter, (b) spring, (c) summer and (d) autumn, each binned to a 2×2 °grid.

shown in fig. 3.4, it shows a later peak over oceans in August–September compared to July–August over land. The mid-latitudes have a sharper peak around these two months, while the tropics show a broader distribution of anvil detections throughout the year.

Figure 3.18 shows the average propagation speed and direction for anvils in the same manner as shown for cores in fig.3.5. In the extra-tropics (>30 °N), there is generally a westerly motion, without the southerly motion seen in the cores. This westerly motion corresponds to the prevailing high-level winds. The change in direc-

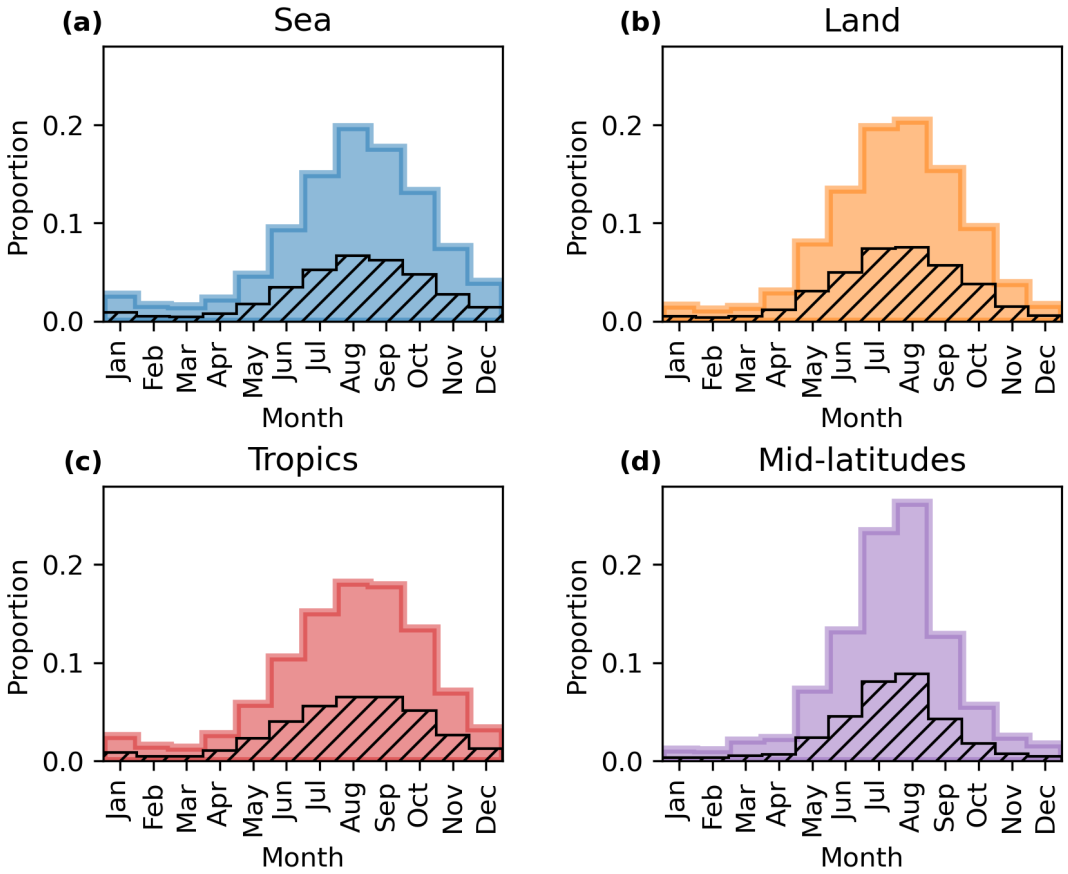


Figure 3.17: Monthly distributions of the proportion of anvils detected each month over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N). The hatched area shows the proportion of the distribution consisting of anvils with multiple cores.

tion and differences in the speed of propagation between anvils and cores indicates a typical shear between the two. Over the tropics (<30 °N) no clear overall motion of anvils is apparent.

Figure 3.19 shows maps of the average anvil properties for the maximum area (a), lifetime (b), BT (c), and number of cores (d) over each 2×2 °grid box. Both the anvil maximum area and lifetime increase towards the northwest corner of the domain. While this may be due to the tendency of MCSs to initiate East of the Rocky mountains [Feng et al., 2019], it may also be a sign of a systematic bias affecting the

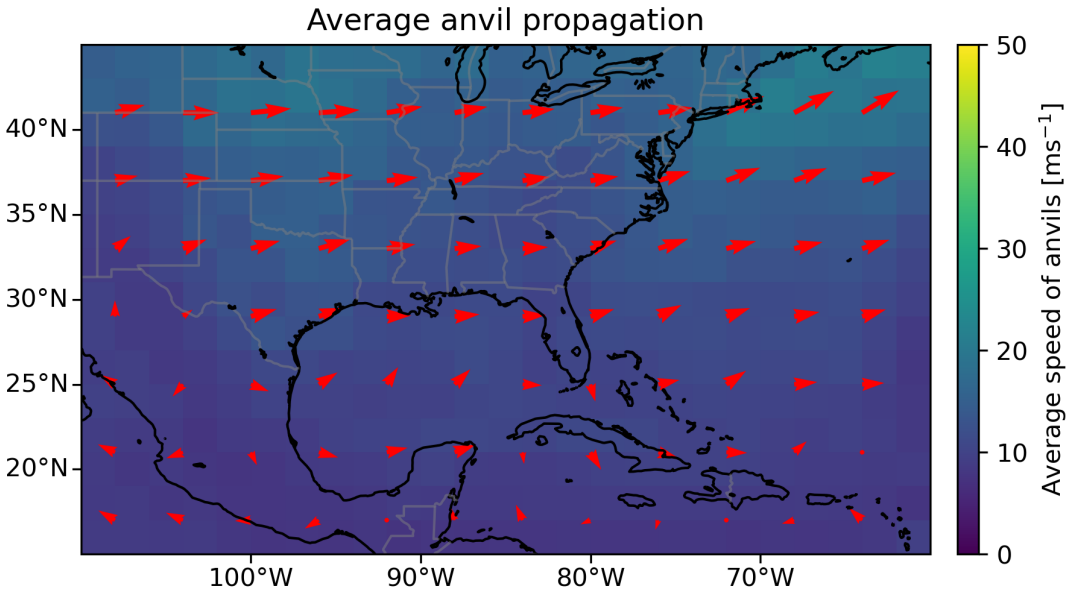


Figure 3.18: The average speed and direction of propagation of cores observed within each 2×2˚ grid box. The colouring shows the average speed of propagation, and the red arrows show the average direction of propagation for each 4×4˚ grid box

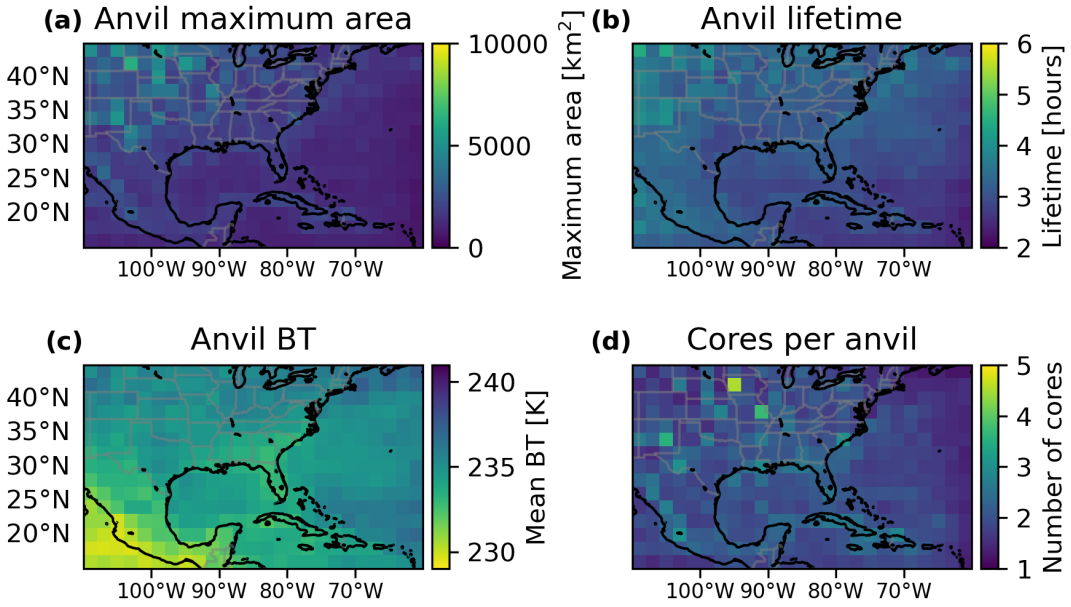


Figure 3.19: Maps showing the spatial changes in the averages of (a) anvil maximum area, (b) lifetime, (c) BT and (d) number of cores, binned to a 2×2˚grid.

anvil areas. Both the increase in average anvil area and lifetime shown in fig. 3.19 a and b appear to correlate with the satellite zenith angle (fig. 3.1). As the minimum core and anvil area requirement in the detection algorithm is determined by the number of pixels, the minimum area in km^2 will increase with the zenith angle. As the anvil lifetime tends to correlate with the maximum area, this zenith angle bias may contribute to the patterns seen in both fig. 3.19 a and b.

The average anvil BT becomes colder towards the southwest of the domain, while the average number of cores per anvil are too variable to see any clear spatial trends. As very large MCSs with many cores are both rare and large outliers, they introduce a large amount of variability into the average number of cores per anvil when gridded by the initiation location.

Figure 3.20 shows the distributions of anvil properties, with the proportion consisting of multiple-core anvils shown by the hatched area. The anvil maximum area distribution has a similar shape for both thick anvils (fig. 3.20 a) and thin anvils (fig. 3.20 b), with the mean shifted to large values for the thin anvil. It should be noted that the thin anvil area includes that of both thick and thin anvil regions, so will always be larger than that of the thick anvil alone. The anvil lifetime distributions for the thick (fig. 3.20 c) and thin (fig. 3.20 d) anvils show a similar relationship, with a shift of the distribution to longer lifetimes for the thin anvil. In both cases, despite the log scaling on the x-axis, there is a long tail towards larger area values and longer lifetimes. Furthermore, this large tail consists primarily of multiple-core systems, as shown by the hatched area, indicating the impact of organisation on the area and lifetime of DCCs.

Figure 3.20 e and fig. 3.20 f show the average and minimum anvil BT over each detected anvil. The minimum anvil BT has a broader distribution than that of the average anvil BT. Multiple core anvils generally have colder anvil BT, particularly so for the minimum BT. The cold tail of the minimum BT, with values less than 200 K,

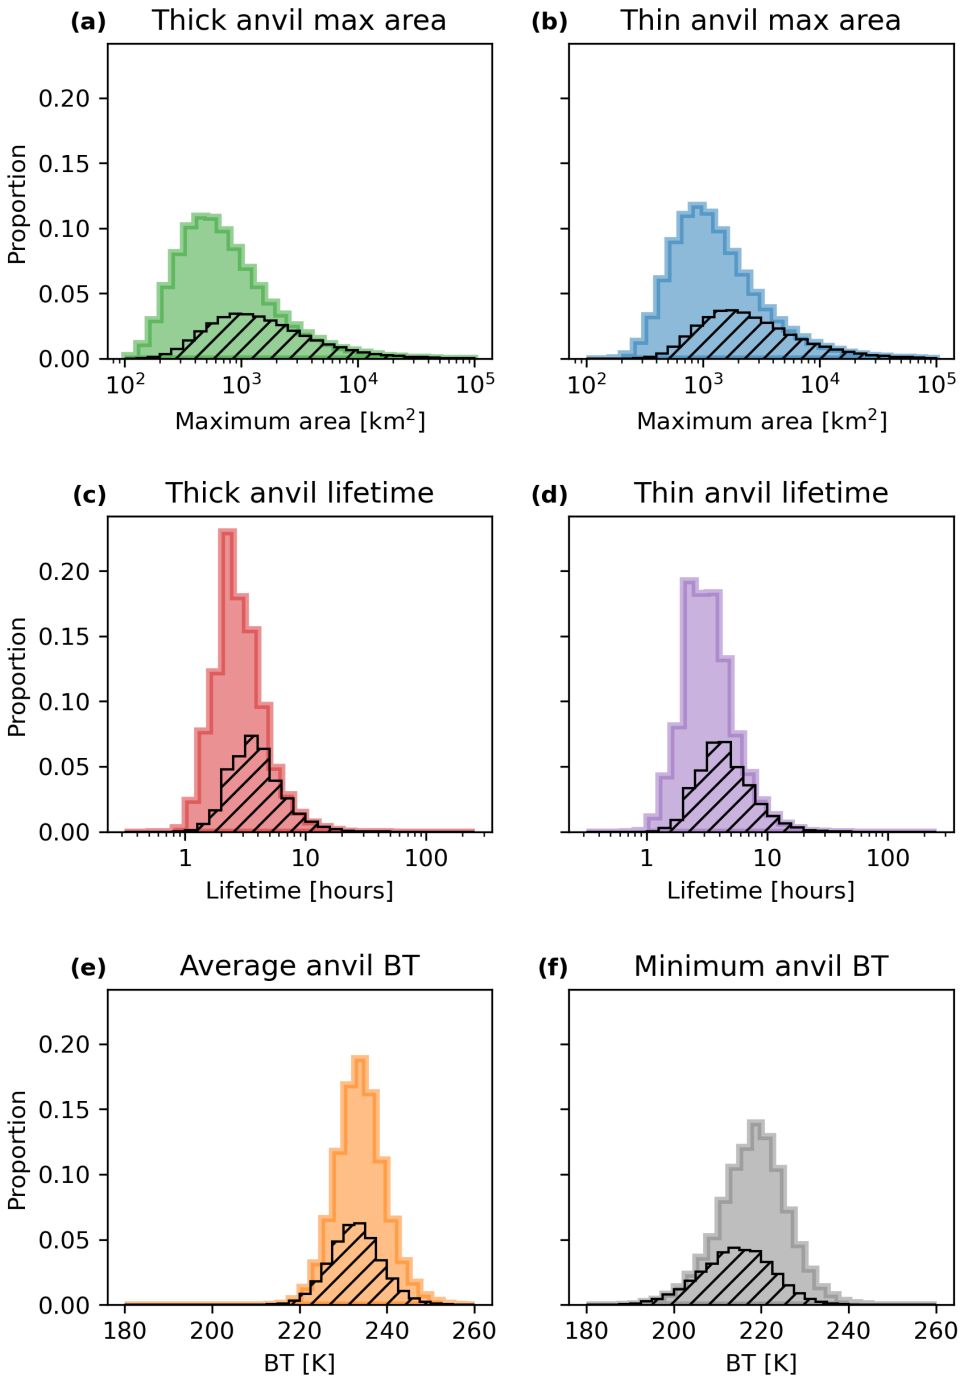


Figure 3.20: Histograms showing the distribution of observed anvil properties, with the hatched area showing the proportion associated with multiple core anvils. (a) The maximum area of the thick anvil cloud; (b) the maximum area of the thin anvil cloud (which includes both the thick and thin anvil regions; (c) the lifetime of the thick anvil and (d) the thin anvil; and (e) the average and (f) minimum observed BT within each anvil.

indicates the presence of overshooting tops within these organised systems.

Figure 3.21 a, b, c and d the average of the thick anvil maximum area, thin anvil maximum area, thick anvil lifetime and thin anvil lifetime for anvils with different numbers of cores. In all cases these areas and lifetimes increase with an increasing number of cores. In particular, the increase in the number of cores has a large impact on the maximum area of anvils, The most organised systems, which contain 10 or more cores, have areas that are more than two orders of magnitude greater on average than the anvils of isolated DCCs.

Figure 3.21 e and f show the distribution of the number of cores associated with each detected anvil, and the proportion of total anvil coverage attributed to anvils with different core counts respectively. Overall, the vast majority of anvils detected are isolated DCCs with only a single core. For anvils with greater than five cores, the number of systems observed drops to such a level that grouping of these DCCs into bins of 6–9 cores and 10 or more is performed to ensure that there is a representative sample size of each group. Despite making up the majority of observed DCCs, single core systems only make up 12% of the total anvil coverage. Instead, despite being few in number, the most organised DCCs with ten or more cores are responsible for the majority of anvil coverage due to their large area and lifetime. The large area and lifetime of these systems—seen in the long tails of those distributions in fig. 3.20—compound to result in this large coverage.

Figure 3.22 shows how the averages of the maximum areas, lifetimes, BT and number of cores vary between different regions. For both isolated and multi-core convection, the average areas and lifetimes shown in fig. 3.22 a and b respectively are larger over land than sea, and in the mid-latitudes than in the tropics. For the anvil BT shown in fig. 3.22 c, both the mean and minimum BT of anvils is colder over land than sea, and also colder in the tropics than in the mid-latitudes. DCCs over land tend to have more cores than those over the sea.

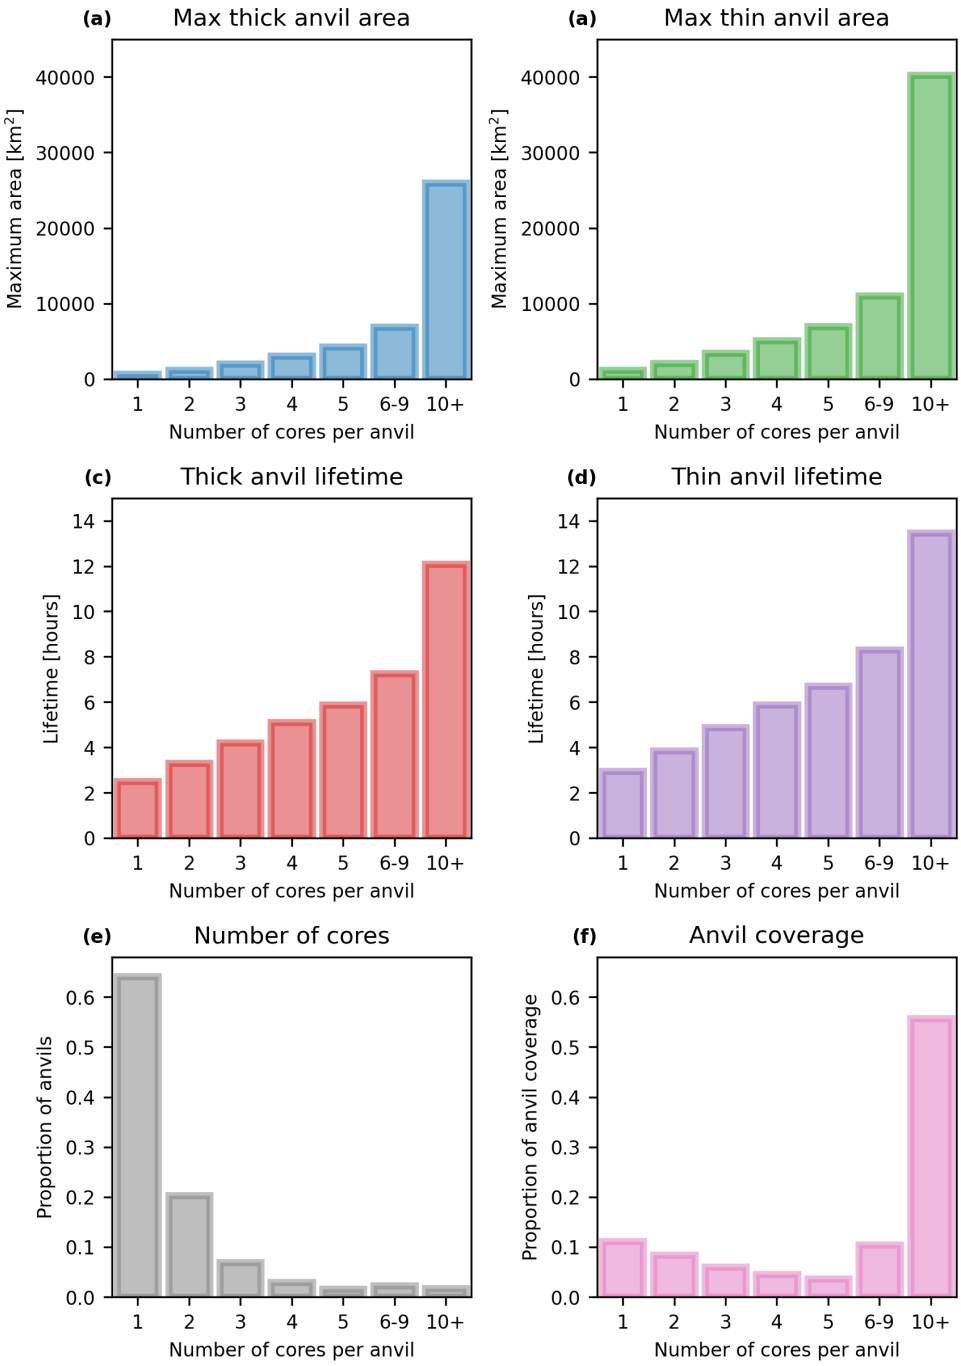


Figure 3.21: The change in the average area and lifetime of anvils with different number of cores. Mean values of (a) maximum thick anvil area, (b) maximum thin anvil area, (c) thick anvil lifetime and (d) thin anvil lifetime all increase with increasing number of cores. (e) The proportion of all anvils with different numbers of cores and (f) the fraction of total anvil coverage attributed to those anvils.

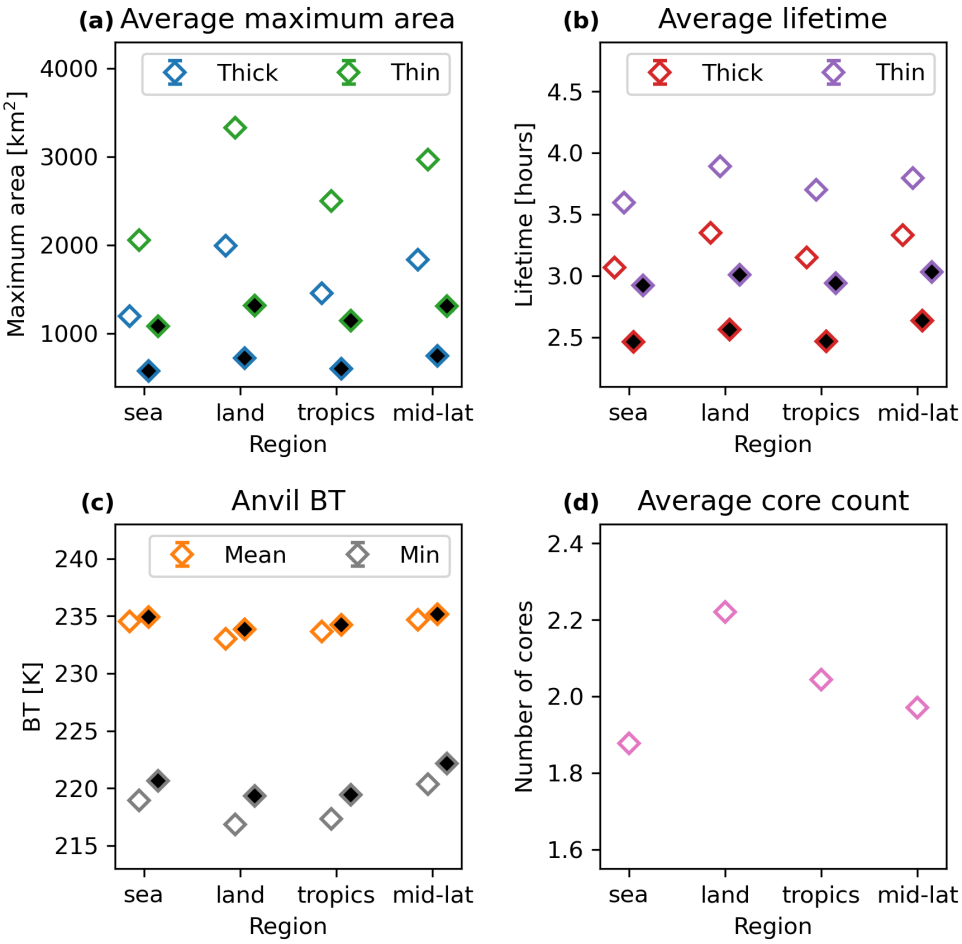


Figure 3.22: The change in the average (a) areas, (b) lifetimes, (c) BT and (d) number of cores of anvils observed over sea, land, tropics and mid-latitudes. The points with no fill show the average across all DCCs, whereas the points with the black fill show the average for isolated DCCs (with one core) only. Error bars are not shown as the standard error of the mean is too small to be visible.

As discussed in section 3.4.1 regarding fig. 3.9, it is expected that there will be some systematic bias in these mean properties, and that this error will be covariant for land and mid-latitude, and for sea and tropics regions due to their overlap. For both the average maximum area and the average lifetime of anvils shown in fig. 3.22 a and b, evidence of this bias is seen as both land and mid-latitude regions display larger areas and longer lifetimes than sea and tropics respectively. On the contrary, similarly to fig. 3.9 d while anvils over land have colder mean and minimum BT than those over the ocean, anvils in the tropics tend to be colder than those in the mid-latitudes. This both indicates that these differences are real, and are as expected as the higher tropopause in the tropics can result in colder anvil clouds. While a difference in the average number of cores is seen between land and ocean, little difference is seen between tropics and mid-latitudes in this regard. The increased tendency of anvils to have multiple cores over land may influence the other differences seen in fig. 3.22, as more organised DCCs tend to have large areas, longer lifetimes and colder CTT. However, these differences also apply to isolated DCCs (those indicated by the black points in fig. 3.22), indicating that they are connected to changes in the convective processes and anvil evolution in the different regions rather than a sampling bias.

3.4.4 Diurnal cycle of anvils

Figure 3.23 shows the average local time of day of anvil detection. Similar to fig. 3.11 a strong land–sea contrast is seen. However, over parts of the Caribbean Sea, the average time of anvil detection is in the afternoon rather than in the morning, which may be linked to DCCs that initiate over land but have anvils which are advected over the ocean. There is also a later average time of detection of anvils over the NGP region, although this shows less of a contrast with surrounding land regions than that of the core initiation times. This indicates that the second peak of core convective activity during the nighttime and early morning in fig. 3.13 is linked to

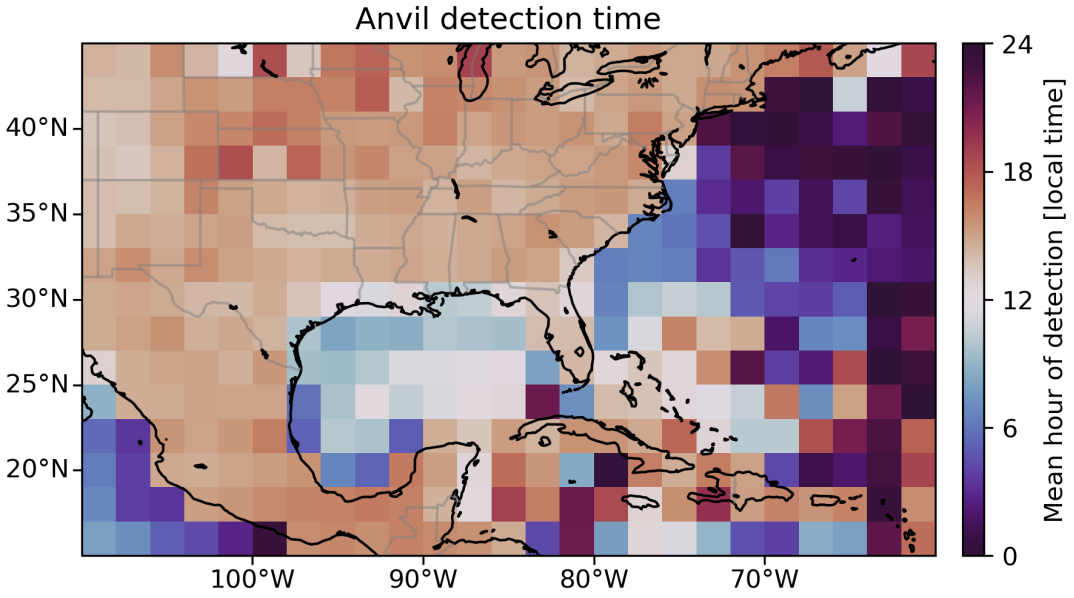


Figure 3.23: The average time of day of detection of anvils observed within each 2×2° grid box, calculated as the circular mean of the local solar time.

long-lived, multiple core systems including MCSs, similar to what was found by Feng et al. [2019].

In fig. 3.24 the diurnal cycles of anvil detections are plotted by region. The distributions generally match those seen of the core detection times shown in fig. 3.12, as the majority of the observed anvils are isolated systems. The diurnal cycle of anvil detections over the sea is much less pronounced however, which may lead to the increased variability seen in fig. 3.23. When comparing the hatched areas, indicating the proportion of the distribution consisting of multi-core anvils, over land, tropics and mid-latitudes the initiation of these organised systems occurs earlier in the day compared to isolated DCCs.

Figure 3.25 shows the average time of observation for anvils. Unlike the previous map of detection time shown in fig. 3.23, in this figure the average of all the time steps at which an anvil is detected is shown. Over the land, the effect is to shift the timing of the maximum to a later point in the diurnal cycle, which is a shift of approximately

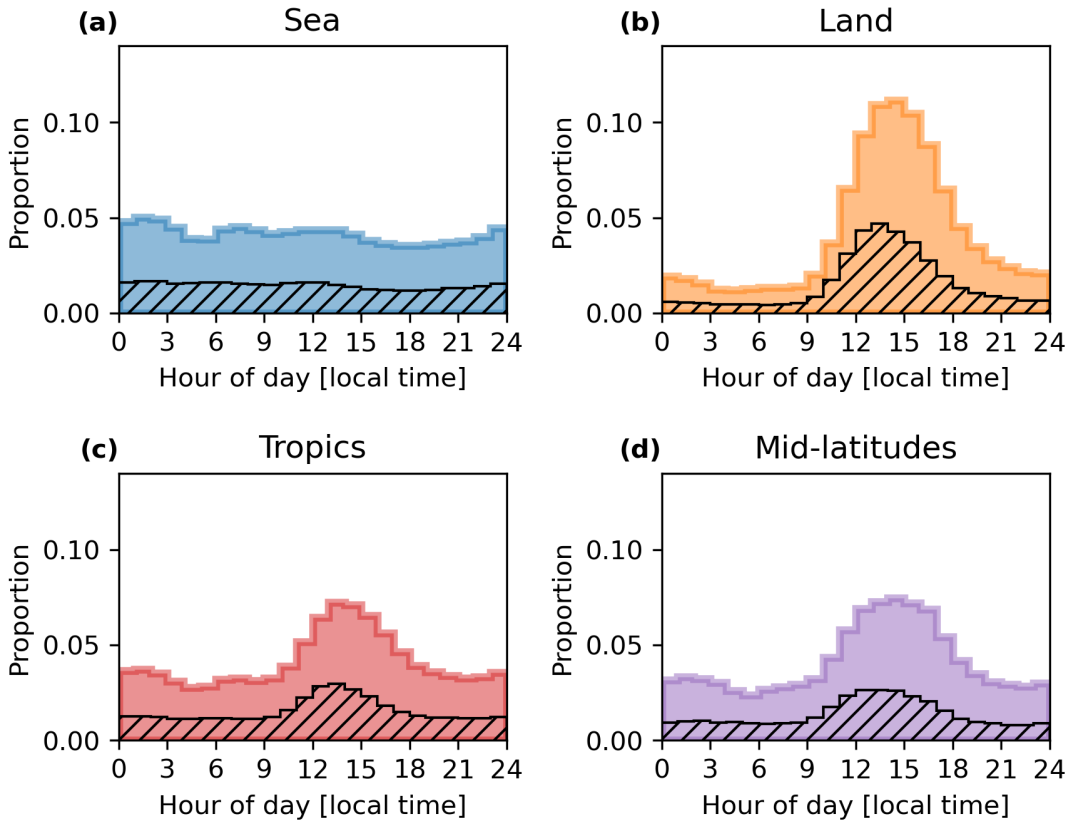


Figure 3.24: The diurnal distributions of the local time of detection for anvils, binned by hour detected over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N). The hatched area shows the proportion of each distribution associated with multi-core anvils.

half the average anvil lifetime. Over sea, however, there is a much larger shift with much of the Caribbean showing a similar time of day to adjacent land regions.

Figure 3.26 shows the diurnal cycle of anvil observations for each of the four regions. For the land, tropics and mid-latitudes regions, there is a shift in the peak of the distribution to 4–5 pm, along with a lengthening of the right tail of the distribution. The distribution of multi-core anvils appears to match that for all DCCs. While in fig. 3.24 the initiation time of these organised anvils tended to be earlier in the day, their longer lifetime means that they tend to exist, on average, at the same times. However, this longer lifetime also means that organised DCCs make

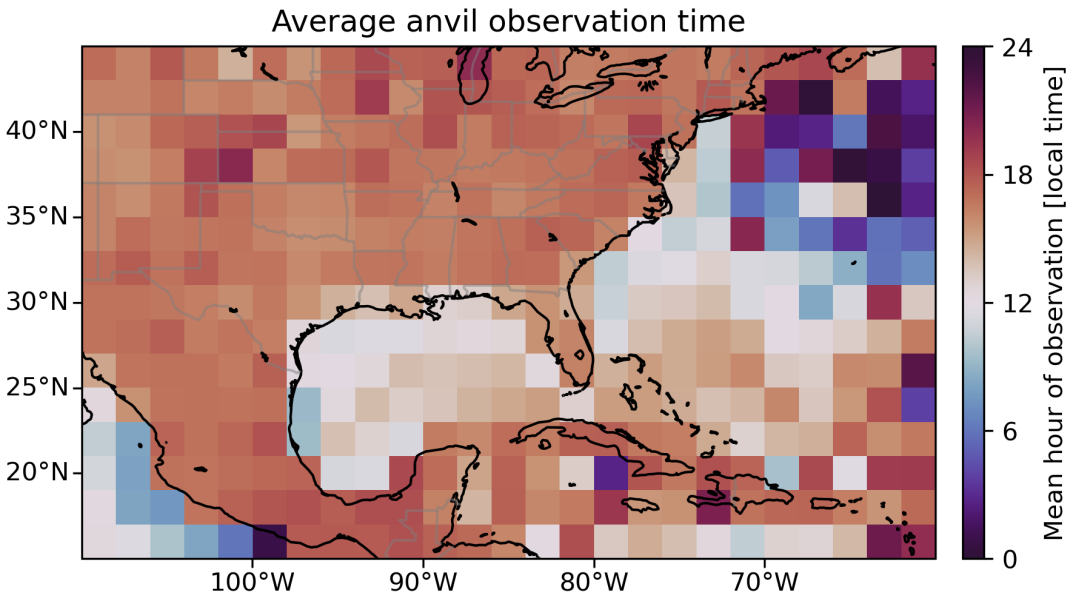


Figure 3.25: The average time of day of observation of anvils observed within each 2×2° grid box, calculated as the circular mean of the local solar time. Unlike fig. 3.23, this is the average of the time of observation at every time step along the anvil lifetime.

up a larger proportion of anvil observations during the nighttime and morning, and fewer during the afternoon during the peak in isolated DCCs. In fig. 3.26 a, the peak of the diurnal distribution over the sea is in the afternoon, with a gradual increase of observations throughout the day. Previous research has shown that solar heating during the daytime increases the area and lifetime of anvils [Gasparini et al., 2022], which may explain why there are more anvils observed during the day, despite more initiations occurring at night.

The difference in the distributions of anvil initiation times and anvil observation times indicates that the lifetime and age of observed anvils change across the diurnal cycle. Figure 3.27 shows the average age of anvils observed during each hour of the diurnal cycle. The increase in the age of anvils observed over sea throughout the daytime provides further evidence for the enhancement of anvil lifetime by SW radiation. For anvil observations over land, there is greater interaction with the

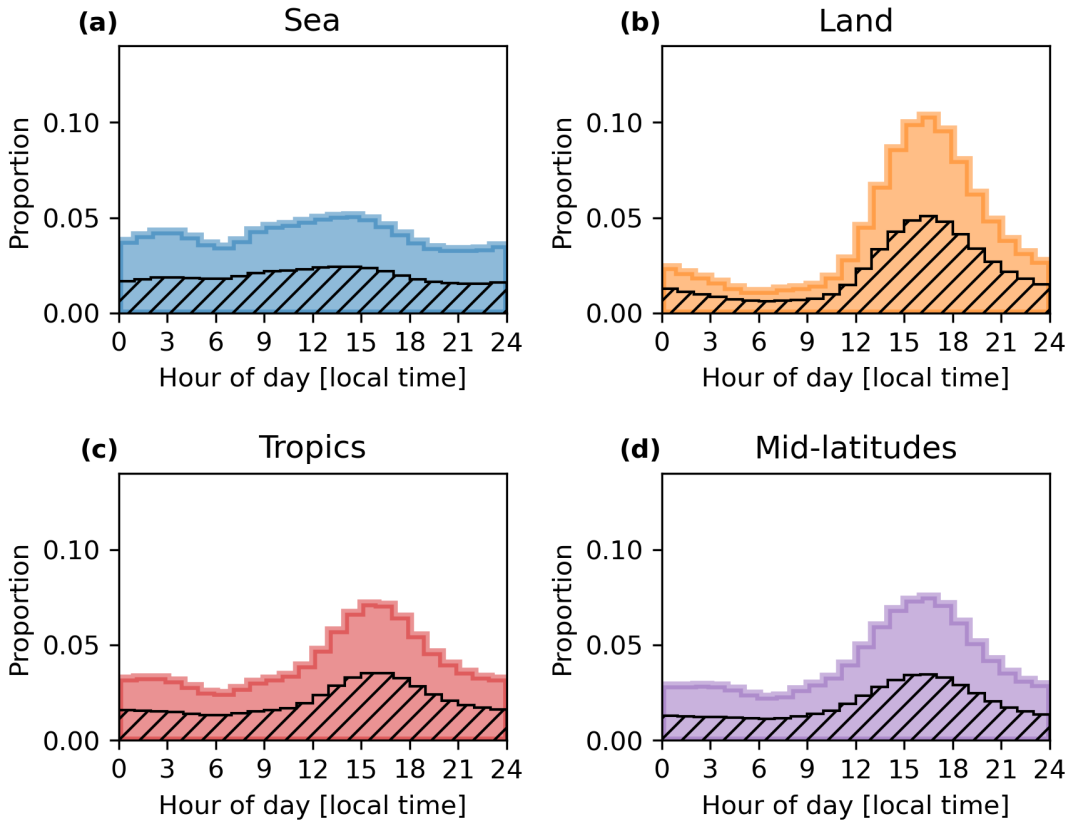


Figure 3.26: The diurnal distributions of the local time of observation for anvils, binned by hour detected over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N). The hatched area shows the proportion of each distribution associated with multi-core anvils.

diurnal cycle of convection. In fig. 3.27 b there is a minima of the anvil age between midday and 1 pm, coinciding with the onset of the afternoon peak of anvil initiations seen in fig. 3.24 b. The age then increases to a maximum at around 10 pm, before decreasing again until midday. The variation in the average age of observed anvils is larger over land than over the ocean. Over the tropics and mid-latitudes however, the combination of opposite signals from sea and land means that there is much less variation of anvil age across either region.

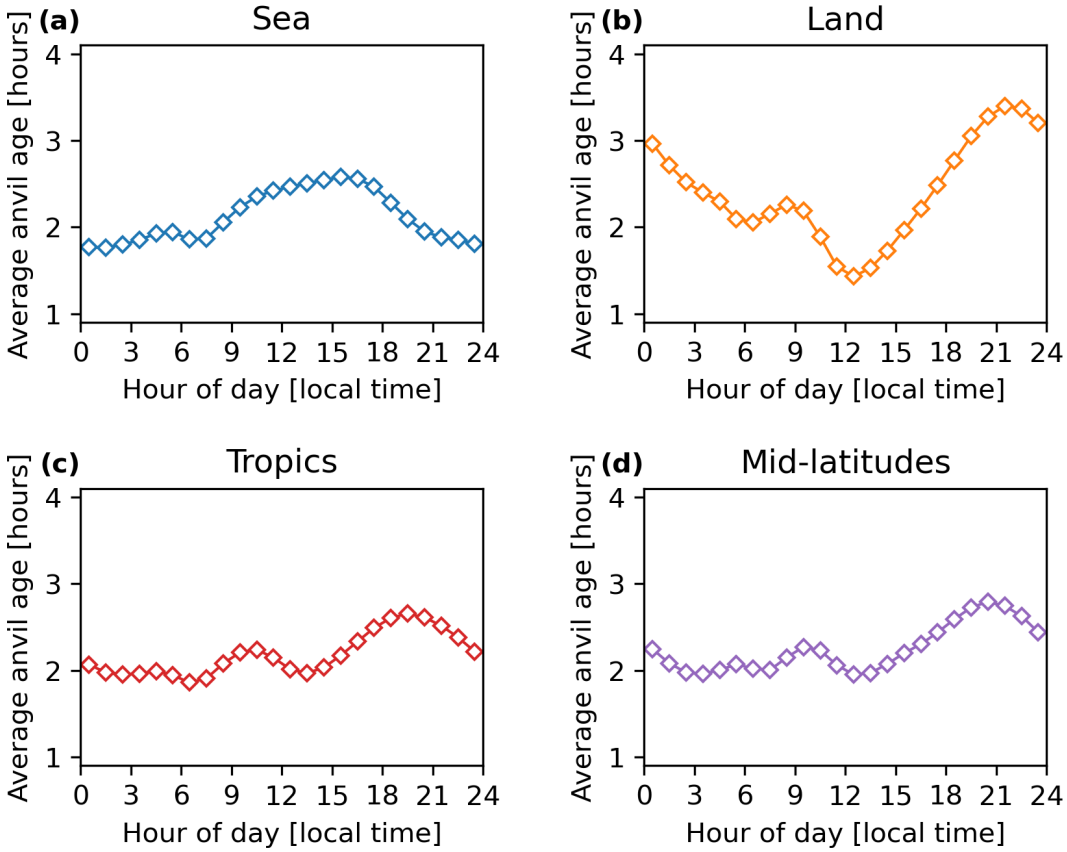


Figure 3.27: The average age of anvils observed over (a) sea, (b) land, (c) tropics (<30 °N) and (d) mid-latitudes (>30 °N) throughout the diurnal cycle. Error bars are not shown as the standard error of the mean is too small to be visible.

3.5 Summary

By applying the *tobac-flow* algorithm to five years of multi-channel BT observations from the GOES-16 ABI CONUS scan region, an extensive dataset of DCC cores and anvils has been produced, including their properties over time. Containing approximately 1.6 million cores and 400 thousand anvils which are considered valid for analysis over their entire lifecycle, this dataset provides valuable information about their properties, distributions and lifecycle in a Lagrangian framework. With five-minute temporal resolution, the properties of developing convective cores can be measured accurately, and then linked to the subsequent evolution of the anvil cloud.

Several key differences in the seasonal, diurnal and regional patterns are seen for DCCs across North America. The spatial distribution of DCC cores depends strongly on the seasonal cycle. There is a large land/sea contrast in the time of initiation, which also displays a sea breeze effect which extends for approximately 200 km inland. The properties of cores themselves, including the lifetime, cooling rate, and time of initiation, have regional dependencies, with more intense convection occurring over land during the daytime and in the tropics.

These differences extend to the anvils produced by these convective cores. The seasonal and diurnal cycles of anvils are closely related to those of the convective cores. Furthermore, there are differences in the properties of anvils between different regions corresponding to differences in convective properties. In particular, DCCs over land have larger areas, longer lifetimes and colder BT than those over the ocean.

While the majority of observed anvil clouds are associated with a single growing core, there are also a large number of multi-core anvils detected in our dataset. The frequency of multi-core anvils as a proportion of all observed anvils increases with latitude. Finally, while the properties of the individual growing cores observed in single- and multi-core anvils show no significant difference, the lifetime and area of anvils show a large increase with the number of cores.

Some of these findings, such as the larger area of DCCs over land, contrast with studies made using polar-orbiting satellites [Ge et al., 2024]. However, our investigation of the age of anvils observed throughout the diurnal cycle provides a mechanism for this discrepancy. The area of DCC anvils evolves as they age [Futyan and Del Genio, 2007]. Impacts of radiative heating and the diurnal cycle act to modify this however [Gasparini et al., 2022]. As a result, the average age of observed anvils varies throughout the day. At the time of the Cloudsat overpass (1:30 pm), the average age of anvils over land is at its minima, while those over the ocean are near their maxima. As a result, the average area of anvils observed at this time over land would

be smaller than those over the ocean, not because of a difference in their processes but because they are being observed earlier in their lifetime before they have reached their maximum area. While at the nighttime overpass time (1:30 am) the average age of anvils over land is older, the reduced number of land DCCs at this time means that any set of snapshot observations will be biased towards younger anvils over land.

By using temporally resolved observations, and placing these in a Lagrangian framework using cloud tracking, these changes in convective cloud behaviour can be accounted for across the diurnal cycle. Furthermore, linking the properties of the developing cores to the resulting anvil clouds over their entire lifetime provides the potential to investigate how the processes of the former affect the latter. Connecting these properties allows investigation of previously uncertain DCC processes, and, in particular, may help shed light on how anvil optical properties, structure and diurnal cycle respond to climate change.

Chapter 4

Contrasting effects of convective intensity and organisation on the structure and lifecycle of DCCs

4.1 Introduction

Understanding how the structure of anvil clouds changes in response to changes in convective behaviour is vital to understanding the radiative impacts and climate feedbacks of DCCs. Around 50% of cirrus clouds in the tropics originate from deep convection [Massie et al., 2002; Luo and Rossow, 2004]. Overall, anvil clouds have a near-neutral radiative impact on the tropics, despite their large SW reflectance and LW absorption [Hartmann et al., 1992; Hartmann, 2016]. Their radiative effect is not homogeneous however; while the optically thick portion of anvil clouds generally have a cooling effect at the ToA, the thin, detrained cirrus has a strong warming effect [Berry and Mace, 2014]. As a result, changes in convective activity that result in a change in the size and structure of anvil clouds may have large impacts on the climate.

The change of anvil cloud area in response to warming—the iris effect [Lindzen et al., 2001; Bony et al., 2016]—is generally considered to have a cooling feedback in response to climate change. However, it is the largest source of uncertainty in cloud–climate feedbacks, with around a third of models predicting a warming response

[Sherwood et al., 2020]. To address this uncertainty, a better understanding of the links between convective processes and anvil properties, changes in anvil structure and the properties of ice clouds are required [Gasparini et al., 2023].

Assessing the extent of thin anvil cirrus is challenging due to the difficulties in observing these clouds using IR radiometers. The low emissivity of thin cirrus clouds means that observed BT is dominated by radiances from the lower atmosphere and surface below. Protopapadaki et al. [2017] used cloud emissivity and cloud top pressure, rather than BT, to detect convective cores, thick and thin anvil cirrus clouds. Data from the AIRS instrument was used to retrieve these emissivity values, as the hyperspectral sounder is more sensitive to thin cirrus than IR radiometers. They then compared the proportion of each anvil cloud consisting of thin anvil to the minimum observed cloud top temperature within the convective core, a proxy for the convective depth which is, in turn, a proxy for convective intensity. It was found that the thin cirrus anvil proportion increased with colder cloud top temperatures, indicating that stronger convection increases the detrainment of thin cirrus.

Takahashi et al. [2017b] built upon this study using collocated measurements from the Cloud Profiling Radar (CPR) aboard CloudSat. The CPR measurements provide the echo top height, measuring the convective depth directly, and also provide another proxy for convective intensity, the echo top distance, which is independent of convective depth [Takahashi and Luo, 2014]. It was again found that the proportion of thin cirrus increased with the echo top height, and also increased with decreasing echo top distance (indicating stronger convection).

It should be noted that both of these studies were performed using observations from polar-orbiting satellites in the A-train, and so do not fully sample the diurnal cycle. While both studies consider the maturity of the DCCs observed, the thin anvil proportion is measured at a single point in time, and so differences in the lifetime of the thick and thin anvil cirrus are not considered. It is expected that the structure

of an anvil will change over its lifetime, with the thin anvil fraction increasing as the anvil dissipates. However, the observed BT of the anvil also changes over its lifetime [Futyan and Del Genio, 2007], and so the use of BT as a proxy for convective intensity is not an independent measurement. While the use of echo top height avoids this confounding issue, these collocated measurements can only be made when the deep convective cloud (DCC) is convectively active and so cannot be used to assess the structure of dissipating anvils.

Chapter 3 presented a large dataset of tracked DCCs, including growing cores, and thick and thin anvils, and showed that this dataset captures relations between convective processes and anvil properties. Using geostationary satellite observations, changes in the structure of DCC anvils can be investigated throughout their lifecycles, along with changes in the lifetime of both thick and thin anvils. Furthermore, the impacts of convective organisation on anvil structure can be investigated by taking into account the number of cores associated with each anvil.

4.2 Data

The dataset of tracked DCCs first presented in chapter 3 is utilised here to investigate how the structure of observed DCCs depends on the properties of their convective cores. Produced using the method detailed in chapter 2, convective cores are tracked during their vertically growing period, and thick and thin anvils are subsequently tracked using two different channel combinations. Changes in the proportion of thick and thin anvil coverage, considering both spatial extent and the lifecycle of the anvil cloud, can be investigated.

Data filtering is applied to detected anvils according to the criteria shown previously in table 3.2. This results in a total of 391,050 anvils considered for analysis, which between them are connected to a total of 792,522 cores. The filtering requires that every anvil in the analysis dataset begins with a valid observed core, ensuring

that the subsequent development of the anvil can be linked to the properties of the core. In addition, it ensures that each anvil is observed for its entire lifetime, and the entire extent of both thick and thin anvil are observed over this period. While this does mean that fewer large anvils are considered valid as their greater extent means that they are more likely to intersect the edges of the domain, it ensures that changes in both thick and thin anvil areas can be considered robust over the entire lifecycle of the analyses anvils.

The dataset contains a range of properties relating to the observed DCCs. This chapter will focus on how the area of the detected cores, thick and thin anvils change with each time step. Two related properties are also analysed; the maximum area of the thick and thin anvil, i.e. the area at the time step when the anvil reaches its maximum observed extent, and the total area which is the sum of the areas observed at each time step over each anvils lifetime. The thin anvil proportion refers to the area of the detected thin anvil divided by the total area of both anvil and core. In general, this is calculated as the total over the entire DCC lifetime, although section 4.4.4 will investigate how this changes over the DCC lifecycle.

There is additional consideration on how the BT of each core and anvil change over time. The area-weighted average of the 10.4 µm BT is taken at each time step for each core, thick and thin anvil. The 10.4 µm LW window BT has a minimal contribution from water in the atmosphere, and so can provide temperature measurements of clouds from near the surface to the top of the atmosphere. It should be noted that the observed BT is affected by the emissivity of the cloud, and so for thin cirrus it will be warmer than the CTT due to the contribution from the atmosphere below. As a result, the mean BT of the thick anvil is used as a proxy for convective cloud height as this avoids a bias from the warmer thin anvil. The (negative) cooling rate of each core is calculated as the largest change in the mean core BT measured over any 15-minute period. This relates to the vertical growth of the DCC core as discussed in

section 2.3.1, with a cooling rate of $1\,\mathrm{K{\cdot}minute^{-1}}$ corresponding to a vertical growth rate at the cloud top of $2.5\,\mathrm{ms^{-1}}$.

4.3 Method

4.3.1 Detection of thick and thin anvils

Observing thin cirrus anvils using satellite BT is difficult due to their low emissivity. Furthermore, BT is affected both by the height of the observed cloud and its optical thickness. As shown in fig. 2.9, the observed BT increases as optical depth decreases. For detection using a fixed threshold, as used in many anvil detection algorithms, this means that the limit of optical depth at which the anvil is detected increases as the height of the cloud decreases. This results in the detected anvil area varying with height, with an anvil at a lower altitude being detected as smaller than one at a higher altitude even if their size and structure are otherwise identical. This problem with detected anvil area using BT thresholds was identified as early as Augustine and Howard [1988], but remains a feature of many detection algorithms to this day (see e.g. section 1.4).

While observing the thin cirrus anvil is challenging, section 2.4.4 showed that thin anvil cirrus can be detected using a combination of the WVD and SWD differences from ABI. Furthermore, the use of channel differences reduces the height dependence of the detection of anvils compared to the use of a single BT. Figure 4.1 shows simulated BT observations using the libRadTran model [Emde et al., 2016], as described in section 2.3. The top panel shows simulation BT observations by ABI at 10.4 µm at heights of 10, 12 and 14 km (shown by the solid, dashed and dotted lines respectively). Also shown by the grey, dashed line is an example of the 241 K BT threshold. At this threshold there is substantial variance in the limit of OD at which anvil clouds are detected, which ranges from 1 at 14 km to 3 at km. This provides clear support

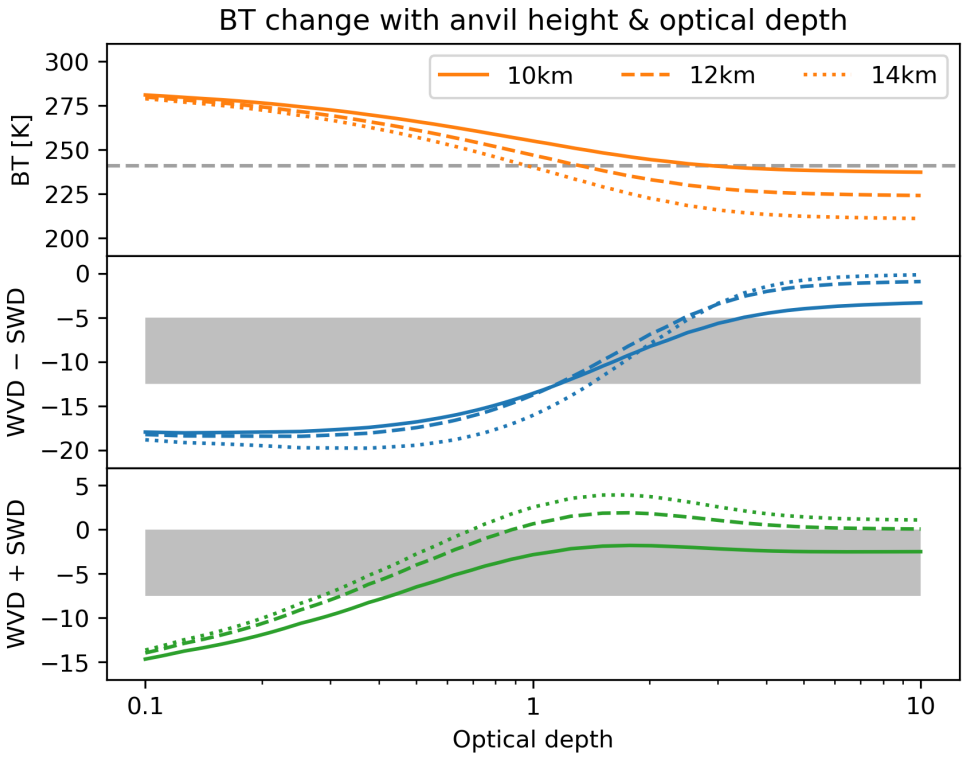


Figure 4.1: Simulated ABI BT observations for DCC anvils at 10 km (solid lines), 12 km (dashed lines) and 14 km (dotted lines) at optical depths between 0.1 and 10. Panels show (top) 10.4 µm BT; (middle) WVD minus SWD, used in thick anvil detection; and (bottom) WVD plus SWD, used in thin anvil detection. The grey dashed line shows the 241 [K] threshold, and the grey shaded regions show the range in which cloud edge detection is performed for thick and thin anvils.

for the argument made by Augustine and Howard [1988] that the anvil area detected using this threshold is "subjective".

The middle and bottom panels in fig. 4.1 show the simulated BT for the combinations of WVD minus SWD and WVD plus SWD respectively across a range of OD. Section 2.4.4 describes how these two channel combinations are used for the detection of thick and thin anvils. In both panels, the grey shaded region shows the BT range in which edge detection is used to find the anvil area. Compared to 10.4 µm BT, the WVD minus SWD combination shows very little variance in the optical depths at which the anvil cloud is detected within the threshold range. Throughout the

threshold range of –5 to –12.5 K the OD corresponding to the same observed BT difference varies by approximately 0.1–0.2 between differing heights, compared to 1–2 for a simple threshold approach on the 10.4 µm BT.

The WVD plus SWD combination is shown in the bottom panel of fig. 4.1. Greater variance between different heights is seen, particularly for the 10 km anvil simulation. However, the OD corresponding to the observed BT differences are very similar throughout the detection range of 0 to –7.5 K for heights of 12 and 14 km. Lidar observations of thin cirrus anvil have shown that the heights of these clouds typically exceed 12 km [Wall et al., 2020; Horner and Gryspeerdt, 2023], and so little difference is expected in the sensitivity to observed anvils.

From fig. 4.1 the limits of detection for thick anvils can be estimated as 2±0.5 OD, and 0.5±0.2 OD for the thin anvil. This means that the thinnest anvil cirrus remains undetected [Berry and Mace, 2014] using this method. This may lead to biases in the calculation of the thin anvil fraction, as while all anvil thicker than the thick anvil threshold will be detected, not all the thin below the threshold will be. This could lead to an underestimate of the thin anvil area and fraction in anvils where the OD distribution shows a large frequency of very thin cirrus. Such thin anvil cirrus will have a small radiative effect on the climate, and, therefore, this bias in the thin anvil fraction may not result in a large uncertainty in the climate impacts due to the low contribution to anvil CRE from the unobserved cirrus.

A limitation of the simple separation of thick and thin anvil is that uniform changes in anvil thickness cannot be distinguished from changes in the anvil thickness distribution. An anvil that is uniformly thicker will result in a larger thick anvil area and a smaller thin anvil fraction. For a straightforward analysis of the impact of changes in anvil structure on the climate this is not a significant difference, as anvil CRE is most sensitive to changes in thickness at OD around 2 [Berry and Mace, 2014], and so the cutoff value for thick anvil is ideal. More detailed analysis of

the mechanisms through which anvil structure changes may, however, require more information about the overall distribution of anvil thickness that is not provided by these metrics. Use of retrieved cloud properties, or colocation with cloud radar instruments that can resolve the vertical structure of anvil clouds, may be important for further understanding of processes affecting anvil thickness.

It should be noted that the simulations shown in fig. 4.1 were calculated using a standard tropical atmospheric profile to best represent the moist environment in which deep convection occurs. While anvils occurring in the extra-tropical parts of the CONUS domain are expected to have lower cloud top heights, there is also expected to be a colder temperature profile. As a result, the observed BT of extra-tropical DCCs behave more like that of the 12 and 14 km simulations shown in fig. 4.1, rather than that at 10 km.

4.3.2 Estimation of convective intensity and organisation

As the vertical velocity or convective mass flux of DCC updrafts cannot be observed directly using geostationary satellite observations, a proxy is needed to estimate their intensity instead. The minimum BT of an observed DCC is commonly used as a proxy for convective intensity, and is used by Protopapadaki et al. [2017] for this purpose. This property is more closely related to the height of the DCC, however. While the height of a DCC is related to the intensity of convection, it is also related to the tropopause height and level of neutral buoyancy which may vary with location and meteorology independent of convective processes. While the minimum BT of observed anvils in this chapter to better compare to previous studies, a more direct proxy for convective intensity is desired.

The primary proxy for convective intensity in this chapter is the maximum cooling rate of the initial observed core for each DCC. As described in section 2.3.1, the core cooling rate is linked to the vertical development of the core over time. A cooling

rate of 1 Kminute^{-1} approximately represents a cloud top vertical velocity of 2.5 ms^{-1}. While this should not be confused with the updraft velocity of the core itself, and is expected to be less due to the effects of entrainment and overturning circulation at the top of the updraft, it is linked to the updraft velocity directly. In section 3.4.1 it was shown that the core cooling rate and minimum core BT are correlated. The maximum cooling rate of an anvil is measured on the initiating core to ensure that there are sufficient observations of the core throughout its growth that are not masked by an anvil above. While the updraft may strengthen later in the development of the DCC, or subsequent cores in a multi-core system have stronger updrafts, this is considered the most direct and robust proxy available for the updraft intensities of DCCs observed by geostationary satellites.

The number of cores associated with each DCC is used as a measure of its organisation, again providing the most direct proxy available. While it may not be possible to observe all the cores associated with larger MCSs due to the anvil blocking observations of cores below, it is still expected that many separate cores associated with such systems are detected. As shown in chapter 3, tracked DCCs with many cores show the properties expected of MCSs.

4.3.3 Lifecycle analysis

To analyse changes in the lifecycle of DCCs, their lifetime is categorised based on the method developed by Futyan and Del Genio [2007]. This method separates the anvil cloud into growing, maturing and dissipating stages based on observations of BT and anvil area. Futyan and Del Genio's method defines the end of the growing phase as the time at which the anvil reaches its minimum observed BT, and the end of the mature phase as the time at which the anvil reaches the maximum observed area. While simple, this approach applies to a wide range of observed DCCs including both isolated and organised convection. Furthermore, by determining the anvil lifecycle

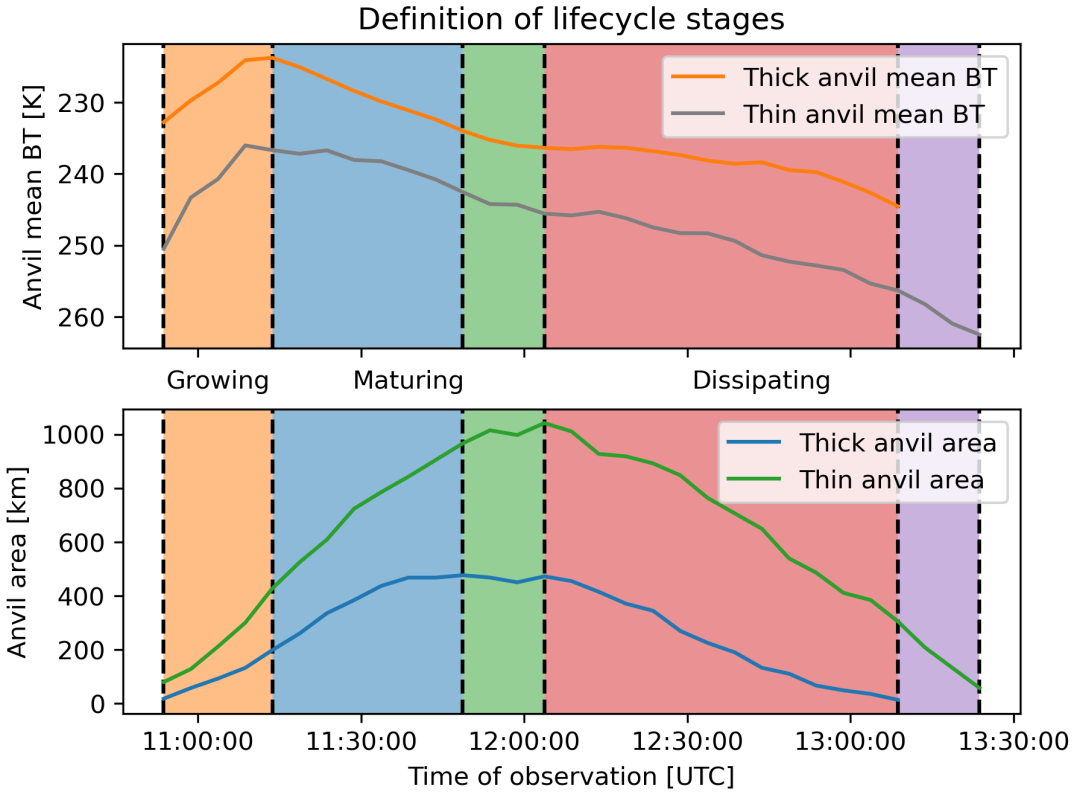


Figure 4.2: Definition of anvil lifecycle stages based on the method of Futyan and Del Genio [2007]. The top panel shows the evolution of the mean BT of the thick and thin anvil regions over the DCC lifecycle, and the bottom panel shows the evolution of the anvil area. The growing phase (orange) is defined as the time between initial detection and the minimum observed BT. The mature phase for the thick (blue) and thin (green) anvil is defined as the time between the minimum BT and maximum thick and thin anvil area respectively. The dissipating phase for the thick (red) and thin (purple) anvils is defined as the time between the anvil maximum area and the final time of detection of the thick and thin anvil respectively.

only in terms of the observed anvil properties, the properties of the detected cores can be treated as independent variables allowing straightforward analysis of their impacts on DCC lifecycle.

Futyan and Del Genio's method is extended in this chapter through consideration of the development of both thick and thin anvil, which results in the lifecycle stages shown in fig. 4.2. The growing stage (orange) is defined as the time from the initial detection of the DCC to the time at which the thick anvil has its coldest average BT. The thick anvil maturation stage (blue) occurs between this time and the time at which the thick anvil reaches its maximum area, and the thin anvil maturing stage (green) occurs until the thin anvil reaches its maximum area. In this stage, the thin anvil continues to expand while the thick anvil begins to shrink. The thick anvil dissipating stage (red) ends at the final detection time of the thick anvil, while the thin anvil dissipating stage (purple) ends at the final detection time of the thin anvil, which is also the final detection time of the tracked DCC. By increasing the granularity of the lifecycle definition in this manner, the impact of convective processes on DCC lifecycle and structure may also be investigated.

4.3.4 Composite analysis

To compare how the structure of tracked DCC evolves over their lifecycle, the area of their cores, thick and thin anvils is interpolated over a normalised lifetime. In this normalised lifetime, 0 is defined as the time of initial core detection, and 1 is the time of the final anvil detection. By normalising their lifetimes in this manner, the change over time of DCCs with different lifetimes can be composited and compared.

An example composite of the average properties of all analysed anvils in this chapter is shown in fig. 4.3. Figure 4.3 a shows the average change in the area of the core (red), thick anvil (orange) and thin anvil (blue) over the normalised lifetime of the DCCs. Figure 4.3 b shows how these areas change as a proportion of the net DCC

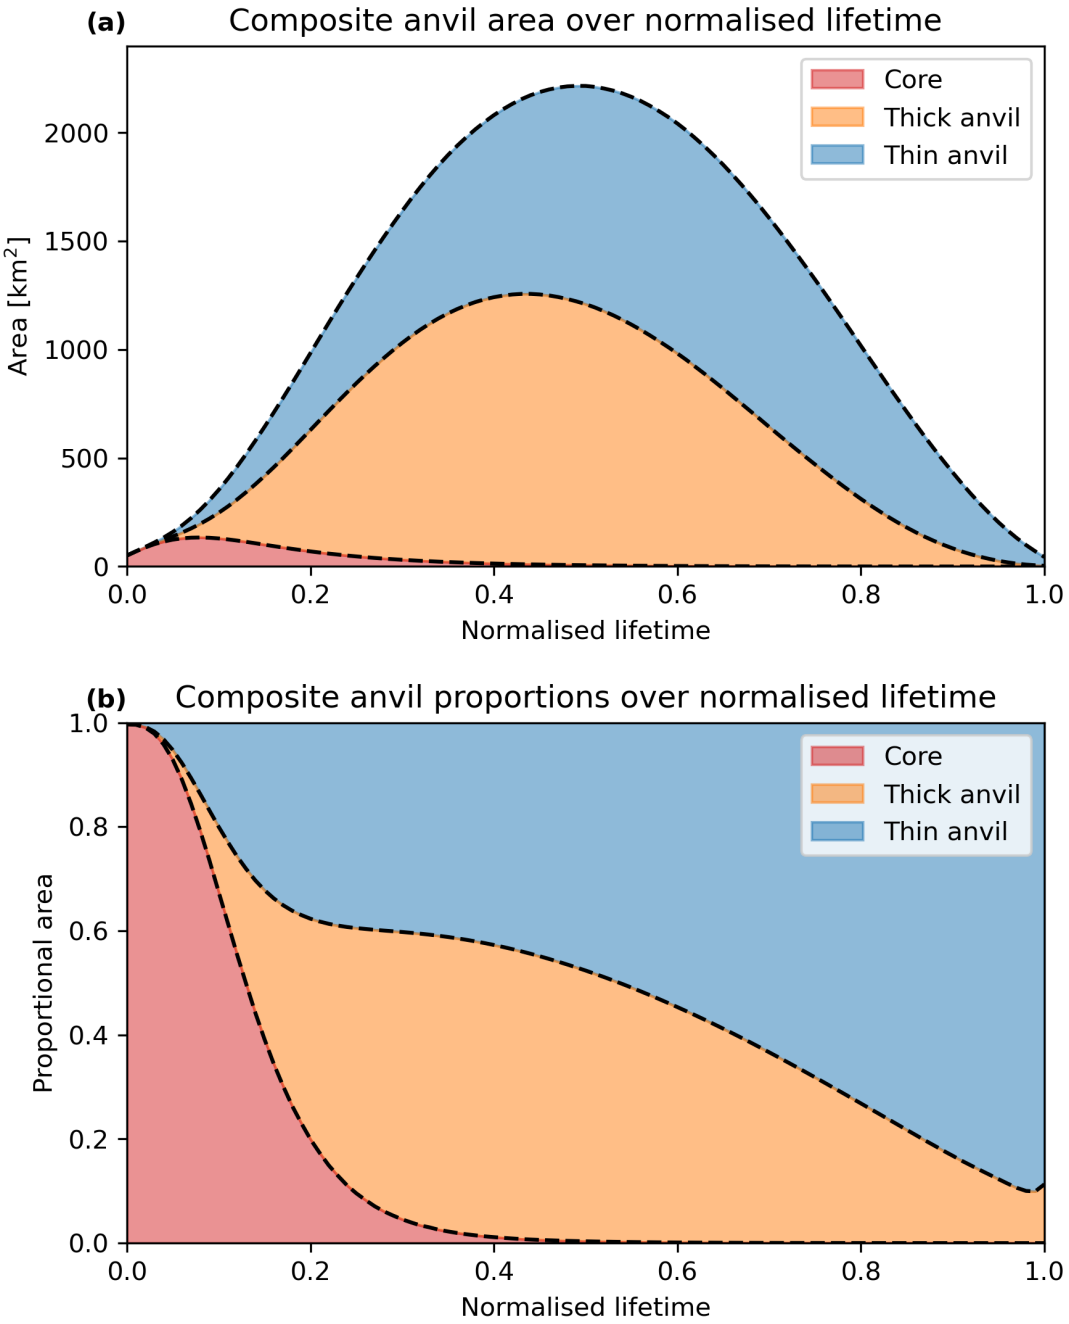


Figure 4.3: The average evolution of composites of cores (red), thick anvils (orange) and thin anvils (blue) over their normalised lifetimes, where 0 is the initial core detection time and 1 is the final anvil detection time. Panel (a) shows the average area of anvils, and (b) shows the proportion of the total area corresponding to each part of the anvil structure.

area. This shows how, on average, the structure of DCCs evolves over their lifetime. Initially, the DCC consists entirely of a growing core, with the initial anvil formation occurring between 10% and 20% of the DCC lifetime. The initial expansion of the anvil is mostly thin, likely because the majority of the mass flux is still travelling upwards, and the area of thick anvil is covered by the developing core. Between 20% and 40% of the normalised lifetime, while the anvil is maturing, the thick and thin anvil expand at similar rates and so the proportion of thin anvil remains consistent. Beyond this point, the thick anvil reaches its maximum area before the thin anvil and dissipates faster, and the fraction of the thin anvil increases until the end of the DCC lifecycle. Note that the average thin anvil fraction does not reach 100% at the end of the normalised lifetime due to the presence of some DCCs where the dissipation of the thick and thin anvil are recorded at the same time.

Through the use of composite analysis, the change in anvil structure throughout the DCC lifecycle can be investigated. This allows the study of responses to prior observed convective processes, rather than limiting our analysis to bulk properties or snapshots.

4.4 Results

4.4.1 Change in thin anvil fraction with intensity and organisation

To begin, the net thin anvil fraction of observed anvils is compared with the mean anvil BT, initial core maximum cooling rate and number of cores. Figure 4.4 a shows how the thin anvil proportion changes with the average thick anvil BT. Note that only the thick anvil BT is considered to avoid a dependence between the anvil BT and the thin anvil area. There is an increase of thin anvil proportion with the average anvil BT, agreeing with the results of Protopapadaki et al. [2017]. Figure 4.4 b compares how the thin anvil fraction changes with the maximum cooling rate of the initial core

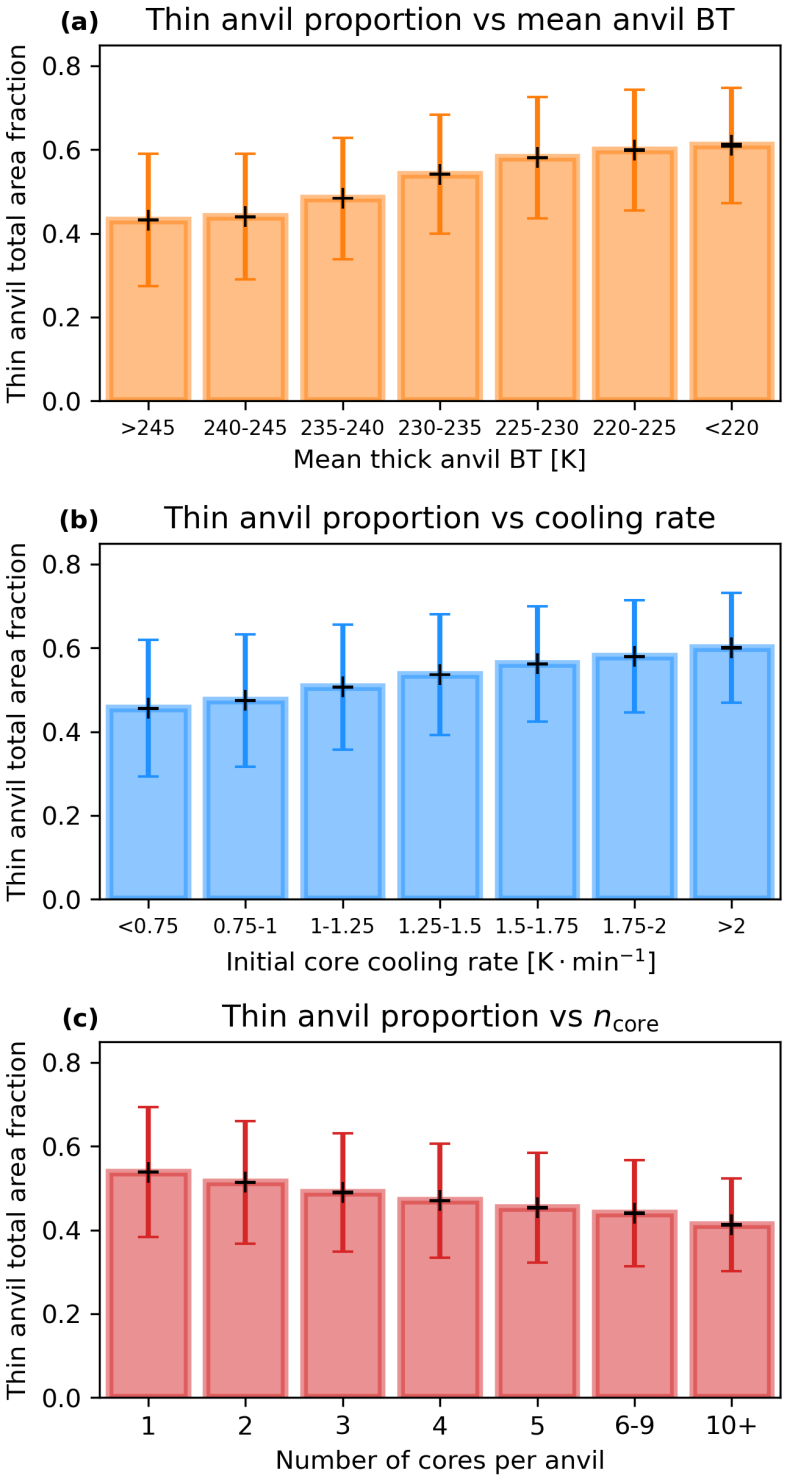


Figure 4.4: Thin anvil proportion of anvils categorised by (a) mean anvil BT, (b) initial core cooling rate, and (c) number of cores. Black error bars show the standard error of the mean, and coloured error bars show the standard deviation.

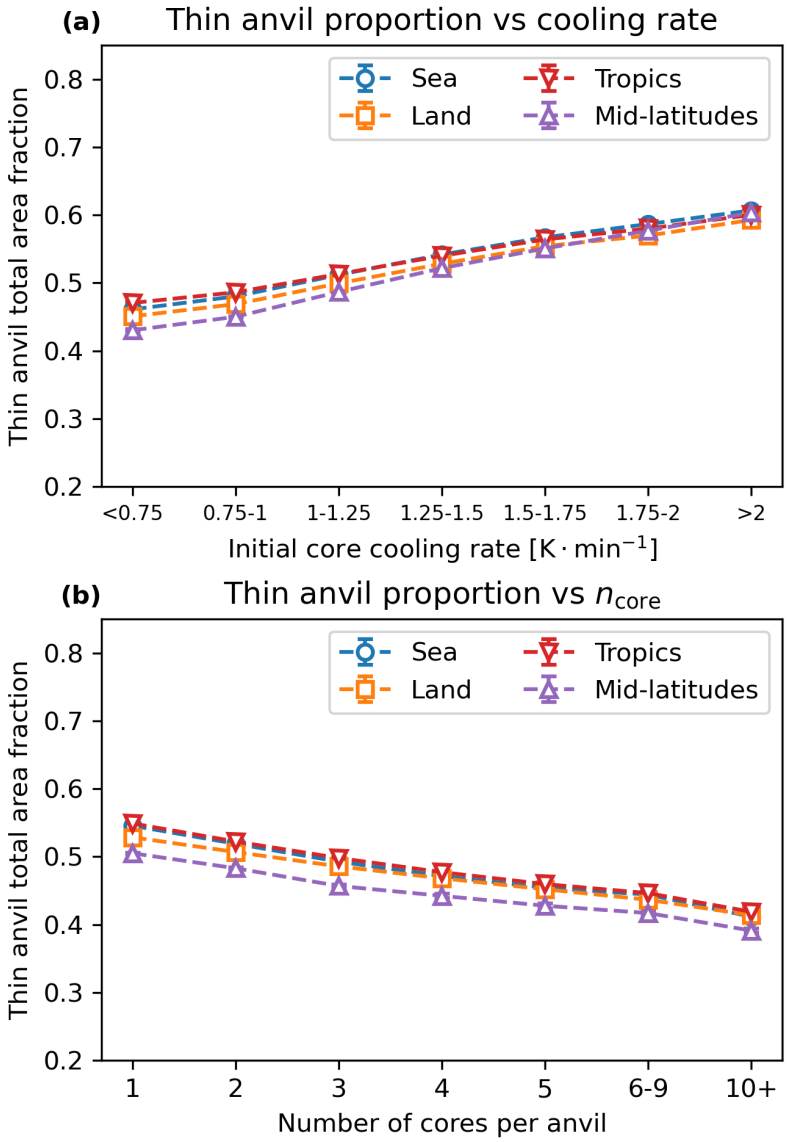


Figure 4.5: A comparison of the change in thin anvil proportion with (a) maximum core cooling rate and (b) number of cores for four different regions: sea (blue diamond markers), land (orange square markers), tropics (<30 °N; red downward triangle markers) and mid-latitudes (>30 °N; purple upward triangle markers). Error bars show the standard error of the mean.

associated with the anvil, and again shows a positive relation between the strength of convection and the thin anvil fraction. While the value of the thin anvil fraction differs from that observed in previous studies, this is expected as both studies use different methods to define thick and thin. In addition, in this chapter, the proportion is measured over the entire DCC lifetime rather than at a single point in time.

In fig. 4.4 c, the thin anvil fraction is compared to the number of cores associated with each anvil. In contrast to fig. 4.4 b and c, there is a decrease in the thin anvil fraction with an increase in the number of cores. As this occurs despite the tendency for multiple-core anvils to have colder average BT, this indicates that further factors are involved in the thin anvil fraction than the temperature and height of the anvil.

To ensure that these relationships are not caused by differences in the properties and distributions of DCCs across different regions, the relationship between cooling rate and number of cores and the thin anvil fraction is further separated into four different regions, shown in fig. 4.5. As with previous studies, the fraction of thin anvil is very similar for DCCs observed over land and ocean. While there is a difference in the thin anvil fraction of approximately 0.05 between DCCs located in the tropics and mid-latitudes (defined as $<30\,^{\circ}$N and $>30\,^{\circ}$N respectively), this difference decreases with increasing convective intensity. Overall, the near identical relationship of thin anvil fraction to intensity and organisation across multiple regions indicates that this is linked to the convective processes themselves, rather than a co-variance with local conditions.

4.4.2 Impact of intensity and organisation on DCC properties

Both the intensity and organisation of DCC cores have impacts on the properties of their anvil clouds. This section will investigate how these processes affect the bulk properties of these anvils; the anvil maximum area and the anvil lifetime for both thick and thin anvils, and the anvil BT. While comparing the bulk properties of these

anvils indicates how their structure changes, it does not fully explain the differences shown in figs. 4.4 and 4.5.

Figure 4.6 shows how the area, lifetime and BT of DCCs change with the core cooling rate. Figure 4.6 a shows that the average thick and thin anvil maximum areas vary very little with cooling rate, except for at the highest observed values of cooling rate. There is however a slight widening in the gap between the thick and thin anvil maximum areas with increasing core cooling rate, as the thin anvil increases in maximum area while the thick anvil remains consistent. Figure 4.6 b shows that while the average thick anvil lifetime remains constant with increasing core cooling rate, the thin anvil lifetime increases. There is also a consistent decrease in both average and minimum thick anvil BT with cooling rate, shown in fig. 4.6 c. This change in anvil BT with the core intensity agrees with the results shown previously in fig. 3.10 b which showed that more intense cores have colder minimum BT.

In fig. 4.7 the change of anvil properties with the number of cores is considered. Figure 4.7 a shows that both the thick and thin anvil maximum area increase substantially with an increasing number of cores (note that it is plotted against a log scale). However, the thin anvil area increases at a slower rate proportional to the thick anvil. Similarly, there is an increase in both the thick and thin anvil lifetime with increasing number of cores shown in fig. 4.7 b. There is a similar rate of the decrease of average anvil BT with the number of cores in fig. 4.7 c as shown in relation to core cooling rate in fig. 4.6 c. The minimum anvil BT decreases at a notably faster rate with the increasing number of cores, however. Although this is indicative of an increased likelihood of overshooting cores in more organised DCCs, it must also be noted that the larger area of these systems also makes it more likely to observe a colder minimum pixel value.

The changes in anvil properties with intensity and convection are intriguing due to both their differences and similarities. Both processes result in anvils with colder BT,

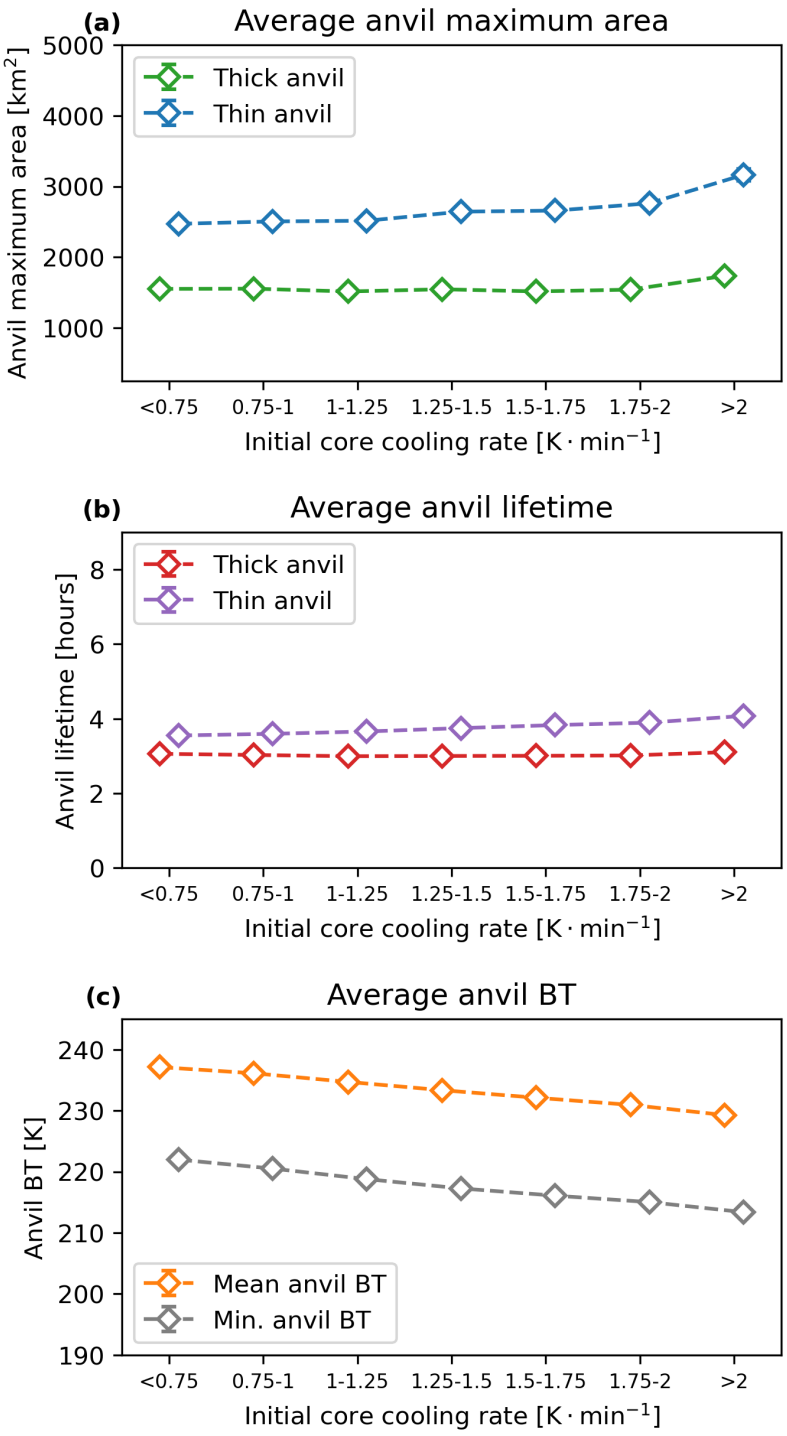


Figure 4.6: The effect of initial core cooling rate on (a) maximum anvil area, (b) anvil lifetime, and (c) mean and minimum BT. Error bars show the standard error of the mean. Points have been staggered to show the thick and thin anvil properties more clearly, but correspond to the same tick marks on the x axis.

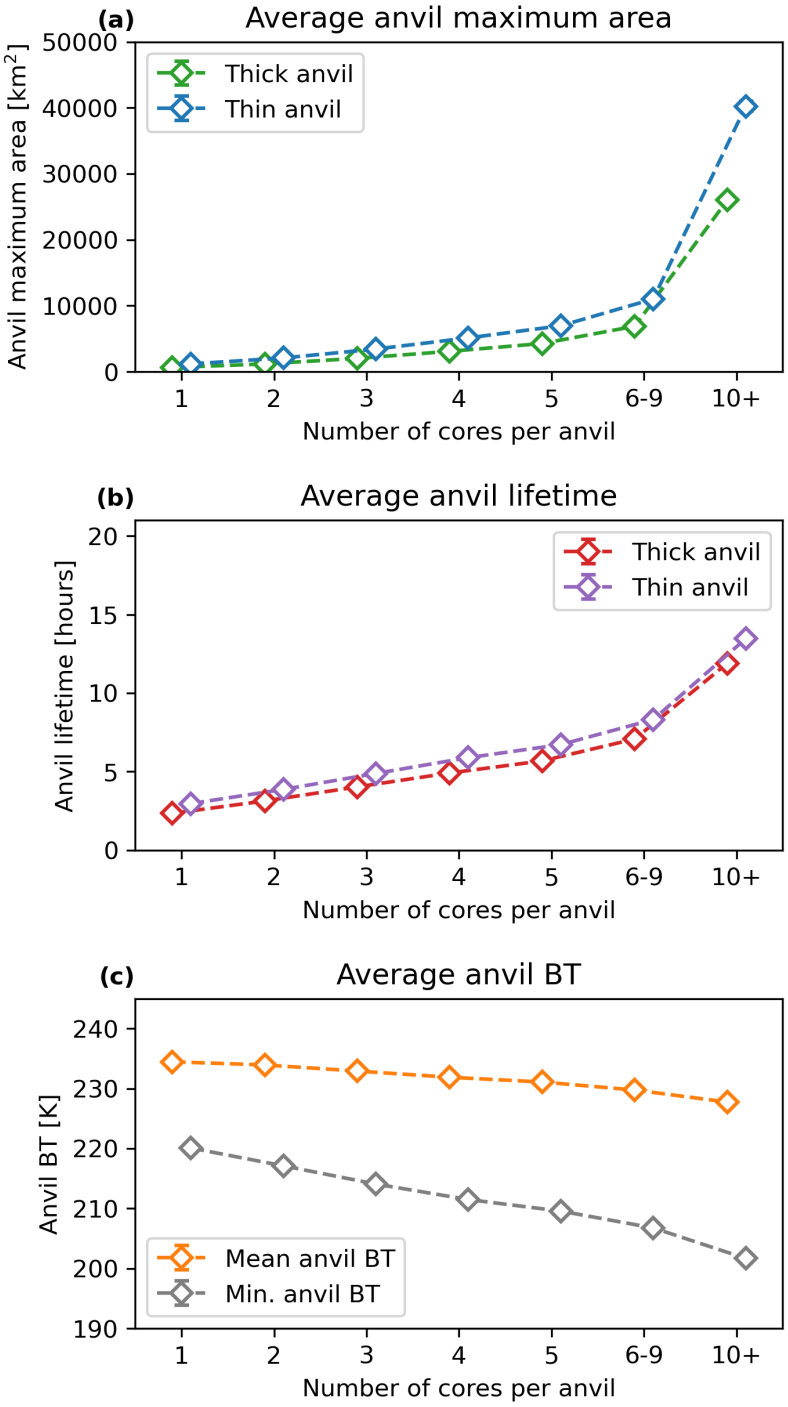


Figure 4.7: The effect of the number of cores on (a) maximum anvil area, (b) anvil lifetime, and (c) mean and minimum BT. Error bars show the standard error of the mean. Points have been staggered to show the thick and thin anvil properties more clearly, but correspond to the same tick marks on the x axis.

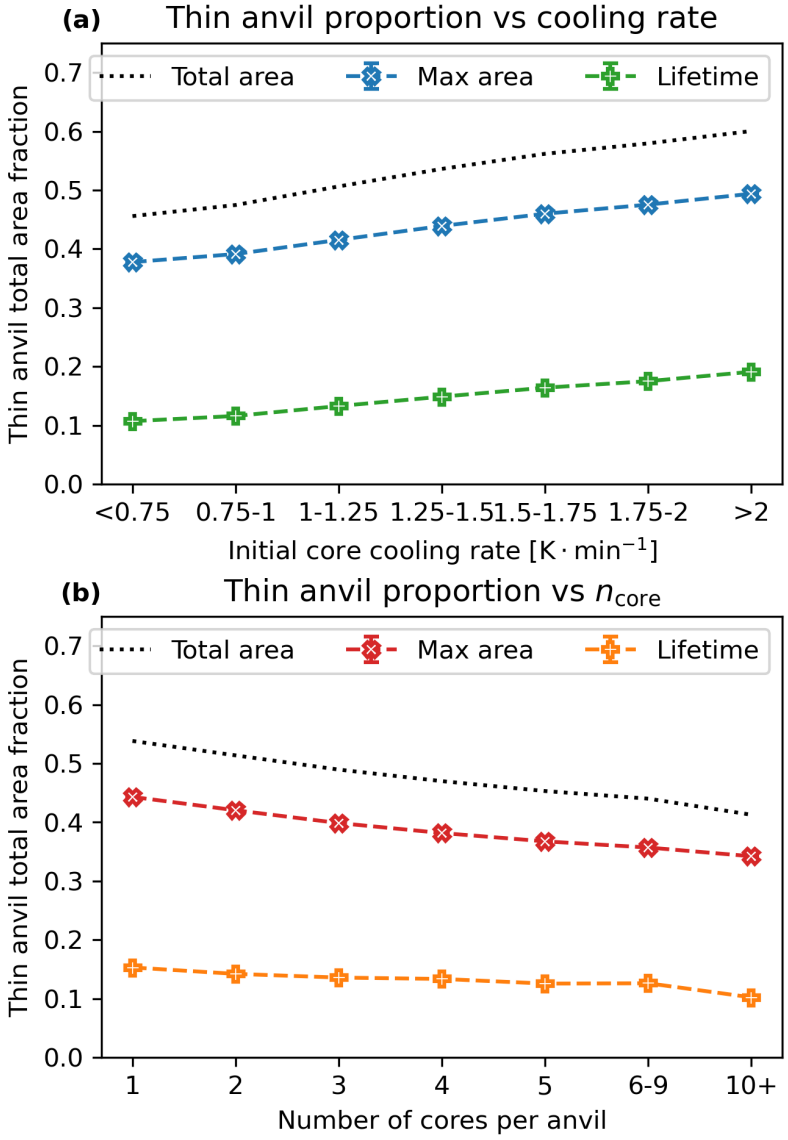


Figure 4.8: A comparison of the proportion of thin anvil measured at the time of maximum anvil area, and the proportional difference in lifetime, compared to the overall thin anvil proportion. Panel (a) shows the change in thin anvil proportion with core cooling rate, and (b) shows the changes with the number of cores. In both panels, the black dotted line shows the total thin anvil proportion over the entire DCC lifecycle. Error bars show the standard error of the mean of the thin anvil fraction.

which is commonly used as a proxy for convective intensity. However, the differences in the changes in anvil areas and lifetimes show that these processes cannot be lumped together so simply.

In fig. 4.8, the proportion of observed thin anvil fraction that can be attributed to the bulk changes in anvil maximum area and lifetime is shown. The proportional difference in the maximum area of the thick and thin anvils represents approximately 80% of the change in thin anvil fraction across both convective intensity and organisation. This shows that while the trend in thin anvil fraction can be captured by snapshot observations from polar-orbiting satellites, they will underestimate the total thin anvil fraction over the entire DCC lifetime. The proportional difference in the thick and thin anvil lifetimes shows a smaller contribution to the net thin anvil fraction. Furthermore, there is a stronger (positive) trend of the lifetime proportion with intensity and a weaker (negative) trend with organisation. Note that while some of the error bars (showing the standard deviation) extend below zero, all values for the thin anvil fraction are positive.

4.4.3 Changes in DCC lifecycle with intensity and organisation

It is clear from the impact on anvil lifetime that both the intensity and organisation of convection affect DCC lifecycle. Furthermore, fig. 4.8 shows that these different processes have different impacts on the lifecycle that warrant further investigation. This section will focus on the effect of both convective intensity and organisation on the different lifecycle stages of DCCs. The modified Futyan and Del Genio method explained in section 4.3.3 is used throughout to separate the anvil lifecycle into growing, mature and dissipating stages.

Figure 4.9 shows the lifetime of the different lifecycle stages for DCCs with increasing core cooling rates. The lifetime of the growing stage decreases from 1 hour to 45 minutes for the most intense convection. While these growing lifetimes are longer

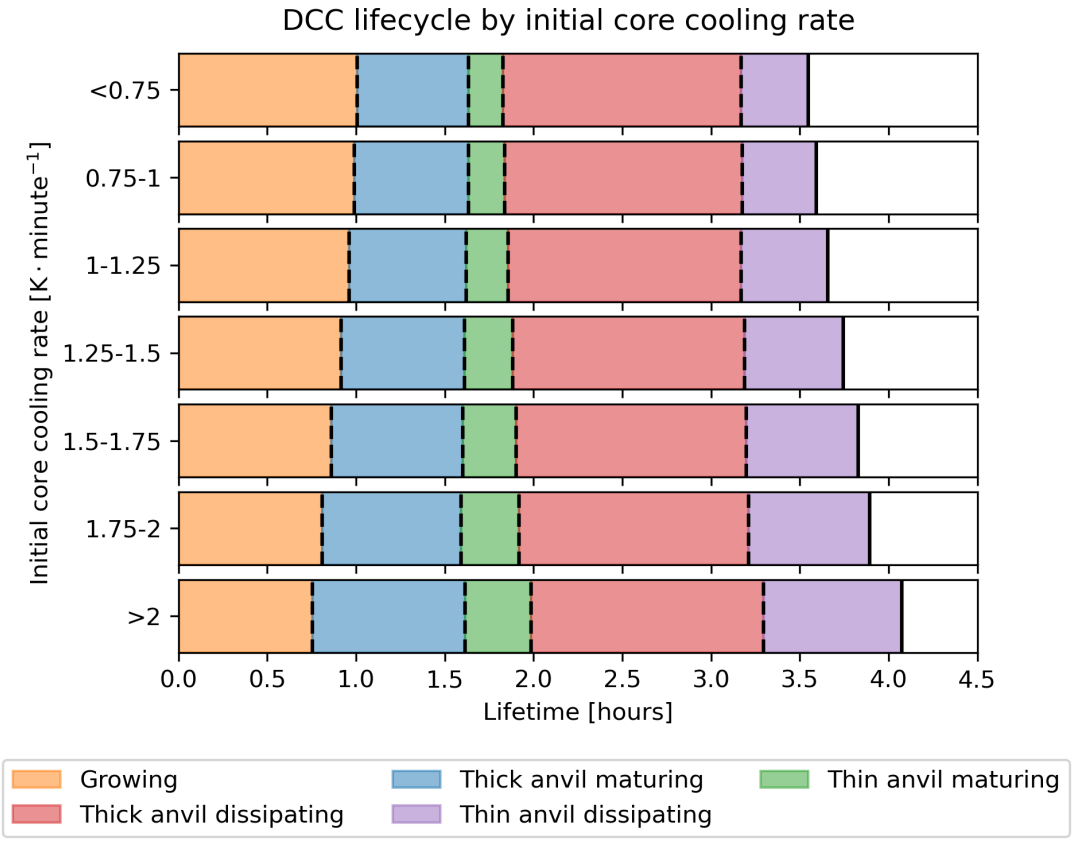


Figure 4.9: The lifetime of each of the lifecycle stages, as defined in section 4.3.3, for anvils categorised by their initial core cooling rate.

than those of the detected cores shown in fig. 3.10, the trend agrees with the shorter lifetime for more intense cores. Despite the colder BT and hence higher cloud top heights exhibited by these DCCs, the stronger vertical growth rates result in an anvil which reaches its minimum BT in a shorter time frame.

The timing of the maximum extent of the thick anvil remains consistent at just over 1.5 hours for all cooling rates, while the time of the maximum thin anvil area increases from 1 hour 45 minutes to 2 hours for the most intense DCCs. Similarly, the timing of the dissipation of the thick remains consistent at just over 3 hours for most cooling rates, while the thin anvil dissipation time increases from 3.5 hours to 4 hours for more intense DCCs. These findings reinforce what was shown in fig. 4.6, that the core cooling rate affects the lifetime of the thin anvil but not the thick anvil.

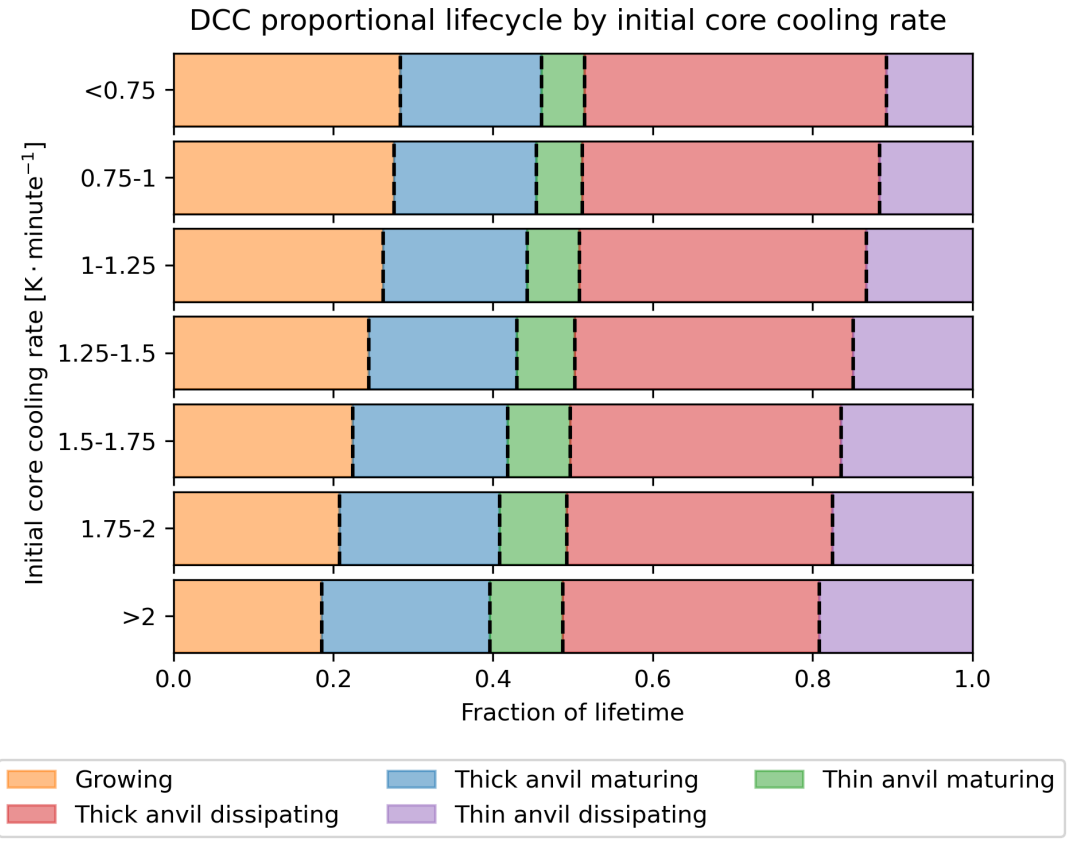


Figure 4.10: The average fraction of anvil lifetime of each of the lifecycle stages defined in section 4.3.3 for anvils categorised by initial core cooling rate. Each bar is normalised such that the total length of the row is equal to the mean thin anvil lifetime of its cooling rate bin.

Figure 4.10 displays the average proportion of anvil lifetime spent in each of the stages shown in fig. 4.9. The anvil lifetime is normalised using the method described in section 4.3.4, where 0 is the time of the initial detection of the DCC and 1 is the time of the final detection of the dissipating thin anvil. By analysing these lifecycle stages as a proportion of the total lifetime, better insight is gained into how changes in lifecycle impact the structure and bulk properties of observed DCCs.

The reduction in the lifetime of the growing phase and increase for thin anvil dissipation for more intense DCCs is again apparent. The time between these two stages, representing the period in which the thick anvil matures and then dissipates, lasts for around 60% of the total anvil lifecycle for all core cooling rates. The shorter

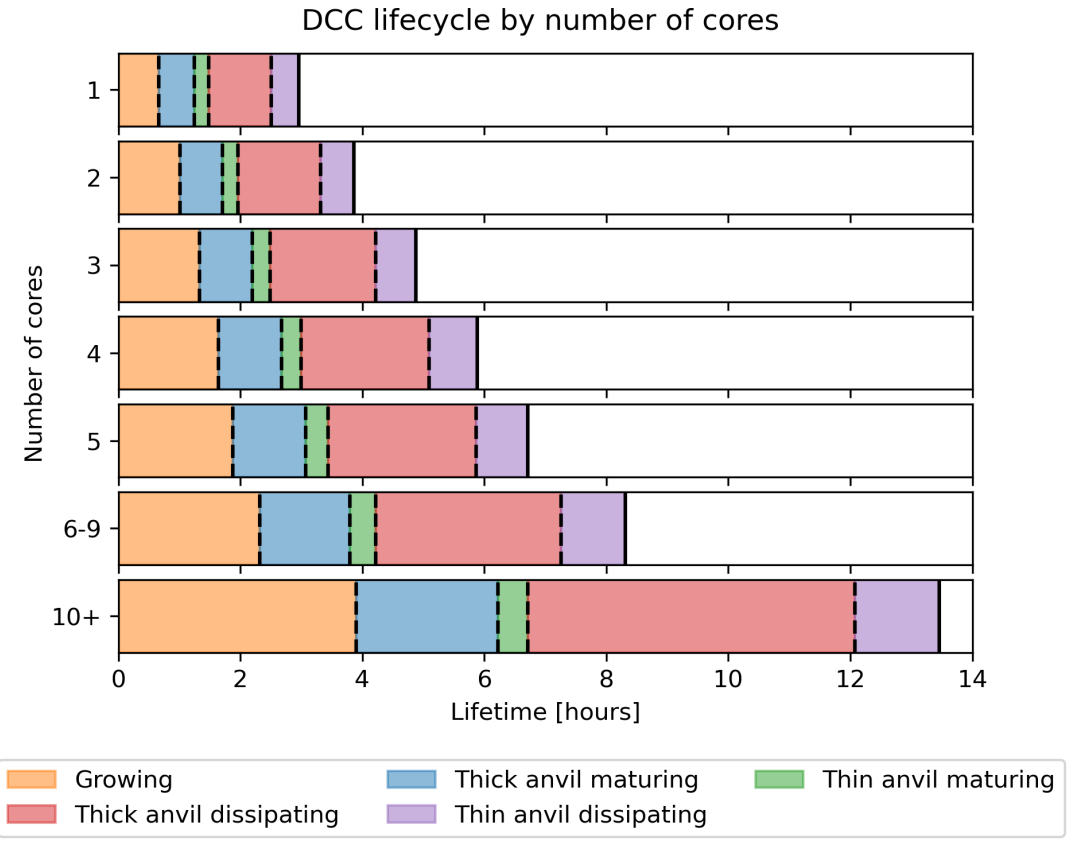


Figure 4.11: The average lifetime of each of the lifecycle stages defined in section 4.3.3 for anvils categorised by number of cores.

growing period and longer thin anvil dissipating period act together to shift the evolution of the thick anvil earlier in the DCC lifecycle for more intense DCCs. For DCCs over land, where there is a pronounced peak in occurrence in the afternoon, it is expected that this shifts the time in which the anvil is at its thickest and has the largest extent earlier in the day, potentially increasing its SW CRE.

Figure 4.11 shows how the timing of the different anvil stages changes with increasing number of cores. It is immediately apparent that the large increase in overall lifetime with the number of cores shown in fig. 4.7 also applies to each of the lifecycle stages. This increase in lifetime occurs to such an extent that the average growing period for DCCs with 10 or more cores lasts longer than the average of the entire lifetime of isolated DCCs. The large spread in lifetimes makes it difficult to compare

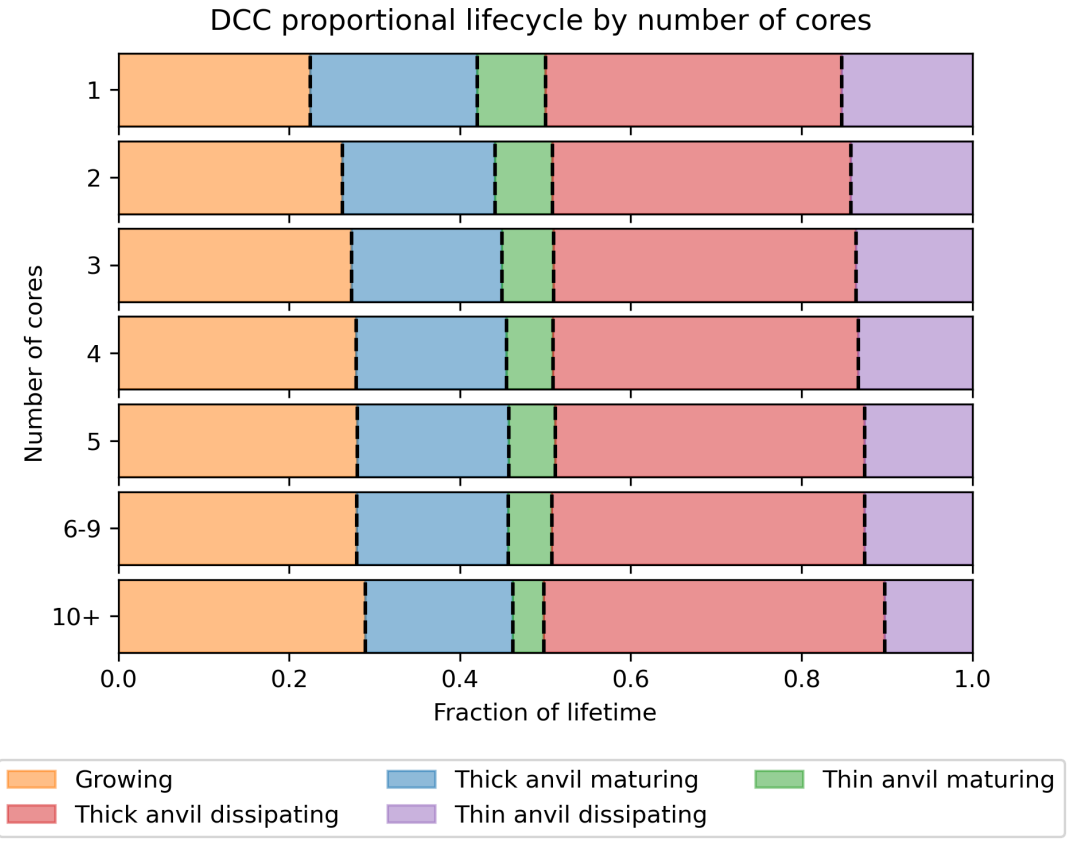


Figure 4.12: The average fraction of anvil lifetime of each of the lifecycle stages defined in section 4.3.3 for anvils categorised by number of cores. Each bar is normalised such that the total length of the row is equal to the mean thin anvil lifetime of its number of cores bin.

how the lifecycle changes with organisation without normalising the lifetimes.

Figure 4.12 shows the lifetime of each lifecycle stage as a proportion of the total lifetime for DCCs with different numbers of cores, in the same manner as used for fig. 4.10. In contrast to the effect of increasing intensity, with increasing organisation the length of the growing phase increases and the length of the thin anvil dissipating phase decreases as a proportion of the overall lifetime. Similar behaviour is seen in the period from the start of the mature phase to the thick anvil dissipation, which remains a consistent length for all core counts. With increasing organisation, this period shifts later in the lifetime of the DCC.

It should be clarified, however, that these differences are only in relation to the

proportion of total lifetime. While more organised systems spend a smaller proportion of their lifetime in the thin anvil dissipating phase, the much greater overall lifetime means that this stage still lasts substantially longer than for deep convective cloud (DCC)s with fewer cores.

4.4.4 Changes in thin anvil fraction over DCC lifecycle

While the contrasting changes in DCC lifecycle with increasing intensity and organisation (in particular the change in the thin anvil dissipating stage) go some way to explaining the difference in the thin anvil fraction, fig. 4.8 showed that the change in lifetime cannot account for the change in thin anvil fraction alone. The contrasting changes in lifecycle do, however, indicate that the two convective processes influence how DCCs evolve over their entire lifetimes. To further investigate how convective processes impact structure and lifecycle, the composite area approach described in section 4.3.4 is used to show how the thin anvil fraction evolves over the anvil lifecycle for DCCs with differing core cooling rates and number of cores.

Figure 4.13 shows composites of the thin anvil fraction for DCCs grouped by core cooling rate, with larger cooling rates shown in darker shades of blue. Each of the composites has a similar shape to the evolution of thin anvil fraction shown in fig. 4.3 b; an initial rapid increase as the anvil develops, a plateau during the maturing stage and a continual increase as the anvil dissipates. Two sections of the lifecycle show divergence in the thin anvil fraction with core cooling rate. Firstly, more intense cores produce a higher thin anvil fraction during the developing phase. This difference in thin anvil fraction is then maintained throughout the mature phase. Secondly, the thin anvil fraction further diverges during the dissipating phase, with the anvils originating from more intense cores having a larger increase in thin anvil fraction. This difference during the dissipating stage may be related to the difference in the thin anvil lifetime seen previously.

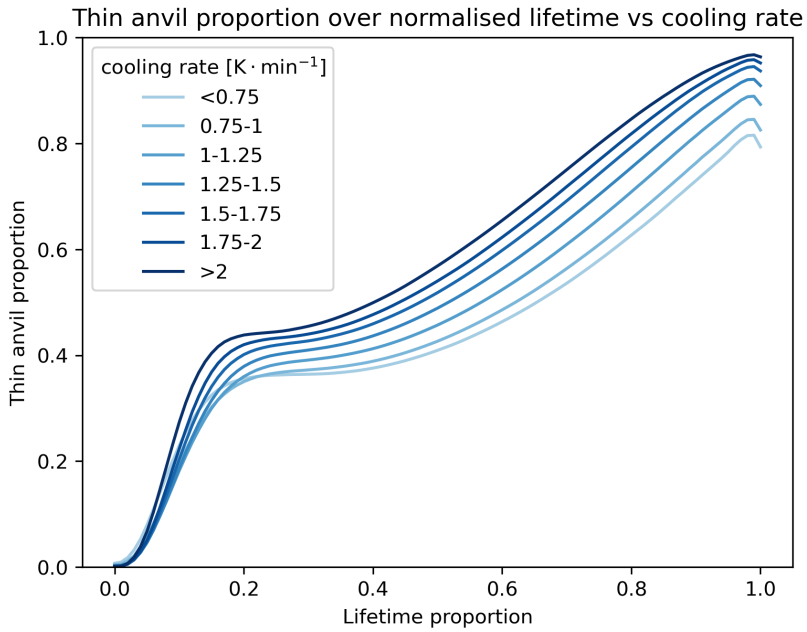


Figure 4.13: Mean thin anvil proportion for composites of anvil area of normalised lifetime categorised by the cooling rate of their initial cores. More intense core cooling rates are shown by the darker blue lines.

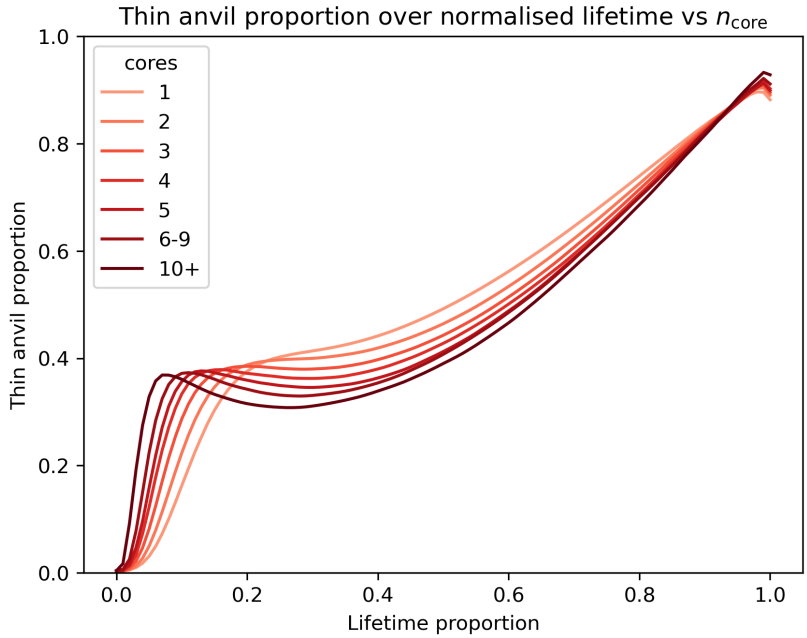


Figure 4.14: Mean thin anvil proportion for composites of anvil area of normalised lifetime categorised by the number of cores. Anvils with a larger number of cores are shown by the darker red lines.

Figure 4.14 compares how the mean thin anvil fraction changes for composites grouped by the number of cores. The composites with a greater number of cores are shown by the darker shades of red. In contrast to fig. 4.13, a difference from the typical shape of the thin anvil fraction established in fig. 4.3 is seen for more organised DCCs. At the end of the initial development phase of the anvil there is a similar thin anvil fraction of around 0.4 across all cases. While the more organised DCCs reach this point faster, this is simply due to the initial anvil time making up a smaller proportion of the total lifetime due to the latter's increase with convective organisation.

After the initial development of the anvil, there is a divergence in the thin anvil fraction during the later growing stage and the mature stage. For the most organised DCCs, there is a reduction in the thin anvil fraction during this period, indicating that the continuous, organised convection from multiple convective cores leads to a thickening of the anvil. This difference in thin anvil remains into the dissipating stage, although the composites all converge to a similar fraction by the end of their lifecycles.

The investigation of changes in thin anvil fraction over DCC lifecycle using composites reveals further complexity to the response of thin anvil structure to convective processes. Not only do changes in intensity and organisation have contrasting effects on the net thin anvil fraction and the lifecycle, but they also impact the anvil structure at different stages throughout the DCC lifetime. The timing of when differences in structure are seen may help investigate the processes through which they occur.

4.5 Summary

Deep convective anvil clouds have a characteristic ToA CRE that is near neutral [Hartmann, 2016]. Anvils are, however, not homogeneous objects, and exhibit structures that transition from thick near the convective cores to thin, detrained cirrus

further away. The tendency of reflectivity to reduce faster than emissivity with reducing cloud optical thickness, combined with the large SW and LW CRE of anvils, means that the CRE varies widely with changing thickness [Berry and Mace, 2014]. Despite this, the response of anvil optical properties to climate change remains largely uncertain, in part due to the lack of data connecting the behaviour of convective cores and anvils Gasparini et al. [2023].

Previous studies have found that the fraction of thin anvil increases with increasing convective intensity [Protopapadaki et al., 2017; Takahashi et al., 2017b]. These studies were performed using sun-synchronous polar-orbiting satellite data, and so could not resolve changes across the lifecycle of DCCs. With the dataset of tracked DCC cores and anvils, changes in the convective properties can be related to changes in structure observed over the entire anvil lifecycle.

Using the initial core cooling rate and number of cores as proxies for the intensity and organisation of convection respectively, a similar increase in thin anvil fraction is found with increasing intensity, but a decrease with increasing organisation. These relationships are found to be robust across a wide range of convective environments. Contrasting impacts on DCC lifecycle are also found as a result of both processes. In more intense DCCs, there is a reduction in the length of the growing phase and an increase in the thin anvil lifetime, whereas the opposite occurs in more organised cases. These corresponding changes in lifecycle indicate that the thin anvil fraction cannot be determined from single snapshots alone.

By analysing composites of the change in thin anvil fraction over the DCC lifecycle, impacts of intensity and organisation are apparent at different periods during the DCC lifetime. The timing of these changes also indicates which processes may be impacting the anvil structure in each case. For the more intense convection, there is an increase in the thin anvil fraction during the initial formation of the anvil, which is then maintained as the anvil matures. Studies of anvil evolution in high-resolution

convective resolving models have shown that, during the initial growth of the anvil, before the development of compensating downdrafts within the core, the divergence of the anvil is directly linked to the updraft velocities within the core [Senf et al., 2018b]. After the formation of compensating downdrafts after around 15–20 minutes, this connection is broken which may explain why this difference only occurs during the very early formation of the anvil.

A second period of increasing divergence in thin anvil fraction is seen during the dissipating stage of the DCCs. As this occurs after the end of the convective updrafts, it cannot be related directly to the convective processes but instead to their impact on the anvil properties. Over the lifetime of anvils, radiative heating of the cloud base and cooling of the cloud top result in circulations within the anvil cloud which have a thinning effect [Gasparini et al., 2019]. Once the anvil is sufficiently thin however (with optical depths around 1), these radiatively driven circulations cease and instead the remainder of the dissipation occurs due to the sedimentation and sublimation of the remaining thin cirrus [Sokol and Hartmann, 2020]. At the higher altitudes and colder temperatures of the anvils produced by more intense convection these processes are substantially slower [Seeley et al., 2019a], potentially resulting in larger extents and lifetimes for the thin, but not thick, anvil.

For more organised DCCs a reduction in the thin anvil fraction is observed not during the initial anvil formation but during the subsequent growing and maturing period of the anvil. Unlike isolated systems, organised DCCs involve the continuous development of new convective cores throughout most of the lifecycle of the more organised DCCs. It is possible that the thickening is a result of the mesoscale circulations associated with large, organised MCSs, which transports additional moisture directly into the anvil where it increases the cloud water content at high levels. Alternatively, convective aggregation may result in moister environments surrounding the cores of these organised DCCs, slowing the dissipation of the thick anvil, and

also drying the surrounding regions resulting in faster decay of the thin anvil further from the core. The convergence of the thin anvil fraction during the dissipating phase of DCCs is indicative of the characteristic decay patterns seen in convective anvils across scales [Roca et al., 2017; Elsaesser et al., 2022].

Care should be taken when interpreting these results to consider covariance in the metrics used to measure convective intensity and organisation with the measured anvil structure. In particular, the occurrence of thicker anvil clouds may result in the detection algorithm considering two neighbouring DCCs as forming a single cloud shield, when otherwise they would be considered isolated. This may result in a confounding factor where the anvil structure is affecting the number of cores, rather than the other way around. However, the time series of thin anvil fraction shown in fig. 4.14 suggests that this is not the case, as on average the anvils of DCCs with different numbers of cores tend to begin with similar thin anvil fraction and only diverged later in their lifetimes. Further analysis into the observations of organised convection and their dynamics could reveal more about the processes affecting the structure of these clouds, as well as examine differences in different modes of organised convection (such as tropical cloud clusters vs linear squall lines) which may show different behaviours regarding their structure.

Further investigation of how these processes impact anvil structure and lifecycle may be difficult due to the lack of time-resolved observational data that can measure these processes. The use of modelling, including both convective resolving models and more idealised experiments, may provide a route to isolate the impacts of each process on DCC evolution [Gasparini et al., 2022]. Lifecycle analysis, as performed in section 4.4.4, provides opportunities for new constraints on the representation of DCCs in convective resolving models. Analysis of km-scale models has revealed responses of anvil structure to changes in convection [Sokol et al., 2024], however as this was performed using bulk analysis it does not provide any insight into the

processes which cause this response, or why models differ from each other. Although geostationary satellite observations cannot observe these processes directly, studying the behaviour of DCCs over time may provide insight into the relative impacts on different processes. For example, in fig. 4.13, changes in the anvil structure are seen at two points during the convective lifecycle. Firstly, changes to the anvil during its initial formation must be linked to the convective dynamics themselves, whereas those during the dissipating phase must instead be a result of in situ anvil processes as by this point in the lifecycle the convective dynamics have stopped. Comparing how anvils in convective resolving models evolve over time to observations may help identify which processes are important to anvil structure and their relative influence, even if these processes cannot be observed themselves.

The final element to consider is how these observed changes in anvil structure may impact the anvil CRE. Considering the thin anvil fraction alone, it could be expected that more intense DCCs are more warming due to their thinner structure, and vice versa for more organised convection. However, this neglects other factors relevant to anvil CRE that are also impacted by convective processes. Both convective intensity and organisation affect the temperature of anvil clouds, resulting in a warming effect for increases in both. Furthermore, the observed changes in the lifecycle and the timing of the maximum extent of the thick anvil within the DCC lifetime may impact the SW CRE, particularly over land due to the pronounced diurnal cycle of convection. Observational studies have shown a link between increasing convective organisation and domain mean ToA cooling [Bony et al., 2020]. While this has traditionally been explained as a result of greater clear sky LW emissions to space [Bony et al., 2016], this could also result from the response of anvil CRE to changes in structure.

Chapter 5

The Lifecycle and Cloud Radiative Effect of Deep Convective Clouds Over Africa

5.1 Introduction

DCC play a key role in the tropical atmosphere over Africa. Forming the ascending branch of the Hadley cells near the equator, DCCs are critical to the circulation and heat transfer of the tropics [Riehl and Malkus, 1958; Weisman, 2015]. The inter-tropical convergence zone (ITCZ) and its location are vital in determining the seasonal cycles of rainfall over central and western Africa [Nicholson, 2009, 2018]. Understanding the behaviour of DCCs over Africa has the potential for major impacts on the atmosphere, weather and society.

DCCs also exert a key influence on the temperature of the tropics through their CRE. Due to their size, height and depth, DCC anvils have large radiative effects in both the SW and LW, with both having average magnitudes that exceed $100\,\mathrm{Wm}^{-2}$ [Hartmann, 2016; Wall and Hartmann, 2018]. Due to the opposite signs of these two components, the average anvil CRE in the tropics is approximately zero [Ramanathan et al., 1989; Hartmann et al., 1992; Stephens et al., 2018]. Much of the focus on the anvil CRE feedback to global warming has been placed on the LW response (see section 1.3, and in particular only consider the net response of anvils to large scale

the tropical atmosphere, rather than the mechanisms affecting individual DCCs. In particular, the sensitivity of anvil CRE to changes in the diurnal cycle of convection has recieved little attention. Previous research has highlighted that changes in the diurnal cycle of convection over Africa may lead to changes in CRE of $\pm 10\,\mathrm{Wm^{-2}}$ [Nowicki and Merchant, 2004]. Further investigation into the response of anvil CRE to changes in the diurnal cycle highlighted the need for cloud tracking approaches to study CRE over the anvil lifetime [Bouniol et al., 2016, 2021].

In this chapter, the novel cloud tracking methodology presented in chapter 2 is used in conjunction with derived all-sky and clear-sky radiative fluxes to characterise the CRE over the lifecycles of individual anvil clouds. This methodology is applied to 4 months of data produced for the European Space Agency (ESA) Cloud-Climate Change Initiative (CCI)+ project over sub-Saharan Africa. This dataset allows us to investigate both the CRE of individual CREs, as well as the net anvil CRE over the entire region. The overall distribution of anvil CRE is found to be determined by the relationship between DCC lifecycle and the diurnal cycle of the SW CRE. This has important implications for the response of DCCs to a changing climate, as previously the impact of changes in the diurnal cycle of anvil CRE has seen little study.

Channel	Wavelength (µm)	Description	Tracking	Retrieval
1	0.64	Visible		✓
2	0.81	NIR		✓
3	1.64	NIR		✓
4	3.92	NIR Window		✓
5	6.25	Upper troposphere WV	✓	✓
6	7.35	Lower troposphere WV	✓	✓
7	8.70	Mid-IR window		
8	9.66	Ozone		
9	10.8	Clean LW window	✓	✓
10	12.0	Dirty LW window	✓	✓
11	13.4	CO_2		✓
12	0.6–0.9	High-resolution visible		

Table 5.1: SEVIRI channels and their use in the DCC tracking algorithm and cloud properties retrieval.

5.2 Data

For this case study, data is used from SEVIRI [Aminou, 2002] aboard the Meteosat Second Generation (MSG) Meteosat-11 satellite, which is in a geostationary orbit above the equator at 0°W. Data from 4 months (May–August 2016) over sub-Saharan Africa (approximately 18 °W–46 °E, 31 °S–15 °N) is used at the full resolution of SEVIRI (3 km at nadir), along with retrieved cloud properties and derived broadband fluxes produced by the ESA Cloud-CCI+ project. BT from SEVIRI is used by the tracking algorithm, and reflectances and BT are used by the cloud retrieval.

SEVIRI is a visible and IR radiometer with a nadir spatial resolution of 3 km and a temporal sampling time of 15 minutes for the full earth disc. SEVIRI has 12 channels across the visible, NIR and thermal-IR spectrum, with one being a high-resolution visible channel with a nadir resolution of 1 km. A brief overview of these channels, along with which are used for tracking DCCs and the cloud properties retrieval, is provided in table 5.1.

An example of observations from SEVIRI is shown in fig. 5.1 for 15:00:00 UTC on 1st June 2016. A visible composite (fig. 5.1 a) is constructed using the 1.64 µm and 0.81 µm near-infrared and 0.64 µm visible channels for the RGB channels respectively. In this composite, ice clouds (which appear cyan) can be seen over central Africa and the southern Atlantic. Figure 5.1 b shows the 10.8 µm brightness temperature for the same scene, showing the coldest temperatures for the high ice clouds over central Africa. Two combinations of channels are used for the detection of anvil clouds. The WVD, shown in fig. 5.1 c, consists of the 6.3 µm BT minus the 7.4 µm BT. The use of these channel combinations for the detection of thick and thin anvil clouds in satellite imagery is detailed in section 2.2.1.1.

Comparing SEVIRI to ABI (fig. 5.2), the 10.8 µm BT (top panel) and WVD (middle panel) show very similar values for both instruments. The 10.8 µm and

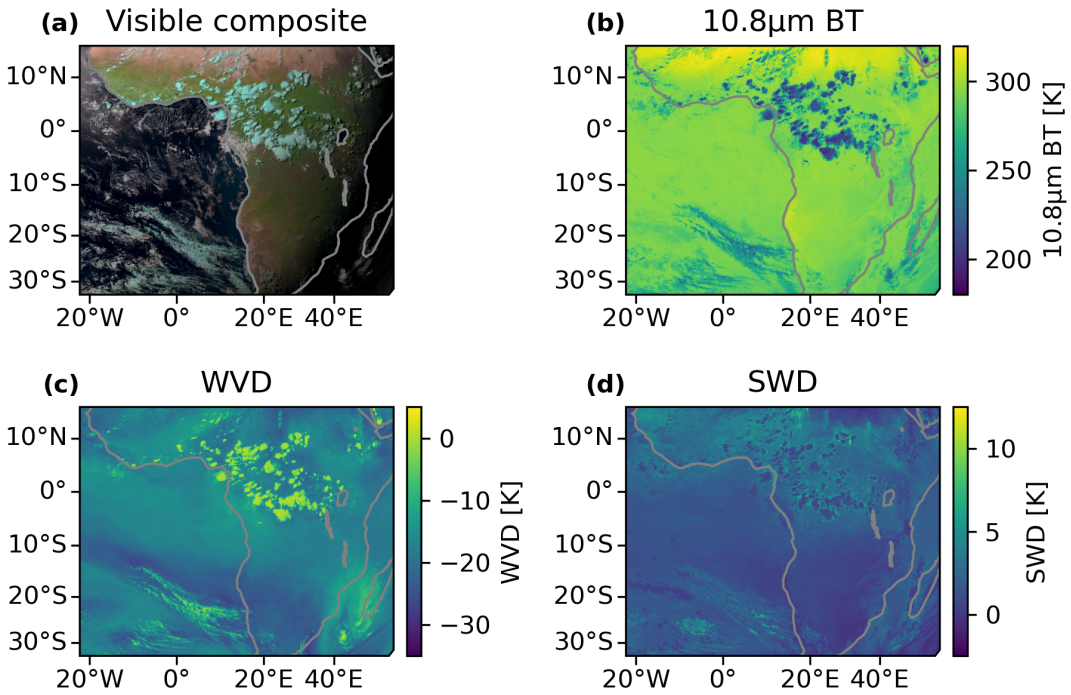


Figure 5.1: Example observations from the Meteosat SEVIRI instrument at 15:00:00 UTC on 2016/6/01. a: A visible composite formed using the 1.6, 0.81 and 0.64 µm channels as the RGB channels respectively, with 10.8 µm BT during the night-time. The scene shows a cluster of cold cloud tops (cyan) over central Africa and the Southern Atlantic. b: 10.8 µm BT. c: WVD formed by the 6.3 µm channel minus the 7.4 µm channel. d: SWD formed by the 10.8 µm channel minus the 12.0 µm channel.

12.0 µm channels of SEVIRI have relatively wide wavebands (see fig. 5.3), and as such are less sensitive to the presence of thin ice clouds (fig. 5.2 c), with a response about half that measured by ABI. Combined with their high NEΔT (see table 2.2 and discussion in section 2.2.1.1), it is found that the detection of thin anvil is unreliable using SEVIRI, and so the thin anvil is not considered within this chapter.

Retrieved cloud properties, including optical thickness, effective radius, liquid/ice water path, CTT and height, are provided by the Community Cloud Retrieval for Climate (CC4CL) algorithm [Sus et al., 2018; McGarragh et al., 2018]. These properties are all retrieved at the same resolution as the input SEVIRI data. Broadband fluxes are derived using the BUGSRad radiative transfer model [Stephens et al., 2001] using

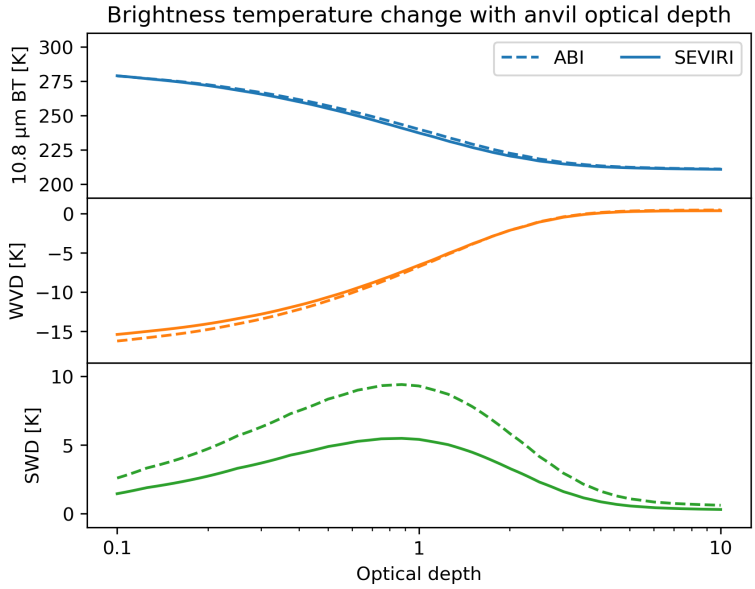


Figure 5.2: Comparison of the sensitivities of ABI (dashed lines) and SEVIRI (solid lines) to anvil clouds of different optical thickness, using the LibRadTran simulation of an anvil at 14 km as described in section 2.3.2.

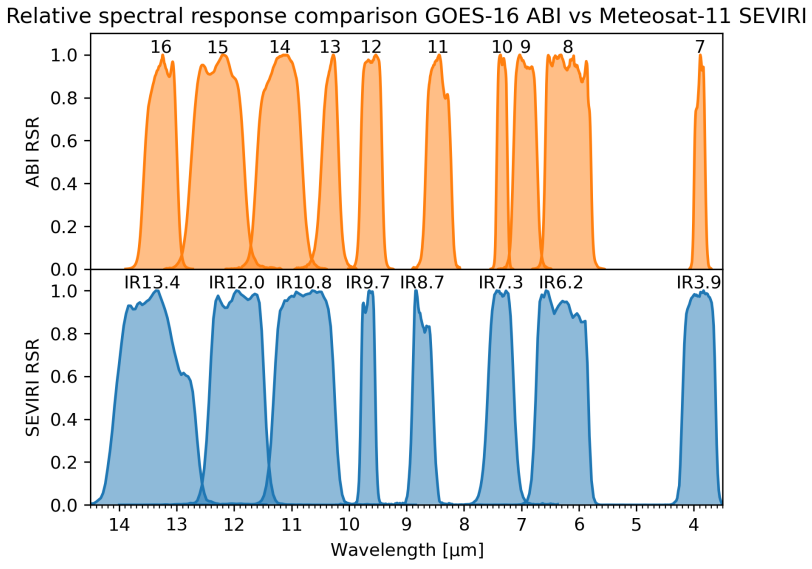


Figure 5.3: Comparison of the relative spectral response (RSR) functions for the GOES-16 ABI and Meteosat-11 SEVIRI thermal IR channels. The LW window channels on ABI (channels 13 and 15) have a wider spacing than those of SEVIRI (channels IR10.8 and IR12.0).

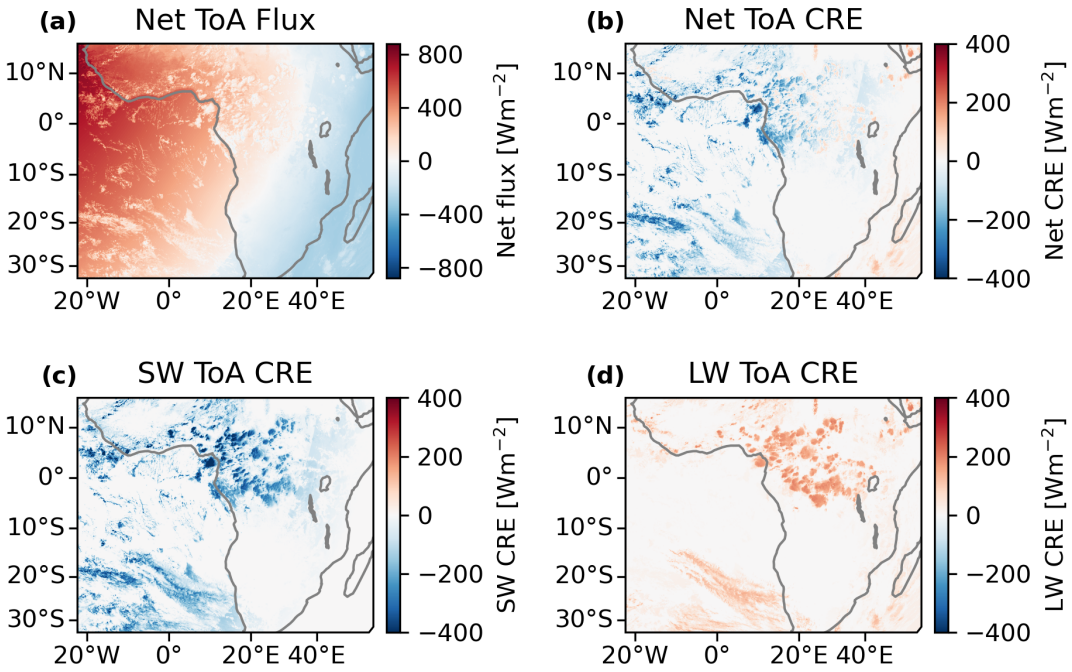


Figure 5.4: An example of the ToA CRE derived using the radiative flux model, for the same time as shown in fig. 5.1 (15:00:00 UTC on 2016/6/01). a: net ToA radiative flux. b: net ToA CRE. c: SW downwards CRE. d: LW downwards CRE.

input cloud properties from the CC4CL retrieval and vertical temperature, moisture and trace gas profiles from ERA-5 [Hersbach et al., 2020]. The BUGSRad model provides ToA and bottom-of-atmosphere LW and SW radiative fluxes for both all-sky and clear-sky conditions. An example of these derived fluxes is shown in fig. 5.4. Figure 5.4 a shows net ToA fluxes, with a net warming during the daytime on the Western side of the image, and a net cooling at night-time on the Eastern side. Figure 5.4 b shows the net ToA CRE, with a net cooling effect during the daytime and warming during the night-time for observed high clouds over central Africa. The SW (fig. 5.4 c) and LW (fig. 5.4 d) components of the CRE show that while the LW, warming component has a smaller magnitude than the day-time, cooling SW CRE, it remains constant during both day- and night-time.

Validation of the SEVIRI broadband fluxes was performed against monthly-mean

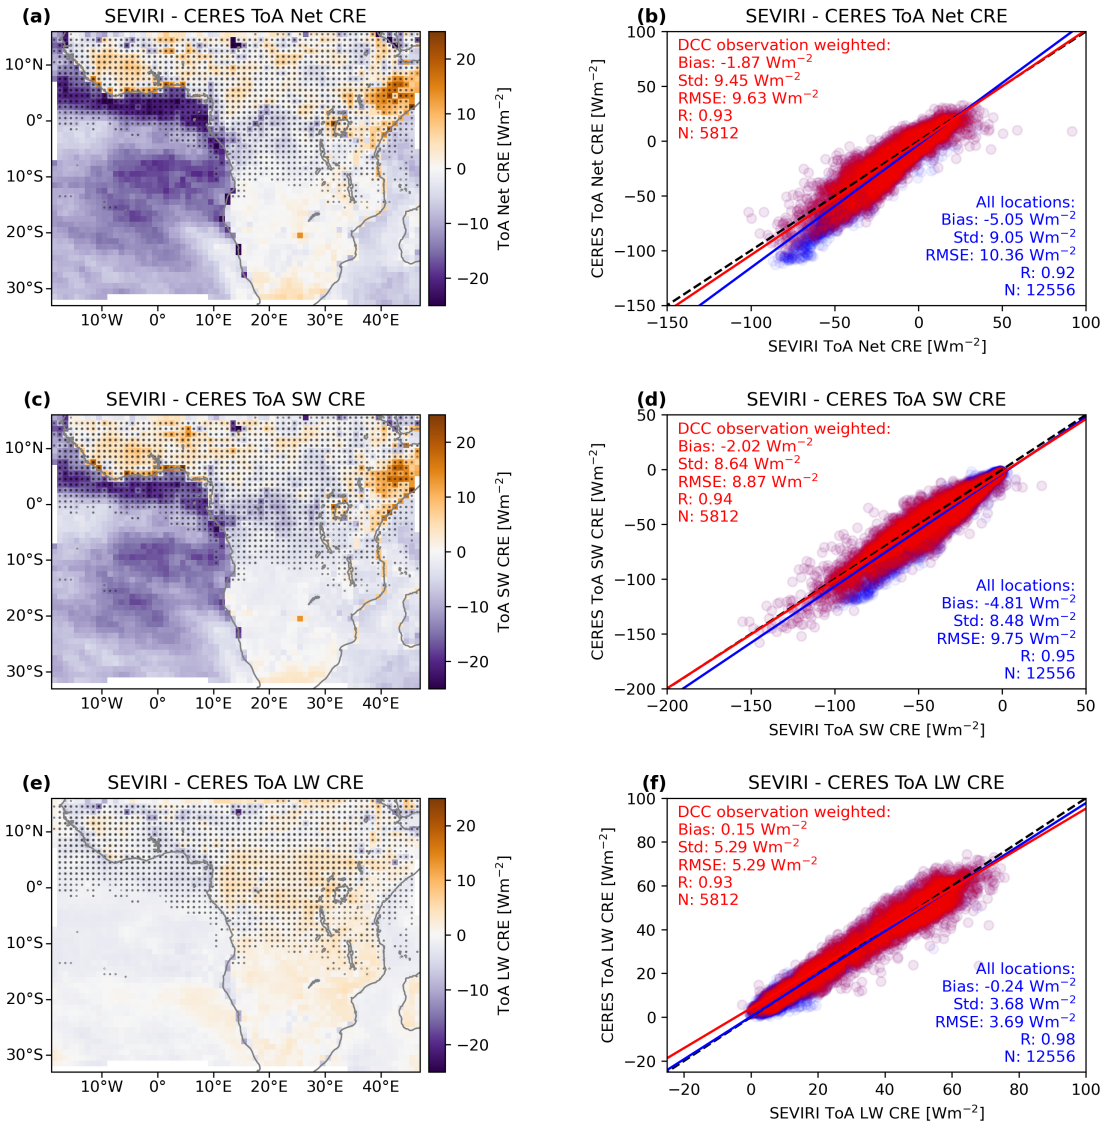


Figure 5.5: Validation of derived broadband fluxes against monthly CERES-EBAF CRE. a.: The mean difference in net ToA CRE by 1×1°grid square. b.: A comparison of observed ToA net CRE for SEVIRI against CERES, with all locations in blue, and those where DCC anvils are observed in red. c.: the mean difference in SW ToA CRE. d.: comparison of SW ToA CRE for SEVIRI and CERES. e.: the mean difference in LW CRE. f.: comparison of LW ToA CRE. The stippling in a, c and e represents the locations in which DCC anvils are observed, with the size of the dots corresponding to the number of observations. The solid lines in b, d and f show the linear regression for all locations (blue) and the locations ofDCC anvils (red) weighted by the number of observations.

observations of ToA broadband CRE from the Clouds and the Earth's Radiant Energy System (CERES) [Loeb et al., 2018] energy balanced and filled (EBAF) climate data record. The results of this validation are shown in fig. 5.5. Monthly mean fluxes were calculated for SEVIRI by first calculating the mean daily fluxes over each 1×1°grid square for days in which over 23 hours of observations are present, and then averaging these daily means over each month. Comparison of the net ToA CRE to CERES revealed a bias of $-1.87\,\mathrm{Wm^{-2}}$ (fig. 5.5 a,b), consisting of a SW bias of $-2.02\,\mathrm{Wm^{-2}}$ (fig. 5.5 c,d) and a LW bias of $+0.15\,\mathrm{Wm^{-2}}$ (Fig 5.5 e,f). Corrections for these biases have been applied uniformly to all further CRE values given in this chapter.

5.3 Method

The detection and tracking of DCCs are performed using the tobac-flow algorithm (see chapter 2, which has been designed specifically to track both isolated and clustered DCCs in geostationary satellite imagery over their entire lifecycle.

Due to the lack of sensitivity of the SEVIRI SWD to thin ice clouds, only the thick portion of the anvil is tracked in this chapter. The WVD channel of SEVIRI is capable of detecting anvils with optical thicknesses of approximately 1.5±0.5 (see fig. 5.6). The anvils tracked in this chapter have a median retrieved minimum optical depth of 1.45, although this value may be biased high as many anvils dissipate at night when accurate satellite retrievals of optical depth are not available. While this sensitivity captures much of the CRE of DCC anvils [Berry and Mace, 2014] the long lifetimes of dissipating thin anvils may have a significant warming contribution to net anvil CRE [Horner and Gryspeerdt, 2023]. As a result, it is expected that the anvil CRE measured in this study is biased low.

An example of the cores and anvils detected by the tobac-flow algorithm is shown in fig. 5.7, at 3-hour intervals. In fig. 5.7 a, a large number of developing cores are seen over central Africa. In fig. 5.7 b, there are more developing cores over Western

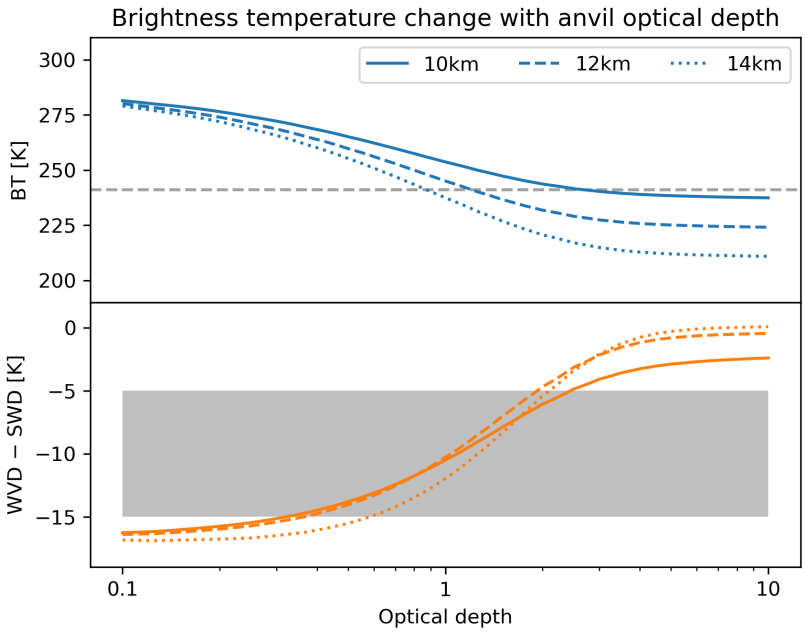


Figure 5.6: Simulated sensitivity of the SEVIRI 10.8 µm BT (top) and WVD minus SWD (bottom) to anvil clouds of varying optical thickness at heights of 10, 12 and 14 km as described in section 2.3.2. The grey dashed line shows the 241 K BT. The grey region in the lower plot shows the range of temperatures in which the edge of the anvil is detected, as described in chapter 2.

Africa as the pattern of initiation has shifted westward with the diurnal cycle. In fig. 5.7 c,d fewer newly developing cores are observed later in the day, but the larger anvil clouds persist into the night-time.

Over the 4 months of the case study a total of 145,463 cores (of which 79,592 are associated with anvil clouds) and 35,941 anvils are tracked. Using the detected regions of both core and anvil components of tracked DCCs, the cloud properties and CRE are calculated for each DCC at each time step from the retrieval and broadband fluxes data. The resulting dataset allows us to analyse the properties of each DCC over their lifetimes from a Lagrangian perspective. While the studied domain contains both land and sea regions, only a small proportion of tracked DCCs occurred over sea (11%), and so analysis of land and oceanic DCCs are not separated in this chapter.

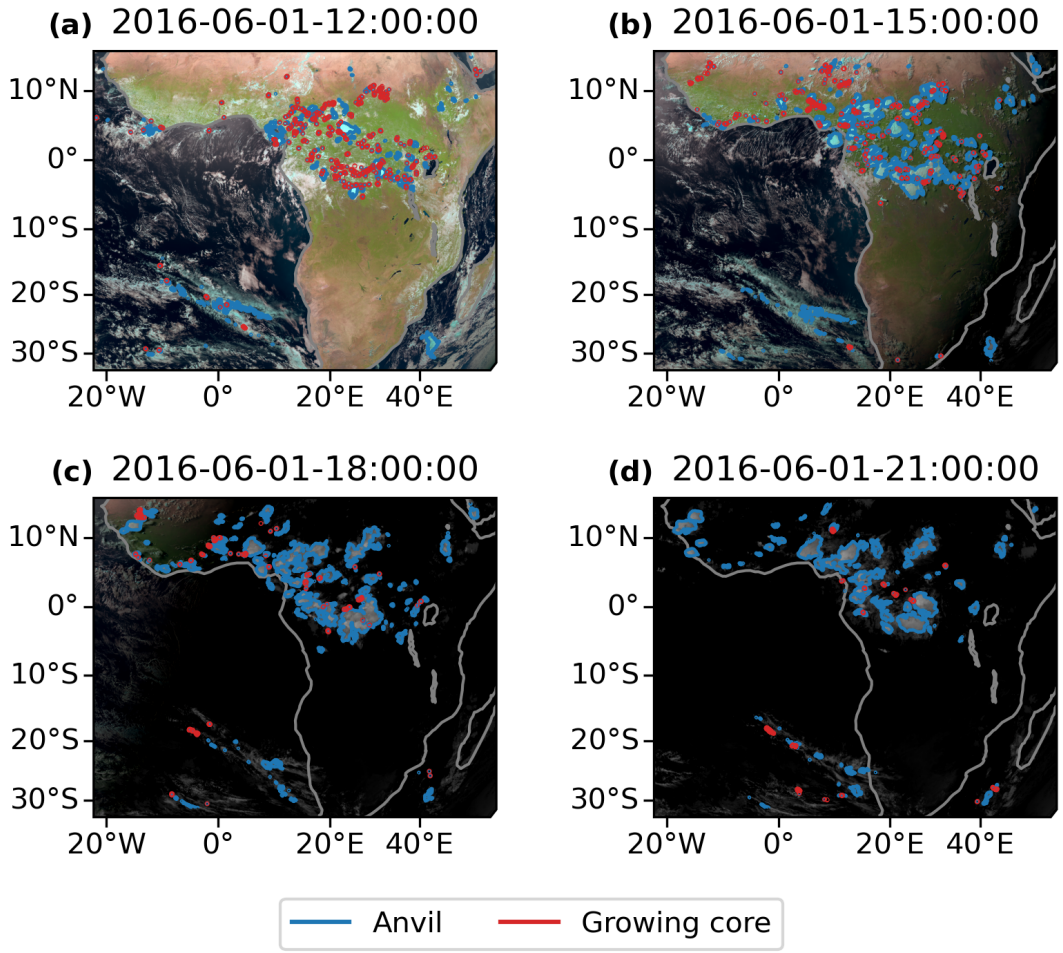


Figure 5.7: An example of the cores (red outline) and anvils (blue outline) detected by tobac-flow plotted over visible composite imagery from SEVIRI, shown at 3-hour time intervals. All times are given in UTC.

5.4 Results

5.4.1 Spatial and temporal distributions

Figure 5.8 a shows the frequency of core detections for each 1×1° grid square throughout the case study. The majority of observed convection occurs over the tropical rainforest regions. During the months of May-August, the inter-tropical convergence zone (ITCZ) is at its northernmost extent over Africa [Nicholson, 2018]. The West African monsoon occurs during these months, with the primary band of convection located between 5-15°N [Nicholson, 2009], which our observations agree with. The

maximum frequency of convection is observed at around 6°N, 12°E over the Western High Plateau of Cameroon, with high frequencies of convection also observed over the Nigerian coastal plains to the West and the Jos Plateau in Northern Nigeria. High rates of convection are also observed over the coastal plains and inland highlands of Guinea, Sierra Leone and Liberia (5–12 °N, 5–15 °W). Almost no convective activity is observed between 10°S and 20°S as, during the Northern hemisphere summer, this is the location of the descending branch of the Hadley cell which suppresses convection.

Figure 5.8 b shows the average time of detection for convection in each $1\times1^\circ$ grid square. The average is calculated as the circular mean of the local solar times of core detection in the grid square. Grid squares with a standard deviation greater than 6 hours (indicating a broad spread of initiation times) are given single hatching, and those with standard deviations greater than 12 hours have cross-hatching. The most notable feature of the time of detection is the clear contrast between land and sea. Convection over the land tends to occur in the afternoon (15:00–18:00), whereas over the ocean it occurs between midnight and early morning (00:00–09:00). Furthermore, convection over land tends to occur in a fairly narrow range of times whereas over the ocean convection occurs throughout the diurnal cycle, resulting in the hatching applied to much of the ocean region. There is also a noticeable lake effect on the time of convection occurring over Lake Victoria (2°S, 34°E) and Lake Tanganyika (7°S, 31°E), with convection typically observed in the early morning.

When comparing the regions of Cameroon and Nigeria (4–10°N, 6–14°E), where the most growing cores were detected in fig. 5.8 a, with the average time of detection in fig. 5.8 b, it appears that the grid squares with more cores also tend to have an earlier average time of detection than the surrounding grid squares. Precipitation over the Nigerian plains and the Jos Plateau is linked to South-westerly winds bringing moist, warm air from the Gulf of Guinea [Vondou et al., 2010]. This warm air may then trigger convection both through the sea breeze effect and orographic lifting when

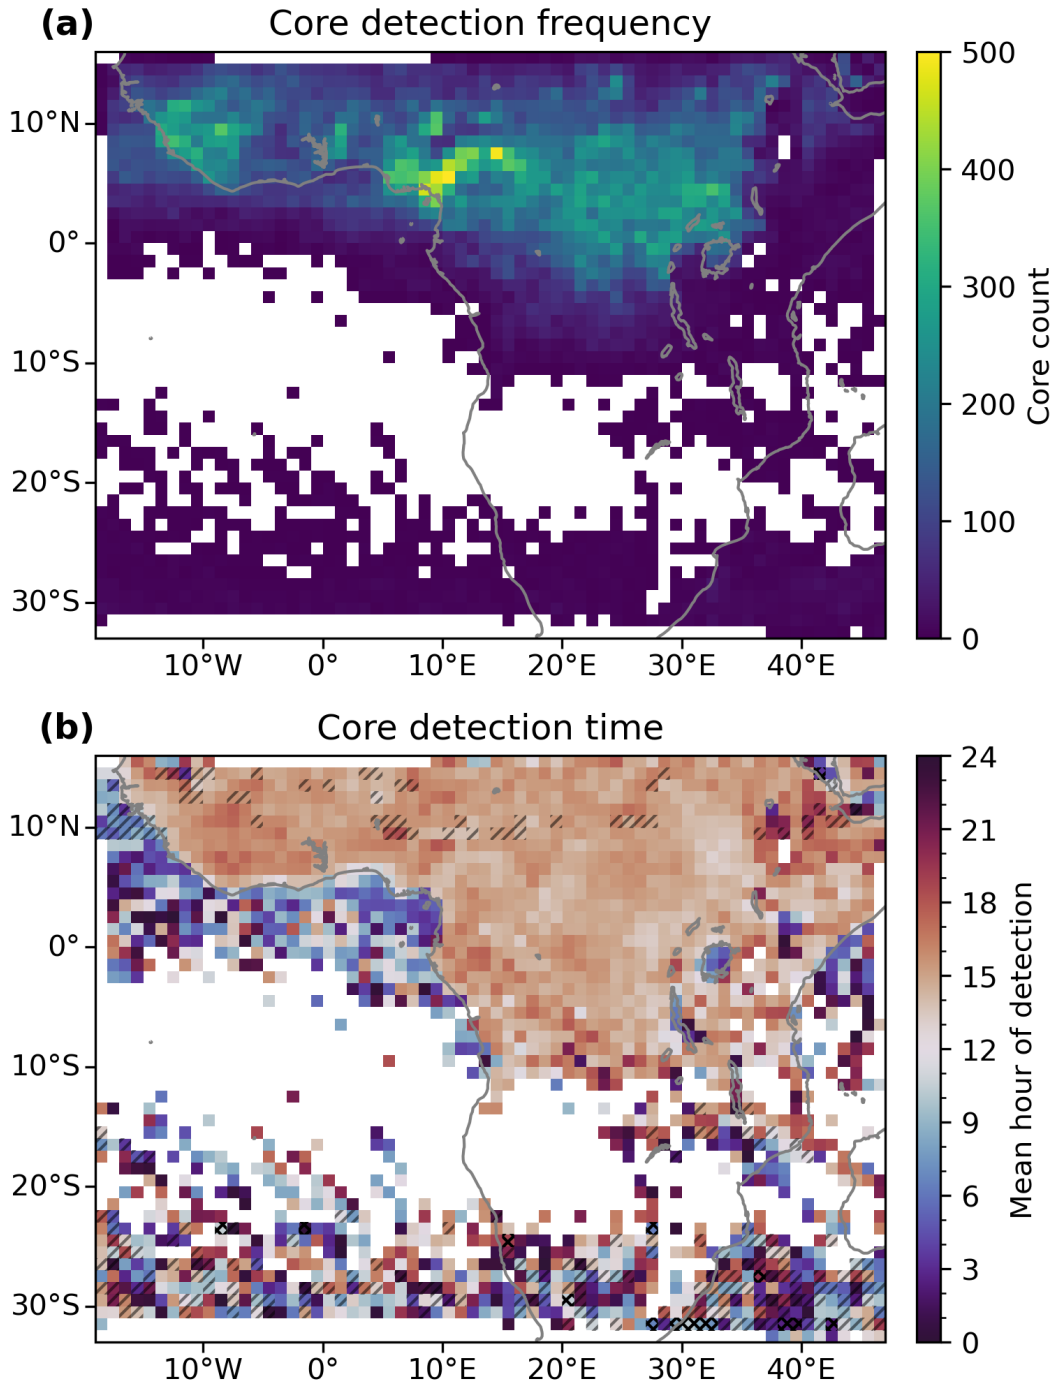


Figure 5.8: a.: The total number of DCC cores detected over the case study for each 1×1° grid box. b.: The average hour of detection for the cores detected in each 1×1°grid box. Grid boxes with a standard deviation greater than 6 hours are single-hatched, and greater than 12 hours cross-hatched.

it reaches the highlands, explaining both the higher frequency and earlier timing of convection compared to surrounding regions. A similar relationship between the high frequency of convection and earlier time of detection is also seen over the coastal region and adjacent highlands of Guinea, Sierra Leone and Liberia (5–12°N, 5–15°W) which may be due to the same mechanism.

It should be noted that due to the method of detection, cores that develop under existing anvils are less likely to be detected than those in clear sky regions. As a result, the occurrence of subsequent developing cores in organised systems may be underestimated, particularly in regions such as the Northern Sahel where a second, night-time peak of precipitation has been observed.

For all further analysis, only cores and anvils that are detected north of 15°S are considered so that the analysis is constrained to tropical DCCs.

5.4.2 Anvil Cloud Properties

To investigate how the behaviour of DCC anvils is affected by their organisation, observed anvils are grouped based on how many cores are associated with them, from isolated DCCs with one core to highly-clustered DCCs (such as tropical cloud clusters and MCSs) with 10 or more cores. Anvils with 6–9 cores, and with 10 or more cores, are grouped together to ensure that these groups have a comparable number of members for analysis.

Figure 5.9 shows properties related to the anvil area and lifetime linked to the number of cores. In fig. 5.9 a the average anvil maximum area is shown for each group. The maximum area increases approximately linearly with the number of cores, with increasingly clustered anvils having increasingly larger maximum areas, and highly clustered anvils having substantially larger anvils. Figure 5.9 b shows the average anvil lifetime compared to the number of cores. While the lifetime also increases with the number of cores, the difference between isolated and highly clustered anvils is

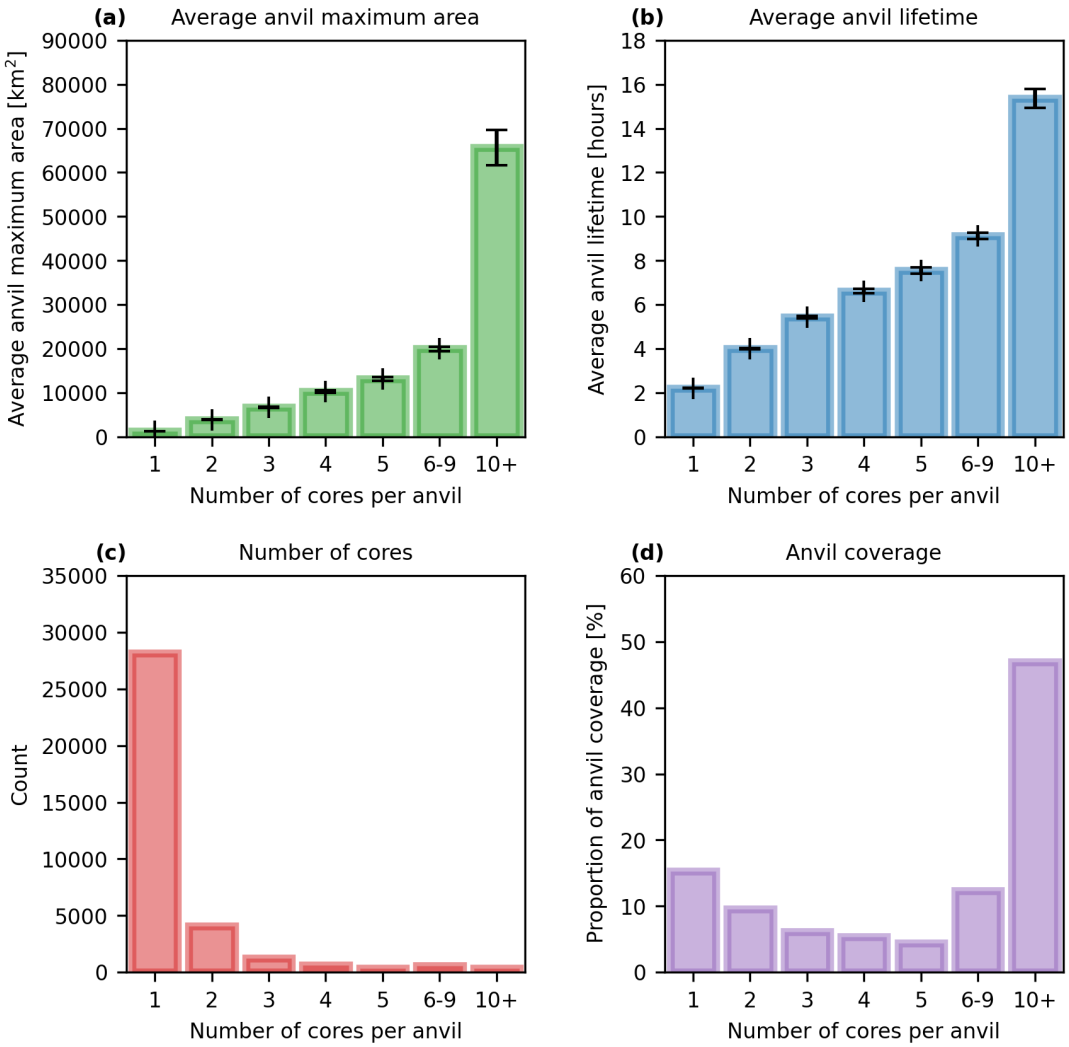


Figure 5.9: Anvil statistics by number of associated cores for a.: average maximum area; b.: average lifetime; c.: the number of observed anvils by number of cores; and d.: percentage of total anvil coverage. Error bars in a and b show the standard error of the mean.

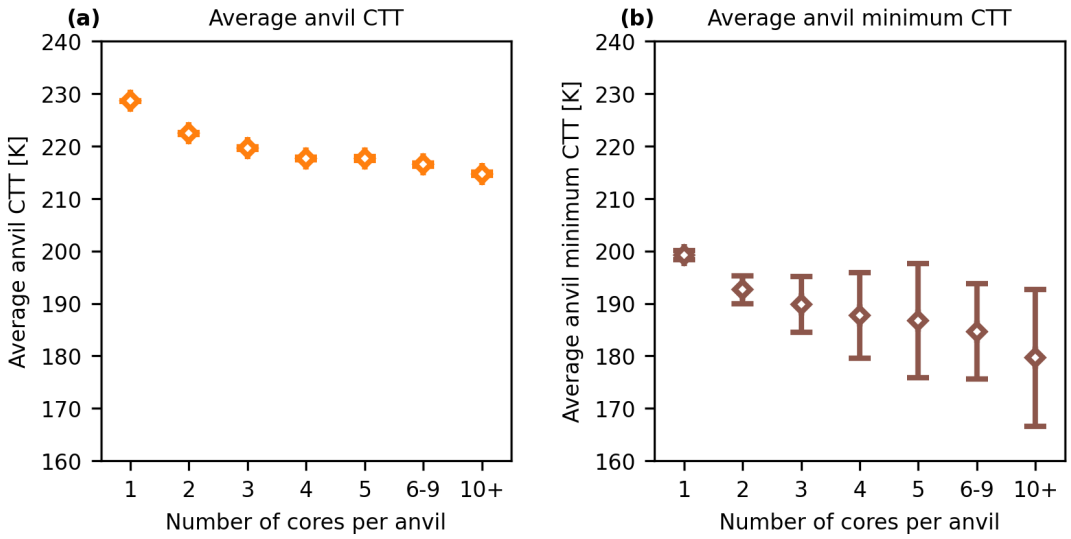


Figure 5.10: Anvil statistics by number of cores for a.: average anvil CTT; and b.: average minimum anvil temperature. Error bars show the standard error of the mean.

proportionately smaller.

Figure 5.9 c shows the number of anvils observed with differing numbers of cores. The vast majority of all anvils observed are isolated DCCs, with over 80% having a single detected core. As the number of cores increases, the number of anvils detected decreases rapidly. However, when considering the large increase in both anvil area and lifetime with the number of cores, the total anvil coverage for highly clustered anvils is much larger (see fig. 5.9 d). Despite their high frequency, isolated DCCs only account for 12% of total anvil coverage, whereas highly clustered (10+ cores) account for over 50%. Previous studies have found that despite being few in number, MCSs account for the majority of precipitation in Western Africa [Vizy and Cook, 2019].

Figure 5.10 a shows the average mean CTT, and fig. 5.10 b the average minimum CTT for anvils with different numbers of cores. While the more clustered anvils have colder average anvil CTT, this decrease plateaus below 220K indicating that the reduction in clear-sky cooling below this temperature may cap the anvil CTT for larger DCCs. The minimum observed CTT within each anvil, however, are colder

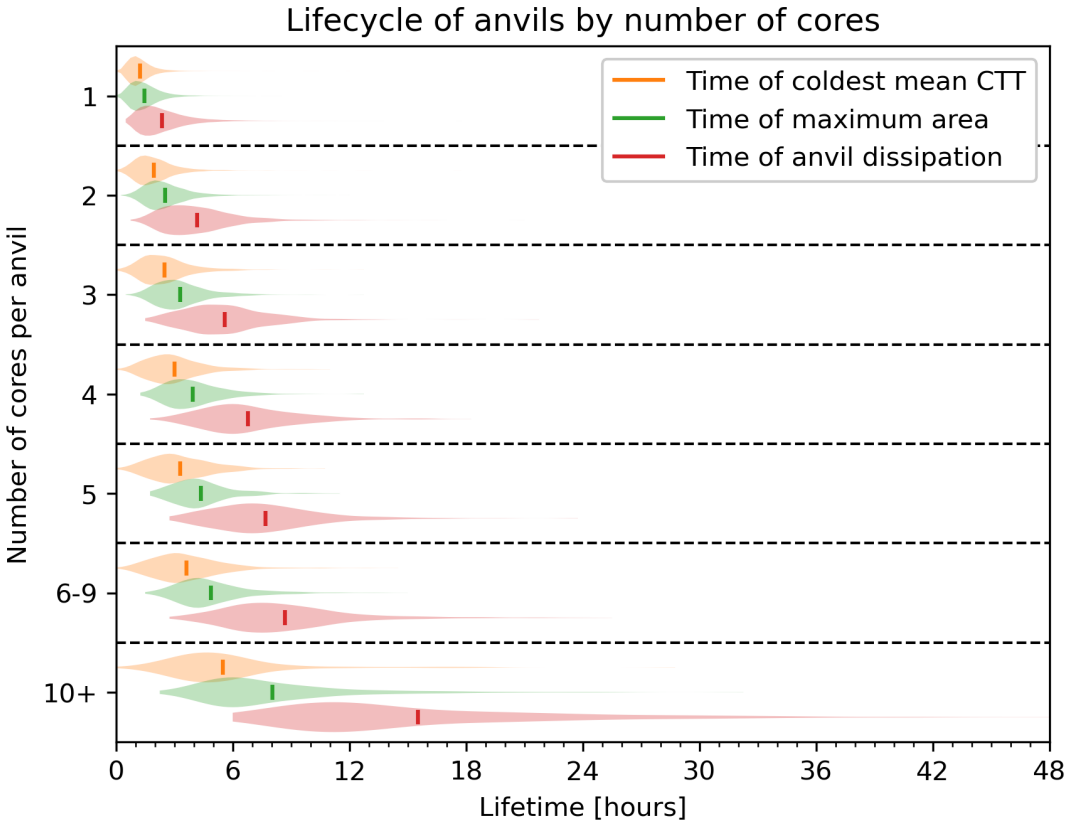


Figure 5.11: The distribution of time to coldest mean anvil CTT (orange), largest anvil area (green) and time of anvil dissipation (red) for anvils grouped by number of cores. The vertical lines show the mean time for each distribution.

and show a greater difference with an increasing number of cores. The most clustered anvils tend to have a minimum CTT of around 180 K, indicating the presence of overshooting tops and the most intense convection. These cold CTT values have large uncertainty due to the high sensor noise-to-signal ratio at these cold temperatures.

Futyan and Del Genio [2007] divide the DCC lifecycle into growing, mature and dissipating phases based on the time of observation of the coldest anvil CTT, maximum anvil area and dissipation of the anvil. Figure 5.11 shows the distribution of the time taken to reach each of these lifecycle milestones for anvils separated by the number of associated cores. For all cases, the average time of minimum anvil CTT occurs before the maximum area, indicating that the anvils continue to grow beyond

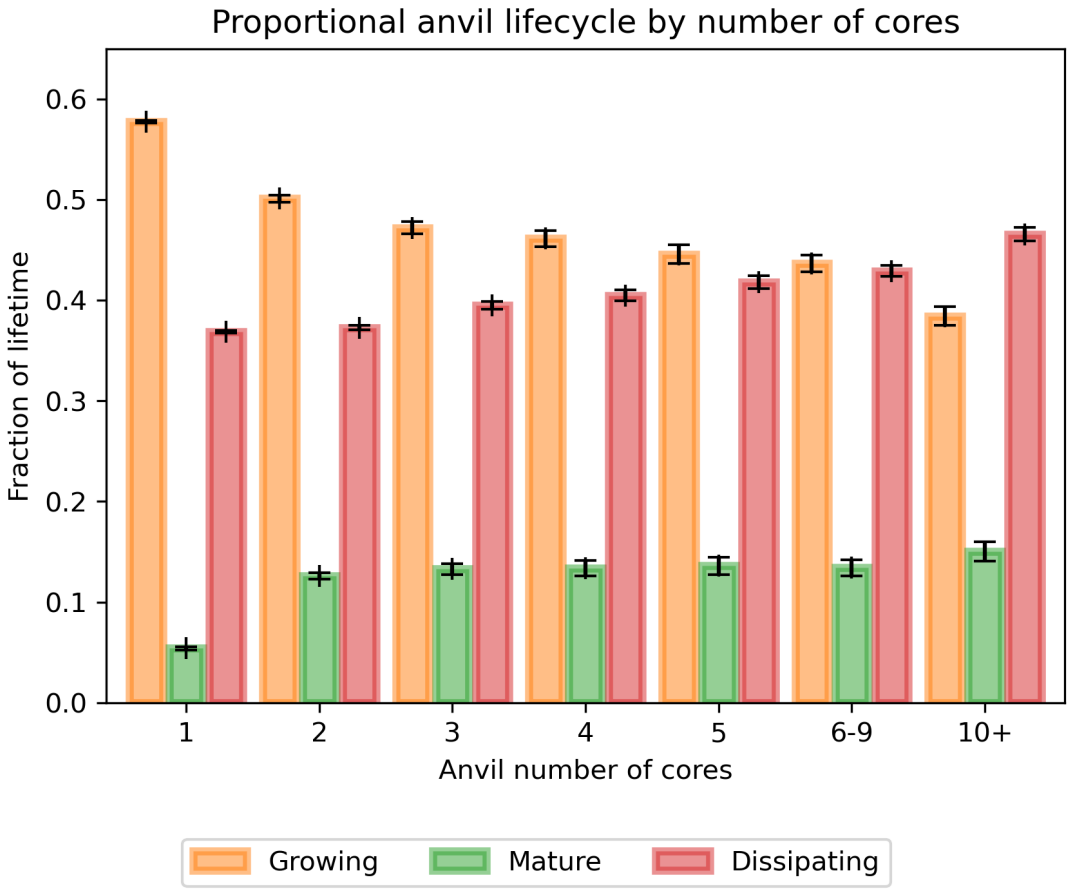


Figure 5.12: The proportion of anvil lifetime spent in the growing (orange), mature (green) and dissipating (red) phase, according to the criteria used by Futyan and Del Genio [2007]

the maximum of convective activity. As the number of cores associated with each anvil increases, the time of the coldest CTT and largest area occur proportionately earlier during the lifetime of the anvil. As a result, these more clustered anvils spend more of their lifetime existing with warming, shrinking anvils than the isolated DCCs.

Figure 5.12, compares the proportion of the overall anvil lifetime spent in each of the lifecycle phases defined by Futyan and Del Genio [2007] to the number of cores associated with the anvil. There is a clear trend that, as the number of cores increases, the proportion of the lifecycle spent in the growing phase decreases, and the proportion spent in the mature and dissipating phases increases. Although this

approach to classifying the lifecycle of anvil clouds is simplistic and does not capture the complexities of large, long-lived DCCs which may go through multiple cycles of growth, dissipation and re-invigoration, it provides a useful simplification for comparing complex DCC lifecycles [Roca et al., 2017]. The time of the coldest average CTT will be when the LW CRE of the anvil cloud is at its greatest, and so can help understand the evolution of the anvil CRE over its lifetime.

5.4.3 Anvil CRE

Using the broadband fluxes data in conjunction with the tracked DCC dataset enables tracking of how the SW, LW and net CRE evolve over the lifetime of each tracked anvil. Figure 5.13 shows the time series of SW, LW and net CRE as well as the cumulative average CRE (the average of net CRE over anvil area and lifetime up until that point in time) for several different anvil lifecycles. Note that all fluxes are ToA and measured in the downward direction, so a positive value represents warming and a negative value represents cooling.

Figure 5.13 a shows the case of an isolated, short-lived DCC. The DCC initiates during the daytime, during which the SW CRE dominates and the net CRE is negative (cooling). However, towards the end of the four-hour lifecycle of the DCC, it transitions to night-time and so while the SW CRE reduces and eventually becomes zero, the LW CRE dominates and the net CRE is positive (warming). While this period of warming moves the cumulative average CRE towards zero, it remains overall negative for the overall lifetime of the DCC both due to the longer period spent during the daytime, and the larger area of the anvil cloud during this period.

Figure 5.13 b shows the case of a longer-lived (22 hours), clustered DCC. It initiates in the morning, and so the SW cooling dominates for the first half of the anvil lifetime. Compared to the isolated DCC, it exists for much longer during the night time, and so the cumulative average becomes positive over the full lifetime of the anvil cloud.

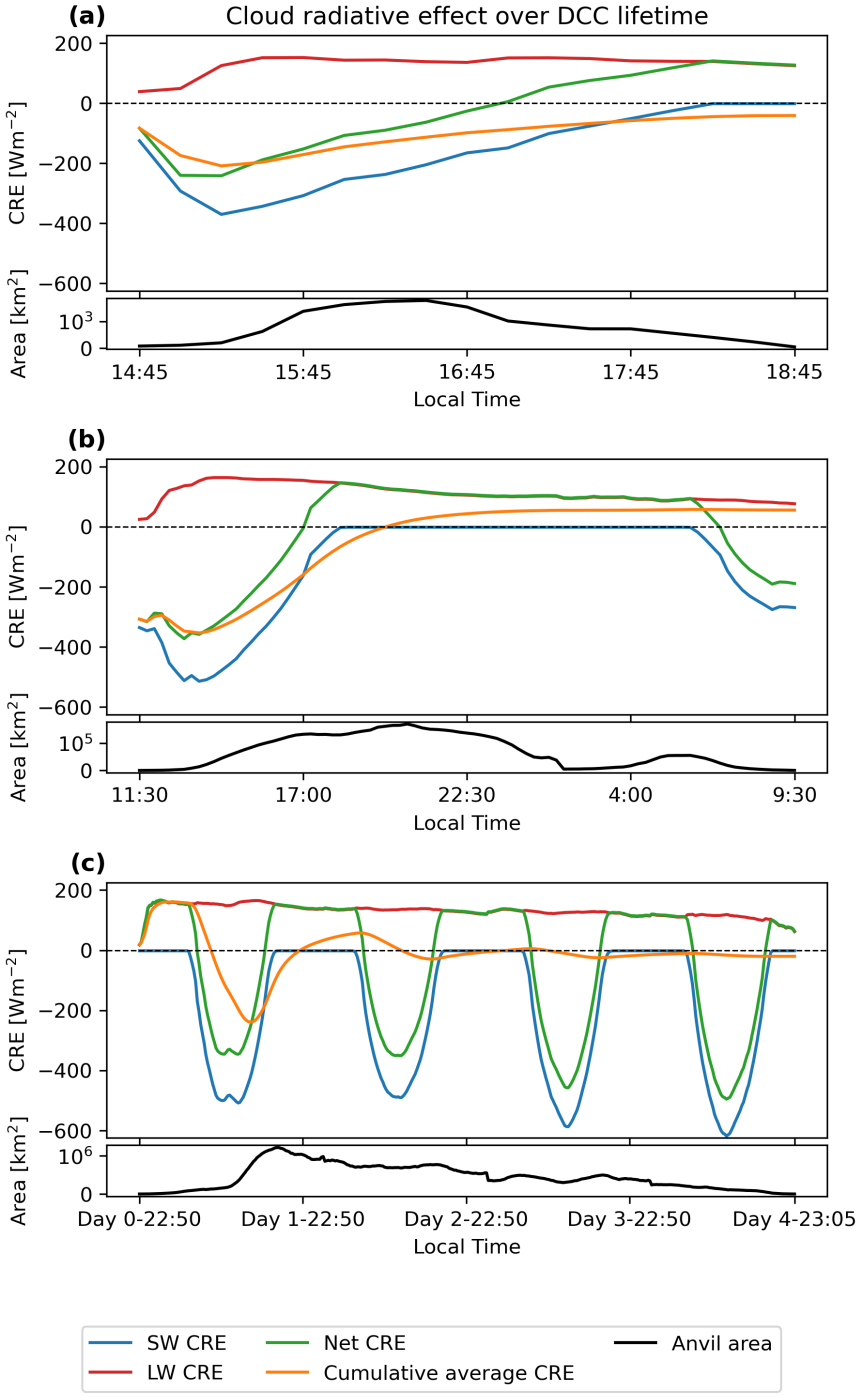


Figure 5.13: Anvil net, LW, and SW CRE, cumulative mean CRE over anvil lifetime and anvil area for a.: an isolated, short-lived (4-hour) DCC; b.: a moderately clustered, 1-day long DCC; and c.: a large, clustered, 4-day long DCC. All times are the local solar time, to the nearest 5-minute interval. The black lines show the change in area of each DCC over their lifecycle.

Figure 5.13 c shows the case of a four-day, highly clustered convective event. In this case, the net CRE alternates between warming and cooling throughout the diurnal cycle. The cumulative CRE also alternates between overall warming and cooling throughout the lifetime of the anvil and results in a small net cooling effect.

In both the longer-lived cases (fig. 5.13 b, c) the LW CRE reduces towards the end of the anvil cloud lifetime. This may be reflective of the findings from fig. 5.11 that the minimum average CTT occurs before the mid-point of the cloud lifecycle for longer-lived systems. This reduction in LW CRE may be due to a thinning of the anvil cloud (allowing increased LW emission from the surface), or due to heating and stabilisation of the upper troposphere by the DCC. In addition, the cumulative radiative cooling of the anvil top may drive subsidence and reduce the cloud-top height of the anvil over time [Sokol and Hartmann, 2020].

Figure 5.14 shows the distribution of net lifetime CRE for all tracked anvils. The overall negative average value of $-0.94 \pm 0.91\,\mathrm{Wm^{-2}}$ is very close to zero considering the large spread in CRE. However, the distribution shows a bimodal structure, with two peaks at around $+100\,\mathrm{Wm^{-2}}$ (warming) and $-180\,\mathrm{Wm^{-2}}$ (cooling). The distribution is coloured according to the mean number of cores associated with the anvils in each bin of the distribution. Both the peaks of the distribution are mainly composed of isolated DCCs which occur during the daytime (negative peak) or night-time (positive peak). The centre of the distribution—with average CREs close to zero—shows a greater number of the clustered DCCs with multiple cores which, due to their longer lifetime, tend to exist during both the day- and night time.

In fig. 5.15 the CRE distribution is broken down into that of the SW (fig. 5.15 a) and LW (fig. 5.15 b) components. The SW CRE shows a similar bimodal distribution to that of the net CRE, whereas the LW distribution shows a normal distribution. The SW CRE has a large peak at $0\,\mathrm{Wm^{-2}}$ for DCCs that occur during the night-time, and a broad peak centred around $-300\,\mathrm{Wm^{-2}}$ consisting of daytime DCCs, with the

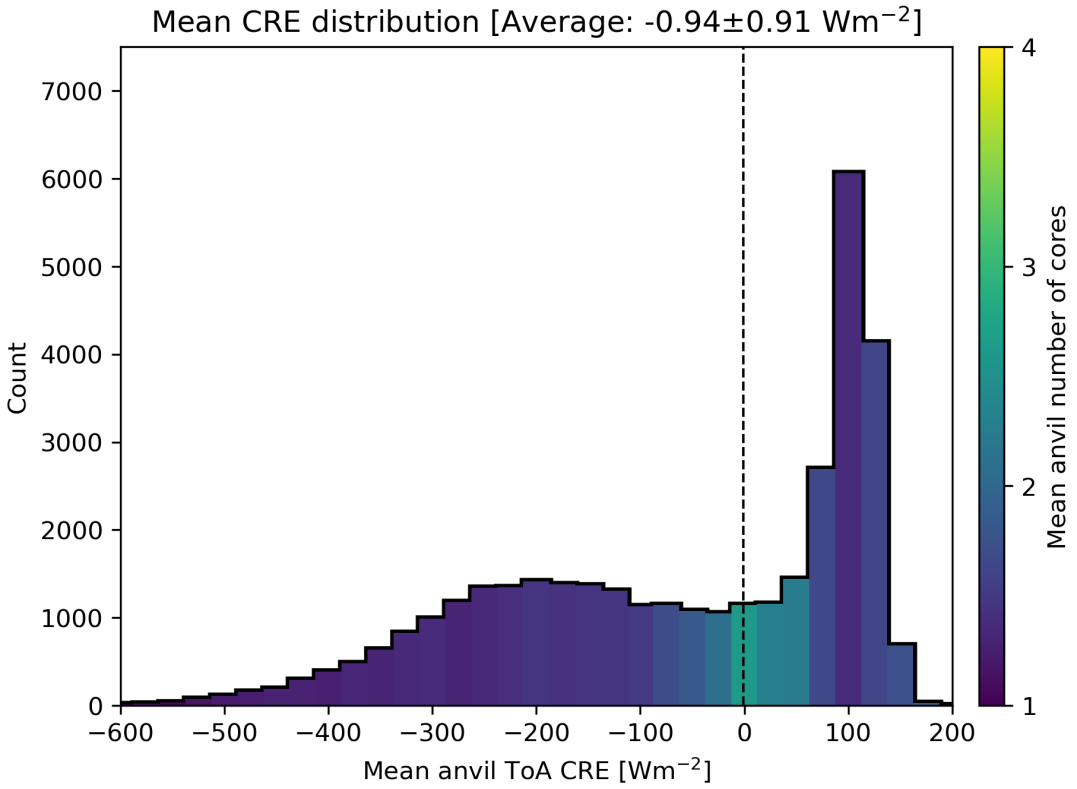


Figure 5.14: The distribution of lifetime anvil CRE for all observed anvils. The mean number of cores per anvil in each bin is indicated by the colour scale. The vertical dashed line shows the integrated mean CRE (over area and lifetime) over all anvils, weighted by the anvil areas ($-0.94 \pm 0.91\,\mathrm{Wm^{-2}}$).

average falling between the two. Note that the average for the LW falls to the right of the peak of the distribution because the average is integrated over the anvil area and lifetime, and the largest and longest-lived anvils tend to have colder CTT and hence larger LW CRE.

Figure 5.16 shows (a) the average instantaneous anvil CRE binned by the time of observation (local solar time) and mean anvil CTT, and (b) the average lifetime anvil CRE binned by time of initial detection (local solar time) and mean anvil CTT. As expected, mean anvil CRE becomes more positive with decreasing CTT due to increased LW warming. However, the diurnal cycle of detection shows a much stronger contrast, with anvils detected during the daytime having a net cooling effect

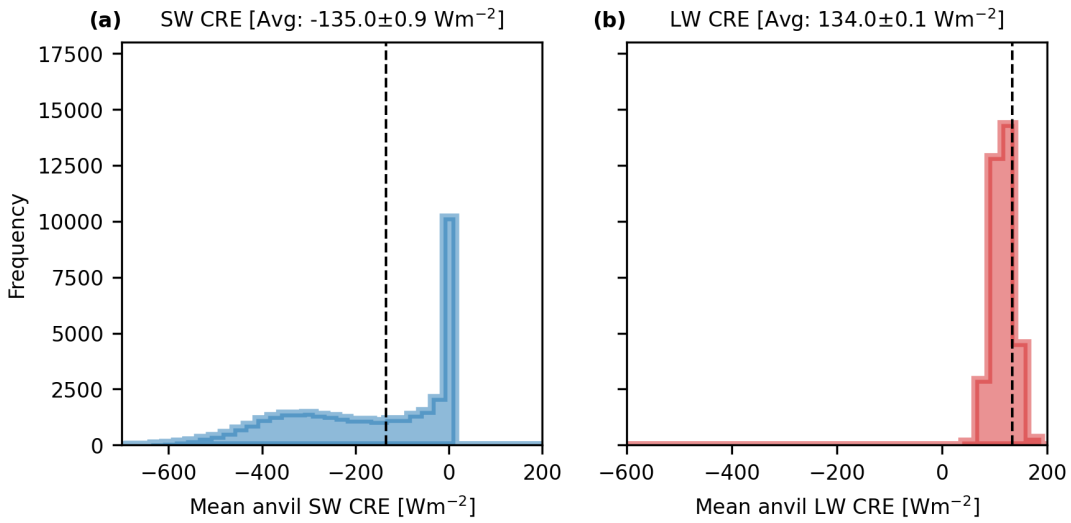


Figure 5.15: The distributions of mean anvil SW CRE (a) and LW CRE (b). The vertical dashed line shows the integrated mean CRE over all anvils (SW: -135.0 ± 0.9 Wm^{-2}, LW: 134.0 ± 0.1 Wm^{-2})

compared to those at night which have a net warming CRE. This diurnal cycle effect is stronger for those anvils with warmer average CTT, generally representing isolated, shorter-lived DCCs, and is weaker for colder anvil CTT. Note also that in fig. 5.16 b that the phase of the diurnal cycle shifts to earlier times of detection as average anvil CTT become colder, as these DCCs tend to have longer lifetimes.

It is apparent from figs. 5.14 and 5.15 that the observed neutral net anvil CRE is not only due to a balance between the SW and LW, but also from a balance of the cooling effect of daytime DCCs and the warming effect of those occurring at night. If the number of DCCs occurring during the daytime were to reduce then there could be a net warming effect without any change to the CRE of individual DCCs. As the diurnal cycle of convection over the ocean is nearly uniform, little impact on anvil CRE should be expected from changes in the time of convective initiation. However, over land, where convective activity is much more common in the afternoon, changes in the diurnal cycle may have a much larger effect on anvil CRE.

Furthermore, fig. 5.16 b highlights that differences in anvil temperature are linked

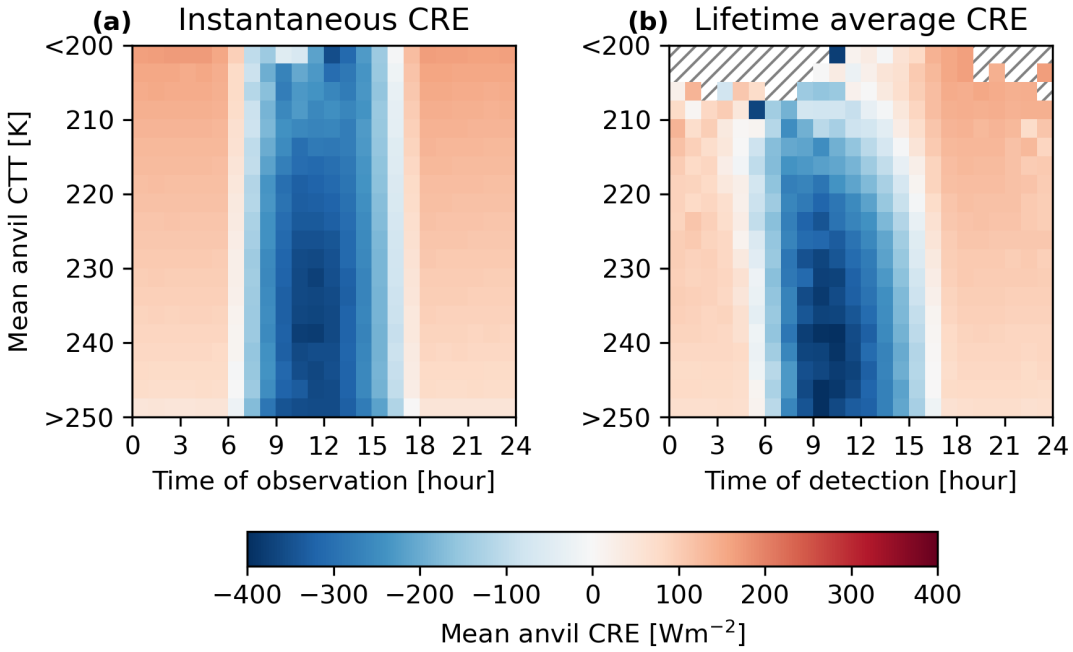


Figure 5.16: (a) Average instantaneous anvil CRE binned by the time of observation (local solar time) and mean anvil CTT. (b) Average lifetime anvil CRE binned by time of initial detection (local solar time) and mean anvil CTT. Hashed regions in (b) show bins in which no anvils were detected.

to the diurnal cycle of anvil CRE as colder anvils tend to have longer lifetimes. As a result, if warming surface temperatures lead to the invigoration of DCCs, the warming effect would be larger than the LW effect from the change in anvil temperature alone. Surface warming may also result in an earlier time of convective initiation, resulting in a cooling feedback.

5.5 Summary

By combining a novel cloud tracking algorithm with a new dataset of derived all-sky and clear-sky fluxes from geostationary satellite observations, DCC anvils were detected and tracked along with their associated cores for both isolated and clustered DCCs, and their properties, lifecycle and CRE investigated. As this study was performed using data from May-August (Northern hemisphere summer), the majority

of convective activity was observed over the Guinea-Congo rainforest and Savanna regions, as the ITCZ is at its northernmost extent.

The degree of convective clustering of each anvil is evaluated by measuring the number of cores it is associated with. As expected, anvils with the greatest number of cores—including MCSs—have larger anvil areas, longer lifetimes and the coldest cloud tops. As a result, despite the majority of observed DCCs being isolated, the highly clustered anvils make up most of the anvil coverage, and so cause most of the anvil impact over this region. In addition, the proportion of the lifecycle spent in the mature and dissipating phases increases with the number of cores, and the proportion spent in the growing phase decreases.

When investigating the net CRE of anvils, although the average CRE across all observed anvils is close to zero, few individual anvils have near zero CRE themselves. There exists a bimodal distribution of anvil CRE, with isolated DCCs that exist during the daytime causing the negative (cooling) peak, and those that exist during the night-time causing the positive (warming) peak. The systems with near zero CRE tend to live longer with more cores, and exist during both the day- and night-time. As a result, when considering the magnitude of the anvil CRE, isolated DCCs have an outsized contribution to the overall average anvil CRE of 21.4% compared to their proportion of all anvil coverage (15.3%) (see fig. 5.17).

The net CRE measured over all anvils ($-0.94 \pm 0.91\,\mathrm{Wm^{-2}}$) is similarly near zero to many previous studies [Ramanathan et al., 1989; Hartmann et al., 1992; Hartmann, 2016, e.g.]. It remains unclear however whether this is an equilibrium state or merely a coincidence. The role of WV in the tropical temperature profile, convective dynamics, the height of the convectively active tropopause and the tropical overturning circulation suggests that there may be a link between these factors. Despite this, no mechanism restoring the tropical anvil CRE to near zero has been discovered. Whether the anvil CRE should itself be considered near-zero is also debated.

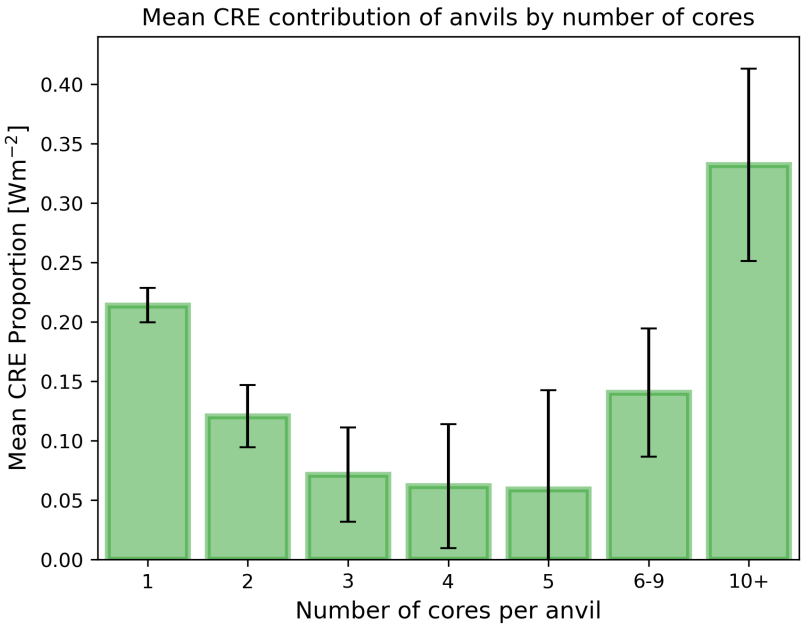


Figure 5.17: The contribution to the net CRE for anvils with differing numbers of cores, which is defined as the sum of the absolute CRE multiplied by anvil area for all anvils with that number of cores, divided by the total for all anvils. Due to the large variance and magnitude of the CRE of isolated DCCs, they have a large impact on the net CRE balance despite their small area.

Stephens et al. [2018] argues that, as tropical DCCs often occur over low-level clouds which have a negative CRE which is cancelled out by the effect of DCCs, meaning that the near zero ToA CRE is actually a result of a positive anvil CRE. This line of reasoning has not become more widespread, however.

The interaction between the diurnal cycle of convection and DCC lifetime plays a key role in the shape of the SW anvil CRE distribution and is important to consider in regard to anvil CRE feedback. As the LW CRE is normally distributed, a response to changing cloud top height or temperature may occur as a shift in the distribution. However, the bimodal distribution of the SW CRE must result in more complex adjustments to shift the overall mean. As the position of the peak at 0 Wm^{-2} relating to night-time DCCs is fixed, to change the overall average SW CRE either the width of the distribution has to increase or decrease, or the number of DCCs

occurring during the day- or night-time has to increase. The former has important implications for the diurnal cycle of temperature in the tropics, and the latter for the diurnal cycle of convection, which, in turn, affects the anvil lifecycle.

Changes in the diurnal cycle of convection may not have a large impact on net anvil CRE over the ocean due to the mostly uniform occurrence of convection throughout the day. Over land, however, the afternoon peak of convection at around 3 pm solar time (see fig. 5.8) coincides with a time at which anvil CRE is very sensitive to shifts in the diurnal cycle (fig. 5.16 b). Furthermore, a reduction or increase in the number of DCCs occurring at a specific time of day may change the net CRE of anvils without any change in the CRE of individual DCCs.

Diagnosing a diurnal-cycle-related anvil cloud feedback in climate models may however be difficult. Beydoun et al. [2021] found that changes in anvil lifetime contributed little to CRE feedbacks in a cloud-resolving radiative-convective-equilibrium model. It is unclear how well the diurnal cycle of convection and convective lifecycle are represented in such a model, although convective-resolving models have been found to model these better than parameterised climate models [Prein et al., 2015; Feng et al., 2023b]. Disentangling the impacts of convective processes and anvil cirrus processes on anvil lifecycle and CRE is also a key challenge. Here, the use of model experiments such as Gasparini et al. [2022] may help understand the impacts of each process on anvil cloud radiative effect (CRE) and the potential climate feedbacks.

Chapter 6

Conclusion

6.1 Summary

DCCs play a key role in the atmosphere, from extreme weather to the global overturning circulation and energy balance. Understanding their response to global warming is vital for climate projection, but DCC feedbacks remain uncertain [Sherwood et al., 2020] and underestimated by current climate models [Hill et al., 2023]. While traditional research on anvil feedbacks has focused on the large scale response of the climate (see section 1.3), a spate of novel research has highlighted the importance of convective processes the CRE of anvils [Raghuraman et al., 2024; Sokol et al., 2024; McKim et al., 2024]. The mechanisms through which convective processes influence anvil feedbacks remain unclear, in part due to a shortage of observational data connecting the properties of anvils to those of their convective cores [Gasparini et al., 2023]. With the growing interest in Lagrangian studies of DCCs [Gasparini et al., 2019; Sokol and Hartmann, 2020; Bouniol et al., 2021], there is a need for improved observational constraints on the properties and behaviours of DCCs. In this thesis, novel cloud tracking techniques have been applied to geostationary satellite imagery to provide a new perspective on the relationship between convective cores and anvils.

Chapter 2 focused on the development of a novel detection and tracking algorithm to better track the behaviour of DCCs. DCC tracking algorithms have been developed for over 60 years, with applications for both weather and climate (see section 1.4).

Despite this, many algorithms still have drawbacks identified decades ago [Augustine and Howard, 1988] that limit their application to developing research areas. Several key requirements were defined to produce tracking data suitable for studying DCCs across a wide range of scales and environments. The algorithm would need to account for both isolated DCCs and MCSs, track both convective cores and anvils over the entirety of their lifetimes, and provide estimates of anvil area that are not dependent on other factors such as the anvil height.

Through theory developed using simple, 1D models (section 2.3), the use of various channel combinations available from modern geostationary imagers was developed to better identify DCCs in observations. To avoid the scaling issue between tracking of isolated systems and MCSs, a semi-Lagrangian scheme was developed that removes the motion dependence from the detection and tracking task. A detection method for growing convective cores based on cloud top cooling rates was developed, providing a good proxy for the vertical velocity of the core. In addition, detection methods for both thick and thin anvils were developed using an edge detection scheme. By avoiding the use of a fixed threshold, this method is more robust to bias than traditional methods, which also enables the detection of anvils both during the mature phase and while dissipating. Evaluation against lightning observations showed that the combination of core and anvil detection provides a good balance of sensitivity and robustness, with an F_1-score of 0.91 providing a high degree of trust in the DCCs detected by the algorithm.

In chapter 3, this tracking algorithm was applied over five years of geostationary satellite observations from GOES-16 ABI, addressing a key shortage of observational data linking core and anvil properties. Analysis of this data shows a number of findings including the spatial distribution and land–sea contrasts of deep convection over North America. Through Lagrangian analysis, the importance of the diurnal cycle both to the properties of DCCs as well as to other studies could be investigated.

In particular, it was found that due to the contrasting diurnal cycles of convection over land and ocean, the average age of observed anvils varies widely throughout the day. This contrast is particularly strong in the early afternoon, when DCCs observed over land tend to be very young while those over the ocean are much older. As this is the observing time for polar orbiting satellite missions such as those of the A-train, it was found that snapshot observations from these sensors may find erroneous differences in the properties of DCCs observed over land and sea due to these contrasting populations at certain times of day. Combining the capabilities of polar-orbiting sensors with cloud tracking, as performed by Elsaesser et al. [2022], appears key to addressing these issues.

Previous research has found that the structure of anvil clouds is related to the convective intensity of their cores [Protopapadaki et al., 2017; Takahashi et al., 2017b]. In chapter 4, the response of anvil cloud structure to both their observed intensity and organisation was evaluated using the dataset produced in the previous chapter. The growing core cooling rate was used as a proxy for convective intensity, and the number of cores as a proxy for organisation, providing more direct proxies than those used in previous studies. In addition, the thin anvil fraction was evaluated over the entire DCC lifetime rather than at a single snapshot, capturing the growth and decay of the anvil.

While the relation between convective intensity and thin anvil fraction was found to be similar to that found in previous studies, increases in organisation had an opposing effect on the anvil structure. The most intense anvils were found to have an increase in the thin anvil fraction of approximately 0.15 over the least intense anvils, whereas the most organised anvils had a fraction 0.12 less than the average for isolated DCCs. While convective organisation has a large impact on both the anvil area and lifetime, intensity was found to have little impact on the thick anvil lifetime and area but does increase that of the thin anvil, similar to findings by Sokol et al. [2024].

Intensity and organisation were also found to have opposing effects on the structure of the DCC lifecycle. While increasing intensity shortened the growing phase of the DCC and lengthened the dissipating phase, organisation had the opposite effect.

To further investigate the cause of these differences, the change in thin anvil fraction over each DCC lifetime was analysed. In more intense anvils, differences were found both during the initial stage of anvil formation with more intense updrafts leading to a thinner anvil (indicating a direct impact of convective dynamics on the anvil structure) as well as during the dissipating phase, where the thin anvils produced by more intense DCCs took longer to dissipate. For more organised anvils, the main differences were observed during the mature phase of the DCC, during which time organised systems showed a thickening of the anvil, possibly due to the impact of mesoscale circulations produced by organised convective systems. While these processes can be observed directly, the use of lifecycle analysis combined with an understanding of how different processes affect DCCs at different stages of their lifetimes provides insight into what effects these processes may have. This lifecycle analysis may also be useful in constraining the behaviour of DCCs in convective resolving models against observations.

The novel approach used in this chapter has allowed new results to be found that were not identified in the previous studies. Explicitly tracking the full lifetime of the anvil cloud results in attributing a larger proportion of thin anvil area than that seen by snapshot observations, and reveals changes caused by an increase in the anvil lifetime. Furthermore, the explicit tracking of convective cores associated with anvils showed the response of thin anvil area to organisation, whereas without this the trends of organised convection would be hidden due to the majority of observed DCCs being isolated systems. In addition, the intensity and organisation of convection are found to have opposing effects on both the lifecycle and structure of DCCs. The thin cirrus produced by DCCs may have a warming effect on the climate, and so understanding

how the area and lifetime of these clouds change with the properties of convection is important to understanding DCC feedbacks. As both the intensity of DCCs and the frequency of organised convective systems are expected to increase with global warming, so further study of the response of thin cirrus to convective activity is vital.

Traditionally, considerations of anvil feedbacks have only considered their bulk properties with a present day net anvil CRE of near zero. In chapter 5, the cloud tracking method developed in chapter 2 is combined with retrieved cloud properties and radiative fluxes from SEVIRI observations over Africa to investigate how the CRE of individual DCCs contributes to the net observed CRE. Together, this allows us to explicitly measure the CRE of DCCs over their lifecycle, and analyse how they change with DCC properties. The distributions of observed DCCs show both similarities and differences from those seen in chapter 3. A similar land–sea contrast was shown in the timing of convective initiation. This contrast was also seen in the time of initiation over the African Great Lakes (particularly Lake Victoria), which may indicate similar mechanisms in action. The small number of DCCs due to the limited extent of the case study prevented further investigation into land–sea contrasts. The observed DCC lifecycle, shown in fig. 5.12, shows notable differences to those found in chapter 4, which may have several reasons. Firstly, there may be a difference in the lifecycle of DCCs between the tropics and extra-tropics. However, this difference may have been due to the lower sensitivity of SEVIRI for observing thin anvils. This lower sensitivity reduces the length of time over which the anvil is observed, increasing the proportion of the lifecycle observed in the growing phase. Furthermore, this effect is expected to be greater for isolated DCCs as the thick anvil dissipates faster than for larger MCSs. Finally, the use of retrieved CTT, rather than observed BT, may cause a difference. An increase in stratospheric WV, along with the thinning of the anvil cloud, will cause the observed BT of the anvil to increase over time. However, the retrieved CTT will remove these effects. As a result, the method of Futyan and

Del Genio [2007] for measuring anvil lifecycle may produce a longer growing phase if CTT is used instead of BT.

The main aim of this chapter was the investigation of how the diurnal cycle affects the CRE of individual DCCs. Historically, studies into the CRE of tropical anvils have averaged over all areas of anvil cloud without tracking individual DCCs [Ramanathan et al., 1989; Hartmann et al., 1992; Stephens et al., 2018]. The overall average CRE of our observed DCCs confirmed the near neutral CRE of tropical anvil clouds with a net value of $-0.94 \pm 0.91\,\mathrm{Wm}^{-2}$. Through the use of tracking, the contribution of individual DCCs to this mean state and the impact of their properties can be studied.

The CRE of DCCs changes over their lifetime both due to the changing properties of the anvil cloud and the timing within the diurnal cycle. While both the LW and SW CRE reduce over the lifetime of the DCC, the SW CRE is also largely dependent on the time within the diurnal cycle. These two effects lead to the bimodal distribution for the CRE of individual DCCs, as shown in fig. 5.14. Three groups of DCCs are identified that contribute to this distribution. The largest group is that of isolated DCCs, which form the two peaks of the distribution. Typically, these either initiate in the early afternoon and have a negative CRE due to the large SW effect during the daytime, or initiate in the early evening and instead have a positive CRE as they exist at night. The second group consists of moderately clustered DCCs which tend to exist for longer than isolated DCCs but less than a day. As these systems tend to initiate in the evening, they mostly exist at night and therefore predominantly have a LW warming effect. Finally, MCSs which exist for multiple days tend to have a neutral CRE as their SW and LW effects average out across the diurnal cycle.

The behaviour of these different DCCs may influence how they respond to climate change. Organised DCCs, including MCSs, may primarily respond through changes in their LW CRE including FAT and anvil area as SW changes tend to average out over their lifecycle. On the other hand, isolated DCCs tend to have CTT warmer than

the top of the convectively active layer (220 K), and so FAT may not apply to them in the same manner. The large variance in the CRE of individual isolated DCCs means that they have a substantially larger influence on the net anvil CRE balance than their area suggests. As a result, the response of these systems to global warming may have a larger effect on the overall anvil feedback than has been previously considered, and, in particular, they may be sensitive to changes in the diurnal cycle.

6.2 Discussion & Future work

In recent years there has been an increased focus on the study of anvil lifecycle [Wall et al., 2018; Sokol and Hartmann, 2020] and radiative effects [Bouniol et al., 2021]. The response of DCCs to climate change is both important and uncertain [Sherwood et al., 2020; Hill et al., 2023]. Improving our understanding of cirrus cloud properties and lifecycles is vital to improving this situation [Sullivan and Voigt, 2021; Gasparini et al., 2023].

The key results of this thesis are in the new observational insights into the lifecycle, structure and CRE of DCCs and their anvils. The development of a novel tracking algorithm allows both isolated and organised DCCs to be tracked across their entire lifetime across a wide range of scales. Furthermore, by tracking both cores and anvil clouds, the behaviours of the convective cores can be linked to the evolution of the anvil clouds, and provide observational evidence for how these properties affect each other.

Chapters 3, and 5 showed significant differences in the lifecycles of DCCs over the land and the ocean. In particular, these differences may influence DCC cores and anvils to different extents. Several processes, occurring over a range of scales, may influence the properties of observed DCCs. Lake effects, such as those seen over the African Great Lakes in chapter 5, may provide a useful case study for understanding these differences, as while their local and surface properties may be similar to that

of oceans, the larger environmental conditions may be more affected by the adjacent land regions.

In chapter 4 contrasting impacts of convective intensity and organisation on the structure and lifecycle of anvil clouds were found. Analysing how the structure of these DCCs evolved over their lifetime provided key insights into which convective and anvil processes lead to the observed changes in thick thin anvils. These contrasts may help further understand the uncertainties in anvil radiative feedbacks, and, in particular, help explain how changes in convection may affect anvil optical properties [McKim et al., 2024]. In addition, it highlights the importance of isolated DCCs to the total anvil CRE, motivating further study into how their behaviour differs from organised convection. Further investigating the contrast between the effects of intensification and organisation on the structure of DCCs anvils may provide insight into the net changes of anvil structure with climate change. However, while with these structural changes alone, more intense DCCs are expected to be warming, and more organised DCCs to be cooling, the changes in BT and the diurnal cycle observed for both of these processes may dominate any changes in CRE.

A key area of opportunity is in combining the investigation of the impacts of intensity and organisation on anvil structure in chapter 4 with the investigation of anvil CRE in chapter 5. While the potential impact of structural changes on anvil CRE was discussed in chapter 4, explicitly resolving the impact on CRE and comparing it to that due to changes in anvil height and diurnal cycle would advance our understanding of these processes collectively. To do so, the issues with tracking thin anvils in SEVIRI observations would need to be overcome. A larger cloud tracking study, combining these elements, investigating how the microphysical, macrophysical properties and structure of DCCs change over their lifecycle and throughout the diurnal cycle, and then demonstrating the effect of these changes on the CRE of the anvil across both thick and thin parts and along its entire lifetime, could greatly reduce

uncertainties in anvil cloud feedbacks.

DCCs play a key role in many atmospheric processes, and their response to global warming is both uncertain [Sherwood et al., 2020] and underestimated by current climate models [Hill et al., 2023]. While the new generation of convective resolving models may help to address this problem [Stevens et al., 2020], these models still do not fully resolve the issues in simulating the dynamics of deep convection [Jeevanjee, 2017] and the behaviour of anvil cirrus [Sullivan and Voigt, 2021]. Observational studies of DCCs using cloud tracking approaches have been highlighted as a key method to addressing some of these issues [Gasparini et al., 2023]. In this thesis, we have developed and utilised a novel cloud tracking algorithm to better investigate the properties of DCC cores and anvils as seen by the latest generation of geostationary weather satellites. Key results from chapters 3, 4 and 5 have shown new findings on the properties, structure and CRE of DCC anvils across their lifecycle, which provide insight into how DCCs may respond to changes in convective processes and climate change. The ability to track the properties of DCCs over their entire lifetimes has been key to these findings. Understanding the lifecycle of DCCs from a Lagrangian perspective has been highlighted as a key need for future satellite missions [van den Heever et al., 2023], convective resolving model development [Prein et al., 2024], and upcoming observation campaigns, and so cloud tracking has a key role to play in future studies of deep convection.

Bibliography

Abercromby, R. (1887). On the identity of cloud forms all over the world; and on the general principles by which their indications must be read. *Quarterly Journal of the Royal Meteorological Society*, 13(62):140–146.

Adams, D. K. and Comrie, A. C. (1997). The North American Monsoon. *Bulletin of the American Meteorological Society*, 78(10):2197–2214.

Agard, V. and Emanuel, K. (2017). Clausius–Clapeyron Scaling of Peak CAPE in Continental Convective Storm Environments. *Journal of the Atmospheric Sciences*, 74(9):3043–3054.

Aggarwal, J. and Nandhakumar, N. (1988). On the computation of motion from sequences of images-A review. *Proceedings of the IEEE*, 76(8):917–935.

Alduchov, O. A. and Eskridge, R. E. (1997). Improved Magnus‘ form approximation of saturation vapor pressure. Technical Report DOE/ER/61011-T6, Department of Commerce, Asheville, NC (United States).

Allen, M. R. and Ingram, W. J. (2002). Constraints on future changes in climate and the hydrologic cycle. *Nature*, 419(6903):228–232.

Aminou, D. M. A. (2002). MSG's SEVIRI Instrument. *ESA bulletin*, 111.

Apke, J. M., Mecikalski, J. R., Bedka, K., McCaul, E. W., Homeyer, C. R., and Jewett, C. P. (2018). Relationships between Deep Convection Updraft Characteristics and Satellite-Based Super Rapid Scan Mesoscale Atmospheric Motion Vector–Derived Flow. *Monthly Weather Review*, 146(10):3461–3480.

Arnaud, Y., Desbois, M., and Maizi, J. (1992). Automatic Tracking and Characterization of African Convective Systems on Meteosat Pictures. *Journal of Applied Meteorology and Climatology*, 31(5):443–453.

Augustine, J. A. and Howard, K. W. (1988). Mesoscale Convective Complexes over the United States during 1985. *Monthly Weather Review*, 116(3):685–701.

Augustine, J. A. and Howard, K. W. (1991). Mesoscale Convective Complexes over the United States during 1986 and 1987. *Monthly Weather Review*, 119(7):1575–1589.

Austin, P. M. (1987). Relation between Measured Radar Reflectivity and Surface Rainfall. *Monthly Weather Review*, 115(5):1053–1070.

Baker, S., Scharstein, D., Lewis, J. P., Roth, S., Black, M. J., and Szeliski, R. (2011). A Database and Evaluation Methodology for Optical Flow. *Int J Comput Vis*, 92(1):1–31.

Bechini, R. and Chandrasekar, V. (2017). An Enhanced Optical Flow Technique for Radar Nowcasting of Precipitation and Winds. *J. Atmos. Oceanic Technol.*, 34(12):2637–2658.

Bedka, K., Brunner, J., Dworak, R., Feltz, W., Otkin, J., and Greenwald, T. (2010). Objective Satellite-Based Detection of Overshooting Tops Using Infrared Window Channel Brightness Temperature Gradients. *Journal of Applied Meteorology and Climatology*, 49(2):181–202.

Bedka, K. M. and Mecikalski, J. R. (2005). Application of Satellite-Derived Atmospheric Motion Vectors for Estimating Mesoscale Flows. *J. Appl. Meteor.*, 44(11):1761–1772.

Bennartz, R. and Schroeder, M. (2012). Convective Activity over Africa and the Tropical Atlantic Inferred from 20 Years of Geostationary Meteosat Infrared Observations. *Journal of Climate*, 25(1):156–169.

Berry, E. and Mace, G. G. (2014). Cloud properties and radiative effects of the Asian summer monsoon derived from A-Train data. *Journal of Geophysical Research: Atmospheres*, 119(15):9492–9508.

Beucler, T. and Cronin, T. (2019). A budget for the size of convective self-aggregation. *Quarterly Journal of the Royal Meteorological Society*, 145(720):947–966.

Beydoun, H., Caldwell, P. M., Hannah, W. M., and Donahue, A. S. (2021). Dissecting Anvil Cloud Response to Sea Surface Warming. *Geophysical Research Letters*, 48(15):e2021GL094049.

Bieniek, A. and Moga, A. (2000). An efficient watershed algorithm based on connected components. *Pattern Recognition*, 33(6):907–916.

Boehm, M. T., Verlinde, J., and Ackerman, T. P. (1999). On the maintenance of high tropical cirrus. *Journal of Geophysical Research: Atmospheres*, 104(D20):24423–24433.

Boer, E. R. and Ramanathan, V. (1997). Lagrangian approach for deriving cloud characteristics from satellite observations and its implications to cloud parameterization. *Journal of Geophysical Research: Atmospheres*, 102(D17):21383–21399.

Bony, S., Semie, A., Kramer, R. J., Soden, B., Tompkins, A. M., and Emanuel, K. A. (2020). Observed Modulation of the Tropical Radiation Budget by Deep Convective Organization and Lower-Tropospheric Stability. *AGU Advances*, 1(3):e2019AV000155.

Bony, S., Stevens, B., Coppin, D., Becker, T., Reed, K. A., Voigt, A., and Medeiros, B. (2016). Thermodynamic control of anvil cloud amount. *Proceedings of the National Academy of Sciences*, 113(32):8927–8932.

Bony, S., Stevens, B., Frierson, D. M. W., Jakob, C., Kageyama, M., Pincus, R., Shepherd, T. G., Sherwood, S. C., Siebesma, A. P., Sobel, A. H., Watanabe, M., and Webb, M. J. (2015). Clouds, circulation and climate sensitivity. *Nature Geoscience*, 8(4):261–268.

Bouguet, J. (1999). Pyramidal implementation of the lucas kanade feature tracker.

Bouniol, D., Roca, R., Fiolleau, T., and Poan, D. E. (2016). Macrophysical, Microphysical, and Radiative Properties of Tropical Mesoscale Convective Systems over Their Life Cycle. *Journal of Climate*, 29(9):3353–3371. Publisher: American Meteorological Society Section: Journal of Climate.

Bouniol, D., Roca, R., Fiolleau, T., and Raberanto, P. (2021). Life Cycle–Resolved Observation of Radiative Properties of Mesoscale Convective Systems. *Journal of Applied Meteorology and Climatology*, 60(8):1091–1104. Publisher: American Meteorological Society Section: Journal of Applied Meteorology and Climatology.

Bowler, N. E. H., Pierce, C. E., and Seed, A. (2004). Development of a precipitation nowcasting algorithm based upon optical flow techniques. *Journal of Hydrology*, 288(1):74–91.

Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.

Bresky, W. and Daniels, J. (2006). The feasibility of an optical flow algorithm for estimating atmospheric motion. In *Proceedings of the Eight International Winds Workshop*, Beijing, China. EUMETSAT.

Bretherton, C. S., Blossey, P. N., and Khairoutdinov, M. (2005). An Energy-Balance Analysis of Deep Convective Self-Aggregation above Uniform SST. *Journal of the Atmospheric Sciences*, 62(12):4273–4292.

Brooks, H. E., Anderson, A. R., Riemann, K., Ebbers, I., and Flachs, H. (2007). Climatological aspects of convective parameters from the NCAR/NCEP reanalysis. *Atmospheric Research*, 83(2):294–305.

Brooks, H. E., Iii, C. A. D., Zhang, X., Chernokulsky, A. M. A., Tochimoto, E., Hanstrum, B., Nascimento, E. d. L., Sills, D. M. L., Antonescu, B., and Barrett, B. (2019). A Century of Progress in Severe Convective Storm Research and Forecasting. *Meteorological Monographs*, 59(1):18.1–18.41.

Brox, T., Bruhn, A., Papenberg, N., and Weickert, J. (2004). High Accuracy Optical Flow Estimation Based on a Theory for Warping. volume 3024, pages 25–36, Berlin, Heidelberg. Springer Berlin Heidelberg.

Bruning, E. C. and MacGorman, D. R. (2013). Theory and Observations of Controls on Lightning Flash Size Spectra. *J. Atmos. Sci.*, 70(12):4012–4029.

Bukowski, J. and van den Heever, S. C. (2021). Direct Radiative Effects in Haboobs. *Journal of Geophysical Research: Atmospheres*, 126(21):e2021JD034814.

Buras, R., Dowling, T., and Emde, C. (2011). New secondary-scattering correction in DISORT with increased efficiency for forward scattering. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 112(12):2028–2034.

Burrows, S. M., McCluskey, C. S., Cornwell, G., Steinke, I., Zhang, K., Zhao, B., Zawadowicz, M., Raman, A., Kulkarni, G., China, S., Zelenyuk, A., and DeMott, P. J. (2022). Ice-Nucleating Particles That Impact Clouds and Climate: Observational and Modeling Research Needs. *Reviews of Geophysics*, 60(2):e2021RG000745.

Carvalho, L. M. V. and Jones, C. (2001). A Satellite Method to Identify Structural Properties of Mesoscale Convective Systems Based on the Maximum Spatial Correlation Tracking Technique (MASCOTTE). *Journal of Applied Meteorology and Climatology*, 40(10):1683–1701.

Ceppi, P., Brient, F., Zelinka, M. D., and Hartmann, D. L. (2017). Cloud feedback mechanisms and their representation in global climate models. *WIREs Climate Change*, 8(4):e465.

Chang, Y. M., Fujita, T. T., and Tecson, J. J. (1973). METRACOM System of Cloud Velocity Determination from Geostationary Satellite Pictures. *SMRP Research Paper*, (110).

Chen, S. S. and Houze, R. A. (1997). Diurnal variation and life-cycle of deep convective systems over the tropical pacific warm pool. *Quarterly Journal of the Royal Meteorological Society*, 123(538):357–388.

Christensen, M. W., Jones, W. K., and Stier, P. (2020). Aerosols enhance cloud lifetime and brightness along the stratus-to-cumulus transition. *Proceedings of the National Academy of Sciences*, 117(30):17591–17598.

Christian, H. J., Blakeslee, R. J., Boccippio, D. J., Boeck, W. L., Buechler, D. E., Driscoll, K. T., Goodman, S. J., Hall, J. M., Koshak, W. J., Mach, D. M., and Stewart, M. F. (2003). Global frequency and distribution of lightning as observed from space by the Optical Transient Detector. *Journal of Geophysical Research: Atmospheres*, 108(D1):ACL 4–1–ACL 4–15.

Churchill, D. D. and Houze, R. A. (1984). Development and Structure of Winter Monsoon Cloud Clusters On 10 December 1978. *Journal of the Atmospheric Sciences*, 41(6):933–960.

Clark, A. J., Bullock, R. G., Jensen, T. L., Xue, M., and Kong, F. (2014). Application of Object-Based Time-Domain Diagnostics for Tracking Precipitation Systems in Convection-Allowing Models. *Weather and Forecasting*, 29(3):517–542.

Coniglio, M. C., Hwang, J. Y., and Stensrud, D. J. (2010). Environmental Factors in the Upscale Growth and Longevity of MCSs Derived from Rapid Update Cycle Analyses. *Monthly Weather Review*, 138(9):3514–3539.

Crane, R. K. (1979). Automatic Cell Detection and Tracking. *IEEE Transactions on Geoscience Electronics*, 17(4):250–262.

Crum, T. D. and Alberty, R. L. (1993). The WSR-88D and the WSR-88D Operational Support Facility. *Bulletin of the American Meteorological Society*, 74(9):1669–1688.

Daniels, J., Velden, C., Bresky, W., Genkova, I., and Wanzong, S. (2016). *Algorithm and Software Development of Atmospheric Motion Vector (amv) Products for the Future Goes-R Advanced Baseline Imager (abi).*

Dawe, J. T. and Austin, P. H. (2012). Statistical analysis of an LES shallow cumulus cloud ensemble using a cloud tracking algorithm. *Atmospheric Chemistry and Physics*, 12(2):1101–1119.

Del Genio, A. D. and Kovari, W. (2002). Climatic Properties of Tropical Precipitating Convection under Varying Environmental Conditions. *Journal of Climate*, 15(18):2597–2615.

Dim, J. R. and Takamura, T. (2013). Alternative Approach for Satellite Cloud Classification: Edge Gradient Application. *Advances in Meteorology*.

Dixon, M. and Wiener, G. (1993). TITAN: Thunderstorm Identification, Tracking, Analysis, and Nowcasting—A Radar-based Methodology. *Journal of Atmospheric and Oceanic Technology*, 10(6):785–797.

Douglas, M. W., Maddox, R. A., Howard, K., and Reyes, S. (1993). The Mexican Monsoon. *Journal of Climate*, 6(8):1665–1677.

Dubuisson, P., Giraud, V., Pelon, J., Cadet, B., and Yang, P. (2008). Sensitivity of Thermal Infrared Radiation at the Top of the Atmosphere and the Surface to Ice Cloud Microphysics. *Journal of Applied Meteorology and Climatology*, 47(10):2545–2560.

Eastman, R. and Wood, R. (2018). The Competing Effects of Stability and Humidity on Subtropical Stratocumulus Entrainment and Cloud Evolution from a Lagrangian Perspective. *Journal of the Atmospheric Sciences*, 75(8):2563–2578.

Elsaesser, G. S., Roca, R., Fiolleau, T., Del Genio, A. D., and Wu, J. (2022). A Simple Model for Tropical Convective Cloud Shield Area Growth and Decay Rates Informed by Geostationary IR, GPM, and Aqua/AIRS Satellite Data. *Journal of Geophysical Research: Atmospheres*, 127(10).

Emanuel, K. A. (1983). The lagrangian parcel dynamics of moist symmetric instability. *J. Atmos. Sci*, 40(10):2368–2376.

Emanuel, K. A. (1994). *Atmospheric Convection*. Oxford University Press.

Emde, C., Buras-Schnell, R., Kylling, A., Mayer, B., Gasteiger, J., Hamann, U., Kylling, J., Richter, B., Pause, C., Dowling, T., and Bugliaro, L. (2016). The libRadtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development*, 9(5):1647–1672.

Endlich, R. M. and Wolf, D. E. (1981). Automatic Cloud Tracking Applied to GOES and METEOSAT Observations. *Journal of Applied Meteorology and Climatology*, 20(3):309–319.

Endlich, R. M., Wolf, D. E., Hall, D. J., and Brain, A. E. (1971). Use of a Pattern Recognition Technique for Determining Cloud Motions from Sequences of Satellite Photographs. *Journal of Applied Meteorology and Climatology*, 10(1):105–117.

Evans, J. L. and Shemo, R. E. (1996). A Procedure for Automated Satellite-Based Identification and Climatology Development of Various Classes of Organized Convection. *Journal of Applied Meteorology and Climatology*, 35(5):638–652.

Farfán, L. M. and Zehnder, J. A. (1994). Moving and Stationary Mesoscale Convective Systems over Northwest Mexico during the Southwest Area Monsoon Project. *Weather and Forecasting*, 9(4):630–639.

Farnebäck, G. (2003). Two-Frame Motion Estimation Based on Polynomial Expansion. In Bigun, J. and Gustavsson, T., editors, *Image Analysis*, Lecture Notes in Computer Science, pages 363–370, Berlin, Heidelberg. Springer.

Feingold, G., Cotton, W. R., Kreidenweis, S. M., and Davis, J. T. (1999). The Impact of Giant Cloud Condensation Nuclei on Drizzle Formation in Stratocumulus: Implications for Cloud Radiative Properties. *Journal of the Atmospheric Sciences*, 56(24):4100–4117.

Feng, Z., Hagos, S., Rowe, A. K., Burleyson, C. D., Martini, M. N., and de Szoeke, S. P. (2015). Mechanisms of convective cloud organization by cold pools over tropical warm ocean during the AMIE/DYNAMO field campaign. *Journal of Advances in Modeling Earth Systems*, 7(2):357–381.

Feng, Z., Hardin, J., Barnes, H. C., Li, J., Leung, L. R., Varble, A., and Zhang, Z. (2023a). PyFLEXTRKR: a flexible feature tracking Python software for convective cloud analysis. *Geoscientific Model Development*, 16(10):2753–2776.

Feng, Z., Houze, R. A., Leung, L. R., Song, F., Hardin, J. C., Wang, J., Gustafson, W. I., and Homeyer, C. R. (2019). Spatiotemporal Characteristics and Large-Scale Environments of Mesoscale Convective Systems East of the Rocky Mountains. *Journal of Climate*, 32(21):7303–7328.

Feng, Z., Leung, L. R., Hardin, J., Terai, C. R., Song, F., and Caldwell, P. (2023b). Mesoscale Convective Systems in DYAMOND Global Convection-Permitting Simulations. *Geophysical Research Letters*, 50(4):e2022GL102603.

Feng, Z., Leung, L. R., Liu, N., Wang, J., Houze Jr, R. A., Li, J., Hardin, J. C., Chen, D., and Guo, J. (2021). A Global High-Resolution Mesoscale Convective System Database Using Satellite-Derived Cloud Tops, Surface Precipitation, and Tracking. *Journal of Geophysical Research: Atmospheres*, 126(8):e2020JD034202.

Field, P. R., Lawson, R. P., Brown, P. R. A., Lloyd, G., Westbrook, C., Moisseev, D., Miltenberger, A., Nenes, A., Blyth, A., Choularton, T., Connolly, P., Buehl, J., Crosier, J., Cui, Z., Dearden, C., DeMott, P., Flossmann, A., Heymsfield, A., Huang, Y., Kalesse, H., Kanji, Z. A., Korolev, A., Kirchgaessner, A., Lasher-Trapp, S., Leisner, T., McFarquhar, G., Phillips, V., Stith, J., and Sullivan, S. (2017). Secondary Ice Production: Current State of the Science and Recommendations for the Future. *Meteorological Monographs*, 58(1):7.1–7.20.

Fiolleau, T. and Roca, R. (2013). An Algorithm for the Detection and Tracking of Tropical Mesoscale Convective Systems Using Infrared Images From Geostationary Satellite. *IEEE Transactions on Geoscience and Remote Sensing*, 51(7):4302–4315.

Fu, Q. (1996). An Accurate Parameterization of the Solar Radiative Properties of Cirrus Clouds for Climate Models. *Journal of Climate*, 9(9):2058–2082.

Fu, Q. (2015). RADIATION TRANSFER IN THE ATMOSPHERE | Cloud-Radiative Processes. In North, G. R., Pyle, J., and Zhang, F., editors, *Encyclopedia of Atmospheric Sciences (Second Edition)*, pages 13–15. Academic Press, Oxford.

Fu, Q., Yang, P., and Sun, W. B. (1998). An Accurate Parameterization of the Infrared Radiative Properties of Cirrus Clouds for Climate Models. *Journal of Climate*, 11(9):2223–2237.

Fueglistaler, S., Dessler, A. E., Dunkerton, T. J., Folkins, I., Fu, Q., and Mote, P. W. (2009). Tropical tropopause layer. *Reviews of Geophysics*, 47(1).

Fujita, T. T. (1969). Present Status of Cloud Velocity Computations from the ATS I and ATS III Satellites. *Cospar Space Research*, IX.

Fujita, T. T., Bradbury, D. L., Murino, C., and Hull, L. (1968). A Study of Mesoscale Cloud Motions Computed from ATS-I and Terrestrial Photographs. *SMRP Research Paper*, (71).

Fujita, T. T., Fearl, E. W., and Shenk, W. E. (1975). Satellite-Tracked Cumulus Velocities. *Journal of Applied Meteorology (1962-1982)*, 14(3):407–413.

Futyan, J. M. and Del Genio, A. D. (2007). Deep Convective System Evolution over Africa and the Tropical Atlantic. *Journal of Climate*, 20(20):5041–5060.

Gasparini, B., Blossey, P. N., Hartmann, D. L., Lin, G., and Fan, J. (2019). What Drives the Life Cycle of Tropical Anvil Clouds? *Journal of Advances in Modeling Earth Systems*, 11(8):2586–2605.

Gasparini, B., Sokol, A. B., Wall, C. J., Hartmann, D. L., and Blossey, P. N. (2022). Diurnal Differences in Tropical Maritime Anvil Cloud Evolution. *Journal of Climate*, 35(5):1655–1677.

Gasparini, B., Sullivan, S. C., Sokol, A. B., Kärcher, B., Jensen, E., and Hartmann, D. L. (2023). Opinion: Tropical cirrus – From micro-scale processes to climate-scale impacts. *EGUsphere*, pages 1–47.

Gasteiger, J., Emde, C., Mayer, B., Buras, R., Buehler, S. A., and Lemke, O. (2014). Representative wavelengths absorption parameterization applied to satellite channels and spectral bands. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 148:99–115.

Ge, J., Hu, X., Mu, Q., Liu, B., Zhu, Z., Du, J., Su, J., Li, Q., and Zhang, C. (2024). Contrasting characteristics of continental and oceanic deep convective systems at different life stages from CloudSat observations. *Atmospheric Research*, 298:107157.

Goodman, S. J., Blakeslee, R. J., Koshak, W. J., Mach, D., Bailey, J., Buechler, D., Carey, L., Schultz, C., Bateman, M., McCaul, E., and Stano, G. (2013). The GOES-R Geostationary Lightning Mapper (GLM). *Atmospheric Research*, 125-126:34–49.

Grabowski, W. W. (2018). Can the Impact of Aerosols on Deep Convection be Isolated from Meteorological Effects in Atmospheric Observations? *Journal of the Atmospheric Sciences*, 75(10):3347–3363.

Grant, L. D. and van den Heever, S. C. (2016). Cold pool dissipation. *Journal of Geophysical Research: Atmospheres*, 121(3):1138–1155.

Gray, W. M. and Jacobson, R. W. (1977). Diurnal Variation of Deep Cumulus Convection. *Monthly Weather Review*, 105(9):1171–1188.

Gryspeerdt, E. (2024). pyLRT: A simple python interface/wrapper for LibRadTran.

Gunshor, M. M., Schmit, T. J., Pogorzala, D. R., Lindstrom, S. S., and Nelson, J. P. (2020). GOES-R series ABI Imagery artifacts. *Journal of Applied Remote Sensing*, 14(3):032411.

Hallett, J. and Mossop, S. C. (1974). Production of secondary ice particles during the riming process. *Nature*, 249(5452):26–28.

Han, L., Fu, S., Zhao, L., Zheng, Y., Wang, H., and Lin, Y. (2009). 3D Convective Storm Identification, Tracking, and Forecasting—An Enhanced TITAN Algorithm. *Journal of Atmospheric and Oceanic Technology*, 26(4):719–732.

Handwerker, J. (2002). Cell tracking with TRACE3D—a new algorithm. *Atmospheric Research*, 61(1):15–34.

Hansen, J. E. and Travis, L. D. (1974). Light scattering in planetary atmospheres. *Space Science Reviews*, 16(4):527–610.

Hartmann, D. L. (2016). Tropical anvil clouds and climate sensitivity. *PNAS*, 113(32):8897–8899.

Hartmann, D. L. and Larson, K. (2002). An important constraint on tropical cloud - climate feedback. *Geophysical Research Letters*, 29(20):12–1–12–4.

Hartmann, D. L., Ockert-Bell, M. E., and Michelsen, M. L. (1992). The Effect of Cloud Type on Earth's Energy Balance: Global Analysis. *Journal of Climate*, 5(11):1281–1304.

Hartung, D. C., Sieglaff, J. M., Cronce, L. M., and Feltz, W. F. (2013). An Intercomparison of UW Cloud-Top Cooling Rates with WSR-88D Radar Data. *Weather and Forecasting*, 28(2):463–480.

Hayden, L., Liu, C., and Liu, N. (2021). Properties of Mesoscale Convective Systems Throughout Their Lifetimes Using IMERG, GPM, WWLLN, and a Simplified Tracking Algorithm. *Journal of Geophysical Research: Atmospheres*, 126(20):e2021JD035264.

Heidinger, A. K. and Pavolonis, M. J. (2009). Gazing at Cirrus Clouds for 25 Years through a Split Window. Part I: Methodology. *Journal of Applied Meteorology and Climatology*, 48(6):1100–1116.

Heidinger, A. K., Pavolonis, M. J., Calvert, C., Hoffman, J., Nebuda, S., Straka, W., Walther, A., and Wanzong, S. (2020). Chapter 6 - ABI Cloud Products from the GOES-R Series. In Goodman, S. J., Schmit, T. J., Daniels, J., and Redmon, R. J., editors, *The GOES-R Series*, pages 43–62. Elsevier.

Heikenfeld, M., Marinescu, P. J., Christensen, M., Watson-Parris, D., Senf, F., van den Heever, S. C., and Stier, P. (2019). tobac 1.2: towards a flexible framework for tracking and analysis of clouds in diverse datasets. *Geoscientific Model Development*, 12(11):4551–4570.

Held, I. M. and Soden, B. J. (2006). Robust Responses of the Hydrological Cycle to Global Warming. *J. Climate*, 19(21):5686–5699.

Hendon, H. H. and Woodberry, K. (1993). The diurnal cycle of tropical convection. *Journal of Geophysical Research: Atmospheres*, 98(D9):16623–16637.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.

Heus, T. and Seifert, A. (2013). Automated tracking of shallow cumulus clouds in large domain, long duration large eddy simulations. *Geoscientific Model Development*, 6(4):1261–1273.

Heymsfield, A. J. and Miloshevich, L. M. (1993). Homogeneous Ice Nucleation and Supercooled Liquid Water in Orographic Wave Clouds. *Journal of the Atmospheric Sciences*, 50(15):2335–2353.

Higgins, R. W. and Shi, W. (2001). Intercomparison of the Principal Modes of Interannual and Intraseasonal Variability of the North American Monsoon System. *Journal of Climate*, 14(3):403–417.

Hildebrandsson, H. H. (1887). Remarks concerning the nomenclature of clouds for ordinary use. *Quarterly Journal of the Royal Meteorological Society*, 13(62):148–154. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/qj.4970136211.

Hill, P. G., Holloway, C. E., Byrne, M. P., Lambert, F. H., and Webb, M. J. (2023). Climate Models Underestimate Dynamic Cloud Feedbacks in the Tropics. *Geophysical Research Letters*, 50(15):e2023GL104573.

Hodges, K. I. (1994). A General Method for Tracking Analysis and Its Application to Meteorological Data. *Monthly Weather Review*, 122(11):2573–2586.

Hodges, K. I. (1995). Feature Tracking on the Unit Sphere. *Monthly Weather Review*, 123(12):3458–3465.

Hodges, K. I. and Thorncroft, C. D. (1997). Distribution and Statistics of African Mesoscale Convective Weather Systems Based on the ISCCP Meteosat Imagery. *Monthly Weather Review*, 125(11):2821–2837.

Holloway, C. E., Wing, A. A., Bony, S., Muller, C., Masunaga, H., L'Ecuyer, T. S., Turner, D. D., and Zuidema, P. (2017). Observing Convective Aggregation. *Surveys in Geophysics*, 38(6):1199–1236.

Hong, G., Heygster, G., Miao, J., and Kunzi, K. (2005). Detection of tropical deep convective clouds from AMSU-B water vapor channels measurements. *Journal of Geophysical Research: Atmospheres*, 110(D5).

Hoose, C. and Möhler, O. (2012). Heterogeneous ice nucleation on atmospheric aerosols: A review of results from laboratory experiments. *Atmospheric Chemistry and Physics*, 12(20):9817–9854.

Horner, G. and Gryspeerdt, E. (2023). The evolution of deep convective systems and their associated cirrus outflows. *Atmospheric Chemistry and Physics*, 23(22):14239–14253.

Hoshen, J. and Kopelman, R. (1976). Percolation and cluster distribution. I. Cluster multiple labeling technique and critical concentration algorithm. *Phys. Rev. B*, 14(8):3438–3445.

Houze, R. A. (2004). Mesoscale convective systems. *Reviews of Geophysics*, 42(4).

Houze, R. A. (2014). Chapter 8 - Cumulonimbus and Severe Storms. In Houze, R. A., editor, *International Geophysics*, volume 104 of *Cloud Dynamics*, pages 187–236. Academic Press.

Houze, R. A. (2018). 100 Years of Research on Mesoscale Convective Systems. *Meteorological Monographs*, 59(1):17.1–17.54.

Hu, Y. X. and Stamnes, K. (1993). An Accurate Parameterization of the Radiative Properties of Water Clouds Suitable for Use in Climate Models. *Journal of Climate*, 6(4):728–742.

Iacovazzi, R. and Wu, X. (2020). GOES-16 Advanced Baseline Imager visible near-infrared channel low-light signal-to-noise ratio. *JARS*, 14(2):026502.

Ickes, L., Welti, A., Hoose, C., and Lohmann, U. (2015). Classical nucleation theory of homogeneous freezing of water: Thermodynamic and kinetic parameters. *Physical Chemistry Chemical Physics*, 17(8):5514–5537.

Igel, M. R., Drager, A. J., and van den Heever, S. C. (2014). A CloudSat cloud object partitioning technique and assessment and integration of deep convective anvil sensitivities to sea surface temperature. *Journal of Geophysical Research: Atmospheres*, 119(17):10515–10535.

Jeevanjee, N. (2017). Vertical Velocity in the Gray Zone. *Journal of Advances in Modeling Earth Systems*, 9(6):2304–2316.

Jeevanjee, N. and Fueglistaler, S. (2020). Simple Spectral Models for Atmospheric Radiative Cooling. *Journal of the Atmospheric Sciences*, 77(2):479–497.

Jeevanjee, N. and Romps, D. M. (2013). Convective self-aggregation, cold pools, and domain size. *Geophysical Research Letters*, 40(5):994–998.

Jirak, I. L. and Cotton, W. R. (2007). Observational Analysis of the Predictability of Mesoscale Convective Systems. *Weather and Forecasting*, 22(4):813–838.

Johnson, J. T., MacKeen, P. L., Witt, A., Mitchell, E. D. W., Stumpf, G. J., Eilts, M. D., and Thomas, K. W. (1998). The Storm Cell Identification and Tracking Algorithm: An Enhanced WSR-88D Algorithm. *Weather and Forecasting*, 13(2):263–276.

Johnson, R. H., Rickenbach, T. M., Rutledge, S. A., Ciesielski, P. E., and Schubert, W. H. (1999). Trimodal Characteristics of Tropical Convection. *Journal of Climate*, 12(8):2397–2418.

Jones, W. K., Christensen, M. W., and Stier, P. (2023a). A semi-Lagrangian method for detecting and tracking deep convective clouds in geostationary satellite observations. *Atmospheric Measurement Techniques*, 16(4):1043–1059.

Jones, W. K., Stengel, M., and Stier, P. (2023b). A Lagrangian Perspective on the Lifecycle and Cloud Radiative Effect of Deep Convective Clouds Over Africa. *EGUsphere*, pages 1–25.

Kalchbrenner, N., Grefenstette, E., and Blunsom, P. (2014). A Convolutional Neural Network for Modelling Sentences. *arXiv:1404.2188 [cs]*.

Kanji, Z. A., Ladino, L. A., Wex, H., Boose, Y., Burkert-Kohn, M., Cziczo, D. J., and Krämer, M. (2017). Overview of Ice Nucleating Particles. *Meteorological Monographs*, 58(1):1.1–1.33.

Kärcher, B. and Lohmann, U. (2002). A Parameterization of cirrus cloud formation: Homogeneous freezing including effects of aerosol size. *Journal of Geophysical Research: Atmospheres*, 107(D23):AAC 9–1–AAC 9–10.

Kärcher, B. and Ström, J. (2003). The roles of dynamical variability and aerosols in cirrus cloud formation. *Atmospheric Chemistry and Physics*, 3(3):823–838.

Keenan, T., Joe, P., Wilson, J., Collier, C., Golding, B., Burgess, D., May, P., Pierce, C., Bally, J., Crook, A., Seed, A., Sills, D., Berry, L., Potts, R., Bell, I., Fox, N., Ebert, E., Eilts, M., O'Loughlin, K., Webb, R., Carbone, R., Browning, K., Roberts, R., and Mueller, C. (2003). The Sydney 2000 World Weather Research Programme Forecast Demonstration Project: Overview and Current Status: Overview and Current Status. *Bulletin of the American Meteorological Society*, 84(8):1041–1054.

Khain, A., Rosenfeld, D., and Pokrovsky, A. (2005). Aerosol impact on the dynamics and microphysics of deep convective clouds. *Quarterly Journal of the Royal Meteorological Society*, 131(611):2639–2663.

Konduru, R. T., Kishtawal, C. M., and Shah, S. (2013). A new perspective on the infrared brightness temperature distribution of the deep convective clouds. *J Earth Syst Sci*, 122(5):1195–1206.

Koop, T., Luo, B., Tsias, A., and Peter, T. (2000). Water activity as the determinant for homogeneous ice nucleation in aqueous solutions. *Nature*, 406(6796):611–614.

Kroeger, T., Timofte, R., Dai, D., and Van Gool, L. (2016). Fast Optical Flow using Dense Inverse Search. arXiv.

Laing, A. G. and Michael Fritsch, J. (1997). The global population of mesoscale convective complexes. *Quarterly Journal of the Royal Meteorological Society*, 123(538):389–405.

Lakshmanan, V. and Smith, T. (2010). An Objective Method of Evaluating and Devising Storm-Tracking Algorithms. *Weather and Forecasting*, 25(2):701–709.

Lakshmanan, V., Smith, T., Stumpf, G., and Hondl, K. (2007). The Warning Decision Support System–Integrated Information. *Weather and Forecasting*, 22(3):596–612.

Leese, J. A., Novak, C. S., and Ray Taylor, V. (1970). The determination of cloud pattern motions from geosynchronous satellite image data. *Pattern Recognition*, 2(4):279–292.

LeMone, M. A. and Zipser, E. J. (1980). Cumulonimbus Vertical Velocity Events in GATE. Part I: Diameter, Intensity and Mass Flux. *Journal of the Atmospheric Sciences*, 37(11):2444–2457.

Li, J., Feng, Z., Qian, Y., and Leung, L. R. (2021). A high-resolution unified observational data product of mesoscale convective systems and isolated deep convection in the United States for 2004–2017. *Earth System Science Data*, 13(2):827–856.

Lilly, D. K. (1988). Cirrus Outflow Dynamics. *Journal of the Atmospheric Sciences*, 45(10):1594–1605.

Lin, B., Wong, T., Wielicki, B. A., and Hu, Y. (2004). Examination of the Decadal Tropical Mean ERBS Nonscanner Radiation Data for the Iris Hypothesis. *Journal of Climate*, 17(6):1239–1246.

Lindsey, D. T., Bikos, D., and Grasso, L. (2018). Using the GOES-16 Split Window Difference to Detect a Boundary prior to Cloud Formation. *Bulletin of the American Meteorological Society*, 99(8):1541–1544.

Lindsey, D. T., Grasso, L., Dostalek, J. F., and Kerkmann, J. (2014). Use of the GOES-R Split-Window Difference to Diagnose Deepening Low-Level Water Vapor. *Journal of Applied Meteorology and Climatology*, 53(8):2005–2016.

Lindzen, R. S., Chou, M.-D., and Hou, A. Y. (2001). Does the Earth Have an Adaptive Infrared Iris? *Bulletin of the American Meteorological Society*, 82(3):417–432.

Liu, H., Koren, I., and Altaratz, O. (2023). Observed decreasing trend in the upper-tropospheric cloud top temperature. *npj Climate and Atmospheric Science*, 6(1):1–8.

Loeb, N. G., Doelling, D. R., Wang, H., Su, W., Nguyen, C., Corbett, J. G., Liang, L., Mitrescu, C., Rose, F. G., and Kato, S. (2018). Clouds and the Earth's Radiant Energy System (CERES) Energy Balanced and Filled (EBAF) Top-of-Atmosphere (TOA) Edition-4.0 Data Product. *Journal of Climate*, 31(2):895–918.

Lucas, B. D. and Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th international joint conference on Artificial intelligence - Volume 2*, IJCAI'81, pages 674–679, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Luo, Z. and Rossow, W. B. (2004). Characterizing Tropical Cirrus Life Cycle, Evolution, and Interaction with Upper-Tropospheric Water Vapor Using Lagrangian Trajectory Analysis of Satellite Observations. *Journal of Climate*, 17(23):4541–4563.

Maddox, R. A. (1980a). Mesoscale Convective Complexes. *Bulletin of the American Meteorological Society*, 61(11):1374–1387.

Maddox, R. A. (1980b). Mesoscale Convective Complexes. *Bulletin of the American Meteorological Society*, 61(11):1374–1387.

Magnus, G. (1844). Versuche über die Spannkräfte des Wasserdampfs. *Annalen der Physik*, 137(2):225–247.

Mapes, B. E. (1993). Gregarious Tropical Convection. *Journal of the Atmospheric Sciences*, 50(13):2026–2037.

Mapes, B. E. (2001). Water's two height scales: The moist adiabat and the radiative troposphere. *Quarterly Journal of the Royal Meteorological Society*, 127(577):2353–2366.

Marinescu, P. J., Heever, S. C. v. d., Heikenfeld, M., Barrett, A. I., Barthlott, C., Hoose, C., Fan, J., Fridlind, A. M., Matsui, T., Miltenberger, A. K., Stier, P., Vie, B., White, B. A., and Zhang, Y. (2021). Impacts of Varying Concentrations of Cloud Condensation Nuclei on Deep Convective Cloud Updrafts—A Multimodel Assessment. *Journal of the Atmospheric Sciences*, 78(4):1147–1172.

Massie, S., Gettelman, A., Randel, W., and Baumgardner, D. (2002). Distribution of tropical cirrus in relation to convection. *Journal of Geophysical Research: Atmospheres*, 107(D21):AAC 19–1–AAC 19–16.

Masunaga, H. and L'Ecuyer, T. S. (2014). A Mechanism of Tropical Convection Inferred from Observed Variability in the Moist Static Energy Budget. *Journal of the Atmospheric Sciences*, 71(10):3747–3766.

Masunaga, H. and Luo, Z. J. (2016). Convective and large-scale mass flux profiles over tropical oceans determined from synergistic analysis of a suite of satellite observations. *Journal of Geophysical Research: Atmospheres*, 121(13):7958–7974.

Matsudo, C. M. and Salio, P. V. (2011). Severe weather reports and proximity to deep convection over Northern Argentina. *Atmospheric Research*, 100(4):523–537.

McGarragh, G. R., Poulsen, C. A., Thomas, G. E., Povey, A. C., Sus, O., Stapelberg, S., Schlundt, C., Proud, S., Christensen, M. W., Stengel, M., Hollmann, R., and Grainger, R. G. (2018). The Community Cloud retrieval for CLimate (CC4CL) Part 2: The optimal estimation approach. *Atmospheric Measurement Techniques*, 11(6):3397–3431.

McKim, B., Bony, S., and Dufresne, J.-L. (2024). Weak anvil cloud area feedback suggested by physical and observational constraints. *Nature Geoscience*, pages 1–6.

Mecikalski, J. R. and Bedka, K. M. (2006). Forecasting Convective Initiation by Monitoring the Evolution of Moving Cumulus in Daytime GOES Imagery. *Monthly Weather Review*, 134(1):49–78.

Menzel, W. P. (2001). Cloud Tracking with Satellite Imagery: From the Pioneering Work of Ted Fujita to the Present. *Bulletin of the American Meteorological Society*, 82(1):33–48.

Morel, C. and Senesi, S. (2002). A climatology of mesoscale convective systems over Europe using satellite infrared imagery. I: Methodology. *Quarterly Journal of the Royal Meteorological Society*, 128(584):1953–1971.

Muller, C. J. and O'Gorman, P. A. (2011). An energetic perspective on the regional response of precipitation to climate change. *Nature Climate Change*, 1(5):266–271.

Müller, R., Haussler, S., and Jerg, M. (2018). The Role of NWP Filter for the Satellite Based Detection of Cumulonimbus Clouds. *Remote Sensing*, 10(3):386.

Müller, R., Haussler, S., Jerg, M., and Heizenreder, D. (2019). A Novel Approach for the Detection of Developing Thunderstorm Cells. *Remote Sensing*, 11(4):443.

Mülmenstädt, J., Sourdeval, O., Delanoë, J., and Quaas, J. (2015). Frequency of occurrence of rain from liquid-, mixed-, and ice-phase clouds derived from A-Train satellite retrievals. *Geophysical Research Letters*, 42(15):6502–6509.

Nakazawa, T. (1988). Tropical Super Clusters within Intraseasonal Variations over the Western Pacific. *Journal of the Meteorological Society of Japan. Ser. II*, 66(6):823–839.

Nicholson, S. E. (2009). A revised picture of the structure of the monsoon and land ITCZ over West Africa. *Climate Dynamics*, 32(7):1155–1171.

Nicholson, S. E. (2018). The ITCZ and the Seasonal Cycle over Equatorial Africa. *Bulletin of the American Meteorological Society*, 99(2):337–348.

Norris, J. R., Allen, R. J., Evan, A. T., Zelinka, M. D., O'Dell, C. W., and Klein, S. A. (2016). Evidence for climate change in the satellite cloud record. *Nature*, 536(7614):72–75.

Nowicki, S. M. J. and Merchant, C. J. (2004). Observations of diurnal and spatial variability of radiative forcing by equatorial deep convective clouds. *Journal of Geophysical Research: Atmospheres*, 109(D11). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2003JD004176.

Núñez Ocasio, K. M., Evans, J. L., and Young, G. S. (2020). Tracking Mesoscale Convective Systems that are Potential Candidates for Tropical Cyclogenesis. *Monthly Weather Review*, 148(2):655–669.

Orville, R. E. and Henderson, R. W. (1984). Absolute Spectral Irradiance Measurements of Lightning from 375 to 880 nm. *Journal of the Atmospheric Sciences*, 41(21):3180–3187.

Park, J. M., van den Heever, S. C., Igel, A. L., Grant, L. D., Johnson, J. S., Saleeby, S. M., Miller, S. D., and Reid, J. S. (2020). Environmental controls on tropical sea breeze convection and resulting aerosol redistribution. *Journal of Geophysical Research: Atmospheres*, 125(6).

Peters, J. M. and Chavas, D. R. (2021). Evaluating the Conservation of Energy Variables in Simulations of Deep Moist Convection. *Journal of the Atmospheric Sciences*, 78(10):3229–3246.

Peterson, M. (2019). Research Applications for the Geostationary Lightning Mapper Operational Lightning Flash Data Product. *J Geophys Res Atmos*, 124(17-18):10205–10231.

Peterson, M. (2020). Removing solar artifacts from Geostationary Lightning Mapper data to document lightning extremes. *J Appl Remote Sens*, 14(3).

Pili, P., Matheson, L., Tranquilli, C., Müller, J., Hewison, T., Carlier, S., Bianchi, S., and Coste, P. (2016). The In-orbit Performance of the Meteosat Second Generation SEVIRI Instruments. In *EUMETSAT Meteorological Satellite Conference*, Darmstadt, Germany.

Pinto, J. O., Grim, J. A., and Steiner, M. (2015). Assessment of the High-Resolution Rapid Refresh Model's Ability to Predict Mesoscale Convective Systems Using Object-Based Evaluation. *Weather and Forecasting*, 30(4):892–913.

Plant, R. S. (2009). Statistical properties of cloud lifecycles in cloud-resolving models. *Atmospheric Chemistry and Physics*, 9(6):2195–2205.

Prein, A. F., Feng, Z., Fiolleau, T., Moon, Z. L., Núñez Ocasio, K. M., Kukulies, J., Roca, R., Varble, A. C., Rehbein, A., Liu, C., Ikeda, K., Mu, Y., and Rasmussen, R. M. (2024). Km-Scale Simulations of Mesoscale Convective Systems Over South America—A Feature Tracker Intercomparison. *Journal of Geophysical Research: Atmospheres*, 129(8):e2023JD040254.

Prein, A. F., Langhans, W., Fosser, G., Ferrone, A., Ban, N., Goergen, K., Keller, M., Tölle, M., Gutjahr, O., Feser, F., Brisson, E., Kollet, S., Schmidli, J., van Lipzig, N. P. M., and Leung, R. (2015). A review on regional convection-permitting climate modeling: Demonstrations, prospects, and challenges. *Reviews of Geophysics*, 53(2):323–361.

Prein, A. F., Rasmussen, R. M., Wang, D., and Giangrande, S. E. (2021). Sensitivity of organized convective storms to model grid spacing in current and future climates. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2195):20190546.

Protopapadaki, S. E., Stubenrauch, C. J., and Feofilov, A. G. (2017). Upper tropospheric cloud systems derived from IR sounders: Properties of cirrus anvils in the tropics. *Atmospheric Chemistry and Physics*, 17(6):3845–3859.

Proud, S. R. and Bachmeier, S. (2021). Record-Low Cloud Temperatures Associated With a Tropical Deep Convective Event. *Geophysical Research Letters*, 48(6):e2020GL092261.

Punge, H. J. and Kunz, M. (2016). Hail observations and hailstorm characteristics in Europe: A review. *Atmospheric Research*, 176-177:159–184.

Pérez, J. S., Meinhardt-Llopis, E., and Facciolo, G. (2013). TV-L1 Optical Flow Estimation. *Image Processing On Line*, 3:137–150.

Raghuraman, S. P., Medeiros, B., and Gettelman, A. (2024). Observational Quantification of Tropical High Cloud Changes and Feedbacks. *Journal of Geophysical Research: Atmospheres*, 129(7):e2023JD039364.

Ramanathan, V., Cess, R. D., Harrison, E. F., Minnis, P., Barkstrom, B. R., Ahmad, E., and Hartmann, D. (1989). Cloud-Radiative Forcing and Climate: Results from the Earth Radiation Budget Experiment. *Science*, 243(4887):57–63.

Raut, B. A., Jackson, R., Picel, M., Collis, S. M., Bergemann, M., and Jakob, C. (2021). An Adaptive Tracking Algorithm for Convection in Simulated and Remote Sensing Data. *Journal of Applied Meteorology and Climatology*, 60(4):513–526.

Riehl, H. and Malkus, J. S. (1958). On the Heat Balance in the Equatorial Trough Zone. *Geophysica*, 6(3–4).

Riehl, H., Yeh, T. C., Malkus, J. S., and la Seur, N. E. (1951). The north-east trade of the Pacific Ocean. *Quarterly Journal of the Royal Meteorological Society*, 77(334):598–626.

Rinehart, R. E. and Garvey, E. T. (1978). Three-dimensional storm motion detection by conventional weather radar. *Nature*, 273(5660):287–289.

Roberts, R. D. and Rutledge, S. (2003). Nowcasting Storm Initiation and Growth Using GOES-8 and WSR-88D Data. *Weather and Forecasting*, 18(4):562–584.

Roca, R., Fiolleau, T., and Bouniol, D. (2017). A Simple Model of the Life Cycle of Mesoscale Convective Systems Cloud Shield in the Tropics. *Journal of Climate*, 30(11):4283–4298.

Rogers, R. R. (1979). *A short course in cloud physics / by R. R. Rogers.* Pergamon Press Oxford ; New York, 2d ed. edition.

Romps, D. M. (2015). MSE Minus CAPE is the True Conserved Variable for an Adiabatically Lifted Parcel. *Journal of the Atmospheric Sciences*, 72(9):3639–3646.

Romps, D. M. and Kuang, Z. (2010). Do Undiluted Convective Plumes Exist in the Upper Tropical Troposphere? *Journal of the Atmospheric Sciences*, 67(2):468–484.

Rosenfeld, D. (1987). Objective Method for Analysis and Tracking of Convective Cells as Seen by Radar. *Journal of Atmospheric and Oceanic Technology*, 4(3):422–434.

Rosenfeld, D., Lohmann, U., Raga, G. B., O'Dowd, C. D., Kulmala, M., Fuzzi, S., Reissell, A., and Andreae, M. O. (2008). Flood or Drought: How Do Aerosols Affect Precipitation? *Science*, 321(5894):1309–1313.

Rosenfeld, D., Wolff, D. B., and Atlas, D. (1993). General Probability-matched Relations between Radar Reflectivity and Rain Rate. *Journal of Applied Meteorology and Climatology*, 32(1):50–72.

Ryan, B. F., Wishart, E. R., and Shaw, D. E. (1976). The Growth Rates and Densities of Ice Crystals between -3°C and -21°C. *Journal of the Atmospheric Sciences*, 33(5):842–850.

Schmetz, J., Tjemkes, S. A., Gube, M., and van de Berg, L. (1997). Monitoring deep convection and convective overshooting with METEOSAT. *Advances in Space Research*, 19(3):433–441.

Schmit, T. J., Griffith, P., Gunshor, M. M., Daniels, J. M., Goodman, S. J., and Lebair, W. J. (2016). A Closer Look at the ABI on the GOES-R Series. *Bull. Amer. Meteor. Soc.*, 98(4):681–698.

Schmit, T. J. and Gunshor, M. M. (2020). Chapter 4 - ABI Imagery from the GOES-R Series. In Goodman, S. J., Schmit, T. J., Daniels, J., and Redmon, R. J., editors, *The GOES-R Series*, pages 23–34. Elsevier.

Schröder, M., König, M., and Schmetz, J. (2009). Deep convection observed by the Spinning Enhanced Visible and Infrared Imager on board Meteosat 8: Spatial distribution and temporal evolution over Africa in summer and winter 2006. *Journal of Geophysical Research: Atmospheres*, 114(D5).

Schulz, H. and Stevens, B. (2018). Observing the Tropical Atmosphere in Moisture Space. *Journal of the Atmospheric Sciences*, 75(10):3313–3330.

Schumacher, C. and Houze, R. A. (2003). Stratiform Rain in the Tropics as Seen by the TRMM Precipitation Radar. *Journal of Climate*, 16(11):1739–1756.

Schumacher, C., Houze, R. A., and Kraucunas, I. (2004). The Tropical Dynamical Response to Latent Heating Estimates Derived from the TRMM Precipitation Radar. *Journal of the Atmospheric Sciences*, 61(12):1341–1358.

Seeley, J. T., Jeevanjee, N., Langhans, W., and Romps, D. M. (2019a). Formation of Tropical Anvil Clouds by Slow Evaporation. *Geophysical Research Letters*, 46(1):492–501.

Seeley, J. T., Jeevanjee, N., and Romps, D. M. (2019b). FAT or FiTT: Are Anvil Clouds or the Tropopause Temperature Invariant? *Geophysical Research Letters*, 46(3):1842–1850.

Seeley, J. T. and Romps, D. M. (2015a). The Effect of Global Warming on Severe Thunderstorms in the United States. *Journal of Climate*, 28(6):2443–2458.

Seeley, J. T. and Romps, D. M. (2015b). Why does tropical convective available potential energy (CAPE) increase with warming? *Geophysical Research Letters*, 42(23):10,429–10,437.

Seeley, J. T. and Romps, D. M. (2016). Tropical cloud buoyancy is the same in a world with or without ice. *Geophysical Research Letters*, 43(7):3572–3579.

Seidel, S. D. and Yang, D. (2022). Temperatures of Anvil Clouds and Radiative Tropopause in a Wide Array of Cloud-Resolving Simulations. *Journal of Climate*, 35(24):8065–8078.

Senf, F., Klocke, D., and Brueck, M. (2018a). Size-Resolved Evaluation of Simulated Deep Tropical Convection. *Mon. Wea. Rev.*, 146(7):2161–2182.

Senf, F., Lenk, S., Deneke, H., and Heinze, R. (2018b). High-resolution simulations of deep convective growth and Meteosat observations.

Sherwood, S. C., Webb, M. J., Annan, J. D., Armour, K. C., Forster, P. M., Hargreaves, J. C., Hegerl, G., Klein, S. A., Marvel, K. D., Rohling, E. J., Watanabe, M., Andrews, T., Braconnot, P., Bretherton, C. S., Foster, G. L., Hausfather, Z., von der Heydt, A. S., Knutti, R., Mauritsen, T., Norris, J. R., Proistosescu, C., Rugenstein, M., Schmidt, G. A., Tokarska, K. B., and Zelinka, M. D. (2020). An Assessment of Earth's Climate Sensitivity Using Multiple Lines of Evidence. *Reviews of Geophysics*, 58(4):e2019RG000678.

Sobel, I. (2014). An Isotropic 3x3 Image Gradient Operator. *Presentation at Stanford A.I. Project 1968*.

Soden, B. J. (2000). The diurnal cycle of convection, clouds, and water vapor in the tropical upper troposphere. *Geophysical Research Letters*, 27(15):2173–2176.

Sokol, A. B. and Hartmann, D. L. (2020). Tropical Anvil Clouds: Radiative Driving Toward a Preferred State. *Journal of Geophysical Research: Atmospheres*, 125(21):e2020JD033107.

Sokol, A. B., Wall, C. J., and Hartmann, D. L. (2024). Greater climate sensitivity implied by anvil cloud thinning. *Nature Geoscience*, pages 1–6.

Song, F., Feng, Z., Leung, L. R., Jr, R. A. H., Wang, J., Hardin, J., and Homeyer, C. R. (2019). Contrasting Spring and Summer Large-Scale Environments Associated with Mesoscale Convective Systems over the U.S. Great Plains. *Journal of Climate*, 32(20):6749–6767.

Steinacker, R., Dorninger, M., Wölfelmaier, F., and Krennert, T. (2000). Automatic Tracking of Convective Cells and Cell Complexes from Lightning and Radar Data. *Meteorology and Atmospheric Physics*, 72(2):101–110.

Stephens, G., Winker, D., Pelon, J., Trepte, C., Vane, D., Yuhas, C., L'Ecuyer, T., and Lebsock, M. (2018). CloudSat and CALIPSO within the A-Train: Ten Years of Actively Observing the Earth System. *Bulletin of the American Meteorological Society*, 99(3):569–581.

Stephens, G. L., Gabriel, P. M., and Partain, P. T. (2001). Parameterization of Atmospheric Radiative Transfer. Part I: Validity of Simple Models. *Journal of the Atmospheric Sciences*, 58(22):3391–3409.

Stevens, B. (2005). ATMOSPHERIC MOIST CONVECTION. *Annual Review of Earth and Planetary Sciences*, 33(Volume 33, 2005):605–643.

Stevens, B., Acquistapace, C., Hansen, A., Heinze, R., Klinger, C., Klocke, D., Rybka, H., Schubotz, W., Windmiller, J., Adamidis, P., Arka, I., Barlakas, V., Biercamp, J., Brueck, M., Brune, S., Buehler, S. A., Burkhardt, U., Cioni, G., Costa-Surós, M., Crewell, S., Crüger, T., Deneke, H., Friederichs, P., Henken, C. C., Hohenegger, C., Jacob, M., Jakub, F., Kalthoff, N., Köhler, M., van Laar, T. W., Li, P., Löhnert, U., Macke, A., Madenach, N., Mayer, B., Nam, C., Naumann, A. K., Peters, K., Poll, S., Quaas, J., Röber, N., Rochetin, N., Scheck, L., Schemann, V., Schnitt, S., Seifert, A., Senf, F., Shapkalijevski, M., Simmer, C., Singh, S., Sourdeval, O., Spickermann, D., Strandgren, J., Tessiot, O., Vercauteren, N., Vial, J., Voigt, A., and Zängl, G. (2020). The Added Value of Large-eddy and Storm-resolving Models for Simulating Clouds and Precipitation. *Journal of the Meteorological Society of Japan. Ser. II*, 98(2):395–435.

Sullivan, S. C. and Voigt, A. (2021). Ice microphysical processes exert a strong control on the simulated radiative energy budget in the tropics. *Communications Earth & Environment*, 2(1):1–8.

Sus, O., Stengel, M., Stapelberg, S., McGarragh, G., Poulsen, C., Povey, A. C., Schlundt, C., Thomas, G., Christensen, M., Proud, S., Jerg, M., Grainger, R., and Hollmann, R. (2018). The Community Cloud retrieval for CLimate (CC4CL) Part 1: A framework

applied to multiple satellite imaging sensors. *Atmospheric Measurement Techniques*, 11(6):3373–3396.

Takahashi, H., Lebsock, M. D., Richardson, M., Marchand, R., and Kay, J. E. (2019). When Will Spaceborne Cloud Radar Detect Upward Shifts in Cloud Heights? *Journal of Geophysical Research: Atmospheres*, 124(13):7270–7285.

Takahashi, H. and Luo, Z. (2012). Where is the level of neutral buoyancy for deep convection? *Geophysical Research Letters*, 39(15).

Takahashi, H. and Luo, Z. J. (2014). Characterizing tropical overshooting deep convection from joint analysis of CloudSat and geostationary satellite observations. *Journal of Geophysical Research: Atmospheres*, 119(1):112–121.

Takahashi, H., Luo, Z. J., and Stephens, G. L. (2017a). Level of neutral buoyancy, deep convective outflow, and convective core: New perspectives based on 5 years of CloudSat data. *Journal of Geophysical Research: Atmospheres*, 122(5):2958–2969.

Takahashi, H., Protopapadaki, S., Stubenrauch, C., Luo, Z. J., and Stephens, G. L. (2017b). Relationships between Convective Strength and Anvil Development based on AIRS-CloudSat Joint Dataset.

Taylor, S., Stier, P., White, B., Finkensieper, S., and Stengel, M. (2017). Evaluating the diurnal cycle in cloud top temperature from SEVIRI. *Atmospheric Chemistry and Physics*, 17(11):7035–7053.

Tecson, J. J. and Fujita, T. T. (1975). Cloud-motion Vectors over the Gate Area Computed by MCIDS and METRACOM Methods. *Satellite and Mesometeorology Research Project Research Paper*, (136).

Thomas, C., Corpetti, T., and Mémin, E. (2010). Data Assimilation for Convective-Cell Tracking on Meteorological Image Sequences. *IEEE Transactions on Geoscience and Remote Sensing*, 48(8):3162–3177.

Trapp, R. J., Diffenbaugh, N. S., Brooks, H. E., Baldwin, M. E., Robinson, E. D., and Pal, J. S. (2007). Changes in severe thunderstorm environment frequency during the 21st century caused by anthropogenically enhanced global radiative forcing. *Proceedings of the National Academy of Sciences*, 104(50):19719–19723.

Trenberth, K. E., Dai, A., Rasmussen, R. M., and Parsons, D. B. (2003). The Changing Character of Precipitation. *Bull. Amer. Meteor. Soc.*, 84(9):1205–1218.

Tripoli, G. J. and Cotton, W. R. (1979). A Numerical Investigation of Several Factors Contributing to the Observed Variable Intensity of Deep Convection over South Florida. *Journal of Applied Meteorology*, 19(9):1037–1063.

Tsakraklides, G. and Evans, J. L. (2003). Global and Regional Diurnal Variations of Organized Convection. *Journal of Climate*, 16(10):1562–1572.

Ullrich, P. A. and Zarzycki, C. M. (2017). TempestExtremes: a framework for scale-insensitive pointwise feature tracking on unstructured grids. *Geoscientific Model Development*, 10(3):1069–1090.

Ullrich, P. A., Zarzycki, C. M., McClenny, E. E., Pinheiro, M. C., Stansfield, A. M., and Reed, K. A. (2021). TempestExtremes v2.1: a community framework for feature detection, tracking, and analysis in large datasets. *Geoscientific Model Development*, 14(8):5023–5048.

V. Hobbs, P. (1993). *Aerosol-Cloud-Climate Interactions*. Academic Press.

van den Heever, S., Haddad, Z., Dolan, B., Freeman, S., Grant, L., Kollias, P., Leung, G., Luo, J., Marinescu, P., Posselt, D., Rasmussen, K., Sai, P., Schulte, R., Stephens, G., Storer, R., and Takahashi, H. (2023). Tropical Convection through the Lens of the INCUS Mission. Vienna, Austria.

Varble, A. (2018). Erroneous Attribution of Deep Convective Invigoration to Aerosol Concentration. *J. Atmos. Sci.*, 75(4):1351–1368.

Varble, A. C., Igel, A. L., Morrison, H., Grabowski, W. W., and Lebo, Z. J. (2023). Opinion: A Critical Evaluation of the Evidence for Aerosol Invigoration of Deep Convection. *EGUsphere*, pages 1–31.

Vecchi, G. A. and Soden, B. J. (2007). Global Warming and the Weakening of the Tropical Circulation. *Journal of Climate*, 20(17):4316–4340.

Vizy, E. K. and Cook, K. H. (2019). Understanding the summertime diurnal cycle of precipitation over sub-Saharan West Africa: regions with daytime rainfall peaks in the absence of significant topographic features. *Climate Dynamics*, 52(5):2903–2922.

Vondou, D. A., Nzeukou, A., Lenouo, A., and Mkankam Kamga, F. (2010). Seasonal variations in the diurnal patterns of convection in CameroonNigeria and their neighboring areas. *Atmospheric Science Letters*, 11(4):290–300.

Waitz, F., Schnaiter, M., Leisner, T., and Järvinen, E. (2022). In situ observation of riming in mixed-phase clouds using the PHIPS probe. *Atmospheric Chemistry and Physics*, 22(11):7087–7103.

Wall, C. J. and Hartmann, D. L. (2018). Balanced Cloud Radiative Effects Across a Range of Dynamical Conditions Over the Tropical West Pacific. *Geophysical Research Letters*, 45(20):11,490–11,498.

Wall, C. J., Hartmann, D. L., Thieman, M. M., Smith, W. L., and Minnis, P. (2018). The Life Cycle of Anvil Clouds and the Top-of-Atmosphere Radiation Balance over the Tropical West Pacific. *Journal of Climate*, 31(24):10059–10080.

Wall, C. J., Norris, J. R., Gasparini, B., Smith, W. L., Thieman, M. M., and Sourdeval, O. (2020). Observational Evidence that Radiative Heating Modifies the Life Cycle of Tropical Anvil Clouds. *Journal of Climate*, 33(20):8621–8640.

Walters, C. K. and Winkler, J. A. (2001). Airflow Configurations of Warm Season Southerly Low-Level Wind Maxima in the Great Plains. Part I: Spatial and Temporal Characteristics and Relationship to Convection. *Weather and Forecasting*, 16(5):513–530.

Wang, D., Jensen, M. P., D'Iorio, J. A., Jozef, G., Giangrande, S. E., Johnson, K. L., Luo, Z. J., Starzec, M., and Mullendore, G. L. (2020). An Observational Comparison of Level of Neutral Buoyancy and Level of Maximum Detrainment in Tropical Deep Convective Clouds. *Journal of Geophysical Research: Atmospheres*, 125(16):e2020JD032637.

Weinzaepfel, P., Revaud, J., Harchaoui, Z., and Schmid, C. (2013). DeepFlow: Large Displacement Optical Flow with Deep Matching. pages 1385–1392.

Weisman, M. L. (2015). MESOSCALE METEOROLOGY | Convective Storms: Overview. In North, G. R., Pyle, J., and Zhang, F., editors, *Encyclopedia of Atmospheric Sciences (Second Edition)*, pages 401–410. Academic Press, Oxford.

Weiss, J. and Zeitler, J. (2008). Supercells of the Serranías del Burro. In *24th Conf. on Severe Local Storms*, volume 17A.4, Savanna, GA. American Meteorological Society.

Westra, S., Fowler, H. J., Evans, J. P., Alexander, L. V., Berg, P., Johnson, F., Kendon, E. J., Lenderink, G., and Roberts, N. M. (2014). Future changes to the intensity and frequency of short-duration extreme rainfall. *Rev. Geophys.*, 52(3):2014RG000464.

Williams, E. R., Geotis, S. G., Renno, N., Rutledge, S. A., Rasmussen, E., and Rickenbach, T. (1992). A Radar and Electrical Study of Tropical "Hot Towers". *J. Atmos. Sci.*, 49(15):1386–1395.

Williams, E. R., Weber, M. E., and Orville, R. E. (1989). The relationship between lightning type and convective state of thunderclouds. *Journal of Geophysical Research: Atmospheres*, 94(D11):13213–13220.

Williams, M. and Houze, R. A. (1987). Satellite-Observed Characteristics of Winter Monsoon Cloud Clusters. *Monthly Weather Review*, 115(2):505–519.

Wilson, J. W., Crook, N. A., Mueller, C. K., Sun, J., and Dixon, M. (1998). Nowcasting Thunderstorms: A Status Report. *Bulletin of the American Meteorological Society*, 79(10):2079–2100.

Wilson, J. W., Ebert, E. E., Saxen, T. R., Roberts, R. D., Mueller, C. K., Sleigh, M., Pierce, C. E., and Seed, A. (2004). Sydney 2000 Forecast Demonstration Project: Convective Storm Nowcasting. *Weather and Forecasting*, 19(1):131–150.

Wilson, J. W. and Schreiber, W. E. (1986). Initiation of Convective Storms at Radar-Observed Boundary-Layer Convergence Lines. *Monthly Weather Review*, 114(12):2516–2536.

Woo, W.-c. and Wong, W.-k. (2017). Operational Application of Optical Flow Techniques to Radar-Based Rainfall Nowcasting. *Atmosphere*, 8(3):48.

Wu, Q., Wang, H.-Q., Lin, Y.-J., Zhuang, Y.-Z., and Zhang, Y. (2016). Deriving AMVs from Geostationary Satellite Images Using Optical Flow Algorithm Based on Polynomial Expansion. *J. Atmos. Oceanic Technol.*, 33(8):1727–1747.

Wulff, J. and Black, M. J. (2015). Efficient sparse-to-dense optical flow estimation using a learned basis and layers. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 120–130.

Wyser, K. (1998). The Effective Radius in Ice Clouds. *Journal of Climate*, 11(7):1793–1802.

Xu, K.-M. and Emanuel, K. A. (1989). Is the Tropical Atmosphere Conditionally Unstable? *Monthly Weather Review*, 117(7):1471–1479.

Yau, K., M. and Rogers, R., R. (1989). *A Short Course in Cloud Physics - 3rd Edition.* Butterworth-Heinemann.

Zach, C., Pock, T., and Bischof, H. (2007). A Duality Based Approach for Realtime TV-L1 Optical Flow: 29th DAGM Symposium on Pattern Recognition. *Proceedings 29th DAGM Symposium "Pattern Recognition"*, pages 214–223.

Zelinka, M. D. and Hartmann, D. L. (2010). Why is longwave cloud feedback positive? *Journal of Geophysical Research: Atmospheres*, 115(D16).

Zhang, G. J. (2009). Effects of entrainment on convective available potential energy and closure assumptions in convection parameterization. *Journal of Geophysical Research: Atmospheres*, 114(D7).

Zhang, X., van der A, R., Ding, J., Eskes, H., van Geffen, J., Yin, Y., Anema, J., Vagasky, C., L. Lapierre, J., and Kuang, X. (2023). Spaceborne Observations of Lightning NO2 in the Arctic. *Environmental Science & Technology*, 57(6):2322–2332.

Zhang, Y., Wistar, S., Piedra-Fernández, J. A., Li, J., Steinberg, M. A., and Wang, J. Z. (2014). Locating visual storm signatures from satellite images. In *2014 IEEE International Conference on Big Data (Big Data)*, pages 711–720.

Zinner, T., Forster, C., Coning, E. d., and Betz, H.-D. (2013). Validation of the Meteosat storm detection and nowcasting system Cb-TRAM with lightning network data – Europe and South Africa. *Atmospheric Measurement Techniques*, 6(6):1567–1583.

Zinner, T., Mannstein, H., and Tafferner, A. (2008). Cb-TRAM: Tracking and monitoring severe convection from onset over rapid development to mature phase using multi-channel Meteosat-8 SEVIRI data. *Meteorology and Atmospheric Physics*, 101(3):191–210.

Zipser, E. J. (1969). The Role of Organized Unsaturated Convective Downdrafts in the Structure and Rapid Decay of an Equatorial Disturbance. *Journal of Applied Meteorology and Climatology*, 8(5):799–814.

Zipser, E. J. and LeMone, M. A. (1980). Cumulonimbus Vertical Velocity Events in GATE. Part II: Synthesis and Model Core Structure. *Journal of the Atmospheric Sciences*, 37(11):2458–2469.

Zipser, E. J. and Lutz, K. R. (1994). The Vertical Profile of Radar Reflectivity of Convective Cells: A Strong Indicator of Storm Intensity and Lightning Probability? *Monthly Weather Review*, 122(8):1751–1759.